Cooperative Exploration in the Electronic Marketplace

David Sarne[1] and Sarit Kraus[1,2]
[1]Department of Computer Science, Bar-Ilan University, Ramat-Gan, 52900 Israel
[2]Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742
{sarned, sarit}@macs.biu.ac.il

Abstract

In this paper we study search strategies of agents that represent buyer agents' coalitions in electronic marketplaces. The representative agents operate in environments where numerous potential complex opportunities can be found. Each opportunity is associated with several different terms and conditions thus differing from other opportunities by its value for the coalition. Given a search cost, the goal of the representative agent is to find the best set of opportunities which fulfills the coalition's demands with the maximum overall utility, to be divided among the coalition members. Given the option of side-payments, this strategy will always be preferred by all coalition members (thus no conflict of interests), regardless of the coalition's payoff division protocol. We analyze the incentive to form such coalitions and extract the optimal search strategy for their representative agents, with a distinction between operating in B2C and C2C markets. Based on our findings we suggest efficient algorithms to be used by the representative agents for calculating their strategy and the appropriate derived expected utilities. A computational-based example is given, illustrating the achieved performance as a function of the heterogeneity level of the coalition's members.

Introduction

Cooperative coordination among agents and coalition formation processes have been widely used as a means for improving the efficiency of task performance and for saving costs in comparison to operating individually (Shehory & Kraus 1998; Tsvetovat *et al.* 2002). Recent research suggests various coalition formation models in different Multi Agent System (MAS) environments (Lermann & Shehory 2000; Shehory & Kraus 1998) and in electronic markets in particular (Breban & Vassileva 2001; Li *et al.* 2003; Tsvetovat *et al.* 2002; Yamamoto & Sycara 2001). In the context of the electronic marketplace, the most common application is a coalition of buyers, given the incentive of obtaining a volume discount according to the size of the coalition (Tsvetovat *et al.* 2002). In this paper we consider an additional benefit of a buyers' coalition, the benefit of improving the search for market opportunities through a coalition, given search costs.

We consider buyer agents to be associated with costs of search when there are no available central mechanisms that can supply full immediate information on the entire market opportunities. These search costs reflect the resources that need to be invested in search activities, such as locating seller agents, interacting with them, analyzing and comparing their offers and negotiating to provide the goods or service. Many authors have argued that advances in communication technologies allow consumers to obtain easier access to more information about price and product offers from alternative suppliers, thus reducing search costs and other market inefficiencies. However the general agreement is that these costs cannot be ignored completely (Bakos 1997). One may argue that the search cost for locating an opportunity is insignificant compared to the product price. Nevertheless, the continuous growth in the number of retailers and virtual stores over the Internet, followed by a phenomenal increase in the number of opportunities available, makes the overall search cost an important parameter affecting the search strategy (Choi & Liu 2000; Kephart & Greenwald 2002).

The existence of search costs, creates a strong incentive for buyer agents to form a coalition, that will conduct a search for potential opportunities jointly. Similar to most coalition formation models (Sandholm *et al.* 1999; Tsvetovat *et al.* 2002), we assume that the agents forming the coalition delegate their search to a representative agent whose search cost is comprised of two main components. The first is the cost of locating new opportunities, which is quite similar to the cost of any single buyer agent in its magnitude. The second is the coalition overhead, derived mainly from the communication and coordination activities (Sarne & Kraus 2003).

Three basic stages are common to all coalition formation models (Sandholm *et al.* 1999; Tsvetovat *et al.* 2002): coalition structure generation (where the agents form/join the coalition), executing the coalition task, and dividing the generated value among the coalition members. Among these stages, most of the proposed models studying coalition formation in electronic markets tend to focus on coalition generation and payoff division, while assuming the agents have complete knowledge concerning market opportunities. Nevertheless, as suggested above, in many scenarios the coalition needs to spend resources for executing its task of obtaining an appropriate set of opportunities for purchasing the product. Thus the problem of finding the optimal search strategy for the coalition, given its structure and the opportunity distribution, cannot be ignored. The representative agent operates in its environment alongside many other agents that represent coalitions differing in their size, their members' utility functions and the products they are seeking. These other coalitions, as well as the different individuals' utility functions play an important role when studying the stability of a coalition and issues of revealing the true utility function (truth telling). The

This research was supported in part by NSF under grant #IIS-0208608.

analysis of these important issues is based on the coalition's utility given any specific self structure (i.e. number of agents it represents and their reported, not necessarily true, utility functions) and the environment it is operating in. Our literature review suggests that mechanisms for finding the optimal search strategy of a coalition were not addressed to date.

