Solving Conflicting Beliefs with a Distributed Belief Revision Approach

Benedita Malheiro[1,2] and Eugénio Oliveira[1,3]

[1] LIACC - Artificial Intelligence and Computer Science Laboratory
http://www.ncc.up.pt/liacc/index.html
[2] ISEP, DEE, Rua António Bernardino de Almeida, 4200 Porto,Portugal
mbnm@fe.up.pt
[3] FEUP, DEEC, Rua dos Bragas, 4099 Porto CODEX, Portugal
eco@fe.up.pt

Abstract. The ability to solve conflicting beliefs is crucial for multi-agent systems where the information is dynamic, incomplete and distributed over a group of autonomous agents. The proposed distributed belief revision approach consists of a distributed truth maintenance system and a set of autonomous belief revision methodologies. The agents have partial views and, frequently, hold disparate beliefs which are automatically detected by system's reason maintenance mechanism. The nature of these conflicts is dynamic and requires adequate methodologies for conflict resolution. The two types of conflicting beliefs addressed in this paper are Context Dependent and Context Independent Conflicts which result, in the first case, from the assignment, by different agents, of opposite belief statuses to the same belief, and, in the latter case, from holding contradictory distinct beliefs.

The belief revision methodology for solving Context Independent Conflicts is, basically, a selection process based on the assessment of the credibility of the opposing belief statuses. The belief revision methodology for solving Context Dependent Conflicts is, essentially, a search process for a consensual alternative based on a "next best" relaxation strategy.

1 Introduction

Multi-Agent Systems (MAS) are a natural environment for the occurrence of information conflicts - frequently, the agents hold contradictory beliefs. These disparate perspectives can either be reconcilable or incompatible, falling, respectively, into the categories of negative and positive conflicts [10]. In particular, this paper is concerned with solving negative conflicts that occur when agents hold contradictory beliefs.

We elected distributed belief revision as the adequate approach for detecting conflicts, maintaining the consistency of the agents' knowledge base and trying to solve the detected conflicts. This view of belief revision as a truth maintenance process followed by a selection mechanism of one (or more) preferred solutions was proposed by [2]. We developed two belief revision methodologies to solve: (i) *Context Independent Conflicts*, which occur when a distributed belief is, simultaneously, *believed* by some agents and *disbelieved* by others; (ii) *Context Dependent Conflicts*, which occur when the agents detect inconsistent distinct beliefs. In the first case, the goal is to choose the most appropriate belief status to adopt, and, in the latter case, to find consensual alternatives to support the affected beliefs. The proposed conflict resolution methodologies, although distinct from the typical negotiation based conflict resolution protocols that perform a distributed search trough the space of possible solutions [7], have identical goals: both try to reach acceptable agreements to all the parties involved by either: (i) choosing between conflicting views, by comparing the reasons behind each stance and choosing the strongest view; or (ii) building a new consensual view, by searching for fully acceptable alternative foundations for the disputed information.

Our agents were remotely inspired by the ARCHON architecture [14], and are structured in two main layers: the intelligent system layer and the cooperation layer. The intelligent system layer is a reason maintenance system which includes an Assumption based Truth Maintenance Systems (ATMS) [1], and contains the individual agent's domain knowledge. The cooperation layer includes a self model, where the agent's intelligent system is described (knowledge and beliefs that the agent has or is expected to have), and an acquaintances model. The implemented distributed reason maintenance methodologies are described in [8], and include communication, representation, evaluation and accommodation policies for local and communicated beliefs.

After this introductory section we describe the distributed and autonomous belief revision methodologies developed. We start by presenting the methodology for solving Context Independent Conflicts and then follow with the methodology for solving the Context Dependent Conflicts. In the forth section we discuss our work by comparing it with related work, and, in the last section we draw our conclusions.

