A Framework for Deliberation and Negotiation among BDI Agents

Sonia V. Rueda *
and
María Vanina Martinez
Departamento de Ciencias e Ingeniería de la Computación
Universidad Nacional del Sur
8000 Bahía Blanca, Argentina
{svr,mvm}@cs.uns.edu.ar

ABSTRACT

This work presents an interaction model among collaborative agents based on dialogues. For the specification we use a formalism, called *Dialogue Games*, which allows to describe the nature of the utterances available in each dialogue. We propose an interaction language which favors deliberation and argumentation-based negotiation among BDI agents. For each locution's specification we establish a set of preconditions, its meaning, the expected responses, and the modifications that it could produce. The interaction protocol and the interaction language as a whole can be used for specifying meaningful interaction between dialogical partners by following the rules of an individual dialogue.

Keywords: Collaboration, Deliberation, Negotiation, Argumentation, Interaction Language, Interaction Protocol, Dialogue Games.

1 INTRODUCTION

Many applications which follow the agents model demand high levels of interaction. The members of the system interact in order to coordinate their actions and to distribute resources and tasks trying to reach a state acceptable to all. The interaction models vary depending on the system's characteristics. If all the members are part of an organization, the relationship among them can be a collaborative one. The group can also be composed of homogeneous or heterogeneous agents. In the former case, all the members share the same view of the world and they have identical capacities. In a heterogeneous group, agents will in general have distinct views of the world and different abilities.

In this work we adopt the BDI model for representing the mental attitudes of each member of the group. The individual knowledge of each agent is conformed by its specific knowledge and the knowledge shared with other members in the group; each agent will reason using the facts that are available to it. As it is proposed in [10], the shared knowledge is distributed among *pairs* of agents; therefore, although each agent's view of the world is consistent, different members of the system can have different views. The group is heterogeneous, and each agent's goals are tied to the abilities of the group. Despite their differences, all of the members in the organization are autonomous and rational entities with a collaborative attitude. Inside a collaborative environment each member tries to build plans using its own set of actions, but when these are not enough it asks for collaboration; in these cases another member of the group may elaborate plans in order to satisfy the request. We can see the interaction between two agents as a *dialogical game*, in the sense that the two participants perform *moves* by taking turns following their own goals and strategies.

Dialogues always start as a *deliberative process* in our approach, that is, starting with an unsolved problem, a set of agents try to find a plan which will include actions from both of them in order to solve the problem. Collaboration yields better opportunities for reaching individual goals, but it may also be the cause of conflicts. If there exist conflicts that prevent agents from collaborating successfully, the dialogue becomes a *negotiation process*. By means of a negotiation process, the agents seek an agreement in which each participant will try to make this agreement be as close as possible to its own interests. In this way, there not only exists a collaborative spirit but also an implicit purpose for maintaining each agent's individual motivations.

Argument-based negotiation is a suitable alternative for modelling situations in which agents have limited information and bounded capacities [9]. During the process, the participants acquire information, but it is also possible for them to reach a point in which they must revise their plans and even modify their

*Partially supported by *Secretaría de Ciencia y Tecnología, Universidad Nacional del Sur*, CONICET, and by the *Agencia Nacional de Promoción Científica y Tecnológica* (PICT 2002 Nro 13096, PICT 2003 15043 and PAV 076)

preferences in order to be able to reach an agreement. In our approach each agent elaborates arguments as part of its own planning processes [3, 13] and to justify its proposals, counter-proposals, and rejections during the negotiation process.

In Section 2 we present an interaction model among collaborative agents based on dialogues. For the specification we use a formalism, called *Dialogue Games*, which allows to describe the nature of the utterances available in each dialogue. In Section 3 we propose an interaction language which favors deliberation and argumentation-based negotiation among BDI agents. For each locution's specification we establish a set of preconditions, its meaning, the expected responses, and the modifications that it could produce. The interaction protocol and the interaction language as a whole can be used for specifying meaningful interaction between dialogical partners by following the rules of an individual dialogue. Finally in Section 4 we discuss conclusions and outline future work.

2 INTERACTION PROTOCOL

A multi-agent system consists of a group of agents that interact with each other. This interaction is generally regarded as the foundation for cooperative and competitive behavior in autonomous agents. The term *interaction protocol* is used in reference to a set of rules that guide interactions. In this work an interaction protocol is implemented by means of *dialogues* between agents. The structure of a dialogical system can be thought of as a dialogical game, in the sense that two participants perform moves by taking turns following their own goals and strategies.

