A Platform for Pervasive Combinatorial Trading With Opportunistic Self-Aggregation*

Antonio Di Ferdinando[†], Alberto Rosi[‡], Ricardo Lent[†], Franco Zambonelli[‡], Erol Gelenbe[†]

Abstract

We describe a prototype of trading system platform populated by agents who autonomously decide to buy and/or sell items according to a set of local needs which arise dynamically (also by possibly accessing information provided by pervasive devices) by in the process of fulfilling a given overall utility. The market has combinatorial nature in a way that items to be traded are combined into packages, in accordance with a principle that drives the nature of many current markets. However, differently from these, items belong to a number of distinct sellers distributed in the platform, and are chosen singularly on the basis of buyers preferences and needs. Agents are thus situation-aware, with sellers coming acquainted of the market demand, and buyers price offers, through a Knowledge Network that drives the way the market balances by gathering the needed information in an autonomous way and taking advantage of pervasive devices. Packaging is realized by agent aggregation into Virtual Sellers, in an autonomous fashion, and we propose an opportunistic policy whereby aggregation is governed by a Combinatorial Auction. The market is studied through proof-of-concept simulation, where the efficiency deriving from the opportunistic aggregation based on Combinatorial Auctions and the influence of contextual self-awareness are studied.

1 Introduction

The increasing interest of enterprises for the Internet is changing the way of conducting business, and determining the growth of virtual markets. In these, the distance between producers, distributors and consumers is seen to collapse [12]. Enormous business possibilities are thus created, and it is natural for the parties involved to take full advantage of them. Among these possibilities, particularly advantageous is the one of selling items in packages in a trend called *Combinatorial Trading* (CT). CT is attractive for both seller and buyer, with the first benefitting of the complementarity of items (which acquire a higher value when in a package), and the second maximizing expressiveness by selecting items in a package according to its very own needs. Examples, even though basic, can be found in today's markets where, for instance, holiday offers combine heterogeneous commodities (flight, accommodation, car rental, etc.) into a single package.

As we testify the growth of virtual markets, we foresee them to enter increasingly important business contexts, and their role to become critical in future economies. Thus, future virtual markets might see scenarios where companies offering diverse transport means engage in partnerships dynamically in order to combine their services in an attempt to predict – and meet – with an increasing degree of accuracy users (dynamic) requirements according to an accurate evaluation of their preferences .

This paper builds in the direction of such future scenarios, and presents a prototype system enabling services capable of satisfying the above needs. In the prototype, sellers proactively build packages of items known to be demanded in the market. Bundling is realized by self-aggregating into *Virtual Seller*s (VSs) whose existence is bound to the existence of the package itself. This process is carried out opportunistically upon detection of buyers requests through an analysis of contextual information provided by a *Knowledge Network* [8] (KN), which can taqke advantage of pervasive devices to detect the needed information and which also allows buyers to dynamically point out the need for specific packages. Thus, the KN drives the way the system reacts to buyers needs in the same way as the nervous system drives reactions in the human body.

The result of these interactions is that agents decide when and/or what to buy and/or sell in the system in an autonomous fashion on the basis of opportunistic behaviour. This latter is also exploited in occasions where sellers receive multiple concurrent requests for items, in which case

*The material here presented is based on an ongoing research partially funded by the CASCADAS project (EU FP6-027807), that authors are glad to acknowledge.

†Intelligent Systems and Networks, Imperial College London, London SW7 2BT

‡Dipartimento di Scienze e Metodi dell'Ingegneria, Università degli studi di Modena e Reggio Emilia

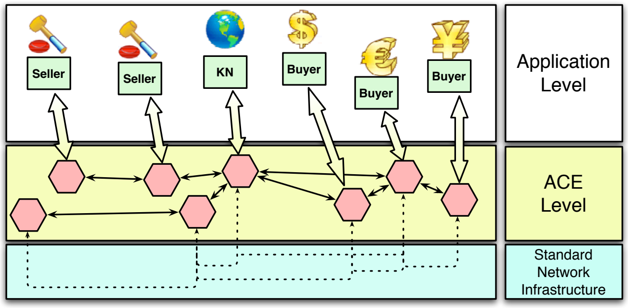

Figure 1. System Architecture.

allocation is decided in a way to enhance economic efficiency. To this end, agents in our system employ a *Combinatorial Auction* [4, 9], and efficiency of this latter method is studied through comparison with the traditional First Come First Served (FCFS) aggregation policy. Our prototype is made pervasive through integration of pervasive devices as constitutive elements for its real-time nature. In fact, our scenario foresees users to interact with the platform, i.e. require and obtain packages, through Bluetooth- GPS- and/or RFID-enabled devices. This information is gathered dynamically and forwarded to the KN through appropriate receptors.

