Towards a Formal Framework for the Search of a Consensus Between Autonomous Agents

Leila Amgoud, Sihem Belabbes, and Henri Prade

IRIT - CNRS, 118 route de Narbonne
31062 Toulouse Cedex 9, France
{amgoud, belabbes, prade}@irit.fr

Abstract. This paper aims at proposing a general formal framework for dialogue between autonomous agents which are looking for a common agreement about a collective choice. The proposed setting has three main components: the agents, their reasoning capabilities, and a protocol. The agents are supposed to maintain beliefs about the environment and the other agents, together with their own goals. The beliefs are more or less certain and the goals may not have equal priority. These agents are supposed to be able to make decisions, to revise their beliefs and to support their points of view by arguments. A general protocol is also proposed. It governs the high-level behaviour of interacting agents. Particularly, it specifies the legal moves in the dialogue. Properties of the framework are studied. This setting is illustrated on an example involving three agents discussing the place and date of their next meeting.

Keywords: Argumentation, Negotiation.

1 Introduction

Roughly speaking, negotiation is a process aiming at finding some compromise or consensus between two or several agents about some matters of collective agreement, such as pricing products, allocating resources, or choosing candidates. Negotiation models have been proposed for the design of systems able to bargain in an optimal way with other agents for, e.g., buying or selling products in e-commerce [6].

Different approaches to automated negotiation have been investigated [11], including *game-theoretic* approaches (which usually assume complete information and unlimited computation capabilities), *heuristic-based* approaches which try to cope with these limitations, and *argumentation-based* approaches [3, 1, 10, 8, 7] which emphasize the importance of exchanging information and explanations between negotiating agents in order to mutually influence their behaviors (e.g. an agent may concede a goal having a small priority). Indeed, the two first types of settings do not allow for the addition of information or for exchanging opinions about offers. Integrating argumentation theory in negotiation provides a good means for supplying additional information and also helps agents to convince each other by adequate arguments during a negotiation dialogue.

S. Parsons et al. (Eds.): ArgMAS 2005, LNAI 4049, pp. 264–278, 2006.

In the present work, we consider agents having knowledge about the environment graded in certainty levels and preferences expressed under the form of more or less important goals. Their reasoning model will be based on an argumentative decision framework, as the one proposed in [5] in order to help agents making decisions about what to say during the dialogue, and to support their behavior by founded reasons, namely "safe arguments". We will focus on negotiation dialogues where autonomous agents try to find a joint compromise about a collective choice that will satisfy at least all their most important goals, according to their most certain pieces of knowledge.

The aim of this paper is to propose a general and formal framework for handling such negotiation dialogues. A protocol specifying rules of interaction between agents is proposed. As the agents negotiate about a set of offers in order to choose the best one from their common point of view, it is assumed that the protocol is run, at most, as many times as there are offers. Indeed, each run of the protocol consists of the discussion of an offer by the agents. If that offer is accepted by all the agents, then the negotiation ends successfully. Otherwise, if at least one agent rejects it strongly and doesn't revise its beliefs in the light of new information, the current offer is (at least temporarily) eliminated and a new one is discussed.

We take an example to illustrate our proposed framework. It consists of three human agents trying to set a date and a place for organizing their next meeting. Thus the offers allow for multiple components (date and place). For simplicity reasons, we consider them as combined offers so that if an agent has a reason to refuse an element of a given offer, it refuses the whole offer. One of the agents starts the dialogue by proposing an offer which can be accepted or rejected. The negotiation goes on until a consensus is found, or stops if it is impossible to satisfy all the most important goals of the agents at the same time.

The remainder of this paper is organized as follows: in section 2 we define the mental states of the agents representing their beliefs and goals. In section 3 we present the argumentative decision framework capturing their reasoning capabilities. Section 4 describes a protocol for multi-agent negotiation dialogues. Section 5 illustrates the argued-decision based approach on an example dealing with the choice of a place and a date to organize a meeting. Section 6 concludes the paper and outlines some possible future work.

2 Mental States and Their Dynamics

As said before, it is supposed that the mental states of each agent are represented by bases modeling beliefs and goals graded in terms of certainty and of importance respectively. Following [4, 12], each agent is equipped with $(2n)$ bases, where n is the number of agents taking part to the negotiation.

Let $\mathcal{L}$ be a propositional language and $Wff(\mathcal{L})$ the set of well-formed formulas built from $\mathcal{L}$. Each agent a_i has the following bases:

$\mathcal{K}_i = \{(k_p^i, \rho_p^i), p = 1, s_k\}$ where $k_p^i \in Wff(\mathcal{L})$, is a knowledge base gathering the information the agent has about the environment. The beliefs can be less or more certain. They are associated with certainty levels ρ_p^i.

$\mathcal{G}_i = \{(g_q^i, \lambda_q^i), q = 1, s_g\}$ where $g_q^i \in Wff(\mathcal{L})$, is a base of goals to pursue. These can have different priority degrees, represented by λ_q^i.

