

# Resolving Goal Conflicts via Negotiation<sup>1</sup>

Katia Sycara

The Robotics Institute, Carnegie Mellon University  
Pittsburgh, PA 15213

## Abstract

In non-cooperative multi-agent planning, resolution of multiple conflicting goals is the result of finding compromise solutions. Previous research has dealt with such multi-agent problems where planning goals are well-specified, subgoals can be enumerated, and the utilities associated with subgoals known. Our research extends the domain of problems to include non-cooperative multi-agent interactions where planning goals are ill-specified, subgoals cannot be enumerated, and the associated utilities are not precisely known. We provide a model of goal conflict resolution through *negotiation* implemented in the PERSUADER, a program that resolves labor disputes. Negotiation is performed through proposal and modification of goal relaxations. Case-Based Reasoning is integrated with the use of multi-attribute utilities to portray tradeoffs and propose novel goal relaxations and compromises. Persuasive arguments are generated and used as a mechanism to dynamically change the agents' utilities so that convergence to an acceptable compromise can be achieved.

## 1. Introduction

Multi-agent planning systems [Konolige 80, Cammarata 83, Rosenschein 86, Fox 84, Durfee 85] emphasize a problem decomposition where individual agents assume responsibility for the generation of their own plans while taking into account the intentions of other agents in the system. AI work has focused primarily on fostering cooperation and avoiding goal conflicts [Georgeff 84]. Most AI research that has dealt with conflicting interactions of conjunctive goals has addressed goal conflict resolution in the limited sense of reordering plan steps so as to avoid having the action of one step invalidate a precondition of a following step (e.g., [Sussman 75, Hammond 86]).

In environments where cooperative behavior of the agents cannot be assumed, goal conflicts have to be resolved by finding *compromises*. This limits the usefulness of traditional planning techniques which assume that conjunctive planning goals must be totally satisfied. Previous work in modeling

interactions among non-cooperative intelligent agents is the work of Rosenschein (e.g., [Rosenstein 86]) who used a game-theoretic approach characterized by payoff matrices that contain the agents' payoffs for each possible outcome of an interaction. Most of the research assumes a single encounter, and no attempt is made to influence other agents' utilities and payoffs, i.e. the dynamics of negotiation is ignored. It is assumed that agents have *common knowledge* of the payoff matrix associated with the interaction, an assumption that is unrealistic considering that the agents are non-cooperative. Another drawback of the approach is that, even if the payoff matrix were known, for large games involving many agents and outcomes, the matrix may quickly become intractable.

Problems with conflicting conjunctive goals fall into two classes. The first contains problems where planning goals are well specified, subgoals can be statically enumerated and the utilities associated with the subgoals are known a priori and are *static*. These assumptions are very limiting considering that the agents may not be cooperative, that during the course of an interaction, new subgoals may be generated and utilities change. Scheduling [Fox 84], resource allocation [Sathi 86] and the multi-agent interactions described by Rosenschein (e.g., [Rosenstein 86]) are members of the first class. For such problems, decision theoretic techniques can in general be used to find an optimal compromise of the conflicting goals.

The second class contains problems where the planning goals are ill-specified, subgoals cannot be enumerated and utilities are not precisely known. Moreover, the problem solving process itself influences the search for a solution. Such problems arise in any complex domain where machines or humans are engaged in group problem solving. In an automated factory, for example, robots compete for limited resources; in design of complex systems, designers responsible for different parts need to find an acceptable overall design. Our research contribution consists in providing a methodology that extends the domain of problems that need compromise resolution of goal conflicts to problems of the second class. For such problems, optimal solutions cannot be found. The best that can be done is to use *negotiation* to find a compromise acceptable to all agents.

