Reaching Agreement Through Argumentation: A Possibilistic Approach

Leila Amgoud and **Henri Prade**

Institut de Recherche en Informatique de Toulouse (IRIT)
118, route de Narbonne, 31062 Toulouse, France
{amgoud, prade}@irit.fr

Abstract

Negotiation plays a key role as a means for sharing information and resources with the aim of looking for a common agreement. This paper proposes a new approach based on possibility theory, which integrates both the merits of argumentation-based negotiation and of heuristic methods looking for making trade-offs. Possibilistic logic is used as a unified setting, which proves to be convenient not only for representing the mental states of the agents (beliefs possibly pervaded with uncertainty, and prioritized goals), but also for revising the belief bases and for describing the decision procedure for selecting a new offer.

Key words: Multi-agents systems, Negotiation, Possibilistic logic, Argumentation.

Introduction

In most agent applications, the autonomous components need to interact with one another because of the inherent interdependencies which exist between them, and negotiation is the predominant mechanism for achieving this by means of an exchange of offers. Agents make offers that they find acceptable and respond to offers made to them.

This paper proposes a negotiation model which integrates both the merits of argumentation and of heuristic methods looking for making trade-offs. We are particularly interested in *deliberative negotiations*. In such negotiations, the agents try to find an agreement on a given subject. For example, let's consider the case of two agents who discuss about the destination of their next holidays. Each agent has a set of goals that it wants to satisfy. Since, the goals of the agents may be conflicting, then each agent try to convince the other agent to change its own goals. Argumentation is the best means to do that.

To build such deliberative negotiation models, one should specify the following parameters:

- **The mental states of the agents:** (their beliefs and goals).
- **Argumentation rules**: Argument generation (what is an argument and how it is built from the mental states of the agent) and argument evaluation (what is a *good / acceptable* argument and what is a *rejected* one).
- **Decision rules:** an agent will have to make a three stage decision. It should: 1) select the content of a move if necessary, 2) decide when a given move may be played and 3) choose the following move to play among all the possible ones.
- **Revision rules:** an agent may revise its beliefs, goals when necessary.

In this paper, possibilistic logic is used as a unified setting, which proves to be convenient not only for *representing* the mental states of the agents (beliefs possibly pervaded with uncertainty, and prioritized goals), but also for *revising* the bases and for describing the *decision* procedure.

This paper is organized as follows: we start by presenting our modeling of the agents' mental states in terms of three possibilistic bases. Then we introduce the different tools needed to build a deliberative negotiation. We present the argumentation framework which will be used to generate and evaluate arguments. We propose different criteria that can be used in the decision process. Finally, we present a revision procedure that will be used by the agents to update their bases. In a next section we describe the protocol used in our model of negotiation and we illustrate the model on an example. We then compare our work to the existing ones and finally we conclude and sugget some perspectives.

The mental states of the agents

Here negotiation dialogues take place between two agents a and $\overline{a}$. However, the proposed approach could be extended to several agents. Each negotiating agent is supposed to have a set $\mathcal{G}$ of *goals* to pursue, a knowledge base, $\mathcal{K}$, gathering the information it has about the environment, and finally a base $\mathcal{GO}$, containing what the agent believes the goals of the other agent are.

$\mathcal{K}$ may be pervaded with uncertainty (the beliefs are more or less certain), and the goals in $\mathcal{G}$ and $\mathcal{GO}$ may not have equal priority. Thus, levels of certainty are assigned to formulas in $\mathcal{K}$, and levels of priority are assigned to the goals. We obtain three possibilistic bases (Dubois, Lang, & Prade 1991) that model graded knowledge and preferences:

$$\mathcal{K} = \{(k_i, \alpha_i), i = 1, \ldots, n\},$$
$$\mathcal{G} = \{(g_j, \beta_j), j = 1, \ldots, m\},$$
$$\mathcal{GO} = \{(go_l, \delta_l), l = 1, \ldots, p\}$$

where k_i, g_j, go_l are propositions or closed first order formulas and $\alpha_i, \beta_j, \delta_l$ are elements of $[0, 1]$, corresponding to uncertainty, or priority levels.
The different bases of agent a will be denoted by: $\mathcal{K}_a$, $\mathcal{G}_a$, $\mathcal{GO}_a$ and those of agent $\overline{a}$ by: $\mathcal{K}_{\overline{a}}$, $\mathcal{G}_{\overline{a}}$, $\mathcal{GO}_{\overline{a}}$. A possibility distribution (Dubois, Lang, & Prade 1991), which expresses the semantics, is associated to each of the three bases: π_K, $\pi_\mathcal{G}$ and $\pi_{\mathcal{GO}}$.
Each agent is allowed to *change* its own goals in $\mathcal{G}$ during a negotiation and it may also *revise* its beliefs in $\mathcal{K}$ or in $\mathcal{G}O$.

In what follows, $\mathcal{O}$ will denote the negotiation object. This may be a quantitative item such as a price (if the two agents are negotiating about the price of a product) or a date. It may be also a symbolic item, for instance a country if the agents are negotiating about a destination for their next holidays.

Let X be a set of possible offers which can be made during a negotiation process by each agent. In fact, they represent the different values that can be assigned to the negotiation object $\mathcal{O}$. In the case of two agents who negotiate about a destination for their next holidays, X will contain a list of countries.

