

Towards argumentation-based contract negotiation

Phan Minh DUNG^a, Phan Minh THANG^a, Francesca TONI^{b,1}

^a *Asian Institute of Technology, Bangkok, Thailand*

^b *Department of Computing, Imperial College London, UK*

Abstract. We present an argumentation-based approach to contract negotiation amongst agents. Contracts are simply viewed as abstract transactions of items between a buyer agent and a seller agent, characterised by a number of features. Agents are equipped with beliefs, goals, and preferences. Goals are classified as either *structural* or *contractual*. In order to agree on a contract, agents engage in a two-phase negotiation process: in the first phase, the buyer agent decides on (a selection of) items fulfilling its structural goals and preferences; in the second phase, the buyer agent decides on a subset of the items identified in the first phase fulfilling its contractual goals and preferences. The first phase is supported by argumentation-based decision making taking preferences into account.

Keywords. Negotiation, Decision-making, Contracts

Introduction

We present an argumentation-based approach to contract negotiation amongst agents. Contracts are simply viewed as abstract transactions of items between a buyer agent and a seller agent, characterised by a number of features. Agents are equipped with beliefs, goals, and preferences. Beliefs are represented as an assumption-based argumentation framework. Goals are literals classified as either *structural* or *contractual*, depending on whether they are about structural, static properties of the item the agents aim at agreeing on to form the contract, or whether they are about features subject to negotiation leading to the agreement of a contract. Preferences are given by numerical rankings on goals.

In order to agree on a contract, agents engage in a two-phase negotiation process: in the first phase, the buyer agent decides on (a selection of) items fulfilling its structural goals and preferences; in the second phase, the buyer agent decides on a subset of the items identified in the first phase fulfilling its contractual goals and preferences. The outcome of the second phase is a set of possible contracts between the buyer and the seller. The first phase is supported by argumentation-based decision making with preferences.

We ground our proposed framework upon a concrete “home-buying” scenario, whereby the buyer agent is looking for a property to buy, and the seller has a number of properties to sell. In this scenario, the structural goals are features of a property such as its location, the number of rooms, etc, and the contractual goals are the price of the

¹Corresponding Author: F. Toni, Imperial College London, UK ; E-mail: ft@doc.ic.ac.uk.

property, the completion date for the sale, etc. A contract is simply the agreement on a concrete property and on a number of features fulfilling all “preferred” goals (according to the given preferences).

The paper is structured as follows. In section 1 we give background on assumption-based argumentation, the form of argumentation we adopt to support the agents’ decision-making during negotiation. In section 2 we define the form of contracts we use. In section 3 we define the internal structure and format of the agents in our framework, based upon assumption-based argumentation and preferences on goals. In section 4 we outline a two-phase negotiation process used by the agents to agree on a contract. In section 5 we discuss relationship to related work and conclude.

1. Background

This section provides the basic background on assumption-based argumentation (ABA), see [3,5,6,4] for details.

An ABA framework is a tuple $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ } \rangle$ where

- $(\mathcal{L}, \mathcal{R})$ is a *deductive system*, consisting of a language \mathcal{L} and a set \mathcal{R} of inference rules,
- $\mathcal{A} \subseteq \mathcal{L}$, referred to as the set of *assumptions*,
- $\bar{\ }$ is a (total) mapping from \mathcal{A} into \mathcal{L} , where \bar{x} is referred to as the *contrary* of x .

We will assume that the inference rules in \mathcal{R} have the syntax $l_0 \leftarrow l_1, \dots, l_n$ (for $n \geq 0$) where $l_i \in \mathcal{L}$. We will refer to l_0 and l_1, \dots, l_n as the *head* and the *body* of the rule, respectively. We will represent $l \leftarrow$ simply as l . As in [5], we will restrict attention to *flat* ABA frameworks, such that if $l \in \mathcal{A}$, then there exists no inference rule of the form $l \leftarrow l_1, \dots, l_n \in \mathcal{R}$, for any $n \geq 0$.

An *argument* in favour of a sentence x in \mathcal{L} supported by a set of assumptions X is a (backward) deduction from x to X , obtained by applying backwards the rules in \mathcal{R} .

