

MARCo: Building an Artificial Conflict Mediator to Support Group Planning Interactions

Patricia Azevedo Tedesco

► **To cite this version:**

Patricia Azevedo Tedesco. MARCo: Building an Artificial Conflict Mediator to Support Group Planning Interactions. International Journal of Artificial Intelligence in Education (IJAIED), 2003, 13, pp.117-155. <hal-00197314>

HAL Id: hal-00197314

<https://telearn.archives-ouvertes.fr/hal-00197314>

Submitted on 14 Dec 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

MARCo: Building an Artificial Conflict Mediator to Support Group Planning Interactions

Patricia Azevedo Tedesco. *Centro de Informática - UFPE. Av. Prof. Luiz Freire, s/n. Cidade Universitária, Recife-PE. CEP: 50740-540, Brazil*
Email: pcart@cin.ufpe.br

Abstract. The emphasis on building co-operative/collaborative environments has brought out the matter of group interactions. This, in its turn, has highlighted the issue of conflicts, inherent to group problem solving. If well employed, conflicts can act as triggers of cognitive changes, and thus help to refine the group's solution to the task. In this paper, we present a computational framework for *detecting* and *mediating Meta-Cognitive* conflicts. The theoretical framework presented here enables a computational system to analyse the ongoing interaction, to detect and mediate conflicts. In order to do so, we consider our model of strategic changes, a model of the group, and the history of the interaction. The objective of the mediation is to suggest courses of action that provoke *articulation* and *reflection*, and thus lead to more refined solutions. In order to diagnose which changes are happening to the group plan we have built a model of strategic operations, which describes what types of changes happen to a plan as well as how changes to one component of the strategy (*beliefs, intentions, ordering relations, context beliefs* and *goals*) impact the other ones. MARCo, our prototype embedding the artificial conflict mediator, shows how our conceptual framework has been put to practical use.

MOTIVATION

People are small group beings. As a matter of fact, we seem to spend most of our daily lives as part of one group or another. In doing so, we have the opportunity to receive feedback, to discuss different ideas, and to get support for our endeavours. It is also the case that we spend a lot of time and effort in making decisions. Decision-making is a complex process, entailing various other subtasks: we need to consider all the different alternatives to our problems, weigh their pros and cons carefully and finally choose the alternative that better satisfies our set of well thought out criteria. This can be quite overwhelming. Thus, research has shown that when making decisions, groups tend to be more effective (they produce better quality solutions) whereas individuals tend to be more efficient (they produce solutions faster).

In this light, we set out to investigate how we could provide support for the consensual decision making process (which is similar to what has been termed collaboration in AIED). Research shows that conflicts are an inherent part of good quality consensual decision-making – they help bring out different insights, ground points of view, and explore different alternatives [Hartley 1997, Lawson and Shen 1998]. A side effect of this process is the occurrence of *reflection* and *articulation*.

Thus, we have explored the idea that it is possible to build an artificial mediator that is able to detect and mediate conflicts, suggesting courses of action that might lead to more refined solutions¹. For the detection phase, an executable model of *Meta-Cognitive* conflict is needed. For the mediation stage, the system needs to know about the changes occurring in the group's plan, group and individual characteristics and the history of the interaction

¹ By more refined we mean solutions that bear indications that *reflection* and *articulation* have occurred. These processes will be discussed in more details later on.

This paper is organised as follows: next section presents an overview of group decision making processes; we then move on to discussing the concept of conflicts, and how they can be used to support group planning interactions. In the following, we present MarCo (Artificial Conflicts Mediator, in Portuguese), our prototype that includes an artificial conflict mediator. MArCo detects and mediates conflicts, with the aim of provoking reflection and articulation in the solution being built. Finally, we describe the results obtained in the experimental study of MArCo, and then point out some directions for further work.

GROUPS AND GROUP INTERACTIONS

Baron, Kerr and Miller [1992] argue that it has been shown that groups can *generally* learn faster, make fewer errors, recall better, make better decisions and be more productive (not only produce more quantity, but also enhance the quality of the production) than individuals working on their own. One explanation for this is what has been called *process gain* – members' interaction brings out insights that none of the group participants had envisaged before. Further explanations for this phenomenon include groups' quicker recognition of incorrect solutions, and their keeping of a better memory of what has happened during the interaction. The authors also enumerate other reasons why groups may be a better option for problem solving: (1) they facilitate higher motivation; (2) they make riskier decisions; (3) involvement in the decision making increases individuals' commitment to implement them; (4) involvement in the decision making facilitates the change of attitude necessary to implement the solution. However, the effectiveness of groups depends on the type of task. According to Steiner [1972] tasks are either divisible² or unitary, and maximising (amount and speed of work) or optimising (quality of work). On divisible, maximising and optimising tasks, groups tend to be a better choice.

Where performance is concerned there are other important factors in group problem-solving: for instance, the quality of communication is fundamental on both simple and complex tasks. Moreover, the presence of well-defined group goals has a big impact on group effectiveness [Weldon and Weingart 1993]. Also, it is not the case that groups are always the best alternative. They are a better option in open-ended tasks. In other cases, individual learning is more effective.

According to [Hartley 1997, Johnson and Johnson 1997] there are several strategies for group decision making (as summarised in Table 1). Groups can, for example, rely on a designated authority, that makes decisions without consulting the group. One of the disadvantages of the method is that it misses out the opportunity of using the group's collective expertise. Alternatively, the authority may take his/her decision after listening to a group discussion. In this case, the quality of the decision can be improved (since group resources were better used) but we might end up with the situation where group members say what they think the authority wants to hear. A decision taken by an expert also fails to involve the group in the process. Moreover, deciding who is the expert not always straightforward.

The other methods of decision-making presented in Table 1 do not involve a central authority. For instance, the group might just average all members' opinions. In this case, opinions of the less knowledgeable may annul opinions of the better informed. Also, although members are consulted, they are not really involved in the decision making process, which may entail that they are not committed to it either. Whenever a group has a lot of decisions to make, it might be a good idea to subgroup, and let minorities take care of decisions separately. On the other end of the spectrum, we have decision making by majority vote. The disadvantage of this method is that it might create an atmosphere of winners and

² Divisible tasks can be decomposed into simpler ones. For example, building software can be seen as a divisible task (different people can build different modules) whereas cutting an orange is a unitary task.

losers within the group, and may foster useless arguments between the parties. Furthermore, members from the losing side may hinder the implementation of the solution.

And last but not least, we have decision by consensus. According to Johnson and Johnson [1997] this is the most effective method, but also takes the most time. Consensus means that everybody in the group agrees with what the decision may be, and is prepared to support it. Thus, members will need time to not only put their points of view across and argue for them, but also to listen carefully while others do the same. One of the issues that must be clear to all participants (when using this approach) is that differences of opinion are a way of gathering new information, generating new ideas, and forcing the group to analyse all the aspects in a solution. In other words, controversy is a tool for reaching better quality solutions.

Table 1 - Strategies for Decision Making (adapted from [Hartley 1997]).

Decision Making Strategy	Advantage	Disadvantage
By authority, no discussion	Speed	Does not use members resources (expertise)
By authority, after discussion	Every member can express their opinion	Members may not be committed to the decision
By Expert	If the expert is really an expert, good quality decision	The most expert member may be difficult to identify
Average members' opinions	Speed	Members may not be committed to the decision
Majority Control	Speed	Minority may feel alienated
Minority Control	Can be useful – in cases where not everyone can attend	Members may not be committed to the decision
Consensus	Members will be committed to the decision	Can take a lot of time, skill and energy

Johnson and Johnson say that consensual decisions are characterised by more conflict, more changes of points of view and longer interactions. These result in members being more confident of the quality of their decision. In this work, we try to support consensual decision making through the use of conflicts that happen during the process. Below we enumerate some guidelines for consensual decision-making (adapted from [Johnson and Johnson 1997, Hartley 1997]):

1. **Avoid blind arguments.** It is important not only to present your point of view as well as possible, but also to listen carefully to others' opinions.
2. **Avoid changing your mind only to avoid conflict and declare goal achieved.** To achieve good quality decisions, the group needs to count on its resources (members' expertise and creativity) as much as possible. Sometimes, controversy plays a great role in bringing group resources to light.
3. **Avoid conflict reducing procedures** (like tossing a coin).
4. **Try to involve as many people as possible in the discussion**, getting them to elaborate on their points of view, and bringing as much insight to the solution process as possible.
5. **Do not assume there are winners and losers.** If the group cannot agree on a solution, look for the next best alternative.

6. *Ground your points of view.*

In our research, we are more interested in the quality of the solutions than on their efficiency. Thus, in building MArCo, we have concentrated in supporting people trying to solve a problem via *consensus*. When the decision making process also involves taking the actions necessary to implement the solution, it is similar to what has been termed collaboration (e.g. [Dillenbourg 1994]). In MArCo we assume that is the case, and focus the mediation process in fostering the situations and a symmetric interaction amongst group members. When mediating this process, we have focused on guidelines 1, 2 and 6 above. The structure of the dialogue games used in the interaction with the prototype helps supporting guideline 1 by making sure that people have at least to listen to others' arguments before being able to agree or disagree with something. Mediation strategies employed by the system at this point encourage people to elaborate as much as possible on their points of view (and their changes), thus supporting guidelines 2 and 6. The general goal is to help participants bring as much as possible into the discussion, as well as to get people to listen to others, and thus reflect on the problem solving process. On future extensions of the prototype we intend to investigate (through the use of social roles played by the group members) how we can ensure that all members are involved in the discussion, and thus support guideline 4. By providing the system with a good model of the task we will be able to support guideline 5 as well. At this point in time, guideline 3 is being supported implicitly, since the environment provided in our prototype tries to promote situations where conflicts would happen naturally.

One of the indications of good quality consensual (and, consequently collaborative) decision-making/problem solving is the presence of *reflection* and *articulation* during the interaction process. Self, Karakirik, Kor, Tedesco and Dimitrova [2000] define *reflection* as being the process of considering our own thought processes, problem solving strategies, knowledge and skills. In short, it is the process of thinking about our own thinking. The authors also point out, in the same lines of Dewey's definition [1933], that *reflection* needs to be motivated. We need to feel we are achieving something either now or in the future (for instance, in the discussion that follows) to engage in the process. Boud, Keogh and Walker [1985b] argue that reflection is the skill of learning in which we tend to be most deficient. Hence the need to facilitate it.

One of the ways of helping people think about their standpoints on various issues is to encourage them to articulate their views. *Articulation* is defined in [Self et al. 2000] as being the act of verbalising the thought processes and problem solving activities that are either happening now or have happened in the past. Such a process has three benefits: (1) the link of language and thought may mean that the act of making something explicit actually helps developing the thought process involved; (2) articulating something may bring its weaknesses and/or misconceptions to the inspection of others (including the articulator) and thus hopefully generate a fruitful discussion on the issue; and (3) being able to provide articulate explanations is worthwhile in itself. The authors then argue that in the case of computer based learning systems there is a necessity to provide support for both *reflection* and *articulation*. The computer support mentioned there is actually a two-way process (as is the case with MArCo). A computational system will be able to support *reflection* provided that it has some feedback (in the form of *articulation* of what is happening) from the user.

ON CONFLICTS AND CHANGE

In this section, we argue that conflicts are not only pervasive to group problem solving interactions, but are also triggers of cognitive changes. Thus, they can be used to provoke *reflection* and *articulation* on the solution process, helping the group to generate more refined solutions. We will begin this discussion by presenting the concept of conflict from different perspectives and then show how they correlate to cognitive change.

Müller and Dieng ([2000], p.1) say that, “If two or more parties are doing something together, there is a high potential for conflict”. Conflicts normally arise in a situation involving at least two parties [Mack and Snider 1971], and necessarily include mutually exclusive, incompatible and opposed values. As evidence shows (e.g. [Easterbrook et al. 1993, Robbins 1998]) conflicts happen through dialogue and are not only inevitable to the group interaction but may also be beneficial to it.

According to Robbins [1998], we need to look at group problem-solving interactions from an *interactionist* point of view. *Interactionism* takes the stand that conflict is not only a positive force in a group, but is also absolutely necessary for a group to perform effectively. Consequently, the advocates of this view also encourage group managers to maintain an ongoing minimum level of conflict, to help keep the group self-critical and creative. Walton [1987] emphasises that conflict, in its different forms, often brings benefits for both its participants and the social system in which it occurs. This is due to three main reasons: (1) a moderate level of conflict may increase motivation and energy; (2) conflict may promote innovation because it highlights different points of view; and (3) participants may improve their understanding of their own standpoints, since they need to articulate their views, providing supporting arguments for them. The author goes on to point out that the presence of a neutral third-party could be useful to help people overcome their fears and suspicions, and turn conflict into a more productive situation. This is one of the issues we will investigate in this work.