Related work on combinatorial trading is mainly in the area of algorithmic e-commerce [4], while relevant literature in the area of Multi-Agent Systems focuses on topics such as dynamics of coalition formation [12] or techniques for task allocation [7]. In advancing these works, this paper makes three significant contributions: it presents an agent-based pervasive system to fulfill the needs of a future scenario, studies the influence of self-awareness in a CT context, and shows effectiveness of CAs to drive the agents' self-aggregation process.

The paper is structured as follows: section 2 describes the system model, also putting emhpasis on its distinctive characteristics. Section 3 provides an example future scenario where our system would naturally suit. Section 4 validates our ideas by describing simulation results, while section 5 draws some conclusions.

2 System Model

General Architecture. The system is populated with users and services, as described in [6]. Both are agents acting to satisfy a set of needs arising dynamically, over time, in the local context as a consequence of actions taken to fulfill an overall utility provided by agent owners. We consider users to be interested in packages composed of items chosen singularly. On the other hand, services own items which can be aggregated, upon request, into packages in order to be sold. Users are thus buyers in the system, while services are sellers. These roles are symmetric in a way that users of certain transactions can be services in other transactions, and vice-versa. The marketplace populated by a number of sellers and bidders, which trade items whose type is chosen among a finite set. Sellers own items are owned in stocks, which are renewed instantaneously only when all items from the previous stock are sold. The length of packages is finite – but variable – and bound to specific needs of the buyer. In particular, buyers might be interested also in packages of length 1, i.e. single items. In this case, we assume them to form a coalition, for the reasons, and according the techniques, described in [12].

All agents in the platform have been developed as *Autonomic Communication Element*s [1] (ACEs), whose framework allows agents to natively handle *self-** capabilities (such as self-configuration, self-healing, etc.) easing the application-level interaction model from the burden of having to take these issues into account. Referring to this figure, role-specific functionalities are defined in the top-most Application Level, where high-level agent behaviour specification allows to determine the nature of each single ACE. Underneath, the network of ACEs provides the ground for the autonomic application behaviour through sophisticated functionalities (such as for instance, reliable contract-based communication and peer discovery, but also agent mobility and aggregation) which, in turn, rely on standard network infrastructures.

Agents interact, in the system, according to the model described in Figure 2. When the need for an item or package arises at some buyer, as per own utility or explicit request by the owner, the agent queries the KN to check availability (interaction(1a) in Figure 2). If the package is already available, then the agent contacts the seller directly and negotiates the price, according to each party's needs and expectations, starting from the base of the proposed price. Otherwise, the buyer notifies its interest in such a package, along with the price offered, and waits for the bundle to be eventually available. The KN gathers this information and processes these notifications so as to build aggregate views on market-wide desirability trends for packages which can be queried by sellers (interaction (1b) in figure). According to these trends, sellers may decide, on the basis of desirability, to build packages or not. When the desirability is high enough and the seller retains worth building the package, owners of items necessary to the package are individuated and requested to aggregate (interaction (2) in the figure). At destination, such requests are handled opportunistically, where a Combinatorial Auction is employed to enhance local economic efficiency. Aggregation requests remain valid for a specified time, and conflicts on sellers trying to build the same package with the same items are solved with the seller having received less consensus, by other sellers, at a certain time dropping the request. This process leads to for-

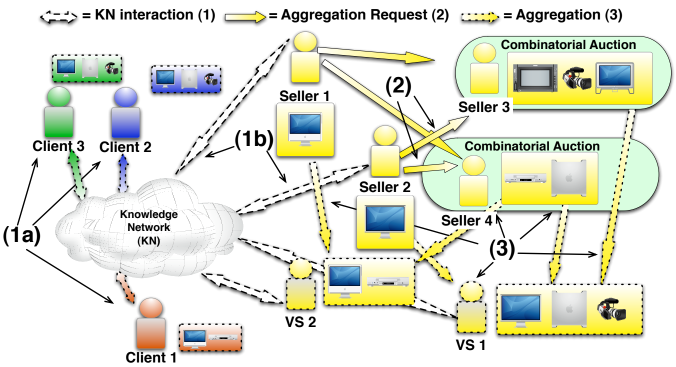

Figure 2. Pervasive Combinatorial Trading System.

