Belief Systems for Conflict Resolution

Massimiliano Garagnani [†] **, Maria Fox** [†] **, Derek P. Long** [†]

Abstract. In a multiagent system, conflicts may arise because of two basic reasons: 1) different agents have contrasting *goals*; 2) different agents have inconsistent *knowledge.* This situation can originate when agents are autonomous and strongly motivated by their own interests, and when heterogeneous agents, with different skills, 'histories' and beliefs, coexist in a dynamic environment.

This paper is focused on the *representation* of two agents' epistemic conflicts, originated by the presence of contrasting *beliefs*. In such a situation, for an agent 'S' to agree on a cooperative plan with another agent 'H' it is necessary to *communicate* and solve the conflicts which could prevent the realisation of the specific plan identified. In order to do this, S must be able to *reason* about H's beliefs and goals. This requires S to maintain an adequate representation of H's mental states and to use this *model* to develop a plan for communication which will persuade H to adopt the cooperative goal.

In this paper we introduce and formalise a *discourse-generation* oriented model of beliefs which uses the concept of 'event' as fundamental unit for the knowledge representation. The basic features of the model defined include a qualitative degrees-of-belief approach to represent *uncertain* beliefs, two-level nested *mutual beliefs* expressions and beliefs *grounding* through *rationality* and *experience*.

1 INTRODUCTION

In a multiagent environment, communication between agents becomes necessary when agents need to *cooperate*. If the system allows the agents to collect incomplete or uncertain information, or to adopt different goals, agents' knowledge and/or beliefs may be inconsistent, and cooperation cannot be relied upon. Hence, any communication between agents which is aimed to organise the execution of a cooperative task will require an initial process of conflicts *check* and *resolution*, and the ability of an agent to understand and reason about another agent's beliefs is vital for the success of this process.

This ability requires an agent to adopt and maintain an *internal model* of another agent's knowledge and beliefs. As Grosz and Sidner have already pointed out in [5], "any model (or theory) of the communication situation must distinguish among beliefs and intentions of different agents".

Although research in modal logic has produced many theoretical results in the field of belief systems (e.g.[6][10]), the applications of these theories to other domains (such as conflictual multiagent systems) have been, to a certain extent, neglected.

In this work, we introduce and formalise a model of beliefs which allows mutual beliefs representation, grounding of beliefs and uncertainty management. The characteristics of the belief system are *discourse-generation* oriented. The model has been developed to enable an agent 'S' (Speaker) to plan for the resolution of possible epistemic conflicts with another agent 'H' (Hearer) through the generation of a *persuasive discourse* (monologue) aimed to convince H to adopt a cooperative goal.

Hence, the architecture described is designed to capture the relevant information which are necessary to construct persuasive arguments, and to make these information easily available and updatable.

[†] Department of Computer Science, Durham University, South Road, Durham, DH1 3LE, UK.
Massimiliano.Garagnani@durham.ac.uk;
D.P.Long@durham.ac.uk; Maria.Fox@durham.ac.uk.

2 BASIC FEATURES

2.1 Event based knowledge

The beliefs representation of the model has, as a fundamental unit, the concept of *event*. The conceptual unit of event is commonly used in semantic networks to express *sentences*, and represents the equivalent of a proposition in predicate logic. For example, the sentence "John likes apples", which in predicate logic can be expressed using a structure of the type *Action(subject, object)* ("Like(John,apples)"), would be represented, as a semantic network event, as in figure 1.

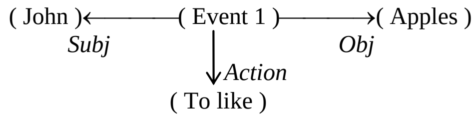

Figure 1 Semantic network for "John likes apples".

The constituent elements of an event are *Subject, Action* and *Object,* which may also be absent (as in "The sun shines"); moreover, if necessary, further features can be added, such as the time or spatial location of the event, or the destination of its object.

