Integrating Parallel Interactions into Cooperative Search

Efrat Manisterski [1]
[1] Department of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
maniste@cs.biu.ac.il

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

Sarit Kraus [1, 3]
[3] Institute for Advanced Computer Studies
University of Maryland
College Park, MD
sarit@cs.biu.ac.il

ABSTRACT

In this paper we incorporate autonomous agents' capability to perform parallel interactions into the cooperative search model, resulting in a new method which outperforms the currently used ones. As a framework for our analysis we use the electronic marketplace, where buyer agents have the incentive to search cooperatively. The new search technique is quite intuitive, however its analysis and the process of extracting the optimal search strategy are associated with several significant complexities. These difficulties are derived mainly from the unbounded search space and simultaneous dual affects of decisions taken in different world states. We provide a comprehensive analysis of the model, highlighting, demonstrating and proving important characteristics of the optimal search strategy. Consequently, we manage to come up with an efficient modular algorithm for extracting the optimal cooperative search strategy for any given environment. A computational based comparative illustration of the system performance using the new search technique versus the traditional methods is given, emphasizing the main differences in the optimal strategy's structure and the advantage of using the proposed model.

Categories and Subject Descriptors

K.4.4 [**Computing Milieux**]: Computers And Society —*Electronic Commerce*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence —*Multiagent systems*

General Terms

Algorithms, Design, Economics, Performance, Theory

Keywords

parallel, cooperative search, search cost

1. INTRODUCTION

Coalition formation is well recognized as a key process in a multi-agent systems, mostly desirable in environments where a group of agents can perform a task more efficiently than any single agent can [10]. In recent years many coalition formation models have been suggested, for various domains [18, 5], particularly for electronic commerce [19, 20, 16]. In the latter context, the most common coalition concerned is a coalition of buyers, derived mainly from the potential of obtaining volume discounts [19, 15] and the ability to search cooperatively for market opportunities in a more efficient manner [16].

The cooperative search incentive derives principally from the existence of search costs found in MAS, reflecting the resources (not necessarily monetary), that need to be invested/consumed while searching for opportunities in the environment (e.g. searching for an opportunity to buy a product in the context of the electronic marketplace) [16]. The scenario of having search costs is common in MAS where the agent needs (for its decision making process) immediate information concerning market opportunities. Given the richness of opportunities and the dynamic and open nature of these environments, central mechanisms are usually incapable of supplying such information with the level of completeness and accuracy required by the agent, certainly not without a cost. Thus the agent needs to spend some of its resources on search related activities. Despite the reduction in the magnitude of these search costs in the electronic commerce era, 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 buyers' search strategy [4, 9, 16, 1].

By forming a coalition and delegating the search task to a representative agent (acting on behalf of the coalition[1]), the cooperative search offers the advantage of sharing, reusing and re-allocating opportunities among the coalition members (e.g. exploiting opportunities which would have been discarded otherwise if each of the agents would have performed an alternative separate search) [16]. Nevertheless, the process of forming and maintaining the coalition induces some overhead, derived mainly from the communication and coordination activities [15], thus the representative agent should set its search strategy in a cost/effective manner.

A classic example of the above in traditional markets is the procurement management officer of a corporation. Instead of having each individual in the cooperation spend time and resources on locating its specific requested supplies, the task is delegated to the procurement management officer. Here, in addition to the price discounts obtained for aggregated demands of identical items, the procurement management officer becomes highly updated with the

[1]As discussed in detail in [16], the use of a representative agent is a common and efficient means for executing the coalition's goal. Furthermore, given the option of side-payments the overall utility maximization strategy taken by the representative agent is always the preferred one by all coalition members (i.e. no conflict of interests), regardless of the pre-set coalition's payoff division protocol.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

different offerings and specific supplies available by the different merchants in the markets. As a result the cost of locating the right deal for each request becomes significantly smaller (in comparison to the equivalent search conducted by each of the individuals).

The basic concepts by which a representative agent should manage the cooperative search, including an analysis and computational means for extracting its optimal search strategies are given in [16]. Nevertheless, the assumption used in that model for constructing the representative agent's strategy is that the representative agent interacts with one seller agent at a time as part of its search process. Such an assumption ignores an inherent strength of autonomous agents, which is their capability to efficiently interact with several other agents in parallel (enhanced by their improved communication capabilities and their ability to process an enormous amount of information in a short time, compared to people). The parallel interaction is preferable when an agent's search cost combines fixed components (e.g. operational costs) and/or non-linear dependency on the number of interactions maintained (e.g. advantage of size). In such cases the adoption of the parallel technique by the representative agent suggests a reduction in the average cost per interaction with seller agents. While the integration of parallel interactions technique into a single search process is quite intuitive, its integration into a coalition's search is not trivial at all. The major difficulty derives from the fact that different coalition members may have heterogeneous multi-attribute utility functions. In this paper we supply a comprehensive analysis of the parallel cooperative search model and present an algorithm which can facilitate the calculation process of the coalition strategy. The new searching technique results in at least as good (and in many environments significantly better) expected utility for the coalition.

Similar to the model introduced in [16], we apply the multi-attribute utility theory (MAUT) [8], 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 utility function. The model is general, though several specific implementation aspects relating to 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, are emphasized. 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.

Notice that among the three basic stages that are common to all coalition formation models [14, 19]: coalition structure generation (where the agents form/join the coalition), executing the coalition task, and dividing the generated value among the coalition members, our focus is on finding the optimal search strategy for the coalition, given its structure and the opportunity distribution. As suggested in [16], 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 individual utility functions play an important role when studying the stability of a coalition and issues of revealing the true utility function (truth telling). The analysis of these important issues is based on the ability to properly derive 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. This paper aims to supply this functionality, laying the foundations and enabling research of many of the important aspects of coalition formation given above in the context of cooperative search (truth telling, stability, payoff division, etc.).