In our framework, negotiation is performed through proposal and modification of goal relaxations. Goal relaxations are alternative ways of achieving a goal. Case-Based Reasoning (CBR) is integrated with use of multi-attribute utilities to portray tradeoffs and propose novel goal relaxations and compromises. The negotiation process itself is

---

<sup>1</sup>The research was funded in part by the ARO contract No. DAAG-29-85-K0023

a search of a dynamic problem space where an agent's beliefs about other agents' beliefs over the cycle of proposals continuously changes the space being searched. What was *not* a solution at one point becomes a solution at a later point<sup>2</sup>. Persuasive arguments are generated and used as a mechanism to *dynamically change* the agents' utilities associated with the conflicting goals so that convergence to an acceptable compromise can be achieved.

Our theory of goal conflict resolution is implemented in the PERSUADER, a program which, acting as a labor mediator, enters in negotiation with each of the parties, the union and company, proposing and modifying compromises and utilities until a final agreement is reached. We used two practicing Federal Mediators as experts: one for development and the other for informal validation<sup>3</sup>. Although the program operates in the domain of labor relations, the techniques it uses are domain independent. The PERSUADER is a complex system. In the allotted space, we can only present a simplified and somewhat high level view of the system.

## 2. Requirements of a Planner for Resolution of Goal Conflicts (RGC)

A planner for RGC via negotiation has as input the conflicting goals of multiple agents and as output a compromise acceptable by the agents. For continuum-valued issues the choices that such a planner has are infinite. Hierarchical decomposition of the problem into smaller subproblems each of which is easier to solve may not be suitable, since a compromise solution may be a "package" whose parts are strongly interconnected and interacting. These difficulties are compounded by the absence of a coherent set of constraints that could guide search through the space of all possible compromises. The problem characteristics impose some requirements on a planner for RGC via negotiation:

- Conflict resolution involving multiple agents with conflicting goals is a cooperative search problem. The planner must guide the agents through a sequence of possible compromises to a final compromise that is acceptable by all. Therefore, a planner for RGC needs to plan in an *iterative* rather than a *one shot* fashion.
- Since the expertise resides in the agents, they have to give feedback to the conflict resolver, about which tradeoffs are acceptable. Hence, a planner for RGC needs to be able to receive *feedback* about the quality of its plan, evaluate it,

---

<sup>2</sup>In labor negotiations, for example, it is unlikely that either party would accept their eventual compromise, if it were presented at the inception of negotiations.

<sup>3</sup>Since Federal Mediators are required by law to destroy all documents and notes pertaining to a case, it was not possible to validate the system by running real cases.

and use it to generate a counterproposal.

- During the course of negotiations, conditions in the world that affect the agents' behavior and goals might change. Therefore, a planner for RGC needs to be able to *react* to the changing planning context.
- A planner for RGC must have a way of *predicting/evaluating* whether each new counterproposal leads toward convergence.
- A planner for RGC needs to have a component that generates *persuasive arguments* to change the parties' utilities.

In the rest of the paper, we present the PERSUADER as a model of RGC via negotiation. It incorporates all the above characteristics.

## 3. Overview of the PERSUADER

The PERSUADER's input is the set of conflicting goals of the company and union, and the dispute context. Its final output is either a single plan in the form of an agreed upon compromise, or an indication of failure if the parties to the dispute did not reach agreement within a particular number of proposals. The PERSUADER's tasks are: (a) propose an initial compromise, (b) repair and improve a rejected compromise, (c) persuade the parties to change their evaluation of a compromise. The PERSUADER views these tasks as planning tasks.

As shown in Figure 1, to perform its tasks, the PERSUADER integrates Case-Based Reasoning (e.g., [Kolodner et al. 85, Hammond 86] (a heuristic technique) and Preference Analysis (a decision theoretic, analytic method). Expert labor negotiators/mediators take into consideration *prevailing practice*, the bargaining behavior of similar disputants, as well as informal notions of the parties' utilities in coming up with an acceptable compromise. In our work, prevailing practice is abstracted to Case-Based Reasoning (CBR), and the parties' utilities are modeled through Preference Analysis. Three foci, contracts, impasses and arguments are used as basis for Case-Based Reasoning at performing the above tasks. The PERSUADER keeps track of compromises that have worked in the past in similar circumstances. The most suitable is retrieved from memory and adapted to fit the current situation. The compromise is then proposed to the parties. If the parties agree, the case memory is updated with a successful episode. If one of the parties disagrees, the PERSUADER either repairs the compromise to better accommodate the rejecting party's utilities, or generates arguments to change the utilities of the disagreeing party with respect to the rejected compromise.