In order to determine whether a conclusion (set of sentences) should be drawn, a set of assumptions needs to be identified providing an “acceptable” support for the conclusion. Various notions of “acceptable” support can be formalised, using a notion of “attack” amongst sets of assumptions whereby X *attacks* Y iff there is an argument in favour of some \bar{x} supported by (a subset of) X where x is in Y . Then, a set of assumptions is deemed

- *admissible*, iff it does not attack itself and it counter-attacks every set of assumptions attacking it;
- *preferred*, iff it is maximally admissible.

We will refer to a preferred set of assumptions as a *preferred extension* of the given ABA framework. We will use the following terminology:

- a preferred extension of $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ } \rangle \cup \{a\}$, for some $a \in \mathcal{A}$, is a preferred extension of $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \bar{\ } \rangle$ containing a ;
- given a preferred extension E and some $l \in \mathcal{L}$, $E \models l$ stands for “there exists a backward deduction for l from some $E' \subseteq E$ ”;
- given a preferred extension E and some $L \subseteq \mathcal{L}$, $E \models L$ stands for $E \models l$ for all $l \in L$.

2. Contracts

We will assume a set of (at least two) agents *Agents*, a set of *Items*, and a set *Attributes* of attributes for the elements of *Items*. Each attribute is associated with a domain of possible values: for each $a \in \text{Attributes}$, this domain is indicated as $D(a)$.

A *contract* is defined simply as a transaction between (some of the) agents, playing different roles. This transaction is characterised by an item and a number of features, possibly including technical aspects and cost of the item. Concretely, we will assume that these features are assignments of values to attributes, for the given item. For simplicity, in this paper we will restrict attention to contracts between two agents playing the role of buyer and seller. Formally, a contract is a tuple $\langle \text{Buyer}, \text{Seller}, \text{Item}, \text{Features} \rangle$ where

- $\text{Buyer}, \text{Seller} \in \text{Agents}$, $\text{Buyer} \neq \text{Seller}$, representing the buyer and seller in the contract;
- $\text{Item} \in \text{Items}$;
- Features is a set of assignments of values to (some of the) attributes: $\text{Features} = \bigcup_{a \in X} \{a = v_a\}$ for some $X \subseteq \text{Attributes}$ and, for each $a \in X$, some $v_a \in D(a)$ representing the value of attribute a for *Item*.

Given a contract with *Features* for attributes in some X , we will consider attributes not in X as irrelevant, in the sense that their values could be any in their domain without altering the worth of the contract.

Attributes may take any number of values. For Boolean attributes, with domain $\{\text{true}, \text{false}\}$, we will represent assignments of attributes to *true* simply by means of the attributes, and assignments of attributes to *false* simply by means of the negation of the attributes. So, for example, $\{a_1 = \text{true}, a_2 = \text{false}, a_3 = \text{true}\}$ will be represented as $\{a_1, \neg a_2, a_3\}$.

In the remainder of the paper, for simplicity, we will assume that $\text{Agents} = \{\beta, \sigma\}$, with β the buyer and σ the seller in every contract.

As an illustrative example, throughout the paper we will use a “home-buying” scenario whereby two agents, a home buyer and an estate agent, engage in negotiations for the purchase of a property. In this scenario, the given set of agents is $\{h_buyer, e_agent\}$, representing a home-buyer and an estate agents respectively (h_buyer is a concrete instance of β and e_agent is a concrete instance of σ). Also, in this scenario, items are properties for sale, and the attributes include *exchange_date* (representing the date when a non-refundable deposit is paid by the buyer to the seller and contracts are exchanged), *completion_date* (representing the date when the final payment is made and the property changes hands) and *price*. An example contract is

$\langle h_buyer, e_agent, house_1, \{completion_date = 10/03/08, price = \$300K\} \rangle$

indicating that h_buyer is committed to purchasing $house_1$ at the price of \$300K and with agreed completion date for the deal on 10 March 2008.

3. Agents' internals

An agent is characterised by its own goals, preferences over goals, and beliefs. Beliefs and goals are expressed in a given logical language \mathcal{L} consisting of literals (we do not impose any restriction on \mathcal{L} , for example it may not include negation and, if it does, it

may not be closed under negation). This language is shared amongst agents. The literals in \mathcal{L} representing possibly desirable properties for agents are called *goal literals*². For example, a goal literal may be that of having a comfortable house (*comfortable_house*), or a house with at least 3 rooms (*number_of_rooms* ≥ 3), or a house costing less than \$ 450K (*price* \leq \$450K). Goal literals may be of two kinds: those concerning the attributes of the items to be bought (for example location, number of rooms, foundations, building permits etc) and goals concerning the contractual features of such items (for example price, time of completion, deposit etc). Beliefs may be about: the items to be traded, norms and conventions governing the agents' behaviour, and issues agents are uncertain about.

