Invited Paper:

Collaboration through Deliberative Dialogues

Sonia V. Rueda svr@cs.uns.edu.ar
Alejandro J. García ajg@cs.uns.edu.ar
Guillermo R. Simari grs@cs.uns.edu.ar

Departamento de Ciencias e Ingeniería de la Computación
Universidad Nacional del Sur
8000 Bahía Blanca, Argentina

Abstract

In a multi-agent system, a group of individuals interact in a social context in order to boost their capabilities and enhance global performance. Each individual's action repertoire may be reduced, but it's social capability allows it to interact with other agents and obtain collaboration.
This work offers an alternative for knowledge representation in a system of collaborative BDI agents and presents an interaction protocol based on dialogues.
The capacity to interact affects the behavioral model of a BDI agent that must consider the possibility of offering and soliciting collaboration. Thus, we propose an algorithm that models the behavior of a collaborative BDI agent.

Keywords: Multi-Agent Systems, Collaborative Agents, Interaction Protocols, Defeasible Logic Programming

1 Introduction

The agent metaphor is suited for modeling dynamic applications subject to unforeseeable changes in which a flexible response is expected. In this work, an agent is conceived as an entity situated in a dynamic environment and characterized by basic properties such as purpose, reactiveness, proactiveness and autonomy.
A rational agent is additionally endowed with some capacity that allows it to select adequate actions regarding its purpose. Autonomy guarantees that the agent can make decisions with a certain level of independence with respect to its environment. Rationality allows it to make *good* decisions regarding the objectives put forth by the purpose while taking the state of the world into consideration.
The BDI model tries to capture the essence of practical reasoning to determine an agent's decision mechanism, and allows the selection at each moment of the action to be performed according to the current goal. An agent's *beliefs* are made up of all its knowledge about itself and its environment. Its desires and intentions refer to the state that it wishes to achieve.

An isolated BDI agent's behavior is determined by its individual motivations, its beliefs about the world, and its own abilities. Autonomy implies control over its behavior and over its internal state. Rationality provides flexibility; the agent reacts to changes that are produced in the world, but it can also take initiative, *i.e.*, it has a proactive attitude. On the other hand, the possibility of achieving its goals is restricted to its own cognitive capability.

In a multi-agent system, a group of individuals interact in a social context in order to boost their individual capabilities and enhance global performance. Each individual's action repertoire may be reduced, but it's social capability allows it to interact with other individuals and obtain collaboration. Thus, each individual tries to build plans using its own action repertoire, but when this repertoire is not sufficient, it requests collaboration and another member of the group can elaborate plans to service the request. Participation in a social context generates opportunities that do not come up in individual work, but it also creates restrictions that need to be considered.

This work offers an alternative for knowledge representation in a system of collaborative BDI agents and presents an interaction protocol based on dialogues. Interaction is initiated through a global requirement put forth by an agent that has formulated a plan for its committed intention but that does not posses the full knowledge needed to execute it. Among the individuals that are willing and able to collaborate, the agent that solicited the dialog chooses one in particular and initiates a deliberative dialog.

The capacity to interact affects the behavior of a BDI agent that must consider the possibility of offering and soliciting collaboration. Thus, an

extension to the algorithm is proposed that models the behavior of a BDI agent introduced in [4]. The extension takes into consideration the roles that an agent may assume in a dialog as the initiator of the interaction or as a collaborator.

2 Practical Reasoning

An agent is a computational entity capable of perceiving, reacting, and acting in an environment. A rational agent adds to its perceptive, reactive, and effectoric capabilities a certain level of *practical reasoning* that allows it to select the most adequate action given its knowledge of its goals and the state of the environment in which it is situated [8]. That is, the cognitive capability of the agent is what will allow it to act adequately in each particular context.
Practical reasoning involves two fundamental processes: *deliberation* and *means-ends reasoning*. Deliberation allows to decide *what* goals will be pursued and requires considering a set of alternative options, selecting some of them, and committing to realizing them. Means-ends reasoning determines *how* the committed goals will be reached, that is it allows to build a plan that will realize the agent's intentions.
An agent modifies its intentions when it decides that it won't be able to reach them, or the reasons for which they were chosen over other alternatives disappear. In this sense, an agent's intentions are linked to its *beliefs* about the world that surrounds it.
The belief, desire, and intentions (BDI) model, based on practical reasoning, provides the essential elements for representing the mental attitude of a rational agent that acts in a dynamic environment subject to sudden and frequent changes. Beliefs make up the agent's knowledge about itself and the environment that surrounds it. Its desires refer to the state that it wishes to reach and represent its motivations. Intentions constitute committed goals that condition and control future activities. The most important characteristics of intentions are:

