Argumentation-based Negotiation with Incomplete Opponent Profiles

Yannis Dimopoulos
Department of Computer Science
University of Cyprus
Nicosia, Cyprus
yannis@cs.ucy.ac.cy

Jean-Guy Mailly
LIPADE
Paris Descartes University
Paris, France
jean-guy.mailly@parisdescartes.fr

Pavlos Moraitis
LIPADE
Paris Descartes University
Paris, France
pavlos@mi.parisdescartes.fr

ABSTRACT

Computational argumentation has taken a predominant place in the modeling of negotiation dialogues over the last years. A competent agent participating in a negotiation process is expected to decide its next move taking into account an, often incomplete, model of its opponent. This work provides a complete computational account of argumentation-based negotiation under incomplete opponent profiles. After the agent identifies its best option, in any state of a negotiation, it looks for suitable arguments that support this option in the theory of its opponent. As the knowledge on the opponent is uncertain, the challenge is to find arguments that, ideally, support the selected option despite the uncertainty. We present a negotiation framework based on these ideas, along with experimental evidence that highlights the advantages of our approach.

KEYWORDS

Argumentation; Automated Negotiation; Multi-Agent Systems

ACM Reference Format:
Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. 2019. Argumentation-based Negotiation with Incomplete Opponent Profiles. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 9 pages.

1 INTRODUCTION

During the last years computational argumentation has taken a predominant place in the modeling of negotiation dialogues (for a survey see [11], [24]). The goal of a negotiation dialogue is to allow interacting agents to resolve conflicts and reach a mutually accepted agreement, which in this work is a mutually accepted offer (e.g. the price of a product, the mode of payment).

In an argumentation-based negotiation (ABN), agents choose offers that are likely to be accepted by the opponent and exchange arguments that support these offers, either based on their own theories (see e.g. [1], [3], [18],[13], [22], [14]), or based on the opponent's profile (e.g. [15], [23], [9]).

The modeling of the opponent profile is an important issue in negotiation dialogues (and more generally other types of dialogue such as persuasion). As explained in [5], although there are important differences between opponent models, there are strong reasons justifying their use, such as the *minimization of negotiation cost*, *the adaptation to the opponent* and *the capacity to reach win-win agreements*, especially in cooperative environments. Learning the opponent profile means learning its acceptance and bidding strategies, the deadlines and its preference profile [5]. In most of the proposed works, the (online) opponent modeling is based on learning techniques (see e.g. [4] for a survey). Apart from the fact that learning the opponent profile with traditional learning techniques is not an easy task, as pointed out in [28], those techniques seem better suited to game-theoretic (or utility-based) negotiations, rather than argumentation-based negotiations. Other works (although they concern persuasion dialogues and legal disputes), have proposed a probabilistic approach for dealing with the uncertainty about the opponent profile. In these works (e.g. [16], [27], [17]), probabilities are used in different ways for finding the arguments that are most likely to be accepted by the opponent. Finally, some works (e.g. [26], [21], [8]) investigate other approaches to modeling the opponent profile in argumentation-based dialogues.

This work advances the state of the art in argumentation-based negotiation by making original contributions to the *opponent modeling*, and the associated *acceptance strategy* (i.e. what offers are most likely to be accepted) as well as *bidding strategy* (i.e. the strategy that an agent applies for choosing the next offer). For opponent modeling, it builds on the work of [10] on control argumentation frameworks (CAFs), a formalism for modeling the uncertainty about the *opponent profile.* More specifically, it borrows the concepts of "on/off" arguments (i.e. arguments we don't know whether they are present or not in a theory), and the three different categories of attacks (i.e. attacks we know their existence and direction, attacks we know the existence but not the direction, attacks we don't know the existence but we know the direction). This allows generating different profiles modeled as completions of the known part of the opponent's theory, and seeking offers that satisfy all possible profiles (or as many as possible). Regarding the *bidding* and *acceptance strategies*, the originality of this work lies in the assumption that in argumentation-based negotiation, a central challenge for an agent is to lead, by means of appropriate arguments, its counter party to change its theory, and eventually accept the offer it proposes, hence influencing its *acceptance strategy.* Thus, in our approach, we propose a *bidding strategy* that relies on the previous assumption. More precisely, the idea is that a proponent agent uses first its own theory for choosing the best offer to propose, but next, it uses the incomplete theory of its opponent to find the arguments to support it. Then, it seeks and puts forward a set of arguments called *control configuration*, that could reinstate the supporting arguments, if these are rejected in the current state of the argumentative negotiation theories of all (or most) of the generated opponent profiles. Once the arguments of the control configuration are inserted in the

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

opponent theory, they would, ideally, allow it to reach an agreement with the proponent, thus they alter its *acceptance decision*.

2 BACKGROUND

2.1 Argumentation Systems

We assume that the reader is familiar with abstract argumentation frameworks as introduced in [12], presented as a pair $\langle A, R\rangle$, where A is a set of *arguments*, and $R \subseteq A \times A$ is an *attack relation*. The relation *a attacks b* is denoted by $a\ R\ b$ or $(a, b) \in R$. Different acceptability semantics were also introduced in this work. Based on the acceptability semantics, we can define the status of any argument, namely *skeptically accepted, credulously accepted* and *rejected arguments*.

2.2 Control Argumentation Frameworks

This section introduces briefly the control argumentation frameworks (CAFs) proposed in [10], and discusses how they capture the knowledge of an agent on its opponents. On a high level, a CAF is an argumentation framework where arguments are divided in three parts, *fixed*, *uncertain* and *control*.

The *fixed* part of the theory concerns the certain knowledge that an agent holds about its opponent. This includes arguments as well as attacks that undoubtedly belong to the argumentation theory of the opponent. For instance, a seller agent knows that the customer agent prefers European cars, that safety is an important issue for it and that it prefers electric or gasoline-powered cars than diesel cars. The *uncertain* part captures the uncertainty about the presence of arguments in a theory (expressed by the "on/off" arguments as shown below), as well as the presence and the direction of attacks between arguments in this theory. It reflects the uncertainty that arises due to lack of complete information on the current state of the world that determines the decisions of the opponent, but also its beliefs and preferences. For example, the seller agent may not know the income of the customer agent, whether a car is a social status symbol for it, the highest price that it is ready to pay, or whether it is willing to pay more if some extras are included, and payment by installments is accepted. Finally, the *control* part contains arguments that can be used against arguments of the fixed or uncertain parts that attack arguments that are in favour of some offer of the proponent. Therefore, the control part serves to ensure that arguments in the fixed part that support some offer of the seller that is not adequate with some certain (i.e. European car) or uncertain (e.g. max price, preferred mode of payment) preferences of the customer, can be accepted under some circumstances. For instance, a control argument could allow a seller agent to propose a car from abroad Europe (which is against the known preference of the customer agent and represented in the fixed part) by proposing some interesting options (e.g. five airbags knowing that safety is an important issue for the customer and also represented in the fixed part) and in a price that is probably higher than the highest price the customer is intended to pay (this is part of the uncertain knowledge) but which allows the seller to accept a payment by installments, if this is the preferred payment mode for the customer (this is also part of the uncertain knowledge).