$\mathcal{GO}_j^i = \{(go_{r,j}^i, \gamma_{r,j}^i), r = 1, s_{go}(j)\}$, where $j \neq i$, $go_{r,j}^i \in Wff(\mathcal{L})$, are $(n-1)$ bases containing what the agent a_i believes the goals of the other agents a_j are. Each of these goals is supposed to have a priority level $\gamma_{r,j}^i$.

$\mathcal{KO}_j^i = \{(ko_{t,j}^i, \delta_{t,j}^i), t = 1, s_{ko}(j)\}$ where $j \neq i$, $ko_{t,j}^i \in Wff(\mathcal{L})$, are $(n-1)$ bases containing what the agent a_i believes the knowledge of the other agents a_j are. Each of these beliefs has a certainty level $\delta_{t,j}^i$.

This latter base is useful only if the agents intend to simulate the reasoning of the other agents. In negotiation dialogues where agents are trying to find a common agreement, it is more important for each agent to consider the beliefs that it has on the other agents'goals rather than those on their knowledge. Indeed, a common agreement can be more easily reached if the agents check that their offers may be consistent with what they believe are the goals of the others. So in what follows, we will omit the use of the bases $\mathcal{KO}_j^i$.

The different certainty levels and priority degrees are assumed to belong to a unique linearly ordered scale T with maximal element denoted by 1 (corresponding to total certainty and full priority) and a minimal element denoted by 0 corresponding to the complete absence of certainty or priority. m will denote the order-reversing map of the scale. In particular, $m(0) = 1$ and $m(1) = 0$.

We shall denote by $\mathcal{K}^*$ and $\mathcal{G}^*$ the corresponding sets of classical propositions when weights are ignored.

3 Argued Decisions

Recently, Amgoud and Prade [5] have proposed a formal framework for making decisions under uncertainty on the bases of arguments that can be built in favor or against a possible choice. Such an approach has two obvious merits. First, decisions can be more easily explained. Moreover, argumentation-based decision is maybe closer to the way humans make decisions than approaches requiring explicit utility functions and uncertainty distributions. Decisions for an agent are computed from stratified knowledge and preference bases in the sense of Section 2. This approach distinguishes between a *pessimistic* attitude, which focuses on the existence of strong arguments that support a decision, and an *optimistic* one, which concentrates on the absence of strong arguments against a considered choice. This approach can be related to the estimation of qualitative pessimistic and optimistic expected utility measures. Indeed, such measures can be obtained from a qualitative plausibility distribution and a qualitative preference profile that can be associated with a stratified knowledge base and with a stratified set of goals [5].

In this paper, we only use the syntactic counterpart of these semantical computations in terms of distribution and profile (which has been proved to be equivalent for selecting best decisions), under its argumentative form. This syntactic approach is now recalled and illustrated on an example.

The idea is that a decision is justified and supported if it leads to the satisfaction of at least the most important goals of the agent, taking into account the most certain part of knowledge. Let $\mathcal{D}$ be the set of all possible decisions, where a decision d is a literal.

Definition 1 (Argument PRO). *An* argument in favor of *a decision d is a triple $A = <S, C, d>$ such that:*

- *$d \in \mathcal{D}$*
- *$S \subseteq \mathcal{K}^*$ and $C \subseteq \mathcal{G}^*$*
- *$S \cup \{d\}$ is consistent*
- *$S \cup \{d\} \vdash C$*
- *S is minimal and C is maximal (for set inclusion) among the sets satisfying the above conditions.*

$S = Support(A)$ is the support *of the argument, C = Consequences(A) its* consequences *(the goals which are reached by the decision d) and $d = Conclusion(A)$ is the conclusion of the argument. The set $\mathcal{A}_P$ gathers all the arguments which can be constructed from $<\mathcal{K}, \mathcal{G}, \mathcal{D}>$.*

Due to the stratification of the bases $\mathcal{K}_i$ and $\mathcal{G}_i$, arguments in favor of a decision are more or less strong for i.

Definition 2 (Strength of an Argument PRO). *Let $A = < S, C, d >$ be an argument in $\mathcal{A}_P$.*
The strength *of A is a pair $<Level_P(A), Weight_P(A)>$ such that:*

- *The* certainty level *of the argument is $Level_P(A) = min\{\rho_i \mid k_i \in S$ and $(k_i, \rho_i) \in \mathcal{K}\}$. If $S = \emptyset$ then $Level_P(A) = 1$.*
- *The* degree of satisfaction *of the argument is $Weight_P(A) = m(\beta)$ with $\beta = max\{\lambda_j \mid (g_j, \lambda_i) \in \mathcal{G}$ and $g_j \notin C\}$. If $\beta = 1$ then $Weight_P(A) = 0$ and if $C = \mathcal{G}^*$ then $Weight_P(A) = 1$.*

Then, strengths of arguments make it possible to compare pairs of arguments as follows:

Definition 3. *Let A and B be two arguments in $\mathcal{A}_P$. A is* preferred *to B, denoted $A \succeq_P B$, iff $min(Level_P(A), Weight_P(A)) \geq min(Level_P(B), Weight_P(B))$.*

Thus arguments are constructed in favor of decisions and those arguments can be compared. Then decisions can also be compared on the basis of the relevant arguments.