It is possible to see the dialectical process as a search process oriented towards finding a situation suitable for both parties. The process is not linear, and therefore the space is not reduced until the solution is reached because it can move and even incorporate new points. In most cases, each agent knows only part of the search space and, within it, there is only a portion which satisfies its expectations. Each agent has a specific set of points within the space of agreements that are acceptable to it. The search is successful when an *agreement space* is reached, that is, there exists a nonempty intersection among the individual spaces. The process ends when the search ends, regardless of its success or failure.

In a simple interaction protocol the agents elaborate, accept, or reject proposals. This approach is not adequate when interaction is viewed as a search process. In this case, the receiver of a proposal must be able not only to accept or reject the proposal, but also to guide the process with its answer. Agents perform proposals and counter-proposals elaborating *arguments* which intend to persuade other agents [4].

Our model proposes that agents' collaboration requests be restricted only to requests for other agents' beliefs, and the possibility of requesting the execution of an action is not considered (at least not directly through a primitive). When an agent requests collaboration for a certain literal p it is indicating that it needs to include p in its own knowledge, and another agent must add it on its behalf. When the dialogue ends successfully, the shared knowledge is modified with the incorporation of new beliefs. The shared knowledge is distributed among pairs of agents, and therefore the modification initially affects only two agents. However, the interaction process may have involved various members of the group. Thus, the language must allow agents to manifest not only their expectation that another member of the group remain committed to the interaction, but also their intention to free such agent from that commitment.

Dialogues and Conversations

Interaction in this model is implemented by means of dialogues between two agents. A *dialogue* is an exchange of speech acts between two speech partners inturn-taking sequence aimed at a collective goal [14]. Each participant may also have its own goals in the dialogue, which should balance with the commitment imposed by the shared goal. Each type of dialogue requires certain level of commitment and argumentation, and each participant has associated a set of propositions which composes its set of agreements. As the dialogue evolves, each interlocutor's set of agreements is modified in order to add or remove propositions.

Starting from a single dialogue, the interaction process derives into a set of dialogues among other pairs of agents, which are conducted in parallel; we use the term *conversation* to represent these sets of dialogues. This last consideration extends the proposal in [11], where agents establish a dialogue with another agent and the interaction process consists of a single dialogue. A conversation evolves through a series of four steps:

- Building the working team
- Social deliberation
- Negotiation
- Dialogue suspension or termination
- Conversations termination

Building the Working team: When an agent elaborates a plan which it cannot execute completely, it performs a global collaboration request. The members of the system which are able to offer collaboration will show their willingness, and those which are not willing will not answer. The agent which performed the global request will build a working team which will include itself and those agents which answered the global request.

Social Deliberation: This step starts when the agent which initiated the conversation performs a specific collaboration request to one of the working team's members, that is, a dialogue is created between them. During the deliberative process, the interlocutors exchange proposals and counter-proposals [**?**]. The dialogue between the two agents continues until

the requested literal is obtained, or one of them decides that it is not able to collaborate. This last situation could arise due to lack of knowledge, or because of conflicts among goals.

Negotiation: If the deliberation step was not able to satisfy the performed request, the agent initiates a negotiation process with one or more individuals of the working team. While negotiating, the agent will insist on its request, compelling the interlocutor to modify its plans, or even demand it to change its goals.

Note that the existence of conflicts between another pair of agents can also prevent the obtention of the requested literal. In these cases, our negotiation model proposes that the agent which made the global request is informed that a conflict with a third agent is preventing the successful collaboration; consequently, the agent must treat directly with the third party in order to resolve the conflict and reach an agreement.

Suspension or Termination of Dialogues: During the conversation, since many agents may be involved in order to satisfy the request, many dialogues may be generated. Some dialogues will conclude rapidly and others may be suspended until certain conflicts are resolved. In either case, each dialogue will eventually end, either successfully or not. The dialogues are suspended due to conflicts among third parties which must be resolved with the agent interested in the request. During the suspension of a dialogue both agents remain committed to eventually obtain the literal requested or dissolve the commitment if the request is no longer reachable.

A dialogue is said to finish successfully if the agents are able to cooperate with each other and the requested literal is added to the knowledge shared between them. A successful end for a dialogue may be reach by means of only deliberation, or deliberation followed by negotiation. In either case, if after these steps the request is not satisfied, the dialogue ends unsuccessfully.

Note that even when a dialogue ends, the corresponding conversation continues until the agent which initiated it decides explicitly to end it. After finishing a dialogue, these agents may start another dialogue in the frame of the same conversation.