Formally, a CAF is defined as follows:

Definition 2.1. Let $\mathcal{L}$ be a language from which we can build arguments, and let $\texttt{Args}(\mathcal{L})$ be the set which contains all those arguments. A Control Argumentation Framework (*CAF*) is a triple $\mathcal{CAF} = \langle F, C, U\rangle$ where F is the *fixed part*, U is the *uncertain part* and C is the *control part* of $\mathcal{CAF}$ with:

- $F = \langle A_F, \rightarrow\rangle$ where A_F is a set of arguments that we know they belong to the system and $\rightarrow \subseteq (A_F \cup A_U) \times (A_F \cup A_U)$ is an attack relation representing a set of attacks for which we are aware both of their existence and their direction.
- $U = \langle A_U, (\rightleftarrows \cup \dashrightarrow)\rangle$ where A_U is a set of arguments for which we are not sure that they belong to the system, $\rightleftarrows \subseteq (((A_U \cup A_F) \times (A_U \cup A_F)) \setminus \rightarrow)$ is an attack relation representing a set of attacks for which we are aware of their existence but not of their direction, and $\dashrightarrow \subseteq (((A_U \cup A_F) \times (A_U \cup A_F)) \setminus \rightarrow)$ is an attack relation representing a set of attacks for which we are not aware of their existence but we are aware of their direction, with $\rightleftarrows \cap \dashrightarrow = \emptyset$.
- $C = \langle A_C, \Rrightarrow\rangle$ where A_C is a set of arguments, called *control arguments*, that the agent can choose to use or not, and $\Rrightarrow \subseteq \{(a_i, a_j) \mid a_i \in A_C,\ a_j \in A_F \cup A_C \cup A_U\}$ is an attack relation.

A_F, A_U and A_C are disjoint subsets of $\texttt{Args}(\mathcal{L})$.

A CAF features a set of distinct attack relations that capture different sorts of information. Its simplest part is $\langle A_F, \rightarrow \cap (A_F \times A_F)\rangle$, which is a classical AF that contains the indisputable knowledge of the agent on its opponent. The idea of CAFs essentially extends this basic argumentation framework with additional attack relations defined on arguments from the sets A_U and A_C. For instance, there is an attack $(a_i, a_j) \in \rightleftarrows$, with $a_i, a_j \in A_F$ when it is certain that the two arguments exist and are in conflict (*e.g.* because they make mutually exclusive claims), but the direction of the attack(s) is unknown (*e.g* because of lack of information on the intrinsic strength of arguments, or on the preference relation between arguments). An attack $(a_i, a_j) \in \rightarrow$, with $a_i \in A_U$ and $a_j \in A_F$, represents a situation where it is unknown whether a_i is present in the system (*e.g.* some of its premises could be false at the current time), but *if* a_i *is in the system*, then a_i definitely attacks a_j.

Central to controllability is the notion of *completion* of a CAF. Intuitively, a completion is a classical AF which is built from the CAF, by choosing one of the possible options for each uncertain argument or attack.

Definition 2.2. [10] Given a CAF $\mathcal{CAF} = \langle F, C, U\rangle$, a completion of $\mathcal{CAF}$ is an AF $\mathcal{AF} = \langle A, R\rangle$, s.t.

- $A = A_F \cup A_C \cup A_{comp}$ where $A_{comp} \subseteq A_U$;
- if $(a, b) \in R$, then $(a, b) \in \rightarrow \cup \rightleftarrows \cup \dashrightarrow \cup \Rrightarrow$;
- if $(a, b) \in \rightarrow$, then $(a, b) \in R$;
- if $(a, b) \in \rightleftarrows$ and $a, b \in A$, then $(a, b) \in R$ or $(b, a) \in R$;
- if $(a, b) \in \Rrightarrow$ and $a, b \in A$, then $(a, b) \in R$.

Note that the definition of a completion leaves the attacks from $\dashrightarrow$ unspecified, as these attacks may not appear in the theory. For some examples of completions the reader can see [10].

Controllability means that we can select a subset $A_{conf} \subseteq A_C$ and the corresponding attacks $\{(a_i, a_j) \in \Rrightarrow \mid a_i \in A_C,\ a_j \in (A_F \cup A_C \cup A_U)\}$ such that whatever the completion of $\mathcal{CAF}$, a given target is always reached. We focus on two kinds of targets: credulous acceptance of a set of arguments (this is reminiscent of extension enforcement [6]), and skeptical acceptance of a set of arguments.

Definition 2.3. [10] A control configuration of a CAF $\mathcal{CAF} = \langle F, C, U\rangle$ is a subset $A_{conf} \subseteq A_C$. Given a set of arguments $T \subseteq A_F$ and a semantics σ, we say that T is skeptically (resp. credulously) reached by the configuration A_{conf} under σ if T is included in every (resp. at least one) σ-extension of every completion of $\mathcal{CAF}' = \langle F, C', U\rangle$, with $C' = \langle A_{conf}, \{(a_i, a_j) \in \Rrightarrow \mid a_i \in A_C,\ a_j \in (A_F \cup A_C \cup A_U)\}\rangle$. We say that $\mathcal{CAF}$ is skeptically (resp. credulously) *controllable* w.r.t. T and σ.

In a nutshell, CAFs are a powerful enabler of advanced negotiation techniques, that blend together a number of desirable features such as the qualitative representation of uncertainty, simultaneous reasoning with different profiles through completions, simultaneous consideration of both certain and uncertain knowledge of the opponent, the use of control arguments (corresponding to a persuasion phase embedded in negotiation, allowing for the reinstatement of rejected arguments), along with a computational model based on QBFs.

3 THE NEGOTIATION FRAMEWORK

This section presents a new argumentation-based negotiation framework that relies on CAFs [10] for representing the incomplete information that agents have about their opponents. Agents communicate through the exchange of messages (or dialogue moves, see e.g. [11]). We assume that agents play the roles of the proponent and opponent in a turn-taking round-based protocol (e.g. similar to the alternating offers protocol of [14]), where a proponent initiates a round and passes the token to its opponent when it is unable to defend an offer rejected by the opponent. The opponent may accept an offer when one of the supporting arguments is an acceptable argument for it, or reject an offer if it cannot accept any of the different supporting arguments sent by the proponent. We build on the works of [1], [14], and in the following, $\mathcal{L}$ denotes a logical language, and $\equiv$ an equivalence relation associated with it. From $\mathcal{L}$, a set $\mathcal{O} = \{o_1, \ldots, o_n\}$ of n *offers* is identified, such that $\nexists o_i, o_j \in \mathcal{O}$ such that $o_i \equiv o_j$. This means that the offers are different. Offers correspond to the different alternatives (e.g. prices for a product) that can be exchanged during the negotiation dialogue. We assume that agents share the same set of offers $\mathcal{O}$ but those offers can be supported by different practical arguments (although not necessarily) in the theories of the negotiating agents. By argument, we mean a *reason* in believing (called epistemic arguments) or doing something (called practical arguments). The set $Args(\mathcal{L})$ is then divided into two subsets: a subset $Args_p(\mathcal{L})$ of practical arguments supporting offers, and a subset $Args_e(\mathcal{L})$ of epistemic arguments supporting beliefs. Thus, $Args(\mathcal{L}) = Args_p(\mathcal{L}) \cup Args_e(\mathcal{L})$. A negotiation theory is therefore represented as follows:

Definition 3.1 (Negotiating agent theory). Let $\mathcal{O}$ be a set of n offers. A *negotiating theory* of an agent α is a tuple $\mathcal{T} = \langle \mathcal{O}, \mathcal{T}^\alpha, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^\alpha\rangle$ with $\mathcal{T}^\alpha = \langle A^\alpha, \rightarrow_\alpha\rangle$ and $\mathcal{CAF}^{\alpha,\beta} = \langle F^{\alpha,\beta}, U^{\alpha,\beta}, C^{\alpha,\beta}\rangle$ and where:

