A Multicriteria Approach for Distributed Planning and Conflict Resolution for Multi-Agent Systems

Pavlos Moraïtis and Alexis Tsoukiàs

LAMSADE, University of Paris Dauphine
75775 Paris Cedex 16, France
e-mail:{moraitis, tsoukias}@lamsade.dauphine.fr

Abstract

A multicriteria approach for distributed planning and conflict resolution for multiagent systems is presented in the paper. A graph representation is used to model the plans of each agent and a multicriteria "best" path procedure can be used in order to obtain the "best" path for each agent considering its private and local goals. In our approach private and local goals do not necessarily coincide. A conflict and/or positive cooperation may occur and a detection procedure is triggered as soon as the agents broadcast their "best plans". A negotiation process is then established if necessary. Such a process iterates the use of the graph representation and of the "best" path algorithm to a level including the agents that have to negotiate. The parameters of the multicriteria model used to evaluate the paths are then negotiated or even the model itself. Some open problems conclude the paper.

Introduction

Distributed planning, conflict resolution and negotiation are central themes of Distributed Artificial Intelligence. This paper treat these three items. Distributed planning may concern either a group of agents pursuing goals which achievement is necessary for the one of a common global goal (multiagent cooperative problem solving) or a group of self-interested agents with different individual goals (planning by multiple agents). In the following, the first type goals will be called "local" goals and the second type ones, "private goals". In the first case, (multiagent cooperative problem solving) agents work as a whole team and are keen to cooperation activity. In the second case (planning by multiple agents) agents are autonomous. There is no overall planning task which has to be solved in close cooperation, each agent wanting to solve its individual problem. In both cases a conflict resolution problem may arise, either between agents' local goals (first case) or between agents' private goals (second case). Conflict resolution may imply a negotiation process among conflicting agents.

In this paper we propose a multicriteria perspective (Vincke, 1992) for distributed planning and conflict resolution, integrating both points of view of distributed planning. We consider that agents accomplish local goals necessary for global goal achievement trying simultaneously to optimize private goals (interests or motivations). Our approach is based on a planning-negotiation-execution (PNE) cycle similar to coordination process of Martial (1992), but in our approach all agents have the same functionalities. So each agent itself conceives, negotiates and executes its plans (for example there is no coordinator agent to which the agents can transfer their plans or executor which is unable to negotiate about plans). The agents have the possibility to develop plans using a multicriteria model, to evaluate their single actions or their plans using a set of criteria, and to measure how well a plan helps them to reach there goal(s) (local and private) (planning stage). Then plans are exchanged among agents and a cycle of negotiation is initiated in order, either to detect harmful interactions and resolve conflicts or to create "positive cooperation" when two agents discover that a collaboration may improve their goals further than acting alone. Conflict resolution is based on a negotiation process between agents, about the scale of the criteria used, the set of evaluation criteria and the proposed plans (negotiation stage). As soon as the negotiation process provides a solution, generated plans are executed (execution stage).

In this paper sections 2 and 3 represent a brief overview of related works in order to situate our work. Section 4 describes our assumptions and section 5 gives an overview of global process of our approach of problem solving. Section 6 presents our multicriteria planning model, by introducing an example used throughout this paper. Section 7 describes our multicriteria negotiation process used for conflict resolution and section 8 compares related works with our. At the end this paper evokes some open problems.

Distributed Planning

Multi-agent Cooperative Problem Solving

In Durfee (1988), Durfee and Lesser, (1987) coordination mechanisms allow the agents to plan coordinated interactions and to modify their plans in order to improve common performance. The plans are exchanged in order to allow agents greater access to the anticipated future behavior of others, avoiding also conflicts. In Ephrati et al.

(1995) optimal global plan (with arbitrary global criteria) is derived by appropriately combining local optimal plans. The synthesis of the global plan is based on the research of a global uility that may differs radically from individual agent utilities. In Ephrati and Rosenschein (1993), agents incrementally construct a plan that brings the group to a state maximizing social welfare. This is the result of a voting procedure between agents, without having to reveal full goals and preference at each step of group planning, in order to choose the next joint action. Sycara (1989) considers non cooperative interactions where cooperation cannot be assumed but needs to be dynamically induced during problem solving. For that, Sycara proposes a process of persuasive argumentation based on negotiation to purposefully modify the plans, goals, and behavior of other agents to increase agent cooperativeness in order to bring about convergence to a global solution. Multicriteria methodology is implicitly used (maximize a global utility function or use of a social welfare choice function), but in a narrow context.