mation of a Virtual Seller, leaded by the seller who initially involved others in the aggregation, if and only if all involved sellers accept the aggregation request (i.e. the package can be built). In this case, the VS is formed (interaction (3)) and availability of the new package is notified at the KN which, in turn, notifies relevant buyers. A point-to-point buyer-VS negotiation starts and involved parties may agree a price for the package. When many buyers compete for the package, the negotiation indeed results in a proper auction, aimed to find the buyer who values the package the best, in a way similar to [5]. Packages, and thus VSs, remain available in the market for a limited time. If not sold within this time, they are removed from the list of available ones, the VS vanishes and sellers are freed from any commitment on the VS itself and their corresponding items in the package.

Indeed, the use of ACEs makes our system a prototype of a truly *autonomic system*. Its distinctive characteristic can be found in its efficiency which, in turn, is achieved with a combination of high degree of accuracy with which market demands are matched by offers and exploitation of opportunistic behaviour. The first feature derives by the platform's pervasive features, while the second is obtained by employing combinatorial dynamics in the self-aggregation model. These are the main novelties of our approach, and the remainder of this section provides a more detailed description of these two features.

Knowledge Networks. The term KN refers to an ontology-based structured collection of knowledge atoms describing specific situations. The KN facilitates the pervasive behaviour in the platform and supports the process of agents acquiring high degrees of situation-awareness in an efficient way, with the effect of refining the degree of accuracy with which the market reacts to buyers needs. Figure 3 shows the archtecture of the KN architecture used in our platform [2]. Two layers compose the architecture of the KN used in our system; the bottom-most contains a sensitive amount of unprocessed data, typically gathered through communication with pervasive devices. Such data is then processed, by agents contained in the top-most layer, in a way to highlight relations and recurrences among raw data. Finally, information is given the shape of knowledge by being structured according the so called *W4 tuples* [3] model, that relies on the consideration that most information can be represented in terms of four "coordinates" – Who, What, Where, When – to enable very expressive and flexible context-aware usage of data. The model also relies on the reasonable assumption that software agents are associated with data sources. Thus, they are in charge of creating and inserting W4 tuples in shared data spaces, called *tuple spaces*, where the "Who" refers to the user, and "What", "Where" and "When" are obtained by pervasive technologies (i.e. Bluetooth, GPS or RFID) that equip devices. Once data is fully organized into W4 tuples, aggregated views of the world emerge from querying the KN about the tuple space of one of more terms. Thus, querying with coordinates "when=*from 9pm to 6am*" and "where=*by bus*", the KN would provide an aggregate view of profiles of people traveling on night time, eventually by bus.

Such information, if accurate and properly exploited, is valuable for agents in the platform. Referring to our case, knowledge of such data enables sellers to meet the market demand with obvious repercussions on their local revenue, while also allowing buyers to benefit of packages built purposely for their needs. Information in the KN is not limited to quantities, but also includes average prices for packages. These are gathered by asking buyers and sellers to specify a referring price when notifying the need and availability, respectively, of a package. These figures are then consulted by both sellers and buyers so that a ground for the expected revenue (for sellers) or cost (for buyers) on a transaction involving a package can be estimated.

Combinatorial Dynamics. In situations where bidders outnumber sellers and the set of items in the market is small, it is foreseeable that sellers will receive high numbers of concurrent requests for items. Furthermore, given that the number of items available at a seller is limited, it is equally reasonable to think that many of these offers will overlap in terms of items requested. Therefore, it becomes of paramount importance to be able to govern such decisions in a way to enhance items allocation. As an example, consider the situation in Figure 4: a seller receives three requests for aggregation, at different times, each leading to formation of a different Virtual Seller. This type of occurrences is traditionally handled in a First Come First Served (FCFS) basis, whereby requests are accepted as they come in a best-effort way to satisfy as many requests as it is possible with the items owned. Referring to figure 4, a FCFS basis would imply accepting the request received at time t_1 which would prevent acceptance of the requests received at times t_2 and t_3 that, in turn, would lead to allocation of a

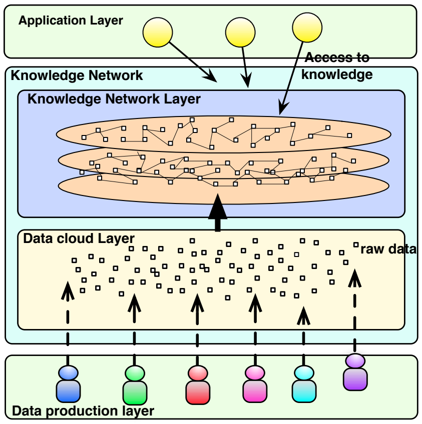

Figure 3. A two layers KN.