Before continuing with our paper we wish to present some definitions:

1. Beliefs are first order logic sentences which where inferred or introduced directly trough perception, communication or assumption;
2. A generic belief ϕ of agent Ag (also referred as agent Ag view of ϕ) is represented by a tuple (also called an ATMS node) $< \phi_{Ag}, \mathcal{E}(\phi_{Ag}), \mathcal{F}(\phi_{Ag}), Ag >$, where ϕ_{Ag} identifies the belief, $\mathcal{E}(\phi_{Ag})$ specifies the belief's endorsement (observed, assumed, communicated or inferred), $\mathcal{F}(\phi_{Ag})$ contains the sets of foundations that support the belief, and Ag identifies the belief's source agent. The belief status is established according to the $\mathcal{F}(\phi_{Ag})$: (i) *Believed*, if $\mathcal{F}(\phi_{Ag}) \neq \{\emptyset\}$; (ii) *Disbelieved*, if $\mathcal{F}(\phi_{Ag}) = \{\emptyset\}$;
3. $\mathcal{F}(\phi_{Ag})$ represents the sets of foundations of ϕ_{Ag}. These sets are made of self-supported beliefs (in our case, observed or assumed beliefs) - for example, $< \alpha_{Ag}, ass, \{\{\alpha_{Ag}\}\}, Ag >$ and $< \beta_{Ag}, obs, \{\{\beta_{Ag}\}\}, Ag >$. The foundations of inferred or communicated beliefs are made of the observed and assumed beliefs that supported the inference or communication of those beliefs - for example, $< \phi_{Ag}, inf, \{\{\alpha_{Ag}, \beta_{Ag}, \ldots\}\}, Ag >$ or $< \phi_{Ag}, com, \{\{\alpha_{Ag}, \delta_{Ag}, \ldots\}\}, Ag >$;

4. A distributed belief ϕ is represented by as many beliefs as there are agents' views - for example, $< \phi_i, \mathcal{E}(\phi_i), \mathcal{F}(\phi_i), Ag_i >, \ldots, < \phi_n, \mathcal{E}(\phi_n), \mathcal{F}(\phi_n), Ag_n >$.

2 Context Independent Conflicts

The Context Independent Conflicts result from the assignment, by different agents, of contradictory belief statuses to the same belief. Every agent maintains not only individual beliefs but also distributed beliefs. While the responsibility for individual beliefs relies on each source agent, the responsibility for distributed beliefs is shared by the agents involved. Three elementary processing criteria for distributed beliefs were implemented to accommodate the multiple views: (i) *The CONsensus (CON) criterion* - The distributed belief will be *Believed*, if all the perspectives of the different agents involved are believed, or *Disbelieved*, otherwise; (ii) *The MAJority (MAJ) criterion* - The distributed belief will be *Believed*, as long as the majority of the perspectives of the different agents involved is believed, and *Disbelieved*, otherwise; (iii) *The At Least One (ALO) criterion* - The distributed belief will be *Believed* as long as at least one of the perspectives of the different agents involved is believed, and *Disbelieved*, otherwise. If the beliefs that constitute a distributed belief are consensual the distributed belief the CON criterion is applied, otherwise the distributed belief results from the resolution of the detected Context Independent Conflict. The methodology we are proposing for solving the Context Independent Conflicts is organized in two steps: first, it establishes the desired outcome of the social conflict episode, and then, it applies the processing criterion that solves the episode accordingly. During the first stage, the conflict episode is analyzed to establish the most reasonable outcome. The necessary knowledge is extracted from data dependent features like the agents' reliability [3] and the endorsements of the beliefs [5], allowing the selection, at runtime, of the most adequate processing criterion. In particular, in our work the reliability of an agent is domain dependent - an agent can be more reliable in some domains than in others. The dynamic selection of the processing criterion is based on assessment of the credibility values associated with each belief status. Two credibility assessment procedures were designed:

> *The Foundations Origin based Procedure (FOR)* - where the credibility of the conflicting perspectives is determined based on the strength of the foundations endorsements (observed foundations are stronger than assumed foundations) and on the domain reliability of the source agents; and
> *The Reliability based Procedure (REL)* - where the credibility of the conflicting views is based on the reliability of the foundations source agents.