Planning by Multiple Agents

Multiagent planning has been approached in several different ways. In such a context separate plans can be combined in a way that avoids interference among the agents executing the plans (Georgeff, 1987). To solve this problem, it is necessary to ascertain, from descriptions of the actions occurring in the individual plans, which actions could interfere with one another and in what manner (Georgeff, 1984). After this has been determined, a coordinated plan that precludes such interference must then be constructed. This plan can be formed by inserting appropriate synchronization actions (inter-agent communications) into the original plans to ensure that only interference-free orderings will be allowed (Georgeff, 1983). In Martial (1992), agents develop their plans autonomously. A communication framework allow agents to exchange their plans and to negotiate about how to resolve the relations (negative or positive) between their plans. Implicitly a multicriteria problem is again settled, but no explicit reference to such a methodology is done.

Conflict Resolution and Negotiation

In distributed planning, agents (cooperative or self-interested) have only a local view of the overall problem, therefore conflicts may exist among agents' subplans and redundant actions may also have been generated. Conflict generation is also linked to the distributed planning process. The problem implies both conflict detection and resolution. Many researchers have developed different conflict resolution strategies. Most of them have proposed negotiation mechanism for conflict resolution. Zlotkin and Rosenshein (1991) describe a negotiation protocol for conflict resolution in noncooperative domains (where "negotiation set" might be empty). They consider that even in a conflict situation, partial cooperative steps can be taken by interacting agents. In Sycara (1989) 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 these of Conry et al. (1988) which have developed a negotiation protocol for cooperatively resolving resource allocation conflicts, Klein (1991) resolving conflicts generated in the cooperative activity of groups of design agents, each with their own areas expertise, Chu-Carroll and Carberry (1995) proposing a plan based model that specifies how the system (as consulting agent) should detect and attempt to resolve conflicts with the executing agent. These conflicts are generated by the proposed actions and the agents' beliefs. The proposed model provide a mechanism by capturing multiagent collaboration in a recursive cycle of Propose-Evaluate-Modify which allow the conflict detection and resolution by modification of actions and/or beliefs. The multicriteria dimension of the negotiation process is basically ignored in all such approaches.

Assumptions

Throughout this paper, we make the following assumptions, some of them being the same as in (Zlotkin and Rosenschein, 91):

- A set of agents $A=\{\alpha_1,...,\alpha_k)$ is given which we consider as actors of a system Φ. All the agents in the system work towards the achievement of a global goal o_Φ ;
- *Complete Knowledge*: each agent knows all relevant information;
- *No History*: There is no consideration given by the agents to the past or the future; each negotiation stands alone;
- *Fixed Local Goals*: the agents can negotiate about the scale of criteria evaluation, actions and plans evaluation criteria choice, proposed plans, private goals, but their local goals remain fixed;
- *Bilateral Negotiation*: negotiation is done between a pair of agents;
- *Not Necessarily Utility Function*, individual or global;
- *Not Necessarily Symmetric Abilities*: all agents are not necessarily able to perform the same set of operations; agents can have complementary abilities;

Planning-Negotiation-Execution Cycle

Our approach is based on a cycle planning-negotiation-execution (PNE) similar to coordination process of Martial (1992). But in contrast to Martial's approach, where agents may have one or several roles among those of coordinator, executor or mediator, in our approach all the agents do have the same functionalities. Each agent itself conceives, negotiates and executes its plans during the (PNE) cycle (for example there is no coordinator agent to which the agents can transfer their plans or executor which is unable to negotiate about plans). Each agent is completely autonomous, do not necessarily share or know the goals of the other agents and global problem solving has to emerge as a result of agent interactions. So agents have to coordinate dynamically their plans in order to achieve any global goal by trying to optimize simultaneously their private goals. First, the agents develop their "best plans" in parallel. Then, they communicate in order to transfer these plans to each other. During the planning stage agents develop best plans for local and private goals. Agents exchange their plans in order to broadcast their intentions to each other and either to detect harmful interactions (conflicts) or create positive cooperation situations. Conflict resolution or positive cooperation generation are carried out during negotiation process. Negotiation process forces the agents to work towards a global coherent solution, even if some of them are "not sincerely cooperative". The parallel execution stage begins, either when there is no objection on generated plans from all the involved agents (all the agents did agree), or when the negotiation's end occurs. In this paper we will present the planning and negotiation stages but not the execution stage and its related problems (i.e. synchronization of executed actions).