Successful plans that satisfy conjunctive goals totally or partially and their justifications are stored in memory so that they can be reused in similar situations. Failures (impasses) and their failure reason (if one can be found) are also stored so

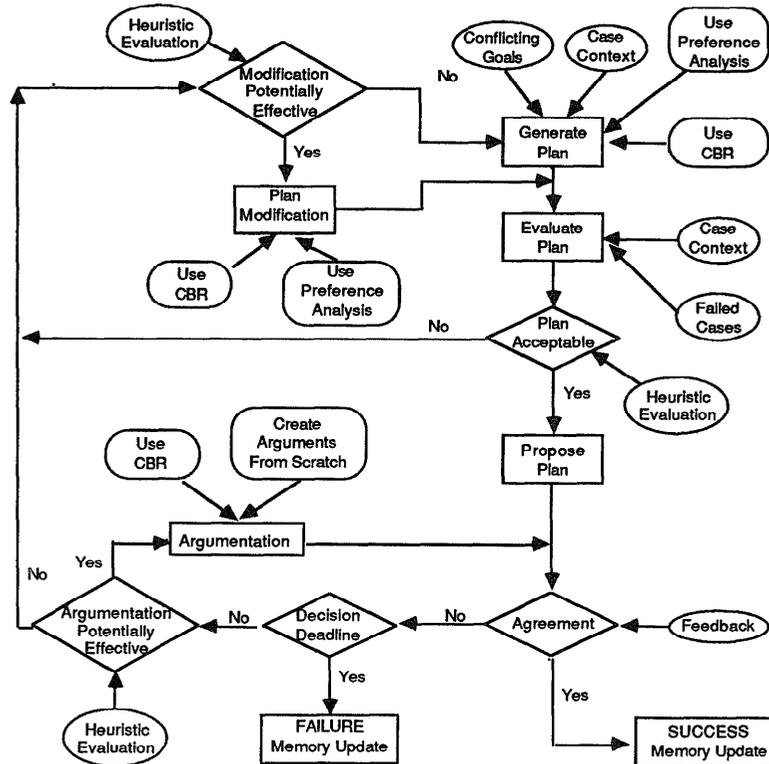


Figure 1: Goal Conflict Resolution via Negotiation

that they can be recalled in situations with similar features to the one where the failure occurred, thus warning the problem solver about potential problems. Unlike other case-based planners (e.g. [Hammond 86]) that only avoid problems that they can anticipate at the beginning of planning, in the PERSUADER, warning/avoidance of problems occurs at each decision point.

The PERSUADER's architecture is particularly suited to negotiation, a task characterized by lack of a strong domain model, many and complex planning steps, and lack of certain or complete knowledge. Case-Based Reasoning (CBR) allows for increased planning efficiency by reusing successful plans and avoiding past mistakes. Re-using successful plans provides a quality of solution unobtainable by traditional planners which are dependent on well-defined goals and operators and strong domain models. Traditional planners (e.g. [Sussman 75]) build plans for each individual goal and then deal (i.e try to avoid) with any interactions as they arise. Such a method is clearly unsuitable for RGC where tradeoffs must be made. CBR can be combined with decision-theoretic techniques, such as Preference Analysis, to improve the quality of solutions.

#### 4. Constructing an Initial Compromise

The PERSUADER uses two methods to construct an initial compromise: CBR and Preference Analysis. We give a brief presentation of the methods. For more detail see [Sycara 87].

CBR consists of the following steps: (a) Retrieve appropriate precedent cases from memory, (b) Select the most appropriate case from those retrieved, (c) Construct a "ballpark" plan based on the selected precedent, (d) Evaluate the "ballpark" plan for applicability to the current case, (e) Adapt the "ballpark" plan to fit the current problem situation.