Termination of Conversations: A conversation among agents in the system ends when all the dialogues that it derived finish. A conversation is said to have finished successfully if the particular performed request is satisfied; this does not guarantee that all the dialogues derived in the conversation have ended successfully. When all the dialogues end the agent which performed the global request explicitly ends the conversation, freeing the members of the working team which will now be able to participate in other conversations.

Dialogue Games

The literature offers different formalisms for specifying interaction protocols in multi-agent systems. No matter which alternative is chosen, it must at least include the following elements:

- Types of participants.
- Interaction states.
- Events which trigger state changes.
- Valid actions given the participant and the state.

Dialogue games are a particular alternative suitable for expressing argumentation. This formalism can be used to specify meaningful interaction between dialogical partners by following the rules of an individual dialogue. The interaction between two or more *players* is defined by means of a formal dialogue game, in which locutions are considered to be moves. The rules specify which locutions are permitted under what circumstances, and which responses are possible. There are different types of dialogue game rules, as proposed in [7]:

- Commencement and termination: define the circumstances under which the dialogue begins and ends.
- Locutions: specify the nature of the utterances permitted in the dialogue.
- Combination: define the dialogical contexts under which a particular locution is allowed.
- Commitment: define the circumstances under which a participant expresses dialogical commitment to a proposition.

This formalism provides a unifying framework that represents different types of dialogues, each of which has a simple semantics. In an interaction protocol based on dialogue games, it is possible to identify appropriate locutions and to define constraints on their utterances. Basically, the goal of a dialogue model is to structure a dialogue into *dialogue acts*, and to find the relationships between dialogue acts -utterances- that explain its coherence. A dialogue is said to be *coherent* if the sequence of utterances performed by the participants builds a dialogue context which represents the set of statements and commitments that were made by them [2].

The proposed interaction protocol requires an interaction language which allows agents to manifest their attitude regarding each particular dialogue, in such a way that the transition from deliberation to negotiation is evident. In the following section we present an interaction language designed for supporting *conversations* among argumentative BDI agents. The language offers a set of primitives suitable for expressing proposals and counter-proposals, providing arguments and expressing the interest level that agents assign to each collaboration request.

3 INTERACTION LANGUAGE

The role of a language for BDI agents is fundamental in allowing them to express their mental attitudes. The purpose of an interaction language, as proposed in [5], is the communication of messages which represent the agent's knowledge and that are interpreted in a well defined manner. Furthermore, these messages

cause certain actions on behalf of both the sender and the receiver.

We propose an interaction language based on *Speech Act Theory*. Speech Act Theory-based languages capture the essential characteristics of human communication, producing a model suitable for the development of artificial agents. One of the essential ideas in the theory is that communication is a special type of *action*, called *speech act* [12].

Speech Act Theory classifies the messages in the following way:

- Affirmative: affirms or establishes the truth of a certain statement.
- Directive: command, request, or suggestion.
- Commitment: proposition by which the performer commits itself with a certain course of action.
- Declarative: triggers, by itself, an action.
- Expressive: expresses feeling or attitudes.

An utterance is a single meaningful unit of communication [2], and it is composed of a *semantic content*, i.e. the information conveyed in it, and a *communicative purpose*.

An interaction language for argumentation-based negotiation among BDI agents is proposed in [10]. In the following we present an extension for this language and we show a dialogue game based specification for it. Along with the preconditions, meaning, responses, and updates for each locution, we present the type of each one according to the Speech Act Theory classification.

Locution Specification

In our work, dialogues are always circumscribed to pairs of agents. Except when the conversation begins or ends, in each interaction primitive the first two agents that appear are the sender and the receiver of the message, and the dialogue involves only them.

In the following let $\mathbf{a}_1$, $\mathbf{a}_2$, and $\mathbf{a}_3$ be agents in the system, p be a literal, and Q be a set of literals.

Building the Working Team

Locution: *Request_coll*($\mathbf{a}_1$).

- **Type:** Directive.
- **Preconditions:** Agent $\mathbf{a}_1$ must need a literal for which it cannot build a warranty from its individual knowledge, nor elaborate a plan that allows it to add the literal to its individual knowledge.
- **Meaning:** Agent $\mathbf{a}_1$ asks the rest of the members in the system which of them are available for considering a collaboration request. Eventually, agent $\mathbf{a}_1$ will utter a *close_conversation*() indicating to all the agents which respond to its global request that the conversation has ended.
- **Response:** Any agent in the system which is available for considering a collaboration request, may respond with an appropriately instantiated *Available*() locution.
- **Updates:** No effects.

