Resource Allocation on Agent Meta-Societies

Alcino Cunha and Orlando Belo

Departamento de Informática, Universidade do Minho
Largo do Paço, 4709 Braga Codex, PORTUGAL
{alcino,obelo}@di.uminho.pt

Abstract. This paper is concerned with the formalization of a automated contracting mechanism that enables a society of cooperative resource allocation agents to negotiate rationally in a self-interested meta-society. Such environments induce agents to adopt different social behaviors according to the negotiation partner. This problem may be solved by taking an economic perspective in all the decisions, namely, by using utility based agents, through the use of marginal utility calculations, and defining dynamically the market extent for a task. The risk attitude and reactivity of each agent can be parameterized in order to achieve different negotiation strategies. The framework presented in this paper can be applied in a wide variety of situations, ranging from electronic commerce on virtual economic markets, to load distribution problems.

1 Introduction

One of the major debates in the multi-agent community is concerned with the social attitude of individual agents [3]. Some researchers argue that agents, being artificial, should behave as cooperative entities whose main objective is to maximize the society's overall goal. Others think that agents should be self-interested, with coordination emerging as a sub-product of their self satisfaction maximization attitude. A similar debate may be found in modern social science between the advocates of the atomistic model of human society and the advocates of the organic model [5]. The former views the society as a group of individuals ruled by the law of maximization of one's rational self-interest, and the latter argues that individuals are products of their society and can only realize themselves inside it.

Such debate could be avoided if we consider any society as a single organism recursively partitioned into several sub-organisms whose behavior is similar to the global organism [4]: from the actions and interactions of a group of societies, a new meta-society emerges which can be viewed as a single entity acting in a coherent manner when reacting to events that occur in its environment. At different levels in this hierarchy of societies, different social attitudes can be found. The study of coalition formation [9][13] provides examples of how selfish agents can eventually coordinate their actions and form an atomistic society that interacts in an also atomistic meta-society.

Resource allocation in a single and homogeneous society is one of the major interest areas of research in multi-agent community [6]. Cooperation mechanisms and protocols for both self-interested and cooperative societies have already been proposed. However, a negotiation mechanism for a meta-society requires that the agents assume, simultaneously, a dual behavior: self-interested when dealing with others agents from external societies and cooperative within their society. This problem has already been addressed, but not solved, in [7]. In this paper we formalize a negotiation mechanism for automated contracting under such circumstances. The presented mechanism works among utility-based agents and is supported by marginal utility calculations in order to guarantee the rational behavior of the society when reacting to its environment. A simpler version was already used in a practical application for load distribution in multi-enterprise environments [2], and extended in a generic framework for resource allocation on multi-agent enterprises [1].

2 A Meta-Society Framework

Each society integrated in a self-interested meta-society is a set of cooperative utility-based agents. A definition of an utility function provides the ability to quantify the monetary compensation that the society is willing to tradeoff for the remote execution of a task. This trading of services and resources for money is typical in negotiation protocols of real economic markets. By using a similar economic approach in the design of the meta-societies, it is possible to enlarge the possibilities of their practical application.

Let $C(t)$ be a function that represents the available cash in an instant t, let $P_i(t)$ be a function that assesses the value of a specific parameter i that influences the level of satisfaction of the society, and let Np denote the number of parameters of the utility function ($C(t)$ is not included in this number), the utility function of a society is given by $U(C(t), P_1(t), ..., P_{Np}(t))$. $C(t)$ must always exist because of the need to quantify the monetary compensations. In the presented equations, we assumed that inflation does not exists. The utility function is replicated in all the agents of the society. The cooperativeness exists because they assess the value of its parameters in the all community and not because all the agents have the same utility function. Let $A_I(P_i(t))$ denote the set of agents where agent I evaluates the level of parameter $P_i(t)$, a set S of agents is a organic society iff

$$\forall I \in S \quad A_I(C(t)) = S \quad \wedge \quad \forall 1 \leq i \leq Np \quad A_I(P_i(t)) = S \tag{1}$$

and is a atomistic society iff

$$\forall I \in S \quad A_I(C(t)) = \{I\} \quad \wedge \quad \forall 1 \leq i \leq Np \quad A_I(P_i(t)) = \{I\}. \tag{2}$$

Similar conditions can be applied to the definition of meta-societies behavior. These conditions only cater for the extreme social attitudes. It is possible to define, at least in theory, other types of society that do not fit in these definitions (such investigation is reserved to future work). In this paper only this two types will be considered. Our meta-society is characterized by condition 2 and its internal societies by condition 1.

Each agent of an organic society is an allocation agent that contains the knowledge related to a set of tasks to be allocated, and has the ability to control a set of resources and services. A task can be executed locally, in a member of the same society, or in an external society when profitable. In the last situation a monetary compensation must be payed. The cooperation process is achieved through the establishment of direct communication channels among agents. When an agent is cooperating within its own society a protocol similar to the Contract Net Protocol (CNP) [11] is used and the negotiation process is task oriented. Each agent can control the answers to its messages through the specification of eligibility conditions and can request remote information in order to update the global society status. When negotiating with external societies, an extension to CNP, similar to the one proposed in [10], is used. The original CNP can not be used in this context because it did not deal, explicitly, with the self-interestness of the agents. The major innovation of this extension is the use of multiple levels of commitment with penalties associated to decommitment.