- An agent frequently selects a course of action geared towards satisfying an intention, which fails. Generally, an agent does not abandon the intention immediately, and it tries again in a different way. An agent is expected to *persist* in its intention and perform a reasonable number of tries.
- Once an agent adopts an intention, this intention will act as a restriction over future practical reasoning. The actions chosen after the moment in which an intention is defined must be consistent with it.
- An agent's intentions are strongly linked to its *beliefs* about the future. A behavior will not be rational if it adopts an intention while believing that it is not possible to achieve.

All of these qualities intervene in the process of practical reasoning, but they frequently collide with each other. An intention may be abandoned after many failed tries or because an agent adopts a new one that is inconsistent with the former. Clearly, having failed a certain number of times the agent will start doubting the feasibility of bringing about its committed goals, and this goal will lose its ties with the agent's beliefs. However, it is difficult to find a balance between constant deliberation while reconsidering intentions, and indefinitely persisting in them.
Formally, the *state* of a BDI agent at a given moment can be represented by a triple $< B, D, I >$ such that $B \subseteq Bel$, $D \subseteq Des$, and $I \subseteq Int$, where Bel is the set of all possible beliefs, Des is the set of all possible desires, and Int is the set of all possible intentions [4].
The agent's belief revision function can be defined as a mapping

$$brf : \wp(Bel) \times P \rightarrow \wp(Bel)$$

that determines a new set of beliefs from the set of perceptions P and the current beliefs. The options generating function maps sets of beliefs and intentions into sets of desires:

$$options : \wp(Bel) \times \wp(Int) \rightarrow \wp(Des)$$

Options must be in some sense *opportunistic* in that they must recognize characteristics of the environment and the agent's knowledge that allow to anticipate success in the planning process if they are chosen as intentions.
The deliberation process can be represented through a function of the form:

$$deliber : \wp(Bel) \times \wp(Des) \times \wp(Int) \rightarrow \wp(Int)$$

that modifies intentions given previous beliefs, desires, and intentions. This function allows abandoning intentions that are not reachable (at least at a reasonable cost), retain those that can still produce a global benefit, and adopt new ones that will allow the agent to come closer to the existing ones or explore new opportunities. That is, new intentions are obtained from previous ones or from new options.
The *plan* function returns a sequence of actions given a set of beliefs and the committed intention:

$$plan : \wp(Int) \times \wp(Bel) \rightarrow \Pi$$

The algorithm proposed in [5] to specify the behavior of a BDI agent is:

$B \leftarrow B_0$
$I \leftarrow I_0$
while true **do**
  get next percept p
  $B \leftarrow brf(B, p)$
  $D \leftarrow options(B, I)$
  $I \leftarrow deliber(B, D, I)$
  $\Pi \leftarrow plan(B, I)$
  $execute(\Pi)$
end while

The sets B_0 and I_0 correspond to the initial beliefs and intentions. In applications in which the environment suffers frequent changes, the algorithm can be progressively refined to incorporate reactivity and intention reconsideration. In this case, the *execute* operation does not cover the entire plan; after executing each individual action there is a perception and an analysis of the necessity of adjusting it. Intention reconsideration allows to take advantage of changes and make the best of new opportunities [5].