- $A^\alpha \subseteq Args(\mathcal{L})$ is a set of arguments s.t. $A^\alpha = A^\alpha_p \cup A^\alpha_e$ where A^α_p is a set of practical arguments, A^α_e a set of epistemic arguments, and $A^\alpha_c \subseteq A^\alpha_e$ is the set of control arguments. For the attack relation it holds $\rightarrow_\alpha = \rightarrow_p \cup \rightarrow_e \cup \rightarrow_m$, with $\rightarrow_p \subseteq A^\alpha_p \times A^\alpha_p$, representing an attack relation for practical arguments, $\rightarrow_e \subseteq A^\alpha_e \times A^\alpha_e$ representing an attack relation for epistemic arguments and $\rightarrow_m \subseteq A^\alpha_e \times A^\alpha_p$ representing an attack relation between epistemic and practical arguments i.e. $(a, \delta) \in \rightarrow_m$, if $a \in A^\alpha_e$ and $\delta \in A^\alpha_p$ (see [2], [14]).
- – $F^{\alpha,\beta} = \langle A^{\alpha,\beta}_F, \rightarrow_{\alpha,\beta}\rangle$ with $A^{\alpha,\beta}_F = A^{\alpha,\beta}_{F_e} \cup A^{\alpha,\beta}_{F_p}$, $\rightarrow_{\alpha,\beta} = \rightarrow^{\alpha,\beta}_e \cup \rightarrow^{\alpha,\beta}_p$ and $\langle A^{\alpha,\beta}_{F_e}, \rightarrow^{\alpha,\beta}_e\rangle$ defining the epistemic arguments subpart s.t. $\rightarrow^{\alpha,\beta}_e \subseteq (A^{\alpha,\beta}_{F_e} \cup A^{\alpha,\beta}_{U_e}) \times (A^{\alpha,\beta}_{F_e} \cup A^{\alpha,\beta}_{U_e})$. The above hold also for the practical arguments subpart. It also holds $A^{\alpha,\beta}_U = A^{\alpha,\beta}_{U_e} \cup A^{\alpha,\beta}_{U_p}$.
  – $U^{\alpha,\beta} = \langle A^{\alpha,\beta}_U, \rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta})\rangle$ with $\rightleftarrows_{\alpha,\beta} = \rightleftarrows_e \cup \rightleftarrows_p$, $\dashrightarrow_{\alpha,\beta} = \dashrightarrow_e \cup \dashrightarrow_p$ and $\langle A^{\alpha,\beta}_{U_e}, \rightleftarrows_e \cup \dashrightarrow_e)\rangle$, $\rightleftarrows_e \subseteq ((A^{\alpha,\beta}_{U_e} \cup A^{\alpha,\beta}_{F_e}) \times (A^{\alpha,\beta}_{U_e} \cup A^{\alpha,\beta}_{F_e}))\setminus \rightarrow^{\alpha,\beta}_e)$, $\dashrightarrow_e \subseteq ((A^{\alpha,\beta}_{U_e} \cup A^{\alpha,\beta}_{F_e}) \times (A^{\alpha,\beta}_{U_e} \cup A^{\alpha,\beta}_{F_e}))\setminus \rightarrow^{\alpha,\beta}_e$, defining the epistemic arguments subpart. The same hold for the practical arguments subpart. $\rightleftarrows_e \cap \dashrightarrow_e = \emptyset$.
  – $C^{\alpha,\beta} = \langle A^\alpha_c, \Rrightarrow\rangle$ where $\Rrightarrow \subseteq \{(a_i, a_j) \mid a_i \in A^\alpha_c$ and $a_j \in A^\alpha_c \cup A^{\alpha,\beta}_{F_e} \cup A^{\alpha,\beta}_{U_e}\} \setminus (\rightarrow^{\alpha,\beta}_e \cup \rightleftarrows_e \cup \dashrightarrow_e))$.
- $\mathcal{F}^\alpha : \mathcal{O} \rightarrow 2^{{A_p}^\alpha}$ s.t $\forall i, j$ with $i \neq j$, $\mathcal{F}^\alpha(o_i) \cap \mathcal{F}^\alpha(o_j) = \emptyset$. Let ${A_p}^\alpha_\mathcal{O} = \cup \mathcal{F}^\alpha(o_i)$ with $i = 1, \ldots, n$. This function returns the practical arguments supporting offers in $\mathcal{O}$.

In the following we present the different procedures that implement the new negotiation framework of this paper.

3.1 Best Offers Selection

Algorithm 1 is the procedure invoked by the proponent agent α in order to compute, first, its best offer, based on its own theory, and it is implemented through function $comp\_next\_offer$. This function looks for the best offer supported by an acceptable practical argument by using a ranking on the supporting arguments based on a partial preorder (other methods can be also applied here). Then, based on its $\mathcal{CAF}^{\alpha,\beta}$, it computes the practical arguments that support this offer in its opponent theory and calls a procedure, implemented by algorithm 2, that selects the supporting argument to be sent. If the proponent agent has no (other) offer to propose, the opponent of the agent is informed by a suitable message (i.e. nothing).

Algorithm 1: choose-best-offer($\mathcal{O}, \mathcal{T}^\alpha, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^{\alpha,\beta}(o)$)

```
o ← comp_next_offer(O, T^α);
if o ≠ ∅ then
    F^{α,β}(o) ← compute_sup_arg(o, A^{α,β}_{F_p} ∪ A^{α,β}_{U_p});
    call choose-support-arg(o, F^{α,β}(o), CAF^{α,β});
else
    message(α, β)=nothing;    send(message(α, β));
```

3.2 Supporting Argument Selection

The algorithm described below, selects (through function $choose-arg$, where the choice can be random, as herein, or based on other methods) the argument that the proponent agent α sends to its opponent agent to support its offer. Moreover, another procedure finds the arguments that defend this supporting argument whenever this argument is currently rejected by the opponent. This task is carried out by the procedure implemented by algorithm 3. If there is no other available argument that supports the current offer, the agent abandons this offer and passes the negotiation token to the opponent agent.

Algorithm 2: choose-support-arg($o, \mathcal{F}^{\alpha,\beta}(o), \mathcal{CAF}^{\alpha,\beta}$)

if $\mathcal{F}^{\alpha,\beta}(o) \neq \emptyset$ **then**
  $\theta \leftarrow$ choose-arg($\mathcal{F}^{\alpha,\beta}(o)$);
  call defend-offer(o, θ, $\mathcal{F}^{\alpha,\beta}(o)$, $\mathcal{CAF}^{\alpha,\beta}$)
else
  O=$O - \{o\}$; $A_p^{\alpha} = A_p^{\alpha} - \mathcal{F}^{\alpha}(o)$;
  message(α, β)=$give\_token$; send(message(α, β));

3.3 The Bidding Strategy

The *bidding strategy* of the proponent agent is implemented by algorithm 3. The main task here is to defend the proposed offer by an argument that (as said before) supports the offer in the opponent's theory. Consider for instance a car seller agent who proposes an expensive luxury SUV of a prestigious brand to a customer who, as the agent understands, seems to afford it. The reason (argument) that the seller agent has chosen this particular car is probably the high sales commission that it brings. However, this is not an argument it can use to convince its customer. The pool of appropriate arguments could include the smooth ride, fast acceleration, high top speed, off-road capabilities, safety features, or even the high social status associated with the brand. In fact, the discovery of those arguments takes place inside algorithms 1 and 2. The role of the bidding strategy algorithm is to determine whether such a supporting argument is already acceptable in the opponent's theory, or to search for a *control configuration* that can defend the selected supporting argument under all possible opponent profiles.