Formally, any agent a is defined as a tuple $\langle G_a, B_a, P_a \rangle$ consisting of

- a *goal-base* $G_a \subseteq \mathcal{L}$ describing the agent's own goals, and consisting of two disjoint subsets: $G_a = G_a^{struct} \cup G_a^{contr}$, where G_a^{struct} are the *structural goals* and G_a^{contr} are the *contractual goals*
- a *belief-base* B_a describing the agent's beliefs
- a *preference-base* P_a describing the agent's preferences over its goals

For simplicity, in this paper we will assume that the seller agent σ only has contractual goals (namely $G_\sigma^{struct} = \{\}$).

We will omit the subscript a from the bases when clear from the context.

The goal-base describes important features of the item the buyer is looking for. The preference-base allows to rank different items according to preferences on their features. The belief-base of both buyer and seller needs to include information about concrete items that can become part of contracts (for example that a given property has 5 rooms), relevant to contractual goals, about norms and conventions used by the agents during the negotiation of the contracts (for example that a seller moving overseas is going to be in a rush to sell, or that a recently-built property with council approval is likely to be safe), and about uncertainties the agents may have during this process (for example that the asking price for a property is too high). Syntactically:

- preferences over goals are expressed by assigning positive integers to them, namely the preference-base is a mapping from the goal-base to the set of natural numbers, ranking the goals so that the higher the number assigned to a goal, the more important the goal
- the belief-base is an ABA framework $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \neg \rangle$ where
 - * $\mathcal{R} = \mathcal{R}_i \cup \mathcal{R}_n \cup \mathcal{R}_c$, where
 - * \mathcal{R}_i represents information about concrete items to be traded
 - * \mathcal{R}_n represents (defeasible) norms
 - * \mathcal{R}_c represents information related to contractual goals
 - * $\mathcal{A} = \mathcal{A}_d \cup \mathcal{A}_c \cup \mathcal{A}_u$ where
 - * \mathcal{A}_d consists of assumptions representing (decisions about) items for transactions, for example *house₁*, *house₂*
 - * \mathcal{A}_c represents "control" assumptions related to defeasible norms (see below)

²We abuse notation here in the sense that something may be desirable for one agent but not for others.

- * \mathcal{A}_u contains assumptions representing the uncertainties about attributes of items to be traded, e.g. whether a given property has a completion certificate

Note that \mathcal{R}_i , \mathcal{R}_c , and \mathcal{A}_d can be obtained directly from information about items to be traded. The rest of the belief base of an agent corresponds to item-independent beliefs held by the agent, and need to be “programmed” into the agent.

Deciding whether or not to start a negotiation about an item in \mathcal{A}_d depends on how this item is evaluated according to the assumptions in \mathcal{A}_u . For example, the lack of a completion certificate is an indication that the house may not be safe. We intend here that the will to (dis)confirm assumptions in \mathcal{A}_u will drive information-seeking steps in the contract negotiation process.

As a simple example of buyer in our running scenario,³ h_buyer may consist of

- $G_{h_buyer}^{struct} = \{own_garden, number_of_rooms \geq 3, safe\}$
 $G_{h_buyer}^{contr} = \{completion_date < 31/05/08, price < \$450K\}$
- $P_{h_buyer}(own_garden) = 2, P_{h_buyer}(number_of_rooms \geq 3) = 2$
 $P_{h_buyer}(completion_date < 31/05/08) = 3$
 $P_{h_buyer}(price < \$450K) = 4$
 $P_{h_buyer}(safe) = 5$
- B_{h_buyer} is $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \neg \rangle$ where
 - * $\mathcal{R} = \mathcal{R}_i \cup \mathcal{R}_n \cup \mathcal{R}_c$ and

$$\mathcal{R}_i = \{ number_of_rooms = 5 \leftarrow house_1, \\ price = \$400K \leftarrow house_2, \\ seller_in_chain \leftarrow house_2 \}$$

$$\mathcal{R}_n = \{ safe \leftarrow council_approval, asm_1, \\ \neg safe \leftarrow weak_foundations, asm_2, \\ council_approval \leftarrow completion_certificate, asm_3, \\ long_time_not_sold \leftarrow price_too_high, asm_4, \\ seller_in_rush \leftarrow seller_moves_overseas, asm_5 \}$$

$$\mathcal{R}_c = \{ long_time_not_sold \leftarrow house_1, \\ seller_moves_overseas \leftarrow house_2, \\ quick_completion \leftarrow seller_moves_overseas, \\ completion_date < now + 60days \leftarrow quick_completion \}$$