2.1 A Behavioral Model for a BDI Agent

An alternative to the view proposed in the previous section is a model in which an agent's desires are not determined from the current beliefs and previous intentions, but are perceived as needs in the environment. That is, the $options(B, I)$ function is eliminated, and in its place desires are updated by a desire revision function that considers perception P:

$$drf : \wp(Des) \times P \rightarrow \wp(Des)$$

The deliberation operation is also modified because instead of considering the intentions previously reached, it selects a new intention considering exclusively the assigned goals and the current beliefs. The algorithm is now:

$B \leftarrow B_0$
$D \leftarrow D_0$
while true **do**
  get next percept p
  $B \leftarrow brf(B, p)$
  $D \leftarrow drf(D, p)$
  $I \leftarrow deliber(B, D)$
  $\Pi \leftarrow plan(B, I)$
  $execute(\Pi, B)$
end while

Once again, the $execute(P, B)$ operation can include an intention reconsideration mechanism such that after the execution of each action the perceived changes in the world are considered, and a decision is made as to if it is possible to continue as planned, if the plan needs modifications, or a different intention is selected.
The BDI model is abstract, it does not specify how an agent's beliefs are represented, nor does it propose a concrete mechanism for deliberating and a specific planning strategy. The selection of intentions can be carried out in different ways, one of which is that all desires have the same importance, while another is to associate a *priority* to each desire.
In this work, the set of goals that an agent perceives and incorporates to its desires corresponds to the actions that it is capable of executing. This does not guarantee that an executable plan can be built because the set of beliefs can be insufficient, in which case we decide that the plan is incomplete. When the execution of a plan finishes successfully, the intention has been reached and becomes a belief.
In the following we briefly describe the formalism adopted for representing a BDI agent's knowledge and for supporting planning. Later on we propose a variation of the algorithm modeling the behavior of a BDI agent adding collaboration.

2.2 Representation of a BDI Agent's Knowledge

The BDI model provides the essential elements for representing the mental attitude of a rational agent that acts in a dynamic environment, subject to sudden and frequent changes. Beliefs make up the agent's knowledge about the world, and its desires and intentions refer to the state that it wishes to reach, as well as representing its motivations and commitments. The effectoric capability is constituted by the set of actions that it is capable of executing.

Definition 2.1 [Argumentative BDI agent] *An argumentatve-BDI agent* **a***will be denoted with the tuple* $\mathbf{a} = \langle \mathcal{B}, \mathcal{D}, \iota, \Gamma \rangle$*, where* $\mathcal{B}$ *represents the agent's beliefs,* $\mathcal{D}$ *its desires,* ι *the agent intention and* Γ *a set of actions that the agent is able to execute.*

The agent's beliefs $\mathcal{B}$ will be represented by a restricted *Defeasible Logic Program* (Φ, Δ), where Φ is a consistent set of facts, and Δ a set of defeasible rules. The interested reader is referred to [2] for full details about DeLP.
Besides its beliefs, desires, and intentions, an agent will have a set of actions Γ that it may use to change its world. The formal definitions that were introduced in [1] are recalled below.
An agent's desires are represented as a set of literals. An intention is a specific literal, selected from the set of desires as a committed goal.

Definition 2.2 [Action] *An action A is an ordered triple* $\langle \mathsf{P}, \mathsf{X}, \mathsf{C} \rangle$*, where* P *is a set of literals representing preconditions for A,* X *is a consistent set of literals representing consequences of executing A, and* C *is a set of constraints of the form not L, where L is a literal. We will denote actions as follows:*

$$\{X_1, \ldots, X_n\} \xleftarrow{A} \{P_1, \ldots, P_m\}, not\ \{C_1, \ldots, C_k\}$$

Notice that the notation not $\{C_1, \ldots, C_k\}$ *represents* $\{not\ C_1, \ldots, not\ C_k\}$*.*

Definition 2.3 [Applicable Action] *Let* $\mathcal{B} = (\Phi, \Delta)$ *be an agent's beliefs and* Γ *the set of actions available to that agent. An action A in* Γ*, defined as before, is applicable if every precondition* P_i *in* P *has a warrant built from* (Φ, Δ) *and every constraint* C_i *in* C *fails to be warranted.*