Locution: *Available*($\mathbf{a}_1, \mathbf{a}_2$).

- **Type:** Commitment.
- **Preconditions:** Participant $\mathbf{a}_2$ must have previously uttered a locution *Request_coll*($\mathbf{a}_2$).
- **Meaning:** Agent $\mathbf{a}_1$ lets agent $\mathbf{a}_2$ know that it is available for considering its collaboration request.
- **Response:** None required.
- **Updates:** Agent $\mathbf{a}_2$, the one which uttered the *Request_coll*($\mathbf{a}_2$) locution, must store the names of all agents that respond to the global request with an *Available*($\mathbf{a}_1, \mathbf{a}_2$) locution, constituting the working team for the conversation, that is the set of all the possible agents that can intervene in the initiated conversation. The agent which utters this locution is committed to this conversation and it will not be able to respond to other requests until $\mathbf{a}_2$ explicitly ends the conversation.

Social Deliberation

Locution: *Request*($\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, p, Q, WT$), where WT is the working team for this conversation.

- **Type:** Directive.
- **Preconditions:** Agent $\mathbf{a}_1$ has a goal for which it has built a plan, but there exists a literal p necessary for this plan which $\mathbf{a}_1$ cannot obtain by itself. That is, it must be impossible for $\mathbf{a}_1$ to build a warranty for p, or elaborate a plan to obtain it.
- **Meaning:** If the third argument is not present the meaning of this locution is that agent $\mathbf{a}_1$ needs p and requests $\mathbf{a}_2$'s collaboration in order to obtain it. In the other case -the third argument is instantiated-, agent $\mathbf{a}_1$ needs p and requests $\mathbf{a}_2$'s collaboration in order to obtain it, but this time the request is on behalf of agent $\mathbf{a}_3$. In both cases argument Q conforms the set of beliefs that must be avoided in the plan for obtaining p.

  There are different reasons for which an agent may utter this locution. If $\mathbf{a}_1$ is the agent which initiated the conversation, it will utter this locution when asking for collaboration to an specific agent $\mathbf{a}_2$ regarding to certain literal p. Also, if $\mathbf{a}_1$ has received a request regarding p from $\mathbf{a}_2$, it may utter this locution as a counter-proposal or as new request soliciting help to a third party. That is, if $\mathbf{a}_1$ needs some literal, let's say t, for the plan it founded for p, it may performed a counter-proposal to $\mathbf{a}_2$, or initiate a new dialogue regarding t with another agent in the working team.
- **Response:** Agent $\mathbf{a}_2$ will try to build a plan for p. It will respond:
  - *Accept*(), if it finds a plan for p and it has all that it needs in order to effectively obtain p;
  - *Unable*(), if it cannot find a plan for p, that is its knowledge and capabilities do no suffice;

- *Reject*(), if it finds a plan for p but some of the preconditions needed -or even p- are in conflict with its own goals or beliefs;
- *Indirect_Reject*(), if it finds a plan for p but it is not able to obtain some of the literals needed in that plan by itself. Agent $\mathbf{a}_2$ asks another agent for help, but this agent cannot do so because it has conflicts.
- *Request*(), if it finds a plan for p but, in order to effectively obtain it, it needs another literal which cannot be obtain by itself. The agent must utter a counter-proposal asking $\mathbf{a}_1$ for help but, if $\mathbf{a}_1$ is not able to help it or if $\mathbf{a}_1$ does not respond, it should choice one WT's member in order to continue the negotiation.

- **Updates:** If it is not uttered as a counter-proposal, then a new dialogue has been created between $\mathbf{a}_1$ and $\mathbf{a}_2$.

  The agent which utter this locution must store the plan for whihc the literal p is necessary. If the locution is uttered as a counter-proposal for literal p, it must store the plan founded for p, along with the plan's requirements -preconditions-, while it waits for help.

Locution: *Accept*$(\mathbf{a}_1, \mathbf{a}_2, p)$.