In particular, the form of event representation of fig. 1 constitutes the basis for the semantic network on which the natural language processing system LOLITA depends [9]. In our model of beliefs (and also in the LOLITA system), the event formalism is recursive, hence events can themselves be subject/object of other events. For example, if E_1 is an event, 'Believes(Agent,E_1)' is the event equivalent to the proposition "Agent believes that E_1".

We considered the *knowledge* of an agent as a collection of events which evolves through time. Each event belonging to this set constitutes a *belief* currently held by the agent. Therefore, an agent's belief system will consist of a set of events.

2.2 Support and grounding of beliefs

In logic, given a proposition *p* and a formal theory, if *p* is true (i.e. it is a theorem of that theory) then *p* must either be a consequence of inference rules and other propositions which are true, or one of the axioms of the theory. Similarly, in a rational belief system, the presence of a belief must be *justified*: if an agent believes E, than either E is a consequence of his reasoning about other beliefs, or it has been accepted as *grounded* belief, i.e. a belief which is not questionable (see also [1],[8]).

If we adopt a 'cynical' point of view, the only beliefs which we do not normally question are those deriving from our direct experience of the physical reality, i.e. our *perceptions* of the external world. For example, if an agent experiences the event E="I see a table", then the agent will strongly believe E.

Therefore, in the model of beliefs, we included explicit notations to represent the two mentioned aspects - *reasoning* and *experience* - of belief justification.

2.2.1 Justification by reasoning

If a belief in E is a consequence of reasoning, then there must be another event E_1 in the beliefs set such that the belief in E_1 *supports* the belief in E. This relationship is formally written as

```
SUPPORT(E1,E)
```

This means that the belief in E_1 is a *sufficient* reason to believe E, but it does not imply that either E_1 or E have to be believed.

It is important to notice that SUPPORT(E_1 ,E) (abbreviated with SUP(E_1,E) from now onwards) is still an event, whose subject and object are, respectively, the event E_1 and the event E. This implies that its presence in the set of beliefs must be justified, either using a further support event (which, in turn, needs a justification), or assuming it to be a grounded belief. This underlines the necessity of using the belief grounding through experience as condition to terminate this potentially infinite recursion of beliefs justification.

2.2.2 Justification by experience

If the agent A believes E because (s)he has experienced it (i.e. (s)he was involved in E, either as a spectator, or as the subject or the object) then the event E does not need further support to be believed, and we shall formally write

```
REAL(A,E)
```

or, equivalently, A_REAL(E). This indicates that E is a *real experience* for the agent A. Even in this case, A_REAL(E) is again an event, but its presence in the set of beliefs does not require any further justification, neither does the presence of the event E.

It should be noticed that if 'E' is a support event, the belief A_REAL(E) should be interpreted as 'inductive experience', rather than real experience. Indeed, a support relation as SUP(E_1,E) cannot be 'experienced' physically, but one can *repeatedly* experience the fact that the happening of E_1 produces always the happening of E, and conclude, through *induction*, that this is a causal relation which holds in general.

2.3 Uncertainty of beliefs

The system of beliefs that we defined allows the inclusion of uncertain beliefs within the model. The model adopts a qualitative tripartite *degrees-of-beliefs* approach (see also [2][3]). We devised a simple model in which every event must be in *one and only one* of the following categories: *Unbelieves, Undecided, Believes*[1]. These degrees are used as predicates (or *modal operators*) to build new events. Moreover, we assumed the two following axioms to hold:

A1) *Unbelieves*(A,E) $\Leftrightarrow$ *Believes* (A, NOT(E))
A2) *Undecided* (A,E) $\Leftrightarrow$ *Undecided*(A, NOT(E))

Hence, the *exhaustive* and *mutually exclusive* opinions which might be held towards a generic event E by an agent A are essentially three: A_BEL(E), A_BEL(NOT(E)) and A_UND(E). For a similar analysis of the degrees of beliefs and supporting arguments, see [11].

In order to separate the model of the speaker agent 'S' from the rest of events which are supposed to be *unknown* by the other agent (the hearer, H), we introduced the modal operator

```
H_UNK(E)
```

where H stands for Hearer. This expression means that (the speaker believes that) the hearer *does not know* about the event E, i.e. E has never been *conceived* as a 'thought' by H's mind.

