An Aggregation-Disaggregation Approach for Automated Negotiation in Multi-Agent Systems

Amal El Fallah Seghrouchni (*), Pavlos Moraïtis (✞,✯) and Alexis Tsoukiàs (✯)

*LIPN-UPRES-A CNRS 7030 – Institut Galilée
99, Av. J.B. Clément – 93400 Villetaneuse, Université Paris13-France
e-mail: elfallah@lipn.univ-paris13.fr

✞ Decision Support Systems Laboratory
Technical University of Crete, 73100 Chania, Crete, Greece
e-mail: moraitis@dias.ergasya.tuc.gr

✯LAMSADE-CNRS, Université Paris-Dauphine
Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16-France
e-mail: {moraitis, Tsoukiàs }@lamsade.dauphine.fr

Abstract: This paper describes an automated negotiation procedure based on an aggregation-disaggregation approach. Negotiation is based on an oriented graph, namely a negotiation graph. In this graph, a node represents a state of the world while an arc represents multiple actions of the involved agents. Agents are equipped with a multi-criteria decision making model. Based on the negotiation graph, an agent can make an offer (a proposal of a path) using his strictly individual preference model (multiple criteria aggregation based decision step), while the same agent can receive the counter-offer of a counterpart. In this paper, such a model takes the form of a multiple-attribute additive value function. On these grounds and using a multiple linear regression model (disaggregation step), agents are able to make a first estimation for the parameters of the preference model of their counterpart (that is the trade-offs and the shape of the value functions). Using such an estimation, agents create an enhanced preference model including the estimated preference model of their counterpart in their own model. Then, they compute a new offer on the basis of the enhanced model. The procedure loops until a consensus is reached, that is all the negotiating agents make the same offer.

1. Introduction

Negotiation has long been recognized as a process of some importance for Multi-Agent area research. The parameters of the Multi-Agent Systems (M.A.S.) consisting of rational agents are becoming ubiquitous. Such agents may be heterogeneous, cooperative or (pure) self-interested. But one of the more important features remains their autonomy to carry out tasks, to make choices and decisions. To fulfil this requirement, it is necessary for agents to be able to develop their proper strategy (i.e. no coordination mechanism can be imposed externally). Consequently, the diversity of strategies may raise conflicts the solving of which requires a negotiation procedure that allows agents to work together and to perform transactions. Usually, negotiation aims to modify the local plans and/or decisions of agents (either cooperative or self-interested) in order to avoid negative (i.e. harmful) interactions and to emphasize the situations where positive (i.e. helpful) interactions are possible.

The design of computational agents needs an automated negotiation the reasons of which may be summarized [Sandholm 99] as follows: 1) Several new applications such as electronic-business are becoming increasingly important and require an operational decision making level; 2) Agents interact in an open environment and may pursue different goals; and 3) The development of virtual enterprises requires negotiation and dynamic alliances. In all this cases, automated negotiation should provide a great help to save labor time in the sense that it may replace, at least partially, human negotiators.

Automated negotiation has long been studied in MAS field and different negotiation mechanisms have been proposed and include: [Chu-Caroll and Carbery 95; Ito and Shintani 97; Klein 91; Jennings et al. 98; Sandholm and Lesser 95; Shehory and Kraus 96; Sycara 89a; Sycara 89b; Zlotkin and Rosenshein 91; Zlotkin and Rosenshein 96]. Generally, the proposed approaches are based on operational research techniques [Kraus 97; Müller 96]. However the multi-criteria dimension of the negotiation process is basically ignored in all such approaches. The main innovation of this paper is to propose an automated negotiation framework based on an aggregation-disaggregation procedure for agents using a multi-criteria decision model. It explores the ideas developed in [Moraïtis and Tsoukiàs 96; El Fallah, Moraïtis and Tsoukiàs 99]. An agent can make an offer using his strictly individual preference model (multiple criteria aggregation based decision step), while the same agent can receive the counter-offer of a counterpart. In this paper, such a model takes the form of a multiple-attribute additive value function. Using a multiple linear regression model (disaggregation step), agents are able to estimate the parameters of the preference model of their counterpart (that is the trade-offs and the shape of the value functions). Based in this estimation, agents generate an enhanced preference model by including the estimated preference model of their counterpart in their own model and compute a new offer on the basis of the enhanced model. The procedure loops until a consensus is reached that is all the negotiating agents make the same offer.