The main contributions of this paper are threefold: First, we formally model and analyze the parallel cooperative search problem of agents operating in a costly environment. This model is a general search model and can be applied in various domains in addition to the electronic marketplace that is being used as a framework for our work. Second, we show that in many environments the parallel cooperative search outperforms the traditional search strategies (either when each agent searches by itself or when using a cooperative sequential search). Furthermore, we draw attention to scenarios where traditional cooperative search is proven to be non-beneficial, however parallel cooperative search is a favorable technique. Finally, we supply an algorithm that facilitates the calculation of the coalition's optimal strategy, and significantly reduces the complexities associated with the attempt to extract this strategy from an appropriate set of equations.

In the following section we address relevant multi-agent and parallel search literature. The parallel cooperative search model assumptions and a formal description are given in sections 3 and 4, respectively. In section 5, we explore the unique characteristics of the representative agent's optimal strategy when using cooperative parallel search, leading to the introduction of an efficient algorithm for extracting it. In section 6 we demonstrate the unique properties of the parallel search method using as a reference two of the classical search models that our suggested model generalizes. We conclude and suggest directions for future research in section 7.

2. RELATED WORK

In many scenarios autonomous agents in multi-agent environments may cooperate in order to perform tasks. The need for cooperation may arise either when an agent is incapable of completing a task by itself or when operating as a group can improve their overall performance [3, 10, 19]. Group based cooperative behavior can be found in various domains, such as solving complex optimization problems [18], military and rescue domains [5], e-business applications [19, 20] and many more. In the electronic market domain, most authors focus on coalitions formed to obtain volume discount [19, 20]. Additional coalition formation models for the electronic marketplace consider extensions of the transaction-oriented coalitions into long-term ones [3], and for large-scale electronic markets [10]. Traditionally, the majority of this research effort has focused on issues concerning optimal division of agents into disjoint exhaustive coalitions [14, 20], division of coalition payoffs [20] and enforcement methods for interaction protocols. Only a few authors have considered the coalition's problem of determining its strategy in the electronic commerce domain, once the coalition is formed [7]. Nevertheless, other than in [16], none of the proposed models have considered a coalition's search in a costly environment, and in particular none of them (including [16]) have made use of the agent's capabilities to maintain parallel interactions.

The problem of a searcher operating in a costly environment, seeking to maximize his long term utility is addressed in classical search theory ([11, 12], and references therein). The three main search models that can be found in the literature are the fixed sample size model, the sequential model and the variable sample size model. In the fixed sample size model, introduced in [17], the searcher draws a single sample where all observations are taken simultaneously. In the sequential search model [11] the searcher draws one observation at a time, allowing multiple search stages. The third search model [2, 6, 13] suggests a combined approach in which several observations may be made at any period. This latter method, which outperforms the other two, is in fact the single agent's equivalent to our cooperative search model considered in this paper. In addition to the establishment of some general properties of a search rule, most of the papers mentioned above considered mainly issues of uncertainty associated with the availability of the inspected offers [6], recall and fallback utilities [13], and the effect of a final decision horizon on the search strategy [2]. Nevertheless, the focus of the above works from the search theory domain is on a single agent's search, and the analysis of a cooperative search is lacking. As shown in [16], the transition from a single agent search to a cooperative search concept as the one we base our model on, is not trivial and encapsulates many complexities resulting in a different strategy structure. In the cooperative search model, the representative agent needs to take into consideration the affect of new opportunities found on any subset of utility functions associated with the different coalition members. Though the process of extracting the optimal strategy is very different from the one used by a single agent that tries to maximize a simple or an aggregated utility function with a single opportunity.

3. THE MODEL

We base our model description and formulation on the definitions given in [16] and extend them to reflect agents' parallel search capabilities. We consider an electronic marketplace where numerous buyer and seller agents can be found. Each agent is 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 price). We assume that while buyer agents are ignorant of individual seller agents' offers, they are acquainted with the overall distribution of opportunities (defining an opportunity as the option to buy the product under specific terms and policies) in the marketplace.

Assuming there are no central mechanisms or mediators which can supply the agents with full immediate information concerning current market opportunities, they need to search for appropriate opportunities to buy their requested product. Throughout the search the buyer agents locate seller agents and learn about their offers by interacting with them. Each buyer agent evaluates opportunities using its own multi-attribute utility function. Buyer agents may have heterogeneous preferences and thus the utility from a given opportunity differs according to the evaluating buyer agent.

In its most basic form, each buyer agent searches in such a way that it interacts with several sellers in parallel at each stage of its search thus learns about a new set of opportunities. Based on the agent's evaluation of the utility that can be gained from each opportunity in the set, the agent makes a decision whether to exploit any of the opportunities it encountered throughout its search (i.e. buy from any of the sellers) or resume its search in a similar method. A decision of resuming search is always accompanied with the number of parallel interactions to be executed next.

The search activity is assumed to be costly [4, 9, 16, 1]. For each search stage in which the buyer agent locates, interacts and evaluates seller agents, the process induces a search cost. This cost is a function of the number of parallel interactions initiated and maintained by the agent. The search cost structure is principally a parameter of the market's liquidity and volatility, and thus shared by all buyer agents operating in the specific marketplace. Recognizing the benefits of a cooperative search, buyer agents, interested in similar products or interchangeable products, may form coalitions [16]. Any coalition formed will always be represented by a representative agent, which conducts its search on behalf of the coalition in a similar method (encountering sellers and accumulating new opportunities). The representative agent's search cost increases both as a function of the number of parallel interactions it forms and the number of buyer agents it represents[2]. 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 that the process induces from this value.

As part of its search process, the representative agent needs to set a strategy for determining, given any set of known opportunities, whether to terminate or resume its search. In the latter case, the agent also needs to determine the number of parallel interactions to be used in the next search round. The optimal strategy is the one maximizing its expected total search utility (opportunities utility minus search costs). 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 [16]. Any of the agents' pre-determined portion of the coalition's utility will increase in its absolute value along with the increase of the net coalition utility, thus the overall utility maximization strategy is the preferred strategy by all agents at every stage of the search.