The methodology for solving Context Independent Conflicts starts by applying the FOR procedure. If the FOR procedure is able to determine the most credible belief status, then the selected processing criterion is applied and the episode is solved. However, if the result of the application of the FOR procedure is a draw between the conflicting perspectives, the Context Independent conflict solving methodology proceeds with the application of the REL procedure. If the REL procedure is able to establish the most credible belief status, then the selected processing criterion is applied and the episode is solved.

The sequential application of these procedures is ordered by the amount of knowledge used to establish the resulting belief status. It starts with the FOR procedure which calculates the credibility of the conflicting perspectives based on the strength of the endorsements and on the domain reliability of the sources (agents) of the foundations. It follows with the REL procedure which computes the credibility of the conflicting beliefs solely based on the domain reliability of the sources of the foundations. We will now describe in detail the procedures for solving Context Independent Conflicts mentioned above.

2.1 The Foundations ORigin based Procedure (FOR)

The individual agents' perspectives are based on their respective sets of foundations (kept by the ATMS modules) which resulted from some process of observation, assumption or external communication. Since communicated beliefs also resulted from some process of observation, assumption or communication of other agents, foundations are, ultimately, composed of just observed or assumed beliefs. Galliers [5] refers to a belief's process of origin as the belief's endorsement.

When a Context Independent Conflict involving a distributed belief ϕ occurs the FOR procedure is invoked and the credibility values for each one of the conflicting views regarding ϕ is computed according to the following formulas:

$$\mathcal{C}(\phi, Bel) = \sum_{Ag=i}^{k} \mathcal{C}(\phi_{Ag}, Bel)/N$$
$$\mathcal{C}(\phi, Dis) = \sum_{Ag=i}^{k} \mathcal{C}(\phi_{Ag}, Dis)/N$$

where N is the number of agents involved in the conflict. The values of the credibility attached to each agent view are equal to the average of the the credibility values of their respective sets of foundations. The credibility of a generic foundation of ϕ - for example, $< \alpha_{Ag}, \mathcal{E}(\alpha_{Ag}), \mathcal{F}(\alpha_{Ag}), Ag >$ which belongs to domain D - depends on the reliability of the foundation's source agent in the specified domain ($\mathcal{R}(D, Ag)$), on the foundation's endorsement ($\mathcal{E}(\alpha_{Ag})$) and on its support set ($\mathcal{F}(\alpha_{Ag})$):

$$\mathcal{C}(\alpha_{Ag}, Bel) = 1 \times \mathcal{R}(D, Ag), \text{ if } \mathcal{F}(\alpha_{Ag}) \neq \{\emptyset\} \text{ and } \mathcal{E}(\alpha_{Ag}) = obs;$$
$$\mathcal{C}(\alpha_{Ag}, Bel) = 1/2 \times \mathcal{R}(D, Ag), \text{ if } \mathcal{F}(\alpha_{Ag}) \neq \{\emptyset\} \text{ and } \mathcal{E}(\alpha_{Ag}) = ass;$$
$$\mathcal{C}(\alpha_{Ag}, Bel) = 0 \text{ if } \mathcal{F}(\alpha_{Ag}) = \{\emptyset\};$$
$$\mathcal{C}(\alpha_{Ag}, Dis) = 1 \times \mathcal{R}(D, Ag), \text{ if } \mathcal{F}(\alpha_{Ag}) = \{\emptyset\} \text{ and } \mathcal{E}(\alpha_{Ag}) = obs;$$
$$\mathcal{C}(\alpha_{Ag}, Dis) = 1/2 \times \mathcal{R}(D, Ag), \text{ if } \mathcal{F}(\alpha_{Ag}) = \{\emptyset\} \text{ and } \mathcal{E}(\alpha_{Ag}) = ass;$$
$$\mathcal{C}(\alpha_{Ag}, Dis) = 0 \text{ if } \mathcal{F}(\alpha_{Ag}) \neq \{\emptyset\};$$