Definition 4. *Let $d, d' \in \mathcal{D}$. d is* preferred *to d', denoted $d \triangleright_P d'$, iff $\exists\ A \in \mathcal{A}_P$, $Conclusion(A) = d$ such that $\forall\ B \in \mathcal{A}_P$, $Conclusion(B) = d'$, then $A \succeq_P B$.*

This decision process is pessimistic in nature since it is based on the idea of making sure that the important goals are reached. An optimistic attitude can be also captured. It focuses on the idea that a decision is all the better as there is no strong argument against it.

Definition 5 (Argument CON). *An* argument against *a decision d is a triple* $A = <S, C, d>$ *such that:*

- $d \in \mathcal{D}$
- $S \subseteq \mathcal{K}^*$ *and* $C \subseteq \mathcal{G}^*$
- $S \cup \{d\}$ *is consistent*
- $\forall\ g_i \in C,\ S \cup \{d\} \vdash \neg g_i$
- *S is minimal and C is maximal (for set inclusion) among the sets satisfying the above conditions.*

$S = Support(A)$ *is the* support *of the argument,* $C = Consequences(A)$ *its* consequences *(the goals which are not satisfied by the decision d), and* $d = Conclusion(A)$ *its conclusion. The set* $\mathcal{A}_O$ *gathers all the arguments which can be constructed from* $<\mathcal{K}, \mathcal{G}, \mathcal{D}>$*.*

Note that the consequences considered here are the negative ones. Again, arguments are more or less strong or weak.

Definition 6 (Weakness of an Argument CON). *Let* $A = < S, C, d >$ *be an argument of* $\mathcal{A}_O$*.*
The weakness *of A is a pair* $<Level_O(A), Weight_O(A)>$ *such that:*

- *The* level *of the argument is* $Level_O(A) = m(\varphi)$ *such that* $\varphi = min\{\rho_i \mid k_i \in S$ *and* $(k_i, \rho_i) \in \mathcal{K}\}$*. If* $S = \emptyset$ *then* $Level_O(A) = 0$*.*
- *The* degree *of the argument is* $Weight_O(A) = m(\beta)$ *such that* $\beta = max\{\lambda_j$ *such that* $g_j \in C$ *and* $(g_j, \lambda_i) \in \mathcal{G}\}$*.*

Once we have defined the arguments and their weaknesses, pairs of arguments can be compared. Clearly, decisions for which all the arguments against it are weak will be preferred, i.e. we are interested in the least weak arguments against a considered decision. This leads to the two following definitions:

Definition 7. *Let A and B be two arguments in* $\mathcal{A}_O$*. A is* preferred *to B, denoted* $A \succeq_O B$*, iff* $max(Level_O(A), Weight_O(A)) \geq max(Level_O(B), Weight_O(B))$*.*

As in the pessimistic case, decisions are compared on the basis of the relevant arguments.

Definition 8. *Let* $d, d' \in \mathcal{D}$*. d is* preferred *to d', denoted* $d \rhd_O d'$*, iff* $\exists\ A \in \mathcal{A}_O$ *with* $Conclusion(A) = d$ *such that* $\forall\ B \in \mathcal{A}_O$ *with* $Conclusion(B) = d'$*, then A is preferred to B.*

Let us illustrate this approach using the two points of view (pessimistic and optimistic) on an example about deciding or not to argue in a multiple agent dialogue for an agent which is not satisfied with the current offer.

Example 1. *The knowledge base is* $\mathcal{K} = \{(a \to suu, 1), (\neg a \to \neg\ suu, 1), (a \to \neg aco, 1), (fco \wedge \neg a \to aco, 1), (sb, 1), (\neg fco \to \neg aco, 1), (sb \to fco, \lambda)\}$ $(0 < \lambda < 1)$ *with the intended meaning:*
suu: saying something unpleasant,
fco: other agents in favor of current offer,
aco: obliged to accept the current offer,
a: argue,
sb: current offer seems beneficial for the other agents.
The base of goals is $\mathcal{G} = \{(\neg aco, 1), (\neg suu, \sigma)\}$ *with* $(0 < \sigma < 1)$.
The agent does not like to say something unpleasant, but it is more important not to be obliged to accept the current offer.
The set of decisions is $\mathcal{D} = \{a, \neg a\}$, *i.e., arguing or not.*

There is one argument in favor of the decision 'a': $< \{a \to \neg aco\}, \{\neg aco\}, a>$. *There is also a unique argument in favor of the decision '*$\neg a'$*':* $<\{\neg a \to \neg suu\}, \{\neg suu\}, \neg a>$.

The level of the argument $<\{a \to \neg aco\}, \{\neg aco\}, a>$ *is 1 whereas its weight is* $m(\sigma)$. *Concerning the argument* $<\{\neg a \to \neg suu\}, \{\neg suu\}, \neg a>$, *its level is 1 and its weight is* $m(1) = 0$.

The argument $<\{a \to \neg aco\}, \{\neg aco\}, a>$ *is preferred to the argument* $<\{\neg a \to \neg suu\}, \{\neg suu\}, \neg a>$.