Algorithm 3: defend-offer($o, \theta, \mathcal{F}^{\alpha,\beta}(o), \mathcal{CAF}^{\alpha,\beta}$)

1 **if** *θ is credulously accepted in all completions of the theory* $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$
2 **then**
3   offer(α, β) = $\langle o, \theta, \langle \emptyset, \emptyset \rangle \rangle$;
4   $\mathcal{F}^{\alpha,\beta}(o) = \mathcal{F}^{\alpha,\beta}(o) - \{\theta\}$;
5   message(α, β)=offer(α, β); send(message(α, β))
6 **else**
7   $S \leftarrow$ comp_contr_conf($\mathcal{CAF}^{\alpha,\beta}$, θ);
8   **if** $S \neq \emptyset$ **then**
9     $\mathcal{R} = \{(a_i, a_j) | a_i \in S, a_j \in A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}\}$;
    offer(α, β) = $\langle o, \langle \theta, \langle S, \mathcal{R} \rangle \rangle \rangle$;
10   message(α, β)=offer(α, β); send(message(α, β))
11   **else**
12     $\mathcal{F}^{\alpha,\beta}(o) = \mathcal{F}^{\alpha,\beta}(o) - \{\theta\}$;
13   call choose-support-arg(o, $\mathcal{F}^{\alpha,\beta}(o)$, $\mathcal{CAF}^{\alpha,\beta}$);

More precisely, acceptance in the context of incomplete theories is based on the notion of *completion* which represents a possible profile (see definition 2.2). The computation in line 1 of the algorithm relies on reasoning with Quantified Boolean Formulas (QBFs), as described in [10], that is carried out by the `quantom` solver [25]. The credulous controllability wrt the theory $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}$ (i.e. arguments in A_c^{α} are not considered in this case) is computed by using the following *Formula* 1:

$$\forall\{on_{x_i} \mid x_i \in A_U^{\alpha,\beta}\}\forall\{att_{x_i,x_j} \mid (x_i,x_j) \in \dashrightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta}\}$$
$$\exists\{acc_{x_i} \mid x_i \in \mathbf{A}\}[\Phi_{st}^{cr}(\mathcal{CAF}, \theta)$$
$$\vee(\bigvee_{(x_i,x_j)\in\rightleftarrows_{\alpha,\beta}}(\neg att_{a_i,a_j} \wedge \neg att_{a_j,a_i}))]$$

where $\mathbf{A} = A_F^{\alpha,\beta} \cup A_{comp}$ with $A_{comp} \subseteq A_U^{\alpha,\beta}$.

The on_{x_i} variable means that the argument x_i currently belongs to the system; it is used for making the differentiation between the completions where x_i is included and those where it is not. Similarly, att_{x_i,x_j} is true when there is an attack from x_i to x_j. This variable has to be true if (x_i, x_j) is a fixed attack of $\mathcal{CAF}$. Otherwise the truth value of this variable allows to make the distinction between the completions where (x_i, x_j) is included and those where it is not. Finally acc_{x_i} is a propositional variable representing the acceptance status of the argument x_i. The propositional matrix $\Phi_{st}^{cr}(\mathcal{CAF}, \theta)$ of the formula is satisfiable when θ belongs to at least one extension of a completion of $\mathcal{CAF}$ (more details about this part are given later). Straightforwardly, the prefix of the formula corresponds to an enumeration of every completion (by the $\forall$ quantifiers); for every such completion, we have to search for at least one extension (represented by the existentially quantified part) such that θ belongs to it.
Now, in case this computation succeeds, θ is acceptable in all possible opponent profiles (completions), and agent α sends to agent β the offer o, along with θ.

In case θ is not acceptable wrt the above theory, agent α reacts as depicted in lines 7-13 of algorithm 3. First, it uses its *CAF* to seek a *control configuration* S, that defends θ. This is again a problem on QBFs that is solved by a call to `quantom` solver (line 7 of the algorithm). However, this time arguments in A_c^{α} are considered and credulous controllability is computed by using the following *Formula* 2:

$$\exists\{on_{x_i} \mid x_i \in A_c^{\alpha}\}\forall\{on_{x_i} \mid x_i \in A_U^{\alpha,\beta}\}\forall\{att_{x_i,x_j} \mid$$
$$(x_i,x_j) \in \dashrightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta}\}\exists\{acc_{x_i} \mid x_i \in \mathbf{A}\}[\Phi_{st}^{cr}(\mathcal{CAF}, \theta)$$
$$\vee(\bigvee_{(x_i,x_j)\in\rightleftarrows_{\alpha,\beta}}(\neg att_{a_i,a_j} \wedge \neg att_{a_j,a_i}))]$$

where $\mathbf{A} = A_F^{\alpha,\beta} \cup A_c^{\alpha} \cup A_{comp}$ with $A_{comp} \subseteq A_U^{\alpha,\beta}$.

Note that this formula is very similar to the previous one. This time, the existential quantifier over the on_{x_i} variables, for $x_i \in A_c^{\alpha}$, corresponds to the search for one control configuration. So the whole formula corresponds to the definition of credulous controllability: the formula is true if there is a control configuration such that, for every completion, θ belongs to at least one extension.

In both above cases we use the formula $\Phi_{st}^{cr}(\mathcal{CAF}, \theta) = \Phi_{st}(\mathcal{CAF}) \wedge acc_\theta$ which is based on

$$\Phi_{st}(\mathcal{CAF}) = \bigwedge_{x_i \in A_F^{\alpha,\beta}}[acc_{x_i} \Leftrightarrow \bigwedge_{x_j \in \mathbf{A}}(att_{x_j,x_i} \Rightarrow \neg acc_{x_j})] \wedge \bigwedge_{x_i \in A_C^{\alpha,\beta} \cup A_U^{\alpha,\beta}}[acc_{x_i} \Leftrightarrow (on_{x_i} \wedge \bigwedge_{x_j \in \mathbf{A}}(att_{x_j,x_i} \Rightarrow \neg acc_{x_j}))] \wedge \bigwedge_{(x_i,x_j) \in \rightarrow_{\alpha,\beta} \cup \Rrightarrow_{\alpha,\beta}} att_{x_i,x_j} \bigwedge_{(x_i,x_j) \in \rightleftarrows_{\alpha,\beta}} att_{x_i,x_j} \vee att_{x_j,x_i} \bigwedge_{(x_i,x_j) \notin \mathbf{R}} \neg att_{x_i,x_j}$$

where $\mathbf{R} = \rightarrow_{\alpha,\beta} \cup \Rrightarrow_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta}$.

Moreover, in the first case, where the control arguments are not used (in *Formula* 1), $\bigwedge_{x_i \in A_C^{\alpha,\beta} \cup A_U^{\alpha,\beta}}$ becomes $\bigwedge_{x_i \in A_U^{\alpha,\beta}}$.

This formula is a generalization of the encoding of stable semantics defined in [7]. When every *att*-variable and every *on*-variable is assigned a truth value, this assignment corresponds to a completion. Then, the consistent truth assignments of the *acc*-variables correspond to the set of stable extensions of the completion. This means that if $\Phi_{st}(\mathcal{CAF}) \wedge acc_\theta$ is satisfiable, then θ belongs to at least one stable extension of the completion which is represented by the *att* and *on*-variables.

Now if in this second case the call succeeds, agent α sends offer o to agent β, along with the supporting argument θ, the set of arguments S, and the associated attacks R. Otherwise, the agent abandons this argument and picks another from $\mathcal{F}^{\alpha,\beta}(o)$ in order to continue defending o. This is done by function *choose-support-arg*. Recall that our approach looks for sets of arguments that are control configurations, i.e. work for all possible profiles of agent β. However, if there is no such solution, the QBF based techniques of `quantom` [25], can find sets of arguments that work for most of these profiles.

In the following we define an operator $\oplus$ that is used in algorithms 4 and 5.

Definition 3.2. Let A_1, A_2, A_3 be sets. We define $(A_1, A_2) \oplus A_3$ as the pair (A'_1, A'_2) such that $A'_1 = A_1 \setminus (A_1 \cap A_3)$ and $A'_2 = A_2 \cup (A_1 \cap A_3)$.

At the beginning of the negotiation each agent has in its theory (i.e. A^α and A^β respectively) only a part of the possible epistemic arguments (wrt a specific application). That means that some arguments are in A^α and not in A^β (and vice-versa). However, when an agent will use arguments (and the associated attacks) that do not belong to the opponent's theory, the opponent agent will add them (as well as the associated attacks) in its own theory, and it will be able to use them from that point onwards in the negotiation. This situation may take place in the algorithms 4 and 5.