The representative agent's goal is to maximize the overall coalition's utility. Assuming side-payments are possible, the overall utility maximization strategy will always be preferred by all coalition members (no conflict of interests), regardless of the coalition's payoff division protocol[1]. We apply the multi-attribute utility theory (MAUT) (Keeney & Raiffa 1976), to analyze preferences with multiple attributes in our agent based search mechanism. This enables a set of preferences to be represented by a numerical utility function. We consider the agents to be heterogeneous, each having its own unique utility function[2]. As we show in the next sections the representative agent's problem is different from any single buyer agent's problem in relation to its complexity, strategy structure and solution methodology. The analysis given relates to two different markets: The B2C (Business-to-Consumer) market, where sellers can supply almost any demanded volume, and the C2C (Consumer-to-Consumer) market, where sellers offer single items for sale. We show that every market suggests a different optimal strategy.

The following sections provide the full model formalization, analysis and algorithms for calculating the representative agent's optimal strategy. Throughout the paper we show the strong incentive of buyer agents to form a coalition, given search costs. For some environments we prove that the representative agent will inevitably increase the sum of the coalition members' utilities. For other environments, with uncertainty concerning the possible improvement, the buyer agents can use our proposed algorithms for calculating the expected cooperative search utility when considering this option. A complementary illustrative evaluation of performance in the different environments is given at the end of paper.

Related Work

The recognition of the advantages encapsulated in teamwork and cooperative behaviors, is the main driving force of many coalition formation models in the area of cooperative game theory and MAS (Lermann & Shehory 2000; Li *et al.* 2003; Shehory & Kraus 1998). Among the various domains investigated, special emphasis is placed on coalitions that can be formed in the electronic marketplace, generally using the advantage of discounts based on volume as the incentive for buyer agents to cooperate (Tsvetovat *et al.* 2002; Yamamoto & Sycara 2001). Additional coalition formation models for the electronic marketplace consider extensions of the transaction-oriented coalitions into long-term ones (Breban & Vassileva 2001), and for large-scale electronic markets (Lermann & Shehory 2000). Traditionally, the majority of this research has focused on issues concerning optimal division of agents into disjoint exhaustive coalitions (Sandholm *et al.* 1999; Yamamoto & Sycara 2001), division of coalition payoffs (Li *et al.* 2003; Yamamoto & Sycara 2001) and enforcement methods for interaction protocols. Only a few authors have considered the coalition representative agent's problem of determining its strategy in the electronic commerce domain, once the coalition is formed (Ito, Ochi, & Shintani 2002). However none have considered the agents' search cost and its influence over their search strategy as part of the problem formulation. In most mechanisms the underlying assumption is that an agent can scan as many opportunities as needed, or simply has a central view of the environment (Ito, Ochi, & Shintani 2002).

The problem of a searcher operating in a costly environment, seeking to maximize his long term utility is addressed in classical search theory ((Lippman & McCall 1976; McMillan & Rothschild 1994), and references therein). Here, various models can be found, where the dominant one is the *sequential* model (Lippman & McCall 1976), a multi-period model allowing only a single observation at each search period. Attempts to adopt the sequential search model into electronic trading environments associated with search costs are suggested in (Choi & Liu 2000; Kephart & Greenwald 2002). Nevertheless, in these works the focus is on a single agent's search, and the analysis of a cooperative search is lacking. As we show in the next sections this transition from a single agent to cooperative search encapsulates many complexities.

The Cooperative Exploration Model

We consider an electronic marketplace with numerous buyer and seller agents, interested in buying or offering to sell various well defined products. A product can be found offered by many different seller agents under various terms and policies (including price). Let $(B_1, B_2, ...B_k)$ be the set of attributes by which any opportunity for purchasing any product in this marketplace can be described. Each attribute B_i can receive any value in the interval $[\underline{b_i}, \bar{b_i}]$. A specific opportunity, $\vec{x_i} = (b_1, b_2, ...b_k)$, is a vector of specific values set by the seller, where b_i is the value of attribute B_i. As an agent conducts its search for a specific product it encounters different sellers and evaluates opportunities for obtaining the product.