For the sake of simplicity, we suppose that the set X is common to the two agents. However, this can be extended to the case where each agent has its own set (X_a, $X_{\overline{a}}$).
Elements x of X are viewed as propositional variables. We denote by $\mathcal{K}_a^x$ the belief state of the agent a once x takes place. In fact, $\mathcal{K}_a^x$ is the projection of $\mathcal{K}_a$ on the beliefs which are related to x.

Tools for deliberative negotiation

The agents have a set of legal moves: they are allowed to make offers, to challenge a given offer, to justify an offer by arguments, to accept or to refuse an offer. An agent can also withdraw from the negotiation. Let $\mathcal{M}$ denote the complete set of moves. $\mathcal{M}$ = {Offer, Challenge, Argue, Accept, Refuse, Withdraw}. Note that in this paper, we consider the minimal set of moves which allows a negotiation dialogue to take place.
In addition to the semantics of the three above bases, each agent maintains three other possibility distribution over $\mathcal{M}$:

- $\pi_\mathcal{P}(m_i) = \alpha_i$ represents the possibility distribution given by the protocol.
- $\pi_\mathcal{D}(m_i) = \beta_i$ represents the possibility distribution given by the agent itself.
- finally their conjunctive combination $\pi_\mathcal{M}(m_i)$ = min($\pi_\mathcal{P}(m_i)$, $\pi_\mathcal{D}(m_i)$).

In this paper, the agent will further consider only the core of $\pi_\mathcal{M}$ i.e. $\{m_i | \pi_\mathcal{M}(m_i) = 1\}$. This last represents the fuzzy set of possible moves at a given step of the negotiation for the agent. The degrees α_i, β_i represent to what extent it is possible for an agent to make the move m_i.
We suppose that a move is either fully possible (α_i, β_i = 1) or impossible (α_i, β_i = 0). However, this work can be extended to take into account values between 0 and 1. This may be possible if we take into account the agent's strategy and the agent's profile. For example, if at a given step of the negotiation an agent is allowed to make either a new offer or a challenge, then if it is a cooperative agent, then making a challenge is more possible than making a new offer.

In what follows, we present the argumentation framework used to generate and evaluate arguments and some decision and revison rules.

The argumentation framework

Argumentation plays a key role in finding a compromise during a negotiation. Indeed, an offer supported by a good argument has a better chance to be accepted by another agent. Argumentation may also lead an agent to change its goals and finally may constrain an agent to respond in a particular way. For example, if an agent receives a threat, this agent may accept the offer even if it is not really acceptable for it.

In addition to explanatory arguments studied in classical argumentation frameworks, works on argumentation-based negotiation, namely in (Kraus, Sycara, & Evenchik 1998), have emphasized different other types of arguments such as *treats*, *rewards*, *appeals*, etc... In (Kraus, Sycara, & Evenchik 1998; Ramchurn, Jennings, & Sierra 2003), these arguments are treated as speech acts with pre-conditions and post-conditions. More recently, in (Amgoud & Prade 2004) we provided a logical framework which encompass the classical argumentation-based framework and handles the new types of arguments. More precisely, we gave the logical definitions of these arguments and their weighting systems. For the sake of simplicity, in this paper, we describe only one kind of arguments: the so-called *explanatory* arguments.

In this section we briefly introduce the argumentation system which forms the backbone of our approach. This is inspired by the work of Dung (Dung 1995) but goes further in dealing with preferences between arguments (further details are available in (Amgoud 1999)).

In what follows, $\vdash$ stands for classical inference and $\equiv$ for logical equivalence. $\mathcal{L}$ is a propositional language.

Definition 1 *An* argument *is a pair* (H, h) *where* h *is a formula of* $\mathcal{L}$ *and* $H = \{\phi_i$ *such that* $(\phi_i, \alpha_i) \in \mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}\}$ *such that:*

- H *is consistent,*
- $H \vdash h$*,*
- H *is minimal (for set inclusion).*

H *is called the* support *of the argument and* h *is its* conclusion. $Size(H)$ *returns the number of formulas in* H *and*

$Size_{\mathcal{GO}}(H)$ *returns the number of formulas of* H *which are in* $\mathcal{GO}$.

In general, since $\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}$ may be inconsistent, arguments in $\mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$, the set of all arguments which can be made from $\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}$, will conflict, and we make this idea precise with the notion of undercutting:

Definition 2 *Let* (H_1, h_1) *and* $(H_2, h_2) \in \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$. (H_1, h_1) undercuts (H_2, h_2) *iff* $\exists h \in H_2$ *such that* $h \equiv \neg h_1$.
In other words, an argument is undercut iff there exists an argument for the negation of an element of its support.

When dealing with threats and rewards, other defeasibility relations are introduced in (Amgoud & Prade 2004).
To capture the fact that some facts are more strongly believed (or desired, or intended, depending on the nature of the facts) we assume that any set of facts has a preference order over it.

Definition 3 *The certainty level of a nonempty subset* H *of* $\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}$, *is* $level(H) = min\{\alpha_i | (\phi_i, \alpha_i) \in \mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}$ *and* $\phi_i \in H\}$.

The certainty levels make it possible to compare different arguments as follows:

Definition 4 *Let* (H_1, h_1) *and* (H_2, h_2) *be two arguments in* $\mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$. (H_1, h_1) *is* preferred *to* (H_2, h_2), *denoted* $(H_1, h_1) \gg (H_2, h_2)$, *iff* $level(H_1) \geq level(H_2)$.