This work discusses multi-agent negotiation in the situations where agents are lead to cooperate in order to achieve a global goal, while simultaneously trying to satisfy as best as possible individual preferences. Our theory, allowing conflict resolution generated by different kind of sources (resource sharing, actions for preferences and/or goal satisfaction, etc.), could be also applied in the case of self-interested agents, considering that goals are independent of the global goal.

The paper is organized as follows. Section 2 presents the multi-criteria problem setting. Section 3 evaluates conventional negotiation model in decision theory and highlights the necessary adaptations for their use in the multi-agent context. Section 4 develops our distributed negotiation procedure. A simple but significant example is introduced in order to make clear our approach. Section 5 discusses and evaluates our approach under some pertinent points of view. Section 6 compares our approach to related work while concluding remarks and future directions are outlined in section 7.

2. The multi-criteria problem setting

In conventional decision theory [Jelassi et al. 90] negotiation is seen as an interactive Multi-Criteria Decision Making procedure [Vanderpooten and Vincke 89] where the exploration of the set of efficient solutions is performed not just by a single decision maker, but by the whole set of participants in the negotiation process. Technically, the procedure is always the same: start with an efficient solution and then move to a next by modifying some parameters as the trade-offs, the zenith and/or nadir points, the shape of an utility function. What changes is the type of interaction, since it does not concern just a decision-maker and an analyst, but several decision-makers. The interaction holds among the different participants each of which may propose a new efficient solution toward a consensus (if ever).

Unlike such a representation, real world negotiation processes do not limit the negotiation to just such issues, but consider more complex objects, such as the set of potential actions, the set of criteria under which actions are evaluated and the negotiation scope itself (possibly). Our claim is that an MAS enabling the participating agents to "negotiate" should not limit itself to the possible negotiation objects, but allow each agent to establish what to negotiate for.

Considering a set **Ag** of agents α_i. An agent, as a dynamic planner [Moraïtis and Tsoukiàs 00], is equipped with the following decision model including his individual preference model $\alpha_i : \langle \mathbf{T_i, A_i, H_i, P_i, G_i, R_i, S_i}, \rangle$ where:

- T_i : a set of tasks to be achieved by the agent (different levels of achievement may be considered);
- A_i : a set of elementary actions available to the agent;
- H_i : a collection of $H_i^j \subseteq A \times A$, binary preference relations on the set A of the type:

$\forall x, y \in A, H_i^j(x, y)$: agent *i*, on dimension *j*, considers action *x* at least as good as action *y*;

- P_i : a set of plans (a sequence of actions) the agent may perform in order to accomplish the tasks;
- G_i : a collection of binary preference relations $G_i^l \subseteq P \times P$, on the set P of the type: $\forall \chi, \psi \in P, G_i^l(\chi, \psi)$: agent *i*, on dimension *l* considers the plan χ at least as good as the plan ψ;
- R_i : an aggregation procedure enabling to establish a global relation H_i and G_i (if it is the case) and to connect the relations H_i^j to the relations G_i^l;
- S_i: a set of states of the world representing the consequences of each elementary action the agent may perform.

Under such a perspective the agent's problem consists in solving a dynamic programming problem, that is to define the "best path" on a graph whose nodes are the states of the world, the arcs are the elementary actions, paths correspond to plans, and H_i and G_i represent the agent's preferences in order to define what is "best".

Moving up an abstraction level, the previous decision model may be extended to a community of agents **Ag** as follows: **Ag:** $\langle \mathbf{T}, \Delta, \mathsf{H}, \mathsf{P}, \Gamma, \Re, \mathcal{S} \rangle$ where:

- T: a set of tasks to be accomplished by the community (different levels of accomplishment may be considered);
- Δ: a set of elementary actions available to the community of agents;
- H: a collection of $H_j \subseteq \Delta\times\Delta$ binary preference relations on the set Δ of the type $\forall$ x, y $\in$ Δ: $H_j(x,y)$: the community, on dimension j, considers action x at least as good as action y;
- P: a set of plans (ordered sets of actions) the community may perform in order to accomplish the tasks belonging to T;
- Γ: is a collection of $G_l \subseteq \mathsf{P}\times\mathsf{P}$ binary preference relations on the set P of the form: $\forall\chi, \psi \in \mathsf{P}$, $G_l(\chi,\psi)$ means that the community, on dimension l, considers plan χ at least as good as plan ψ;
- $\Re$: is an aggregation procedure enabling to establish a global relation H and Γ (if it is the case) and to connect the relations H_j to the relations G_l;
- $\mathcal{S}$: is a set of states of the world representing the consequences of each elementary action the community may perform.