Within this procedure, the credibility granted, *a priori*, to observed foundations and assumed foundations was, respectively, 1 and 1/2. The credibility of each foundation is also affected by the reliability of the origin agent (source agent) for the domain under consideration. As a result, each perspective is assigned a credibility value equal to the average of the credibility values of its support foundations. The credibility of any perspective has, then, a value between 0 and 1. A credibility value of 1 means that perspective is 100% credible (it solely depends on observed foundations generated by agents fully reliable on the data domain), whereas a credibility value of 0 means that no credibility whatsoever is associated with the perspective. Semantically, the 1 and 1/2 values granted to observed and assumed foundations have the following meaning: evidences corroborated by perception are 100% credible, whereas assumptions have a 50% chance of being confirmed or contradicted through perception.

The FOR procedure calculates the credibility attached to the each one of the conflicting belief status (*Believed* and *Disbelieved*), and chooses the multiple perspective processing criterion whose outcome results in most credible belief status. If the most credible belief status is: (i) *Disbelieved*, then the CON criterion is applied to the episode of the conflict; (ii) *Believed*, then, if the majority of the perspectives are in favor of believing in the belief the MAJ criterion is applied; else the ALO criterion is applied to the episode of the conflict. The agents' reliability on the domain under consideration is affected by the outcome of the Context Independent Conflict episodes processed so far. An episode winning agent increases its reliability in the specified domain, while an episode loosing agent decreases its reliability in the specified domain. At launch time, the agents are assigned a reliability value of 1 to every knowledge domain, which, during runtime, may vary between 0 and 1. If the agent's view won the conflict episode of domain D then

> $\mathcal{R}(D, Ag) = \mathcal{R}(D, Ag) \times (1 + N_w/N) \times f_{norm}$, where N_w, N and f_{norm} represent, respectively, the number of agents who won the episode, the total number of agents involved, and a normalization factor needed to keep the resulting values within the interval [0,1];

If agent Ag view lost a Context Independent Conflict episode of domain D then

> $\mathcal{R}(D, Ag) = \mathcal{R}(D, Ag) \times (1 - N_l/N) \times f_{norm}$, where N_l, N and f_{norm} represent, respectively, the number of agents who lost the episode, the total number of agents involved in the conflict episode and a normalization factor needed to keep the resulting values within the interval [0,1].

A domain reliability value of 1 means that the agent has been fully reliable, and a value near 0 means that the agent has been less than reliable.

Example Suppose a multi-agent system composed of three agents, $Agent_i$, $Agent_j$ and $Agent_k$, with the following knowledge bases:

> $Agent_i$ has observed $\alpha(a)$, assumed $\beta(a)$ and has two knowledge production rules[1], $r_{i1} : \alpha(X) \wedge \beta(X) \rightarrow \phi(X)$ and $r_{i2} : \delta(X) \wedge \phi(X) \rightarrow \psi(X)$:
> $< \alpha_i(a), obs, \{\{\alpha_i(a)\}\}, Agent_i >$;
> $< \beta_i(a), ass, \{\{\beta_i(a)\}\}, Agent_i >$;

---

[1] rules are also represented as beliefs which may be activated or inhibited (believed/disbelieved).

$< r_{i1}, obs, \{\{r_{i1}\}\}, Agent_i >$;
$< r_{i2}, obs, \{\{r_{i2}\}\}, Agent_i >$;
$Agent_j$ has observed $\alpha(a)$, assumed $\delta(a)$ and has one knowledge production rule, $r_{j1} : \alpha(X) \wedge \delta(X) \rightarrow \phi(X)$:
$< \alpha_j(a), obs, \{\{\alpha_j(a)\}\}, Agent_j >$;
$< \delta_j(a), ass, \{\{\delta_j(a)\}\}, Agent_j >$;
$< r_{j1}, obs, \{\{r_{j1}\}\}, Agent_j >$;
$Agent_k$ has observed $\phi(a)$:
$< \phi_k(a), obs, \{\{\phi_k(a)\}\}, Agent_k >$.