We assume that while the buyer agents are ignorant of individual seller agents' offers, they are acquainted with the overall distribution of opportunities in the marketplace which can be described using the p.d.f. $f(\vec{x})$ and c.d.f. $F(\vec{x})$, or $P(\vec{x})$ for discrete environments[3]. Buyer agents can learn about specific opportunities in the market by locating seller agents and interacting with them. Each buyer agent, A_j, evaluates an opportunity, $\vec{x_i}$, by using its own multi-attribute utility function, $U_j(\vec{x_i})$, defined over the set of attributes $(B_1, B_2, ...B_k)$. Buyer agents might have heterogeneous preferences and thus the utility from a given opportunity differs according to the evaluating buyer agent. We assume that all agents share a consistent preference for each attribute, i.e. given attribute B_i, either $\frac{dU_j(\vec{x})}{dB_i} \geq 0$ for any buyer agent A_j or $\frac{dU_j(\vec{x})}{dB_i} \leq 0$ for any buyer agent A_j. For simplicity, and without loss of generality, in this paper we consider agents with utility functions increasing in the value of any of their attributes.

We also assume there are no central mechanisms or mediators which can supply the agents with full immediate informa-

[1] The payoff division protocol defines the portions each coalition member receives from the overall coalition utility, thus improving the overall coalition utility is always preferred.

[2] Product attributes (performance and functionality), terms and policies (concerning warranties, return policy, payment policy, delivery time and policy, etc.), as well as reputation and trust factors are among the attributes influencing a buyer agent's utility.

[3] There are many ways an agent can learn about market opportunities' distribution - using market indicators, spectator agents, etc.

tion concerning current market opportunities, thus the agents need to search for opportunities. In its most basic form, each buyer agent searches sequentially, in such a way that it locates a seller and learns about a new opportunity $\vec{x_i}$ at each search stage. Based on the agent's evaluation of the utility that will be gained from this opportunity, $U_j(\vec{x_i})$, the agent makes a decision whether to exploit the current opportunity (i.e. buy from this seller) or resume its search in a similar method.

The search activity is assumed to be costly. For each search stage in which the buyer agent locates, interacts and evaluates a new seller agent (resulting in a new opportunity), the process induces a specific search cost c. This cost is principally a parameter of the market's liquidity and volatility, and thus shared by all buyer agents operating in the specific marketplace. We assume a buyer agent's utility from a given opportunity may be interpreted into monetary terms. Thus the utilities are additive and the total search utility can be obtained by subtracting the search cost from this value. Recognizing the benefits of a cooperative search, buyer agents, interested in similar products or interchangeable products, may form coalitions. Each new coalition formed is handled by a representative agent. The representative agent also conducts its search sequentially, scanning opportunities one by one. As suggested in the introduction, in addition to the fixed search component c, the representative agent incurs a certain overhead (e.g. communicating with the coalition members during the search), which is a parameter of the coalition size. We denote the total search cost associated with each additional search stage of the representative agent c_N (where N is the number of buyers represented by the coalition) satisfying: $\frac{dc_N}{dN} \geq 0$. We do not limit the decision horizon nor the required number of search rounds. This will be determined by the agent given the search cost and the opportunities it encounters.

We define the search strategy of a representative agent as a function S, assigning a binary decision (terminate or resume search), to any known set of opportunities ($S(\vec{x_1}, ..., \vec{x_m}) \rightarrow (terminate, resume)$). Given the representative agent's goal of maximizing the overall coalition utility, its decision is not influenced by the payoff division protocol, nor by coalition stability considerations, but rather influences these two factors. The overall coalition expected utility is obtained by subtracting the coalition cost of the search from the aggregated utilities of the coalition members at the end of the process. As suggested earlier, since any agent's portion of the total net payoff is pre-determined, and increases with the increase of the net payoff, the overall utility maximization strategy is the preferred strategy by all agents at every stage of the search.

A Single Agent's Search

Before analyzing the cooperative search, we wish to recall some known results for the problem of a single searcher that can be found in classical search theory literature (Lippman & McCall 1976). Given the electronic marketplace environment, the single agent's search strategy is stationary (i.e. it does not change from one search stage to another). The agent sets a lower limit stopping rule and terminates the search when reaching an opportunity that yields a utility greater than or equal to this limit. The agent's expected utility, denoted $V(U)$, when using a limit of value U_{rv} can be expressed as:

$$V(U_{rv}) = \frac{-c + \int_{U(\vec{y}) \geq U_{rv}} U(\vec{y}) f(\vec{y}) d\vec{y}}{\int_{U(\vec{y}) \geq U_{rv}} f(\vec{y}) d\vec{y}} \tag{1}$$

The agent's optimal lower limit satisfies $V(U_{rv}) = U_{rv}$, thus the optimal expected utility can be derived by substituting the equality in (1). Notice that these results are valid regardless of the market the agent is operating in (B2C or C2C) since it searches for a single item.