3.4 The Acceptance Strategy

This section discusses Algorithm 4, that implements the *acceptance strategy* of an agent. Upon receiving an offer and its supporting arguments (and the associated attacks) sent by a proponent agent, the algorithm updates the theory as well as the $\mathcal{CAF}$ of the receiving agent by integrating the supporting arguments, the defending arguments (i.e. the control configuration), and the associated attacks into both theories (i.e. the receiving agent own theory and its $\mathcal{CAF}$). Then, the receiver agent either accepts the offer (i.e. if the supporting arguments are acceptable) and informs the proponent accordingly, or sends to the proponent the reasons for rejecting its offer.

Algorithm 4: decide-upon-offer($\mathcal{T}^\alpha$, $\mathcal{CAF}^{\alpha,\beta}$, offer($\beta, \alpha$))

```
⟨o, θ, ⟨S, R⟩⟩=offer(β, α);
if S ≠ ∅ then
    T^α=(A^α ∪ S, →_α ∪ R);
    (A_U^{α,β}, A_F^{α,β}) = (A_U^{α,β}, A_F^{α,β}) ⊕ S;
    (⇢_{α,β}, →_{α,β}) = (⇢_{α,β}, →_{α,β}) ⊕ R;
    (⇄_{α,β}, →_{α,β}) = (⇄_{α,β}, →_{α,β}) ⊕ R
if θ is a credulous conclusion of theory T^α then
    message(α, β)=Accept(o);
    send(message(α, β))
else
    Compute Q ⊆ E where E is an extension of T^α and Q is the set
    of arguments from which θ is reachable in the attack graph;
    Reasons={(p, θ)|(p, θ) ∈→_α and p ∈ Q };
    message(α, β)=Reject(o, θ, ⟨Q, Reasons⟩);
    send(message(α, β));
```

3.5 The Negotiation Protocol

The algorithm 5 described below implements the core procedure that drives the overall negotiation between the two negotiating agents through the necessary updates of their negotiation theories and calls to appropriate functions. The first part of algorithm (lines 1-2) implements the behavior of an agent when it is the proposer of the first offer, whereas the second part (lines 3-24) is concerned with its reaction when it receives an answer from another agent (i.e. the opponent). While the first part is straightforward as it concerns the selection of the best offer to propose, the second part is more involved and breaks down to several subcases. Those cases concern different situations that may arise during a negotiation, such as the rejection of an offer by the opponent, the acceptance of an offer (that terminates the negotiation with an agreement), the situation where the opponent informs that it has no other offer to propose, the situation where the opponent responds that it has no offer to propose too in a received similar message by the (proponent) agent (this ends the negotiation without agreement), the situation where an agent informs that it gives the token, and the situation where an offer is received and the receiver agent has to decide upon its acceptance or rejection. The example below explains how the protocol works.

3.6 A Negotiation Example

In the following we run an example of negotiation for illustrating our framework. Figure 1 presents the agents α and β theories (before (a) and after the negotiation (b)) and their associated CAF respectively. Green arguments (resp. attacks) represent certain arguments (resp. attacks), red arguments (resp. attacks) represent uncertain arguments (resp. attacks) and blue arguments (resp. attacks) represent control arguments (resp. attacks). Thus in the current example we have $A_p^\alpha = \{X\}$ and $A_e^\alpha = \{B, E, K\}$ for agent α and $A_p^\beta = \{Y\}$ and $A_e^\beta = \{B, E, D, F\}$ for agent β. The arguments $\{D, F\}$ are ignored by agent α. We have also the common set of offers $O^\alpha = O^\beta = \{o\}$.

Algorithm 5: Procedure negotiate($\langle O, \mathcal{T}^{\alpha}, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^{\alpha}\rangle$)

1 **if** *agent α proposes first* **then**
2 call choose-best-offer($O, \mathcal{T}^{\alpha}, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^{\alpha,\beta}(o)$);
3 **while** *true* **do**
4 get message(β, α);
5 **switch** *message(β, α)* **do**
6 **case** *Reject($o, \theta, \langle Q, Reasons\rangle$)*
7 $(A_U^{\alpha,\beta}, A_F^{\alpha,\beta}) = (A_U^{\alpha,\beta}, A_F^{\alpha,\beta}) \oplus Q$;
8 $(\dashrightarrow_{\alpha,\beta}, \rightarrow_{\alpha,\beta}) = (\dashrightarrow_{\alpha,\beta}, \rightarrow_{\alpha,\beta}) \oplus Reasons$;
9 $(\rightleftarrows_{\alpha,\beta}, \rightarrow_{\alpha,\beta}) = (\rightleftarrows_{\alpha,\beta}, \rightarrow_{\alpha,\beta}) \oplus Reasons$;
10 call defend-offer($o, \theta, \mathcal{F}^{\alpha,\beta}(o), \mathcal{CAF}^{\alpha,\beta}$);
11 **case** *Accept(o)*
12 End of negotiation with agreement on offer o
13 **case** *nothing*
14 **if** $O \neq \emptyset$ **then**
15 call choose-best-offer($O, \mathcal{T}^{\alpha}, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^{\alpha,\beta}(o)$);
16 **else**
17 answer(α, β)=*nothing_too*;
18 send(answer(α, β))
19 **case** *nothing_too*
20 End of negotiation without agreement
21 **case** *give_token*
22 call choose-best-offer($O, \mathcal{T}^{\alpha}, \mathcal{CAF}^{\alpha,\beta}, \mathcal{F}^{\alpha,\beta}(o)$);
23 **case** *offer(β, α)=$\langle o, \langle \theta, \langle S, \mathcal{R}\rangle\rangle\rangle$*
24 call decide-upon-offer($\mathcal{T}^{\alpha}, \mathcal{CAF}^{\alpha,\beta}$, offer($\beta, \alpha$));
25

$\mathcal{F}^{\alpha}(o) = \{X\}$ and $\mathcal{F}^{\beta}(o) = \{Y\}$ represent the practical arguments supporting offer o in the agents α and β theories respectively. For their CAF we have $\mathcal{F}(o)^{\alpha,\beta} = \{Y\}$ and $\mathcal{F}(o)^{\beta,\alpha} = \{X\}$ respectively. Regarding the uncertainty, for $\mathcal{CAF}^{\alpha,\beta}$ we have $A_{U_e}^{\alpha,\beta}$=$\{B\}$, $\dashrightarrow_{\alpha,\beta}$=$\{(E,Y)\}$ and for $\mathcal{CAF}^{\beta,\alpha}$ we have $A_{U_e}^{\beta,\alpha}$=$\{E\}$, $\dashrightarrow_{\beta,\alpha}$=$\{(B,X)\}$, $\rightleftarrows_{\beta,\alpha}$=$\{(K,E),(E,K)\}$, $\Rrightarrow_{\beta,\alpha}$=$\{(F,E),(D,B)\}$ and control arguments $A_c^{\beta} = \{D,F\}$.