Example 1 *Let* $\mathcal{K} = \{(\neg a, 0.7), (a, 0.5), (a \rightarrow b, 0.5), (\neg b, 0.2)\}$. *Let's suppose that* $\mathcal{GO} = \emptyset$. *Now,* $(\{\neg a\}, \neg a)$ *and* $(\{a, a \rightarrow b\}, b)$ *are two arguments of* $\mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$. *The argument* $(\{\neg a\}, \neg a)$ *undercuts* $(\{a, a \rightarrow b\}, b)$. *The certainty level of* $\{a, a \rightarrow b\}$ *is 0.5 whereas the certainty level of* $\{\neg a\}$ *is 0.7, and so* $(\{\neg a\}, \neg a\}) \gg (\{a, a \rightarrow b\}, b)$.

The preference order makes it possible to distinguish different types of relation between arguments:

Definition 5 *Let* A, B *be two arguments of* $\mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$ *and* $S \subseteq \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$.

- *B* strongly undercuts *A iff B undercuts A and not* $A \gg B$.
- *If B undercuts A then A* defends itself *against B iff* $A \gg B$.
- $\mathcal{S}$ defends A *if there is some argument in* $\mathcal{S}$ *which strongly undercuts every argument* B *where* B *undercuts* A *and* A *cannot defend itself against* B.

Henceforth, $C_{Undercut,\gg}$ will gather all non-undercut arguments and arguments defending themselves against all their undercutting arguments.

In (Amgoud & Cayrol 1998), it was shown that the set $\underline{\mathcal{S}}$ of acceptable arguments of the argumentation system $\langle \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$, Undercut, $\gg \rangle$ is the least fixpoint of a function $\mathcal{F}$:

$$\mathcal{S} \subseteq \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}) \text{ and}$$
$$\mathcal{F}(\mathcal{S}) = \{(H,h) \in \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}) \text{ st. } (H,h) \text{ is defended by } \mathcal{S}\}$$

Definition 6 *The set of* acceptable *arguments for an argumentation system* $\langle \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO}), Undercut, \gg \rangle$ *is:*

$$\underline{\mathcal{S}} = \bigcup \mathcal{F}_{i \geq 0}(\emptyset)$$
$$= C_{Undercut,\gg} \cup [\bigcup \mathcal{F}_{i \geq 1}(C_{Undercut,\gg})]$$

An argument is acceptable if it is a member of the acceptable set.

Example 2 *(follows Example 1) The argument* $(\{\neg a\}, \neg a)$ *is in* $C_{Undercut,\gg}$ *because it is preferred (according to* $Pref$*) to the unique undercutting argument* $(\{a\}, a)$. *Consequently,* $(\{\neg a\}, \neg a)$ *is in* $\underline{\mathcal{S}}$. *The argument* $(\{\neg b\}, \neg b)$ *is undercut by* $(\{a, a \rightarrow b\}, b)$ *and does not defend itself. On the contrary,* $(\{\neg a\}, \neg a)$ *undercuts* $(\{a, a \rightarrow b\}, b)$ *and* $(\{\neg a\}, \neg a) \gg (\{a, a \rightarrow b\}, b)$. *Therefore,* $C_{Undercut,\gg}$ *defends* $(\{\neg b\}, \neg b)$ *and consequently* $(\{\neg b\}, \neg b) \in \underline{\mathcal{S}}$.

The decision rules

During a deliberative negotiation process, an agent will have to follow a three stage decision process for choosing next moves and their contents if any (arguments or offers):

Level 1: to select the content of a move if necessary. In fact this concerns only the argue and offer moves.

Level 2: to decide when a given move may be played. This consists of determining $\pi_{\mathcal{D}}$ and thus $\pi_{\mathcal{M}}$.

Level 3: to choose the following move to play within the core of $\pi_{\mathcal{M}}$.

In what follows, we will present different criteria for each decision stage. Let us start with level 1 stage.

Decision criteria at level 1

Argument selection: As said before, during a negotiation different types of arguments may be exchanged (threats, rewards, etc...). In (Kraus, Sycara, & Evenchik 1998), all argument types have been ordered from the weaker ones to the most agressive ones. Threats are considered as the most severe ones and explanatory arguments as the less severe ones. This ordering should be refined. Indeed, if for instance an agent chooses to present an explanatory argument, it should select among all the availble ones, the argument he will present. The same thing should be done for each argument type.
Since we have presented only the explanatory arguments, in what follows we will present some criteria for choosing between several such arguments.
We can imagine that the smallest arguments are preferred in order to restrict the exposure to defeaters. In what follows, $Pref$ will denote the preference relation used to select an

argument.

Definition 7 (Criterion of total size) *Let* (H_1, h) *and* $(H_2, h) \in \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$. (H_1, h) *Pref* (H_2, h) *iff* $Size(H_1) < Size(H_2)$.

Another criterion consists of privileging the arguments which use as much as possible information from $\mathcal{GO}$. In fact, an argument which uses elements from $\mathcal{GO}$ has a better chance to be accepted by the other agent.

Definition 8 (Criterion of partial size) *Let* (H_1, h) *and* $(H_2, h) \in \mathcal{A}(\mathcal{K} \cup \mathcal{G} \cup \mathcal{GO})$. (H_1, h) *Pref* (H_2, h) *iff* $Size_{GO}(H_1) > Size_{GO}(H_2)$.