4. PROBLEM FORMULATION

Let $B = (B_1, B_2, ..., B_k)$ be the set of the attributes defining any of the potentially available opportunities in the market, where each attribute B_i can be assigned a value from the finite set $(b^i_{min}, ..., b^i_{max})$. An opportunity's type is defined by the vector $\overrightarrow{o_i} = (b_1, b_2, ...b_k)$, assigning a value b_i to each specific attribute B_i. We denote by O the space of potential opportunity types the coalition may encounter. The opportunity types' distribution in the marketplace, is denoted by the probability function $p(\overrightarrow{o})$.

We consider a coalition $A = \{A_1, A_2, ..., A_m\}$ of a general size, where A_j is the $j - th$ buyer agent in the coalition. Each buyer agent, A_j, evaluates different opportunities using a utility function $U_j : O \to R$, where $U_j(\overrightarrow{o})$ is the agent's utility from opportunity type $\overrightarrow{o}$. The search cost associated with having a coalition of m agents maintaining w simultaneous interactions with seller agents over a search round is denoted by the function $c(w, m)$, satisfying $\lim_{w\to\infty} c(w, m) = \infty$ and $\frac{dc(w,m)}{dw} > 0, \frac{dc(w,m)}{dm} > 0$.

The set of opportunities known to the representative agent at a given stage of its search is denoted by θ_{known}. It is possible to have several opportunities of the same type in θ_{known}, as similar opportunities can be offered by several seller agents or a single seller agent may offer opportunities of a specific type with a quantity greater than one unit.

5. ANALYSIS

5.1 The search strategy

Let Θ be the collection of all possible sets of opportunities. Consider a function $alloc : \Theta \to O^m$ that maps a given set of opportunities θ to the coalition members in A (i.e. an allocation)[3] in a way that the aggregated agents' utility with such allocation is maximized[4]. Let $alloc(\theta) = (\overrightarrow{y_1}, ...\overrightarrow{y_m})$, $\overrightarrow{y_i} \in (\theta \cup \{\emptyset\})$ be the allocation, resulting from operating the function $alloc$ over the set θ, where $\overrightarrow{y_i}$ denotes the opportunity associated with agent A_i and $y_i = \emptyset$ denotes that no opportunity was allocated to agent A_i. The immediate utility for the coalition if it terminates the search at the

[2]The reasoning for correlating the coalition's search cost with the number of members being represented is mainly associated with some coordination overhead. See [16] for details.

[3]In B2C markets the same opportunity may be allocated to more than one agent, while in C2C markets each opportunity is restricted to only one agent.

[4]If there is more than one allocation that maximizes the overall coalition utility then the function $alloc$ will always choose one of them according to a pre-defined ordering.

current point given a set of known opportunities θ_{known}, is denoted as $U^s(\theta_{known})$ and can be calculated as:

$$U^s(\theta_{known}) = \sum_{j=1}^{m} U_j(\overrightarrow{y_j}) \quad (1)$$

where $U_j(\emptyset) = 0, \forall j$.

Notice that up until this point, the world states space on which the agent needed to define its strategy (reflected by the set Θ) was potentially infinite. Nevertheless, since there are only m agents represented by the coalition, any additional occurrence of an opportunity of type o_i in θ_{known} beyond the first m occurrences can neither improve the overall coalition utility nor the representative agent's strategy. As a result, the representative agent can reduce the set of known opportunities according to which it determines its strategy to a subset, s, of θ_{known} where each opportunity of a given type appears at most m times. We refer to s as a $state$, and define S as the set of all potential states. Given the $state$ definition, we define a strategy as a function $x : S \rightarrow N$, where $x(s) = 0$ if the agent decides to terminate its search; otherwise $x(s)$ is the number of parallel searches the representative agent should maintain next, when in state s. We denote the optimal strategy by x^*.

We define $V(s, w)$ as the expected utility when using w parallel searches when in state s (assuming any search decision taken at a future state $s' \neq s$ will make use of the optimal number of parallel searches). The term $V(s, 0)$ denotes the immediate utility obtained, if the agent decides to terminate the search at state s, thus: $V(s, 0) = U^s(s)$. The value w ($w \in N$, $w \geq 0$) that maximizes the coalition expected utility $V(s, w)$, is denoted as $x^*(s)$:

$$x^*(s) = \arg\max_{w} V(s, w) \quad (2)$$

In order to formulate the appropriate equation for $V(s, w)$ (from which $x^*(s)$ can be derived) we make use of several additional notations and definitions. Consider a search round in which the representative agent interacts simultaneously with w seller agents, yielding a set $\theta_w = \{\overrightarrow{o_1}, ..., \overrightarrow{o_w}\}, \vec{o_i} \in O$ of opportunities. Let Θ_w be the collection of all w-sized sets of opportunities that can be produced in the environment the representative agent is operating. We denote by $p_w(\theta_w)$ the probability of encountering a specific set of opportunities θ_w, when maintaining w random interactions with seller agents.

We define $next\_s(s, \theta_w)$ as a function that returns the new state s' of a representative agent, upon encountering the set θ_w given its initial state s.[5] For a given number of simultaneous interactions, w, and a given state s, let Θ^s_w be the collection of all w-sized sets of opportunities, θ_w, that change the agent's current state (formally stated as: $\Theta^s_w = \{\theta_w | \theta_w \in \Theta_w \ and \ next\_s(s, \theta_w) \neq s\}$). We denote the complementary set of Θ^s_w by $\bar{\Theta}^s_w$ (the set that includes all w-sized sets of opportunities θ_w that does not change the agent's current state). The expected utility when using w parallel searches while in state s, $V(s, w)$, can now be expressed as ($\forall w > 0$):

$$V(s, w) = -c(w, m) + \sum_{\theta_w \in \Theta^s_w} p_w(\theta_w) V(s', x^*(s')) + \sum_{\theta_w \in \bar{\Theta}^s_w} p_w(\theta_w) V(s, w) \quad (3)$$

where $s' = next\_s(s, \theta_w)$. This is derived from the stationary nature of the problem - if no better state was reached, the search resumes using the same strategy, yielding the same expected utility. After some simple math manipulations of the above term we obtain:

$$V(s, w) = \frac{-c(w, m) + \sum_{\theta_w \in \Theta^s_w} p_w(\theta_w) V(s', x^*(s'))}{\sum_{\theta_w \in \Theta^s_w} p_w(\theta_w)} \quad (4)$$

[5]The set returned by $next\_s$ is the union of opportunities in s and θ_w, where each opportunity type appears at most m times.