The negotiation starts with agent α as proponent (see Fig. 1 (a)) by invoking algorithm 5. Following line 2 there is a call of algorithm 1. This algorithm computes the next (best) offer (line 1) to propose that is supported by an acceptable argument. In our example there is offer o but the supporting argument X is rejected as it is attacked by arguments B and E that belong into the two stable extensions namely $\{B,K\}$ and $\{B,E\}$. Agent α has no offer to propose to agent β and following line 6 it prepares a $message(\alpha,\beta) = nothing$ and sends it to agent β. Agent β acts now as proponent (see Fig. 1, (a)). By using algorithm 5 (line 13) it checks whether $O^{\beta} \neq \emptyset$ (line 14) which is the case and calls algorithm 1. This algorithm computes (as previously) the next (best) offer (line 1) that is supported by an acceptable argument. In the current situation we have the offer o which is now supported by the acceptable argument Y as it belongs to the (only) stable extension $\{Y,D,F\}$. Then (line 3) it computes the supporting practical arguments in the uncertain theory of agent α namely $\mathcal{F}(o)^{\beta,\alpha} = \{X\}$ by using its CAF. Then (line 4) there is a call of algorithm 2. This algorithm allows to choose a supporting argument (line 2). In our case there is only one the argument X. Then there is a call (line 3) of algorithm 3. This algorithm allows to check firstly (line 1) whether X is credulously accepted in the uncertain theory of agent α without the use of a control configuration (see *Formula* 1). Argument X is attacked by the uncertain argument E (i.e. see attack (E,X)). That means that there is a completion (or profile) where this argument is present in the theory. Moreover the type of uncertain attack between arguments K and E informs us that an attack is indeed present but the direction is unknown. That means that there are two completions (profiles) (among the three possible ones) where we have $\{(K,E),(E,K)\}$ and $\{(E,K)\}$ as possible attacks. In one of these completions argument E defends itself against the attack from K and in the other it attacks K. Therefore in both cases E will be an acceptable argument and X will be rejected (as there is no defence against this attack). Argument X is also attacked by argument B through the uncertain attack (B,X). That means that there is a completion (profile) where this attack is present in the theory and in that case X will also be rejected as B is an acceptable argument and there is no defence for X against the attack (B,X). Therefore X cannot be accepted without the use of a control configuration. By looking at the real theory (green arguments) of agent α we may observe that the profile with the attacks $\{(K,E),(E,K)\}$ is the right one but agent β ignores this information. Then the algorithm tries to check whether it can find (see *Formula* 2) a control configuration S (line 7). As we may observe such a set exists (see line 9) that can defend X no matter the real profile (i.e. for all the completions) of agent α. More precisely we have $S = \{D,F\}$ and $\mathcal{R}$=$\{(F,E),(D,B)\}$ and an offer(β,α)=$\langle o, \langle X, \langle\{D,F\}, \{(F,E),(D,B)\}\rangle\rangle\rangle$ is built. Then, following line 10, a message(β,α)=offer(β,α) is prepared and sent to agent α. Agent α acts as receiver now. By using algorithm 5 (see line 23) it calls algorithm 4 (see line 24). By using algorithm 4 agent α updates its theory and $\mathcal{CAF}$(see lines 3-6), by using $S = \{D,F\}$ and $\mathcal{R}$=$\{(F,E),(D,B)\}$ (see Fig. 1 (b)). Then it checks whether it can accept X (see line 7). As shown in Figure 1 (b), the integration of agent's β control (blue) arguments $\{D,F\}$ (and the associated attacks) in agent's α theory (see green arguments and attacks in Fig. 1 (b)), allows this agent to accept argument X as $\{X,D,F,K\}$ is a stable extension and therefore to accept offer o. Thus, following lines 8-9 it prepares a message(α,β)=accept(o) and sends it to agent β. Agent β acts as receiver by using algorithm 5 (see line 11) and the negotiation ends successfully (line 12) with an agreement on offer o.

4 EXPERIMENTAL EVALUATION

The proposed framework has been implemented by using the JADE (http://jade.tilab.com/) platform and evaluated on negotiations with random argumentation theories.

4.1 Random Theory Generation

The experimental evaluation of the proposed framework is based on a system, implemented in Java, that generates pairs of random negotiation theories and associated CAFs, with different user specified characteristics.

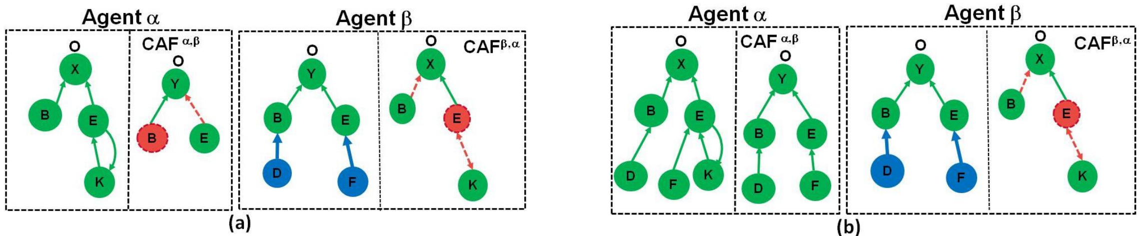

Figure 1: The theories of agents α and β before (a) and after (b) the negotiation and their respective CAFs.

Each negotiation experiment involves a pair of random theories $\mathcal{T}^{\alpha} = \langle A^{\alpha}, \rightarrow_{\alpha} \rangle$ and $\mathcal{T}^{\beta} = \langle A^{\beta}, \rightarrow_{\beta} \rangle$ that share a common part, i.e. there exists $N_{\alpha,\beta} = \langle A^{N_{\alpha,\beta}}, \rightarrow_{N_{\alpha,\beta}} \rangle$, such that $A^{N_{\alpha,\beta}} = A^{\alpha} \cap A^{\beta}$ and $(a,b) \in \rightarrow_{N_{\alpha,\beta}}$ iff $(a,b) \in \rightarrow_{\alpha} \cap \rightarrow_{\beta}$. Moreover, control arguments are only attacked by other control arguments, i.e. $((A^{\alpha} \setminus A^{\alpha}_{c}) \times A^{\alpha}_{c}) \cap \rightarrow_{\alpha} = \emptyset$.
The structure of the generated theories depends on a number of user supplied parameter values that are explained briefly below.

The user inputs the number of epistemic, practical and control arguments of theories $\mathcal{T}^{\alpha}$ and $\mathcal{T}^{\beta}$, as well as their `density`, defined as the ratio of attacks present in the theory to the number of all possible attacks between the arguments of the theory. Moreover, the instance generation system receives as input the number of epistemic, practical and control arguments of the shared part $N_{\alpha,\beta}$.

From theory $\mathcal{T}^{\beta}$, the CAF $\mathcal{CAF}^{\alpha,\beta} = \langle \langle A^{\alpha,\beta}_{F}, \rightarrow_{\alpha,\beta} \rangle, \langle A^{\alpha,\beta}_{U}, \rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta} \rangle, \langle A^{\alpha}_{c}, \Rrightarrow \rangle \rangle$ is built (similarly for $\mathcal{T}^{\alpha}$ and $\mathcal{CAF}^{\beta,\alpha}$), which is the theory that agent α holds about agent β. $\mathcal{CAF}^{\alpha,\beta}$ satisfies the following conditions (a) $A^{\alpha,\beta}_{F} \cup A^{\alpha,\beta}_{U} = A^{\beta} \cup A^{\alpha}_{p}$, (b) $A^{\beta}_{p} \subseteq A^{\alpha,\beta}_{F}$.

The attack relation $\rightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta}$ of $\mathcal{CAF}^{\alpha,\beta}$, is generated so that it satisfies the following conditions:
a) $\rightarrow_{\alpha,\beta} \subseteq \rightarrow_{\beta}$,
b) $\rightarrow_{\beta} \cap (A^{\alpha,\beta}_{F} \times A^{\alpha,\beta}_{F}) \subseteq \rightarrow_{\alpha,\beta}$,
c) $(\rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta}) \subseteq (\rightarrow_{\beta} \setminus \rightarrow_{\alpha,\beta})$, and
d) $\rightleftarrows_{\alpha,\beta} \cap \dashrightarrow_{\alpha,\beta} = \emptyset$.

The main consequence of the above requirements is that the attack relation of $\mathcal{CAF}^{\alpha,\beta}$ is a subset of the attack relation of $\mathcal{T}^{\beta}$. The rationale for this restriction, in this initial experimental evaluation, is to focus on negotiation experiments where agents possess an "accurate" model of their opponent. One way to formalize the model accuracy is via the above relation between individual theories and CAFs. Moreover, it is interesting to study how the framework behaves when this restriction is removed. Indeed, the next section provides initial evidence that the method of this paper can cope with the relaxation of this restriction.