Definition 2.4 [Action Effect] *Let* $\mathcal{B} = (\Phi, \Delta)$ *be an agent's beliefs and* Γ *the set of actions available to that agent. Let A be an applicable action in* Γ *defined as before. The effect of executing A is the revision of* Φ *by* X*, i.e.,* $\Phi^{*\mathsf{X}} = \Phi^{*\{X_1,\ldots,X_n\}}$*. Revision will consist of removing any literal in* Φ *that is complementary of any literal in* X *and then adding* X *to the resulting set. Formally:*

$$\Phi^{*\mathsf{X}} = \Phi^{*\{X_1,\ldots,X_n\}} = (\Phi - \overline{\mathsf{X}}) \cup \mathsf{X}$$

where $\overline{\mathsf{X}}$ *is the set of complements of members of* X*.*

In [1] we have shown that the interaction between actions and the defeasible argumentation formalism is twofold. On one hand, as stated by Definition 2.3, defeasible argumentation is used for testing preconditions and constraints through the warrant notion. On the other hand, actions may be used by agents in order to change the world (actually the set Φ) and then have a warrant for a literal L that has no warrant from the current knowledge base (Φ, Δ).

2.3 Planning with Defeasible Argumentation

When an agent **a** adopts an intention ι, and ι is not warranted by (Φ, Δ) agent **a** looks for an action A in Γ that modifies the beliefs in such a way that an argument without defeaters that supports ι can be obtained. If A cannot be executed because its preconditions are not warranted, **a** elaborates a sequence of actions that allows it to establish them. The complete sequence, including A, will make up a plan for ι.
Each action in the plan can modify the beliefs by adding literals that allow to build new arguments. When the execution of intention ι is complete, it is added to the set $\mathcal{B}$ of beliefs.
The selection of actions is not a trivial task, and the classic issues that arise in planning are reflected in the argumentative scheme. In a sequence of actions [A_1,A_2], A_2 can be applicable according to the initial state Φ, but not after A_1 is executed. Therefore, the consequences of A_1 can modify Φ in such a way that A_2's preconditions are not warranted, or a constraint is warranted.
The execution of A_1 can add literals that allow to build new defeaters for the preconditions of A_2 and eliminate literals that allow the construction of warrants for the preconditions of A_2. The criterion proposed in [1] for selecting actions is to minimize the unwanted changes in the preconditions and constraints.
We have stated that in this work an agent only adopts intentions for which it can build at least one plan. That is, there is at least one action among whose post-conditions is the literal that corresponds to the intention. However, it is possible that the preconditions of this action are not satisfied considering the agent's set of beliefs, and that the agent cannot elaborate a plan that will allow it to generate them. In this case, the plan will be incomplete. If the agent is isolated, an incomplete plan cannot be executed.
When an agent is part of a system, and its capabilities are insufficient for completing a plan, it can request collaboration. However, its actions can interfere with the plans of other members. If its attitude is cooperative, it will provide collaboration when it receives requests, and will request authorization to make changes that may affect others.

3 Collaboration among BDI Agents

An individual agent's plan is built considering only its beliefs, intentions, and abilities. When an agent is part of a system, it can interact with other members of the system in which it participates to make use of its beliefs and abilities.
Interaction allows more possibilities of reaching the committed goals, but it also requires considering the impact that the social behavior has over knowledge representation, the deliberative mechanism, the planning process, and each individual agent's behavior model.
In this work, interaction is structured in *dialogues* between agents that seek to collaborate in order to augment their individual capabilities. The behavior of the interlocutors is still rational, that is, as before each individual must balance the effort that it allocates to elaborate an effective plan, in

relation to the time that it employs in executing it. Furthermore, in a dialogue, each agent allocates resources for requesting and providing collaboration; once again it is important to maintain a balance between interaction and the execution of plans.

3.1 Dialogues in a Cooperative Context

A dialogue is a sequence of locutions that are exchanged between two interlocutors that share the same objective and try to follow a turn taking model. The objective is to reach an *agreement*; each participant has associated a set of propositions that are adequate to it and consistent with its own goals. The dialogue can be thought of as a *dialogical game* in which the players apply different strategies in trying to balance the intention of reaching their goals and the commitment to collaborate.
When the dialogue evolves, the set of agreements of each interlocutor is modified to incorporate or remove propositions. There are different ways of dialoguing, and dialogues can be grouped into types according to different criteria, but all of them require a certain level of commitment and argumentation.
One criterion for classifying dialogues is to identify the goals and initial situations that regulate them in order to characterize each particular type and distinguish it from the rest, following the proposals of [].
The following list presents some types of very general dialogues, without intending to offer a complete taxonomy but only to identify the models that are significant for supporting communication and coordination [9].