Whilst it is somewhat easy to pinpoint situations where we engage in conflicting interactions (e.g. a war, an argument about where to go for lunch, a lawsuit), it can be quite hard to characterise the conflict process. Since conflicts have been found in so many levels, there are several definitions that can be found in the literature. For example, the Pocket Oxford English Dictionary says:

Conflict – n. 1.a. A state of opposition; b. fight, struggle; 2. Clashing of opposing interests, etc.

There is a lot of research in Psychology, Management and Social Sciences that tries to define the *process* of conflict and analyse its effects on group interactions. Although these definitions generally have a very broad character, they can help us to identify some process characteristics, and thus help us to arrive at a computational definition of it. In [Robbins 1998], for example, we find conflict being defined as a *process* that begins when one party perceives that the other has frustrated (or negatively affected) some concern of his/hers. Their definition approximately expresses our view. To us, negatively influencing (or frustrating) our concerns entails that the other party is trying to convince us of a point of view (that contradicts our own) about a particular topic.

Some authors (e.g. [Müller and Dieng 2000, Castelfranchi 2000]) present definitions of conflict that are divided into two broad categories:

7. **Conflict among Multiple Parties** (and in the specific case when there are only two parties involved, this has been called *dyadic* conflict).
8. **Conflicts within Oneself** (called *cognitive dissonance* by the Encyclopaedia Britannica).

In our research, we focus on multi-party conflicts. Also, the process of *cognitive dissonance* is considered here to be rather a *disagreement* between one’s previously held beliefs and newly arrived information. Nevertheless, we do consider the case where we have conversations (or conflicts) with ourselves (for instance, weighing pros and cons of a given issue, as considered in [Burton 1998]).

Within the Social Sciences, conflicts have also been analysed from different perspectives. Analogously to the case with Psychology, the earlier literature concentrates on the dysfunctional side of conflict, and then evolves to analyse their benefits to the group. For example, Coser [1956] says that conflict pushes the social systems towards innovation and creativity. Through conflicts we can revitalise the existent norms, and create new ones.

Piaget, in his earlier writings, had already argued that peer interaction is important to cognitive development – because it fosters conflict. These ideas were later elaborated by Doise and Mugny [1984]. In their work, they set out to investigate how social interaction (i.e. interaction with peers) affects the

individual's development. The basic idea is that cognitive development is boosted when peers are faced with different perspectives on the solution of a problem (the *centrations*) and arrive at a more refined solution. In other words, their argument is that, when two peers have contradictory viewpoints about the problem at hand, they have two ways out of the conflicting situation:

By conforming to the other member's opinion. This can be seen in situations where there is an authority figure on the pair. This conformity illustrates the cases where members are looking for a way of eliminating the conflict; or

By integrating the contradictory viewpoints. If that is the case, they show that there is cognitive change. However, Doise and Mugny do not explain how the resolution of the conflicting situation takes place. Joiner's work, discussed below, tries to fill in this gap.

Crook [1994] argues that the presence of a peer during the problem-solving (learning) process encourages people to articulate their ideas. This stems from the process of conflict, inevitable in such situations. According to him, conflicting interactions prompt dialogue moves of justification and negotiation, which, in their turn, provoke cognitive changes. He goes on to argue that this is a line of reasoning closely related to Piaget, who said that, in peer interactions, we need to look at "the shock of our thought coming into contact with that of others" ([Piaget 1928]). The idea is that demands for justification that appear in a conflict will provoke *reflection* on our own thought. Furthermore, the fact that we encounter conflicting ideas precipitates cognitive restructuring [Piaget 1928]. This is the rationale of the building of our conflict model.

Joiner [1995] argues that certain ways of solving some types of conflict will lead to learning. Moreover, since the theory of socio cognitive conflict does not really elaborate on how the conflict resolution process provokes *cognitive change*, Joiner investigated the process of resolving inter-individual conflicts. He proposes a dialogue model of the processes relating to conflict resolution, and attempts to explain the cognitive changes that happen in joint planning. Joiner's dialogue model is based on the idea of focus. To model inter-individual conflicts he defines three types of *foci*:

Dialogue focus: relates to the representation of objects and events mentioned explicitly in the discourse and that are relevant to the current topic of conversation.

Task focus: refers to the knowledge people think might be relevant to what is being talked about.

Task representation: is the representation of the objects, beliefs and actions the participant has about the task. Items in Joiner's task representation have an associated degree of confidence.

Within Joiner's framework, conflicts are detected whenever participants perceive contradictions between their task and dialogue foci. Joiner claims that the resolution of these contradictions might lead to learning. We try to take this a step further, presenting a model that considers other levels of cognition.

Several authors have looked at the occurrence of conflicts in collaborative (or co-operative) interactions aiming at analysing the quality of these. For example, Mühlenbrock and Hoppe [1999] claim that their framework can be used for the assessment of group interactions. Their approach is based on a protocol of user actions (e.g. aggregation, conflict creation) in a shared card workspace. The authors also claim that although the assessment of the interaction is their main goal, the system could be extended by rules to remedy frequent conflict situations, unbalanced co-operations, and so on. Their framework, however, only takes the concept of conflicts to the domain level, missing out on using the other cognitive levels.

Soller, Linton, Goodman and Lesgold [1999] present a series of studies that focus on analysing the quality of collaboration in groups using computer mediated communication. The experiments in question were done using the Collaborative Learning Interface and the Collaborative OMT editor, which are both sentence-opener based interfaces. The sentence openers used were divided into the categories of the Collaborative Learning Conversation Skills Network (presented in [Soller et al. 1998]), namely: Active

Learning, Conversation, and Creative Conflict. Their studies have shown that skills found in the Creative Conflict category were significantly present in productive groups.

Quignard and Baker [1999] aim to model the cognitive changes produced as a result of argumentative interactions. In order to do so, they have established some design conditions and modelling constraints (e.g. the concepts of the task have to be understood by all the participants; participants' knowledge need to be collected just before and immediately after the interaction) to build a computer-mediated environment that can help provoke argumentation. Their approach considers mainly belief conflicts. The results of the reported experimental studies imply that argumentative interactions can create critical awareness of the problem-solving domain in the students. As a result, they can lead to learning. Having presented a brief overview of the research in conflicts, we will now move on to defining what the conflict process means to us, and how we have implemented the computational conflict model.

What do we mean by Conflict?

Before we actually design a computational model of conflicts, however, it is important to define what we mean by a conflict. For the purposes of this work, a *conflict* is the final stage of a three-phase process:

In the first stage we have a ***difference of views***: when the agents have got different views about a certain topic, but have not yet communicated them to their partners in the discussion.

After that we have a ***disagreement***: where the agents inform each other about their inconsistent views, but a discussion over them does not necessarily follow.

Finally, we have a ***conflict***: when the agents involved not only inform their partners about their inconsistent views, but also try to convince the others about their own points of view.

The idea we are trying to convey here is that by trying to convince our partner in the conversation of our own point of view, we have to begin by elaborating on our own mental state. And this can only be achieved through *reflection* (both on our own mental state and that of our partner). Secondly, to make ourselves understood (and to sound convincing) we also need to *articulate* our points of view. And, in doing so, we might provoke cognitive changes both in our own minds and in our partner as well, and thus lead to more refined solutions to our joint problems.

The Categorisation of Meta-Cognitive Conflicts

By distinguishing between disagreements and conflicts we want to make sure that interventions will only be made in situations with more potential for *reflection* and *articulation*. This line of reasoning will play an important role on how the system decides if/how to intervene in conflicting situations. Let us now present our categorisation of the conflicts happening in a group planning situation. Since the categorisation and conflict model have been presented in full detail in [Tedesco and Self 2000a, Tedesco 2001a], we will just briefly describe them here.

Non task-related or social conflicts: are the ones related to the social roles and positions of the members within the group.

Task-related conflicts can be further divided into:

Belief conflicts: are about the domain under discussion. Here we are considering conflicts in the two lower levels of the DORMORBILE framework [Self 1995]. In DORMORBILE, the agents' knowledge is described in four levels. The *Domain level* concerns facts related to the actual domain of discussion. Next, we have the *Reasoning level*, that describes the process of problem-solving on the domain. In this level, Self defines the concept of *reasoners* - functions that when applied to domain facts, will produce other facts in the domain. The *monitors* that compose

the *Monitoring* level are functions that when applied to the *Domain* and *Reasoning* levels will produce new domain facts and *reasoners* (thus, monitors help us devise strategies and heuristics); and *Reflection* (refers to choosing amongst the various possible strategies to solve a problem, as well as to creating new possible heuristics). Conflicts at the *Domain* and *Reasoning* levels have been extensively investigated (for example, see [Gärdenfors 1992]). Some authors (e.g. [Joiner 1995]) talk about conflicts in the focus of the dialogue. In our framework, these would be classified as *Belief Conflicts* as well. In this work, we take that model a step further, modelling conflicts on other levels of cognition (according to DORMORBILE).

Contextual conflicts: these relate to defining what exactly is the problem being solved, and what are its restrictions and criteria for success. From our point of view, *Contextual Conflicts* are a specialisation of *Belief Conflicts*.

Meta-Cognitive conflicts: these correspond to the two upper levels of the DORMORBILE framework (*Monitoring* and *Reflection*) and are about how to build strategies to solve a given problem, and how to choose among different strategies. They can be subdivided into the following categories:

Reflector Conflicts: can be exemplified/triggered by the questions “How can we set about building our strategies?” or “How can we choose between two apparently similar strategies?”

Intention Conflicts: refer to the actions we intend to do in order to complete our plan. The question “Which steps make up our strategy to solve the problem?” gives us a starting point to understand what is involved in them.

Ordering Conflicts: are those that happen when we argue about the organisation of our intentions in a plan.

Goal Conflicts: these can be divided into the categories below.

Goal Definition Conflicts: sometimes, because the model of the problem is not very well defined, members of the group get into conflicts about what is the goal they are trying to achieve.

Goal Achievement Conflicts: can be exemplified by the questions: “Have we achieved it?”, or, “Is our plan ready?”.

The Computational Model

Since the first step when mediating a situation is to detect its occurrence, we needed to have a computational model of conflicting situations to be able to eventually build an artificial mediator that can recognise these situations. Thus, we have built a *Belief-Desires-Intentions (BDI)* model. The choice of a *BDI* approach was mainly due to the mapping that can be drawn between the levels of the DORMORBILE framework and the mental attitudes, and also to the parallel we could draw with the various models of multi-agent interaction that can be found in the literature.

In this work, we wanted to model the cognitive states of intelligent agents that are capable of joint planning. The idea was to be able to build an executable model of their reasoning, and thus be able to identify the cognitive changes that happened along the problem solving process. In order to do so, we have constructed a BGI (beliefs, goals, intentions) model that considers the following mental attitudes: (1) *cognitive attitudes*, represented by beliefs in the domain, contextual beliefs – those that represent constraints of our problems (for example, if we are going shopping and need to be back by five o’clock, we cannot intend to do something that will take us until seven) and reflectors; (2) *affective*, represented by

the agents' goals, that can be seen as a consistent set of the agents' desires; and (3) *conative*, represented by the agents' commitment to take the actions necessary to achieve their goals.

In our models of *meta-cognitive* conflicts and *strategic changes* we have used a propositional logic (as used in [Galliers 1989]), together with the operators presented in table 2 below. The role of the formalisation in this thesis is to provide us with a way to make our models executable. In other words, the formalisation provided us with the basis for implementing MArCo's reasoning mechanism.

In this work, we make a distinction between *conflicts* and *mutual conflicts*. Agent x is in conflict with y about s (Conflict (x y s)) if x believes that y has got a different view about s , and intends to convince y to adopt his/her view. However, x 's beliefs about y may be wrong (i.e. x might have misunderstood y and there is not actually a conflict), and thus the conflict situation might generate confusion. On the other hand, *mutual* conflicts (corresponding to the notion presented in the previous section) have more potential for *reflection* and *articulation*. Besides presenting a better picture of possible conflicting interactions, this distinction between conflicts and mutual ones also has the advantage of making it simpler for us to extend the two-party case to the n-party case.

Table 2 – Operators Used and their Meanings

Operator	Meaning
Bel	Indicates agents' beliefs. For example, $Bel(x, a)$, indicates that agent x believes sentence a .
Intend	Represents agents' intentions to do a certain action. $Intend(x, action)$ means that agent x intends to do $action$.
Goal	Represents agents' goals. $Goal(x, s)$ means that agent x has a goal of achieving s .
Context	Represents agents' beliefs about the constraints of the problem.
Reflector	Represents agents' reflectors (as explained in the previous section). Thus $Reflector(x, a)$ means that x will use a to reason about his/her planning.
◇	Represents the modal operator eventually, as defined in [Galliers 1989].