- * $\mathcal{A} = \mathcal{A}_d \cup \mathcal{A}_u \cup \mathcal{A}_c$ and
 $\mathcal{A}_d = \{house_1, house_2\}$
 $\mathcal{A}_c = \{asm_1, asm_2, asm_3, asm_4, asm_5\}$
 $\mathcal{A}_u = \{\neg price_too_high, price_too_high, \neg seller_in_rush, , \\ \neg council_approval \}$

³We omit to exemplify the seller for lack of space.

$$\begin{aligned}
* \quad & \overline{house_1} = house_2, \overline{house_2} = house_1, \\
& \overline{asm_1} = \neg safe, \overline{asm_2} = safe, \overline{asm_3} = \neg council_approval, \\
& \overline{asm_4} = short_time_not_sold, \overline{asm_5} = \neg seller_in_rush, \\
& \overline{price_too_high} = \neg price_too_high, \\
& \overline{\neg price_too_high} = price_too_high, \\
& \overline{\neg seller_in_rush} = seller_in_rush, \\
& \overline{\neg council_approval} = council_approval.
\end{aligned}$$

Note that there are different kinds of uncertainties. Some are directly related to the structural properties of items to be traded, like the lack of a council approval. Others are related either to the contractual properties of the items, like *price_too_high*, or to the behaviors of the other agent, like *seller_in_rush*. Note also that assumptions in \mathcal{A}_u are not ordinary assumptions, in that they would not be normally assumed by the agent unless explicit information is obtained. For example, the agent would not ordinarily assume that the seller is in a rush: it will want to check this. Indeed, assumptions in \mathcal{A}_u are meant to be used to start information-seeking dialogues in the negotiation process. This kind of dialogues will be ignored in this paper.

The assumptions in \mathcal{A}_c are used to reflect the defeasible nature of the corresponding rules and their potential to give rise to an inconsistency. These control assumptions can also be useful to resolve conflicts between conflicting information, e.g. originating from different sources of information. For example, the first rule above could be replaced by

$$\begin{aligned}
number_of_rooms = 5 & \leftarrow house_1, asm_{11} \\
number_of_rooms = 4 & \leftarrow house_1, asm_{12}
\end{aligned}$$

with asm_{11} , asm_{12} additional assumptions and $\overline{asm_{11}} = number_of_rooms = 4$ and $\overline{asm_{12}} = number_of_rooms = 5$. This would reflect that the agent has been exposed to two conflicting pieces of information (possibly from two different sources of information), that *house_1* has both 5 and 4 rooms, and would need to decide (by selecting one of asm_{11} or asm_{12} , or neither) which of the two it will make use of.

4. Negotiation process

The decision making of a buyer can be structured into two phases. In the first phase the agent evaluates the items that are available, according to their attributes, to determine whether and how they satisfy its needs. In our running example, the agent would have to decide which houses satisfy its goals about location, safety etc. In the second phase, a negotiation will be conducted with the seller for items (e.g. houses) that have passed the first phase. These phases would benefit from information-seeking, but we ignore this here for simplicity. We focus instead on the decision-making mechanisms needed to support the two phases.

The principle for decision making for both phases is that higher-ranked goals should be pursued at the expense of lower-ranked goals, and thus choices enforcing higher-ranked goals should be preferred to those enforcing lower-ranked goals.

Choices in the first phase are items (e.g. houses) that are available for transactions. Choices in the second phases are possible deals that the buyer and seller could struck (e.g. prices, deposits etc).

Choices are compared based on how they satisfy the goals. A *goal state* (or simply a *state* for short) may be seen as a set of goal literals. Abstractly, this set is intended to be the set of all goals satisfied by a choice. We will see some concrete realisations of this abstract notion in sections 4.2 and 4.3 below. But first, in section 4.1, we will provide a generic way of comparing goal states by taking into account preferences amongst goal literals in them.