To retrieve a set of cases similar to the current one, the PERSUADER uses a set of salient features of the domain (e.g., industry, geographical location) as memory probes. An evaluation function based on a prioritization of the features is used to select the best (most similar case). Knowledge is extracted from the solution part (the contract) of the selected case, and adjusted through standard adjustments to form the "ballpark" solution which is further adapted to the current case. A final check for unforeseen problems is then performed through *intentional reminding* [Schank 82] of failures. The conjunction of the solution's features are used as indices to retrieve failures that have the same features as the contemplated compromise. If an associated repair is stored along with the retrieved failure, the planner can apply the repair to the compromise.

Consider, for example, the PERSUADER trying to find a compromise for the Getaway transit transport company and its union. The union wants 15% wage increase, 7% increase in pensions, and no subcontracting. The company wants no wage increase, no pension increase, and unlimited subcontracting. The PERSUADER searches memory for similar past contracts. The most similar are contracts of competitors. Out of those, it selects the contract of Bluchound transit company since its location (Alabama, a southeastern state) is similar to Getaway's (Georgia). This contract provided 12% wage increase, 5% pension increase and unlimited subcontracting.

Next, the Bluchound contract needs adjustment. This is done using known heuristic modifications in labor mediation, namely adjustments with respect to the competitors' position in industry, and area wage differentials between Alabama and Georgia. These adjustments result in 11% wage increase, 4% increase in pensions and limited subcontracting only when extra work is available.

The PERSUADER now adapts the ballpark compromise to the current situation. Checking the financial situation of the company, it finds out that Getaway has suffered 4% losses in the past three years. It searches memory for similar cases, selects the most similar and applies the heuristic used in that case.

```
Searching memory with index
TRANSIT-TRANSPORT and CONTINUOUS-LOSS
3 cases found
Select case2
since it is same area, same company size
Apply heuristic used in this case
Decrease wage increases by half percentage of losses
Wage increase becomes 9%
```

Before proposing the updated compromise, the PERSUADER searches memory to discover potential problems with the contemplated subcontracting language (with indices "failure", "subcontracting language, "limited to extra work"). It retrieves a case where the union had filed a grievance protesting that the company, having extra work, resorted to subcontracting for long periods of time instead of hiring more workers. The arbitrator in that case did not vindicate the union because no time limitation was written in the contract but proposed that the union get a time limitation for its next contract.

```
Searching memory with index FAILURE,
SUBCONT-LANG, LIMITED-EXTRA-WORK
1 case found
Apply repair used in this case
Put time limit in subcontract language
```

If previous similar cases are not available, the PERSUADER uses *Preference Analysis* [Sycara 87, Sycara 88] to find suitable compromises. Preference Analysis is based on *Multi-Attribute Utility Theory* [Keeney 76]. The utilities of each agent associated with each of the conflicting goals are used to rank possible compromises. The PERSUADER derives the utilities of the agents by using utilities of similar agents that it has encountered. If no such past experiences are available, the PERSUADER uses a set of domain-specific heuristics to select the appropriate utilities from a set of utility

functions that it knows about. The linear combination of the utilities associated with an agent's goals forms the *payoff* of the agent with respect to a compromise. The compromise that the planner selects to propose is the one that *maximizes the joint payoff of the agents and minimizes the payoff difference*. This criterion combines maximal gains with equity. For more details, see [Sycara 88].

## 5. Reactive Phase

In iterative planning, such as RGC through negotiation, feedback from the environment may inform the planner that his plan is unacceptable. The PERSUADER has two ways of reacting to negative feedback (rejection of the solution/plan): changing the rejecting agent's evaluation of the plan through persuasive argumentation [Sycara 85, Sycara 87], and modifying/repairing the plan so that it will be more acceptable. Persuasive argumentation is tried first, since, if the objecting agents can be convinced to relax their utilities and accept the compromise, then a successful resolution has been found. If, on the other hand, a rejected compromise is modified/repared, the repair may make it objectionable to agents that had agreed before.