This situation might occur when the speaker has some knowledge which (s)he believes the hearer can not possibly have. The symmetrical form, S_UNK(E) (namely, "Speaker unknows E"), could be present *only* as one of the hearer's beliefs, i.e. H_BEL(S_UNK(E)). In fact, the expression A_BEL(A_UNK(E)) is contradictory for any agent A and for any event E: as long as the agent A simply 'thinks' about the concept A_UNK(E), this becomes immediately false by definition.

Finally, another axiom of equivalence, analogous to A2, has been introduced for this last predicate:

A3) *Unknown* (A, E) $\Leftrightarrow$ *Unknown* (A, NOT(E))

This corresponds to the assumption that if an agent A has never thought about an event E, then neither can (s)he have conceived its negation, and viceversa.

3 LANGUAGE FORMALISATION

3.1 The *outer* language

The syntax for the well-formed belief expressions of the system is defined by the following BNF production rules, where the initial symbol is <WFE> ('Well Formed Expression'):

```
<WFE> ::= <Pred>(<A>) | NOT (<Pred>(<A>))
<PRED>::= S_BEL | S_UND | H_BEL | H_UND
                                | H_UNK
<A>   ::= <E> | NOT (<E>)
<E>   ::= S_REAL(<Es>)| H_REAL(<Es>)| <Es>
<Es>  ::= SUP (<En>, <En>) | <Ei>
<En>  ::= <Es> | NOT (<Es>)
<Ei>  ::= E0 | E1 | E2 | E3 | ... | En
```

Before analysing the characteristics of this syntax and explaining why it generates an *outer* language, let's see, with an example, how it can be used to describe a simple conflictual situation.

Given the three possible opinions which an agent A can hold towards an event 'E' (namely, A_BEL(E), A_BEL(NOT(E)) and A_UND(E)), an epistemic conflict between a speaker (S) and a hearer (H) can be defined as follows:

Definition Given a set S of expressions generated by <WFE>, S contains a *conflict of beliefs* iff $\exists$ an event 'E' such that either {S_BEL(E),H_BEL(NOT(E))} or {S_BEL(NOT(E)),H_BEL(E)} is a subset of S.

Let's consider, for simplicity, a situation in which conflicts between two agents, say Sally (S) and Harry (H), arise because of their disagreement about one event, which we shall identify with E_1. Let E_1 be "The best direction for us to go is left", E_2 be "There is an apple tree at our left", and E_3 be "The road turns to the right". Let's assume Sally to have actually already *located* the apple tree through her sensorial perceptions (e.g. sight), and thus to believe in E_2 as a *real* experience. On the other hand, the agent Harry has not seen the tree yet, and still believes that the best thing to do in order to find some food is to follow the road, which, at the point where they are, turns to the right.

The other important facts which need to be specified consist of the attitudes of the two agents towards the three following support relationships:

A) SUP(E_2,E_1), expressing the fact that the presence of an apple tree is a *sufficient* reason for them to believe that they should move towards it. It is plausible to assume that both of the agents believe this relation if, for

[1] LOLITA's plausible inference engine works with a further refinement of these levels of certainty: {none, veryLow, low, medium, high, veryHigh, total}.

example, they are *cooperating* in a situation in which their common goal is to find some food to eat.

B) SUP(E_3,NOT(E_1)), representing the fact that if the road turns to right, it would not be sensible to move left.

C) SUP(E_2,NOT(SUP(E_3,NOT(E_1)))), representing the fact that the belief in the presence of an apple tree at their left *invalidates* the support relation B.