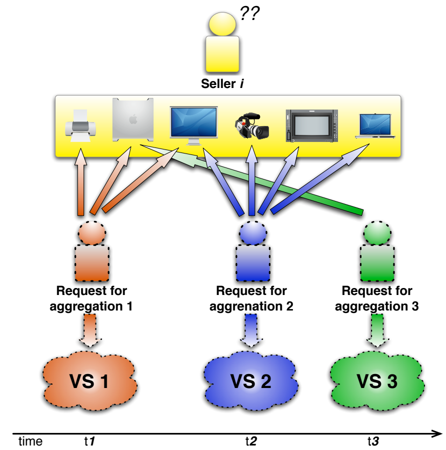

Figure 4. Combinatorial Auction Scenario.

higher number of items.

Combinatorial Auctions seem to provide an ideal ground for solving such type of problems. In the classical scenario, bidders submit offers for a subset of items owned by a seller. This, in turn, receives a number of offers, each on a different subset of items, and needs to decide the combination of offers that enhances economic efficiency. This latter problem is known in relevant literature as the *Winner Determination Problem*[4] (WDP), and is known to be $\mathcal{NP}$-complete [10].Thus, a solution can be found in polynomial time, or approximated, only if the number of items is small at the seller. The WDP poses serious limitations to the applicability of CAs as solution for combinatorial problems in scalable systems (such as Multi-Agent Systems), and for this reason many approaches found in literature make assumptions in the direction of reducing complexity [11].

In our case, we propose a variation whereby items are spread among a set of sellers which are distributed in the platform. This allows our assumption on these having a small amount of items not to loose generality, so that each seller is faced the WDP over a small number of items. As a result, the complexity of the WDP can be reduced, and sellers can decide a nearly optimal allocation. In detail, sellers in our system set a timeout within which they expect to receive requests for aggregation. Upon expiration, requests received are handled and the number of items allocated in each of the possible subsets of offers is calculated. Then, the optimal possible combination of offers is determined as the one that maximizes such number. With respect to the example in Figure 4, application of a CA would result in allocation of items to the requests received at t_2 and t_3, which would altogether allow an efficient allocation of items while uncompromising final revenue.

3 Example: A Future Transport Scenario

It is becoming increasingly common, for transport companies, to keep information about clients and their mobility patterns. This is done (through fidelity cards, online surveys or smart tickets) to the extent of characterizing a variety of attributes such as clients satisfaction levels, mobility trends, preferences and service *pros* and *cons*.

An example of such trend is the London Transport System, whose authority in 2003 decided to equip tickets with *Radio-Frequency Identification* (RFID) tags so as to capture sensible information on the transport system (e.g. most crowded stations, clients mobility and hourly flow patterns, etc.) and focus improvement interventions on relevant areas according to information gathered.

A hypothetical next evolutionary step might build in the direction of broadening the scale of such information and introduce a personal "travel profiler" (in the form of badge or device) that replaces travel tickets altogether while collecting information on personal mobility patterns in a profile. Such type of tool would constitute the terminal for our pervasive combinatorial trading system, which would allow users needs to be accounted as travel prerequisites and be autonomically translated into effective pervasive travel possibilities – built around patterns extracted from the profiler – to be bought instantaneously.

Suppose a number of travel agents and (travel) service providers populate a pervasive combinatorial transport market as sellers. Whenever needed, a user notifies, through the profiler or any device, the need for a specific travel arrangement and complements this information with the own profiler ID. This information is routed to the Knowledge Network, where stigmergic interactions among constitutive elements (as in figure 3) generate knowledge – on the basis of information in the profiler – such as, for instance, preferred leave and arrival time and venues, favorite and worse travel mean and average ideal price etc. This information flows into aggregate views of single and/or group profiles, which are made available to sellers that proactively start building packages accordingly. This determines availability of a number of VSs offering packages needed. Each offer, through VSs, is notified to the device as a possible travel option, and the original user request is finally answered with a set of these.