- *Persuasive dialogue*
- *Negotiation*
- *Inquiry*
- *Information Seeking*
- *Deliberation*

All of the models mentioned have a characteristic in common, and that is that the interlocutors act in a context of trust and cooperation. Even when they may disagree, the relation is not antagonistic, there are no attacks and the realization of the goals are not evaluated in terms of individual defeats or victories.
The initial situation in a *persuasive dialogue* is the conflict between different points of view, and the objective is that one of the parts change its stance through the dialectic process. In the simplest case, one of the parts has a positive take on a given proposition, while the other has a negative view of the same proposition.

The goal of a *negotiation dialogue* is to reach a shared agreement. Each participant tries to make this agreement be as close as possible to its own interests. Even though the spirit of collaboration exists, there is also an implicit purpose of maintaining individual motivations. The dialectic process can be thought of as a search process oriented towards finding a situation that is adequate for both parts.

Inquiry is a type of dialogue oriented towards proving propositions to answer a question or solve a problem that is recognized and shared. The dialectical process is very similar to that of a persuasive dialogue, and in both cases if the dialogue is successful, the participants adopt the same conclusion. The difference lies in the initial state, because an inquiry does not start with a conflict between two points of view, but with an open problem.

Information seeking begins when an agent needs information about a particular subject. Contrary to an inquiry, the relation between the participants is asymmetric because one of them manifests its ignorance and expects to become richer with the knowledge of the other, without requiring a proof that backs the response.

A *deliberative* dialogue begins, like an inquiry, with an open problem. However, in this case it has a practical nature, and the goal is a decision on how to act. In the deliberation process, each participant can try to influence the final solution, sharing in this case characteristics of persuasive dialogues.

It is clear that the five types of dialogues described are not separated by absolute limits, and on occasions a dialogical game can have elements of two or more types. Frequently, a dialogical game that initially corresponds to a certain type is transformed into a different one as a response to the evolution of the context in which it develops. The transition is not always clear, and often a dialogue becomes embedded in another one that precedes it and later follows it.

The importance of identifying what type a particular dialogue fits is related to the need for recognizing the level of commitment of each part with respect to the goal, and identifying what situations determine that the process finishes successfully or not. In this work, the interaction between agents corresponds to a deliberative dialogue in which the interlocutors exchange proposals and counterproposals trying to make effective collaboration.

3.2 Knowledge Representation in a Social Context

A BDI agent is characterized by exhibiting flexible behavior and being capable of taking actions autonomously according to its mental attitudes. In a multi-agent system, the representation of cognitive capability covers two important aspects. First, the relation between knowledge and action, that is, how an agent's knowledge about its environment and its own goals determines its behavior. Second, how knowledge influences the interaction with other agents.
In a social context, each individual acts and interacts considering what it knows about the world, including what it knows about what the rest knows [6]. In applications in which agents have conflicting goals, this aspect is fundamental. If a seller agent offers a product that has a defect, for instance a house with problems in the rain gutters, it may conceal this situation and act in a certain manner as long as the possible buyer is not aware of the situation. However, if the seller *knows* that the possible buyer *knows* about the state of the gutters, its attitude may be different. The knowledge about the world that is maintained by the agents in a system can then be organized in different ways. The weakest way is that beliefs are distributed between the members of the system, each of which maintains its own set of beliefs. In this model, each agent may perceive different aspects of the real world and may also have a different view of a particular aspect. Because of this consideration, it is probable that the beliefs of a member of the group are inconsistent with those of the rest.
The strictest alternative is to maintain the set of beliefs shared by all the members of the system. Any modification of the knowledge will affect everyone. However, joining all of the agents' beliefs invalidates the advantages of a system divided into modules. Thus, the entire model could be reduced to an only agent.
An intermediate position, according to what is proposed in [3], is to maintain part of the *distributed shared beliefs* between pairs of agents. When an agent needs collaboration, it interacts with another member of the group by means of a dialogue. If the request can be serviced, the shared beliefs between the two are modified. Furthermore, there exists a set of facts known to all members of the group that make up the global shared beliefs, and whose modification involves general consensus.
In the proposed model, individual beliefs of each agent are composed by their specific beliefs, the shared beliefs with each of the other members of the group, and the global shared beliefs. Every set of individual beliefs is consistent.