¿From a pessimistic point of view, decision a *is preferred to the decision* $\neg a$ *since* $<\{a \to \neg aco\}, \{\neg aco\}, a>$ *is preferred to* $<\{\neg a \to \neg suu\}, \{\neg suu\}, \neg a>$.

Let us examine the optimistic point of view. There is one argument against the decision 'a': $<\{a \to suu\}, \{\neg suu\}, a>$. *There is also a unique argument against the decision* $\neg a$*:* $<\{sb, sb \to fco, fco \wedge \neg a \to aco\}, \{\neg aco\}, \neg a>$.

The level of the argument $<\{a \to suu\}, \{\neg suu\}, a>$ *is 0 whereas its degree is* $m(\sigma)$. *Concerning the argument* $<\{sb, sb \to fco, fco \wedge \neg a \to aco\}, \{\neg aco\}, \neg a>$, *its level is* $m(\lambda)$, *and its degree is 0.*

Then the comparison of the two arguments amounts to compare $m(\sigma)$ *with* $m(\lambda)$.

The final recommended decision with the optimistic approach depends on this comparison.

This argumentation system will be used to take decisions about the offers to propose in a negotiation dialogue. The following definition is the same as Definition 1 where the decision d is about offers.

Definition 9 (Argument for an offer). *An* argument *in favor of an offer* x *is a triple* $A = <S, C, x>$ *such that:*

- $x \in X$
- $S \subseteq \mathcal{K}^*$ *and* $C \subseteq \mathcal{G}^*$
- $S(x)$ *is consistent*

- $S(x) \vdash C(x)$
- *S is minimal and C is maximal (for set inclusion) among the sets satisfying the above conditions.*

X is the set of offers, $S = Support(A)$, $C = Consequences(A)$ (the goals which are satisfied by the offer x) and $x = Conclusion(A)$. $S(x)$ (resp. $C(x)$) denotes the belief state (resp. the preference state) when an offer x takes place.

Example 2. *The example is about an agent wanting to propose an offer corresponding to its desired place for holidays.*
The set of available offers is $X = \{Tunisia, Italy\}$.
Its knowledge base is:

$\mathcal{K} = \{(Sunny(Tunisia), 1), (\neg Cheap(Italy), \beta), (Sunny(x) \rightarrow Cheap(x), 1)\}$.
Its preferences base is: $\mathcal{G} = \{(Cheap(x), 1)\}$.

The decision to take by the agent is whether to offer Tunisia or Italy. Following the last definition, it has an argument in favor of Tunisia:

$A =< \{Sunny(Tunisia), Sunny(x) \rightarrow cheap(x)\}, cheap(Tunisia), tunisia >$.
It has no argument in favor of Italy (it violates its goal which is very important). So this agent will offer Tunisia.

4 The Negotiation Protocol

4.1 Formal Setting

In this section, we propose a formal protocol handling negotiation dialogues between many agents ($n \geq 2$). Agents having to discuss several offers, the protocol is supposed to be run as many times as there are non-discussed offers, and such that a common agreement is still not found. The agents take turns to start new runs of the protocol and only one offer is discussed at each run.

A negotiation interaction protocol is a tuple ⟨ Objective, Agents, Object, Acts, Replies, Wff-Moves, Dialogue, Result⟩ such that:

Objective is the aim of the dialogue which is to find an acceptable offer.
Agents is the set of agents taking part to the dialogue, $Ag = \{a_0, \ldots, a_{n-1}\}$.
Object is the subject of the dialogue. It is a multi-issue one, denoted by the tuple $\langle O_1, \ldots, O_m \rangle$, $m \geq 1$. Each O_i is a variable taking its values in a set T_i.
Let X be the set of all possible offers, its elements are $x = \langle x_1, \ldots, x_m \rangle$ with $x_i \in T_i$.
Acts is the set of possible negotiation speech acts: *Acts* = {Offer, Challenge, Argue, Accept, Refuse, Withdraw, Say nothing}.
Replies: Acts $\longrightarrow$ Power(Acts), is a mapping that associates to each speech act its possible replies.

- $Replies(Offer)$ = {Accept, Refuse, Challenge}
- $Replies(Challenge)$ = {Argue}
- $Replies(Argue)$ = {Accept, Challenge, Argue}
- $Replies(Accept)$ = {Accept, Challenge, Argue, Withdraw}
- $Replies(Refuse)$ = {Accept, Challenge, Argue, Withdraw}
- $Replies(Withdraw) = \emptyset$

Well-founded moves $=\{M_0, \ldots, M_p\}$ is a set of tuples $M_k = \langle S_k, H_k, Move_k \rangle$, such that:
- $S_k \in Agents$, the agent which plays the move is given by the function $Speaker(M_k) = S_k$.
- $H_k \subseteq Agents \setminus \{S_k\}$, the set of agents to which the move is addressed is given by the function $Hearer(M_k) = H_k$.
- $Move_k = Act_k(c_k)$ is the uttered move where Act_k is a speech act applied to a content c_k.