Under a conventional negotiation scheme the only object on which the negotiation may hold are the parameters defining $\Re$. In such a case it is necessary to consider:

$T = \cup_i T_i$ and $\Delta = \cup_i A_i$

It is clear that such a perspective is very reductive with respect to the negotiation requirements of MAS. Moreover the existence of MAS level may enable actions not foreseen on a single agent level and modify the way by which plans are evaluated (i.e. each agent G_l).

Under such a perspective we claim that the negotiation objects in MAS include:

- the establishment of $\Re$ and its parameters, considering T, and Δ fixed;
- the establishment of Γ possibly modifying each agent G_i;

- the establishment of P possibly modifying each agent A_i, T_i and H_i.

This paper presents a procedure concerning the first among the above negotiation objects. In fact, although it concerns the most commonly explored problem, it turns out that the extension of conventional negotiation models in the context of MAS is far than trivial.

3. Conventional Negotiation Models in Decision Theory

Conventional negotiation models [Jelassi et al. 90] imply the existence of a facilitator, that is an agent who tries to identify a compromise among the feasible solutions proposed by the participants to the negotiation process. In fact such models extend in the frame of negotiation well known multi-criteria decision making methods [Steuer 86]. The negotiation procedure in this case holds as follows:

- suppose *n* agents α_i each of them having preferences on a set *A* of alternatives a_j (we make no particular hypothesis on *A* since it could be discrete or continuous, and on the nature of the preferences of the agent);
- each agent communicates to the facilitator his/her preferences;
- the facilitator computes a compromise using a multi-criteria decision making tool (where each agent is considered as a criterion) and submits it to the agents;
- if there is consensus on the compromise proposed, then the negotiation process ends;
- else the facilitator tunes the parameters of the model (asking the agents some further improvement on their preferences) until a consensus is reached.

In order to facilitate readers comprehension and to simplify the presentation, we will place ourselves in a normative frame of perfectly rational agents (although technically this is not a limitation). Under such an hypothesis:

- each agent preferences can be represented by a value function $g_j(x)$, $x \in A$, $g_j: A \rightarrow R$ defining a consequence set X_j image of the function g_j;
- the negotiation space is therefore defined by the set $X_1 x X_2 \ldots x X_n$ (the Cartesian product of all consequence spaces);
- the search of a compromise considers as a starting point an element of the set of efficient (non dominated) solutions and uses as an exploration tool a compromise function (a pseudo-concave value function [Zionts and Wallenius 83] or as a scalarizing function [Wierzbicki 82] or as a Tchebychev distance [Vanderpooten and Vincke 89];
- the parameters which normally have to be tuned in order to find a consensual compromise are:
    - o the ideal and anti-ideal point of the negotiation space;
    - o the trade-offs or importance parameters or scaling constants among the agents;
    - o the shape of the compromise search function.

Such hypotheses immediately call for multi-objective programming procedures where each agent is replaced by an objective. Two observations are possible:

1. Each agent on his/her turn may have a multi-criteria evaluation model by which his/her preferences are elaborated. The simple version of the above procedure will just not consider this second layer since there is no formal link between the way by which an agent elaborates his/her preferences and the way by which these are considered in the negotiation procedure;

2. From an algorithmic point of view, the class of interactive procedures enabling the exploration of a multi-dimensional consequence space can be divided in two categories [Vanderpooten and Vincke 89]:

- o strictly monotone, thus convergent, in the sense that any compromise solution defines a non return point for each consequence dimension [Ozernoy and Gaft 77];
- o learning, thus not necessarily convergent, in the sense that to each step of the algorithm, it is possible to come back to earlier preferences.

It may be worth to notice that intermediate procedures have been proposed in the literature [Vanderpooten 89].

