

A Generic Framework for Argumentation-Based Negotiation

Markus M. Geipel¹ and Gerhard Weiss²

¹ Chair of Systems Design, ETH Zurich, Kreuzplatz 5, 8032 Zurich,
Switzerland, mgeipel@ethz.ch

² Software Competence Center Hagenberg GmbH, Softwarepark 21, 4232 Hagenberg,
Austria, gerhard.weiss@scch.at

Abstract. Past years have witnessed a growing interest in automated negotiation as a coordination mechanism for interacting agents. This paper presents a generic, problem- and domain-independent framework for argumentation-based negotiation that covers both essential agent-internal and external components relevant to automated negotiation. This framework, called Negotiation Situation Specification Scheme (N3S), is both suited as a guideline for implementing negotiation scenarios as well as integrating available approaches that address selective aspects of negotiation. In particular, N3S contributes to the state of the art in automated negotiation by identifying and relating basic argument types and negotiation stages in a structured and formal way.

1 Introduction

Agents are autonomous entities that are situated in an environment and capable of flexible action [30,31]. Agents typically have conflicting interests, incompatible goals and limited capabilities, and so there is a need for principles and mechanisms that enable agents to coordinate themselves. In the multi-agent area various such principles and mechanisms have been proposed, including partial global planning [3], contracting [1,29], commitments and conventions [6], GPGP/TAEMS [10], auctioning, voting, and many others. Recent years have witnessed a growing interest in automated negotiation as a coordination mechanism especially for complex applications [11,14,19].

Negotiation is a communication-based, knowledge-intensive process during which agents try to come to mutually acceptable agreements by exchanging appropriate arguments that influence – convince, persuade, etc. – the respective other agents [13,21]. Many approaches to automated negotiation have been described in the literature. In [20] these approaches are divided into three groups: the more traditional game-theoretic and heuristics-based approaches, and argumentation-based approaches. Compared to each other, the argumentation-based approaches are better suited for applications in which agents have incomplete or inconsistent information about each other and the environment in which they act. Argumentation-based negotiation (ABN), however, has many facets and is a quite complex mechanism. As a consequence,

and as has been already argued in [20], there is a strong need for a framework that identifies the essential components that are needed to conduct automated negotiation. This paper presents a generic argumentation framework, called Negotiation Situation Specification Scheme (N3S), that we developed in response to this need. N3S, which in part is inspired by the real-world negotiation method proposed in [5], covers negotiation both from an agent-internal perspective and an external perspective. While the former perspective focuses on the individual agents' cognitive and societal abilities needed for negotiation, the latter perspective concentrates on the communication and argumentation processes in which the agents are involved during negotiation. Specifically, N3S addresses the open issue of structuring the overall negotiation and argumentation process by identifying argument types and relating them to different negotiation stages.

The paper is structured as follows. Section 2 describes selected related work. Sections 3 and 4 introduce the basic agent-internal elements and the external perspective of N3S, respectively. Section 5 illustrates the overall flow of argumentation induced by N3S. Finally, Section 6 summarizes the key features of N3S, compares N3S to related approaches, and shows important research issues raised by N3S.

2 Related Work

Several frameworks for ABN have been proposed in the literature. Four of the most prominent representatives of these frameworks are overviewed below; others are, for instance, described in [15,16,26,27].

The interest-based negotiation framework [17] builds on the observation that in human negotiation an agreement is often easier to achieve through argumentation about interests (i.e., desires and goals adopted to fulfill the desires) rather than positions – while positions may be fully incompatible, some of the interests underlying incompatible positions may even be fully shared by the negotiating parties. As a consequence, the concept of conflicts is essential to this framework, and different conflicts (namely, conflicts w.r.t. resources, goals, desires and beliefs) and conflict resolution methods (namely, concession making, exploration of alternative offers/goals, and persuasion) are distinguished. Specifically, three types of arguments are identified, including arguments for beliefs, arguments for desires, and arguments for plans. Other key elements of the framework are a communication language and protocol and a methodology for negotiation strategy design.

The trust/persuasion-based negotiation framework [22] concentrates on the reduction of multiple sources of uncertainty a negotiating agent typically is confronted with. Specifically, the framework includes two main components. First, an interaction-based trust model, called CREDIT, that allows to reduce uncertainty regarding the reliability and honesty of negotiation partners. CREDIT aims to assess the trustworthiness of other agents by taking into account own experience from direct interactions and experience gathered by other agents.