4.1. Comparing Choices

As an example of the decision making principle given earlier, consider goals g_0, g_1, g_2 in the goal-base G such that $P(g_0) = P(g_1) = 2$ and $P(g_2) = 1$, where P is the agent's preference-base. Consider states $s_1 = \{g_0, g_1\}$, $s_2 = \{g_0, g_2\}$, $s_3 = \{g_1, g_2\}$. Then, s_1 is preferred to both s_2, s_3 whereas s_2, s_3 are incomparable and thus equally preferred. Formally:

Definition 4.1 Let s, s' be states. We say that s is preferred to s' , denoted by $s \sqsupseteq s'$, iff

1. there exists a goal g that is satisfied in s but not in s' , and
2. for each goal g' , if $P(g') \geq P(g)$ and g' is satisfied in s' then g' is also satisfied in s .

It follows:

Proposition 4.1 The preference relation \sqsupseteq is a partial order.

Often decisions need to be made even though the satisfaction of some goals is undetermined. For example, our buyer may want to buy a home that is situated in a quiet neighbourhood. But one of the properties on offer is located in an area where a project to build a new airport is under consideration by the government, and there are strong arguments for and against the airport. Some other property also under the buyer's consideration does not have council approval, and so may be unsafe. Deciding which of these two properties to buy amounts to comparing the preference over two sets of states. For example, in the case of the first property, a state with the airport being built and a state without the airport being built. This uncertainty is represented by allowing two assumptions *airport* and \neg *airport* in \mathcal{A}_u . Due to the presence of uncertainties, comparing items means comparing sets of states. A possible notion for this comparison is the following:

Definition 4.2 Let S be a nonempty set of states. The *min-state* of S , denoted by $\min(S)$, is a state such that for each goal $g \in G$, g is satisfied in $\min(S)$ iff g is satisfied in each state in S . Let S, S' be sets of goal states. S is said to be *minmax-preferred* to S' if $\min(S)$ is preferred to $\min(S')$.

4.2. First phase: Decision making for the buyer

The items to be chosen (e.g. houses) are represented by assumptions in \mathcal{A}_d . In this setting, goal states are concretely defined as follows:

Definition 4.3 A *structural (goal) state* is a maximal consistent ⁴ set of goal literals from G^{struct} .

Until the next section 4.3, we will refer to structural states simply as states.

Choices determine states as follows:

Definition 4.4 Let s be a goal state, $d \in \mathcal{A}_d$ and $g \in G^{struct}$. We say that

- s is *credulously satisfied* by d if there is a preferred extension E of $B \cup \{d\}$ such that $E \models s$
- g is *skeptically satisfied* by d if, for each preferred extension E of $B \cup \{d\}$, $E \models g$
- s is *skeptically satisfied* by d if for each goal $g \in G^{struct}$, g is skeptically satisfied by d iff $g \in s$

It is not difficult to see that there is exactly one state that is skeptically satisfied by any given d .

For $d \in \mathcal{A}_d$, the *characteristic set of goal states* of d , denoted by CS_d , consists of all goal states s credulously satisfied by d .

Definition 4.5 Given $d_0, d_1 \in \mathcal{A}_d$:

- d_0 is said to be *minmax preferred* to d_1 if CS_{d_0} is minmax preferred to CS_{d_1}
- d_0 is said to be *skeptically preferred* to d_1 if the unique goal state that is skeptically satisfied by d_0 is preferred (with respect to \sqsupseteq) to the unique goal state that is skeptically satisfied by d_1

The following result links our notion of (minmax) preference between states (the characteristic sets given by decisions) and our argumentation-based notion of (skeptical) preference between decisions:

Proposition 4.2 Let $d_0, d_1 \in \mathcal{A}_d$. d_0 is minmax preferred to d_1 iff d_0 is skeptically preferred to d_1 .

Then, decision-making in the first-phase amounts to choosing any “most skeptically preferred” decision. We will refer to these decisions as the *most favored items*.

4.3. Second phase: How Should a Fair Negotiator Proceed ?

In this phase, buyer and seller make decisions by negotiation to agree on a contract. After the first phase, the buyer decides to start negotiation with the seller on (one of) the most favored items. Each agent (the buyer and the seller), ranks these items according to how they satisfy the contractual goals in G^{contr} . Note that the second phase is only concerned with the contractual goals - of both agents (the structural goals of the buyer have all been taken into account in the first phase, and we are assuming that the seller has no structural goals).