$Agent_i$ is interested in receiving information regarding $\alpha(X)$, $\beta(X)$, $\delta(X)$, $\phi(X)$ and $\psi(X)$, $Agent_j$ is interested in $\alpha(X)$, $\delta(X)$ and $\phi(X)$, and $Agent_k$ is interested in any information on $\phi(X)$. The agents after sharing their results (according to their expressed interests) end up with the following new beliefs:

$Agent_i$ has received $\alpha_j(a)$, $\delta_j(a)$, r_{j1}, $\phi_k(a)$, and $\phi_j(a)$, and has inferred $\phi_i(a)$ and $\psi_i(a)$:
$< \alpha_j(a), com, \{\{\alpha_j(a)\}\}, Agent_j >$;
$< \delta_j(a), com, \{\{\delta_j(a)\}\}, Agent_j >$;
$< r_{j1}, com, \{\{r_{j1}\}\}, Agent_j >$;
$< \phi_k(a), com, \{\{\phi_k(a)\}\}, Agent_k >$;
$< \phi_i(a), inf, \{\{\alpha_i(a), \alpha_j(a), \beta_j(a), r_{i1}\}\}, Agent_i >$;
$< \phi_j(a), com, \{\{\alpha_i(a), \alpha_j(a), \delta_j(a), r_{j1}\}\}, Agent_j >$;
$< \psi_i(a), inf, \{\{\alpha_i(a), \alpha_j(a), \beta_i(a), \delta_j(a), r_{i1}, r_{j1}, \phi_k(a), r_{i2}\}\}, Agent_i >$;
$Agent_j$ has received $\alpha_i(a)$, $\beta_i(a)$, $\phi_k(a)$, r_{i1}, and $\phi_i(a)$, and has inferred $\phi_j(a)$:
$< \alpha_i(a), com, \{\{\alpha_i(a)\}\}, Agent_i >$;
$< \beta_i(a), com, \{\{\beta_i(a)\}\}, Agent_i >$;
$< r_{i1}, com, \{\{r_{i1}\}\}, Agent_i >$;
$< \phi_k(a), com, \{\{\phi_k(a)\}\}, Agent_k >$;
$< \phi_j(a), inf, \{\{\alpha_i(a), \alpha_j(a), \delta_j(a), r_{j1}\}\}, Agent_j >$;
$< \phi_i(a), com, \{\{\alpha_i(a), \alpha_j(a), \beta_i(a), r_{i1}\}\}, Agent_i >$;
$Agent_k$ has received $\alpha_i(a)$, $\alpha_j(a)$, $\beta_i(a)$, $\delta_j(a)$, r_{i1}, r_{j1}, $\phi_i(a)$, $\phi_j(a)$ r_{i1}, $\phi_i(a)$, and has inferred $\phi_j(a)$:
$< \alpha_i(a), com, \{\{\alpha_i(a)\}\}, Agent_i >$;
$< \alpha_j(a), com, \{\{\alpha_j(a)\}\}, Agent_j >$;
$< \beta_i(a), com, \{\{\beta_i(a)\}\}, Agent_i >$;
$< \delta_j(a), com, \{\{\delta_j(a)\}\}, Agent_j >$;
$< r_{i1}, com, \{\{r_{i1}\}\}, Agent_i >$;
$< r_{j1}, com, \{\{r_{j1}\}\}, Agent_j >$;
$< \phi_i(a), com, \{\{\alpha_i(a), \alpha_j(a), \beta_i(a), r_{i1}\}\}, Agent_i >$;
$< \phi_j(a), com, \{\{\alpha_i(a), \alpha_j(a), \delta_j(a), r_{j1}\}\}, Agent_j >$;