Note that the two above criteria may lead to contradictory results as shown by the following example:

Example 3 *Let* $\mathcal{K} = \{a \rightarrow b, d \rightarrow b\}$, $\mathcal{GO} = \{a, c, c \rightarrow d\}$ *and* $\mathcal{G} = \emptyset$. *From* $\mathcal{K}$ *and* $\mathcal{GO}$, *two arguments in favour of* b *can be constructed:* A : $(\{a, a \rightarrow b\}, b)$ *and B:* $(\{c, c \rightarrow d, c \rightarrow d\}, b)$. *According to the total size criterion, A Pref B since* $Size(\{a, a \rightarrow b\}) = 2$ *whereas* $Size(\{c, c \rightarrow d, c \rightarrow d\}) = 3$. *However, according to partial size criterion, we have B Pref A since* $Size_{GO}(\{a, a \rightarrow b\}) = 1$ *whereas* $Size_{GO}(\{c, c \rightarrow d, c \rightarrow d\}) = 2$.

The agents will thus use the criterion of partial size and if this last returns the same preference for two arguments, then the second criterion will be used to refine the decision.

Offer selection: Selection of offers is the main decision making process that directs the progress of deliberative negotiation and influences its outcomes. It involves search for prospective solutions from the individual areas of interest that move the parties towards an agreement within the common area of interest.

In (Dubois *et al.* 1999b) two qualitative criteria expressed in the setting of possibility theory, an optimistic one and a pessimistic one, are proposed to compute optimal decisions.

Definition 9 (Pessimistic criterion) $N_{\mathcal{K}^x}(\mathcal{G}) = min_\omega \ max(\pi_{\mathcal{G}}(\omega), 1 - \pi_{\mathcal{K}^x}(\omega))$.

An agent is supposed to be able to evaluate to what extent it is certain that its set of prioritized goals is satisfied (on the basis of its beliefs about the state of the world) and assuming that an offer x takes place. This is evaluated in the possibilistic setting by the inclusion degree of the fuzzy set of plausible states of the world into the set of worlds which satisfy the goals at a high degree. $N_{\mathcal{K}^x}(\mathcal{G}) = 1$ iff $\nexists\, \omega$ such that $\pi_{\mathcal{K}^x}(\omega) > 0$ and $\pi_{\mathcal{G}}(\omega) < 1$. $N_{\mathcal{K}^x}(\mathcal{G}) > 0$ iff $\nexists\, \omega$ such that $\pi_{\mathcal{K}^x}(\omega) = 1$ and $\pi_{\mathcal{G}}(\omega) = 0$. Thus, the pessimistic criterion is all the greater as there exists no world with high plausibility and poor satisfaction degree.

Definition 10 (Optimistic criterion) $\Pi_{\mathcal{K}^x}(\mathcal{G}) = max_\omega \ min(\pi_{\mathcal{G}}(\omega), \pi_{\mathcal{K}^x}(\omega))$.

This criterion only checks the *existence* of a world where the goals are highly satisfied and which is highly plausible according to agent's beliefs, when an offer x takes place. However, it leaves room for other worlds which are both highly plausible and very unsatisfactory.

The above indices (here w.r.t the offer x to make) have been axiomatically justified in the framework of qualitative decision making (Dubois *et al.* 1999b; Dubois, Prade, & Sabbadin 2001). The choice of an x based on the above indices is here given in semantic terms (i.e. in terms of possibility distributions). It can be also directly obtained from the syntactic possibilistic logic bases $\mathcal{K}$ and $\mathcal{G}$ (Dubois *et al.* 1999a).
An agent can also compute to what extent it is possible that an offer x is acceptable for the other agent, according to its belief by: $\Pi_{\mathcal{K}^x}(\mathcal{GO}) = max_\omega \ min(\pi_{\mathcal{GO}}(\omega), \pi_{\mathcal{K}^x}(\omega))$.
The way the above decision criteria are combined for selecting an offer depends on the agent's attitude.

Definition 11 *An agent may have two "offer" attitudes:*

- *A* cooperative *agent takes into account the preferences of the other agent when it suggests a new offer. In this case, among the elements of* X, *an agent selects* x *that maximises* $\mu(x) = min(N_{\mathcal{K}^x}(\mathcal{G}), \Pi_{\mathcal{K}^x}(\mathcal{G}O))$.
- *A* non cooperative *selects an* x *that maximises* $\mu(x) = N_{\mathcal{K}^x}(\mathcal{G})$, *ignoring the known preferences of the other agent.*

Note that we would substitute $N_{\mathcal{K}^x}(\mathcal{G})$ *by* $\Pi_{\mathcal{K}^x}(\mathcal{G})$ *in the case of an optimistic agent. So we will get several kinds of agents:* pessimist and cooperative, pessimist and non cooperative, optimist and cooperative, optimist and non cooperative.

When agents are cooperative, it is generally easier to find a compromise or to reach an agreement. However, that compromise may not be an optimal solution for both agents. This case arises when an agent is misled by what it believes the goals of the other agent are, as illustrated by the following example.

Example 4 *Let* a *and* $\overline{a}$ *be two agents negotiating about an object* i *(Note that* i *may be a price, a destination, etc...). Let's suppose that the set* X *contains only two offers* x_1 *and* x_2 *(*$X = \{x_1, x_2\}$*). For agent* a, x_1 *is as good as* x_2*. Moreover, it believes that agent* $\overline{a}$ *prefers* x_1 *over* x_2*. Agent* $\overline{a}$ *prefers* x_2 *over* x_1 *and it believes that* x_2 *is inacceptable for the agent* a*. If agent* a *starts by making the offer* x_1 *then the agent* $\overline{a}$ *will accept it. However, it is obvious that* x_2 *is the best solution for the two agents.*

Decision criteria at level 2
At each step of a deliberative negotiation, an agent decides which are the moves (among those allowed by the protocol) that can be played. For example, an agent can make an offer if it can find an offer x which satisfies one of the above criteria. Similarly, an agent can make an argue move if it has an *acceptable* argument. An agent accepts only the offers which satisfy its goals.