⁴A goal state is inconsistent if, for some atom g , it contains both g and its negation $\neg g$; a goal state is consistent if it is not inconsistent. Note that we do not impose that \mathcal{L} is closed under negation, and in particular \mathcal{L} could be a set of atoms. In this special case, any set of goal atoms would be a goal state.

Definition 4.6 A *contractual (goal) state* is a maximal consistent set of goal literals from G^{contr} .

In our home-buying example, a contractual state consists of a price, a deposit, time for completion and several add-ons items like washing-machines, curtains etc. We assume that the set of contractual states is finite and is known to both buyer and seller.

The preference of an agent α (which may be the buyer β or the seller σ) between contractual states can be represented as a total preorder \preceq_α ⁵, where, given contractual states t and t' , $t \preceq_\alpha t'$ states that t is preferred to t' (for α). As \preceq_α , we can choose any pre-order consistent with the partial order \preceq obtained as in definition 4.1 for goals and preferences of α .

For simplicity, we assume that both buyer and seller know each other's preferences between contractual states. We also assume that each agent α possesses an *evaluation function* λ_α that assigns to each structural state s a contractual state $\lambda_\alpha(s)$ representing the "value" of s , such that if s is preferred to s' (as in definition 4.1) then $\lambda_\alpha(s)$ is preferred to $\lambda_\alpha(s')$.

For the buyer agent β , $\lambda_\beta(s)$ represents the "reservation" value of s , i.e. the maximal offers the buyer could make (for the features affecting the contractual goals, that will determine the contract). For the seller agent σ , $\lambda_\sigma(s)$ represents the "reservation" value of s , i.e. the minimal offers the sellers could accept (for the features affecting the contractual goals, that will determine the contract).

From now on, we assume that the agents are negotiating about one of the most favored items characterized by structural state s . The possible deals (contracts) between the buyer and the seller are characterized respectively by the sets

- $PD_\beta = \{t \mid t \text{ is a contractual state and } t \preceq_\beta \lambda_\beta(s)\}$,
- $PD_\sigma = \{t \mid t \text{ is a contractual state and } t \preceq_\sigma \lambda_\sigma(s)\}$.

If $PD_\beta \cap PD_\sigma \neq \emptyset$ then a deal is possible. We define the *negotiation set* as $NS = PD_\beta \cap PD_\sigma$. We assume that the agents are rational in the sense that they would not accept a deal that is not Pareto-optimal, defined below:

Definition 4.7 Let t, t' be contractual states. We say that:

- t is *strictly preferred to* t' for agent α if $t \preceq_\alpha t'$ and $t' \not\preceq_\alpha t$
- t *dominates* t' if t is preferred to t' for both seller and buyer (i.e. $t \preceq_\beta t'$ and $t \preceq_\sigma t'$) and, for at least one of these agents, t is strictly preferred to t'
- t is *Pareto-optimal* if it is not dominated by any other contractual state

The agents bargain by successively putting forward offers. A negotiation is defined as a sequence of alternating offers and counter-offers from the negotiation set NS between the buyer and the seller. Offers and counter-offers are represented by contractual states. An agent could accept an offer or reject it and then make a counter-offer. We assume that our agents are honest and do not go back on their offers. This implies that when an agent makes a new offer, it should be at least as preferred to its opponent as the one it has made previously.

⁵A total preorder \preceq on a set T is a reflexive and transitive relation such that for each two elements t, t' from T , either $t \preceq t'$ or $t' \preceq t$.

Reciprocity is a key principle in negotiation. There is no meaningful negotiation without reciprocity. An agent is said to adhere to the *principle of reciprocity* if, whenever the other agent has made a concession, it will reciprocate by conceding as well. We say that an agent *concedes* if its new offer is strictly preferred to its opponent than the one it made previously. Otherwise the agent is said to *stand still*. Agents do not have unlimited time for negotiation. Hence practical agents will terminate a negotiation when they see no prospect for a successful conclusion for it. This happens when both agents refuse to concede/reciprocate.

Offers and counter-offers may be seen as steps in a negotiation. A negotiation *terminates in failure* at step n if both agents stand still at steps $n, n - 1, n - 2$ ⁶. It is understood that a failure is worse than any agreement for both agents.