From the point of view of the speaker (Sally), who will try to convince Harry to move left, the conflictual situation could be formulated as follows, using the <WFE> language and abbreviating the expression “NOT( )” with “¬”:

S_REAL(E_3)	H_REAL(E_3)
S_BEL(E_3)	H_BEL(E_3)
S_BEL(¬SUP(E_3,¬E_1))	H_BEL(SUP(E_3, ¬E_1))
S_REAL(E_2)	H_UNK(S_REAL(E_2))
	H_UNK(¬S_REAL(E_2))
S_BEL(E_2)	H_UND(E_2)
	H_UND(¬E_2)
S_BEL(SUP(E_2,¬SUP(E_3,¬E_1)))	
	H_BEL(SUP(E_2,¬SUP(E_3,¬E_1)))
S_BEL(SUP(E_2,E_1))	H_BEL(SUP(E_2,E_1))
S_BEL(E_1)	H_BEL(¬E_1)

It should be noted that all the beliefs of the right column, i.e. those concerning Harry, should be thought as having the predicate S_BEL as prefix, as this list of propositions constitutes Sally’s set of beliefs.

Notice that whilst both of S and H believe the supports A and C, only Harry still believes B, whereas Sally believes in the opposite thesis: S_BEL(¬SUP(E_3,¬E_1)).

Graphically, this situation can be represented as in fig. 2, where the bold circles around E_2 and E_3 mean that these are real event for S, and the short arrows, going from E_2 to E_1 and from E_3 to ¬E_1, represent, respectively, the support relationships A and B.

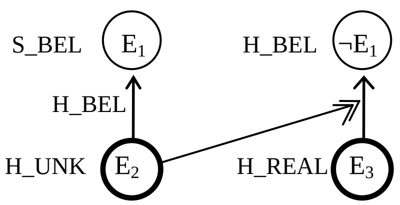

Figure 2 Graphical representation of the conflictual situation.

The arrow with ‘double tip’, representing the support C, expresses the fact that E_2 supports the negation of the event pointed (which, in this case, is a support) and can be defined as *attack* relationship. This relation is believed by both of the agents. The expressions adjacent to the four events indicate the main H’s and S’s attitudes towards the basic elements of the picture.

Finally, it should be noticed that, with respect to E_2, although Harry does not know that this event is a *real* experience for Sally (H_UNK(S_REAL(E_2)), the event E_2 itself is not unknown to him. This means that Harry has *considered*, at least once, in the past, the event E_2, but, at the moment, is “distracted”, and his attitude towards E_2 is of *indecision* (H_UND(E_2)).

3.2 Redundancy of the *outer* language

The example underlines an important point: there are many propositions (such as NOT(H_BEL(E_2)), NOT(H_BEL(E_1)), NOT(S_BEL(NOT(E_1)), etc.) which, although intuitively *true* and syntactically correct <WFE> expressions, have not been included in the list of beliefs. In fact, they should not, for they can be *deduced* directly from the other beliefs explicitly stated.

In other words, the system should be able to ‘evaluate’ these propositions as ‘true’ even if they are not present in the current set of beliefs. Furthermore, even the expression S_BEL(E_2), which has been stated explicitly, could have been deduced as a direct consequence of the belief S_REAL(E_2), which implies that S considers E_2 a real experience, and thus should strongly believe it.

Finally, because of the axiom A2 and A3 introduced before, the two expressions H_UND(E_2), H_UND(NOT(E_2)) are to be considered equivalent, and so are H_UNK(E), H_UNK(NOT(E)), where E = S_REAL(E_2). In other words, the language defined by <WFE> is, in many cases, ‘redundant’.

In order to avoid double checks, or the management of equivalent expressions, we adopted - for each case of ambiguity - one of the two forms as *canonical*, reducing to this one all the occurrences of the other, as the syntactical transformations below show.

T1) S_UND(NOT(e)) ⇔ S_UND(e)
T2) H_UND(NOT(e)) ⇔ H_UND(e)
T3) H_UNK(NOT(e)) ⇔ H_UNK(e)
T4) S_BEL(a) ⇔ a

Each left-hand side constitutes a non-canonical form which is automatically transformed into the corresponding (canonical) right-hand side.

In addition, to avoid the need to state ‘superfluous’ beliefs which are either i) a *direct* consequence of other beliefs, or ii) beliefs whose truth can be *deduced* from the presence/absence of related beliefs, we introduced, respectively, i) a set of *inference rules,* and ii) a list of *boolean functions* associated to specific expressions.