The Multicriteria Planning Model

Definition 1-*Local and Private Goals*

To each agent a set of objectives can be associated under the form of a vector $o=\{o_1,...,o_n\}$ where o_j can be a local or private goal and can be described by either a numerical value or whatever quantifiable event or by a qualitative (symbolic) description. A local goal is necessary for global goal achievement and one or more private goals express(s) agent's motivations. We denote as o_q the set of goals of agent α_q (o_{qj} being the j-th component).

In both domains of research on distributed planning discussed above, agents have either a global common goal or individuals goals. In our approach, we consider that agents accomplish *local goals* necessary for global goal achievement trying simultaneously to optimize *private goals.* The last ones can be inconsistent with their own local goals or with the private goals of other agents. Conflicts may also exist between local goals and resources sharing. We consider a situation where agents have to cooperate to achieve a global goal simultaneously trying to take out as much personal benefit as possible from this situation.

This situation is presented in real world and specially in social organizations where a cooperative work process is performed by individuals with individual interests and motivations (March and Simon, 1958). The individuals involved in cooperative production may have conflicting economic interests and ideological allegiances (Schmidt, 1991). Cyert and March, (1963) consider that an enterprise is made of general individual's coalitions with their own goals. So enterprise's goals appear as the result of a negotiation process between coalitions and between individuals.

The private goals of an agent can be inconsistent with its local goal in the sense where the best actions allowing the local goal optimization for the global good are not necessarily the ones which optimize its motivations. In other words, agent would not have done the same actions for its local goal if it was completely "devoted" to the global success (personal interest sacrifice in front of the collective one's) and if it had no "back thoughts" on what it can win from its cooperation with the others.

An Example

In this section, we present an example which points out our multicriteria approach and the chosen context. Let us consider a room with a bookcase (C), a table (A) and a heavy piece of furniture (B). This state of the world is presented as: (IN(A), IN(B), IN(C)). Two agents α_1 and α_2 have, as a common goal, to empty this room. This new state of the world will be presented as: (OUT(A), OUT(B), OUT(C)).

The world may be described by the following relations:

- Object characteristics: heavy(furniture); light(table); light(bookcase);
- Possible actions associated on agents abilities: wait(z); transport(α_1,x)$\wedge$light(x);takedown(α_1,x);transport(α_2,x) $\wedge$(light(x);transport(α_2,x)$\wedge$heavy(x);transport-r(α_2, x, y) $\wedge$ light(x)$\wedge$light(y)
- Private agents' goals: p-goal(α_1, max-profit); p-goal(α_2, min-time);
- In this example, local goals are not clearly distinct. Therefore, we do consider that each agent has to use its abilities in order to achieve the global goal, as far as they have complementary abilities.
- Relation between actions: before(takedown(bookcase), transport(bookcase));

We assume that when α_1 performs an action there is a profit of 1 money unit and when α_2 performs an action loses 1 time unit.

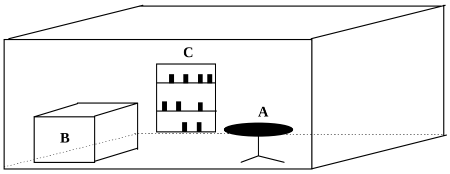

Figure 1: Nil state of the world

Definition 2-*Actions*

For each agent α_q a set of possible actions A^q can be associated $A^q=\{a_{q1},...,a_{qm}\}$ which can be performed in order to enable the agents to attain its goals.

Definition 3-*Criteria*

Consequently each agent is equipped with a set of criteria $G^q=\{g_{q1},...,g_{qn}\}$ each of them corresponding to a goal. In order words g_{qj} (a_{qi}) measures how well the action a_{qi} helps the agent α_q to reach its goal o_q.