At some point, $Agent_k$ realizes, via observation, that $\phi(a)$ is no longer believed, i.e, $< \phi_k(a), obs, \{\emptyset\}, Agent_k >$. A first episode of a Context Independent Conflict regarding $\phi(a)$ is declared: $\phi_i(a)$ and $\phi_j(a)$ are believed while $\phi_k(a)$ is disbelieved. In the case of our conflict, the credibility values assigned to the believed status is obtained through the following expressions:

$$\mathcal{C}(\phi_i(a), Bel) = (\mathcal{C}(\alpha_i(a), Bel) + \mathcal{C}(\alpha_j(a), Bel) + \mathcal{C}(\beta_i(a), Bel) + \mathcal{C}(r_{i1}, Bel))/4$$
$$\mathcal{C}(\phi_j(a), Bel) = (\mathcal{C}(\alpha_i(a), Bel) + \mathcal{C}(\alpha_j(a), Bel) + \mathcal{C}(\delta_j(a), Bel) + \mathcal{C}(r_{j1}, Bel))/4$$
$$\mathcal{C}(\phi_k(a), Bel) = \mathcal{C}(\phi_k(a), Bel)$$

As the default reliability values assigned to every agent domain was 1, the credibility values associated to the conflicting perspectives are:

$$\begin{array}{lll} \mathcal{C}(\phi_i(a), Bel) = (1 + 1 + 1/2 + 1)/4 & \quad and \quad & \mathcal{C}(\phi_i(a), Dis) = 0 \\ \mathcal{C}(\phi_j(a), Bel) = (1 + 1 + 1/2 + 1)/4 & \quad and \quad & \mathcal{C}(\phi_j(a), Dis) = 0 \\ \mathcal{C}(\phi_k(a), Bel) = 0 & \quad and \quad & \mathcal{C}(\phi_k(a), Dis) = 1 \end{array}$$

resulting in

$$\mathcal{C}(\phi(a), Bel) = 7/12 \qquad and \qquad \mathcal{C}(\phi(a), Dis) = 1/3.$$

Since $\mathcal{C}(\phi(a), Bel) > \mathcal{C}(\phi(a), Dis)$, the multi-agent system decides to believe in $\phi(a)$. The first episode of the conflict regarding $\phi(a)$ was successfully solved through the application of the FOR procedure, and the processing criterion that must be applied to generate the adequate social outcome of this conflict episode is the MAJ criterion. Finally, the conflict domain reliability values of the agents involved in the conflict episode are updated accordingly. So $Agent_i$, $Agent_j$ and $Agent_k$ updated credibility values for the domain under consideration (D) are:

$\mathcal{R}(D, Agent_i) = 1 \times (1 + 2/3)/(1 + 2/3)$, i.e., $\mathcal{R}(D, Agent_i) = 1$,
$\mathcal{R}(D, Agent_j) = 1 \times (1 + 2/3)/(1 + 2/3)$, i.e., $\mathcal{R}(D, Agent_j) = 1$, and
$\mathcal{R}(D, Agent_k) = 1 \times (1 - 1/3)/(1 + 2/3)$, i.e., $\mathcal{R}(D, Agent_k) = 6/15$.

2.2 The RELiability based Procedure (REL)

The REL procedure assigns each conflicting view a credibility value equal to the average of the reliability values of the source agents of its foundations. The credibility associated with the different perspectives that contribute to each belief status are added and the REL procedure chooses processing criterion whose outcome results in adopting the most credible belief status. If the most credible belief status is: (i) *Disbelieved*, then the CON criterion is applied to the episode of the conflict; (ii) *Believed*, then, if the majority of the perspectives are in favor of believing in the belief the MAJ criterion is applied, else the ALO criterion is applied to the episode of the conflict. The reliability of the agents in the domain conflict is also affected by the outcome of the conflict episodes solved so far. Episode winning agents increase their reliability in the conflict domain, while episode loosing agents decrease their reliability in the conflict domain (see previous sub-section).