As with the individual agent theories $\mathcal{T}^{\alpha}$ and $\mathcal{T}^{\beta}$, the random instance generation software accepts as input a number of parameter values that determine various features of the CAFs of the agents. Most of them concern the uncertainty of an agent profile on its opponent, as captured by the corresponding CAF. The first is parameter `rateUncertArgs` that defines the ratio of uncertain arguments to all (fixed and uncertain) arguments of the theory. That is, `rateUncertArgs`= $|A^{\alpha,\beta}_{U}|/|A^{\alpha,\beta}_{F} \cup A^{\alpha,\beta}_{U}|$ for agent α, and similarly for agent β.

Other parameters of the system include `rateUncertAtt`, that defines the ratio of uncertain attacks over all attacks, as well as `rateUndirAtt` that defines the ratio of undirected attacks to all attacks. That is, `rateUncertAtt`=$| \dashrightarrow_{\alpha,\beta} |/| \rightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta} |$, and `rateUndirAtt`=$| \rightleftarrows_{\alpha,\beta} |/| \rightarrow_{\alpha,\beta} \cup \rightleftarrows_{\alpha,\beta} \cup \dashrightarrow_{\alpha,\beta} |$.
Moreover, parameter `densContrAtt` defines the ratio of attacks from the control arguments of the agent to the arguments of its opponent that are included in its CAF to all possible such attacks from control arguments. For instance, `densContrAtt=0.1` for $\mathcal{CAF}^{\alpha,\beta}$, means that 10% of all possible attacks from arguments of A^{α}_{c} to arguments in $A^{\alpha,\beta}_{F} \cup A^{\alpha,\beta}_{U}$ are included in the particular $\mathcal{CAF}^{\alpha,\beta}$.
Finally, the instance generation system receives as input the number of offers, i.e. $|O^{\alpha}|$ and $|O^{\beta}|$, as well as the number of practical arguments that support each offer.

4.2 Experimental Results

This section reports on selected results of the experimental evaluation of the framework. As the negotiation theory generation system accepts several parameter values, it is outside the scope of this work to provide exhaustive experimental results for all possible value combinations. Instead, we present results for selected runs that reveal important factors that influence the working of the negotiation algorithm, and highlight its merits and limitations. In all experiments we fix $|A^{\alpha}| = |A^{\beta}| = 40$, $|A^{\alpha}_{p}| = |A^{\beta}_{p}| = 6$, $|O^{\alpha}| = |O^{\beta}| = 4$ and $A^{\alpha}_{c} \cap A^{N_{\alpha,\beta}} = A^{\beta}_{c} \cap A^{N_{\alpha,\beta}} = \emptyset$.

The experimental evaluation is centered around 12 sets of agent theories, and associated CAFs, that differ in the uncertainty of these CAFs and the size of the shared part of agent theories. More specifically, four (4) combinations of parameter values concerning the CAFs are considered, same for both agents. The first combination, abbreviated as `comb1`, is determined by the values `rateUncertArgs`= 0, `rateUndirAtt`= 0, `rateUncertAtt`= 0 which correspond to the case where both agents have complete knowledge of their opponent. Then, `comb2` is defined by the values `rateUncertArgs`= 0.10, `rateUndirAtt`= 0.5, `rateUncertAtt`= 0.5. Moreover, the third combination `comb3` is `rateUncertArgs`= 0.25, `rateUndirAtt`= 0.125, `rateUncertAtt`= 0.125. Finally, the last combination `comb4` is `rateUncertArgs`= 0.50, `rateUndirAtt`= 0.25, `rateUncertAtt`= 0.25 and is the case where the agents have the highest uncertainty about their opponents among all the experiments.
Each of the above set of values for the 3 CAF parameters is combined with one of the three possible values $\{0.25, 0.5, 0.75\}$ for the ratio $|A^{N_{\alpha,\beta}}|/|A^{\alpha}|$ that capture different degrees of similarity between agent theories.

Each row of Tables 1 and 2 presents the *agreement rate*, (i.e. ratio of the number of negotiations terminated with agreement over the total number of negotiations) of 600 negotiations consisting of 50 randomly generated experiments for each of the 12 parameter values combinations described above. Therefore, each experiment is an amalgamation of negotiation theories of various types as far as the values of the 12 value parameters is concerned. Each row of Table 1 corresponds to an experiment (600 negotiations) where the number of control arguments is shown in the `numContrArg` column, whereas the value of parameter `dens ContrAtt` in the corresponding column. The last two columns refer to the agreement rates achieved when the `density` of the individual theories of the agents participating in the negotiations is fixed to 0.15 (column "Agree 0.15") and 0.2 (column "Agree 0.2") respectively. The first row corresponds to the case where none of the agents has any control arguments.

The first conclusion that can be readily drawn from Table 1 is that the presence of control arguments increases significantly the number of negotiations that terminate with agreement. Indeed, for theories with density 0.15 (column "Agree 0.15"), the agreement rate almost doubles from 0.23 to 0.44 for cases where there are relatively few control arguments and attacks from those arguments, and triples to 0.65 in the experiments with the highest number of control arguments and attacks.

Similar are the results when the density of the individual theories of the participating agents is set to 0.2 (column "Agree 0.2"). Observe that the slight increase of the density leads to a decrease in the rate of agreements in all cases. However, again the presence of control arguments increases the agreement rate from 0.16 to as high as 0.56.

`numContrArg`	`densContrAtt`	Agree 0.15	Agree 0.2
0	0	0.23	0.16
3	0.03	0.44	0.39
3	0.05	0.46	0.44
3	0.1	0.57	0.49
3	0.2	0.60	0.52
6	0.03	0.58	0.52
6	0.05	0.59	0.50
6	0.1	0.65	0.56
6	0.2	0.58	0.54

Table 1: Agreement rate for negotiations with individual theories of density 0.15 and 0.20

Recall that the negotiation experiments are generated so that $A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta} = A^\beta \cup A_p^\alpha$ i.e. agent α CAF about β contains all the arguments of its opponents. In the experiments of Table 2 this assumption is removed by allowing agent β to possess arguments that are not part of the CAF of agent α. The number of these arguments is determined by the value of parameter `unknown` defined as $|(A^\beta - (A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta}))|/|(A_F^{\alpha,\beta} \cup A_U^{\alpha,\beta})|$. In the experiments of Table 2 this value is set to 0.25 with the effect of a decrease in the agreement rate when compared to the case with no unknown arguments. This decrease was less significant for theories with more control attacks.

`numContrArg`	`densContrAtt`	Agreement
3	0.03	0.32
3	0.05	0.37
3	0.1	0.42
3	0.2	0.43
6	0.03	0.45
6	0.05	0.43
6	0.1	0.55
6	0.2	0.57

Table 2: Agreement rate for negotiations with individual theories of density 0.15 and `unknown`= 0.25

The experimental evaluation leads to a number of general conclusions. The first is that, not surprisingly, the effectiveness of the approach wrt the rate of agreements depends on a number of parameters including the density of the individual theories, the number of attacks from control arguments, etc. Moreover, other experiments not reported here, have shown that the agreement rate also depends on the size of the shared part $N_{\alpha,\beta}$. In all cases it seems that, for "reasonably good" opponent profiles, the method leads to a significant increase in the number of negotiations that terminate with agreement.

5 RELATED WORK AND CONCLUSIONS

In this paper we presented an original argumentation-based negotiation framework that exploits a recent work proposed in [10] on control argumentation frameworks for modeling the uncertainty about the opponent profile and also the acceptance and bidding strategies of the negotiating agents. Compared to previous works proposed in the literature on argumentation-based negotiation (see e.g. [1],[3], [18],[13], [22],[14], [19],[20]) this new framework introduces and combines together a number of original ideas, with most notable a qualitative representation of uncertainty that enables simultaneous consideration of several different profiles, the bidding strategy that allows an agent to use arguments that do not belong to its theory, along with the notion of control arguments that facilitates persuasion and utilizes arguments that defend against all the possible attacks at once, hence minimizing the number of exchanged messages. We consider that our work generalizes the works proposed in [15],[9]. Our work is also different from the work proposed in [23] where the agents have an incomplete theory on the opponent which evolves based on the information contained in the exchanged offers during the negotiation through classical belief revision. The bidding strategy also used in this work is different to ours. Our experimental results have shown that the outcome of an argumentation-based negotiation dialogue depends on different parameters of the argumentation theories of the agents but in all cases the use of control arguments seems to have a positive impact on the number of agreements.