Dialogue is a finite non-empty sequence of well-founded moves $\mathcal{D} = \{M_0, \ldots, M_p\}$ such that:
- $M_0 = \langle S_0, H_0, offer(x) \rangle$: each dialogue starts with an offer $x \in X$
- $Move_k \neq offer(x), \forall k \neq 0$ and $\forall x \in X$: only one offer is proposed during the dialogue at the first move
- $Speaker(M_k) = a_{k\ modulo\ n}$: the agents take turns during the dialogue.
- $Speaker(M_k) \notin Hearer(M_k)$. This condition forbids an agent to address a move to itself.
- $Hearer(M_0) = a_j, \forall j \neq i$: the agent a_i which utters the first move addresses it to all the agents.
- For each pair of tuples M_k, M_h, $k \neq h$, if $S_k = S_h$ then $Move_k \neq Move_h$. This condition forbids an agent to repeat a move that it has already played.

These conditions guarantee that the dialogue $\mathcal{D}$ is ***non circular***.

Result: $\mathcal{D} \longrightarrow \{success, failure\}$, is a mapping which returns the result of the dialogue.
- $Result(\mathcal{D}) = success$ if the preferences of the agents are satisfied by the current offer.
- $Result(\mathcal{D}) = failure$ if the most important preferences of at least one agent are violated by the current offer.

This protocol is based on dialogue games. Each agent is equipped with a *commitment store* (CS) [9] containing the set of facts it is committed to during the dialogue.

Using the idea introduced in [2] of decomposing the agents' commitments store (CS) into many components, we suppose that each agent's CS has the structure

$$CS = \langle \mathcal{S}, \mathcal{A}, \mathcal{C} \rangle$$

with:

$CS.\mathcal{S}$ contains the offers proposed by the agent and those it has accepted ($CS.\mathcal{S} \subseteq X$),
$CS.\mathcal{A}$ is the set of arguments presented by the agent ($CS.\mathcal{A} \subseteq Arg(\mathcal{L})$), where $Arg(\mathcal{L})$) is the set of all arguments we can construct from $\mathcal{L}$,
$CS.\mathcal{C}$ is the set of challenges made by the agent.

At the first run of the protocol, all the CS are empty. This is not the case when the protocol is run again. Indeed, agents must keep their previous commitments to avoid to repeat what they have already uttered during previous runs of the protocol.

4.2 Conditions on the Negotiation Acts

In what follows, we specify for each act its pre-conditions and post-conditions (effects). For the agents' commitments (CS), we only specify the changes to effect. We suppose that agent a_i addresses a move to the $(n-1)$ other agents.

Offer(x) where $x \in X$. It's the basic move in negotiation. The idea is that an agent chooses an offer x for which there are the strongest supporting arguments (w.r.t. $\mathcal{G}_i$). Since the agent is *cooperative* (it tries to satisfy its own goals taking into account the goals of the other agents), this offer x is the also the one for which there exists no strong argument against it (using $\mathcal{GO}^i_j$ instead of $\mathcal{G}_i$).

Pre-conditions: Among the elements of X, choose x which is preferred to any $x' \in X$ such that $x' \neq x$, in the sense of definition 4, provided that there is no strong argument against the offer x (i.e. with a weakness degree equal to 0) where $\mathcal{G}_i$ is changed into $\mathcal{GO}^i_j$, $\forall j \neq i$ in definition 8.
Post-conditions: $CS.\mathcal{S}_t(a_i) = CS.\mathcal{S}_{t-1}(a_i) \cup \{x\}$.

Challenge(x) where $x \in X$. This move incites the agent which receives it to give an argument in favor of the offer x. An agent asks for an argument when this offer is not acceptable for it and it knows that there are still non-rejected offers.

Pre-conditions: $\exists x' \in X$ such that x' is preferred to x w.r.t. definition 4.
Post-conditions: $CS.\mathcal{C}_t(a_i) = CS.\mathcal{C}_{t-1}(a_i) \cup \{x\}$: the agent a_i which played the move $Challenge(x)$ keeps it in its CS.

Challenge(y) where $y \in Wff(\mathcal{L})$. This move incites the agent which receives it to give an argument in favor of the proposition y.

Pre-conditions: There is no condition.
Post-conditions: $CS.\mathcal{C}_t(a_i) = CS.\mathcal{C}_{t-1}(a_i) \cup \{y\}$: the agent a_i which played the move $Challenge(y)$ keeps it in its CS.

Argue(S) with $S = \{(k_p, \alpha_p), p = 1, s\} \subseteq \mathcal{K}_i$ is a set of formulas representing the support of an argument given by agent a_i. In [5], it is shown how to compute and evaluate *acceptable* arguments.

Pre-conditions: S is acceptable.
Post-conditions: $CS.\mathcal{A}_t(a_i) = CS.\mathcal{A}_{t-1}(a_i) \cup S$. If S is acceptable (according to the definition given in [5]), the agents a_j revise their base $\mathcal{K}_j$ into a new base $(\mathcal{K}_j)^*(S)$.

Withdraw: An agent can withdraw from the negotiation if it hasn't any acceptable offer to propose.