It is easy to observe that several among the previous assumptions do not hold in reality and that could generate different problems in the frame of MAS which has to implement a negotiation procedure. We list among other:

- the existence of a facilitator contradicts the distributed nature of decision capabilities in MAS (in reality also the existence of a facilitator is observed only in very specific negotiation processes);
- agents do learn during a negotiation procedure so that it makes no sense to consider strictly monotone compromise procedures;
- as already mentioned [El Fallah, Moraïtis and Tsoukiàs 99] the negotiation objects cannot be limited to the tuning of the compromise search procedure, but may concern the negotiation model and purpose;
- it cannot be neglected the fact that each agent owns an individual (often multi-dimensional) evaluation model since such a model contains the reasons under which the agent's preferences have been elaborated.

The negotiation procedure proposed here tries to replie (at least partially) to some of the above problems. It is a distributed procedure where:

- No facilitator is necessary;
- Each agent is endowed with specific learning capabilities;
- A two-layers evaluation model is considered for each agent.

4. A New Distributed Negotiation Procedure

Intuitively speaking, the idea is that, during a negotiation process each participant in making an offer (that is making a choice) tries to take into account the preferences of his/her counterpart. However, such preferences are initially unknown and are revealed gradually during the negotiation process through the counter-offers of the counterpart. Offers and counter-offers are based on a collective efficient plans graph, called negotiation graph, representing the possible plans the community may perform (i.e. the set P). Each negotiating agent creates such a graph using his individual efficient plans, and those of his counterpart, which are exchanged between them. The computation of individual efficient plans is based on the set P_i of possible plans each agent may perform and the collection of binary preference relations, the set G_i [for more details, see Moraïtis and Tsoukiàs , 96; 00]. The computation of such a set P may not be just the union of each P_i (the community graph is not necessarily the merging of each agent graph). In fact, some actions in some A_i may disappear, some new actions may enter directly in Δ, the way by which each action is evaluated by each agent (the H_i) may be modified. Then the community has to evaluate P. The final decision is taken at the end of the negotiation (the definition of $\Re$).

4.1 An Example

Let us show how our negotiation procedure works on the following example where conflicts arise among agents' preferences. Let us consider an empty (EM(R)) room (R) which must be restored (RS), (painted and plumbed), and equipped (EQ(R)) with a bookcase (A) full of books (B). RS(R) is the goal of α_1 while EQ(R) is the one of α_2. For RS(R) accomplishment, agent α_1 has to paint and to plumb the room. For EQ(R) accomplishment, agent α_2 has to assemble bookcase, to move it inside of the room (the order of these two actions execution has no importance) and to put the books in the bookcase. We can resume the situation as follows:

- Initial state of the world: (EM(R), takedown(A), OUT(A), OUT(B));
- Final state of the world to be obtained (common goal): (RS(R), EQ(R));
- Agent's goals: RS(R) for α_1, EQ(R) for α_2
- A Possible actions of α_1: to-paint(x, y) : the agent x paints room y; to-plumb(x, y): agent x plumbs the room y; wait (x): agent x waits;
- Possible actions of α_2: move(x, y, z, w), to-put-on(x, y, y'); agent x puts an object y on object y', to-assemble(x, y): the agent x assembles the object y, wait (x);
- Relations between actions: before(to-assemble(x, A), to-put-on(x, B, A)), before(to-paint(x, R), move(x, A, OUT, R));
- Agents' preferences: (max-profit, p), (min-time, t). We assume that actions to-paint(x, y), to-plumb(x, y), to assemble(x, y) leave a profit of 2 units while they generate a loss of 1 time unit, and actions move(x, y, z, w), put-on(x, y, y'), leave a profit of 1 unit while they generate a loss of 1 time unit. Action wait (x) generates a loss of 1 time and 1 profit units.

Agents generate possible and efficient paths using their multi-criteria model. Then they exchange individual efficient plans, which represent the efficient ways for each agent to reach his goal, trying to satisfy his preferences at the best. Hence, each agent creates a collective efficient plans graph [Moraïtis and Tsoukiàs 96] compatible with the community decision model presented in §2, called also "the negotiation graph".