Definition 12 (Agent possibility distribution) *1. If $\forall$ x $\mu(x)$ = 0 or $X = \emptyset$, then $\pi_{\mathcal{D}}(Withdraw)$ = 1.*
2. If $\exists S \in \mathcal{A}(\mathcal{K} \cup \mathcal{GO})$ such that S is acceptable *then $\pi_{\mathcal{D}}(Argue(S))$ = 1.*
3. If S is an acceptable argument for the agent, $\pi_{\mathcal{D}}(Accept(S))$ = 1.
4. If $\exists$ x which maximises the chosen criterion $\mu(x)$ then $\pi_{\mathcal{D}}(Offer(x))$ = 1.
5. If x maximises $N_{\mathcal{K}^x}(\mathcal{G})$ among the elements of X then $\pi_{\mathcal{D}}(Accept(x))$ = 1.

Decision criteria at level 3 Once the agent has defined the list of the possible moves at level 2, it should select among those having $\pi_{\mathcal{M}}(m_i)$ = min($\pi_{\mathcal{P}}(m_i)$, $\pi_{\mathcal{D}}(m_i)$) = 1 the best one to play and this is a strategic matter. Let's take the example of two agents a and $\overline{a}$ who negotiate about a destination for their next holidays. We can imagine that the agent a proposes a country which is not acceptable for the other agent. Let's suppose that at the issue of level 2, the agent $\overline{a}$ has two possibilities: either to propose a new offer or to challenge the current offer made by a. If $\overline{a}$ is cooperative he will prefer to make a challenge rather than suggesting a new offer.

The revision rules

An agent may be led to revise its beliefs about the world ($\mathcal{K}$), or its goals ($\mathcal{G}$) or in case of a cooperative agent, it may also revise $\mathcal{GO}$. There are two situations where a revision may take place: when an agent receives an offer or when it receives an argument.

Concerning belief revision, the procedure that we present here is based on the possibilistic revision framework presented in (Dubois & Prade 1997). When an agent a receives an *acceptable* argument <H, h> (h is a belief or an offer) then it will revise its base $\mathcal{K}_a$ into a new base $(\mathcal{K}_a)^*(H)$: H is forced to hold after revision in the base $\mathcal{K}_a$. More precisely:

Definition 13 *$\pi_{(\mathcal{K}_a)^*(H)}(\omega)$ = 1 if $\omega \in \{\omega \models H$ and $\pi_{\mathcal{K}a}(\omega)$ is maximal over $H\}$ and $\pi_{(\mathcal{K}_a)^*(H)}(\omega) = \pi_{\mathcal{K}a}(\omega)$ otherwise.*

See (Dubois & Prade 1997) for the syntactic counter-part of this definition.

If the received argument is *against* a goal g of the agent, then the agent will give up that goal by putting its degree of priority to 0.

Definition 14 *If $(S, \neg g) \in \mathcal{A}(\mathcal{K}_a \cup \mathcal{G}_a \cup \mathcal{GO}_a \cup S)$ and (g, α) $\in \mathcal{G}_a$ and $(S, \neg g)$ is* acceptable*, then set α = 0.*

As said before, cooperative agents may also revise what they believe about the goals pursued by the other agents. This case occurs mainly when an agent receives an offer from another agent. It seems natural to assume that the set of worlds that are assumed to be somewhat fully satisfactory for the other agent have a non-empty intersection with the set of worlds with the highest plausibility according to what the agent knows. Letting $Core(\mathcal{K}_a) = \{\omega$ st. $\pi_{\mathcal{K}_a}(\omega) = 1\}$, this consistency condition writes $max_{\omega \in Core(\mathcal{K}_a)} \pi_{\mathcal{GO}_a}(\omega)$ = 1. If this condition no longer holds for $\pi^x_{\mathcal{K}a}$, a revision similar to the one given in definition 13 can take place, where the models of $Core(\mathcal{K}_a)$ which maximise $\pi_{\mathcal{GO}_a}$ are put to 1 in the revised $\pi'_{\mathcal{GO}_a}$.

The negotiation protocol

We consider only negotiation between two agents a and $\overline{a}$. However, this work may be extended to several agents. The protocol used supposes the following conditions:

1. an agent cannot address a move to itself.
2. the two agents take turns.
3. agent a begins the negotiation by making an *offer*.
4. an agent is not allowed to provide a move which has been already provided by another agent or by itself. This guarantees non circular dialogues.
5. Any rejected offer is removed from the set of possible offers X.

In the following, we give for each move, the next legal moves.

Offer(x) : $\pi_{\mathcal{P}}(Accept) = \pi_{\mathcal{P}}(Refuse) = \pi_{\mathcal{P}}(Challenge) = \pi_{\mathcal{P}}(Offer)$ = 1 and $\pi_{\mathcal{P}}(Argue) = \pi_{\mathcal{P}}(Withdraw)$ = 0.

Argue(S) : $\pi_{\mathcal{P}}(Accept) = \pi_{\mathcal{P}}(Offer) = \pi_{\mathcal{P}}(Challenge) = \pi_{\mathcal{P}}(Argue) = 1$ and $\pi_{\mathcal{P}}(Refuse) = \pi_{\mathcal{P}}(Withdraw) = 0$.