The above result is important as it supplies a good benchmark for the cooperative search performance. Adding a specific agent to the coalition is favorable, as long as the derived increase in the overall coalition utility is greater than the agent's expected utility obtained by searching by itself.

Before continuing our analysis, we would like to point out that the calculation of the above cost function (and those given in the next section) might not be trivial or may contain non-integrated terms. This can be resolved using simple approximation methodologies which are proved to produce accurate results for functions of these types (Sarne & Spiegler 2005).

Cooperative Search Strategy

The major benefit of the cooperative search is the potential of exploiting opportunities which would have been discarded in any of the alternative single agents' separate searches for the benefit of other agents in the coalition. In fact, the probability that any new opportunity encountered will improve the coalition utility increases as the variance between the different coalition member's utility functions increases. Furthermore, for the specific case when there are no coalition creation and maintenance costs, the coalitional search will always be preferable, as the following proposition suggests.

Proposition 1 *When the coalition does not produce any overhead search costs (i.e. $c_N = c$), then the search through a coalition of heterogeneous agents will always improve the overall utility (the sum of members' utilities).*

Proof: Extract the expected number of opportunities to be scanned when each coalition member conducts the search by itself (notice that according to (1) this number is (random) geometric and the probability of success is $\int_{U(\vec{y}) \geq U_{rv}} f(\vec{y}) d\vec{y}$) and locate this number of opportunities. The assign the buyer agent the one opportunity from which it will gain the highest utility. This should result in an equal utility for the two methods (as the cost and expected utilities from the best opportunity for each agent are equal). However, the result for each buyer agent in the coalition can now be improved by possibly using some of the opportunities discarded along the above process for other agents with different utility functions. □

The major complexity of finding the optimal strategy for the representative agent lies in the difference between this form of search and the traditional single agent's search. The main difference is that the representative agent seeks an optimal *set* of opportunities (matched to the different members in a certain order) rather than a single opportunity. Each member in the set has a different affect on the overall coalition utility, and since new opportunities are being evaluated sequentially, the optimal strategy must be defined as a function of the currently known opportunities (in comparison to the stationary fixed limit-based strategy in a single agent's search). Furthermore, in this case the determination of whether a search through a coalition is beneficial, in terms of the overall members' utility (given a coalition overhead component in the search cost), is not trivial. Here, the difference between the expected utilities obtained using the two methods (cooperative search and aggregated separate autonomous searches) depends on the opportunities' distribution and the similarity level among the coalition members' utility functions (in addition to the search cost parameter). In the following section we describe how the representative agent's strategy can be constructed as a function of the market it is operating in.

B2C Markets

Operating in a B2C market, where sellers can supply almost any demanded volume, the representative agent can assign any new opportunity encountered to any number of buyer agents in the coalition simultaneously. Therefore in addition to the benefit of better exploiting any encountered opportunity, any new opportunity being evaluated has a much greater potential affect over the coalition's overall utility.

Consider a representative agent representing N agents with different utility functions $U_1, U_2, ..., U_N$ operating in a B2C marketplace. Throughout its search, the representative agent maintains the set of opportunities $O = (\vec{x_1}, \vec{x_2}, ..., \vec{x_N})$ where $\vec{x_j}$ is the opportunity with the maximum utility for agent A_j from all opportunities known to the representative agent at the current point. Notice that the same opportunity might be reused for more than one agent in the set. This is inevitably the case during the first N search rounds, and possibly throughout the entire search. In each stage of the search, the representative agent considers the structure of its set O as an input for its decision. Let $V(\vec{x_1}, \vec{x_2}, ..., \vec{x_N})$ denote the coalition's expected utility from the point where the representative agent is acquainted with $O = (\vec{x_1}, \vec{x_2}, ..., \vec{x_N})$, then in case the representative agent has decided to resume its search, we obtain:

$$V(\vec{x_1}, \vec{x_2}, ..., \vec{x_N}) = \frac{-c_N + \int\limits_{U_i(\vec{y}) \geq U_i(\vec{x_i}), i \in (1,...,N)} V(\vec{x}_1^*, \vec{x}_2^*, ..., \vec{x}_N^*) f(\vec{y}) d\vec{y}}{\int\limits_{U_i(\vec{y}) \geq U_i(\vec{x_i}), i \in (1,...,N)} f(\vec{y}) d\vec{y}} \quad (2)$$

where $\vec{x}_i^* = arg \max_{\vec{x} \in \vec{x_i}, \vec{y}} U_i(\vec{x})$; $i = 1, ..., N$.