- **Type:** Commitment.
- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Request*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, or *Demand*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution. Agent $\mathbf{a}_1$ must have p in its individual knowledge, have a warranty for it, or be able to build a plan for obtaining it. This locution cannot be uttered after an *Unable*$(\mathbf{a}_1, \mathbf{a}_2, p)$ or after a *Reject*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution. It is necessary for agent $\mathbf{a}_1$ to utter a *Done*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution after having uttered this locution. At this moment, the interaction is not a turn-taking dialogue, agent $\mathbf{a}_1$ will utter two locutions in a row.
- **Meaning:** Agent $\mathbf{a}_1$ informs agent $\mathbf{a}_2$ that it is able to collaborate because it knows p or because it can build a plan for it. Furthermore, there are no conflicts in adding p to the shared knowledge.
- **Response:** None required.
- **Updates:** Agent $\mathbf{a}_1$ is committed to eventually obtain p to satisfy $\mathbf{a}_2$'s request.

Locution: *Reject*$(\mathbf{a}_1, \mathbf{a}_2, p)$.

- **Type:** Affirmative.
- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Request*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, or *Demand*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution. There must exist conflicts between $\mathbf{a}_1$'s plans or goals and $\mathbf{a}_2$'s request about literal p. This locution cannot be uttered after an *Unable*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution.
- **Meaning:** Agent $\mathbf{a}_1$ informs agent $\mathbf{a}_2$ that there exists a conflict between its own plans and p.
- **Response:** If the locution corresponds to a *Request*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution, agent $\mathbf{a}_2$ may respond with *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$. If the locution correspond to an *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution, agent $\mathbf{a}_2$ may respond with *Demand*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$.
- **Updates:** Agent $\mathbf{a}_1$ has rejected $\mathbf{a}_2$'s proposal because there exist conflicts between the plan it found for p and its own goals or beliefs. However, it must store this plan because it is possible for $\mathbf{a}_2$ to insist on p; in this case, $\mathbf{a}_1$ will need the plan in order to revise it against its goals and beliefs.

Locution: *Indirect_Reject*$(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, p, Q)$.

- **Type:** Affirmative.
- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Request*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$, or *Demand*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution. There must exist literals in Q that $\mathbf{a}_1$ needs in order to build a plan or a warranty for p but agent $\mathbf{a}_3$ has conflicts their addition to the shared knowledge or with helping agent $\mathbf{a}_1$ in obtaining them. This locution cannot be uttered after an *Unable*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution.
- **Meaning:** Agent $\mathbf{a}_1$ informs agent $\mathbf{a}_2$ that it needs the beliefs contained in set Q in order to obtain p, but there exists a conflict between these beliefs and agent $\mathbf{a}_3$'s individual knowledge.
- **Response:** If agent $\mathbf{a}_2$ decides to continue the negotiation with agent $\mathbf{a}_3$ by itself then it may respond with a *Still_Int*$(\mathbf{a}_2, \mathbf{a}_1, p)$ locution commiting $\mathbf{a}_1$'s availability for accepting the request.
- **Updates:** Agent $\mathbf{a}_1$ has rejected $\mathbf{a}_2$'s proposal because there exist conflicts between the plan it found for p and the goals or beliefs of another agent. However, it must store this plan because it is possible for $\mathbf{a}_2$ to insist on p; then, $\mathbf{a}_1$ will need the plan for revising it against its goals and beliefs.

Negotiation

Locution: *Insist*$(\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, p, WT)$, where WT is the working team for this conversation.

- **Type:** Directive.
- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Reject*$(\mathbf{a}_2, \mathbf{a}_1, p)$ or *Indirect_Reject*$(\mathbf{a}_3, \mathbf{a}_1, \mathbf{a}_2, p, Q)$ locution to agent $\mathbf{a}_1$ referring to a *Request*$(\mathbf{a}_1, \mathbf{a}_2, p, WT)$ locution for literal p. In this case, agent $\mathbf{a}_1$ should have revised its initial plan for which it needed p. If $\mathbf{a}_1$ could not find a way to avoid needing p, it should insist about obtaining p.
- **Meaning:** If the third argument is not present the meaning of this locution is that agent $\mathbf{a}_1$ asks agent $\mathbf{a}_2$ to revise its plans in order to avoid any

conflict with p. In the other case -the third argument is instantiated-, agent $\mathbf{a}_1$ needs p and insists on $\mathbf{a}_2$'s collaboration in order to obtain it, but this time the request is on behalf of agent $\mathbf{a}_3$.