Pre-conditions: $\forall x \in X$, there is an argument with maximal strength against x, or $(X = \emptyset)$.
Post-conditions: *(Result*$(\mathcal{D})$ *= failure)* and $\forall i, CS_t(a_i) = \emptyset$. As soon as an agent withdraws, the negotiation ends and all the commitment stores are emptied.
We suppose the dialogue ends this way because we aim to find a compromise between the n agents taking part to the negotiation.

Accept(x) where $x \in X$. This move is played when the offer x is acceptable for the agent.

Pre-conditions: The offer x is the most preferred decision in X in the sense of definition 4.
Post-conditions: $CS.\mathcal{S}_t(a_i) = CS.\mathcal{S}_{t-1}(a_i) \cup \{x\}$.
If $x \in CS.\mathcal{S}(a_i)$, $\forall i$, then *Result*$(\mathcal{D})$ = *success*, i.e if all the agents accept the offer x, the negotiation ends with x as compromise.

Accept(S) $S \subset Wff(\mathcal{L})$.

Pre-conditions: S is acceptable for a_i.
Post-conditions: $CS.\mathcal{A}_t(a_i) = CS.\mathcal{A}_{t-1}(a_i) \cup S$.

Refuse(x) where $x \in X$. An agent refuses an offer if it is not acceptable for it.

Pre-conditions: There exists an argument in the sense of definition 5 against x.
Post-conditions: If $\forall a_j, \nexists(S, x)$, i.e. if there not exist any acceptable argument for x then $X = X \setminus \{x\}$. A rejected offer is removed from the set X. *Result*$(\mathcal{D})$ = *failure.*

Say nothing: This move allows an agent to miss its turn if it has already accepted the current offer, or it has no argument to present. This move has no effect on the dialogue.

4.3 Properties of the Negotiation Protocol

Property 1 (Termination). *Any negotiation between n agents managed by our protocol ends, either with* $\text{Result}(\mathcal{D}) = \text{success}$ *or* $\text{Result}(\mathcal{D}) = \text{failure}$.

Property 2 (Optimal outcome). *If the agents do not misrepresent the preferences of the other agents ($\mathcal{GO}^i_j$), then the compromise found is an offer x which is preferred to any other offer $x' \in X$ in the sense of definition 4, for all the agents.*

5 Example of Deliberative Choice

We illustrate our negotiation protocol through an example of dialogue between three agents: Mary, John and Peter, partners on a common project aiming at setting a town and a date for their next meeting. The negotiation object O is in this case the couple $(Town, Date)$ denoted $\langle t, d \rangle$, where t is for the town and d the date.

Suppose that the set of offers is $X = \{(V, E), (L, S), (V, J)\}$, i.e. the meeting will take part either in Valencia (denoted V), at one of the dates respectively denoted E and J; or in London (denoted L) at the date denoted S.

In what follows, we use the following scale $T = \{a, b, c, d\}$ with the condition $a > b > c > d$. We recall that m is the order reversing map on the scale T such that $m(a) = d$ and $m(b) = c$.

Suppose Mary has the following *beliefs*:

$\mathcal{K}_0$ = {(disposable(V,E), 1), (disposable(t,d) → meet(t,d), 1), (free(V,E),1), (¬ free(L,S) ,1), (disposable(t,J),1)}.

The *goals* of Mary are to meet her partners in any town and at any date, provided that accommodations are free. This can be written: $\mathcal{G}_0$ = {(meet, 1),(free, b)}.

Where "*meet*" is a short for $(meet(V, E) \vee meet(L, S) \vee meet(V, J))$. "*free*" is defined the same way. We use this type of abbreviation in what follows.

Suppose John's beliefs are: $\mathcal{K}_1 = \{(hot(V, d), a), (\neg hot(L, S), 1),$ (disposable $(L, S), 1),$

$(disposable(t, d) \rightarrow meet(t, d), 1), (meet(V, J) \rightarrow work\_saturday, 1)\}$.

His goals are to meet his partners in any town and at any date, and that this town must be not hot at this date. We write:

$\mathcal{G}_1 = \{(meet, 1), (\neg hot, c)\}$.

Finally we suppose Peter's beliefs are:
$\mathcal{K}_2 = \{(\neg meet(V, E), 1), (\forall d \neq E, meet(V, d), 1),$
$(disposable(t, d) \rightarrow meet(t, d), 1), (disposable(V, J), b),$
$(manager, 1), (manager \rightarrow work\_saturday, 1)\}$.

His goals are to meet his partners and to don't work on Saturday. We write:
$\mathcal{G}_2 = \{(meet, 1), (\neg work\_saturday, d)\}$.

For simplicity, we suppose that Mary, John and Peter ignore the preferences of each other. This means that $\mathcal{GO}_j^i = \emptyset$, $\forall i, j$.

In what follows, we illustrate the dialogue between the agents and give the moves played by each agent.

First run of the protocol
Mary starts the dialogue by proposing an offer.