Otherwise, if the agent decides to terminate its search, then the immediate utility obtained is:

$$V_{terminate}(\vec{x_1}, \vec{x_2}, ..., \vec{x_N}) = \sum_{i=1}^{N} U_i(\vec{x_i}) \quad (3)$$

The representative agent will terminate its search at any given stage if $V_{terminate} \geq V(\vec{x_1}, \vec{x_2}, ..., \vec{x_N})$, i.e. upon adding an opportunity to the set O, where, given the new set created, the immediate utility that can be obtained is greater than or equal to the expected utility when resuming the search.

Figure 1(a) illustrates the space of opportunities that improves the coalition utility, given the set of known opportunities O. The sketch describes a given set $O^i = (\vec{x}_1^{\,i}, \vec{x}_2^{\,i}, \vec{x}_3^{\,i})$ of three different opportunities (thus correlated with $N \geq 3$ buyer agents). The lower left area represents all opportunities that do not suggest any improvement in the coalition utility. When reaching any of the opportunities in this area, the representative agent will certainly resume its search without changing the set O^i. The right upper area represents all opportunities that will undoubtedly improve the coalition utility, and thus will enter the set O^i, replacing one or more of the opportunities (hense creating a new set O^j). Given the new set O^j the representative agent will need to decide whether to resume or terminate its search at this point. The four rectangular middle areas represent opportunities for which a calculation is required to determine their affect.

Recalling the recursive form of (2), the following theorem suggests an important simplification of the calculation of the coalition's expected utility if the search is resumed.

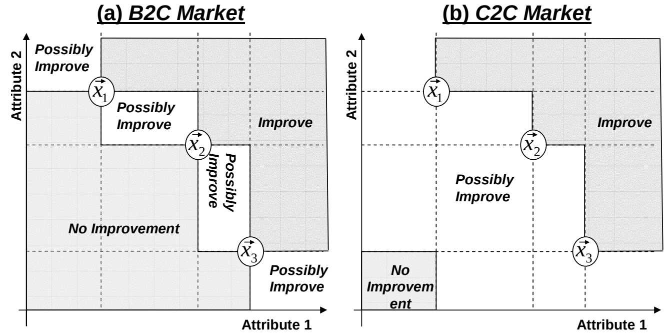

Figure 1: Utility improvement opportunities

Theorem 1 *If the representative agent's optimal strategy given set O^i is to continue its search, then so is its strategy given set O^j where $\vec{x}_v^{\,j}[m] \leq \vec{x}_l^{\,i}[m]$ $\forall$ $v = 1, ..., N$; $m = 1, .., k$; $l \in (1, ..., N)$.*

Proof: Suppose the representative agent, when possessing set O^j adopts the same strategy as when possessing O^i (i.e. taking the same decision whether to terminate or resume the search, for any new opportunity, as if having O^i). The difference in the agent's expected utility when adopting a new opportunity $\vec{y}$, starting with O^j is always greater than in the case of starting with O^i. Thus the sum of utility improvements until termination of the search when starting with O^j is always greater than when starting with O^i. Since the agent in the case of O^j mimics the strategy taken for O^i then the search costs along the search are also equal. Thus if the agent has an incentive to continue its search given set O^i then certainly it has such an incentive when given O^j. $\square$

Considering figure 1(a), assuming the agent's strategy given set O^i is to resume its search, then the strategy of the agent given any set O^j of opportunities in the left bottom area of the graph is undoubtedly to resume the search.

Theorem 1 can be utilized to suggest an efficient algorithm which facilitates the calculation of a representative agent's optimal strategy. The algorithm starts with a set of known opportunities, and returns the appropriate set of decisions to be taken given any future sets of opportunities.

Algorithm 1 *An algorithm for finding the optimal strategy of an agent representing a coalition of heterogeneous buyer agents in the B2C marketplace, given a set of known opportunities.*
Input: *Ψ - a set of product attributes; $\theta_{market}[]$ - a collection of all feasible market opportunities; P[] - probabilities array, where P[i] is the probability of reaching opportunity $\theta_{market}[i]$ in a search stage; N - number of agents in the coalition; U[] - utilities array, where U[i] is the utility function of agent A_i; $\Omega_{known}[]$ - a collection of known opportunities.*
Output: *$Strategy[]$ - a collection of opportunities sets and the appropriate action to take when reaching each set (terminate\resume).*
Internal parameters: *O_{known} - stores the effective subset of Ω_{known} for the given coalition; $\Omega_{scanned}[]$ - a collection of opportunity sets for which a strategy was already determined; $\Omega_{potential}[]$ - a collection of opportunity sets to be scanned at the current stage.*