Notice that in the case where no new better state can be reached, the denominator becomes zero, and $V(s, w) = -\infty$. This is quite straightforward as the representative agent infinitely maintains a costly search. Here, the coalition's optimal strategy is inevitably to terminate the search. This important characteristic will be used later for designing the proposed solution algorithms.

At this point, one may attempt to compute the coalition's strategy x^* by solving a set of equations of types 1, 2 and 4. Nevertheless, this solution approach is accompanied by many inherent complexities, derived from the structure of the equations. First, notice that equation 4 is a recursive equation and one needs to know the optimal strategy taken in future states s' when extracting the optimal strategy of a given state s. Second, the computation of $V(s, w)$ in equation (4) is exponential in the number of parallel searches, w, used (affecting the number of sets in Θ^s_w, both in the denominator and numerator). Last, according to the above formulation, the potential number of parallel searches that may be used is not bounded, thus reaching a local maximum does not guarantee that a higher utility cannot be obtained.

In the next subsection we present a comprehensive analysis of the problem, emphasizing some unique characteristics of the representative agent's optimal strategy. These findings lead to an algorithm with a polynomial complexity (in the number of potential interactions, w) for computing $V(s, w)$ (which is the key component for computing $x(s)$).

5.2 Analysis

We begin our analysis by facilitating the state definition in a way that allows us to consider simple divisions of the search space into improving and non-improving areas with regard to the coalition's utility. For this purpose we introduce the concept of equivalence between different sets of opportunities within the context of cooperative search. We say that two sets of opportunities $\theta', \theta'' \in \Theta$ are equivalent sets $\theta' \equiv \theta''$, if the following hold: (a) $U^s(\theta'') = U^s(\theta')$; and (b) $U^s(\theta' \cup \theta) = U^s(\theta'' \cup \theta)$ for any set of opportunities $\theta \in \Theta$ that the agent may encounter in the future. For any two equivalent sets θ', θ'', the representative agent is indifferent to knowing the opportunities in θ' and the opportunities in θ''. This is because any set of opportunities the representative agent will encounter in the future results in a similar utility thus the overall coalition's utility will be the same in both cases. Moreover, since the agent's decisions are merely determined by the overall coalition utility and in both cases similar utilities are reached with similar probabilities, the agent will use the same number of parallel interactions for both sets.

Notice that according to the definition above $equivalent$ is a transitive relation ($\theta' \equiv \theta'', \theta'' \equiv \theta''' \rightarrow \theta' \equiv \theta'''$). Moreover, $\theta' \equiv \theta''$ implies that $(\theta' \cup \theta) \equiv (\theta'' \cup \theta), \forall \theta \in \Theta$. Given an allocation $\ell = (\overrightarrow{y_1}, ..., \overrightarrow{y_m})$ of a set θ, we denote the set of opportunities that appear in ℓ by $\{\ell\}$.

THEOREM 1. *Any set of opportunities θ is equivalent to the set of opportunities returned by the function $alloc(\theta)$. Formally stated: $\theta \equiv \{alloc(\theta)\}$.*

The above theorem is certainly non-trivial, particularly when considering the C2C market, where the optimal allocation is permutation based.

An immediate result obtained from theorem 1 above is that the representative agent's strategy is affected only by the subset of θ_{known} defined by $\{alloc(\theta_{known})\}$. Therefore we can now redefine $state$ to be the set of opportunities that are members in $alloc(\theta_{known})$. This latter definition significantly simplifies our analysis and enables us to suggest an efficient algorithm for extracting x^*. It is notable that the computation method used by the function $alloc$ is market-dependent. While in B2C markets the function assigns each agent with the opportunity that maximizes its utility, $\vec{y_j}^* = \arg\max_{\vec{y}\in\theta} U_j(\vec{y}), j = 1, ..., m$, in C2C markets $alloc$ can be computed by solving a maximum weighted matching in a bipartite graph[6].

Given the analysis above, the set of all possible states, S, is given by $S = \{s|\exists\theta \in \Theta \; where \; s = \{alloc(\theta)\}\}$. Also, we can now denote by $S_s = (s_1, ...s_{|S|})$ the states constituting S ordered by their utilities, where s_1 is the state with the highest utility in S. Notice that now the function $next\_s(s, \theta_w)$ simply returns the set of opportunities in $alloc(s \cup \theta_w)$.

The above changes in the definition of a $state$ does not affect the definition of Θ^s_w and $\bar{\Theta}^s_w$. We denote by $p^{stay}(s, w)$ the probability the agent stays at the same state s after conducting w parallel searches. This can be calculated as the probability of having none of the encountered w opportunities change the representative agent's state:

$$p^{stay}(s, w) = (p^{stay}(s, 1))^w = (\sum_{\{\vec{o}\}\in\bar{\Theta}^s_1} p(o))^w \tag{5}$$

The term $1-p^{stay}(s, w)$ can now be used as a better structured representation of the element $\sum_{\theta_w\in\Theta^s_w} p_w(\theta_w)$ that appears in equation (4). Similarly, we may consider the use of $V^{new}(s, k)$, defined as the coalition's expected utility obtained by potentially reaching new states (e.g. different than s) after maintaining k parallel interactions, while using the optimal strategy $x^*(s')$ for each new state s'. The term $V^{new}(s, k)$ does not take into account the cost associated with the current k interactions, however it does consider the search costs associated with any further search stages, originated in new (better) states. Notice that if the agent stays in state s during all k searches $v^{new}(s, k)$ is equal to zero. The term $V^{new}(s, w)$ is actually a representation of the element $\sum_{\theta_w\in\Theta^s_w} p_w(\theta_w)V(s', x^*(s'))$ in equation (4). Therefore, equation (4) can now be formulated as:

$$V(s, w) = \frac{-c(w, m) + V^{new}(s, w)}{1 - p^{stay}(s, w)} \tag{6}$$

The use of $V^{new}(s, w)$ in the above equation is a primary concept used in the algorithm suggested in the next section. Nevertheless, in order to extract $x^*(s)$ it is essential to supply the representative agent with an upper bound, w^s_{max}, for the optimal number of parallel searches to be used when in state s. The following proposition suggests an efficient bound for $x^*(s)$.