3.2.1 Inference rules

The list of inference rules adopted is shown in the list below:

```
I1) S_R(e)        |- e
I2) H_R(e)        |- e, H_BEL(H_R(e))
I3) NOT(H_R(e))   |- H_BEL(NOT(H_R(e))
I4) H_BEL(S_R(e)) |- H_BEL(e)
I5) H_BEL(H_R(e)) |- H_BEL(e)
```

It should be noted that S_REAL and H_REAL have been abbreviated with S_R and H_R, while the variable ‘e’ represents any event generated by the symbol <Es> of the <WFE> syntax. For example, the rule I3 means that if in the current set there is a belief which matches the left hand side NOT(H_REAL(e)), then the belief H_BEL(NOT(H_REAL(e))) should be added to the set.

All the inference rules should - in principle - be applied recursively to the set of beliefs every time this is modified, to produce a new, larger set of beliefs, containing the original one[2].

Considering the syntax of the <WFE> expressions and the characteristics of the five inference rules adopted, is not difficult to prove the following result:

Definition The extension *Ext(S)* of a set *S* of <WFE> expressions is the set of <WFE> expressions obtained applying recursively the inference rules I1-I5 to the set *S.*

Theorem 1 Given a finite set of <WFE> expressions *S,* the set *Ext(S)* is finite.

[2] Nevertheless, to eliminate this computational load, the system implemented actually performs this operation only once, as a pre-compilation process.

Proof First of all, it should be noticed that each of the two arguments of a SUP( ) event is allowed to be another SUP( ) event, yielding to expressions with an arbitrarily high level of nested supports. Since this infinite recursion is syntactically permitted only in the case of a SUP( ) event, there are only two ways of obtaining an infinite set of beliefs $\subseteq$ <WFE>: 1) to generate all the possible SUP( ) expressions; 2) to generate an infinite number of primitive events <Ei> = $E_0, E_1, ..., E_n,...$.

However, considering the inference rules I1-I5, none of them increases the level of support nesting or introduces new primitive events Ei, and, therefore, none of them can be used to extend a finite initial set of <WFE> expressions into an infinite one.

#

3.2.2 Boolean functions

Let's consider, as an example, the belief S_UND(E_4). The truth of this belief can be determined simply assuming it to hold iff neither S_BEL(E_4) nor S_BEL(NOT(E_4)) holds. Thus, in general, any belief such as S_UND(e) does not need to be explicitly stated, but can simply be associated with a boolean function in the following manner:

$$\text{SUND(e)} \Leftrightarrow \neg\ (\text{S\_BEL(e)} \vee \text{S\_BEL(NOT(e))})$$

The symbol '¬' should be read as the normal negation '*not*' in classical logic. The boolean expression at the right-hand side will return *False* only when one of the two propositions S_BEL(e), S_BEL(NOT(e)) is present in the current set of beliefs.

This mechanism, which we defined in general as *belief by default,* allows the association of a specific expression with a boolean function whose truth depends on the presence/absence of other beliefs. It is important to underline that without the introduction of the above mechanism, it would not have been possible, for the system, to represent 'completely' a beliefs situation. To show this, let's consider the following example.

Let S = {S_BEL(E_1), S_BEL(E_2)} be the (extended) set of beliefs currently held by a beliefs system which does not make use of boolean functions. The system's answer to a 'query' such as "Does S believe SUP(E_1,E_2)?" would reasonably be 'False', as - according to the Closed World Assumption (CWA) - since the proposition S_BEL(SUP(E_1,E_2)) is not explicitly asserted in the set of beliefs, then it should be considered false.

The same answer would have been given had the query consisted of the proposition S_BEL(NOT(SUP(E_1,E_2))). Since the three mutually exclusive and exhaustive opinions that the speaker can have towards any event 'E' are S_BEL(E), S_BEL(NOT(E)) or S_UND(E), one would expect that, by exclusion, for the event E = SUP(E_1,E_2) it must be S_UND(E). Unfortunately, because of the CWA, the system would return 'False' even for S_UND(SUP(E_1,E_2)).