So, if we consider *Mental_Attitude* to be the set of mental attitudes present in a task related conflicts (i.e. $Mental_Attitude = def \{Bel, Intend, Goal, Context, Reflector\}$), and *mental_att* to be a generic mental attitude such that $mental_att \in Mental_Attitude$, then a generic conflict between agents x and y about a sentence s is defined as:

$$Conflict(x\ y\ s) = def (mental_att(x\ s) \wedge Bel(x\ mental_att(y\ \neg s)) \wedge Intend(x\ \diamond(mental_att(y\ s))))$$

Thus, a generic mutual conflict between agents x and y about sentences s and t (when $t \wedge s$ is false) is defined as:

$$Mutual_Conflict(x\ y\ s) = def Conflict(x\ y\ s) \wedge Conflict(y\ x\ t).$$

This relation between s and t is valid for all other definitions of *Mutual-Conflicts*. Just as an example, a Belief Conflict is defined as:

$$Bel_Conflict(x\ y\ s) = def (Bel(x\ s) \wedge Bel(x\ Bel(y\ \neg s)) \wedge Intend(x\ \diamond(Bel(y\ s)))).$$

A Mutual Belief Conflict is defined as $Mutual_Bel_Conflict(x\ y\ s) = def Bel_Conflict(x\ y\ s) \wedge Bel_Conflict(y\ x\ t)$.

At this point, some clarifications on the nature of the mental attitude *intention* are needed. According to Rao and Georgeff [1995] intentions are what move an agent to act. They represent the current course of action and thus guide the agents' future activities. We follow this same line of reasoning, and thus represent the steps (what is to be done in a plan) as intentions. So, when we say that agent x intends to change the attitude of y in what concerns some sentence, we mean that x is committed to engage into some action (for example, reflecting upon and articulating his/her arguments) that will lead to y 's changing his/her attitude.

Since conflicts are triggers of cognitive changes, another way of improving the efficacy of the mediation process is to try and identify the changes happening to the group plan, and try to guide the solution process towards more refined decisions. In order to do so, we have provided MArCo's mediator with a model of strategic changes. The changes model is based on the literature about belief changes and extends the current planning operations models (e.g. [Grosz and Kraus 2000]) to include all the changes we have found in MArCo explicit, and to clarify the interconnections that exist among the strategic components. Since the model has been described in full detail elsewhere [Tedesco and Self 2000b, Tedesco 2001a] we will not present it here. Just to give the reader an idea, Figure 7 presents a classification of the possible group plan operations.




Figure 1- The Classification of Strategic Operations

Having presented the two major models in MArCo, we will now move on to present other theoretical considerations that were made during the design of the prototype.

THEORETICAL ISSUES UNDERLYING THE DESIGN OF MARCO

The Chosen Domain

As we have discussed previously, the main goal of this project is to investigate how *Meta-Cognitive* conflicts can be used to support group planning interactions. Consequently, the choice of a planning domain seemed to be the obvious one when designing the prototype. We also wanted a domain that provided graphical representations of plans, thus making the dialogue analysis task somewhat simpler.

Thus, we chose to use Operational Research (OR) as our implementation domain. According to Winston [1991], the field of OR applies a scientific approach to planning. It aims at determining how to best design/operate a system, which usually involves the allocation of scarce resources. In other words, the field is about planning the design and use of a system. Within OR, various different techniques can be used to accomplish such a task. As remarked by Yang [1997], PERT graphs are a good, easy to understand representation of a plan structure. Consequently, in MArCo we have used PERT/CPM graphs to represent our strategies. For our purposes, the use of PERT/CPM graphs has yet another advantage: it is a simple technique. As a result, MArCo's users will not have to spend too much time learning it. They can concentrate on the planning interaction instead. In this section, discuss PERT/CPM graphs in more detail.

PERT/CPM graphs

Both *CPM* and *PERT* are used to aid in the scheduling of large projects that consist of several interdependent activities. According to Hillier and Lieberman [1995] a PERT type system is used to aid in planning and control. Whenever the duration of the activities is known for sure, *CPM* can be applied to determine the length of a project, as well as how long each of the activities in the project can be delayed without delaying project completion. When the duration of the activities is uncertain, then *PERT* is used to calculate the probability that a project will be completed by a certain deadline. Both techniques were developed at the end of the 1950's and have been used in a variety of problems.

At this stage, all the problems in MArCo have activities with known durations. This is because we do not want to make users spend much effort in the calculations related to *PERT*. We would rather have our users concentrating as much as possible on the planning aspects of the problem at hand. In order to be able to apply either of the two techniques we need a list of the activities that compose the project along with their durations. In typical PERT/CPM problems the predecessors of each activity are also given as part of the question. In MArCo, this list has been taken out of the problems, to provide more scope for *reflection* and better simulate a real-life strategy.

Let us now show an example of a planning problem, together with its graphical representation. Figure 2 represents a typical problem, and is captured from MArCo:

Problem Statement
_ □ ×

Widget Co. wants to introduce product P. P is produced by assembling one unit of product P1 and one of P2. Before the production begins on either P1 or P2, materials are purchased and workers are trained. Before P1 and P2 can be assembled into P3, P2 must be inspected. A List of Activities and their durations is given below. Draw a project diagram for this project. Find out if the project can be completed within 40 days.

Activity	Duration
A - Train Workers	6 days
B - Purchase Materials	9 days
C - Produce P1	8 days
D - Produce P2	7 days
E - Inspect P2	10 days
F - Assembly P1 and P2	12 days

Figure 2 - A Sample Problem.

To give you a better idea of how a network model of this problem would look like, let us present its graphical solution (this graph can be found in [Winston 1991] p. 400).




Figure 3 - A Project Network for the Problem in figure 2.

In this section, we have briefly presented PERT/CPM graphs. In the following, we will discuss the role of the mediator in full detail.

The Role of the Mediator

MArCo's in-built mediator aims at facilitating *reflection* and *articulation* on the plan being built. More specifically, the aim is to "break up" conflict situations, prompting the participants to think about their mental attitudes before the conflict escalates to a personal level. There is no requirement that people follow the mediator's advice. Even if they choose not to do so, we believe that the pause in the interaction might be beneficial to the *reflection* process.

According to Moore [1987], mediation is the intervention by an acceptable, impartial and neutral third party in a dispute or negotiation. The mediator should have no authority, and should assist the disputant parties in voluntarily reaching their own mutually acceptable settlement of the issues in dispute.

Silbey and Merry [1986] (as cited in [Burton and Dukes 1990]) have identified two main types of mediation: "agreement centred" and "relationship centred". In the first case, the mediator's main goal is to reach agreements tolerable to all parties involved (this is the approach adopted by Kolodner and Simpson's MEDIATOR). The second type of mediation focuses on creating relationships that will allow for the development and implementation of solutions. This is the line of mediation adopted in MArCo.

There is a lot of work concerning conflict management and its importance in organisations (e.g. [Tjosvold 1993, Walton 1987]). In this literature, we can find two main approaches: (1) *conflict resolution*, which aims at eliminating the original differences; and (2) *conflict control*, which tries to reduce the negative consequences of conflict. In MArCo, we have adopted *conflict control* approach. Thus, the mediator will not ask participants to settle for agreements, even though sometimes they will need to be reached for the sake of completing a problem. This is due to the fact that, for the purposes of this work, the process of conflict control and the *reflection* and *articulation* it entails are much more relevant than the conflict resolution itself.

Kolodner and Simpson's MEDIATOR [1989] is an example of an AI work that fits under the *conflict resolution* trend. Its domain is common-sense advice for the resolution of resource disputes. The MEDIATOR attempts to use case-based reasoning to derive a compromise agreement.

There are other systems that use a mediator to "coach" students into considering different alternatives in their solution processes. One example of these is the work of Constantino-González and Suthers [2000] where they have built a collaborative learning environment (Coler) that aims at teaching Entity-Relationship database modelling. Coler's coach tries to identify significant differences (in what concerns domain knowledge) between the private and shared workspaces. It also monitors the amount of participation from the student, but not the quality of their contributions (differently from what we are trying to do in MArCo). The coach decides on how/when to intervene according to a decision tree that lines up events happening and intervention strategies.

It is fair to say that MArCo has a different character from most of the coach systems, since we are not trying to teach any particular domain. The idea here is to use what has happened in the interaction and what can be inferred from it (for example, the history of the dialogue, group and individual models) to get people to think about the solution process. One way of doing this is to point out areas that people can reflect on and/or alternative courses of action. Upon detecting a conflict, the mediator takes one (or more) of the following courses of action:

tells the participants that a conflict has been detected, and asks them if they can offer alternatives to their current views.

checks if the group has already defined the goals and context of their strategy. If not, the mediator reminds the group of the fact, and asks them whether they would like to do so. This is based on the idea that context/goals are determinants of a strategy. Consequently the group can probably achieve more refined solutions if these components are explicit in the solution process.

checks if any of the participants contradicting what was previously held in his/her model. If that is the case, the mediator asks the participant to elaborate on his/her change of mind. The rationale for this mediation strategy is that there must be a reason for such a change of mind and that the group would probably benefit from the explanation, whether it is from receiving it or from giving it;

during the preliminary experiment, there were some examples of group agreements that were immediately followed by a member making an utterance that would contradict it. Based on the same rationale for (3), the mediator will check if that is the case, and if so, will ask the group member to elaborate on his/her change of mind.

At the next level of refinement, MArCo's mediator should be able to diagnose which strategic change is taking place at a given conflicting interaction, simulate the effects of each members contribution, and then advise the group to think about the change that would result in a more refined solution. This is the issue we are pursuing now. During the discussion of the experimental studies, we will present users' perceptions of the effects the mediation had on their interaction and their suggestions for further improvements.

The Model of the Group and of the Individuals

In order to be more effective, MArCo's mediator needs to draw on group's characteristics and knowledge. But although we need group models to cater for collaborative/co-operative interactions, we cannot forget that groups are made of individuals that need to taken care of as well. In MArCo we have both models – empty when the interaction starts and updated as it proceeds.

Even though this situation is changing, it is still true that the vast majority of CSCLEs (Computer Supported Collaborative Learning Environments) do not have a well-defined model of the group interaction. Hoppe [1995] says that we could think about different ways of caring for the social aspects of learning. We could, for example, embed our individual systems into social contexts, or build software that could somehow follow a social interaction. With the demand for collaborative/co-operative systems, especially in geographically distant situations, the second alternative needs to be better explored. This is the case in MArCo.

Hoppe [Hoppe 1995, Hoppe and Plötzner 1999] tries to address the second question, to an extent. His work uses individualised models to inform and parameterise human-human co-operation without really monitoring the interaction. In MArCo, we need to monitor the interaction, in order to pick up conflicts as they happen, and to intervene more effectively. In Tedesco [2001a, 2001b] a review of the streams of research in group modelling is presented, and MArCo's group and individual models are discussed in length. In this paper, we will only cast a brief look at MArCo's models.

MArCo's Group model (GM) was built based on Paiva's recommendations [1997]. From her point of view, a group model should have the following components:

1. **Group beliefs:** are inferred from the group actions as well as from the ongoing discussion.
2. **Group actions:** are the ones the system assumes to have been done by the group as a whole, or, alternatively, by an individual that cannot be identified.
3. **Group misconceptions:** are those diagnosed from the group actions.
4. **Group differences:** may be a way of settling the negotiation process. For instance, many times people tend to accept other's suggestions based on their expertise. They can be obtained by comparing the individual models, and used as a way of providing feedback to the group. Hoppe [Hoppe and Plötzner 1999], has used a few basic parameters (*competence, has difficulty, competitive*) to differentiate between learners, and thus provide better feedback.

5. **Group conflicts:** are used in Paiva's model to point potential triggers of cognitive changes out to the system. From our point of view, these conflicts are *differences*, and although we agree that they are good pointers for a system to act upon, they might not have as much potential for *reflection* and *articulation*.

So, let us now discuss MArCo's Group and Individual Models. Both models have been based on what is discussed in the literature, and have been adapted to cater for the knowledge the system needs to have when mediating for *Meta-Cognitive Conflicts*. Thus a typical Individual Model (IM) in MArCo consists of the following sets of mental attitudes:

- ❖ **Domain Beliefs:** these correspond to the first two levels of Self's DORMORBILE framework [Self 1995]. In our PERT/CPM case, domain beliefs will appear mostly when we are discussing the context of our problem, the ordering in which things should be done, and the justifications for that ordering. Because *contexts* are one of the defining issues in a strategy we are marking the beliefs relating to the context with "context". So, examples of domain beliefs of Agent X³ in the problem presented in figure 2 are as follows:

Belief_domain(X, parallel(produce P1, produce P2)), meaning that X believes that P1 and P2 can be produced in parallel.