PROPOSITION 5.1. *For each state, s_i, an upper bound, $w^{s_i}_{max}$, to $x^*(s_i)$ can be calculated using $w^{s_i}_{max} = \lceil w \rceil$, where w is the solution of the following equation:*

$$c(w, m) = (U^s(s_1) - U^s(s_i))(1 - p^{stay}(s_i, w)) \tag{7}$$

The suggested bound is valid simply because for every value of w greater than $w^{s_i}_{max}$ the search cost associated with the following immediate search round is greater than any possible future improvement in the coalition's utility. Later on, we show that the above upper bound value is a by product of the main loop in our proposed algorithm, thus it does not even need to be directly calculated.

[6]For a set of opportunities $\theta \in \Theta$ found in the C2C market we construct a graph $G_\theta = (V_1, V_2, E)$, where each vertex of V_1 corresponds to an agent in A and each vertex in $V2$ corresponds to an opportunity $\vec{o} \in \theta$. An edge is connected between each agent A_j in V_1 and each opportunity $\vec{o}$ in V_2. The weight of such an edge is the utility for agent A_j from opportunity $\vec{o}$, $U_j(\vec{o})$. Here $alloc(\theta) = (\vec{y_1}, .., \vec{y_m})$, where $\{(A_1, \vec{y_1}), .., (A_m, \vec{y_m})\}$ is a maximum weighted matching in G_θ.

5.3 Algorithmic Approach

Recall that when attempting to solve the problem as a set of equations (see section 5.1) the potential number of parallel searches that may be used is unbounded. Furthermore even if we do manage to come up with a bound for $x^*(s)$ then the calculation of $V(s, w)$ (from which $x^*(s)$ can be derived) is exponential in the number of parallel interactions used, w. Our analysis, which is based on the restructuring of the different elements composing $V(s, w)$, allows us to bypass these two main complexities through the introduction of a finite algorithm with a polynomial computational complexity in w that will necessarily identify the optimal strategy for the representative agent.

In order to efficiently compute $V^{new}(s, w)$ in equation (4) we consider the w simultaneous interactions as w sequential interactions, associated with no search cost. This fully complies with the definition of $V^{new}(s, w)$ as given above (as the cost of the w interactions is already considered). The justification for the above representation method is given in the following lemma 1 which follows directly from theorem 1.

LEMMA 1. *A new state reached by obtaining a new set of opportunities is equivalent to a state reached by sequentially obtaining pairwise disjoint subsets of this set. Formally stated, given a set θ_w and any number of subsets $\theta^1_{w_1}, ..., \theta^r_{w_r}$, $\theta^i_{w_i} \subseteq \theta_w$, $\theta^i_{w_i} \cap \theta^j_{w_j} = \emptyset, i \neq j$, $\theta^1_{w_1} \cup ... \cup \theta^r_{w_r} = \theta_w$ and an initial state s, then the following holds:*

$$ns(s, \theta_w) \equiv ns(ns(...ns(ns(s, \theta^1_{w_1}), \theta^2_{w_2})..., \theta^{r-1}_{w_{r-1}}), \theta^r_{w_r}) \tag{8}$$

where $ns(s, \theta^i_{w_i}) = next\_state(s, \theta^i_{w_i})$. The proof is given in the full paper.

A specific case of the above lemma 1 is where each subset consists of a single opportunity. Thus the calculation of $V^{new}(s, k)$ can recursively rely on the values of $V^{new}(s', k-1)$ where s' represents any of the new states reached after obtaining one additional opportunity from O. Therefore in order to compute $V^{new}(s, k)$, we merely consider the expected utility after conducting one interaction from the planned k interactions. Here, with a probability of $p^{stay}(s, 1)$ the agent remains in the same state s, where the expected utility (having $k-1$ more interactions to go) is $V^{new}(s, k-1)$ (Recall that $V^{new}(s, 0) = 0$ according to the definition of this function that was given in the previous section). Otherwise, if a new state s' is reached after the single interaction, then there is the possibility of reaching further new states in the remaining $k-1$ interactions (taking these states' utility into consideration by the term $V^{new}(s', k-1)$) or with a probability of $p^{stay}(s', k-1)$ the agent will remain in the new state s' even after the additional $k-1$ interactions (in which case the utility is the one obtained by resuming the search from this state using the optimal strategy, $V(s', x^*(s'))$). The above description is encapsulated in the following recursive equation:

$$V^{new}(s, k) = p^{stay}(s, 1)V^{new}(s, k-1) + \tag{9}$$
$$+ \sum_{\{\vec{o}\}\in\Theta^s_1} p(\vec{o})(V^{new}(s', k-1) + p^{stay}(s', k-1)V(s', x^*(s')))$$

where $s' = next\_state(s, \overrightarrow{o})$. Notice that when repeating the calculation using the above equation with increasing k value, starting from $k = 1$, each iteration includes only a single unknown parameter, $(V^{new}(s, k))$.

The above analysis leads to the following algorithm for computing the agent's optimal strategy.