- **Response:** Agent $\mathbf{a}_2$ will revise its plans in order to obtain p. It will respond:
  - *Accept*(), if it finds another plan for p and it has all that it needs in order to effectively obtain it;
  - *Reject*(), if it cannot find another plan for p or it find one but some of the preconditions needed -or even p- are in conflict with its own goals or beliefs;
  - *Indirect_Reject*(), if it finds another plan for p but it is not able to obtain some of the literals needed in that plan by itself. Agent $\mathbf{a}_2$ asks another agent for help, but this agent cannot do so because it has conflicts.
  - *Request*(), if it finds a plan for p but, in order to effectively obtain it, it needs another literal which cannot be obtain by itself. The agent must ask $\mathbf{a}_1$ for help but, if $\mathbf{a}_1$ is not able to help it, it should choice another WT's member in order to continue the negotiation.
- **Updates:** No effects.

Locution: *Demand*($\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3, p, WT$), where WT is the working team for this conversation.

- **Type:** Directive.
- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Reject*($\mathbf{a}_2, \mathbf{a}_1, p$) or *Indirect_Reject*($\mathbf{a}_3, \mathbf{a}_1, \mathbf{a}_2, p, Q$) locution to agent $\mathbf{a}_1$ referring to an *Insist*($\mathbf{a}_1, \mathbf{a}_2, p, WT$) locution for literal p. In this case, agent $\mathbf{a}_1$ should have revised its goals; if it still needs p after the revision, it should demand collaboration for obtaining p.
- **Meaning:** If the third argument is not present, the meaning of this locution is that agent $\mathbf{a}_1$ order agent $\mathbf{a}_2$ to revise its goals in order to avoid any conflict with p. In the other case -the third argument is instantiated-, agent $\mathbf{a}_1$ needs p and demands $\mathbf{a}_2$'s collaboration in order to obtain it but this time the request is on behalf of agent $\mathbf{a}_3$.
- **Response:** Agent $\mathbf{a}_2$ will revise its goals so it can obtain p. It will respond:
  - *Accept*(), if its goals are no longer in conflict with p and it has all that it needs in order to effectively obtain p;
  - *Reject*(), if after having revised its goals, it has conflicts in obtaining p;
  - *Indirect_Reject*(), if after having revised its goals, it needs help from another agent but this agent cannot do so because it has conflicts.
  - *Request*(), if it finds a plan for p but, in order to effectively obtain it, it needs another literal which cannot be obtain by itself. The agent must ask $\mathbf{a}_1$ for help but, if $\mathbf{a}_1$ is not able to help it, it should choice one WT's member in order to continue the negotiation.
- **Updates:** No effects.

Suspension or Termination of Dialogues

Locution: *Still_Int*($\mathbf{a}_1, \mathbf{a}_2, p$).

- **Type:** Commitment.
- **Preconditions:** There must exists an open dialogue between agents $\mathbf{a}_1$ and $\mathbf{a}_2$. In this dialogue, agent $\mathbf{a}_1$ must have requested collaboration, regarding p, to agent $\mathbf{a}_2$, but it is not possible for agent $\mathbf{a}_2$ to accept it because it needs some elements that are not available to it and another agent is preventing its obtention. That is, agent $\mathbf{a}_2$ must have uttered an *Indirect_Reject*($\mathbf{a}_2, \mathbf{a}_1, \mathbf{a}_3, p, Q$) with respect a request for p performed by $\mathbf{a}_1$.
- **Meaning:** Agent $\mathbf{a}_1$ lets $\mathbf{a}_2$ know that it is still interested in agent $\mathbf{a}_2$'s help in obtaining p. Thus, agent $\mathbf{a}_2$ is committed to reserve its availability for obtaining p and to not change anything in its individual knowledge that could prevent it from obtaining p. Agent $\mathbf{a}_1$ uttered this locution because it will suspend the dialogue with $\mathbf{a}_2$ in order to begin a new dialogue with another agent which is preventing them to reach an agreement. Also, agent $\mathbf{a}_1$ assumes the responsibility of informing whether it will actually make the request or free $\mathbf{a}_2$ from the commitment.
- **Response:** None required.
- **Updates:** The internal state of both $\mathbf{a}_1$ and $\mathbf{a}_2$ changes. Agent $\mathbf{a}_2$ becomes committed to being available for obtaining p, that is, agent $\mathbf{a}_2$ is not allowed to change anything in its individual knowledge that could prevent it from obtaining p in the future.

Locution: *Free*($\mathbf{a}_1, \mathbf{a}_2, p$).