Accept(x) : $\pi_{\mathcal{P}}(Refuse)$ = 0 and $\forall\, m_i \neq Refuse, \pi_{\mathcal{P}}(m_i)$ = 1.

Accept(S) : $\pi_{\mathcal{P}}(Refuse)$ = 0 and $\forall$ $m_i \neq Refuse$, $\pi_{\mathcal{P}}(m_i)$ = 1.

Refuse(x) : $\pi_{\mathcal{P}}(Refuse) = 0$ and $\forall m_i \neq Refuse, \pi_{\mathcal{P}}(m_i) = 1$.

Withdraw : $\forall m_i \in \mathcal{M}, \pi_{\mathcal{P}}(m_i) = 0$.

Challenge(x) : $\pi_{\mathcal{P}}(Argue) = 1$ and $\forall$ $m_i \neq Argue$, $\pi_{\mathcal{P}}(m_i)$ = 0.

Property 1 *Any negotiation dialogue between two agents a and $\overline{a}$ will terminate. Moreover, termination takes place with either an* accept(x) *move or a* withdraw *move.*

Property 2 *A* compromise *x found by the agents a and $\overline{a}$ maximizes $min(N_{\mathcal{K}^x_a}(\mathcal{G}_a), N_{\mathcal{K}^x_{\overline{a}}}(\mathcal{G}_{\overline{a}}))$, provided that agents do not misrepresent the preferences of the other in $\mathcal{GO}$.*

Illustrative example

As an illustration of the approach, we consider the example of Peter and Mary who discuss about the place of their next holidays.

Peter's goals are a place which is cheap and preferably sunny. This can be encoded by a possibilistic base like: $\mathcal{G}_{Peter}$ = $\{(Cheap(x), 1), (Sunny(x), \alpha)\}$ with $0 < \alpha < 1$.
In terms of the associated possibility distribution, this means that any place which is not cheap is impossible for Peter, any cheap but not sunny place is possible only at a level 1 - α, and any cheap and sunny place is is fully satisfactory for him.
Peter's beliefs are that Tunisia is certainly cheap and that Italy is likely to be not cheap. This is encoded by the following base: $\mathcal{K}_{Peter}$ = $\{(Cheap(Tunisia), 1), (\neg Cheap(Italy), \beta)\}$.

Mary definitly wants a sunny place, hopefully cheap and preferably not too warm. This can be encoded by a base like: $\mathcal{G}_{Mary}$ = $\{(Sunny(x), 1), (Cheap(x), \epsilon), (\neg Too - warm(x), \delta)\}$ such that $\epsilon > \delta$. Mary's beliefs are that Tunisia is sunny, that Italy is sunny, cheap and not too warm. Her belief base is as follows: $\mathcal{K}_{Mary}$ = $\{(Sunny(Tunisia), 1), (Sunny(Italy), 1), (Cheap(Italy), 1), (\neg Too - warm(Italy), 1)\}$.

For the sake of simplicity, we suppose that both Peter and Mary know nothing about the preferences of the other. This means that the possibility distributions $\pi_{\mathcal{GO}_{Peter}}$ = $\pi_{\mathcal{GO}_{Mary}}$ = 1 for any interpretation. We suppose also that they are both *pessimist* and that X = $\{Tunisia, Italy\}$.

Assume that Mary makes the first offer, say Italy ($Offer(Italy)$), because Italy maximises $\mu_(x)$ = $N_{\mathcal{K}^x}(\mathcal{G})$. Indeed, $N_{\mathcal{K}^{Italy}_{Mary}}(\mathcal{G}_{Mary})$ = 1 with $\mathcal{K}^{Italy}_{Mary}$ = $\{(Sunny(Italy), 1), (\neg Too - warm(Italy), 1), (Cheap(Italy), 1)\}$ and $N_{\mathcal{K}^{Tunisia}_{Mary}}(\mathcal{G}_{Mary}) = 1 - \epsilon$ with $\mathcal{K}^{Tunisia}_{Mary}$ = $\{(Sunny(Tunisia), 1)\}$. However, Peter finds the following values: $N_{\mathcal{K}^{Italy}_{Peter}}(\mathcal{G}_{Peter}) = 0$ with $\mathcal{K}^{Italy}_{Peter}$ = $\{(\neg Cheap(Italy), \beta)\}$ and $N_{\mathcal{K}^{Tunisia}_{Peter}}(\mathcal{G}_{Peter}) = 1 - \alpha$ with $\mathcal{K}^{Tunisia}_{Peter}$ = $\{(Cheap(Tunisia), 1)\}$.
Peter cannot then accept the offer of Mary because he has a better offer. Peter chooses to challenge the offer of Mary and obliges then Mary to justify her offer by an argument. The argument of Mary is the following one: she believes that Italy is sunny, not too-warm and cheap. Peter revises his belief base by integrating Mary's argument except the information $Cheap(Italy)$. The new base of Peter is:

$$\mathcal{K}_{Peter} = \{(Cheap(Tunisia), 1), (\neg Cheap(Italy), \beta), (Sunny(Italy), 1), (\neg Too - warm(Italy), 1)\}.$$

Peter presents a counter-argument which says that he is sure that Italy is not cheap and Tunisia is cheap. Mary revises her belief base and accepts the counter-argument of Peter. Her new base is the following one:

$$\mathcal{K}_{Mary} = \{(Sunny(Tunisia), 1), (Sunny(Italy), 1), (\neg Too - warm(Italy), 1), (Cheap(Tunisia), 1), (\neg Cheap(Italy), \beta)\}.$$