Context(X, before(purchase_materials, (produce P1, produce P2))), meaning that from the problem restrictions, X believes that the purchasing of materials must be done before the production of P1 and P2.

- **Intentions:** represent the actual steps of a strategy. For instance, when a pair is cooking dinner together and one proposes to chop the vegetables, this is represented in their cooking dinner strategy as the intention of chopping vegetables. Since MArCo is an artificial environment, we do not have the actual execution stage (i.e. we do not really chop the vegetables), but every time we propose something, the system interprets it as our intention to do it.
- ❖ **Goals:** are the other defining factor in a strategy (together with the *context*). They determine the use of our reflectors, in Self's terminology, and as such are defined separately. Furthermore, if the system realises that the group has not yet defined their goal, it will encourage them to do so, in an attempt to help the group make more objective decisions concerning the strategy being built. If we consider the problem in figure 6, a goal of agent X would be represented as follows:

Goal(X, finish(project, 20)), meaning that X would like to finish the project at time 20.

- ❖ **Reflectors:** can be seen as meta-attitudes. In other words, they represent things we do before we are actually able to talk about our own strategy, or, alternatively, they represent our thoughts about how we should set out to build/choose our strategies. For instance, suppose that for the problem presented in figure 2, the group had come up with two different ways (*strategy₁* and *strategy₂*) of solving the problem. Now, imagine that agent X had suggested that the group should choose *strategy₁* since it is more coherent with their goals. Such a reflector is represented in MArCo as:

Reflector(X, strategy₁ ⇒ achieved(goal))

Every time a group member makes a contribution, his/her IM is updated. Whenever the contribution represents a group agreement, the same happens to the GM. MArCo's GM consists of the same sets of attitudes as the IMs, plus the roles played by the group members during the dialogue. The following roles

³ By agent X we mean any of the human group members.

are defined in MArCo: *leader, follower, outspoken, shy, reflector, actor*. Roles are attributed to members according to the number and type of contributions they make.

Recording the roles played by the group members serves a twofold purpose: first, it can serve as an indication to the mediator on how to pose its interventions (for instance, after learning that one of the participants is a reflector, it could ask him/her for his/her opinion about something next time the opportunity arises); secondly, it can serve as an indication of how well the collaboration/co-operation process is going, as well as of how effective the mediator's strategies are.

Having looked at the importance of group models, and their use in MArCo, let us now move on to analysing the model of the interaction. Although the dialogue is not the focus of the work, it does play an important role in our framework.

The Interaction

It is possible to find several examples of systems that, like MArCo, need to monitor the interaction, but want to do so with a minimal analysis of its contents. In some cases, systems use the form of the dialogue to enforce some desired characteristics (as is the case with [Pilkington et al. 1992], where the authors have built a system to teach argumentation skills). There are three main categories of such dialogue systems that relate to MArCo:

Those using Sentence Openers: In general, these are systems concerned with analysing the collaboration process. One of the first uses of Sentence Openers was by McManus and Aiken [1995] who implemented the "Group Leader", embedded in an ICLS (Intelligent Collaborative Learning System) which aimed to teach students collaborative skills. Students' utterances were made by using sentence openers coupled with their contents. Baker and Lund [1996] have studied the usage of two interfaces for their C-CHENE system: one where students could use a chat tool to communicate freely while collaboratively solving physics problems, and another where students needed to use a set of Sentence Openers when communicating. They found that the use of the second interface helped students focus on the task level.

Those using Dialogue Games – Dialogue Games was a term used by Levin and Moore [1977] to characterise some patterns in Natural Language Dialogue. Dimitrova's STyLE-OLM [2001] and Burton's Clarissa [1998] are examples of systems that took this approach.

Those relying on the analysis of the actions rather than on the dialogue: One example of this approach is Mühlenbrock and Hoppe's system [1999, 2001] that works as a plug-in to their cardboard interface. In that case, the system analyses only the actions (e.g. changing cards places in the puzzle, removing cards) of users working collaboratively to solve a cards puzzle.

In her simulation based data, Pilkington [1999] has observed that the two main types of DGs seen to provoke *reflection* were: (1) inquiry (where we have a player asking the other questions); and (2) the one where the roles of challenger and defender take place (here we are talking about the case where players are challenging each other views). As we will discuss below, these are two of the main games adopted in MArCo. The DGs chosen in our system aim at making group planning interaction possible, and also at enabling explanation requests and offers. The idea in our interaction model is that different games deal with the various strategic components. The moves that compose the DGs in MArCo are drawn from the DISCOUNT scheme [Pilkington 1999], and are mapped on to the mental attitudes that compose a strategy.

Since it is not an active partner in the dialogue, MArCo only needs to keep track of the moves used in the DGs and the focus of the game. When playing the games in MArCo, participants are presented with some Sentence Openers every time they choose a dialogue move. This was included in the dialogue to make it look more "natural".

There are four DGs in MArCo: (1) **Inquiry**, played when a user wants some information about a topic concerning the current plan; (2) **Statement**, when a user wants to inform the other about his/her beliefs; (3) **Proposal**, played when a user wants to add steps to the plan; (4) **Suggestion**, played when they want to discuss meta-entities in the plan. For instance, when discussing goals and context (the two entities that define a plan) users play the Suggestion game. In MArCo, the usage of the different DGs not only helps the mapping of users' utterances onto the BGI attitudes that compose a plan but it also helps users to become aware of the different components of our strategies. The correspondence between the DGs and the strategic components is as follows:

- ❖ **Statement Game**: used to discuss our beliefs in the domain. In the implementation, the moves of this game are mapped onto domain beliefs, denoted by *beldom* in our prototype.
- ❖ **Proposal Game**: this game is used to discuss the agent's intentions. The utterances belonging to this game are mapped onto the predicate *intend*.
- ❖ **Suggest Game**: used to discuss *Goals*, *Context* beliefs and *Reflectors*. In the implementation, if the utterances of this game contain the predicate *goal*, they are mapped onto the agents' *goals*. If they contain *context*, they are mapped on the agents' *context*. Otherwise, they are mapped onto the agents' reflectors, denoted by *belref* in the implementation.

Below we will show a sketch of each of the games definitions, together with the moves that compose it.

Table 3. Dialogue Games and their Moves

Dialogue Game	Parameters	Available Moves
Inquiry	One Inquirer and One responder Inquirer aims at finding more information about a given topic.	Inquire, Statement, Challenge, Prompt, Clarify, Critique
Statement	One Informer and One responder Informer wants to convince responder of his/her beliefs	All of the above, with exception of the <i>inquire</i>
Proposal	One Proposer and One Responder. Proposer's main goal is to convince the responder to accept his/her intentions.	Propose, Counter-Propose, Statement, Challenge, Clarify, and Critique
Suggest	One suggester and One responder. Suggester's goal is to convince the reponder to accept his/her intentions.	Suggest, Counter-Suggest, Statement, Prompt, Clarify and Critique

The Model of the Task

At this stage of the work, MArCo's model of the task consists of a *consistency checker*, that scans the contents input by the user and checks them for logical inconsistencies according to time sequencing operators (inherent to planning domains, and described below) and the logical operators AND, OR, NOT and IMPLIES. Logical operators obey a propositional logic semantics.

The time sequencing operators used in our prototype have been derived from the empirical investigation described in [Tedesco and Self 2000a], and from our analyses of sample PERT/CPM problems. They can be divided into three main categories:

Table 4. Language Operators and Their Categories

Category	Available Operators
Ordering of Activities	Before (A, B) – indicates that activity A should be done before B in our strategy; After (A, B) – indicates the opposite situation; Parallel (A, B) – meaning that A and B can be done in parallel.
Time Duration of Activities	Starts (A, T) – indicates that A starts at time T; Lasts (A, T) – indicates that A lasts for T time units; Ends (A,T) – indicates that A ends at time T; Done (A) – indicates that A is completed.
Time Markers	Start_Project(A) – indicates that A is the first activity in the project (i.e. the lowest in the strategies' Ordering Relation, according to the definition presented in section 3.5) of the current plan. End_Project(A) – indicates that A is the last in the current plan (the highest in the Ordering Relation).

THE SYSTEM

MArCo consists of a Java based distributed interface that communicates with a Prolog Server.⁴ The distributed interface allows geographically separated group members to build a plan together. The Server, implemented in Sicstus Prolog, enables MArCo to reason over the interaction taking place at the client interfaces.

In the next section we will discuss MArCo's architecture. After that, we will present an interaction example, to help make its functioning somewhat clearer. Those less interested in implementation details should skip the next section.

MArCo's Architecture

When they log onto MArCo, users are presented with the window shown in Figure 6. Once the log on process is completed, group members are taken to MarCo's interface, where they are presented with a common workspace where their plan will be built. Users are able to contribute to the discussion via the dialogue and/or graph tools available at the interface. Once a user finishes her/his utterance, s/he presses the *Submit your Contribution* button on the interface (shown in figure 7). The contribution is then sent to the *Dialogue Record* (present at the interface) and to the *Dialogue Processor* on the Prolog Server.

⁴ The client interfaces can be used either over a local network or over the Web.




Figure 4 - MArCo's main Components and the flow of information.

The *Dialogue Processor* is responsible for three main tasks, as follows.

6. Receiving users' utterances and mapping them into the BGI attitudes that describe their strategies. In order to model the agents' cognition when planning together we have adopted a Beliefs-Goals-Intentions model. To us, these attitudes can be inferred from the agents' dialogue. From our point of view, Beliefs-Goals-Intentions are directly related to the components of our plans, as follows:

Intentions represent the actions the agents are committed to carrying out;

Goals represent the state of affairs the agents would like to bring about;

Beliefs represent one of the following (depends on the context and contents of the utterance):

Constraints to our problems – in this case, we have called them *contextual beliefs*;

The ordering of our intentions, and also, the justifications for them – in this case, they are called *domain beliefs*;

The agents' thoughts about how to proceed with the planning task –the *reflectors*.

7. Storing and manipulating the state transition automata that represents the dialogue games (Figure 6 shows a state transition automaton that amalgamates all the dialogue games available in MArCo); and
8. Holding the text templates used by the *Mediator* to intervene in the discussion. The arrows indicate the flow of information amongst the different components.




Figure 5 - The Dialogue Processor in detail.

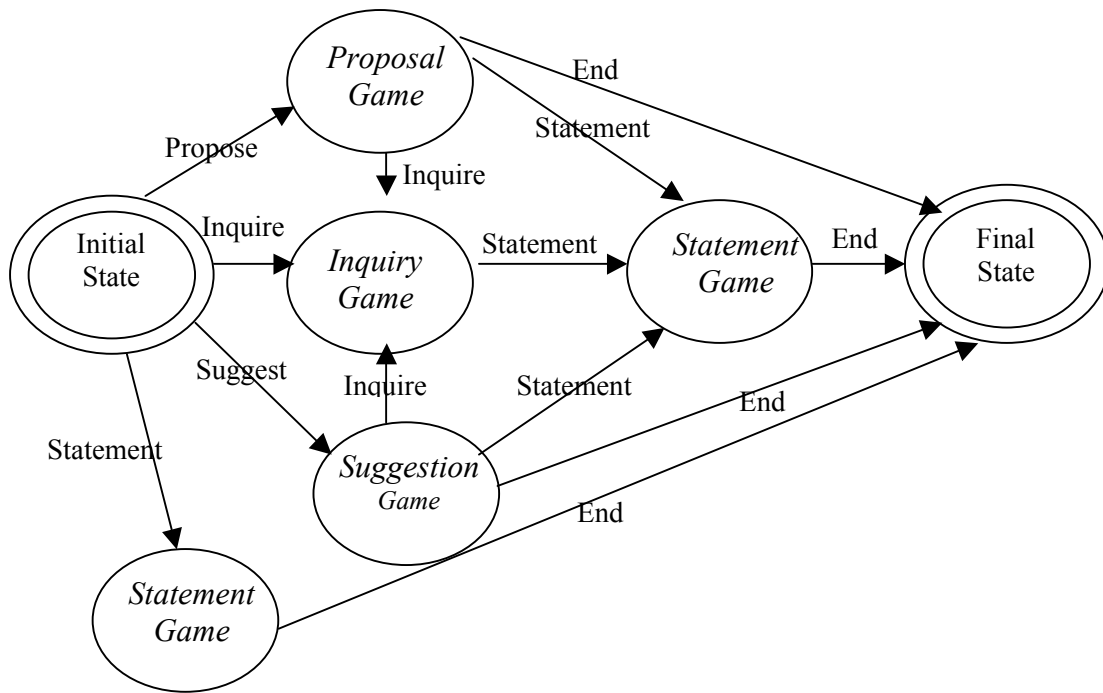


Figure 6 – Dialogue Games Available in MArCo.