The addition of the above proposition to S would merely shift the problem, as one could interrogate the system again with SUP(E_1, SUP(E_1,E_2)), and so forth, by increasingly high levels of nesting of support.

Therefore, the adoption of the default mechanism is necessary in order to describe an *infinite* number of beliefs using only a *finite* set of expressions, and, thus, to represent completely and correctly the expected belief state of the agent (speaker).

3.3 The *inner* language

The addition of the above features (syntactical transformations of equivalent expressions, extension of the beliefs set through recursive application of inference rules and association of beliefs to boolean functions of other beliefs) leads to a necessary distinction between the language <WFE> and the simpler '*inner'* language obtained removing all of the superfluous expressions from <WFE>. This simplified, unequivocal inner language should be used to state explicitly the expressions contained in the set of beliefs, while the *outer* <WFE> syntax should be adopted only as a language to 'query' the belief system.

The BNF definition of this unambiguous, simpler inner language is given below by the set of production rules for the initial symbol <WFF>.

```
<WFF> ::= <A> | H_BEL(<A>) | H_UND(<Eh>)
<Eh>  ::= S_REAL(<Es>) | <Es>
```

Notice that the non-terminal symbols <A> and <Es> are assumed to be those defined in <WFE>. As expected, no expressions such as S_UND( ) or H_UND(NOT( )) need to be specified when using this syntax.

4 PLANNING COMMUNICATION FOR CONFLICTS RESOLUTION

The model of beliefs hitherto described has been used as a basic component of a discourse planning system which generates *plans* for *persuasive* communication.

During the process of planning a discourse, the set of beliefs held by the system represents the current *mental state* of the speaker, and is treated by the planner as the normal *state* containing the current description of the world. The aim of the planner is to find a successful plan, i.e. a sequence of actions (*operators*) that transform the initial state (set of beliefs) into a new mental state, in which the specific (communicative) goal assigned is achieved.

Each operator takes part in the transformation *adding* and *deleting* beliefs from the set as a direct consequence of the specific 'speech act' that it performs.

In the previous example, in which S and H had conflictual beliefs regarding the event E_1, the plan which the discourse planning system - developed integrating the model of beliefs with a planner based on the recent 'graphplan' technique [7] - produced in order to achieve the goal 'H_BEL(E1)' is showed in fig. 3.[3]

The propositions written at the right side of the steps show the hearer's change of attitude towards the relevant events.

In the first step, the opinion towards the event E_2 changes from H_UND(E_2) to H_BEL(E_2). This is realised by the operator *assert-real,* whose task consists of asserting that the event E_2 is a real experience for the speaker S. This assertion is expected to produce two results on the hearer beliefs: first of all, since the event S_REAL(E_2) has been formulated and communicated, it will not be anymore an *unkown* event. In case of a generic event *e*, the new attitude of H towards it would be H_UND(*e*). However, since E_2 is claimed by S to be a *real* experience, H (who is supposed to assume the speaker to be *sincere*) shall believe that S is actually experiencing (or has experienced) E_2 as part of the reality. Hence, the second result of this assertion is that H is expected to *believe* S_REAL(E_2), and, as a consequence of the inference rule I4 (indicated with the note [←I4]), he is also expected to believe E_2 itself.

Notice that, in a real situation, the validity of the event E_2 could be verified by H through his own physical perceptions, without having to rely on S's sincerity: H could simply turn his head, and check if the apple tree is actually there. Nonetheless, this is not always possible: in this case, for example, Harry might have some problems with long-distance vision.

The second step uses the belief (not shown in fig. 3 but held by both H and S) which concerns the support relationship between E_2 and E_1. As indicated in figure 2, H and S believe that SUP(E_2,E_1), i.e. that believing in E_2 is a sufficient reason to

3 For a more detailed description of the the integration of the system within the planner , see [4].

believe E_1. Therefore, the belief H_BEL(NOT(E_1)) becomes uncertain - H_UND(E_1) - through the application of the *attack* operator, which points out to H that believing in E_2 and in the support relationship SUP(E_2,E_1) gives credibility to E_1, and therefore 'undermines' H's certainty about the belief in NOT(E_1).