Now, Peter decides to make a new offer which is Tunisia. Mary will accept it because Italy est less satisfactory for her than Tunisia. Indeed, $N_{\mathcal{K}^{Italy}_{Mary}}(\mathcal{G}_{Mary}) = 1 - \epsilon$ and $N_{\mathcal{K}^{Tunisia}_{Mary}}(\mathcal{G}_{Mary}) = 1 - \delta$. Since $\epsilon > \delta$, then $N_{\mathcal{K}^{Tunisia}_{Mary}}(\mathcal{G}_{Mary}) > N_{\mathcal{K}^{Italy}_{Mary}}(\mathcal{G}_{Mary})$. The dialogue stops since Peter and Mary have found Tunisia as a compromise.

Related works

Works in multi-agents negotiation can be roughly divided into two categories. The first one has mainly focused on the numerical computation of trade-offs in terms of utilities, and the search for concessions which still preserve the possibility of reaching preferred states of affairs e.g.(Luo *et al.* 2003; Wong & Lau 2000). This type of approaches often uses heuristic strategies and does not incorporate mechanisms for modeling persuasion processes. Recently, a second line of research (Amgoud, Parsons, & Maudet 2000; Sycara 1990; Sierra *et al.* 1997) has focused on the necessity of supporting offers by arguments during a negotiation. Indeed, an offer supported by a good argument has a better chance to be accepted by an agent, and also may lead an agent to revise its goals. These works have mainly introduced a protocol for handling arguments. In (Amgoud, Parsons, & Maudet 2000), a formal model of reasoning shows how arguments are constructed from the knowledge bases of the agents and how these arguments are evaluated. However, these approaches have some limitations. Indeed, the first category of approaches, although effective for finding compromises, does not leave much room for the exchange of arguments and information. They have been centred on the trading of offers and the only feedback that can be made to an offer is another offer. In these approaches, it is hard to change the set of issues under negotiation. On the contrary, argumentation-based approaches allow additional information, over the offers to be exchanged. However, in these approaches, it is not clear how the goals are handled and updated if necessary, how an agent chooses an offer which of course should satisfy its goals and how a compromise can be searched for. Our approach integrates both the merits of argumentation and of heuristic methods looking for making trade-offs. We have shown how the two above approaches may be combined in a unique model. Possibilistic logic is used as a unified setting, which proves to be convenient not only for *representing* the mental states of the agents, but also for *revising* the belief bases and for describing the *decision* procedure for selecting a new offer.

Conclusion

This paper has introduced a negotiation framework based on possibilistic logic. The basic idea is to represent the beliefs, the preferences and the argumentation, decision and revision processes in a unified framework. The negotiation moves consider a subset of those in (Amgoud, Maudet, & Parsons

2000), including additional moves which simplify the handling of negotiation dialogues. Unlike (Amgoud, Maudet, & Parsons 2000), we have shown how these can be operationalized in terms of a possibilistic base. Each move has a degree of possibility to be made. The highest degree (1) means that it is possible for the agent to make that move. However, a degree (0) means that that move cannot be played. The resulting set of moves makes it possible to capture the kind of negotiation exchanges proposed in (Amgoud, Parsons, & Maudet 2000; Sierra *et al.* 1997) as the minimum suitable set for argumentation-based negotiation, and to engage in the kind of negotiations discussed in (Parsons & Jennings 1996). Thus these moves seem adequate for supporting negotiations based on argumentation too. Our approach is not only equal in scope to those in (Parsons & Jennings 1996; Sierra *et al.* 1997) (and indeed other argument-based approaches) but goes somewhat beyond them in directly relating the arguments to the negotiation through the operationalisation of the dialogue moves. As a result the moves are intimately connected to the arguments that an agent makes and receives.

Appendix: Basics of possibilistic logic

Let $\mathcal{L}$ be a propositional language over a finite alphabet $\mathcal{P}$ and Ω be the set of classical interpretations for $\mathcal{L}$. Let ϕ be a formula, $[\phi]$ denotes the set of all its models. $\omega \models \phi$ means that ω is a model of ϕ.

The representation tool used in this paper is the necessity-valued possiblistic logic ((Dubois, Lang, & Prade 1991; 1993)).

At the semantic level, the basic notion in possibilistic logic is the *possibility distribution* denoted by π, which is a function from Ω to $[0, 1]$. $\pi(\omega)$ represents the degree of compatibility of the interpretation ω with the available beliefs about the environment if we represent uncertain knowledge, or a degree of satisfaction if we represent preferences. When modeling beliefs, $\pi(\omega)$ = 1 means that it is completely possible that ω is the real world, $1 > \pi(\omega) > 0$ means that it is only possible that ω is the real world and $\pi(\omega)$ = 0 means that it is certain that ω is not the real world.

From a possibility distribution π, two measures can be defined for a given formula ϕ: a *possibility* degree and a *certainty* (or *necessity*) degree:

- The *possibility* degree of ϕ, $\Pi(\phi) = max\{\pi(\omega) : \omega \in [\phi]\}$, evaluates to what extent ϕ is coherent with the available beliefs encoded by π.
- The *certainty* degree (or necessity degree) $N(\phi)$ = 1 - $\Pi(\neg\phi)$ evaluates to what extent ϕ may be infered from the available beliefs. Thus $N(\phi)$ = inf$\{1 - \pi(\omega), \omega \models \neg\phi\}$.