- **Type:** Declarative.
- **Preconditions:** Agent $\mathbf{a}_1$ must have previously uttered a *Still_int*($\mathbf{a}_1, \mathbf{a}_2, p$) locution committing agent $\mathbf{a}_2$ to the request for literal p.
- **Meaning:** Agent $\mathbf{a}_1$ indicates to $\mathbf{a}_2$ that it is free from the request done for p.
- **Response:** None required.
- **Updates:** Agent $\mathbf{a}_2$ is no longer engaged with agent $\mathbf{a}_1$ regarding the collaboration for obtaining literal p, and the dialogue between them has finished.

Locution: *Unable*($\mathbf{a}_1, \mathbf{a}_2, p$).

- **Type:** Affirmative.

- **Preconditions:** Agent $\mathbf{a}_2$ must have previously uttered a *Request*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution. Agent $\mathbf{a}_1$ must not be able to build a plan or obtain a warranty for p. *Unable()* cannot be uttered after an *Accept*$(\mathbf{a}_1, \mathbf{a}_2, p)$ or a *Reject*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution.
- **Meaning:** Agent $\mathbf{a}_1$ informs agent $\mathbf{a}_2$ that it is not capable of obtaining p.
- **Response:** None required.
- **Updates:** With this utterance, the dialogue between $\mathbf{a}_1$ and $\mathbf{a}_2$ has finished unsuccessfully. After uttering this locution, $\mathbf{a}_2$ should not utter an *Insist*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ or *Demand*$(\mathbf{a}_2, \mathbf{a}_1, p, WT)$ locution because $\mathbf{a}_1$ will never be capable of accepting its request.

Locution: *Done*$(\mathbf{a}_1, \mathbf{a}_2, p)$.

- **Type:** Affirmative.
- **Preconditions:** Agent $\mathbf{a}_1$ must have uttered an *Accept*$(\mathbf{a}_1, \mathbf{a}_2, p)$ locution immediately before this utterance.
- **Meaning:** Agent $\mathbf{a}_1$ informs agent $\mathbf{a}_2$ that it has performed all the necessary actions in order to obtain p and that the literal has been added to the shared knowledge.
- **Response:** None required.
- **Updates:** The shared knowledge between agent $\mathbf{a}_1$ and $\mathbf{a}_2$ is updated with the addition of literal p. After uttering this locution, the dialogue between $\mathbf{a}_1$ and $\mathbf{a}_2$ has finished successfully.

Termination of Conversations

Locution: *Close_conversation*$(\mathbf{a}_1, WT)$, WT is the working team for this conversation.

- **Type:** Declarative.
- **Preconditions:** Agent $\mathbf{a}_1$ must have uttered a *Request_coll*$(\mathbf{a}_1)$ locution.
- **Meaning:** Agent $\mathbf{a}_1$ announces to all participants of the working team that the conversation has ended. This locution may be uttered at any time following the *Request_coll*$(\mathbf{a}_1)$.
- **Response:** None.
- **Updates:** All of WT's members become free to participate in other conversations.

Commencement and Termination Rules

A conversation among agents in the system starts when:

- An agent performs a global collaboration request, and
- one or more agents show their willingness to collaborate.

The conversation starts as a deliberative process, but if the agent which initiated it performs specific collaboration requests to every member of the working team and in every case there exist conflicts, the conversation turns into a negotiation.

An explicit description of the individual decision mechanism which leads each member of the system to ask for and offer collaboration is outside the scope of this work. In fact, each agent may use its own mechanism, but once the conversation begins the protocol establishes the expected behavior for the agent which initiated it and for all the agents which showed willingness to collaborate.

A conversation is always ended by the agent which performed the initial request; this agent bases its decision of ending the conversation on one of the following facts:

- The agent obtained what it needed,
- the social deliberation and negotiation steps were unsuccessful, or
- the agent modified its committed goals or plans due to changes in the environment unrelated to the conversation.

A conversation is said to have finished successfully if the particular performed request is satisfied; note that this does not guarantee that all the dialogues derived in the conversation have ended successfully. When the conversation ends, all the members of the working team are freed from the commitments established by it.

Combination and Commitment Rules

The basic rule of combination in each dialogue is the turn-taking scheme. However, there exist situations in which this scheme is broken. For example, when an agent accepts to collaborate with some specific request, it is committed to execute the plan it found for the literal requested, and then to inform that the execution is completed. Another instance occurs when an agent requests, by means of the *Still_int()* locution, that another agent maintain its plan while the conflict with a third party is trying to be resolved. The agent which uttered the *Still_int()* locution is committed to free the other agent in the future, or to retry the request if the conflict is solved. The last form of commitment inside the protocol is carried out by the agents which answered the initial global request. Showing its willingness, each agent commits itself to being able to participate -if requested to do so- in a deliberative dialogue, which may become a negotiation dialogue.