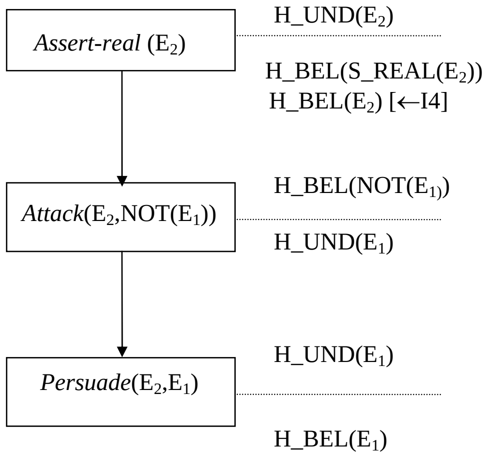

Figure 3 A simple discourse plan..

The final step reuses the same support relation: as effects of the first two steps, we have H_UND(E_1), H_BEL(E_2). Thus, from H_BEL(SUP(E_2,E_1)) and H_BEL(E_2), the *persuade* operator simply applies the *Modus Ponens* rule, and deduces H_BEL(E_1).

Notice that this last operator could not have been applied directly as second step, because it requires the hearer to be currently undecided about the event-conclusion, E_1. This produced the previous process of 'spoiling' of the belief in NOT(E_1) through the *attack* operator.

A possible natural language translation of this simple plan could correspond to the following monologue, directed from Sally to Harry:

"I can see an apple tree at our left".

"Thus, you should not anymore be convinced of the fact that we should *not* move left".

"Indeed, since you believe that there is an apple tree at our left, *and* you believe that this supports the thesis that we should move left, *then* you'll have to agree that we should move left."

The schemes of the three operators employed to solve this simple conflictual situation are shown in figure 4. It has to be noticed that even if the *precondition* list of an operator is allowed to contain any set of *outer* <WFE> expressions - which identify the situation in which the operator can be applied - the *add* and *delete* lists have to specify only the *essential* <WFF> (inner) beliefs, as the missing propositions and implicit consequences are added *automatically* before the planning process.

Assert-real(e::<Es>)

Preconditions:	S_REAL(e)
Add list:	H_BEL(S_REAL(e))
Delete list:	

Attack(e_1::<A>, e_2::<A>)

Preconditions:	H_BEL(NOT(e_2)), H_BEL(e_1), H_BEL(SUP(e_1 ,e_2))
Add list:	H_UND(e_2)
Delete list:	H_BEL(NOT(e_2))

Persuade(e_1::<A>, e_2::<A>)

Preconditions:	H_UND(e_2) H_BEL(e_1), H_BEL(SUP(e_1 ,e_2))
Add list:	H_BEL(e_2)
Delete list:	H_UND(e_2)

Figure 4 *Assert-real( )*, *Attack( )* and *Persuade( )* operators' schemes.

In fact, a pre-processing algorithm, which operates as interface between the belief system and the planner (allowing the use of these two modules in sequence), has been implemented to automatically produce a *completed* version of each operator scheme, through a *pre-compilation* of the effects of the inference rules, of the equivalent expressions and of the default mechanism into the operators themselves. The same pre-compilation is performed on the initial state, which is transformed into a (finite) set of <WFE> expressions containing *all* the outer language beliefs which would have been evaluated *true* by the beliefs system[4].

The automatically completed and extended version of the original operator *Assert-real( )* is shown in figure 5.

Assert-real(e::<Es>)

Preconditions:	SR(e)	
Add list:	HBEL(SR(e))	HBEL(e)
	NOT(HBEL(NOT(SR(e))))	NOT(HBEL(NOT(e)))
	NOT(HUND(SR(e)))	NOT(HUND(e))
	NOT(HUND(NOT(SR(e))))	NOT(HUND(NOT(e)))
	NOT(HUNK(SR(e)))	NOT(HUNK(e))
	NOT(HUNK(NOT(SR(e))))	NOT(HUNK(NOT(e)))
Delete list:	NOT(HBEL(SR(e)))	NOT(HBEL(e))
	HBEL(NOT(SR(e)))	HBEL(NOT(e))
	HUND(SR(e))	HUND(e)
	HUND(NOT(SR(e)))	HUND(NOT(e))
	HUNK(SR(e)))	HUNK(e)
	HUNK(NOT(SR(e)))	HUNK(NOT(e))

Figure 5 Extended and completed *Assert-real( )* operator.