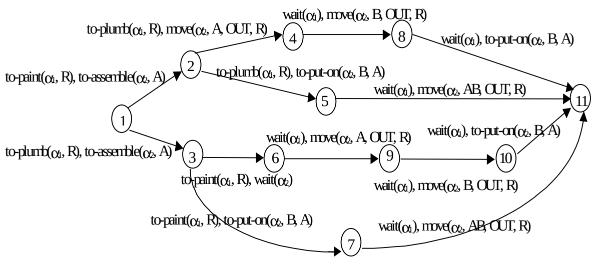

Figure 1: The negotiation graph

Figure 1 represents the negotiation graph which contains all the possible ways to achieve the common goal and it is considered as the starting point of the negotiation procedure.

4.2 The Negotiation Protocol

A simplified procedure representing the negotiation protocol is as follows:

begin

1. Creation of the negotiation graph (the set P) by each agent

2. Evaluation of P

loop

- **2.1** An agent makes an offer using his strictly individual preference model presented in § 2;
- **2.2** The same agent receives the counter-offer of the counterpart;
- **2.3** The agent tries to estimate the counterpart's preference model;
- **2.4** The agent incorporates the counterpart's model in his own model and makes a new offer;
- **2.5** The process goes back to point 2.1 and loops until a consensus is reached;

end loop

end

More formally consider the following problem setting: There exist two agents α_1 and α_2 and a set of potential alternatives A. Elements of A can be different configurations of the same object (for instance a plan is composed from different actions in different sequences). In the context of this paper, it corresponds to the set P. Further on, we consider that the agents share the set of criteria under which they evaluate the set *A* (but not necessarily the evaluations). This is the set $G_1 \cup G_2$. More precisely in our example agents share the same criteria, i.e. time and profit.

Each agent is equipped with an individual preference model (presented in §2) on the set P. For simplicity, we assume that such a model takes the form of a multi-attribute additive value function of the type: $\sum_j p^i_j v^i_j(x)$, where p^i_j are the trade-offs among the different attributes and v^i_j are the value functions associated to each attribute (for agent α_i). We call $U_i(p)=\sum_j p^i_j v^i_j(x), \forall x \in p \wedge p \in P$, the additive utility functions. So, in our case we have for example for the path(1-3-7-11), $U_{\alpha_1}=(0.5*3)_{profit}+(0.5*3)_{time}=3$, and $U_{\alpha_2}=(0.5*4)_{profit}+(0.5*3)_{time}=3,5$ if we consider that the trade-off among profit and time is 0.5 for both agents.

The step 2 of the procedure previously introduced can now be represented as follows:

loop

2.1 agent α_1 (arbitrary) makes an offer x_1^* such that: $x_1^*=\max_{x \in \Delta} \sum_j p^1_j v^1_j(x)$. This offer corresponds to the choice of a path (plan) $P_1 \in P$;

2.2 agent α_2 makes an offer x_2^* such that: $x_2^*=\max_{x \in \Delta} \sum_j p^2_j v^2_j(x)$. This offer corresponds to the choice of another path (plan) $P_2 \in P$;

2.3 knowing agent's α_2 counter-offer, agent α_1 can establish that: $\forall x \in A, \sum_j p^2_j v^2_j(x_2^*) > \sum_j p^2_j v^2_j(x)$;

2.4 on these grounds and using a multiple linear regression model agent α_1 is able to make a first estimation of the parameters of the preference model of agent α_2 (that is the trade-offs and the shape of the value functions);

2.5 using such estimation, agent α_1 creates an enhanced preference model including the estimated preference model of agent α_2 in its own model; (actually agent α_2 estimated value function will become a criterion to add to agent's α_1 preference model)

2.6 agent α_1 goes back to the first step (2.1) and computes a new offer on the basis of the enhanced model, the procedure loops until a consensus is reached (agent α_2 makes the same offer as agent α_1);

end loop

5. Discussion

1. Applicability: How general is the proposed model and how applicable is it? We claim that our approach includes several different modeling possibilities.

1.1. It is not necessary that all agents have a multi-attribute model. If an agent has a mono-dimensional preference model the task of estimating such a model is simplified. The hypothesis about the existence of multi-attribute models is just a generalization.

1.2. The additive value function representation of the preference model is just a technical choice. In fact we could consider any type of preference model (with or without a numerical representation) and any choice procedure (for instance based on majority rules and outranking relations, for details see [Vincke 92]). In such a case different parameters have to be estimated in the learning part of the procedure.