Definition 3.1 [**Multiagent System**] *A system composed of collaborative BDI agents is described as* $MS = \langle A, K, KP \rangle$*, where* A *is a set of agents,* K *is the set of global shared beliefs, and* $KP = \{(\mathbf{a}_i, \mathbf{a}_j, K_{ij})\}$ *represents the family of sets of shared beliefs between every pair* $\mathbf{a}_i, \mathbf{a}_j$ *of agents in* A*.*

In the BDI model, the cognitive capability of an agent includes its beliefs, but also its desires and intentions. In this work, the presence of other individuals in the environment affects the set of desires that now will not only include the goals perceived as necessary, but also the *collaboration requests*. Thus, the set $\mathcal{D}$ of each individual will be composed by the goals that correspond to goals perceived in the world and to collaboration requests of other members of the system. That is, $\mathcal{D} = \mathcal{M} \cup \mathcal{R}$, where $\mathcal{M}$ are the goals and $\mathcal{R}$ are the received collaboration requests.
The desire perception function $drf(D, p)$ is replaced by a pair of functions $rrf(R, p)$ and $grf(M, p)$ that repeatedly perceive collaboration requests and goals, respectively. During the deliberation process the collaboration requests and the goals perceived in the world will compete with each other.

3.3 A Dialogical Protocol for Collaborative Agents

The behavior of a collaborative BDI agent is more complex than that of an isolated agent. In this work, every member of a system acts motivated by its goals, but also considers the collaboration requests that it receives. An individual requests collaboration when its beliefs are insufficient for executing the plans that allow it to reach its goals. Thus, the existence of other members is going to allow it to request and obtain collaboration, but it will also require a certain capability for interaction.
When an agent $\mathbf{a}_1$ requires a literal p, and cannot find a warrant for it nor elaborate a sequence of actions that allows it to add it, it performs a *call for dialogue* requesting collaboration in relation to p.
Each of the other members of the system perceives the global request for collaboration, and two situations can arise: an agent has available an action among whose post-conditions is p, or not. In the first case, it incorporates the request into its set of desires, while in the latter it rejects the request because none of the actions that make up its effectoric capabilities allows it to service the request, and therefore cannot elaborate a plan for p.
If a request is selected as an intention, the agent notifies its availability to collaborate and waits

for a specific request. For each request that is not adopted as an intention, the agents must notify that they are not available. In the call, an agent initiates the interaction through a global request for the rest of the members of the system, as shown in Figure 1. Figure 2 shows the sequence diagram in UML for the call for dialogue.

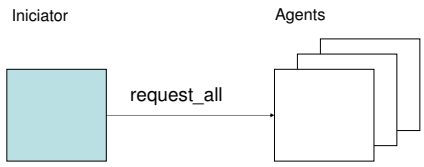

Figure 1: *Roles Diagram for call for dialogue.*

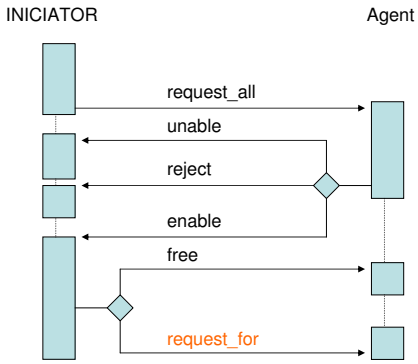

Figure 2: *Sequence diagram for call dialogue.*

As shown in the diagram, every agent that receives the global request responds according to one of the following situations:

a) It does not adopt p as a committed intention, and rejects the request.

b) None of the actions that make up its effectoric capabilities allows it to service the request, and therefore cannot elaborate a plan for p; this is informed.

c) Adopts p as a committed intention, and indicates its availability to collaborate.