ACKNOWLEDGMENTS

The authors would like to thank their students Toufik Ider and Mickael Lafages for their excellent work in the implementation of the proposed framework. They would also like to thank the anonymous reviewers for their valuable and constructive comments.

REFERENCES

[1] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. 2007. A unified and general framework for argumentation-based negotiation. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems, (AAMAS07)*. 158. http://doi.acm.org/10.1145/1329125.1329317

[2] Leila Amgoud, Yannis Dimopoulos, and Pavlos Moraitis. 2008. Making Decisions through Preference-Based Argumentation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference, KR 2008, Sydney, Australia, September 16-19, 2008*. 113–123.

[3] Leila Amgoud and Souhila Kaci. 2006. On the Study of Negotiation Strategies. In *Agent Communication II, International Workshops on Agent Communication, AC 2005 and AC 2006, Utrecht, Netherlands, July 25, 2005 and Hakodate, Japan, May 9, 2006, Selected and Revised Papers*. 150–163.

[4] Tim Baarslag, Mark J. C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems* 30, 5 (2016), 849–898. https://doi.org/10.1007/s10458-015-9309-1

[5] Tim Baarslag, Mark J. C. Hendrikx, Koen V. Hindriks, and Catholijn M. Jonker. 2016. A Survey of Opponent Modeling Techniques in Automated Negotiation. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls (Eds.). ACM, 575–576.

[6] Ringo Baumann and Gerhard Brewka. 2010. Expanding Argumentation Frameworks: Enforcing and Monotonicity Results. In *Computational Models of Argument: Proceedings of COMMA 2010, Desenzano del Garda, Italy, September 8-10, 2010*. 75–86.

[7] Philippe Besnard and Sylvie Doutre. 2004. Checking the acceptability of a set of arguments. In *10th International Workshop on Non-Monotonic Reasoning (NMR 2004), Whistler, Canada, June 6-8, 2004, Proceedings*. 59–64.

[8] Elizabeth Black and Katie Atkinson. 2011. Choosing persuasive arguments for action. In *10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), Taipei, Taiwan, May 2-6, 2011, Volume 1-3*. 905–912. http://portal.acm.org/citation.cfm?id=2034398&CFID=69154334&CFTOKEN=45298625

[9] Elise Bonzon, Yannis Dimopoulos, and Pavlos Moraitis. 2012. Knowing each other in argumentation-based negotiation. In *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012, Valencia, Spain, June 4-8, 2012 (3 Volumes)*. 1413–1414.

[10] Yannis Dimopoulos, Jean-Guy Mailly, and Pavlos Moraitis. 2018. Control Argumentation Frameworks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence AAAI 2018, New Orleans,, USA, February 2-7, 2018*. 4678–4685.

[11] Yannis Dimopoulos and Pavlos Moraitis. 2014. Advances in Argumentation-based Negotiation. In *Negotiation and Argumentation in Multi-Agent Systems: Fundamentals, Theories, Systems and Applications*. 82–125.

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

[13] Phan Minh Dung, Phan Minh Thang, and Francesca Toni. 2008. Towards argumentation-based contract negotiation. In *Computational Models of Argument: Proceedings of COMMA 2008, Toulouse, France, May 28-30, 2008*. 134–146.

[14] Nabila Hadidi, Yannis Dimopoulos, and Pavlos Moraitis. 2010. Argumentative alternating offers. In *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada, May 10-14, 2010, Volume 1-3*. 441–448. https://dl.acm.org/citation.cfm?id=1838268

[15] Nabila Hadidi, Yannis Dimopoulos, and Pavlos Moraitis. 2012. Tactics and Concessions for Argumentation-based Negotiation. In *Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012*. 285–296. https://doi.org/10.3233/978-1-61499-111-3-285

[16] Christos Hadjinikolis, Yiannis Siantos, Sanjay Modgil, Elizabeth Black, and Peter McBurney. 2013. Opponent Modelling in Persuasion Dialogues. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. 164–170.

[17] Anthony Hunter. 2015. Modelling the Persuadee in Asymmetric Argumentation Dialogues for Persuasion. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*. 3055–3061.

[18] Antonis C. Kakas and Pavlos Moraitis. 2006. Adaptive agent negotiation via argumentation. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*. 384–391. https://doi.org/10.1145/1160633.1160701

[19] Omar Marey, Jamal Bentahar, Ehsan Khosrowshahi Asl, Mohamed Mbarki, and Rachida Dssouli. 2014. Agents' Uncertainty in Argumentation-based Negotiation: Classification and Implementation. In *Proceedings of the 5th International Conference on Ambient Systems, Networks and Technologies (ANT 2014)*. 61–68. https://doi.org/10.1016/j.procs.2014.05.398

[20] Ariel Monteserin and Analía Amandi. 2013. A reinforcement learning approach to improve the argument selection effectiveness in argumentation-based negotiation. *Expert Syst. Appl.* 40, 6 (2013), 2182–2188. https://doi.org/10.1016/j.eswa.2012.10.045

[21] Nir Oren and Timothy J. Norman. 2009. Arguing Using Opponent Models. In *Argumentation in Multi-Agent Systems, 6th International Workshop, ArgMAS 2009, Budapest, Hungary, May 12, 2009. Revised Selected and Invited Papers*. 160–174. https://doi.org/10.1007/978-3-642-12805-9_10

[22] Simon Parsons, Carles Sierra, and Nicholas R. Jennings. 1998. Agents That Reason and Negotiate by Arguing. *J. Log. Comput.* 8, 3 (1998), 261–292. https://doi.org/10.1093/logcom/8.3.261

[23] Pablo Pilotti, Ana Casali, and Carlos Iván Chesñevar. 2015. A Belief Revision Approach for Argumentation-Based Negotiation Agents. *Applied Mathematics and Computer Science* 25, 3 (2015), 455–470. http://www.degruyter.com/view/j/amcs.2015.25.issue-3/amcs-2015-0034/amcs-2015-0034.xml

[24] Iyad Rahwan, Sarvapali D. Ramchurn, Nicholas R. Jennings, Peter McBurney, Simon Parsons, and Liz Sonenberg. 2003. Argumentation-based negotiation. *Knowledge Eng. Review* 18, 4 (2003), 343–375. https://doi.org/10.1017/S0269888904000098

[25] Sven Reimer, Matthias Sauer, Paolo Marin, and Bernd Becker. 2014. QBF with Soft Variables. In *14th International Workshop on Automated Verification of Critical Systems (AVOCS)*.

[26] Tjitze Rienstra, Matthias Thimm, and Nir Oren. 2013. Opponent Models with Uncertainty for Strategic Argumentation. In *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*. 332–338.

[27] Régis Riveret, Antonino Rotolo, Giovanni Sartor, Henry Prakken, and Bram Roth. 2007. Success chances in argument games: a probabilistic approach to legal disputes. In *JURIX 2007: The Twentieth Annual Conference on Legal Knowledge and Information Systems, Leiden, The Netherlands, 2007*, Arno R. Lodder and Laurens Mommers (Eds.), Vol. 165. IOS Press, 99–108.

[28] Farhad Zafari and Faria Nassiri Mofakham. 2017. POPPONENT: Highly accurate, individually and socially efficient opponent preference model in bilateral multi issue negotiations (Extended Abstract). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*. 5100–5104.