A negotiation *terminates successfully* when one of the agents accepts an offer. An agent α accepts an offer from the other agent if it is preferred to α to the one proposed by α itself before. For example, consider the following negotiation between β and σ :

- Step 1: β puts forward an offer of 10.
Here, 10 can be seen as the price that β is prepared to pay for the item at stake, characterised by the contractual state s .
- Step 2: σ makes a counter-offer of 12.
Here, 12 can be seen as the price that σ is prepared to accept as a payment for the item.
- Step 3: β makes a further counter-offer of 11.
Namely, β increases the amount it is willing to pay for the item, in other words, it concedes.
- Step 4: σ makes a counter-offer of 11.
Namely, σ decreases the amount it is willing to accept for payment for the item, in other words, it concedes.
- Step 5: β accepts the offer.
Indeed, the offer by σ at step 4 is preferred to β (by being =) to the offer it made at step 3.

An agent α is said to be *fair* if it adheres to the principle of reciprocity. Formally, this means that whenever α has to move at step n , it will concede or accept if the number of concessions made by the other agent $\bar{\alpha}$ up to step $n - 1$ is more than the number of concessions made by α up to step $n - 1$. Note that for ease of reference, we refer to the opponent of α as $\bar{\alpha}$.

Due to the finiteness assumption of the set of contractual states, the negotiation set is also finite. Hence it is immediate that

Theorem 4.1 Every negotiation terminates.

A strategy is defined as a mapping assigning to each history of negotiation an offer. We are now interested in strategies for fair agents that ensure an efficient and stable outcome in the sense of a Nash equilibrium.

⁶This means that when an agent stands still and in the next move its opponent also stands still then the first agent has to concede if it does not want to terminate the negotiation in failure.

Definition 4.8 A contractual state t' is said to be a *minimal concession of agent α* wrt t , if t' is strictly preferred to t for $\bar{\alpha}$ and for each contractual state r , if r is strictly preferred to t for $\bar{\alpha}$ then r is preferred to t' for $\bar{\alpha}$.

An agent *concedes minimally* at step i if it offers at step i a contractual state t that is a minimal concession wrt the offer the agent made at step $i - 2$. The *minimal concession strategy* calls for agents

1. to start the bargain with their best state and
2. to concede minimally if the opponent has conceded in the previous step or it is making a move in the third step of the negotiation, and
3. to stand still if the opponent stands still in previous step.

Note that the third step in the negotiation has a special status, in that if no concession is made at that step the negotiation stops.

It is obvious that the minimal concession strategy adheres to the reciprocity principle. Hence the minimal concession strategy is permissible for fair agents.

It is not difficult to see

Proposition 4.3 If both agents use the minimal concession strategy then they terminate successfully.

A strategy is said to be in *symmetric Nash equilibrium* if under the assumption that one agent uses this strategy the other agent can not do better by not using this strategy.

Theorem 4.2 The minimal concession strategy is in symmetric Nash equilibrium.

Proof Sketch Let st be the minimal concession strategy and suppose that agent α is using st and the other agent $\bar{\alpha}$ is using st' that is different from st . If the negotiation terminates in failure then it is clear that the outcome is worse for $\bar{\alpha}$ in comparison to the choice of using st . Suppose now that the negotiation terminates with an agreement t . Because α uses the minimal concession strategy, if $\bar{\alpha}$ stands still in one step, the negotiation will terminate in failure. Therefore we can conclude that there is no stand-still step according to st' . Let t_0 be the agreement if both parties use the minimal concession strategy. We want to show that $t_0 \succeq_{\bar{\alpha}} t$. From the definition of the minimal concession strategy, it follows that no agent stands still in this negotiation. This implies that $\bar{\alpha}$ in many steps makes a bigger concession than a minimal one. It follows then that $t_0 \succeq_{\bar{\alpha}} t$.

The Nash equilibrium of the minimal concession strategy means that when a fair agent is using the minimal strategy, the other agent is doing best by also using this strategy. In other words, the minimal concession strategy is an efficient and stable strategy for fair agents.

5. Conclusions

We have outlined a two-phase negotiation process whereby two agents, a buyer and a seller, aim at agreeing on an item fulfilling all “preferred” goals of the agents. These

goals are classified as structural and contractual. We have focused on covering the full negotiation life-cycle, from identifying items to be negotiated upon to conducting the actual negotiation for (contractual) features of these items. We have worked out how argumentation in general and assumption-based argumentation in particular can support the first phase. We have also proven several results on the outcome of the two-phase negotiation process, and defined a strategy for agents allowing them to achieve Nash equilibria.