The dialogue can finish successfully or not; it finishes successfully if the agents reach an agreement and the requested literal is added to the knowledge shared between them. Also, if a negotiation step was necessary, there are two possible outcomes: if after subsequent request, insistence, and demand they reach an agreement, then the dialogue ends successfully and the literal is added to the knowledge shared by both agents. On the other hand, it is possible for them not to reach an agreement, a situation which can arise due either to non-resolved conflicts with respect to goals or plans, or to lack of information or capabilities of the agent which is trying to help. In any of these cases the dialogue ends unsuccessfully.

4 CONCLUSION AND FUTURE WORK

Social work allows members in a community to resolve problems that they could not face individually, but must be able to *interact* in order to do so. In this work we proposed an interaction language based on Speech Act theory, where the basic characteristics of human communication are captured and represented in a model suitable for artificial agents. We also present an interaction protocol which gives a set of rules for determining the structures of the dialogues generated during the negotiation process.

Our future work is oriented towards the analysis of alternatives in the specification of conversations. Graphic specification languages are particularly attractive because they allow making the connection among dialogues explicit. UML is currently one of the most powerful graphic design languages for describing software systems, and it provides activity diagrams that can be used for specifying the interaction among the agents in a system. Computations are expressed in terms of states and the progression through them, and the main components used in the description of interaction protocols are action states, activity states, and transitions.

A large body of research proposes an extension of UML, increasing its expressive power in order to support concepts which are specifically oriented towards interaction among agents; other authors consider that it is important to maintain only one general graphic language. Our proposal is to model conversations by means of activity diagrams in UML in order to reflect the structure of the set of dialogues and their interactions.

REFERENCES

[1] Bratman, M., Israel, D., and Pollack, M. Plan and resource-bounded practical reasoning. *Computational Intelligence 4* (1988), 349–355.

[2] Dastani, M. Negotiation protocols and dialogue games. In *Proceedings of the Fifth International Conference on Autonomous Agents* (Montreal, Canada, 2001), J. P. Müller, E. Andre, S. Sen, and C. Frasson, Eds., ACM Press, pp. 180–181.

[3] García, A. J., and Simari, G. R. Defeasible logic programming: An argumentative approach. In *Theory and Practice of Logic Programming* (2004), pp. 95–138.

[4] Jennings, N. R., Parsons, S., Noriega, P., and Sierra, C. On argumentation-based negotiation. In *Proc. of the Int. Workshop on Multi-Agent Systems* (Boston, USA, 1998).

[5] Labrou, Y. *Semantics for an Agent Communication Language*. PhD thesis, Department of Computer Science and Electrical Engineering, Univ. of Maryland and Baltimore County, 1996.

[6] Levin, J. A., and Moore, J. A. Dialogue-games: Metacommuncation structures for natural language interaction. *Cognitive Science 1*, 4 (Oct. 1977), 395–420. (Also in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, eds, pp. 385–397, Morgan Kaufmann, 1988.).

[7] Mcburney, P., Eijk, R. M. V., Parsons, S., and Amgoud, L. A dialogue game protocol for agent purchase negotiations. *Autonomous Agents and Multi-Agent Systems 7*, 3 (2003), 235–273.

[8] Parsons, S., Sierra, C., and Jennings, N. Agents that reason and negotiate by arguing. *Journal of Logic and Computation 8*, 3 (June 1998), 261–292.

[9] Rahwan, I., Ramchurn, S., Jennings, N., McBurney, P., Parsons, S., and Sonenberg, L. Argumentation-based negotiation, 2004.

[10] Rueda, S., García, A., and Simari, G. Alternativas para la representación del conocimiento compartido entre agentes colaborativos. In *Proc. of the VIII CACIC* (2004), pp. 1508–1519.

[11] Rueda, S. V., García, A. J., and Simari, G. R. Argument-based negotiation among bdi agents. *Journal of Computer Science and Technology 2*, 7 (2002), 1–8.

[12] Searle, J. R. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, United Kingdom, 1969.

[13] Simari, and Loui. A mathematical treatment of defeasible reasoning and its implementation. *AIJ: Artificial Intelligence 53* (1992).

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

[15] Wooldridge, M. Intelligent agents. In *Multiagent Systems A Modern Approach to Distributed Artificial Intelligence*, G. Weiss, Ed. MIT Press, 2000, pp. 27–77.