It is possible to characterize the behavior of the agents according to two parameters: risk attitude and reactivity. The risk attitude is determined by α ($0 < \alpha < 1$). A low value of α stands for a risk-seeking agent, that, for instance, will expect higher returns from contract announces. A high value of α characterizes a risk-averse agent. The reactivity is determined by R, the size of the agent memory. A low value for R will make the agent more reactive because its memory is smaller and only keeps record of the more recent experiences. One can also see α as defining the agent exploratory behavior. In fact, a risk-seeking agent sends announces to larger groups of agents and societies and thus searches a larger task allocation solution space.

The presented automated contracting mechanism is based on marginal utility calculations. It provides rational behavior to an organic society by guaranteeing that all its agents make decisions that maximize the expected value for the marginal utility. To avoid local optima, tasks can be clustered and negotiated together and lower decommitment penalties can be used in order to search a larger task allocation space [10].

3 Automated Contracting Through Negotiation

When an agent receives a new task, either created in its society or as result of a contract with another society, it must decide if it will be executed locally or if the task will be announced for remote execution. In the latter case it is also necessary to decide if the task will be announced only to its society or if it will be negotiated abroad. Additionally, it is necessary to define precisely to which agents and societies the announce of the task will be sent. This decision process depends mainly on three factors: the cost of executing locally the task, the communication costs of sending the announces and the results of past interactions. The former can also be stated as the influence of the task execution in the society global level of satisfaction, or, more precisely, the marginal utility of adding the task to the local plan of the agent.

To handle transparently such kind of decisions, a new mechanism is presented: the dynamic definition of the task market extent. Every time a new task allocation process begins, the allocation agent determines the set of agents and societies to which the contract will be announced. If the determined set is empty then the task is executed locally. If it contains only agents from the same society then the cooperation process will be very simple, being the task assigned without negotiation to the agent which offers lower execution costs. If other societies are present in the set, then an iterative negotiation process will begin, in order to determine if the remote execution of the task is profitable.

The market of a set of tasks T is defined as $M(T)$. This set results from the maximization of the expected return of an announcement and is determined as

$$\max_{M(T) \subseteq Others} \mathcal{E}(\Delta U^{M(T)}(T)), \tag{3}$$

where $Others$ is a set which contains the remaining allocation agents of the announcing agent society and all the external societies known to it; $\mathcal{E}(\Delta U^{M(T)}(T))$ is the expected change in utility if the tasks in T are announced to the set of agents and societies $M(T)$.

Let $\mathcal{E}(\Delta U^{O}_{with\_off}(T))$ be the expected change in utility if the set of agents O make a valid offer and $\Delta U^{M(T)}_{comm}(T)$ be the change in utility caused by the communication costs of sending announces to all the agents in $M(T)$, then for any task set T and for any set of entities $M(T)$, $\mathcal{E}(\Delta U^{M(T)}(T))$ is bounded below by

$$\mathcal{E}^{-}(\Delta U^{M(T)}(T)) = \min_{O \subseteq M(T)} (\mathcal{E}(\Delta U^{O}_{with\_off}(T)) + \Delta U^{M(T)}_{comm}(T)),$$

and above by

$$\mathcal{E}^{+}(\Delta U^{M(T)}(T)) = \max_{O \subseteq M(T)} (\mathcal{E}(\Delta U^{O}_{with\_off}(T)) + \Delta U^{M(T)}_{comm}(T)).$$

Between $\mathcal{E}^{-}(\Delta U^{M(T)}(T))$ and $\mathcal{E}^{+}(\Delta U^{M(T)}(T))$ there is a wide range of values that could be used to estimate $\mathcal{E}(\Delta U^{M(T)}(T))$. Each agent must decide which of these values will be used in its calculations. Since a risk-seeking agent expects a higher return from the announce, the expected change in utility, if T is announced to $M(T)$, can be determined by

$$\mathcal{E}(\Delta U^{M(T)}(T)) = \mathcal{E}^{+}(\Delta U^{M(T)}(T)) - \alpha \times (\mathcal{E}^{+}(\Delta U^{M(T)}(T)) - \mathcal{E}^{-}(\Delta U^{M(T)}(T))).$$