1. Set $O_{known} = (\vec{x_1}, ..., \vec{x_N})$ where:
$\vec{x}_i = argmax_{\vec{x} \in \Omega_{known}}(U[i](\vec{x}))$, $i = 1, ..., N$
2. Set $\Omega_{scanned}[]$=null;
3. Set $\Omega_{potential}[] = (O_1, ..., O_m)$, $O_i \notin \Omega_{scanned} \forall i = 1, ..., m$ where for each $O_i[k]$ $(i \in (1, ..., m),\ k \in (1, ..., N))$ $\exists attr_z \in \Psi$ such that $\forall j \in (1, ..., N) k \neq j$ holds $O_i[k].attr_z > O_i[j].attr_z$, and for each $O_i, O_j \in \Omega_{potential}[]$ $\exists k \in 1, ..., N$ and $(attr_{z_1} \in \Psi, ..., attr_{z_N} \in \Psi)$ such that $O_i[k].attr_{z_l} > O_j[l].attr_{z_l}$ $l = 1, ..., N$.

4. For each member in $\Omega_{potential}$ store the maximum between the expected utility from terminating and from resuming the search (using equations (3) and (2)), assuming the strategy of all other members of $\Omega_{potential}$ is to terminate search.
5. Add O_i associated with the highest utility in $\Omega_{potential}[]$ to $Strategy[]$, along with its expected utility and optimal act (terminate or resume) when reaching this set.
6. If O_i added in step 5 is O_{known} then return $Strategy[]$ and exit.
7. Goto step 3

Theorem 2 *Algorithm 1 is finite and upon termination returns the optimal strategy of the representative agent given a set of known opportunities.*

Proof: The first time the algorithm reaches step 4, it will add the opportunity set yielding the highest utility for the coalition to the array $\Omega_{scanned}[]$. Obviously, the decision of the representative agent when having this set is to terminate the search (having no future opportunity outperforming the current one). Now consider any other execution step 4. The expected utility of the selected set to be entered into $\Omega_{scanned}[]$ will always be smaller than the expected utility of formerly marked opportunity sets. Thus the strategy that was set for formerly entered sets will not be affected by this new set, even if this set improves its expected utility by adopting a search continuation strategy. As we always move sets into $\Omega_{scanned}[]$ and never take any out, and $\theta_{market}[]$ is a finite set, the algorithm will terminate, reaching O_{known} in a finite time. □

Notice that in step 4 of the algorithm, the calculation according to (2) and (3) is immediate, as for each potential future state, the expected utility is either the sum of the different utility functions or has already calculated in former steps.

Algorithm 1 starts with a given set of known opportunities. Nevertheless, since the first opportunity encountered is feasible for all coalition members, the performance (expected utility) of the representative agent when starting the search from scratch (i.e. with zero known opportunities) is given by:

$$V = -c_N + \int_{\vec{y}} V(\vec{y}, \vec{y}, ..., \vec{y}) f(\vec{y}) d\vec{y} \tag{4}$$

C2C Markets

Here the representative agent needs to supply the coalition N different opportunities (as sellers offer single items). Thus an opportunity which potentially improves the utility of several buyer agents in the coalition can be applied only to one of them. Hence the expected performance of the coalition in C2C markets is lower than in B2C markets.

The analysis methodology for this case is quite similar to the one used in the previous section. Upon terminating the search, the representative agent's immediate utility is defined by (3). However restricting each opportunity to a single coalition member, requires modification of (2):

$$V(\vec{x_1}, \vec{x_2}, ..., \vec{x_N}) = -c_N + \int_{\vec{y}} V(\vec{x}^*_1, \vec{x}^*_2, ..., \vec{x}^*_N) f(\vec{y}) d\vec{y} \tag{5}$$

Where the set $(\vec{x}^*_1, \vec{x}^*_2, ..., \vec{x}^*_N)$ is a sorted subset of $(\vec{x}_1, \vec{x}_2, ..., \vec{x}_N, \vec{y})$ yielding the maximum utilities' sum, given the different agents' utility functions.