Tables 1 and 2 summarize the possible actions for each agent and how they are evaluated following the criteria modeling their private goals.

α_1	TR(1, A)	TR(1, C)	TD(1, C)
C_{11}	1	1	1

Table 1: Agent α_1 actions and evaluations

α_2	TR(2, A)	TR(2, B)	TR(2, C)	TR-R(2, A, C)
C_{21}	1	1	1	1

Table 2: Agent α_2 actions and evaluations

Definition 4-*Plans*

An agent has to perform a sequence of actions which may have a cumulative effect (of different nature) which may affect one or more goals. Of course different sequences of actions are possible and each of them corresponds to different levels of attainment of the goals. We call any possible sequence of actions a "plan" and we denote it as $p=<a_i...,a_k>$. Not all plans are individually feasible because some actions may have a condition which has to be fulfilled by another agent.

Definition 5-*Best Plans*

A "best" plan is the sequence of actions that enables an agent to attain its goals to the best possible level. From a technical point of view the difficulty consists to the fact that the agent can evaluate any single action through its criteria, while it needs to evaluate plans (which are defined by the actions which belong to them). In the decision theory this problem is known as the "fragmented alternatives evaluation problem" (see Roy, 1985). Unfortunately such a problem has been studied very little in literature. In the following we will use some ideas introduced in Dellacroce et al. (1996).

Consider a directed graph $\Gamma=(A, N)$, A being the arcs and N being the nodes. We associate to each action a_i an arc and to each node a "state of the world" which describes the state of the agent confronted to its goals. Each "state" of the world enables only a limited set of successive actions which are the arcs departing from the node. The source node corresponds to a "nil state of the world" where no action has been performed and which is considered the worst state. The intermediate nodes account for the eventual cumulative effect of a sequence of actions in any way such accumulation is measured. More formally we associate to each node n^k an ordered couple $<d^k, s^k>$ where d^k is the description of the state of the world and s^k is a vector $[s^k_1 ... s^k_n]$ representing the level of attainment of the agent's goals. Such a level depends from the sequence of actions previously performed, that is from the different paths leading to such a node. In other words s^k_j represents the j-th evaluation criterion of the path leading to the node n^k. Under such a representation the choice of the "best" plan corresponds to the identification of the "best path" on the graph Γ.

Figures 2 and 3 present the graphs of possible paths for agents α_1 and α_2 where the best paths are in dotted lines.