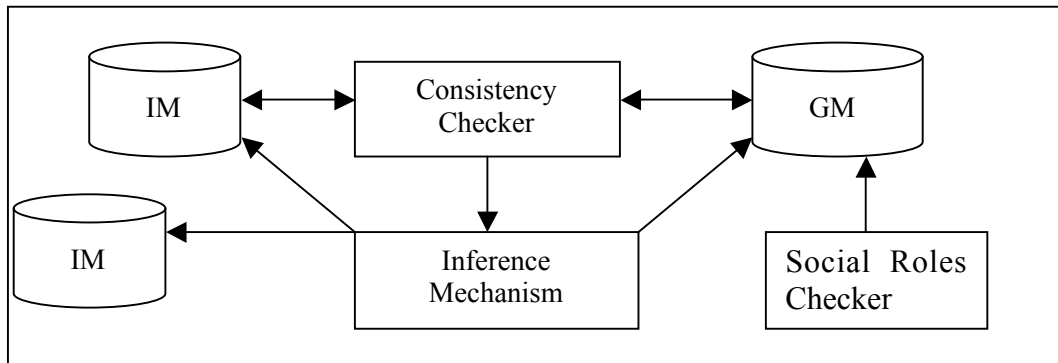


Figure 7 – The Maintainer in Detail.

Once the mapping of the utterances is completed, the *Maintainer* (presented in detail in Figure 7) proceeds to update the individual and group models, discussed previously. IMs are updated whenever a member makes a contribution. If an individual says something that is inconsistent with what was previously held in her/his model, the system assumes that s/he has changed her/his mind, and revises the relevant set of attitudes in the model, in order to eliminate inconsistencies. It uses *modus ponens* as the inference rule. For instance, let us consider our *user1*. Now, suppose we are adding *s* to one of the attitude sets in *user1*'s IM. Thus, the checker looks if the relevant set of *user1*'s IM contains any information of the form $\neg s$ or $t \Rightarrow \neg s$. If that is the case it revises the set, eliminating the inconsistencies. Once this step is completed, the *Consistency Checker* calls the *Inference Mechanism*.

The *Inference Mechanism* is responsible for inferring other mental attitudes related to the information just being added to the models, as well as any logical implications of this new information. These inferences are based on Cohen and Levesque's [1990] axioms for beliefs, goals and intentions. If any attitude can be inferred, the *Consistency Checker* is called in again, to try to ensure that the models are kept in a consistent state.

Whenever we have an agreement by the majority, the *Maintainer* updates the GM, including the agreement in the respective attitude set. Whenever the group finishes a dialogue game, the system updates the social roles played by its members.




Figure 8 – MArCo's interface.

MArCo's *Mediator* has a two-fold role. It is responsible for (1) detecting conflicts; and (2) for deciding on how to intervene in such situations, always aiming at supporting the solution process via *reflection* and *articulation*. The *Mediator's* reasoning is based on information obtained from the other components of the system. It has an intervention rule base that implements the mediation strategies discussed earlier in this paper. In order to perform its reasoning, the *Mediator* draws information from the group and individual models, the history of the dialogue, our models of conflicts and of strategic changes and a model of the task. Whenever the *Mediator* decides to intervene, it sends its interventions to the *Dialogue Processor*, which, in its turn, broadcasts them to the clients' interfaces.

In order to be able to do its reasoning, the mediator analyses information from the dialogue history, both IM and GM, as well as the consistency checker. The *Mediator's* conflict detection mechanism works as follows: the *Mediator* keeps observing the dialogue, and whenever it finds a disagreement about a given focus it adds that fact to its knowledge base. For instance, consider the excerpt shown below.

```
1.A.I Propose we do.before(a, (b,c) ).
2.B.I disagree, I think we should do.before(a,c) .
```

At this point, MArCo's mediator signals an intention disagreement (because of *A* and *B* declaring inconsistent intentions about activities *a*, *b* and *c*), and asserts the following fact:

Disagreement([A,B], before(a,c), [a,b,c], int) – which means that agents *A* and *B* had an intention disagreement about the focus *[a,b,c]*, and the last utterance in the process was *before(a,c)*. Not all

disagreements/conflicts are as explicit as the one portrayed above. For example, suppose that at some stage of the dialogue between *user1* and *user2*, *user1* proposed that they did *before(a,b)*, to which *user2* replied that she did not agree. Now, suppose that further down the dialogue, *user2* comes back to activities *a* and *b*, and proposes that they do *parallel(a,b)*. The *Mediator*, by analysing the *Dialogue History*, finds that *user2* is still declaring opposing intentions to those of *user1*, and signals an intention disagreement, asserting the following fact:

disagreement([user1,user2], parallel(a,b), [a,b], intention). (♥)

Now, let us go back to considering the excerpt in figure above. If, further down the dialogue, the mediator notices users *A* and *B* disagreeing about their intentions about the same focus (*[a, b, c]*), or about a focus that contains part of it (say, for instance, about *[a, b, d]*), it assumes that *A* and *B* were still trying to convince each other about their intentions regarding *a* and *b* and signals an intention conflict. The same applies for the example above, because the *Mediator* finds (♥) to be the second intention disagreement between *user1* and *user2* about *[a,b]* and signals an intention disagreement. Of course this scenario applies to other mental attitudes as well. In order to help the reader understand the functioning of MarCo, next section presents an example interaction. The example discussed was collected during the preliminary investigation, translated into MarCo's notation by hand. The process was based on the literature available and the results were discussed with colleagues, for validation. The correspondence between dialogue moves and mental attitudes was based on the analysis of the DISCOUNT scheme [Pilkington 1999], and is made as follows:

Statements: these moves are mapped onto domain beliefs. So, for instance, if *user1* states that *before(a,b)*, the interpreter maps this onto *beldom(user1, before(a,b),[a,b])*. The parameters in this expression mean that *user1* holds the domain belief, *before(a,b)* is the sentence *user1* believes in, and *[a,b]* is the focus of this particular belief. If the contents part of the utterance includes a reference to the context, the interpreter maps it onto *context(user1, contents, focus)*. The same applies for utterances that refer to goals.

Statements: these moves are mapped onto domain beliefs. So, for instance, if *user1* states that *before(a,b)*, the interpreter maps this onto *beldom(user1, before(a,b),[a,b])*. The parameters in this expression mean that *user1* holds the domain belief, *before(a,b)* is the sentence *user1* believes in, and *[a,b]* is the focus of this particular belief. If the contents part of the utterance includes a reference to the context, the interpreter maps it onto *context(user1, contents, focus)*. The same applies for utterances that refer to goals.

Proposal: these moves are mapped onto intentions. In other words, the system assumes that if somebody proposes that activity *a* is done, it is either because s/he has the intention to do it, or the intention that it gets done. When a user makes a counterproposal, MArCo assumes that s/he does not intend to do what the previous proposal referred to, and either intends to do something else or intends that something else gets done instead.

Suggests: are the moves used to talk about the strategy on a meta-level. Thus, they can be either employed to talk about reflectors, or are used to talk about goals and context. We have decided to use *suggest* to talk about context and goals because even though, to us, they are part of our strategies, context and goals can be used as criteria to distinguish between two different strategies to solve a problem. Consequently, they can be thought of as reflectors. Analogously to the propose/counter-propose case, counter-suggests are used to disagree with the previous suggestion, and inform our listeners of our different reflectors.

An Interaction Example

Before we show more of an interaction example, some details about the language we have used need to be discussed. When composing formulas, participants use the operators described in section in Table 4. In our implementation, these operators are used in a prefix form. So, for example, the expression *after(a, b)* means that activity *a* is done after activity *b*; *lasts(a, 10)* means that activity *a* will last for 10 unit of times and so forth. The logical operator *implies* will appear throughout the excerpts and, depending on the context, can denote a causal link. For example, *implies(a, done(b))* is used to mean that activity *a* needs activity *b* to be done first. In other situations activities may be grouped in an operator, as is the case in *before((a, b), c)*. In this case, *(a, b)* means that activity *a* is followed immediately by activity *b*. Thus, *before((a, b), c)* means that activity *a*, followed immediately by activity *b*, is done before *c*. In other words, activity *c* is done after *a* and *b*. In order to facilitate comprehension, we will provide natural language interpretations of the more complex expressions.

Let us now go back to the example. To begin with, let us assume we are dealing with learners *A* and *B*. Their individual models will be henceforth referred to IA and IB, and the group model will be called GM, as usual. At the beginning of the interaction, the available dialogue moves are *inquire*, *statement*, *propose*, and *suggest*, each corresponding to the initial moves of the four available dialogue games. *A* and *B* start their dialogue by discussing the constraints of the problem, trying to decide when Christmas sales start. The beginning of their discussion is shown in the excerpt below

```
1. Inquire(A, starts(xmas_sales, *)).
2. Statement(B, starts(xmas_sales, december)).
```

When *A* submits utterance 1 to MArCo, it is sent to the *Dialogue Processor*. Once it is received, it is added to the *Dialogue History*. No mapping onto mental attitudes is made, since *A* is not yet making any declarations about her mental attitudes. When utterance 2 is received by the *Dialogue Processor*, the *Dialogue Interpreter* determines its focus, maps it onto a domain belief, and send the following message to the *Maintainer*.

Beldom(B, starts(xmas_sales, december), [xmas_sales, december]).()*

The *Maintainer* then checks if there is any domain belief in IB that is inconsistent with this new sentence. Because the model is empty, a message is sent to the *Inference Mechanism*, telling it to determine which attitudes can be inferred from (*). Since no inference can be made⁵, nothing is added to IB. At this point, the dialogue moves available at the *Dialogue Moves* window at the interface are: *challenge*, *prompt*, *clarify* and *critique*. Below we present the remainder of the current dialogue game:

```
3. Challenge(A, starts(xmas_sales, november)).
4. Critique(B, agree).
```

Upon receiving utterance 3, the *Dialogue Interpreter* maps it onto two new domain beliefs:

(a) *beldom(A, not(starts(xmas_sales, december)), [xmas_sales, december]).*

(b) *beldom(A, starts(xmas_sales, november), [xmas_sales, november]).*

Once this mapping is ready, a message is sent to the *Maintainer*, so IA can be updated. Since IA is empty, there is no need to revise IA's belief set. Also, at this stage the *Inference Mechanism* cannot infer any other attitudes. Thus it adds both (a) and (b) to IA. After that, information is passed onto the

⁵ Neither from the IB's belief set (via modus ponens), nor from the other attitude sets.

Mediator, who, by analysing the pattern of the dialogue (a statement, followed by a challenge) detects a belief disagreement, and adds the following fact to its knowledge base:

Disagreement([B,A], starts(xmas_sales,november), [xmas_sales,november], beldom)

On utterance 4, *B* declares her change of mind, and agrees with *A* that Christmas sales actually starts in November. Thus, her critique is mapped by the *Dialogue Interpreter* onto:

*Beldom(B, starts(xmas_sales, november), [xmas_sales, november]).(**)*

When the *Maintainer* receives this message, it adds it to the belief set in IB. After that the *Consistency Checker* finds that IB is now inconsistent, because of the presence of both (*) and (**). It then proceeds to revise IB with respect to (**). Because there are no further logical implications of (*) in IB, the *Consistency Checker* only needs to eliminate (*) from its belief set. Once this is done, it sends the results to the *Inference Mechanism*, which, at this point is unable to make any further inferences.

At the end of the dialogue game, with both *A* and *B* agreeing about the fact that Christmas sales start in November, the *Maintainer* updates the GM with this information, as well as looking at the roles played by *A* and *B*. In this case, since both came up with different alternatives to the solution, and took half of the available turns MArCo registers them as being *outspoken*.

The dialogue moves shown in below present *A* and *B* continuing their discussion about the time limits of their project. In order to show how we have adapted our natural language data to the formal notation needed by MArCo, we will present the natural language utterance (in bold), followed by the formal notation.

I guess that Christmas sales start at mid-November, or something like that? Let's say it begins on the first of November.

5. Suggest(A, context(starts(xmas_sales, first_november))).

Yes. Let's just invent a date...

6. Critique(B, agree).

So. July. Let's suppose it is the 1st of July.