Notice that in this case the space of opportunities for which we cannot determine a-priori whether adopting them will result in a utility improvement or not, increases significantly (a simple illustration is given in figure 1(b) for the same settings used to create the equivalent B2C example). Furthermore, the determination process for each of these opportunities in this undetermined space becomes more complex. While in the case of B2C markets the adoption of a new opportunity could have been evaluated simply by checking if it can improve any of the coalition members' utilities, here a permutation analysis is required. The representative agent must extract the entire set of permutations and check the overall utility obtained from each permutation, since new opportunities that do not increase any of the personal members' utilities might be adopted and placed instead of existing opportunities in the set $(\vec{x_1}, \vec{x_2}, ..., \vec{x_N})$. Such a scenario is certainly non-intuitive and thus will be demonstrated by a simple example. Consider a representative agent, representing two buyer agents A_1 and A_2 with different utility functions. Suppose that the current set O contains two opportunities, $\vec{x_1}$ and $\vec{x_2}$, with the associated utilities $U_{A_1}(\vec{x_1}) = 80$; $U_{A_1}(\vec{x_2}) = 20$; $U_{A_2}(\vec{x_1}) = 20$; $U_{A_2}(\vec{x_2}) = 10$. Obviously the representative agent will assign $\vec{x_1}$ to A_1 and $\vec{x_2}$ to A_2 (with an overall utility of 90). Now assume, in its next search round the representative agent found an opportunity $\vec{x_3}$ with the associated utilities $U_{A_1}(\vec{x_3}) = 75$; $U_{A_2}(\vec{x_3}) = 5$. Notice the new opportunity $\vec{x_3}$ does not improve either of the two agents' utilities. Nevertheless, by considering the permutation where $\vec{x_3}$ is assigned to A_1 and $\vec{x_1}$ is assigned to A_2, we obtain a better overall utility (95).

The expected utility when starting the search from scratch in the case of operating in a C2C market is given by:

$$V = -Nc_N + \int_{\vec{y_1}} ... \int_{\vec{y_N}} V(\Theta(\vec{y_1}, ..., \vec{y_N}) f(\vec{y_N}) d\vec{y_N} ... f(\vec{y_1}) d\vec{y_1} \tag{6}$$

where $\Theta(\vec{y_1}, ..., \vec{y_N})$ is the permutation that maximizes the coalition's expected utility given the opportunities $\vec{y_1}, ..., \vec{y_N}$. The element Nc_N in the above equation is associated with N opportunities that need to be acquired as a starting point.

We propose a modification of Algorithm 1 for the C2C market. Here, the assignment of opportunities to the different coalition members, given a set of opportunities, is made through a permutation analysis and the expected utility when continuing the search is calculated according to (5) (instead of (2)). In addition, the set of possible candidates for the collection $\Omega_{potential}[]$ is expanded as demonstrated in figure 1(b).

Comparative Illustration

In this section we aim to illustrate and give the reader a general perspective of the performance (coalition expected utility) of representative agents when searching from scratch in the different scenarios described in the former sections.

For this purpose we used an environment where the opportunities space is based on two attributes B_1 and B_2, with possible values for each attribute uniformly distributed in the range (1,...,5) using discrete values. Two agent types were considered: The first agent was associated with a utility function $U_1(\vec{x}) = B_1 + B_2$ and the second agent with $U_2(\vec{x}) = 2(1-\alpha)B_1 + 2\alpha B_2$. The expected utilities were calculated with respect to the value α (indicating the level of agents' similarity) along the horizontal axis. The search cost of any single agent was taken to be $c = 0.2$ and for a coalition $c_N = c \cdot ln(N+1)\ \forall (N > 1)$. Curves 1 and 2 of figure 2, depict the expected utility of agents A_1 and A_2, respectively, when searching autonomously by themselves. Notice that the expected utility of agent A_2 increases as the ratio of the two attribute coefficients in its utility function increases. This observation can be explained by the increase in variance

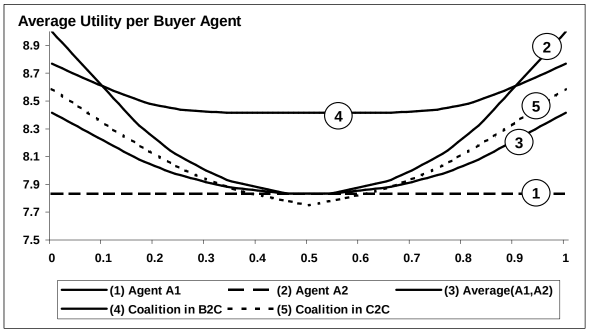

Figure 2: Coalition Performance

of the opportunities' utility as parameter α becomes closer to the extremes. Such an increase results in a greater incentive to continue the search, as the possible improvement in the agent's expected utility increases. Curve 3 is the average of curves 1 and 2, reflecting the average expected utility of the two agents when each agent searches autonomously by itself. This should be the point of reference when evaluating any coalition search performance. Curves 4 and 5 represent the average expected utility per participating agent when both agents are searching through a coalition in B2C and C2C markets, respectively. As expected the coalition search in B2C markets results in a better expected utility compared to the same coalition search in C2C markets. Notice that in this scenario the average expected utility of the coalition operating in the B2C market outperforms the average expected utility of the self search case, though this is not a general observation. The expected utility of the coalition that searches in the C2C market, outperforms the self search option for the non-central values of α where the coalition has an advantage of reusing opportunities for the benefit of other coalition members.