### 5.1. Generating Arguments

The PERSUADER's planning goal during argument generation is to change the belief structure of another agent, the *persuadee*, with respect to a proposed compromise. "Convincing" someone can be modeled as increasing the payoff that the compromise gives him. Hence, the task of a persuader can be viewed as finding the most effective argument that will increase the agent's payoff. Since an agent's payoff can be approximated by a linear combination of his utilities, his payoff can be increased by (a) changing the importance (coefficient) the agent attaches to an issue, and (b) by changing the utility value of an issue. These constitute a persuader's argumentation goals. In labor mediation, the mediator is the persuader and the union or company the persuadee.

The PERSUADER's model of a persuadee's goals is a directed acyclic graph, called a *belief structure*. It is searched and updated during argument generation. The nodes are goals with the associated importance, utility value, and desired direction of change (increase or decrease). The arcs represent the percent contribution of a goal to each of its ancestor goals. For example, an increase in wages contributes to increases in total company labor costs. In contrast, the subgoal of decreasing employment contributes to a decrease in labor cost. The argument, addressed to a union that has refused a proposed wage increase, "If the company is forced to grant higher wage increases, then it will decrease employment" is meant to decrease the importance the union attaches to wage increases by pointing out unpleasant consequences for the union of forcing an unwanted by the company wage increase.

To generate the above argument, the PERSUADER matches the wage goal in the company's belief graph. It propagates the wage increase that the union wants to force to the ancestors of the wage goal (e.g., economic concessions, total labor cost, production cost, profits). Children of these nodes might indicate subgoals that the company can fulfill to counteract the wage increase. Such a counteracting action that violates a union goal that is more important than the union wage increase constitutes an argument that is aimed at reducing the importance that the union attaches to wage increase.

```
Importance for wage-goal1 is 6 for union1
Searching company1 goal-graph...
A increase in wage-goal1 by company1
will result in a increase in econ-concessions1,
labor-cost1, production-cost1
To compensate, company1
can decrease fringe-benefits1, employment1
which violates goals for union1
Importance of fringe-benefits1 is 4 for union1
Importance of employment1 is 8 for union1
Importance of employment1 > importance of wage-goal1
One possible argument found
```

## 5.2. Repairing Rejected Compromises

When a rejected solution needs to be improved, the PERSUADER ascertains from the rejecting agent's feedback the most objectionable goal, the reason for the rejection and the importance the agent attaches to the goal. The rejected goal and reason are used as probes to select impasses with the same stated impasse goal and impasse cause as in the present impasse to supply improvements. If no appropriate impasses can be found, the PERSUADER uses standard heuristics that it knows about.

The PERSUADER's strategy for repair is explanation-based where the explanation (reason for rejection) is supplied by the rejecting agent. This is realistic for complex domains where there is no strong domain model, hence automatic explanation methods are not applicable. The PERSUADER thus uses a combination of similarity-based retrieval (during initial compromise generation) and an explanation-based retrieval (during repair of a failed compromise).

For example, confronted with the company's objection that it cannot afford the proposed "economic package", the PERSUADER recalls impasses that have failed for the same reason and examines the associated repairs.

```
Searching memory with index FAILURE,
ECON-PKGE, INABILITY-TO-PAY
5 impasses found
Select impasse1
since it is same industry, same area
Looking at the repair
"pass the extra cost to the consumer" in impasse1
Since demand for product city transit service
is INELASTIC, repair seems applicable
```

In multivariate planning there are many ways a plan could be modified/repared. A planner seeks not only a plausible

repair but one that with some confidence *improves* the rejected plan. The criterion of plan improvement that the PERSUADER uses is whether the contemplated repair increases the rejecting agents' payoff more than it might decrease the payoff of the agents who have agreed to the compromise. Without an ability to predict which repair has a chance of being accepted, the planner could propose repairs that do not converge.