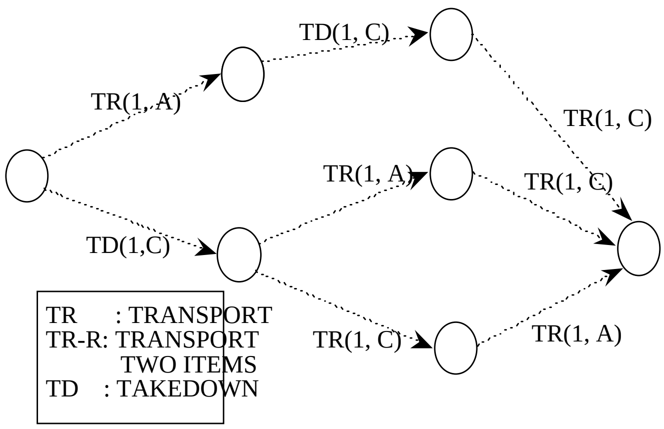

Figure 2: Agent α_1 possible and best paths

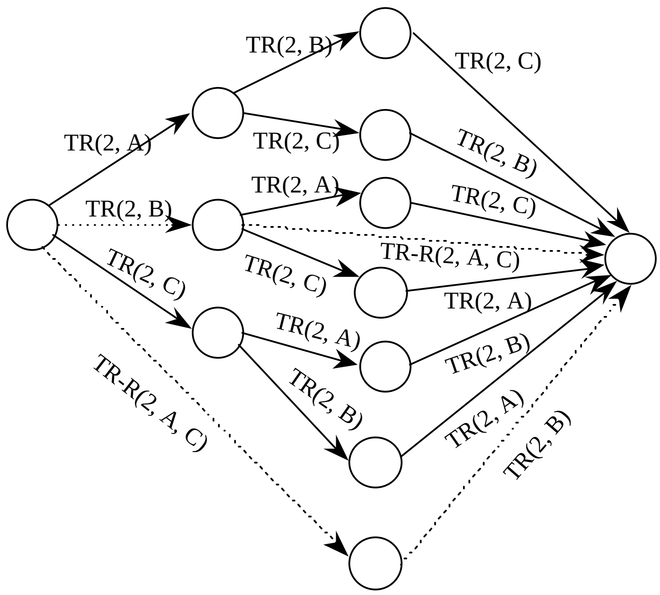

Figure 3: Agent α_2 possible and best paths

All possible paths for agent α_1 are equivalent, the sink node been labeled in all cases as couple $<d^s, s^s>$ where: d^s=(OUT(A), OUT(C), IN(B)) and s^s_1 =3. On the other hand is interesting to observe that no paths of α_2 graphs is feasible since it can not perform TR(2, C) or TR-R(2, A, C) before TD(1, C). Under such conditions the best paths are two (see figure 3) leading to a sink node labeled $<d^s, s^s>$ where d^s=(OUT(A), OUT(C), OUT (B)) and s^s_1=2. Only agent α_2 sink node corresponds to the system's goal, but such an attainment is conditioned by α_1 which is the only agent that can perform TD(1, C).

In this example both paths are computed using the usual shortest (longest) path algorithms (see Deo and Pang, 1984) since the cumulative effect of the actions is additive. Under the more general setting the information available on each arc is a vector of not necessary quantified (or cardinal) evaluations, that is the best path choice is a multicriteria ordinal "best" path problem (this is one of the reasons for which we do not assume that necessarily exists an utility function). This is also a problem which has been studied very little in literature. Hansen (1980) and Henig (1994) present interesting surveys for the multicriteria shortest path problem (but the information is considered implicitly cardinal, therefore efficient paths can be calculated), while Dellacroce et al., (1996), study some new classes of "ordinal best paths", but not necessary under a multicriteria setting. In the following we will not continue on the problem of the best path, since it is out of the scope of the paper. We will make the assumption that each agent is equipped with the necessary algorithms, although large part of them is in NP. An agent α_q is therefore able to compute a best plan which we denote as p^*_q

Definition 6-*Coordination Problem*

A coordination problem occurs when the feasible paths of any agent do not lead to a node labeled with the system's goal.

In our example both agents have a coordination problem since their feasible paths do not solve the problem of getting the room empty.

Definition 7-*Conflict Situation*

Intuitively conflict situation occurs when an action performed by an agent impedes an action to be performed by another agent or when induces a worsening to another agent's goals. More formally consider the system Φ. To each agent we can associate a graph Γ. Given any two agents α_x and α_y each of them may compute its best paths p^*_x and p^*_y .

A conflict situation ($p^*_x \perp p^*_y$) may occur when:

- $\exists\, a_i \in p^*_x$, $a_j \in p^*_y$: φ_{a_i} (x, k, ...) $\wedge \psi_{a_j}$ (y, k, ...)

In other words there exist two actions (a_i , a_j) belonging to the best paths of α_x and α_y respectively, such that the associated predicates φ_{a_i} and ψ_{a_j} (each action is a n-ary predicate) have a common extension of variables. In our example the problem exist among the actions TR(1, A), TR(1, C), TR-R(2, A, C).

- $\exists\, a_i \in p^*_x \exists\, n^k \in p^*_x \cap p^*_y \exists\, s^k_j$ $s^{kx}_{jy} < s^{ky}_{jy}$

where s^{kx}_{jy} indicates the level of the j-th goal of agent α_y of the node n^k when reached by p^*_x and s^{ky}_{jy} when the same node is reached by p^*_y , < standing for a strict preference relation. In other words the same state of the world can be reached under a descriptive point of view, but with different levels of goals attainment.

The problem arise in our example since the state of the world labeled for instance (OUT (A) and OUT(C)) can be reached with different goal attainment for the two agents.