Mary: The next meeting should be in Valencia during the conference ECAI. $Offer(V, E)$.
Pre-condition: (V,E) is the most preferred decision for Mary.
Post-condition: $CS.\mathcal{S}(Mary) = \{(V, E)\}$.
John: Why? $Challenge(V, E)$.
Pre-condition: For John, there exists another decision which is preferred to (V,E).
Post-condition: $CS.\mathcal{C}(John) = \{(V, E)\}$.
Peter: What are the advantages? $Challenge(V, E)$.
Pre-condition: For Peter, this decision violates his most important goal.
Post-condition: $CS.\mathcal{C}(Peter) = \{(V, E)\}$.
Mary: I think we can meet as soon as it will be during ECAI. $Argue(meet(V, E))$.
Pre-condition: The argument is acceptable.
Post-condition: $CS.\mathcal{A}(Mary) = \{disposable(V, E), disposable(V, E) \rightarrow meet(V, E)\}$.
John: I refuse Valencia because it is hot. $Argue(hot(V, d))$.
Pre-condition: $\{hot(V, d)\}$ is an acceptable argument.
Post-condition: $CS.\mathcal{A}(John) = \{hot(V, d)\}$.
Peter: For my part, I will not be able to meet you. $Argue(\neg meet(V, E))$.
Pre-condition: $\{\neg meet(V, E)\}$ is an acceptable argument.
Post-condition: $CS.\mathcal{A}(Peter) = \{\neg meet(V, E)\}$.
Mary: Nevertheless the accommodation will be free. $Argue(free(V, E))$.
Pre-condition: $\{free(V, E)\}$ is an acceptable argument.
Post-condition: $CS.\mathcal{A}(Mary) = CS.\mathcal{A}(Mary) \cup \{free(V, E)\}$.
John: It still doesn't fit me. $Refuse(V, E)$.
Pre-condition: the offer violates one of his goals.
Peter: Neither do I. $Refuse(V, E)$.
Pre-condition: the offer violates his most important goal.
Post-condition: *Result($\mathcal{D}$) = failure.*
$X = X \setminus \{(V, E)\}$ and all the CS are emptied except the components of the arguments.

***Second run of the protocol*:** It is started by John.

John: What about London in September ? $Offer(L, S)$.
Pre-condition: (L,S) is the most preferred decision for John.
Post-condition: $CS.\mathcal{S}(John) = \{(L, S)\}$.
Peter: I refuse. $Refuse(L, S)$.
Pre-condition: this offer violates his most important goal.
Mary: John, what are your arguments in favor of your offer ? $Challenge(L, S)$.
Pre-condition: (L,S) is not the preferred decision for Mary.
Post-condition: $CS.\mathcal{C}(Mary) = \{(L, S)\}$.

John: London is not hot and I will be able to meet you.
$Argue(\neg hot(L,S), meet(L,S))$.
Pre-condition: The argument is acceptable.
Post-condition: $CS.\mathcal{A}(John) = CS.\mathcal{A}(John) \cup \{\neg hot(L,S), disposable(L,S), disposable(L,S) \rightarrow meet(L,S)\}$.
Peter: But we have already said that the meeting should be in Valencia !
$Argue(\forall d \neq E, meet(V,d))$.
Pre-condition: $\{\forall d \neq E, meet(V,d)\}$ is an acceptable argument.
Post-condition: $CS.\mathcal{A}(Peter) = CS.\mathcal{A}(Peter) \cup \{\forall d \neq E, meet(V,d)\}$.
Mary: I can't pay the accommodation. $Argue(\neg free(L,S))$.
Pre-condition: $\{\neg free(L,S)\}$ is acceptable.
Post-condition: $CS.\mathcal{A}(Mary) = CS.\mathcal{A}(Mary) \cup \{\neg free(L,S)\}$.
John: I forgot that we have chosen Valencia to organize the meeting. $Accept(\forall d \neq E, meet(V,d))$.
Pre-condition: $\{\forall d \neq E, meet(V,d)\}$ is an acceptable argument.
Post-condition: $CS.\mathcal{A}(John) = CS.\mathcal{A}(John) \cup \{\forall d \neq E, meet(V,d)\}$.
John revises its belief base which becomes: $\mathcal{K}_1$ = {(hot(V,d), 0.9), (¬ hot(L,S), 1), (disposable(L,S), 1), (disposable(t,d) → meet(t,d), 1), (meet(V,J) → work_saturday, 1), (∀ d ≠ E, meet(V,d), 1)}.
Peter: I have nothing more to say. *Say_nothing*.
Mary: I also forgot this. $Accept(\forall d \neq E, meet(V,d))$.
Pre-condition: $\{\forall d \neq E, meet(V,d)\}$ is acceptable.
Post-condition: $CS.\mathcal{A}(Mary) = CS.\mathcal{A}(Mary) \cup \{\forall d \neq E, meet(V,d)\}$.
Mary revises its belief base which becomes:

$\mathcal{K}_0$ = {(disposable(V,E), 1), (disposable(t,d) → meet(t,d), 1), (free(V,E),1), (¬ free(L,S) ,1), (disposable(t,J),1), (∀ d ≠ E, meet(V,d), 1)}.