Obviously a complete analysis concerning the magnitude of improvement achieved in the cooperative search as a function of the different model parameters would require further detailed and more comprehensive scenarios and environments. Nevertheless, such an analysis is beyond the scope of this paper, as our main focus is on the introduction of the general model and its unique solution method.

Discussion and Conclusions

The analysis given suggests that in many markets associated with a search cost, autonomous agents have a strong incentive to search as a coalition. Nevertheless, the search strategy the coalition adopts is different in its structure from the optimal strategy used by a single buyer agent. Based on the proposed analysis, the representative agent can calculate its optimal strategy given the utility functions of the coalition members and the specific environment in which it operates.

In addition to the theoretical importance of the results we see a potential for an actual implementation. Given today's electronic markets, the most applicable one for a cooperative search seems to reside in B2C markets. In fact, the basic infrastructure of group buying can be found today in eCommerce sites, however it is based on an already known prenegotiated specific opportunity. Additionally, in B2C markets the concept of volume discount is applicable, increasing the incentive to cooperate (in this case the agents will benefit both from reducing the search cost and the discount price). In C2C markets products highly vary in brand, quality and age, thus it is more difficult to formulate the buyer agents' utility function. Nevertheless, we do believe that agents in future C2C markets will also benefit from such models.

As suggested in the introduction, other than extracting the optimal search strategy for the coalition representative agent, the coalition formation process involves many issues that were not included in the scope of this paper. The two most important issues are coalition stability, given the MAS settings and the division of payoffs (in terms of side payments) between the coalition members. The latter issue also affects the agents' incentive for revealing their true utility functions. The optimal representative agent's strategy and its associated expected utility are important inputs for the analysis of these two issues. The rich literature in the area of game theory and MAS research (Sandholm *et al.* 1999; Yamamoto & Sycara 2001; Li *et al.* 2003), can enable us to adopt many results and ideas to resolve these issues for the model presented herein. A promising mechanism that we are currently evaluating in terms of coalition stability is the division of payoffs according to a member's contribution to the overall coalition utility.

References

Bakos, Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Sci.* 42:1676–92.

Breban, S., and Vassileva, J. 2001. Long-term coalitions for the electronic marketplace. In *B. Spencer, ed., Proc. of E-Commerce Applications Workshop*.

Choi, S., and Liu, J. 2000. Optimal time-constrained trading strategies for autonomous agents. In *Proc. of MAMA'2000*.

Ito, T.; Ochi, H.; and Shintani, T. 2002. A group-buy protocol based on coalition formation for agent-mediated e-commerce. *IJCIS* 3(1):11–20.

Keeney, R., and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley & Sons.

Kephart, J., and Greenwald, A. 2002. Shopbot economics. *JAAMAS* 5(3):255–287.

Lermann, K., and Shehory, O. 2000. Coalition formation for large scale electronic markets. In *Proc. of ICMAS'2000*, 216–222.

Li, C.; Chawla, S.; Rajan, U.; and Sycara, K. 2003. Mechanisms for coalition formation and cost sharing in an electronic marketplace. In *Proc. of EC03*, 68 – 77.

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 *Handbook of Game Theory with Economic App.* 905–927.

Sandholm, T.; Larson, K.; Andersson, M. R.; Shehory, O.; and Tohme, F. 1999. Coalition structure generation with worst case guarantees. *AIJ* 111:209–238.

Sarne, D., and Kraus, S. 2003. Search for coalition formation in costly environments. In *Proc. of CIA03*, 117–136.

Sarne, D., and Spiegler, I. 2005. Approximating optimal search strategy in emarketplaces. Working Paper, The Recanati Graduate School of Business Adm., Tel-Aviv Uni.

Shehory, O., and Kraus, S. 1998. Methods for task allocation via agent coalition formation. *AIJ* 101:165–200.

Tsvetovat, N.; Sycara, K. P.; Chen, Y.; and Ying, J. 2002. Customer coalitions in electronic markets. In *Proc. of AMEC-2000*, 121–138.

Yamamoto, J., and Sycara, K. 2001. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proc. of AA01*, 576–583.