In other words, the necessity measure of a formula ϕ equals to the complement of the highest *possibility* attached to an interpretation ω which falsifies ϕ.

At the syntactic level, a possibilistic formula is of the form (p, α), where $p \in \mathcal{L}$ is a classical closed formula and α is a certainty or priority level, which may belong to a qualitative scale made of a finite number of grades which are linealy ordered, or may be a numerical degree belonging in the unit interval [0,1]. In the following we use a numerical encoding of the levels. The intuitive meaning of (p, α) is that the truth of p is certain at least to a degree α (i.e. $N(p) \geq \alpha$, where N is a necessity measure reflecting the available information).

The semantics of a set of classical formulae Δ is defined by $M(\Delta) \subset \Omega$ that satisfies all formulae in Δ. Each $\omega \in M(\Delta)$ is called a *model*. For a set $\mathcal{B} = \{(p_i, \alpha_i), i = 1, \ldots, n\}$, the semantics is expressed by a *possibility distribution* π over Ω that characterises the fuzzy set of models $M(\mathcal{B})$. $\mathcal{B}$ induces a preference ordering over Ω via $\pi_{\mathcal{B}}$ as defined in (Dubois, Lang, & Prade 1993):

Definition 15 $\forall\, \omega \in \Omega$,

$$\pi_{\mathcal{B}}(\omega) = \begin{cases} 1 if \omega \models p_i, \forall i \\ 1 - max\{\alpha_i, \omega \not\models p_i\}, otherwise \end{cases}$$

$\pi_{\mathcal{B}}(\omega)$ is all the greater as ω is not a counter-model of a formula p_i having a high weight α_i. It can be checked that $N(p_i) \geq \alpha_i$ where N is necessity measure defined from $\pi_{\mathcal{B}}$. $\Pi_{\mathcal{B}}$ and $N_{\mathcal{B}}$ denote the possibility and the necessity measures associated with a base $\mathcal{B}$.

Acknowledgments

This work was supported by the Commission of the European Communities under contract IST-2004-002307, ASPIC project "Argumentation Service Platform with Integrated Components".

References

Amgoud, L., and Cayrol, C. 1998. On the acceptability of arguments in preference-based argumentation framework. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, 1–7.

Amgoud, L., and Prade, H. 2004. Generation and evaluation of different types of arguments in negotiation. Technical report, Institut de Recherche en Informatique de Toulouse (I.R.I.T.).

Amgoud, L.; Maudet, N.; and Parsons, S. 2000. Modelling dialogues using argumentation. In *Proceedings of the International Conference on Multi-Agent Systems*, 31–38.

Amgoud, L.; Parsons, S.; and Maudet, N. 2000. Arguments, dialogue, and negotiation. In *Proc. of the 14th ECAI*.

Amgoud, L. 1999. *Contribution a l'integration des préferences dans le raisonnement argumentatif*. Ph.D. Dissertation, Université Paul Sabatier, Toulouse.

Dubois, D., and Prade, H. 1997. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. Approx. Reasoning* 17:295–324.

Dubois, D.; Berre, D. L.; Prade, H.; and Sabbadin, R. 1999a. *Using possibilistic logic for modeling qualitative decision: ATMS-based algorithms*, volume 37. Fundamenta Informaticae.

Dubois, D.; Godo, L.; Prade, H.; and Zapico, A. 1999b. *On the possibilistic decision model: from decision under* *uncertainty to case-based decision*, volume 7:6. Int. J. of Uncert., Fuzziness and Knowledge-Based Syst.

Dubois, D.; Lang, J.; and Prade, H. 1991. A brief overview of possibilistic logic. In *Proc. of Symb. and Quanti. Approaches to Uncert., ECSQAU'91. LNCS 548.*, 53–57.

Dubois, D.; Lang, J.; and Prade, H. 1993. Possibilistic logic. *Handbook of Logic in Artificial Intelligence and Logic Programming* 3:439–513.

Dubois, D.; Prade, H.; and Sabbadin, R. 2001. *Decision-theoretic foundations of qualitative possibility theory*, volume 128. European Journal of Operational Research.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77:321–357.

Kraus, S.; Sycara, K.; and Evenchik, A. 1998. *Reaching agreements through argumentation: a logical model and implementation*, volume 104. Journal of Artificial Intelligence.

Luo, X.; Jennings, N.; Shadbolt, N.; fung Leung, H.; and man Lee, J. H. 2003. A fuzzy constraint based model for bilateral, multi-issue negotiations in semi-competitive environments. In *To appear in Artificial Intelligence*.

Parsons, S., and Jennings, N. R. 1996. Negotiation through argumentation—a preliminary report. In *Proceedings of the 2nd International Conference on Multi Agent Systems*, 267–274.

Ramchurn, S. D.; Jennings, N.; and Sierra, C. 2003. Persuasive negotiation for autonomous agents: a rhetorical approach. In *IJCAI Workshop on Computational Models of Natural Arguments*.

Sierra, C.; Jennings, N.; Noriega, P.; and Parsons, S. 1997. A framework for argumentation-based negotiation. In *4th Workshop ATAL*, 167–182.

Sycara, K. 1990. Persuasive argumentation in negotiation. *Theory and Decision* 28:203–242.

Wong, O., and Lau, R. 2000. Possibilistic reasoning for intelligent payment agents. In *Proc. of the 2nd Workshop on AI in Electronic Commerce AIEC*, 170–180.