Definition 8-Positive Cooperation Situation

A positive cooperation situation occurs when at least two agents discover that the collaboration may improve their goals further than acting alone:

- $\exists\, a_i \in \Gamma_x \exists\, p^{**}_y \in \Gamma_\Phi$ $p^{**}_y > p^*_y$

< standing for a strict preference relation. The condition represents the situation where exists an action among the possible actions of α_x such that it exists a path p^{**}_y for α_y which is better than its best path computed ignoring the other agents.

The positive cooperation situation will not be detailed in this paper, but we consider it will be treated in a similar way than conflict resolution, by negotiation process presented in §7.

Negotiation Process and Conflict Resolution

In the following we will consider only the case where the negotiating agents are "sincerely cooperative" in the sense that they effectively look for a compromise, being available, if necessary, to lose something on their private goals in order to reach the system goal. Generally this may be not always the case. Opposite to "sincerely cooperative" agent, we consider a "not sincerely cooperative" agent, the one which can "bluff" refusing any concession, hoping to force a "sincerely cooperative" agent to do as many concessions as possible and trying to delay its concession as much as possible. It is possible to foresee a negotiation limit due to resources bounds where the system randomly chooses a solution therefore randomly penalizing some agents. This is a kind of "game rule" accepted by each agent, entering the system.

When situations of coordination, conflict and/or positive cooperation occur among two different agents (which have a complete knowledge of Γ_Φ and of their respective best paths) a negotiation process is established. In order to model negotiation we will again use a graph representation and a multicriteria model. The negotiation process is then structured as follows.

- Establish a negotiation graph concerning the negotiating agents. In such a graph nodes represent always states of the world labeled $<d^k, s^k>$ where s^k is now the vector of goals of the negotiating agents. Arcs represent n-uples of actions performed in a parallel way by the negotiating agents. Such a graph is constructed using the sub-graphs of best paths of its negotiating agent and solves the coordination problem and the first type conflict situations (common actions) if any. In our example the negotiation graph is represented in figure 4.

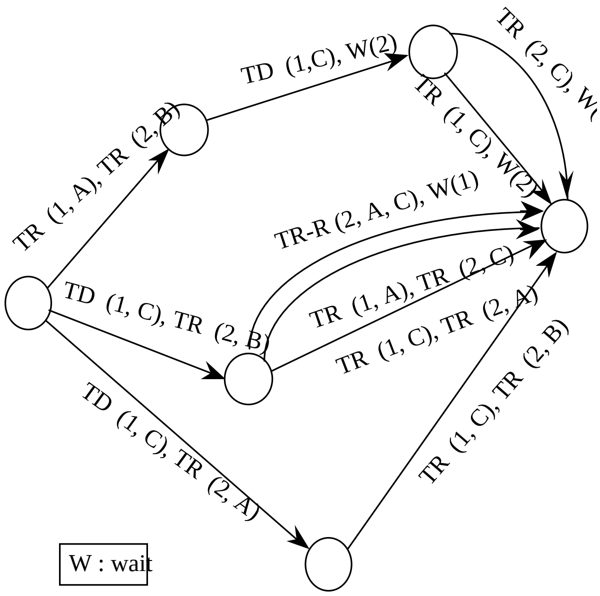

Figure 4: The negotiation graph for agents α_1, α_2

- A new multicriteria model can be settled using the union of the set of criteria of each negotiating agent. There exists two possibilities: the first is to use an hierarchical model of preference aggregation (each agent criteria aggregated to a single criterion and so on), the second is to use a flat model considering all the criteria contemporaneously. The choice depends on the nature of the criteria and the agent's preferences. The first negotiation step consists in trying to define a compromise solution among the efficient paths of the negotiation graph. Different procedures can be used as establishing a global utility function (if the agents accept establish trade-offs among their criteria), go through direct pairwise comparisons (if the agents accept to simply compare their criteria) and so on. In our example the efficient paths are:

1.(TR(1, A),TR(2, B)),(TD(1, C),W(2)),(TR(1, C),W(2))

2. (TD(1, C),TR(2, B)), (TR(1, A), TR(2, C))

3. (TD(1, C),TR(2, B)), (TR(1, C), TR(2, A))

4. (TD(1, C),TR(2, A)), (TR(1, C), TR(2, B))

For example if a trade-off is accepted such that each time unit is equivalent to 0.2 unit of profit the path 1 is the best one.