Interaction with those agents that reject the request or inform that they are not available to service it ends directly. Agent $\mathbf{a}_1$ selects an agent among those that have manifested their availability, say $\mathbf{a}_2$, with which it initiates a deliberative dialogue through a specific request for p.
Agent $\mathbf{a}_2$ may have a complete plan for p, or may request collaboration with respect to another literal q. In this case, it performs a counterproposal to agent $\mathbf{a}_1$ with respect to q. Both agents then exchange proposals and counterproposals, trying to make effective collaboration, until the dialogue ends. Figure 3 shows the sequence diagram in UML for the deliberative dialogue.

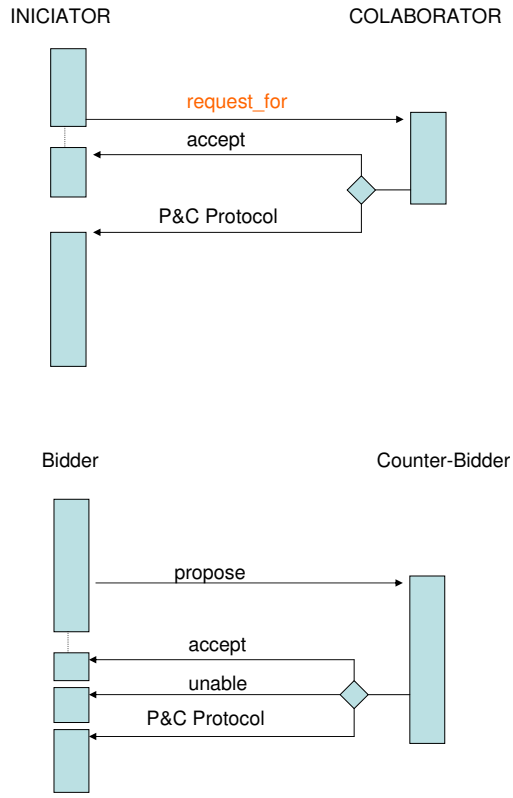

Figure 3: *Deliberative Dialogue* and *P&C Protocol.*

When an agent receives a counterproposal it responds:

a) Indicating that it cannot elaborate a plan that allows it to satisfy the request

b) Accepting the request

c) Performing a counterproposal

The behavior of a BDI agent must consider the two roles with which it can interact. That is, as initiator of a dialogue or as collaborator. As we will see in the following, the algorithm is extended to reflect that an agent can accept to participate in a dialogue when it receives a collaboration request, and also perform itself a global request to initiate a dialogue.

3.4 Behavioral Model for a Collaborative BDI Agent

The possibility of obtaining collaboration will allow every agent to elaborate plans that will require the application of knowledge from the other

member of the system. In the following algorithm, when an agent adopts an intention that corresponds to an individual goal and elaborates a plan that is not complete, it initiates a call for dialogue by means of the *request* operation, which may end successfully or not.
In the case in which the committed intention corresponds to a request for collaboration, the agent adopts the opposite role in the dialogue. That is, it manifests its disposition to collaborate by means of the *dialogue* operation, and waits for the start of a deliberative dialogue.

$B \leftarrow B_0$
$D \leftarrow D_0$
while true **do**
  get next percept p
  $B \leftarrow brf(B, p)$ {belief revision}
  $M \leftarrow grf(G, p)$ {goal revision}
  $R \leftarrow rrf(R, p)$ {collaboration request revision}
  $I \leftarrow deliber(B, M, R)$
  $P \leftarrow plan(\Pi, I)$
  if $i \in R$ **then**
    $OK \leftarrow dialogue()$
  else
    if $incomplete(\Pi)$ **then**
      $OK \leftarrow request()$
    end if
  end if
  $check(B, \Pi, OK)$
  if OK **then**
    $execute(\Pi, B)$
  end if
end while