Note that, for simplicity, the expression S_REAL( ) has been abbreviated with SR( ), while H_BEL, H_UND and H_UNK have been reduced to HBEL, HUND, HUNK.

5 CONCLUSIONS

In a multiagent system in which agents are driven by their own intentions and beliefs, the management and resolution of epistemic conflicts amongst agents demands the development of rich communication in order to allow an agent to *convince* another agent to adopt a cooperative goal.

In this paper, we have described the characteristics of the two-languages (inner/outer) belief system which has been designed to represent conflictual beliefs situations between two agents, a speaker and a hearer. This structure has actually been implemented and used as a base to build a discourse planning system which can generate plans for epistemic conflicts resolution through *persuasive* communication.

---

[4] The actual set of beliefs evaluated true by the system contains an inifinite number of propositions, due to the arbitrary depth of nesting of the support events. This set has been reduced to a finite set eliminating all the events containing a SUP( ) expression with level of nesting greater than the maximum level present in the initial state.

It is important to note that the generation of a discourse plan constitutes only a fragment of the whole process of conflicts resolution, which should also take into account the ‘answer’ given by the hearer in reply to the message generated by the speaker. This answer, which might itself consist of a discourse, has to be analysed by the speaker and used to *feed-back* the belief system, which should be checked against its contents in order to reveal a possible failure of part of the communication plan.

If some of the expected effects of the communication have not been achieved, it is necessary to perform a new planning process, followed by a new communication, which will produce, in turn, a new reply from the hearer. This sequence of alternate ‘monologues’, exchanged between H and S, will eventually lead to the development of an effective *cooperative dialogue*, aimed to the resolution of the conflicts which prevent the agreement on a common goal.

6 BIBLIOGRAPHY

[1] Ackermann, R.J. (1972) *Belief and Knowledge*, Anchor, New York.

[2] Forrest, P. (1986) *The Dynamics of Belief*, Basil Blackwell.

[3] Fox, J. (1986) “Knowledge, Decision Making and Uncertainty”, in Gale, W.A., ed. *Artificial Intelligence and Statistics*, Addison-Wesley, Reading, MA.

[4] Garagnani, M. (1997) “Belief Modelling for Discourse Plans”, *Proceedings of the 16th Workshop of the UK Planning and Scheduling Special Interest Group 1997*.

[5] Grosz, B.J., Sidner, C.L. (1990) “Plans for discourse”, in Cohen P., Morgan J., Pollack M.E., eds. *Intentions in Communication* (MIT Press, Cambridge, MA), pp.417-444.

[6] Halpern, J.Y., Moses, Y. (1992) “A guide to completeness and complexity for modal logics of knowledge and belief”, *Artificial Intelligence* **54**:319-379.

[7] Koehler, J., Nebel, B., Hoffmann, J., Dimopoulos, Y. (1997) “Extending Planning Graphs to an ADL Subset”, *Proceedings ECP*-97.

[8] Reed, C., Fox, M., Long, D.P., Garagnani, M. (1996) “Persuasion as a Form of Inter-Agent Negotiation”, *Lecture Notes in Artificial Intelligence* **1286**:120-136.

[9] Smith, M.H., Garigliano, R., Morgan, R. (1994) “Generation in the LOLITA system: an Engineering Approach”, *Proceedings of the 7th International NL Generation Workshop*, Maine.

[10] Smith, R.S. (1991) “Modal logic”, *Artificial Intelligence Review* **5**, pp.5-34.

[11] Young, R.M., Moore, J.D., Pollack, M.E. (1994) “Towards a Principled Representation of discourse Plans”, *Proceedings of the Sixteenth Annual Meeting of the Cognitive Science Society*, Atlanta, GA.