7. Suggest(B, context(starts(project, first_july))).

Yes, 1st of July.

8. Critique(A, agree).

So, the project lasts four months... Sixteen weeks?

9. Inquire(B, lasts(project, *)).

Let's say it is seventeen weeks.

10. Statement(A, lasts(project, 17weeks)).

Sorry?

11. Prompt(B).

It is seventeen weeks. From the 1st of July to the 1st of November – July, August and October are 31 days. So, approximately seventeen weeks.

12. Statement(A, implies(and(starts(project, first_july), starts(xmas_sales, first_november)), lasts(project, 17weeks))).

Hum, Ok.

13. Critique(B, agree).

At the end of the dialogue game, the *Maintainer* adds to IA, IB and GM the fact that they have agreed on: (1) context(starts(xmas_sales, first_november)), (2) context(starts(project, first_july)), and (3) context(lasts(project, 17weeks)).

Whenever the participants make a contribution that can be mapped directly onto the components of a strategy, the *Maintainer* updates their private and/or group strategies, via the predicate build_strategy.

Thus, it can identify which changes are being made. Thus, for example, at this point, the *Maintainer* has added to the GM that its group strategy is:

Strategy(grp, Set_of_Intentions, Ordering, [context(starts(xmas_sales,first_november)), context(starts(project,first_july), context(lasts(project,17weeks))], Goal).

Utterance 12 in the dialogue above shows evidence of a reflective activity, where *B* presents reasons for choosing such a context of the problem. Before we proceed, let us have a look at the contents of IA, IB and GM.

Table 5 - The state of IA, IB and GM, after utterance 13.

IA	IB	GM
Beldom (A, (not (starts (xmas_sales, december))))	Beldom(starts(xmas_sales, december))	
Beldom (A, starts (xmas_sales, november))	Beldom (B, starts (xmas_sales, november))	Beldom(grp, starts(xmas_sales,november))
Context (A,starts (xmas_sales, first_november))	Context(starts(xmas_sales, first_november))	Context(grp, starts(xmas_sales, first_november))
Context (A, starts (project, first_july))	Context (B, starts (project, first_july))	Context(grp, starts (project, first_july))
Beldom (A, lasts(project, 17weeks) (♣))		
Beldom(A, implies(and (starts (project,first_july), starts (xmas_sales,first_november)), lasts(project, 17weeks)))	Beldom(B, implies(and(starts (project,first_july),starts (xmas_sales,first_november)), lasts(project, 17weeks)))	Beldom(grp, implies (and (starts (project,first_july), starts (xmas_sales,first_november)), lasts(project, 17weeks)))
	Beldom(B, lasts(project,17weeks)) (♦)	Beldom(grp, lasts(project, 17weeks)) (♦)

Table 5 presents the state of the IA, IB and GM at this stage of the interaction. It should be noted that when *B* agreed with *A* (in utterance 4), her previous belief about Christmas sales was deleted from IB by the *Consistency Checker*. Also, the difference between (♣) and (♦) in IA and IB is that while (♣) was added to the model when user *A* stated it, (♦) was added to IB by the *Inference Mechanism*, just after *B* agreed to *A*'s utterance. The same reasoning applies to the addition of (♦) to the GM.

Having determined their problem context, the pair then goes on to discuss how to organise the remaining activities. The next few dialogue moves show a sample of how that dialogue went.

Yes. I mean, the distribution and the... These two distribution things, they can only be done if the whole thing is produced, actually.
 14.Propose (A, before (production, distribution)).
Yeah.
 15.Critique (B, agree).
I mean, probably together with advertising, this (distribution) can go in parallel, but this is definitely something you can't do without having produced the whole thing.
 16.Propose (A, parallel (distribution, advertising)).
Yes. So you mean that this will be done by the last week of October?

17.Critique(B, agree).

Yes. Last two weeks of October, because on the first of November they must start to sell it.

18.Propose(A, ends((l,k), first_november)).

Ah, Yes. I was confused by the 17 weeks.

19.Critique(B, agree).

And advertising campaign watch-for and hard-sell go in parallel.

20.Propose(B, parallel(m,n)).

Yes.

21.Critique(A, agree).

So, you can actually do the hard-sell campaign only if things are already in the shop. Is that true?

22.Inquire(B, (l,n)).

Yeah. You can do this is parallel, anyway.

23.Statement(A, parallel(l,n)).

Yes.

24.Critique(B, agree).

It is like you do the hard-sell campaign for two weeks and then the things are in the shops, so people can buy them.

25.Propose(A, before(n,l)).

Yeah.

26.Critique(B, agree).

Production run is the precedent activity, 8 weeks, can't be done in parallel with distribution?

27.Inquire(B, (j,k)).

I mean, this (production run) has to run before the distribution.

28.Statement(A, before(j,k)).

Hum Hum.

29.Critique(B, agree).

Ok. And in front of this production run we need one week training of labour, I think, because they can't do it.

30.Propose(A, before(i,j)).

Training of labour yes, before production run.

31.Critique(B, agree).

In the dialogue moves presented above, the roles played by the participants did not really change, with *A* playing the *leader* (characterised by her making most of the proposals/statements in the game, and her opinions being accepted by *B*) and *B* playing the role of the *follower* most of the time. Thus, there were no conflicts or mediator interventions. One idea worth investigating is to determine which parameters could regulate the timing of the *Mediator's* intervention, and the social roles played by the participants is definitely an interesting avenue to be explored. In this way, we could improve the scope of the *Mediator's* actions, making it intervene in more situations (instead of only intervening on conflict detection). For instance, in the case shown in above, there are not many indications of reflective activities. Thus, (an improved version of) the *Mediator* could have, after utterance 31, for example, intervened and asked *B* to elaborate on the reasons for the agreement.

In order to exemplify the inferences carried out by the *Inference Mechanism*, let us look at utterance 14 above. The *Inference Mechanism* starts by adding *before(production, distribution)* to the set of intentions in IA. It then tries to find any logical implications (via *modus ponens*) of the new sentence being added. Next it adds the following sentences to IA (according to the inference axioms):
Beldom(A, before(production,distribution),[production,distribution]) –

By axiom I3, it can infer that

$Goal(A, \phi)$ where $before(production, distribution) \mathbf{a} \phi$ ⁶

Once IA is in a consistent state, it is time to look at A's strategy. Strategies are updated by the *Inference Mechanism*. At this point, A's strategy is described as follows:

$Strategy(A, [production, distribution], [before(production, distribution)], [context(starts(xmas_sales, first_november)), context(starts(project, first_july), context(lasts(project, 17weeks))], Goal]$.

Since we are adding an intention to A's strategy, MArCo records this operation as being a *strategic expansion*.

Utterances 23 and 25 indicate some reflective activity done by A, showing her change of mind. When utterance 25 is mapped onto the BDI model, it results in:

$Intend(A, before(n, l), [n, l])$. (I)

When the inference checker works with (I) above, it infers (by axiom I4):

$beldom(A, before(n, l), [n, l])$ – which is inconsistent with $beldom(A, parallel(n, l), [n, l])$, that was added to IA when utterance 23 was mapped. Thus, a revision is performed in IA's belief set, to bring it back to a consistent state.

Let us now fast forward this dialogue to the excerpt where we get some indication of reflective activities, and the explicit occurrence of a conflict.

Utterances 33 and 34 indicate a reflective activity, when B is asked to provide more information about the ordering of the activities being discussed. This leads to B articulating some causal links in her reasoning, when she says that in order to do *h* and then *i*, we need to have done *d* first. As we have discussed in the beginning of this section, during the experiments users have used implies to express "depends on". So, in utterance 34, $implies(before(h, i), done(d))$, means that doing *h* before *i* depends on *d* being done.

Contributions 36 to 40 show an explicit conflict when we have two consecutive disagreements about the same focus (activities *b*, *d* and *f*). After the conflict, we can see some indication of reflection, when user B declares that they needed to complete activity *b* in order to be able to do activity *d*. Figure 9 shows the state of the interaction at the point of disagreement (just after utterance 39) with the mediator's intervention, and then, with A's disagreement (utterance 40), the mediator's full-fledged intervention (which occurred when the intention conflict was detected). Figure 9 also shows the state of the graph at that point (this particular pair had chosen to build it backwards, having discussed about the time restrictions of the problem first).

⁶ By $before(production, distribution) \mathbf{a} \phi$ we mean that the goal ϕ can be derived from the contents.

I think that retooling can be done in parallel with training of labour.

32.propose(B, parallel(h, i)).

But this means that people will be trained with the old machines somehow. Is that true?

33.prompt(A).

I would say that retooling just...ah, because here they manufacture the moulds and they will fix the moulds and set the machines. This can't be done in parallel, because let's say the work force has to be trained with them, so it can't be done in parallel.

34.statement(B, implies(before(h, i), done(d))).

Yes. Ok, so...

35.critique(A, agree).

(dialogue proceeds here, adding more activities to the strategy, and drawing the graph)

Yes, manufacture of moulds. This has to be before the retooling and takes two weeks.

36.propose(B, before(d, h)).

Yes.

37.critique(A, agree).

Manufacture of moulds before the retooling and maybe before (A) the trial manufacture.

38.propose(A, before(d, f)).

Yes, but it is impossible.

39.counterpropose(B, and(before(b, d), before(d, f))).

Yes, it is possible.

40.critique(A, disagree).

Yes, it is impossible because you have detailed design of toy and moulds. How can you manufacture them without the design?

41.statement(B, implies(done(b), d)).

Yes, that is true.

42.critique(A, agree).

So, detailed design and manufacture together take five weeks.

43.statement(A, lasts((b, d), 5 weeks)).

Yes, but

44.critique(B, agree).

It is five weeks and then trial manufacture.

45.statement(B, and(before(b, d), before(d, f))).

No, this is not possible, because it is a week too much.

46.challenge(A, implies(and(before(b, d), before(d, f)), lasts(project(18weeks)))).

Yes, Yes.

47.critique(B, agree).

The following dialogue re-discusses the context of the problem.

Yes, ok, we will start two weeks late.

48.statement(A, starts(xmas_sales, 15th_november)).

Yes, it is Brazil, let's say.

49.critique(B, agree).




Figure 9 - Screenshot of MArCo, when the intention conflict was detected.

When the mediator intervenes, after utterance 40, it tries to provoke *reflection* and *articulation* by first reminding the pair that they had not yet discussed the goals of their strategy (which not only gives them an opportunity to think about the solution, but also can lead to a better defined plan) and asks both to elaborate on why they think their solutions should be adopted. In order to generate the intervention, the *Mediator* sends the names of the participants, the fact that they had not yet discussed their goals and the request to elaborate on previous views to the *Template Holder*. The *Holder* then fills the relevant template in and sends it to the interface to be displayed.

Utterances 41 to 47 discuss the causal links between activities *b*, *d*, and *f*. For example, in (40), *B* states that we need to have done *b* (detailed design of moulds) in order to be able to do *d* (manufacture of moulds). With *A*'s agreement (on utterance 46), the pair goes on to analyse the implications of this ordering to the context of the problem (the project lasting 17 weeks, and thus finishing on the 1st of November), and discover that the context can be redefined without any loss to the problem.

THE EXPERIMENTAL STUDY OF MARCO

Our study of MArCo consisted of two stages of *formative* evaluations. Firstly, we wanted to find out whether the functionalities of the interface were sufficient to allow people to communicate and build their

plans together. In the second stage, our main goal was to investigate whether people's awareness of conflicts influenced the quality of the results they achieved. Moreover, we wanted to see if the mediator's intervention helped participants to reflect on and articulate their ideas about the solution. We have also interviewed our participants, if they were aware that conflicts were taking place, and how adequate/effective the mediator's interactions were.

When the prototype was implemented, we asked some colleagues to help us find out more about MARCo's interface. Our aim was to check whether it was possible to communicate through it. Besides analysing the usability of the interface, we also wanted to find bugs in the system. For instance, we were looking at issues of graph manipulation (is it possible to manipulate the graph with the tools provided? Do we need more functionalities to do it?), of text formulas usage (do we need more assistance when composing them? Is the need to use a formal language too big a complication? Will that actually prevent people from communicating?), and of dialogue moves approach (are we missing dialogue moves?). This study showed that the overall functionality of MARCo seemed fine.

For the second stage of the experimental study we have observed pairs of users interacting through MARCo. Each user worked on a different machine. Due to logistics, both participants worked in the same room, and were asked not to talk to each other during the experiment. However, users were allowed to ask questions to us (the observer) at any time.