3 Context Dependent Conflicts

The detection of contradictory distinct beliefs within the system triggers the reason maintenance mechanism and, as a result, previously believed conclusions may become disbelieved. Although this activity is essential to the maintenance of well founded beliefs, the system should make an effort to try to believe in its conclusions as much as possible. To solve this type of conflicts the system needs to know how to provide alternative support to the invalidated conclusions. This search for "next best" solutions is a relaxation mechanism called *Preference ORder based procedure (POR)*.

Each knowledge domain has lists of ordered candidate attribute values for some domain concepts. These lists contain the sets of possible instances ordered by preference (the first element is the best candidate, the second element is the second best, and so forth) for the attribute values of the specified concepts. When a Context Dependent Conflict occurs it means that the originally built support (containing the current instances of the attributes) became invalid. The search for alternative support based the "next best" strategy will be triggered only when the system concludes its inability to fulfill its goals with its current knowledge. The search is achieved by looking for the "next best" instances for the foundations of the invalidated belief, which, if founded, will provide a new valid support for the affected concept. The preference order values are affected by the proposing agents' reliability in the domain under consideration. In the case of the foundations maintained by a single agent, the process is just a next best candidate retrieval operation. In the case of foundations maintained by groups of agents, a consensual candidate has to be found. If the gathered proposals are: (i)*Identical*, then the new foundation has been found; (ii)*Different*, then the agents that proposed higher preference order candidates generate new next best proposals, until they run out of candidates or a consensual candidate is found. The alternative foundations found through this procedure are then assumed by the system. The credibility measure of the resulting new foundations is a function of the lowest preference order candidate used and of the involved agents' reliability. The agents' reliability values of the domain under consideration are not affected by the Context Dependent Conflicts resolution activity.

4 Discussion

The autonomous belief revision methodologies presented above have been inspired by the work of several authors. The idea of using endorsements for determining preferred revisions was first proposed by Cohen [4] and, later, by Galliers in [5]. However, while Galliers proposes a several types of endorsements according to the process of origin of the foundational beliefs (communicated, given, assumed, etc.), we claim that they can be reduced to two processes: perception or assumption. We base this decision on the fact that all agent's beliefs result from some process of observation, assumption or external communication, and since communicated perspectives also resulted from some process of observation, assumption or communication of other agents, ultimately, the foundations set of any belief is solely made of observed and assumed beliefs. Similarly, Gaspar [6] determines the preferred belief revisions according to a belief change function which is based on a set of basic principles and heuristic criteria (sincerity of the sending agent, confidence and credulity of the receiving agent) that allow the agents to establish preferences between contradictory beliefs. Beliefs are associated to information topics and different credibility values are assigned to the agents in these topics. More recently, Dragoni and Giorgini [3] proposed the use of a belief function formalism to establish the credibility of the beliefs involved in conflicts according to the reliability of the source of belief and to the involved beliefs credibility. Each agent is assigned a global reliability value which is updated by the belief function formalism after each conflict. Our agents have domain dependent reliability values because, like Gaspar, we believe that agents tend to be more reliable in some domains than others and that the assignment of a global agent reliability (like Dragoni and Giorgini do) would mask these different expertise levels. We also guarantee, unlike Dragoni and Giorgini, that communicated foundations are only affected by the reliability of the foundation source agent, and not by the reliability of agent that communicated the foundation (which may be different).