1.3. The multiple linear regression based learning procedure is again a technical choice. Any learning procedure based on genetic algorithms, neural networks, rough sets etc., could apply provided it fits the type of parameters to be estimated. The association of additive value functions and linear regression learning procedures is very effective [Jacquet-Lagreze and Siskos 82], but by no mean should be considered as an unique solution.

1.4. Last, but not least, our approach is not limited to only two agents negotiations, since the procedure can be generalized to include n participants. This is expensive from an efficiency point of view, but the more agents negotiate the longest the process becomes.

Therefore our approach covers a wide range of modeling possibilities at least as far as the preference model and the learning capabilities are concerned. It is clear however, that the choice of an appropriate preference model and associated learning procedure has to be extremely accurate.

2. Distributivity: How distributed is the proposed approach? In fact each agent conserves his(her) own preference model to which may add or modify criteria and parameters. Therefore it remains with a complete autonomous decision capacity. If the other agents are silent (s)he might be able to make a decision and implement it. A critical point concerns the hypothesis concerning the share of the criteria used by all agents. This is a limitation in the present state of our approach.

3. Convergence. Does the procedure guarantees to always achieve a solution? Generally speaking the answer is negative. This is due to the fact that an agent may make ``strange" and apparently "inconsistent" offers during the process, obliging the learning procedure to very bad estimations and therefore to a poor enhancement of the counterpart's preference model.

Theoretically there is no way to get of this situation, unless the liberty of each agent in making an offer is limited by his/her previous offers (but this is a limitation to agent's autonomy). Further on, whatever is the preference model adopted, Arrow's theorem [Arrow 63] always hold constraining the choice procedure to violate one of the axioms (Pareto optimality, independence, non dictatorship) or to not guarantee the identification of a consensual solution. In practice it is possible to endow each agent with a memory of the previous transactions so that the update of its preference model might be easier. Further on it is possible to endow each agent with a conflict-resolution reasoning schema (e.g. a temporal priority). On the other hand it should be noticed that in presence of rational agents and cooperative assumptions the more the agents negotiate the more they learn about their counterpart's preference model and very soon they become able to have a perfect estimation of it. Under this point of view a consensus will be reached in a finite number of steps, as soon as the agents will build a preference model that coincides.

6. Related Work

Research in negotiation models has been developed following different directions. Our approach, compared to the existing ones can emphasize the following features.

First we can distinguish between two main classes of models: centralized or distributed models. The centralized models imply the existence of a special agent who negotiates with other agents (e.g. the coordinator in [Martial 92], the persuader system mediator in [Sycara 89b], etc.). Several criticisms could be made to the centralized approach, at least from the efficiency point of view. Our approach takes into account the intrinsic feature of M.A.S. that is the real distribution of agents, their autonomy (i.e. no external strategy can be imposed to agents), their proper motivations, and their internal decision mechanism.

When the negotiation process is ditributed among agents, (e.g. [Zlotkin and Rosenshein 91]; [Ito and Shintati 97], etc.), only two agents may be involved at once. Our procedure is more general since it can be applied to any sub-set of agents.

Many interesting frameworks have been proposed in negotiation for conflict resolution and coordination. [Zlotkin and Rosenshein 91] describe a negotiation protocol for conflict resolution in non-cooperative domains. They consider that even in a conflict situation, partial cooperative steps can be taken by interacting agents. They also propose [Rosenshein and Zlotkin 94] a “monotonic concession protocol” where each agent must “improve” his offer in every step of negotiation. In [Sycara 89a], negotiation is an iterative process involving identification of potential interactions between non-fully cooperative agents, either through communication or by reasoning about the current states and intentions of other agents. This process allows the modification of intentions of these agents in order to avoid harmful interactions or create cooperative situations. Other works are those of [Klein 91] resolving conflicts generated in the cooperative activity of groups of design agents, each with his own area of expertise. In [Chu-Carroll and Carberry 95], the authors propose a plan based model that specifies how the system (as consulting agent) should detect and attempt to resolve conflicts with the executing agent, [Sandholm and Lesser 95] proposing an automated negotiation protocol for self-interested agents whose rationality is bounded by computational complexity. The above works consider agents either cooperative or self-interested. The procedure we proposed includes both the two cases. In fact, in our framework, agents are lead to cooperate in order to achieve a global goal, while simultaneously trying to satisfy as best as possible individual preferences. However agents can be considered as self-interested if individual goals are independent of a global goal. Our theory can be also applied in this situation, because it solves potential conflict generated by different kind of sources (resource sharing, actions for preferences and/or goal satisfaction, etc.). Agents being pragmatic are aware that the best results for their proper goals require to avoid onflicts.