We have made a number of simplifying assumptions. First, both agents are supposed to be honest and open. The seller agent is supposed to have no structural goals. We have ignored the need for information-seeking in both phases. In the future, we plan to extend this work by dropping these assumptions. We also plan to define a communication machinery to support our strategy and protocol.

We have illustrated our approach using a simple home-buying scenario. We believe that our approach could be fruitfully defined for other scenarios too, for example e-business scenarios like the ones studied in the ARGUGRID project ⁷. We plan to study these other scenarios in the future.

The first phase is supported by a decision-making mechanism using argumentation and preferences. A number of such decision-making mechanisms exist, e.g. [8,10,9,2]. In this paper, we have provided an argument-based framework that can deal with decision making, uncertainties and negotiation but we have restricted ourselves only to a simple and ideal case where we assume that the agents are honest and open to each other.

The second phase could also be supported by argumentation. The use of argumentation here could be beneficial also to support resolution of disputes over contracts. We plan to explore this in the future.

Several works exist on argumentation-based negotiation [11]. For example, [12] propose a protocol and a communication language for dealing with refusals in negotiation. It would be useful to see how this protocol and communication language may be used to support the two-phase negotiation framework we have defined. Also, [1] presents an abstract negotiation framework whereby agents use abstract argumentation internally and with each other. Our framework instead is tailored to the specific case of contract negotiation and assumes a very concrete and structured underlying argumentation framework.

Our minimal concession strategy for fair agents is inspired by the monotonic concession protocol of [14], though it differs from it in significant ways. In our framework the agent moves alternatively where in [14] they move simultaneously. The condition for terminating the negotiation is also different. As a result, the minimal concession strategy is a symmetric Nash equilibrium in our framework while the corresponding strategy in [14] is not. Other work exists on deploying the minimal concession strategy within multi-agent systems, e.g. [7,13], looking at negotiation amongst multiple agents. In this paper we have considered just two agents, and focused instead on the full negotiation process, from the identification of issues to bargain about to the actual bargaining, thus linking argumentation-based decision making to the monotonic concession protocol.

⁷www.argugrid.eu

Acknowledgements

This research was partially funded by the EU ARGUGRID project. We thank Kevin Ashley for some useful suggestions for future work.

References

- [1] L. Amgoud, Y. Dimopolous, and P. Moraitis. A unified and general framework for argumentation-based negotiation. In *Proc. AAMAS'2007*, 2007.
- [2] K. Atkinson and T. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 171(10–15):855–874, 2007.
- [3] A. Bondarenko, P.M. Dung, R.A. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1-2):63–101, 1997.
- [4] P.M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
- [5] P.M. Dung, R.A. Kowalski, and F. Toni. Dialectic proof procedures for assumption-based, admissible argumentation. *Artificial Intelligence*, 170:114–159, 2006.
- [6] P.M. Dung, P. Mancarella, and F. Toni. Computing ideal sceptical argumentation. *Artificial Intelligence, Special Issue on Argumentation in Artificial Intelligence*, 171(10–15):642–674, 2007.
- [7] U. Endriss. Monotonic concession protocols for multilateral negotiation. In P. Stone and G. Weiss, editors, *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*, pages 392–399. ACM Press, May 2006.
- [8] A.C. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In *Proc. AAMAS'03*, pages 883–890, 2003.
- [9] M. Morge and P. Mancarella. The hedgehog and the fox: An argumentation-based decision support system. In *Proc. ArgMAS*, 2007.
- [10] I. Rahwan and L. Amgoud. An argumentation-based approach for practical reasoning. In *Proc. AAMAS'06*, pages 347–354. ACM Press, 2006.
- [11] I. Rahwan, S. Ramchurn, N. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4):343 – 375, 2003.
- [12] J. van Veenen and H. Prakken. A protocol for arguing about rejections in negotiation. In *Proc. ArgMAS'06*, volume 4049 of *LNAI*, pages 138–153. Springer, 2006.
- [13] D. Zhang. Reasoning about bargaining situations. In *Procs AAAI-07*, pages 154–159, 2007.
- [14] G. Zlotkin and J. S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proc. IJCAI*, pages 912–917, 1989.