ALGORITHM 5.1. *An algorithm for computing the optimal search strategy x^*. **Input:** O - set of potential opportunity types in the*

market; $p(\overrightarrow{o})$ *- opportunity types' probability function;* m *- coalition's size;* $U_j(),\ j = 1, ..., m$ *- coalition members' utility functions;* $c(w, m)$ *- search cost function;*
Output: $x^*(s)\ \forall s \in S$ *- the representative agent's optimal strategy.*

1. *Build the set of ordered states* S_s*;*
2. *For i=1 to* $|S|$ $\{$
3. *Set* $V(s_i, 0) = U^s(s_i)$ *using equation 1*
4. *Set* $w = 1$*;*
5. *While* $c(w, m) \leq (U^s(s_1) - U^s(s_{|S|}))$ $\{$
6. *Compute* $V^{new}(s_i, w)$ *using equation (9).*
7. *Compute* $V(s_i, w)$ *using equation (6); w++;* $\}$
8. *Set* $x^*(s_i) = \arg\max_{w' \in (0,..w)} V(s_i, w')$ $\}$
9. *Return(*$x^*(s_i),\ i = 1, ..., s_{|S|}$*)*

Notice that at each stage of its execution, algorithm 5.1 reuses components computed in earlier stages. For example, $V^{new}(s, w)$ appears both in the computation of $V(s, w)$ (equation (6)), in the computation of $V(s', w + 1)$ (equation (9)), where $\exists \vec{o}' \in O$ such that $next\_s(s', \vec{o}') = s$ and in the computation $V(s'', w+1)$ equation (9), where $\exists \vec{o}'' \in O$ such that $next\_s(s'', \vec{o}'') = s$. Storing in memory the result for each such computational element for the purpose of reusing it at later stages significantly improves the efficiency of the algorithm. This is accomplished by using two matrixes V and V^{new} of size $|S| \times (w_{max}^{s_{|S|}} + 1)$, where the corresponding $V(s, w)$ and $V^{new}(s, w)$ values are stored for each pair (s, w), representing a state and the correlated result for each number of simultaneous interactions used for calculation. Additionally, we store x^* values for reusing $x^*(s')$ in the computation of $v^{new}(s, w)$.

THEOREM 2. *Algorithm 5.1 returns the optimal strategy of the representative agent in a polynomial time of* $w_{max}^{s_{|S|}}$*.*

PROOF. In order to prove that the algorithm is polynomial in $w_{max}^{s_{|S|}}$, it is sufficient to prove that for each state s the computation of steps 3-8 are polynomial in $w_{max}^{s_{|S|}}$. Notice that the first elements being calculated are for state s_1 (i.e. the one with the maximum utility) according to the loop in step 2. Here, as explained in section 5.1, the expected utility from any strategy in which the search is resumed (i.e. using $w \geq 1$) is $V(s, w) = -\infty$ (formally, since $\Theta_w^{s_1} = \emptyset\ \forall w$ then $p^{stay}(s_1, w) = 1$ and $V^{new}(s_1, w) = ... = V^{new}(s_1, 1) = 0$, thus $V(s_1, w) = -\infty\ \forall w \in (1, .., w_{max}^{s_{|S|}})$). Thus the optimal strategy in state s_1 is to terminate the search, i.e. $x^*(s_1) = 0$ and $V(s_1, 0) = U^s(s_1)$. For any other state, s, when reaching step 6 of the algorithm, the agent has already computed both $V(s', x^*(s'))$ and $V^{new}(s', w)\ \forall w \in (0, ..w_{max}^{s_{|S|}})$ for all potential future new state s' of the state s. This is due to the fact that all future states of a state s appear before s in the set of ordered states S_s (either because having a higher utility, or equal utility yet sorted before s according to the function $alloc$). In addition, for any number of parallel interactions $w \geq 1$ the agent has already computed $V(s, w - 1)$. Therefore the computation time in step 6 does not depend on w. Then, when reaching step 7 of the algorithm, the numerator computation takes a constant time and the computation time of the denominator is polynomial in w. In step 8 the agent chooses the maximum value between $w_{max}^{s_{|S|}}$ values that have already been computed, therefore it is polynomial in $w_{max}^{s_{|S|}}$. The algorithm uses $w_{max}^{s_{|S|}}$ as an upper bound for the optimal number of interactions $x^*(s_i), \forall i = 1, ..., |S|$ (Notice that for $s_{|S|}$ the term $(1 - p^{stay}(s, w))$ in equation 7 is 1). This bound is valid since according to equation 7 the value of $w_{max}^{s_i}$ decreases as a function of i. □

Further significant improvement of the above algorithm's performance can be achieved by calculating and using the specific upper bound, $w_{max}^{s_i}$, correlated with each state, s_i in step 5, according to equation 7. This might require at some points re-computation of $V(s_i, w)$ for w values that were not used in former algorithm execution stages; however the total number of calculations for each state s_i will significantly decrease[7].

Before we demonstrate some important properties of our search model, we wish to emphasize that the cooperative parallel search is a generalization of both the single agent parallel search and the cooperative sequential search models[8]. it is clear that the algorithm suggested in section 5.3 results in the same strategy as in the cooperative sequential search and in the single agent's parallel search, for the specific cases in which $\forall s_i\ w_{max}^{s_i} = 1$ parallel interactions or $m = 1$ agents, respectively.

Finally we wish to note the case where the agents are fully homogeneous (in terms of their utility functions) and operate in B2C markets. Here, we can prove that the optimal strategy of the representative agent is stationary (i.e. does not change according to the current state). Furthermore, we show that the stationary strategy characteristic holds not only for fully homogeneous agents but also when the agents have correlated preferences. We say that two agents A_i and A_j have correlated preferences if agent A_i prefers $\vec{o}'$ over $\vec{o}''$ if and only if agent A_j prefers opportunity $\vec{o}'$ over $\vec{o}''$, (i.e $\forall \vec{o}', \vec{o}'' \in O\ U_j(\vec{o}') \leq U_j(\vec{o}'') \leftrightarrow U_i(\vec{o}') \leq U_i(\vec{o}'')$).

THEOREM 3. *In B2C market if all agents have correlated preferences, the search strategy is based on a reservation value*[9] U_{rv}*. In such scenario the number of parallel interactions the representative agent uses according to the optimal strategy (in case it resumes the search) is fixed during the search (*$\forall s \in S$ *such that* $U^s(s) < U_{rv}$ *exists* $x^*(s) = w_{fixed}$*).*

PROOF. In the above scenario, the search can be represented as the search of a single agent with a utility function which equals the sum of the different agents' utilities, $U = U_1 + U_2 + ... + U_m$. In such case after terminating the search, each of the coalition members will always be assigned with the same opportunity. Therefore, the search strategy is reservation value based, and the search will be terminated upon reaching an opportunity with a utility exceeding some reservation value U_{rv}. Since the probability of reaching such opportunity does not depend on the agent's state, the number of parallel interactions used throughout the search is fixed. □

6. ILLUSTRATIVE EXAMPLES

Having an efficient means for calculating the representative agent's optimal strategy when using parallel search, we can now demonstrate some specific properties of this search method. As a reference we use the single agent's parallel search and the cooperative sequential search models introduced in [2, 6, 13] and [16], respectively.