The Participants

We have investigated the use of MARCo with seven pairs of people. Part of the reason to work with pairs was that it would be easier to co-ordinate times and availability. The underlying mechanisms of the system are extensible for groups of more than two people, though. Out of fourteen participants, twelve were postgraduates, taking different courses at the University. The remaining two were Spanish language teachers. All of the participants had some computer skills, even though the levels of their knowledge varied. Out of the twelve postgraduates, six were CBL's PhD students, two were MEds, two were computing MScs, and the other two were PhD students at the School of Geography with significant computer skills. In fact, MARCo is intended to be used by mature students, working to complete their projects. Thus, we have tried to choose participants that fit the description as closely as possible. We were also concerned with the used of logics for the interaction. Thus, we also tried to find participants with some logics background, in order to lessen their (possible) difficulties at interacting, and make interesting interactions possible. At the time of writing, we are investigating other possible interaction mechanisms.

Five of these pairs used the prototype with the mediator turned on (pairs 2, 3, 4, 6 and 7), and two used it without (pairs 1 and 5). Participants' interactions were recorded in a log file, to be later analysed. Our presence in the room during the interactions allowed us to record interesting comments and questions that happened during the problem solving. It was also important in case we needed to fix something that escaped our testing of the prototype. While we recognise that perhaps participants would feel more at ease if the experiment was observed by someone else, logistics at the time made it difficult.

After participants had finished with their task (it was left to the pairs to decide when they had finished), they were interviewed. This happened after the participant's interactions finished, in the same room. We had the users answer a semi-structured interview, and recorded their answers for posterior analysis.

Preliminary Comments on our Findings

During the experiment, we were able to find evidence of all the conflicts predicted in our computational model. For the sake of space, we will only show a short example here. We refer the interested reader to [Tedesco 2001a]. In the log files of the experimental study, we have found several

evidences of strategic operations, mainly revisions and additions. No contractions were found, due to the fact that participants were asked to put all the activities in the graph representation (this being a requirement of PERT graphs). Nevertheless, two of the pairs asked the observer if they could leave activities out of the graph. One example of a concatenation is found in figure 10 below.

1.A4.I Propose we do. and((after(j,h),after(n,m)), (after(k,j), after(l,k))).
Let's do activities j after h, n after m,k after j, and l after k.
 2.B4.Ok, I agree.

Figure 10 - Example of a Strategic Concatenation.

At this point, the mediator could have intervened, by asking B4 why s/he agreed with all the proposed steps, or by asking A4 to explain the orderings of the steps, to make clear to B4 what the other components of his strategy were (since they were not explicitly mentioned in the excerpt). The mediator's role, in this case, could be to try to make sure that both partners understood completely the concatenation. In the case of a conflict, the mediator could then have asked A4 to elaborate on different parts of the strategy, or asked B4 to explain to which parts he disagreed, and why.

After investigating whether the events predicted in our models actually occurred, we wanted to find if our approach was useful. In other words, we wanted to investigate if conflicts are actually triggers of changes, and if by detecting and mediating them, we can see indications of *reflection* and *articulation*. We were particularly interested in analysing the occurrences of reflective activities within the context of conflicts, and their occurrence after the intervention of the mediator. Our hypothesis was that, in general, we can find a fair percentage of reflective activities happening after an intervention. If that is the case because the mediator interrupts people and gives them space to think in an interaction that could otherwise be too fast or because the mediator's interventions are relevant to the situation is another issue that we would like to investigate in the near future. Whatever the case, it gives us grounds to argue that the presence of the mediator helps to promote reflection. Thus, before showing evidence of reflection, we will define which reflective activities we are considering in the study.

Some dialogue studies (for instance, [de Vicente et al., 1999, Katz et al. 2000]) remark that interactions that are more successful as far as learning goes, involve *clarify*, *challenge* and *justify* moves. From our point of view, such moves are considered successful because they indicate *reflection* on the part of the individuals. Thus, we have considered these moves as indicators of reflective activities. In the context of MArCo, they are indicated by the following main factors:

The user changes his/her mind. Whatever the reason, we will assume that changes of mind are due to user's reflection. For instance, figure 11⁷ depicts users A and B discussing how to draw the graph of the plan they have devised. Utterances 6 and 7 show A suggesting something and then contradicting herself, due to some reflection that occurred (she presents reasons for the change of mind). Most times, these changes of mind reflect a *strategic revision*.

⁷ These dialogues belong to the data collected in the experiment described in [Tedesco and Self 1998]. To protect participants privacy, we have decided to call them A and B, regardless of the pair they belong to.

1.A.Whatever we want, right? The planning goes back, if we do this..
 2.B.Yes, let's just do the graph?
 3.B.Something symbolic?
 4.A.Can we do like this?
 5.O. (Confirms)
 6.A.all right, let's do it? Why not?
 7.A.Because here it says weeks, so we have like a timetable. If it goes back, it's gonna be confusing.
 8.B.Yes...

Figure 11 - Example of a Reflective Activity - Change of Mind.

The user elaborates on what was previously on his/her (or the group) model. This can be done by, for example making causal links between pieces of information held or introducing new lines of reasoning that include the topic of the conversation. To make matters clearer, let us suppose that *user1* said that she believes there will be a traffic jam on the way to Porto de Galinhas (a beach outside Recife). Following that, she informs *user2* that she thinks there will be traffic because she has heard on the news that a nasty traffic accident happened on that road just before they set off to the beach.

The user goes back to what was previously discussed. This indicates that s/he has found out that things were not too clear in the beginning, and s/he wants to ground his/her point of view. The excerpt presented in Figure 12 shows A and B going back to discuss the context of their problem, in the light of the graph that they have built.

1.B. Instead of beginning selling at 1st of November, we begin later, because it was our 1st assumption, that is not very strict.
 2.A. Yes.

Figure 12- Reflective Activity - Return to a previously discussed topic.

The user challenges an utterance that has been made, or inquires about its causes. We consider the act of asking for justifications/explanations to be reflective. We assume that when asking for justifications, participants might either have other ideas in mind, or not have been able to work out the reasons for what has been said. The same line of reasoning applies for challenges.

Let us now examine some evidence of reflective activities found in the experimental study. The excerpt presented in Figure 13 happened during a disagreement between users A4 and B4. It represents an instance of a reflective activity where users are trying to convince each other by elaborating on the causal links between the activities in question. In utterance 2, A4 states that he thinks we need to have done *c* to be able to do *e*. In the follow-up, B4 then states that if they want to spend only 4 weeks with *e*, they need to do it in parallel with *c*.

1.B4.I Propose we do.parallel(c,e) .
 2.A4.I think that.implies(c,e) .
 3.B4.I think that.implies(lastse(4),parallel(c,e)) .
For e to take four weeks, c needs to be done in parallel with e.

Figure 13 - Example of a Reflective Activity during a Disagreement.

The excerpt in Figure 14 shows another example of how conflicts can be triggers of changes, and how a mediator can take advantage of such situations to try to provoke *reflection*. Utterances (1) and (2) present a disagreement about the ordering of activity *d* (that was a follow-up of a previous discussion) and the mediator's intervention. When intervening the mediator tries to provoke some *reflection* and *articulation* on the plan being constructed. In this case, after the intervention, we can see B7 explaining that *d* needed *b* to be completed.

1.A7.I Propose we do.after(d,c) .
 2.B7.I disagree, I think we should do.after(d,b) .
 3.System.statement.There Was a Disagreement About last Proposal. Would you like to say why you disagree?
 An Intention Conflict was Detected.
 Do you not think you should define your goals first, to guide the problem solution?
 B7, can you elaborate on why you think you are correct?
 A7, can you elaborate on why you think you are correct?
 A difference of views was detected here!
 4.B7.I think that.implies(d,done(b)) .
In order to do d, we need b to have been done.
 5.A7.Ok,I agree.

Figure 14 - Example of a Reflective Activity during a Conflict.

Summary of the Findings

There are other pieces of evidence similar to the ones discussed in the log files of the experiment. Taking into consideration both mediated and non-mediated cases, we found nine Belief Conflicts, five Ordering conflicts and twelve Intention Conflicts. Interestingly, most of the belief conflicts were found in pairs 2 and 3 (the ones that had the majority of conflicts during the interaction), and are mostly caused by the members arguing about justifications for their attitudes. It is also true that reflective utterances (clarifications, prompts, inquires, and justifications), as well as other reflective activities (defined above) appear more in the cases where conflicts were detected, as we thought.

In total, we have found fifty-eight reflective utterances (i.e. those indicating evidence of reflective activities). Out of these fifty-eight, twenty happened outside the context of a conflict (for example, when participants explained the reasons for their attitudes to their partners). There were thirty-eight (65.5%) reflective utterances in the context of conflicts and/or disagreements (these correspond to individuals grounding their points of view as well as challenging their partners' positions). This corroborates our hypothesis that conflicts can be promoters of change, and present a good opportunity for us to facilitate *reflection* and *articulation*.

Many of the reflective activities in the context appear after a conflict or disagreement was pointed out, indicating that the mediator's intervention was adequate. In the mediated pairs, there were thirty-three reflective utterances in a conflict/disagreement context. Out of these, twenty-one utterances happened after a conflict was detected. Table 5 shows the comparison between reflective activities during a conflict/disagreement and the ones happening after a conflict was detected.

Pairs 2 and 3 were the ones that engaged into conflicts more frequently. The basic difference between them was that pair 2 found that the mediator's interventions were sometimes too long to read. Pair 3, on the other hand, took their time to consider the mediator's interventions more carefully. Pair 4 had a fairly short interaction (both participants were pressed for time). With pair 6, there was quite a definite role division between the two participants— one took the leader role and the other followed most times. There are not many conflicts or disagreements in this pair's dialogue. As a result, their reflective activities are mostly related to moves of clarification/justification. These results indicate that the mediator could be extended to intervene in situations where roles are too static, to explore further the potential for *reflection* and *articulation*.

Table 6 - Comparison between Reflective Activities found during and after conflicts in the Mediated Pairs.

Pair Number	Reflective Activities in the Context of Conflict/Disagreements	Reflective Activities After Conflict/Disagreements
Pair 2	11	5 (45%)
Pair 3	11	11 (100%)
Pair 4	3	1 (33%)
Pair 6	1	1 (100%)
Pair 7	7	3 (43%)

For comparison purposes, let us discuss the results of the non-mediated pairs. In the dialogue of pair 1, who did not have the mediator turned on, we only found one intention and one belief disagreements, as well as one intention conflict (we speculate that this was partly due to the fact that one of the participants was much more experienced than the other) there were only four reflective activities detected. Three of these occurred within the context of their disagreements. Most correspond to clarification moves prompted by the less experienced participant. The resulting explanations from the expert are quite short - and not questioned by the novice. Pair 5, also not mediated, had seven reflective activities, two of which happened in the context of their conflicts/disagreements. Both of these two reflective activities were by members trying to put their point across during a disagreement.

During the debriefing sessions, the several issues about the mediation process came up. All of the participants said that they were aware of the conflicts happening, and were trying to better put their point across. They remarked that their awareness was due to the nature of the domain, and to the dialogue game structure, that made things more obvious to them. They said that this state of affairs might change in more complex domains (for instance, in a software design meeting), which would make the idea of having the artificial mediator even more interesting.

Participants remarked that mediator's interventions were adequate, and did help them think about what was going on in the discussion. They said that even though sometimes they thought the intervention was too long, the pause in the interaction gave them time to stop and think matters through. All the pairs, including the non-mediated ones, liked the idea of having the artificial mediator. One participant remarked that even though the mediator did not make her change her mind, its presence helped her plan her next moves more carefully, deciding how to say things better. When asked what they would do to improve the usefulness of the mediator, users made the following points: (1) the mediator should be more

clear when pointing out group agreements; (2) use of stronger statements when referring to group actions; (3) point out the usefulness of the group agreements button more often; (4) including tools that helped in the planning process (like a time counter) and making the mediator as informed as possible (as in able to predict, for example, which solution will help them get to goal).

The general assessment of the approach seems to be quite positive, with participants of the experiment comfortable with the idea of working with an artificial mediator. Participants' experiences have led us to think that the approach is useful, and with improvements, it can certainly be used in different settings and domains. The suggestions received have been recorded and are being considered for further extensions of the prototype. We believe that MArCo can be extended fairly easily to other domains. However, in order to enable the mediator to make more interesting interventions, we should also provide MArCo with some knowledge of the intended domain.

CONCLUSIONS AND FURTHER WORK

In this paper, we have presented a summary of the work carried out in the building of MArCo. We have explored the idea that it is possible to use the conflicts inherent to group interactions to foster *reflection* and *articulation*, and thus get more refined solutions. In order to do so, we have built computational models of conflicts and of the strategic changes they bring about. In order to make mediation possible, we have also analysed the issues of interaction, and of group and individual models.