Other interesting works in negotiation are also those of [Jennings et al. 98] where agents make proposals and counter-proposals by including arguments in order to persuade opponents to change their stance and [Kraus and Wilkenfeld 93] presenting a strategic model for negotiation of alternatives offers which takes into account the effect of time on the negotiation process. Finally [Faratin et al. 98] propose a range of strategies and tactics that agents can employ to generate initial offers, evaluate proposals and offer counter proposals. Offers and counter-offers are generated by linear combination of tactics, which are simple functions. Tactics generate an offer, or counter-offer, for a single component of the negotiation object using a single criterion (time, resources, etc). Different weights in the linear combination allow the varying importance of the criteria to be modeled.

However, the multi-criteria dimension of the negotiation process is basically ignored in all such approaches and this is a main difference with our work. [Faratin et al. 98] have in their works some common elements with our approach. Nevertheless there exists many important differences including: 1) the possibility of our approach to incorporate any preference model (multi-attribute or mono-dimensional preference model) for offers and counter-offers generation; 2) the use of learning procedures (in the case of this paper a multiple linear regression model) allowing each agent to estimate the parameters of the preference model of a counterpart; 3) as well as the possibility to include it (i.e. the preference model of a counterpart) in its own model in order to make a new offer.

7. Conclusion

This paper has presented an automated negotiation procedure based on an aggregation-disaggregation approach. The main innovation consists of introducing the multi-criteria dimension in the agents preference model, which is used to generate offers and counter offers (aggregation step), as well as, learning procedures enabling agents to make an estimation of their counterpart preference model (disaggregation step) and include it in their own models. In this paper the negotiation object is the establishment of $\Re$ (aggregation procedure for the community of agents). Work in progress aims to extend the negotiation objects by considering: 1) the establishment of Γ (collection of binary preference relations for the community of agents), possibly modifying each agent´s G_i (binary preferences relations on the set of his plans), and 2) the establishment of P (the set of plans the community may perform), possibly modifying each agent´s A_i (set of elementary actions), T_i (set of tasks) and H_i (set of binary preferences on his actions). This can be viewed as negotiating the negotiation model itself.

REFERENCES

[Arrow 63] Arrow K., (1963), Social Choice and Individual Values, J. Wiley, New York.

[Chu-Caroll and Carbery 95] Chu-Carroll, J., and Carberry, S. (1995), Communication for Conflict Resolution in MultiAgent Collaborative Planning. In Proceedings ICMAS-95, 49-56, San Francisco.

[El Fallah, Moraïtis and Tsoukiàs 99] El Fallah Seghrouchni A., Moraïtis P. and Tsoukiàs A. (1999): Multi-Criteria Negotiation for Agent-Based Decision makers. DSI'99. 5th International Conference on Decision Science Institute. Athens. June. 99.

[Faratin et al. 98] Faratin P., Sierra C. and Jennings N.R. (1998). Negotiation Decision Functions for Autonomous Agents. In Int. Journal of Robotics and Autonomous Systems 24 (3-4) 159-182.

[Ito and Shintani 97] Ito, T. and Shintani, T. (1997), Persuation among agents: An Approach to Implementing a Group Decision Support System Based on Multi-Agent Negotiation. In Proceedings of IJCAI'97 Nagoya, JAPAN, 593-597.

[Jacquet-Lagrèze and Siskos 82] Jacquet-Lagrèze E. and Siskos J.(1982). Assessing a set of additive utility functions for multicriteria decision making: the UTA method, European Journal of Operational Research. 10: 151-164.

[Jelassi et al. 90] Jelassi T., Kersten G., Zionts S. (1990), An introduction to Group Decision and Negotiation Support, in: C. Bana e Costa (ed.), Readings in Multicriteria Decision Aid, Springer Verlag, Berlin, 537 - 568.