- A second negotiation step, in the case the first fails to find a compromise solution, is to re-discuss the model enhancing or contracting the criteria set. Under the new model a negotiation as in the first step can hold. In our example each agent could introduce a cost function such that the following tables 3 and 4 hold. Such a cost function could be introduced because the agents realize to use a common resource which is limited.

α_1	TR(1, A)	TR(1, C)	TD(1, C)
C_{11}	1	1	1
C_{12}	2	1	1

Table 3: New criteria set for agent α_1

α_2	TR(2, A)	TR(2, B)	TR(2, C)	TR-R(2, A, C)
C_{21}	1	1	1	1
C_{22}	0,5	1	0,2	1

Table 4: New criteria set for agent α_2

- A third step, in the case the two previous fail, is to change again the model introducing new actions which were not previously considered. If such a situation occurs the criteria set has also to be redefined. The negotiation goes back to the first step and the whole process restarts. In our example a new action that can be introduced is that an agent can quit the system before the system reaches the final state. In this case new path are added including such actions.

- The stoping condition is either an agreement reached among the negotiating agents in the sense that a compromise solution is accepted or the negotiation limits are exceeded and the system randomly imposes a solution (an efficient path of the negotiation graph).

Related work

In distributed planning field the other works consider either that agents are completely cooperative or that agents are self-motivated without global goal. The difference with our work is that we consider that agents accomplish local goals necessary for global goal achievement trying simultaneously to optimize private goals. More precisely, the originality of our work compared to the multiagent cooperative problem solving (Durfee and Lesser, 87; Decker, 95; Ephrati, Pollack and Rosenschein, 95) where agents identify their private goals with their local goals, is that our agents make the difference between these two types of goals and can take out individual benefice, as much as possible, from a situation where they work for a common goal. Compared to (Ephrati and Rosenshein, 93b) where non-benevolent agents should be motivated to contribute to the global goal (should be paid for their work) in our work, agents can simultaneously take in to account several private goals. The level of satisfaction of these goals depends on their negotiation ability with others. Another interesting point related to the other works (for example the cost in Zlotkin et Rosenschein, 91; Ephrati and Rosenschein, 93b), when the best plan is computed according to a single criterion, is that in our work, this computation is made according to several criteria. In the conflict resolution field, conflicts and positive interactions are addressed within a unified framework as in (Ephrati and Rosenshein, 93b). Compared to Sycara(89) and Martial (92) where negotiation is made through a persuader or coordinator, in our work, our agents do negotiate directly. Grosz and Kraus (93) develop a Shared-Plan model for collaborative activities where conflicting intentions are avoided instead of trying to resolve them. In (Zlotkin et Rosenshein 89; Zlotkin et Rosenshein, 91) where agents are considered either as fully-cooperative or not-at-all cooperative. The most important aspect of our work is that while other authors choose to impose a priori a multicritera model (implicitly or explicitly) we consider that this can be negotiated (for instance accept or not trade-offs). Moreover in our approach the multicriteria model is not fixed, since everything can be negotiated including the criteria set and the set of possible actions. Finally the graph representation gives an explicit representation of the agents difficulty to compare plans which correspond to paths.

Conclusion and open problems

In this paper a multicriteria approach for distributed planning and conflict resolution in a multiagent system context is outlined. The basic innovation consists in introducing a multicriteria graph representation of agent plans and the use of a multicriteria model for negotiation purposes. Other innovations are that private and local goals do not necessarily coincide in a cooperative work and that agents can negotiate everything including the criteria set of multicriteria model used to evaluate the paths, the set of possible actions or even the model itself. Of course the problems open by our approach are more than the ones solved. The question of computing the best paths in the individual graphs and the negotiation graph is by-passed, while it represents a big technical problem since very few efficient algorithms exist in the literature for this purpose. Another point related to our approach, which will be introduced in our future work, is the problem of interleaving of planning and execution in the case of a not foreseen event arrival (for instance in our example the arrival of a new agent). And this in order to preserve, by the new planning-negotiation-execution cycle, the effects of already executed actions.