## 6. Concluding Remarks

We have presented the PERSUADER as a model of RGC among conjunctive goals through negotiation for problems with ill-specified goals, non-enumerable subgoals and unknown associated utilities. The PERSUADER plans iteratively by interacting with the agents, using their feedback in refining and repairing compromises, and in generating persuasive arguments. The PERSUADER plans for labor mediation, a domain full of uncertain knowledge and changing circumstances. Efficient planning is provided by:

- having good criteria to evaluate compromises. In our model, these criteria are provided through Preference Analysis.
- avoiding bad proposals. In our model, this is done through CBR and Preference Analysis.
- having models of other agents' intentions. In our model this is done through a memory for cases.

The integration of analytic and heuristic methods makes the PERSUADER robust and flexible. It does not break down when heuristic methods fail. Moreover, it has the flexibility to use whichever method is more natural to the particular problem solving stage it is engaged in.

## References

- [Cammarata 83] S. Cammarata, D. McArthur, R. Steeb. *Strategies of Cooperation in Distributed Problem Solving*. Tech Report N-2031-ARPA, The Rand Corporation, 1983.
- [Durfee 85] E. Durfee, V. Lesser, and D. Corkill. *Coherent Cooperation Among Communicating Problem Solvers*. Technical Report, Department of Computer Science and Information Science, University of Massachusetts - Amherst, Massachusetts 01003, September, 1985.
- [Fox 84] Fox, M.S. and Smith, S.F. ISIS: A Knowledge-Based System for Factory Scheduling. *Expert Systems* 1:25-49, 1984.
- [Georgeff 84] Georgeff, M.A. Theory of action for multi-agent planning. In *AAAI-84*, pages 121-125. AAAI, Austin, TX, 1984.

- [Hammond 86] Hammond, K.J.  
CHEF: A model of case-based planning.  
In *AAAI-86*, pages 267-271. AAAI,  
Philadelphia, PA, 1986.
- [Keeney 76] Keeney, R.L. and Raiffa, H.  
*Decisions with Multiple Objectives*.  
John Wiley and Sons, New York, 1976.
- [Kolodner et al. 85] Kolodner, J.L., Simpson, R.L., and Sycara-  
Cyranski, K.  
A process model of case-based reasoning in  
problem solving.  
In *IJCAI-85*, pages 284-290. IJCAI, Los  
Angeles, CA, 1985.
- [Konolige 80] Konolige, K, and Nilsson, N. J.  
Multiple-Agent Planning Systems.  
In *AAAI-80*, pages 138-142. AAAI,  
Stanford University, Palo Alto, CA.,  
1980.
- [Rosenschein 86] J. Rosenschein.  
*Rational Interaction: Cooperation Among  
Intelligent Agents*.  
PhD thesis, Department of Computer  
Science, Stanford University, 1986.  
Also appeared as Technical Report STAN-  
CS-85-1081.
- [Sathi 86] Sathi, A., Morton, T.E., and Roth, S.  
Callisto: An Intelligent Project  
Management System.  
*AI Magazine* 7:34-57, 1986.
- [Schank 82] Schank, R.C.  
*Dynamic Memory*.  
Cambridge University Press, Cambridge,  
1982.
- [Sussman 75] Sussman, G.  
*A computer model of skill acquisition*.  
American Elsevier, New York, 1975.
- [Sycara 85] Sycara-Cyranski, K.  
Arguments of persuasion in labor  
mediation.  
In *IJCAI-85*, pages 294-296. IJCAI, Los  
Angeles, CA, 1985.
- [Sycara 87] Sycara, K.  
*Resolving Adversarial Conflicts: An  
Approach Integrating Case-Based and  
Analytic Methods*.  
PhD thesis, School of Information and  
Computer Science Georgia Institute of  
Technology, 1987.
- [Sycara 88] Sycara, K.  
Utility Theory in Conflict Resolution.  
*Annals of Operations Research* 12:65-84,  
1988.