When taking the cooperative sequential search as a baseline, in many cases the use of the cooperative parallel search results in a significantly improved joint expected utility for the agents constituting the coalition (in the worst case scenario, the optimal cooperative search strategy will yield equal utility). When considering

[7]The extent of the achieved improvement is highly correlated with the specific environment in which the representative agent is operating.

[8]Notice that in this context the single agent sequential search is a specific case of the single agent parallel search, where the agent interacts with a single seller agent at a time.

[9]See [16] for definition and discussion concerning reservation value based strategies in cooperative search.

single agents' parallel search, the cooperative parallel search has the potential to produce a significantly better expected utility (in comparison to the aggregated utility of the single searches) however, the decision of which of the two methods to use depends on the amount of coalition overhead costs induced by the cooperative search.

Figure 1 illustrates the above, depicting the expected utility per agent in any of the search methods[10] in the C2C market (left hand-side) and the B2C market (right hand-side) as a parameter of the similarity level (α) between the utility functions of the agents constituting the coalition. The results are based on the following environment, which was used originally for evaluating the performance of the cooperative sequential search [16]:

ENVIRONMENT 1. *A coalition of two agents, A_1 and A_2, searching for opportunities defined by two attributes, B_1 and B_2, where each attribute can have a value from the discrete range $(1, ..., 5)$ with an equal probability for each of the values. The agents are heterogeneous in respect to the way they evaluate each potential opportunity: agent A_1 is associated with the utility function $U_1(\overrightarrow{o}) = B_1 + B_2$ while agent A_2 with the utility function $U_2(\vec{x}) = 2(1-\alpha)B_1 + 2\alpha B_2$. Thus the parameter α indicates the level of agents' similarity/heterogeneity. The search cost of any single agent was taken to be $c(w, 1) = 0.15 + 0.05w$ and for a coalition $c(w, m) = c(w, 1)ln(m + 1)$, $\forall(m > 1)$.*

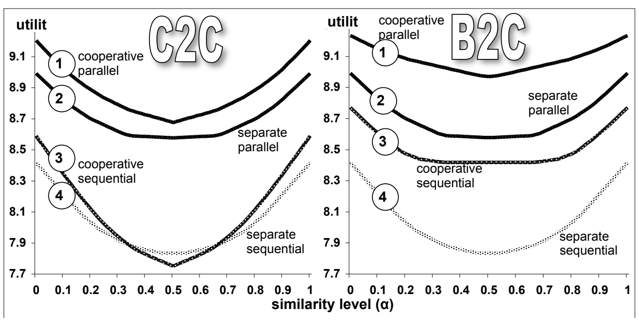

Figure 1: Average utility per buyer agent for different models in different markets

Curve 1 in each graph depicts the average expected utility when the two agents form a coalition, making use of the suggested parallel cooperative search. As expected our model outperforms the cooperative sequential search model (represented by curve 3) in terms of the expected utility for the agents. The other two curves describe the average expected utility of the agents when each is searching separately using the parallel (represented by curve 2) and the sequential (represented by curve 4) methods. In this specific environment the use of the cooperative parallel search also outperforms the single agent parallel search model though it is not always the case. Notice that the results obtained for the cooperative parallel search are consistent with a general characteristic of cooperative search [16] in a way that the usage of the method in the B2C market results with a better expected utility than in the C2C market. In the case of the separated single searches the market type does not affect the strategy structure nor the expected utility.

Figure 1 also reflects an interesting insight which contradicts an important strategy domination relationship found between single and cooperative sequential search techniques of fully homogeneous agents (i.e. with the same utility function, as in the case where $\alpha = 0.5$ in our environment). While for the sequential search the use of single agents' searches always outperforms cooperative search in C2C markets (when considering fully homogeneous agents) [16], here we have a true evidence that when parallel search is concerned, the cooperative search technique may outperform the aggregated result of the single homogeneous agents' search.

[10]For the cooperative models the average expected utility per coalition member measure is used.

Next we introduce and make use of a simpler sample environment for demonstrating some additional properties of the cooperative parallel search.

ENVIRONMENT 2. *Similar to environment 1 above, except for the following changes: (1) each attribute has only two possible values $(1, 2)$ with an equal probability of $\frac{1}{2}$; (2) The utility functions used are $U_1(\overrightarrow{o}) = 1.9B_1 + 0.1B_2$ and $U_2(\overrightarrow{o}) = 0.1B_1 + 1.9B_2$. The search cost of any single agent was assumed to be $c(w, 1) = 0.5 + 0.05w$ and for a coalition $c(w, m) = c(w, 1) * ln(m + 1)$, $\forall(m > 1)$.*

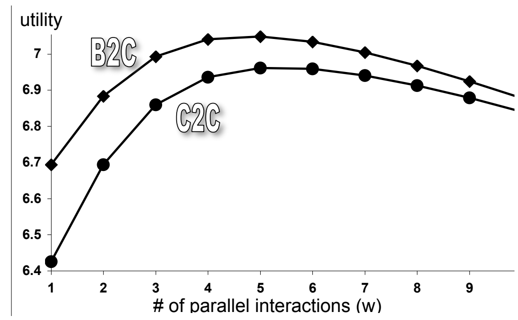

Figure 2: Coalition's overall utility as a function of w

Figure 2 depicts the expected coalition's overall utility with respect to the number of searches conducted at the beginning of the search (i.e. the first search stage, before the agent knows of any of the opportunities), assuming that in all the other states the agent uses the optimal number of parallel interactions, $x^*(s)$. Here, we can see the affect of two conflicting forces: as the number of parallel interactions the agent uses in this stage increases, the probability of associating a better opportunity with any of the two coalition members increases, however the overall search cost associated with the search stage increases. From the figure, we learn that the optimal number of parallel interactions to be used in this stage is $x*(\emptyset) = 5$.