Acknowledgments

The first author was with the support of a CNET FRANCE TELECOM CTI project N° 904-15-4371-123 which is gratefully acknowledged.

References

Chu-Carroll, J., and Carberry, S. 1995. Communication for Conflict Resolution in MultiAgent Collaborative Planning. In Proceedings of the First International Conference on Multiagent Systems (ICMAS-95), 49-56, San Francisco.

Conry, S.E., Meyer, R.A., and Lesser, V.R. 1988. Multistage negotiation in distributed planning. In Bond, A.H., and Gasser, L., (eds.), *Readings in Distributed Artificial Intelligence*, 367-383, San Mateo, California: Morgan Kaufmann Publishers, Inc.

Cyert, R.M., and March, J.G., 1963. *A Behavioral Theory of the Firm.*. Englewood Cliffs, New York: Prentice Hall.

Decker, K.S. 1995. Designing a Family of Coordination Algorithms. In Proceedings of the First International Conference on Multiagent Systems (ICMAS-95), 73-80, San Francisco.

Dellacroce F., Paschos V., and Tsoukiàs A. The shortest path problem under qualitative information, forthcoming.

Deo N., and Pang C.Y. 1984. Shortest path algorithms: a taxonomy and annotation, *Networks*, 14: 275 - 323.

Durfee, E.H. 1988. *Coordination of Distributed Problem Solvers*. Boston: Kluwer.

Durfee, E.H., and Lesser, V.R. 1987. Using partial global plans to coordinate distributed problem solvers. In

Proceedings of the 10th International Joint Conference on Artificial Intelligence, 875-883, Milan.

Ephrati, E., Pollack, M.E., and Rosenschein, J.S. 1995. A Tractable Heuristic that Maximizes Global Utility through Local Plan Combination. In Proceedings of the First International Conference on Multiagent Systems (ICMAS-95), 94-101, San Fransisco.

Ephrati, E., and Rosenschein, J.S. 1993. A Framework for the Interleaving of Execution and Planning for Dynamic Tasks by Multiple Agents, In Proccedings of MAAMAW'93, 137-153, Neuchâtel, Switzerland.

Georgeff, M.P. 1983. Communication and interaction in multiagent planning. In Proceedings of Third National Conference on Artificial Intelligence, 125-129, Washington, DC.

Georgeff, M.P. 1984. A Theory of action for multiagent planning. In Proceedings of Fourth National Conference on Artificial Intelligence, 121-125, Austin, Texas.

Georgeff, M.P. 1987. Planning. In *Readings in Planning*, 5-25.

Grosz, B., and Kraus, S. 1993. Collaborative plans for group activities. In Proceedings of the 13th International Joint Conference on Artificial Intelligence, 367-373, Chembéry, France.

Hansen P. 1980. Bicriterion path problems. In G. Fandel, T. Gal, (eds.), *Multiple Criteria Decision Making: Theory and Applications*, LNEMS 177, 109-127, Heidelberg: Springer-Verlag.

Henig M. 1994. Efficient Interactive Methods for a Class of Multiattribute Shortest Path Problems, *Management Science*, 40: 891 - 897.

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

March, J.G., and Simon, H.A. 1958. *Organizations*, New York: Wiley.

Martial, v.F. 1992. *Coordinating Plans of Autonomous Agents*, Heidelberg: Springer-Verlag.

Roy B., 1985. *Méthodologie multicritère d'aide à la décision*, Paris: Economica.

Schmidt, K. 1991. Cooperative Work: A Conceptual Framework. In J. Rasmussen, B. Brehmer and J. Leplat (Eds.), *Distributed Decision Making: Cognitive Models for Cooperative Work*, 75-110, New York: John Wiley.

Sycara, K. 1989. Argumentation: Planning Other Agents'Plans. In Proceedings of the 11th International Joint Conference on Artificial Intelligence, 517-523, Detroit, USA.

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

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.

Zlotkin, G., and Rosenschein, J.S. 1989. Negotiation and Task Sharing Among Autonomous Agents in Cooperative Domains. In Proceedings of the 11th International Joint Conference on Artificial Intelligence, 912-917, Detroit, USA.