[Jennings et al. 98] Jennings N.R., Parsons S., Noriega P., and Sierra C. (1998). On Argumentation-Based Negotiation. In Proceedings of Int. Workshop on Multi-Agent Systems, Boston, USA.

[Klein 91] Klein, M. (1991), Supporting conflict resolution in cooperative design systems, IEEE Transactions on Systems, Man, and Cybernetics, 21, (6): 1379-1389.

[Kraus 97] Kraus, S. (1997), Negotiation and cooperation in multi-agent domains. Artificial Intelligence 94, 79-97.

[Kraus and Wilkenfeld 93] Kraus S. and Wilkenfeld J. (1993). A Strategic Negotiations Model with Applications to an International Crisis. IEEE Transactions on Systems, Man, and Cybernetics, 23(1):313-323 1993.

[Martial 92] von Martial, F. (1992). Coordinating Plans of Autonomous Agents, Heidelberg: Springer-Verlag.

[Moraïtis and Tsoukiàs 00] Moraïtis, P., and Tsoukiàs, A. (2000). Graph Based Representation of Dynamic Planning. To appear in Proceedings of ECAI-00.

[Moraïtis and Tsoukiàs 96] Moraitis P., Tsoukiàs A. (1996), A Multi-Criteria Approach for Distributed Planning and Conflict Resolution for Multi-Agent Systems. In Proceedings of ICMAS-96, Kyoto, Japan, 213-219.

[Müller 96] Müller H.J. (1996), Negotiation Principles. In Foundations of Distributed Artificial Intelligence, O´Hare and Jennings (Eds.), 211-229.

[Rosenshein and Zlotkin 94] Rosenschein, J. and Zlotkin, G. (1994). Rules of Encounter – Designing Conventions for automated negotiation. AI magazine 15(3), pp . 29-46.

[Sandholm 99] Sandholm, T.W. (1999). Distributed Rational Decision Making. Chapter 5 in "Multiagent Systems: A modern approach to distributed artificial intelligence". G. Weiss (Eds.). MIT press, 201-258.

[Sandholm and Lesser 95] Sandholm T. and Lesser V. (1995). Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework. In Proceedings of ICMAS -95, San Fransisco, pp. 328-335.

[Shehory and Kraus 96] Shehory, O., and Kraus, S. (1996). Emergent cooperative goal-satisfaction in large scale automated-agent systems. In Proceedings ECAI-96, Budapest, 544-548.

[Steuer 86] Steuer R. (1986). *Multiple Criteria Optimization*, J.Wiley, New York.

[Sycara 89a] Sycara, K. (1989a). Argumentation: Planning Other Agents' Plans. In Proceedings of IJCAI'86, Detroit, USA, 517-523.

[Sycara 89b] Sycara, K. (1989b). Multiagent compromise via negotiation. In L.Gasser and M. Huhns, (eds.), Distributed Artificial Intelligence II, Pitman/Morgan Kaufman, 1989.

[Vanderpooten and Vincke 89] Vanderpooten D. and Vincke Ph., (1989). Description and Analysis of some representative interactive multicriteria procedures, Mathematical and Computer Modelling, vol. 12, 1221-1238.

[Vanderpooten 89] Vanderpooten, D. (1989). The interactive approach in MCDA: a technical framework and some basic conceptions, Mathematical and Computer Modelling 12: 1213-1220

[Vincke 92] Vincke, P. 1992. *Multicriteria Decision Aid*. New York: John Wiley

[Zionts and Wallenius 83] Zionts, S. and Wallenius, J. (1983). An interactive multiple objective linear programming method for a class of underlying nonlinear utility functions, Management Science 29: 519-529.

[Zlotkin and Rosenschein 96] Zlotkin, G., and Rosenschein, J.S. (1996), Mechanism design for automated negotiation and its application to task oriented domains, Artificial Intelligence 86, 195-244.

[Zlotkin and Rosenschein 91] Zlotkin, G., and Rosenschein, J.S. (1991). Cooperation and Conflict Resolution via Negotiation Among Autonomous Agents in Noncooperative Domains. *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 21, NO. 6, November/December, 1317-1324.

[Wierzbicki 82] Wierzbicki, A. (1982). A mathematical basis for satisficing decision making, Mathematical Modelling 3: 391-405.