Other authors, like [12], [11] or [13]), use argumentation based negotiation as a methodology for solving inter-agent conflicting objectives. However, there are some important differences - our goal is to solve information conflicts and not to solve conflicts among agents' objectives. In fact, our methodologies can be regarded as argumentation based. In particular, the procedures used to solve Context Independent Conflicts can be regarded as the final stage of an implicit argumentation based negotiation process. Furthermore, in our approach beliefs are, by default, exchanged together with their sets of foundations. These sets of foundations constitute the full arguments' list in favor of the communicated belief, which include the rules used to infer the communicated belief - as suggested by [11]. So, when a Context Independent conflict occurs, we are one step ahead since we have already exchanged the existing arguments for each conflicting perspective - the FOR and REL procedures are the last stage needed for establishing which is the most credible set of foundations or, in other words, which is the strongest argument.

5 Conclusion

The belief revision methods we have designed for the resolution of the identified types of negative conflicts are based on data dependent features or explicit knowledge rather than belief semantics. The data features used to solve conflicting beliefs have been previously proposed by other authors (source agent reliability by Dragoni and Giorgini [3]), (endorsements by Galliers [5]), and (specification of preferences between beliefs by Gaspar [6]). However, we believe that we are combining them in a novel and more efficient manner.

These methodologies have been implemented and are being tested in a project location application described in [9]. They attempt to solve the detected conflicting beliefs but cannot, beforehand, guarantee whether their effort will be successful or not.

References

1. de Kleer, J. 1986. An Assumption-Based Truth Maintenance System. Artificial Intelligence 28(2):127-224.
2. Dragoni, A. F. Giorgini, P. and Puliti, P. 1994. Distributed Belief Revision versus Distributed Truth Maintenance. In Proceedings of the Sixth IEEE Conference on Tools with Artificial Intelligence, IEEE Computer Press.
3. Dragoni, A. F. and Giorgini, P. 1997. Belief Revision through the Belief Function Formalism in a Multi-Agent Environment. Third International Workshop on Agent Theories, Architectures and Languages, LNAI N. 1193, Springer-Verlag, 1997.
4. Cohen, P. R. 1985. Heuristic Reasoning about Uncertainty: an Artificial Intelligence Approach. Pitman, Boston, 1985.
5. Galliers, R. 1992. Autonomous belief revision and communication, Gärdenfors, P. ed. Belief Revision. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, N. 29.
6. Gaspar, G. 1991. Communication and Belief Changes in a Society of Agents: Towards a Formal Model of an Autonomous Agent, Demazeau, Y. and Müller, J. P. eds. Decentralized A. I., Vol. 2, North-Holland Elsevier Science Publishers.
7. Jennings, N. R., Parsons, S., Noriega, P. and Sierra, C. 1998. On Argumentation-Based Negotiation. In Proceedings of the International Workshop on Multi-Agent Systems, Boston, USA.
8. Malheiro, B. and Oliveira, E. 1996. Consistency and Context Management in a Multi-Agent Belief Revision Tesbed. Wooldridge, M., Müller, J. P. and Tambe, M. eds. Agent Theories, Architectures and Languages, Springer-Verlag, Lecture Notes in AI, Vol. 1037.
9. Malheiro, B., Oliveira, E. 1997. Environmental Decision Support: a Multi-Agent Approach, The First International Conference on Autonomous Agents (Agents'97), Marina del Rey, California, USA.
10. Oliveira, E., Mouta, F. and Rocha, A. P. 1993. Negotiation and Conflict Resolution within a Community of Cooperative Agents. In Proceedings of the International Symposium on autonomous Decentralized Systems, Kawasaki, Japan.
11. Parsons, S., Sierra, C. and Jennings, N. R. 1998. Agents that reason and negotiate by arguing, Journal of Logic and computation, 8(3):261-292.
12. Sycara, K. 1990. Persuasive Argumentation in negotiation, Theory and Decision. 28:203-242.
13. Vreeswijk, A. W. 1997. Abstract argumentation systems. Artificial Intelligence 90(1-2).
14. Wittig, T. ed. 1992. ARCHON: An Architecture for Multi-Agent Systems, Ellis Horwood.