$\mathcal{E}(\Delta U^{O}_{with\_off}(T))$ is determined by

$$\mathcal{E}(\Delta U^{O}_{with\_off}(T)) = \begin{cases} \Delta U_{add}(T) & \text{if } O = \emptyset \\ \max_{Winner \in O} (\mathcal{E}(\Delta U^{Winner}_{ex}(T)) + \Delta U^{Winner}_{send}(T) \\ \quad + \Delta U^{O-\{Winner\}}_{dec}(T)) & \text{otherwise,} \end{cases} \tag{4}$$

where $\mathcal{E}(\Delta U_{ex}^{Winner}(T))$, $\Delta U_{send}^{Winner}(T)$, $\Delta U_{dec}^{O-\{Winner\}}(T)$ and $\Delta U_{add}(T)$ are, respectively, the expected change in utility, if the task is executed by $Winner$, the change in utility caused by the communication costs of sending the task to $Winner$, the change in utility caused by the penalties that must be payed to the remaining agents, whose offers must be rejected, and the marginal utility of adding the tasks in T to the local plan of the announcing agent. When $O = \emptyset$, the task will not be announced but executed locally. This is the reason why the expected change in utility when $O = \emptyset$ is equal to $\Delta U_{add}(T)$. This also guarantees that equation 3 returns an empty set if the announcement is not profitable.

The values of $\Delta U_{comm}^{M(T)}(T)$ and $\Delta U_{send}^{Winner}(T)$ do not require further formalization because they simply accounts for the communication monetary costs, which are very well defined. The determination of $\Delta U_{add}(T)$ is achieved through similar calculations and is partially based on the specification presented in [8]. The values of $\mathcal{E}(\Delta U_{ex}^{Winner}(T))$ and $\Delta U_{dec}^{O-\{Winner\}}(T)$ are estimated based on past values stored in the agent memory and thus are directly influenced by R, the parameter that defines the agent reactivity.

The transparency concerning the social attitude of the agents in $M(T)$ is achieved because, when an agent belongs to the same society, parameters like $\Delta U_{send}(T)$ or $\Delta U_{dec}(T)$ have values very close to zero or even zero. This characteristic guarantees that equation 3 implicitly gives priority to the agents of the same society when building $M(T)$. However, it is important to remember that all the decisions are taken in an economic perspective and try to maximize the expected utility of the consequent action. So, if the execution in an external society is expected to be more profitable, even with all the costs associated to the remote negotiation, that society will be included in $M(T)$.

To complete the specification of the automated contracting strategy it is also necessary to define how agents will assess the global value of the various parameters which contribute to the society utility function. The adopted solution is based on the periodic exchange of the utility function parameters between the allocation agents and was presented in detail in [1].

4 Conclusions

This paper presents the formalization of an automated contracting mechanism that can be used to build a rational cooperative society of utility based agents that interacts coherently in a self-interested meta-society. The use of a negotiation mechanism based on marginal utility calculations and the dynamic definition of the market extent for an announce provide the agents with the ability to manage transparently the dual behavior that they must have in such environments. The presented approach also allows the implementation of different negotiation strategies inside the same society just by changing the parameters that define the risk attitude and reactivity of its agents. Such possibility facilitates the task of the system designer when setting up the system. Since almost every computational problem can be modeled in a resource allocation fashion

[12], the presented framework can be applied in a wide variety of situations, ranging from load balancing problems to distributed factory production plans.

References

1. Alcino Cunha and Orlando Belo. An electronic commerce framework for resource allocation among multi-agent entreprises. In *Proceedings of The 10th International FLAIRS Conference (FLAIRS97)*, pages 362–366, Florida, USA, May 1997.
2. Alcino Cunha and Orlando Belo. A multi-agent based approach for load distribution in multi-enterprise environments. In *IASTED International Conference on Applied Informatics (AI-97)*, Innsbruck, Austria, February 1997.
3. Edmund H. Durfee. What your computer really needs to know, you learned in kindergarten. In *Proceedings of the 10th National Conference on Artificial Intelligence*, pages 858–864, July 1992.
4. Brian R. Gaines. The collective stance in modelling expertise in individuals and organizations. Technical report, Knowledge Science Institute, University of Calgary, Calgary, Alberta, Canada T2N 1N4, August 1995.
5. Robert Hollinger. *Postmodernism and the social sciences: A thematic approach*, volume 4 of *Contemporary social theory*. Sage Publications, Inc., 1994.
6. Victor R. Lesser. Multiagent systems: An emerging subdiscipline of ai. *ACM Computing Surveys*, 27(3):340–342, September 1995.
7. Tuomas Sandholm and Victor R. Lesser. On automated contracting in multi-enterprise manufacturing. In *Improving Manufacturing Performance in a Distributed Enterprise: Advanced Systems and Tools*, pages 33–42, Edimburgh, Scotland, 1995.
8. Tuomas W. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, Washington D.C., 1993.
9. Tuomas W. Sandholm and Victor R. Lesser. Coalition formation among bounded rational agents. In *14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 662–669, Montereal, Canada, 1995.
10. Tuomas W. Sandholm and Victor R. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *First International Conference on Multiagent Systems (ICMAS - 95)*, pages 328–335, San Francisco, 1995.
11. Reid G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, December 1980.
12. Michael P. Wellman. The economic approach to artificial intelligence. *ACM Computing Surveys*, 27(3):340–342, September 1995.
13. Gilad Zlotkin and Jeffrey S. Rosenschein. Coaliation, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *The AAAI 1994 Spring Symposium on Software Agents*, pages 87–94, Stanford, California, March 1994.