Result($\mathcal{D}$) = failure, $X = X \setminus \{(L,S)\} = \{(V,J)\}$ and all the CS are emptied except the components of the arguments.

***Third run of the protocol*:** It is Peter's turn to propose an offer.

Peter: What about reorganizing the sessions to satisfy every body? Offer(V,J).
Pre-condition: the decision (V,J) is the most preferred one.
Post-condition: $CS.\mathcal{S}(Peter) = \{(V,J)\}$.
Mary: It fits me. $Accept(V,J)$.
Pre-condition: (V,J) is the most preferred decision for her.
Post-condition: $CS.\mathcal{S}(Mary) = \{(V,J)\}$.
John: Not me ! $Refuse(V,J)$.
Pre-condition: the decision (V,J) violates one of his goals.
Peter: John, what doesn't fit you ?
$Challenge(Refuse(V,J))$.
Pre-condition: There aren't.
Post-condition: $CS.\mathcal{C}(Peter) = \{(V,J)\}$,
Mary: I have nothing to say. *Say_nothing*.

John: If we organize the sessions this way, the managers would have to work on Saturday. $Argue(work\_saturday)$.
Pre-condition: the argument is acceptable.
Post-condition: $CS.\mathcal{A}$(John) = CS.$\mathcal{A}$(John) ∪ {meet(V,J), meet(V,J) → work_saturday}.

Peter: The managers can make the effort of working on Saturday. $Argue(manager, manager \rightarrow work\_saturday)$.
Pre-condition: Peter has an acceptable argument to convince John: $\{manager, manager \rightarrow work\_saturday\}$.
Post-condition: $CS.\mathcal{A}$(Peter) = $CS.\mathcal{A}$(Peter) ∪ {manager, manager → work_saturday}.

Mary: I have nothing to say. $Say\_nothing$.

John: I think you don't let me any choice ! Accept(manager, manager → work_saturday).
Pre-condition: The argument is acceptable.
Post-condition: $CS.\mathcal{A}$(John) = $CS.\mathcal{A}$(John) ∪ {manager, manager → work_saturday}.

Furthermore, the offer (V, J) is the most preferred one in X in the sense of definition 4.

In other words, all the agents have accepted the offer (V, J) and *Result($\mathcal{D}$) = success*.

The negotiation dialogue ends with a compromise found by the agents to organize their meeting: in Valencia at the date J.

6 Conclusion

This paper has proposed a general formal framework for handling negotiation dialogues where autonomous agents aim at finding a common agreement about a collective choice. The agents are equipped with knowledge bases graded in certainty levels and gathering what they know about the environment, and with preference bases representing their more or less important goals.

The reasoning model of the agents is captured by a formal decision framework. The basic idea is that an agent utters and accepts offers which are supported by strong arguments. Similarly, agents refuse or challenge offers for which there exists at least one strong argument against them.

The interaction between agents is captured by a protocol which is run at most as many times as there non discussed offers, and such that at each run only one offer is discussed. If it is accepted by all the agents, then an agreement is found. In the opposite case, it is removed from the set of offers and another one is proposed.

In future work, we plan to propose a protocol less restrictive by considering stratified sets to store the rejected offers. A level of rejection will be computed to allow the affectation of the offers to the different sets. The last set in the stratification will gather the offers which are definitively rejected, i.e. those which are impossible. Once all the offers are studied without finding an acceptable one, the agents negotiate again on the set gathering the less rejected offers and proceed the same way. This requires that the agents revise their bases by being less demanding regarding their preferences.

References

1. L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In *Proc. 4th Intl. Conf. on Multi-Agent Systems (ICMAS'00)*, pages 31–38, Boston, 2000.
2. L. Amgoud, N.Maudet, and S. Parsons. An argumentation-based semantics for agent communication languages. In *Proc. 15th Eur. Conf. on Artificial Intelligence (ECAI'02)*, pages 38–42, Lyon, 2002.
3. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In *Proc. 14th Eur. Conf. on Artificial Intelligence (ECAI'00)*, pages 338–342, Berlin, 2000.
4. L. Amgoud and H. Prade. Reaching agreement through argumentation: A possibilistic approach. In *Proc. 9th Intl. Conf. on the Principles of Knowledge Representation and Reasoning (KR'04)*, pages 175–182, Whistler, 2004.
5. L. Amgoud and H. Prade. Using arguments for making decisions: A possibilistic logic approach. In *Proc. 20th Conf. on Uncertainty in Artificial Intelligence (UAI'04)*, pages 10–17, Banff, 2004.
6. S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda based framework for multi-issues negotiation. *Artificial Intelligence*, 152:1–45, 2004.
7. A. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proc. 2nd Intl. Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS'03)*, pages 883–890, Melbourne, 2003. ACM Press.
8. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104(1–2):1–69, 1998.
9. J. MacKenzie. Question-begging in non-cumulative systems. *Journal of philosophical logic*, 8:117–133, 1979.
10. S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.
11. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4):343–375, 2003.
12. C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In *Proc. 4th Intl. Workshop on Agent Theories, Architectures and Languages (ATAL'97)*, pages 167–182, Providence, 1997.