Once eventually chosen the ideal option, the user can buy it directly through the device with the result of his/her travel profiles being enabled on the corresponding travel means and venues. Finally, as the travel starts, the profiler is used as an electronic ticket and the user's profile is updated so so as to be more accurately processed next time another arrangement is needed. It is worth noting that, although apparently similar, the type of offer described here differ from the ones that can be found in today's markets in a way that the choice of constituting elements is based on user's preference validated by a profile of dynamic generation. This might also be reflected on the price associated to each option, whereas combinations expected to be less attractive might be associated with smaller costs, and would give the user incentives to select them, while still offering service providers possibilities to widen customership to users who would not otherwise consider such options.

4 Simulation Results

System performances have been evaluated, through proof-of-concept simulations, to the extent of testing our ideas and studying the overall efficiency of the market proposed. We put emphasis on two aspects: (i) influence of opportunistic behaviour as brought by the use of Combinatorial Auctions in the self-aggregation process; (ii) financial spin-off in the market as influenced by the accuracy of the information in the Knowledge Network.

Opportunistic Behaviour: opportunistic behaviour was evaluated by simulating a seller, owning a randomly generated number of items, receiving a number of requests for aggregation. A comparison on the number of items allocated with both the FCFS and CA aggregation policies is performed, and results are shown. The simulation model implies the use of a virtual slot of time of fixed length requests fall into. Thus, the FCFS policy implies the seller to accept requests instantaneously in the order they come (given the random nature of requests, this equals to allocating items on a random basis), while the CA policy involves the seller to wait for the end of the slot before deciding the best possible combination of offers as the subset of requests that maximizes the number of items allocated. Item types are uniformly selected within a set of 5, while the number of items owned by the seller is randomly generated in the range $[100, 200]$; bundles have randomly generated length, while the number of requested items in each request is randomly generated in the range $[0, 200]$. In the graphs soon to be described, sampling points contain an average value calculated with 100,000 executions.

Figure 5 shows the rate at which items are allocated, indicated as item allocation rate in the vertical axis, as a function of the rate at which requests are made, indicated in the horizontal axis as allocation request rate. The two lines in the graph denote results of the use of the classical FCFS policy and the one based on CAs. When the number of requests is small, policies are shown to behave similarly. However, as the number of requests increases, the opportunistic optimization introduced by the CAs allows for a higher number of items to be allocated, as compared to the FCFS policy. This trend seems to become further established as the aggregation rate increases. The graph in Figure 6 shows the ratio of successful aggregations, indicated as request success ration in the vertical axis, as a function of the request rate, indicated as allocation request rate in the horizontal axis. Again, both policies are found to perform similarly for small request rates, and thus they are capable of accepting all requests for aggregation. However, the trend highlighted previously seems to be confirmed as when more requests are received, the policy based on CAs outperforms the policy based on FCFS by allowing a higher rate of requests to be accepted.

Pervasive Features: the second set of simulations aimed to study the influence of pervasive features on the platform's economic efficiency in the following way: a market demand is randomly generated through a set of pairs $(package, dl)$, where $package$ is the actual package generated, of random length, and dl is its desirability level, i.e. the number of buyers in need of the corresponding package. Introduction of such information in the KN takes place by associating two values (α, γ) representing the probabilities to store accurate information for $package$ and dl respectively. When relevant information results to be inaccurate, as per such probabilities, it is substituted with a newly generated data before being included in the KN. Sellers are then assumed to build packages as per information in the KN, and finally the ratio $(sold/produced)$ is calculated and observed assuming packages being equally priced. Figure 7 shows results of this simulation. The ratio above is expressed as a function of α, and lines in the graph show results for different values of γ. The graph clearly shows that as α and γ tend to 1, i.e. information is more accurate, the ratio increases. More packages are thus sold, the market offer better meets to demand and the economic efficiency of the market improves.

5 Conclusions

We have presented a prototype system for pervasive combinatorial trading populated by agents who act autonomously as bidders and/or sellers to the extent of satisfying the needs of their own local utility. We have described the interaction model, along with its distinctive characteris-

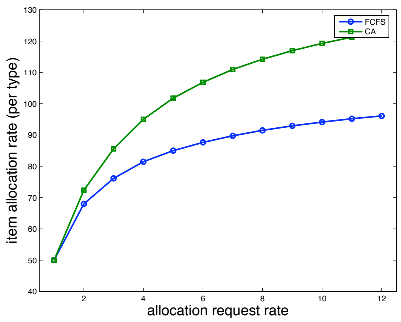

Figure 5. Item type allocation rate.