The intention adopted in the deliberative process may correspond to a goal or to a collaboration request; in both cases the agent elaborates a plan. If the intention corresponds to a collaboration request, it notifies its disposition to collaborate and if it receives a specific request a dialogue will be generated between the agent that initiated the interaction and the collaborator.
In the case in which the intention corresponds to a perceived goal, the agent elaborates a plan that may not be complete, that is, its beliefs are not sufficient to satisfy the preconditions of the plan. It then initiates the interaction with other members of the group by means of a global collaboration request.
The *check* function requires perceiving the world again and deciding if the plan is complete, considering the current state. In the case in which this is so, the plan is executed. Once again, the *execute* operation can be extended to allow that after the execution of each action the state of the world to be perceived, and the necessary adjustments made.
As before, when an agent completes the execution of a plan elaborated for reaching a committed intention, its beliefs are modified. If the intention corresponds to a desire perceived in the world, the set of global shared beliefs is updated. In the case in which the intention corresponds to a request for collaboration, the knowledge shared between the interlocutors is modified.
In both cases, the execution of the plan can modify other beliefs, apart from those linked to the committed intention. If the modification affects the set of beliefs shared with another member of the system, both must agree on the change. If the belief is part of the global knowledge, all agents must authorize the modification.

4 Conclusions and Future Work

In a multi-agent system, a group of autonomous entities can collaborate in solving a problem that escapes each individual's capabilities. In a cooperative environment, each BDI agent's reasoning involves deliberation and planning, but a certain *social commitment* is added. In this work, social commitment allows each member of the system to provide and request collaboration by participating in deliberative dialogues.
Individual autonomy and social commitment are two properties that oppose each other. The former allows each agent to *act* guided by its own goals, independently of the environment to which it belongs. Social commitment restricts the way in which agents *interact* and influences individual behavior, limiting the level of autonomy.
As in a human organization, the specification of norms, roles, and relations allows to balance autonomy with the social commitment that is imposed when being part of a group. Individual behavior and interaction are determined by a relation between the roles that the interlocutors have during the dialogue. That is, the structure in which the roles are related within the organization restricts the possible locutions in certain moments of the dialogue. Autonomy is still an important quality, but constrained by the behavior expected of the role in the context that determines the structure of relations.
Our future work is then oriented towards the introduction of norms, roles, and relations in the agent system representation, and the modification of each agent's individual behavior and the interaction protocol. The evolution of each particular dialogue will be determined by the interaction protocol and the relation between the specific

roles that the interlocutors occupy. The initiation of a dialogue depends on the individual behavior of the agent that requests collaboration, and that of the one that provides it.

5 Acknowledgements

This research was partially supported by the *Department of Computer Science and Engineering, Secretaría General de Ciencia y Tecnología, Universidad Nacional del Sur*, *Consejo Nacional de Investigaciones Científica y Técnicas* (CONICET), and *Agencia Nacional de Promoción Científica y Tecnológica* (ANPCyT), PICT 15043, PAV076, Argentina.

References

[1] Simari, G., García, A., Capobianco, M. *Actions, Planning and Defeasible Reasoning*, Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR 2004), pp. 377–384, Whistler, Canada, June 6-8, 2004.

[2] García, A., Simari, G. *Defeasible Logic Programming: An Argumentative Approach*, Theory Practice of Logic Programming, vol 4, **1**, pp. 95–138, 2004.

[3] Rueda S., García A., Simari, G. *Alternativas para la representacion del conocimiento compartido entre agentes colaborativos*, II Workshop en Agentes y Sistemas Inteligentes (WASI), VIII CACIC, Argentina, pp. 1508–1919, 2004.

[4] Wooldridge, M., *Intelligent Agents*, in Multiagent Systems A Modern Approach to Distributed Artificial Intelligence. Editor G. Weiss. The MIT Press, 2000.

[5] Wooldridge, M., *Reasoning about Rational Agents*, The MIT Press, 2000.

[6] Fagin, R., Halpern J., Moses, Y., Vardi, M. *Reasoning About Knowledge*, The MIT Press, Cambridge, MA 1995.

[7] Rao, A. S., and Georgeff M. P., *BDI Agents: From Theory to Practice*, In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95) pp.312-319 San Francisco, USA, 1995.

[8] Bratman, M.E., Israel D.J., Pollack, M.E., *Plans and resource-bounded practical reasoning*, Computational Intelligence 4: 349,355 1988.

[9] Walton, D. N., Krabbe E.C., *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, State University of New York Press, Albany, NY, 1995.