There are many possible extensions to the work proposed in MArCo. In fact, the first author is working on a research project funded by CNPq that aims at investigating the following:

Group characteristics that can be used to support *Focus Groups* and *e-training* systems, and how we can enable intelligent agents to negotiate such a model. This investigation is currently being carried out as part of the AMADeUs Project, that aims at creating virtual environments for teaching maths. Within the AMADeUs framework, both students and teachers can interact with their peers. Learners to learn collaboratively, and teachers to share experiences and plan courses together. In this case, Individual and Group Agents negotiate group formation and monitoring.

Interaction mechanisms that are more natural - and how we can give the mediator other clues as to how it can intervene more effectively;

Extending the mediator in MArCo, so that it can not only include other functionalities, but also be more effective in other domains.

Including other agents that interact with the user (on the lines of learning companions) and that can play different roles in the discussion (expert, devil's advocate and so forth).

The approach implemented in MArCo brings contributions not only to Computer Science, but also to AIED (where we add to the trend of research working on supporting reflection, and also to the trend that works on monitoring/supporting the collaborative process). The presence of artificial mediators can be seen as yet another way to support electronic meetings. The use of dialogue games helps people to be aware of the rationale for their decisions, and thus be able to reach better-informed solutions. These two issues also represent contributions to Computer Supported Co-operative Work. The novelty of the approach lies in its definition of a metacognitive conflicts model (that can also be used in DAI situations) and the framework for conflict mediation, that embeds several issues fundamental to group interactions.

References

- Baker, M. J., and Lund, K. (1996). "Flexibly Structuring the Interaction in a CSCL Environment". In Brna, P., Paiva, A. and Self, J. (eds.), *Proceedings of the European Conference on Artificial Intelligence in Education*, Edições Colibri, 401-408
- Baron, R. S., Kerr, N. L., and Miller, N. (1992). *Group Process, Group Decision, Group Action*, Open University Press
- Boud, D., Keogh, R. and Walker, D., (1985). "Promoting Reflection in Learning: a Model". In Boud, D., Keogh, R. and Walker, D.(eds.) *Reflection: Turning Experience into Learning*, Kogan Page, 18-40
- Burgoon, M., Hunsaker, F. G. and Dawson, E. J. (1994). *Human Communication*. Sage Publications
- Burton, M., (1998). "Computer Modelling of Dialogue Roles in Collaborative Learning Activities". Unpublished Doctoral Dissertation, Computer Based Learning Unit, University of Leeds.
- Castelfranchi, C., (2000). "Conflict Ontology". In Müller, H.J., and Dieng, R. (eds), *Computational Conflicts: Conflict Modeling for Distributed Intelligent Systems*, Springer Verlag, 21-40
- Cohen, P. R., and Levesque, H. J., (1990). "Intention is Choice with Commitment", *Artificial Intelligence*, 42, 213-261
- Constantino-González, M.A., and Suthers, D.D. (2000). "A Coached Collaborative Learning Environment for Enitty-Relationship Modeling". In Gauthier, G., Frasson, C., and VanLehn, K. (eds.) *Proceedings of the 5th International Conference Intelligent Tutoring Systems (ITS'2000)*; Lecture Notes in Computer Science, 1839, Springer Verlag, 324-333
- Crook, C. (1994). *Computers and the Collaborative Experience of Learning*, Routledge
- de Vicente, A., Bower, A. and Pain, H. (1999). "Initial Impressions on using the DISCOUNT scheme". In Pilkington, R., McKendree, J., Pain, H., and Brna, P. (eds.) *Proceedings of the Workshop on Analysing Educational Dialogue Interaction*. Workshop at the AIED'99, 9th International Conference on Artificial Intelligence in Education, Le Mans, France, 87-94
- Dewey, J. (1933). *How we Think: a restatement of the relation of reflective thinking to the educative process*, Heath.
- Dimitrova, V., (2001). "Interactive Open Learner Modeling". Unpublished Doctoral Dissertation. Computer Based Learning Unit, The University of Leeds
- Doise, W., and Mugny, G., (1984) *The Social Development of the Intellect*, International Series in Experimental Social Psychology, vol. 10, Pergamon Press
- Easterbrook, S. M., Beck, E. E., Goodlet, J. S., Plowman, L., Sharples, M. and Wood, C. C. (1993). "A Survey of Empirical Studies of Conflict". In Easterbrook, S. (ed.) *CSCW: Cooperation or Conflict?*. Diaper, D. and Sanger, C. (eds.) Computer-Supported Cooperative Work series, Springer-Verlag, 1-68
- Galliers, J. R., (1989). "A Theoretical Framework for Computer Models of Cooperative Dialogue, Acknowledging Multi-Agent Conflict", PhD Thesis, Open University.
- Hartley, P. (1997). *Group Communication*, Routledge
- Hillier, F. S., and Lieberman, G. J., (1995). *Introduction to Operations Research*, McGraw-Hill International Editions
- Hoppe, H. U. (1995). "The Use of Multiple Student Models to Parameterise Group Learning". In Greer, J. (ed.) *Artificial Intelligence in Education 1995*", AACE
- Hoppe, H.U., and Ploetzner, R. (1999). "Can Analytic Models Support Learning In Groups?". In Dillenbourg, P. (ed.) *Collaborative Learning: Cognitive and Computational Approaches*". Bennett, N., DeCorte, E. Vosniadou, S. and Mandl, H. (eds.) *Advances in Learning and Instruction Series*. Pergamon Press, 103-121
- Johnson, D.W., and Johnson, F. P., (1997). *Joining Together: Group Theory and Group Skills*", Allyn and Bacon.

- Joiner, R. (1995). "The Negotiation of Dialogue Focus: An Investigation of Dialogue Processes in Joint Planning in a Computer Based Task". In Claire O'Malley (ed.) *Computer Supported Collaborative Learning*, Nato ASI Series, 203-222
- Katz, S., O'Donnell, G., and Kay, H. (2000). "An Approach to Analyzing the Role and Structure of Reflective Dialogue". *International Journal of Artificial Intelligence in Education*, 11, 320-343
- Kolodner, J. L. and Simpson, R. L. (1989). "The MEDIATOR: Analysis of an Early Case-Based Problem Solver". *Cognitive Science*, 13, Ablex Publishing Corporation, 507-549
- Lawson, R. B., and Shen, Z. (1998). *Organizational Psychology: Foundations and Applications*, Oxford University Press
- Levin, J. A., and Moore, J. A., (1977). "Dialogue-Games: Metacommunication structures for Natural Language Interaction". *Cognitive Science*, 395-420
- McManus, M. M. and Aiken, R. M. (1995). "Monitoring Computer-Based Collaborative Problem Solving". *Journal of Artificial Intelligence in Education*, 6(4), 307-336[Moore 1987] Moore, C. W. (1987). *The Mediation Process: Practical Strategies for Resolving Conflict*, Jossey-Bass
- Mühlenbrock, M., and Hoppe, H.U., (1999), "Computer Supported Interaction Analysis of Group Problem Solving". In Hoadley, C. and Roschelle, J. (eds.); *Proceedings of the Conference Computer Support for Collaborative Learning, CSCL '99*, 398-405
- Mühlenbrock, M., and Hoppe, H. U., (2001). "A Collaboration Monitor for Shared Workspaces". Accepted for presentation at AIED'2001, San Antonio, Texas
- Müller, H. J., and Dieng, R., (2000). "On Conflicts in General and their use in AI in Particular". In Müller, H. J., and Dieng, R. (eds) *Computational Conflicts: Conflict Modeling for Distributed Intelligent Systems*, Springer Verlag, 1-20
- Paiva, A., (1997). "Learner Modelling for Collaborative Learning Environments". In du Boulay, B. and Mizoguchi, R. (eds.) *Proceedings of the 8th World Conference on Artificial Intelligence in Education – Knowledge and Media in Learning Systems*. IOS Press, 215-222
- Pilkington, R. M., Hartley, J. R., and Hintze, D. (1992). "Learning to Argue and Arguing to Learn: An Interface for Computer-Based Dialogue Games". *Journal of Artificial Intelligence in Education*, 3(3), 275-295
- Pilkington, R. M., (1999). *Analysing Educational Discourse: the DISCOUNT Scheme*. Technical Report 99/2, Computer Based Learning Unit, The University of Leeds, ISBN: 1901418022
- Quignard, M. and Baker, M., (1999). "Favouring Modellable Computer-Mediated Argumentative Dialogue in Collaborative Problem Solving Situations". Lajoie, S., and Vivet, M., *Artificial Intelligence in Education – Open Learning Environments: New Computer Technologies to Support Learning, Exploration and Collaboration*, IOS Press, 129-136
- Rao, A. and Georgeff, M. P., (1995). "BDI Agents: From Theory to Practice". In *Proceedings of the First International Conference in Multi-Agent Systems – ICMAS'95*. Also available as Technical Note 56, AAIL, April 1995.
- Robbins, S. P., (1998). *Organizational Behaviour: Concepts, Controversies and Applications*, London, Prentice-Hall
- Self, J. A., (1995). "Dormobile: a Vehicle for Metacognition". In Chan, T.W. and Self, J. A. (eds.) *Emerging Computer Technologies in Education*, Charlottesville, AACE
- Self, J., Karakirik, E., Kor, A., Tedesco, P. and Dimitrova, V. (2000). "Computer-Based Strategies for Articulate Reflection (and Reflective Articulation)". In Young, S. S., Greer, J., Maurer, H. and Chee, Y. S. (eds.) *Proceedings of the International Conference on Computers in Education/International Conference on Computer Assisted Instruction*, National Tsing Ha University, Taiwan, 3-12
- Silbey, S. S. and Merry, S. E. (1986). "Mediator Settlement Strategies", *Law and Policy*, 8(1),7-32
- Soller, A., Goodman, B., Linton, F. and Gaimari, R., (1998). "Promoting Effective Peer Interaction in an Intelligent Collaborative Learning System". In Goettl. B.P., Halff, H.M., Redfield, C.L., and Shutte,

- V.J. (eds.), *Proceedings of the 4th International Conference on Intelligent Tutoring Systems, ITS'98*, Lecture Notes in Computer Science 1452, Springer Verlag, 186-195
- Soller, A., Linton, F., Goodman, B., and Lesgold, A., (1999). "Toward Intelligent Analysis and Support of Collaborative Learning Interaction". In Lajoie, S. P., and Vivet, M., (eds.) *Artificial Intelligence in Education – Open Learning Environments: New Computer Technologies to Support Learning, Exploration and Collaboration*", IOS Press, 75-83
- Tedesco, P. and Self, J.A. (1998). *Towards a computational model of meta-cognitive conflicts: a preliminary investigation*. Technical Report 98/28, Computer Based Learning Unit, University of Leeds
- Tedesco, P. and Self, J. A. (2000a). "Using Meta-Cognitive Conflicts to Support Group Problem-Solving". In Gauthier, G., Frasson, C., and VanLehn, K. (eds.) *Proceedings of the 5th International Conference Intelligent Tutoring Systems (ITS'2000)*, Lecture Notes in Computer Science, 1839, Springer Verlag, 232-241.
- Tedesco, P. and Self, J.A. (2000b). "MArCo: using Using Meta-Cognitive Conflicts to Provoke Strategic Changes". In Sichman, J. and Monard, M.C. (eds.) *Proceedings of the International Joint Conference 7th Ibero-American Conference in AI 15th Brazilian Symposium on AI - IBERAMIA-SBIA'2000*, Lecture Notes in Artificial Intelligence 1952, Springer Verlag, 186-195
- Tedesco, P. (2001a) *Mediating Meta-Cognitive Conflicts in Group Planning Situations*. Unpublished Doctoral Dissertation, CBLU,, University of Leeds.
- Tedesco, P. (2001b) "Using Group Models to support Group Planning Interactions in MArCo, paper accepted at the XII Brazilian Symposium on Informatics in Education, Vitória, Espírito Santo, November 21-23
- Tjosvold, D. (1993). *Learning to Manage Conflict – Getting People to Work Together Productively*, Lexington Books
- Walton, R. E. (1987). *Managing Conflict: Interpersonal Dialogue and Third-Party Roles*. Addison-Wesley Publishing Company.
- Weldon, E. and Weingart, L. R. (1993). "Group Goals and Group Performance". *British Journal of Social Psychology*, 32, 307-334
- Winston, W. L., (1991). *Operations Research – Applications and Algorithms*, PWS-Kent Publishing Company
- Yang, Q. (1997). *Intelligent Planning – a Decomposition and Abstraction Based Approach*. Springer Verlag