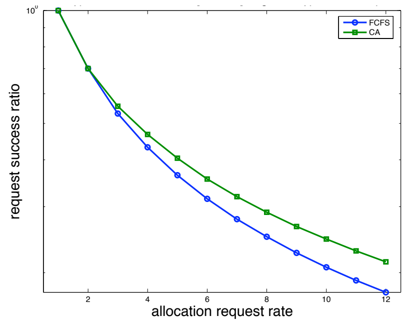

Figure 6. Aggregation success ratio.

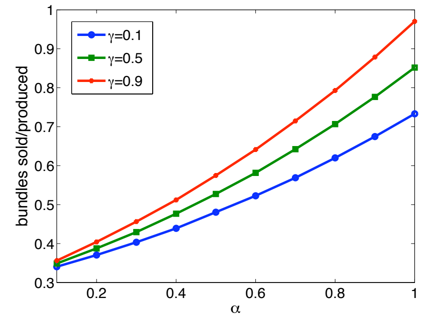

Figure 7. Influence of the KN.

tics of opportunistic behaviour and self-awareness features. The former is exploited through a self-aggregation model based on Combinatorial Auctions, while the latter allows the market demand to be balanced with an appropriate offer. In particular, this is made possible by the use of a Knowledge Network, that gathers information and processes into aggregate flows to usable in the market.

We have estimated efficiency of this market by studying results of simulations on the advantages of opportunistic self-aggregation, through comparison with the classical First Come First Served aggregation policy, and economic spin-off deriving from the accuracy of the KN. Both simulations show positive results, and our future work builds in the direction of finalizing the prototype and take real-world measurements.

References

[1] IP CASCADAS Component-ware for Autonomic Situation-aware Communications, and Dynamically Adaptable Services. http://www.cascadas-project.org.

[2] Castelli G., Mamei M., Zambonelli F. ”Engineering contextual knowledge for autonomic pervasive services”, *Source Information and Software Technology archive*, vol. 50 , No 1-2, 36-50, 2008.

[3] Castelli G., Rosi A., Mamei M. and Zambonelli F. “A Simple Model and Infrastructure for Context-Aware Browsing of the World”. In *Proc. of the 5th IEEE International Conference on Pervasive Computing and Communications* (PERCOM 2007), New York. March 2007.

[4] Cramton P., Shoham Y and Steinberg R (eds). *Combinatorial Auctions*, MIT Press, 2006.

[5] Di Ferdinando A., Lent R. and Gelenbe E. “A Framework for Autonomic Networked Auctions”. In *Proc. of the 1st International Conference on Autonomic Computing and Communication Systems* (AUTONOMICS07), Rome, October 2007.

[6] Gelenbe E. “Users and Services in Intelligent Networks”. In *Lecture Notes in Computer Science*, 3837 (3045), 2005.

[7] Guo W., Nygard K. E. and Kamel A. “Combinatorial Trading Mechanism for Task Allocation.”. In *Proc. of the 14th International Conference on Computer Applications in Industry and Engineering* (CAINE-2001), Las Vegas, Nov. 2001.

[8] Mulvenna M.D., Zambonelli F., Curran K and Nugent C.D. “Knowledge Networks: the Nervous System of an Autonomic Communication Infrastructure”. In *Proc. of the 2nd IFIP TC6 International Workshop on Autonomic Communication* (WAC 2005), Athens, Oct. 2005.

[9] Rassenti S.J., Smith V.L and Bulfin R.L. “A Combinatorial Auction Mechanism for Airport Time Slot Allocation”, *Bell Journal of Economics*, vol. 13, No. 2, 402-417, 1982.

[10] Sandholm T. “An algorithm for winner determination in combinatorial auctions”. In *Proc. of the 16th Joint Conference on Articial Intelligence* (IJCAI-99), Stockholm, Aug.1999.

[11] Shi M.-Y. and Hu S.-H. “Modeling Negotiation in Combinatorial Auctions Based on Multi-agent”. In *Lecture Notes in Computer Science*, vol 4088/2006, 644-649. July 2006.

[12] Ye Y. and Tu Y. “Dynamics of Coalition Formation in Combinatorial Trading”. In *Proc. of the 18th International Joint Conference on Artificial Intelligence* (IJCAI-03), Acapulco, Aug. 2003.