An additional important characteristic of the cooperative parallel search we wish to emphasize concerns the number of parallel interactions used as part of the optimal search strategy along the search. While in a single agent's parallel search the search strategy is stationary (i.e. the number of parallel interactions used does not change along the search process) in our model the number of simultaneous interactions along the search that needs to be maintained depends on the agent's state (i.e. the set of known opportunities known to the representative agent). This is demonstrated in the directed acyclic graph (DAG) given in figure 3, which represents the feasible transitions from one state to other states in a B2C market described by environment 2. The state is determined according to the relevant set of known opportunities, correlated with the definition given in section 5.2. The relatively large number of parallel interactions used in the optimal strategy, allows the representative agent to reach all better states that are defined in this example. Notice that when reaching states $\{(1, 2), (2, 1)\}$ and $\{(2, 2), (2, 2)\}$ the optimal strategy of the agent is to terminate the search. For comparison purposes, notice that in any of the single agents' separate search (i.e. single agent parallel search model) the optimal strategy is to constantly use 4 parallel interactions (as long as resuming search).

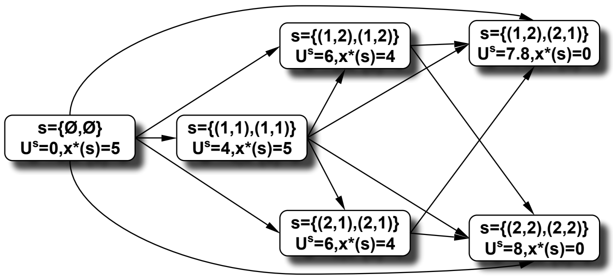

Figure 3: **Potential transitions between states in a simple B2C market**

7. DISCUSSION AND CONCLUSIONS

The capability of using parallel interactions as part of a search process is inherent in the infrastructure of autonomous information agents. When using the cooperative parallel search, the representative agent should use a new strategy, different in its structure in comparison to the optimal strategy used in the cooperative sequential search and in a single agent's parallel search. As expected, the use of the new model has the potential of significantly improving the coalition's expected utility as demonstrated in the previous section. Furthermore, we emphasize that the coalition's expected utility will never decrease when using our proposed mechanism in comparison to pure sequential cooperative search. This is mainly because in the case where maintaining more than a single interaction in some of the world states is not favorable the suggested algorithm will converge to one interaction at a time strategy, as used in the sequential cooperative search (which is a specific case of our model). Obviously if the search cost is linear and depends solely on the number of interactions being maintained, then the sequential cooperative search is the dominating strategy. Nevertheless, a scenario in which the representative agent's search cost combines additional fixed components and/or non-linear dependency on the number of interactions maintained are much more realistic [16]. In those scenarios the parallel cooperative search yields great benefits for searchers.

The novelty of the analysis given in this paper is twofold. First it supplies us with a better understanding of the opportunities space, dividing it into improving and non-improving areas, thus instead of having dual simultaneous dependencies between states we can now define a single directional dependency for each pair of states. Second, it supplies a bound for the optimum number of parallel interactions that the representative agent uses in each state in its optimal strategy. These two features allow us to overcome the main complexities associated with the attempt to solve the problem as a set of equations and offer a finite algorithm with a polynomial (rather than current exponential) computational complexity that inevitably reaches the optimal strategy. The model does not depend on attribute preferences or a specific cost function and with appropriate modification of the allocation function it can be generalized to additional markets other than C2C and B2C.

It is notable that the same considerations discussed in [16] concerning the potential for an actual implementation of cooperative search hold for the parallel (general) model presented in this paper. Similarly, future research in the area of parallel cooperative search should extend the scope of research to include additional topics associated with the coalition formation process such as the coalition stability, given the MAS settings and the division of payoffs (in terms of side payments) between the coalition members.

8. ACKNOWLEDGMENTS

This work was supported in part by NSF #IIS208608 and ISF #8008. In memory of Matat Rosenfeld-Adler.

9. REFERENCES

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

[2] J. Benhabib and C. Bull. Job search: The choice of intensity. *J. of Political Economy*, 91:747–764, 1983.

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

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

[5] M. Dias. *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. PhD thesis, Robotics Institute, CMU, 2004.

[6] S. Gal, M. Landsberger, and B. Levykson. A compound strategy for search in the labor market. *International Economic Review*, 22(3):597–608, 1981.

[7] T. Ito, H. Ochi, and T. Shintani. A group-buy protocol based on coalition formation for agent-mediated e-commerce. *International Journal of Computer and Information Science (IJCIS)*, 3(1):11–20, 2002.

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

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

[10] K. Lermann and O. Shehory. Coalition formation for large scale electronic markets. In *Proceedings of ICMAS'2000*, pages 216–222, Boston, 2000.

[11] S. Lippman and J. McCall. The economics of job search: A survey. *Economic Inquiry*, 14:155–189, 1976.

[12] 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.

[13] P. Morgan. Search and optimal sample size. *Review of Economic Studies*, 50(4):659–675, 1983.

[14] T. Sandholm, K. Larson, M. R. Andersson, O. Shehory, and F. Tohme. Coalition structure generation with worst case guarantees. *Artificial Int. Journal*, 111:209–238, 1999.

[15] D. Sarne and S. Kraus. The search for coalition formation in costly environments. In *Proc. of CIA*, pages 117–136, 2003.

[16] D. Sarne and S. Kraus. Cooperative exploration in the electronic marketplace. In *Proceedings of AAAI 2005*, pages 158–163, 2005.

[17] G. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961.

[18] S. Talukdar, L. Baerentzen, A. Gove, and P. S. de Souza. Asynchronous teams: Cooperation schemes for autonomous agents. *Heuristics*, 4(4):295–321, 1998.

[19] N. Tsvetovat, K. Sycara, Y. Chen, and J. Ying. Customer coalitions in electronic markets. In *Proceedings of AMEC 2000*, pages 121–138, Barcelona, 2000.

[20] J. Yamamoto and K. Sycara. A stable and efficient buyer coalition formation scheme for e-marketplaces. In *Proc. of the Fifth International Conference on Autonomous Agents*, pages 576–583, 2001.