Second, a model of persuasive negotiation (i.e., negotiation based on rewards that are either given or asked from one agent to another) that allows to reduce uncertainty regarding the preferences and action sets of negotiation partners during bargaining. This model is complemented by a protocol for persuasive argumentation and a strategic reasoning mechanism for generating persuasive arguments. The framework distinguishes two broad classes of illocutions to be used in persuasive negotiation, namely, negotiation illocutions (*propose* and *accept*) and persuasive illocutions (*reward* and *ask_reward*).

The logic-based negotiation framework proposed in [9] distinguishes six argument types known from human negotiation: threats to produce goal adoption or abandonment; promises of future reward; appeal to past reward; appeal to precedents as counterexamples; appeal to prevailing practice; and appeal to self-interest. In the case multiple arguments are available to an agent, he applies them successively in a predefined order (which appears to be plausible but not necessarily optimal). The framework emphasizes agent-internal elements, and does not provide explicit communication protocols.

The formal framework suggested in [28] is intended to cover negotiation scenarios in which an agent tries to convince another agent to execute a particular task on his behalf. Negotiation, considered as a sequence of offers and counter offers containing values for the issues under negotiation, is achieved through the exchange of illocutions in a shared communication language. Two types of negotiation-relevant illocutionary acts are distinguished: acts used to make offers (*offer*, *request*, *accept*, *reject*, *withdraw*), and acts used for argumentative persuasion (*appeal*, *threaten*, *reward*). The framework allows to specify attack and support relationships among arguments as well as authority relationships between agents. Based on these relationships, arguments are evaluated and generated.

While the available ABN frameworks provide guidance regarding the types of arguments and illocutions needed during negotiation, they do not structure the overall argumentation process per se and it remains unclear whether and how the different arguments and illocutions correspond to specific phases of negotiation. Such a correspondence, however, is essential to the realization of complex automated negotiation scenarios.

3 Agent-Internal Elements of N3S

N3S distinguishes five agent-internal elements relevant for automated negotiation: options, constraints, interests, utilities and arguments. These are described below. In the following, let $A = \{a_1, \dots, a_n\}$ denote the set of agents participating in the negotiation process.

3.1 Options

N3S assumes that each agent a has his own set of options which he deems feasible. This set is denoted by $O_{feasible}^{(a)} \subseteq \mathbb{O}$ where \mathbb{O} is the set of all possible options.

According to N3S, an option defines a world state, no matter which modeling perspective is applied. In the case of a task-oriented perspective of the world, each option $o \in O_{feasible}^{(a)}$ defines an ordered set of task allocations, that is, a collection of tasks $T = \langle T_1, \dots, T_n \rangle$ where $\forall i T_i$ is the task assigned to $a_i \in A$.

Example 1. Two agents, a_1 and a_2 , negotiate the assignment of the tasks “collect data from database” ($T_{collect}$) and “analyze data” ($T_{analyze}$). An option of a_1 may be to take care of data collection while a_2 takes care of data analysis, i.e., $\langle \{T_{analyze}\}, \{T_{report}\} \rangle$.

In value-oriented domains, each option $o \in O_{feasible}^{(a)}$ defines a set of values of parameters describing the world or a part of it.

Example 2. An agent a_b wants to buy a car from a sales agent a_s , where one of his buying options specifies the following car attribute values (“AC” stands for “Air Conditioning”): ($price \leq 50000, brand \in \{\text{Mercedes}, \text{BMW}\}, hasAC = \text{true}, topSpeed \geq 200$) (assuming prices in € and speed in km/h throughout this paper).

An important subclass of options distinguished by N3S are *non-negotiated options*, that is, options that can be achieved by an agent a even *without* agreement of any of the agents in $A \setminus a$. Note that each non-negotiated option, $o_{nn}^{(a)}$, of an agent a is feasible by definition, that is, $o_{nn}^{(a)} \in O_{feasible}^{(a)}$. In this paper, $O_{nn}^{(a)}$ denotes the set of all non-negotiated options of agent a . (The utility of the best $o_{nn}^{(a)}$ defines the acceptance bottom line for the agent (see 4.2), and non-negotiated options play an important role in the argumentation process (see 3.4).)

Example 3. Consider the car dealer scenario from Example 2. If a_s is not the only sales agent out there, a_b can also buy a car from another sales agent a_{s2} . The options a_{s2} proposes to a_b are a_b 's non-negotiated options in his negotiation with a_s .

Negotiation is a chance for the agent to ameliorate his situation by cooperation, but not every negotiation necessarily generates acceptable options. If agents agree upon an option, they commit themselves to achieving a specific world state. In the sequel, o^* denotes the option upon which the agents finally agree. A negotiation is said to fail, if $o^* = \epsilon$ (i.e., o^* is empty).

3.2 Constraints

A constraint defines an option or a set of options which an agent considers as impossible to realize (under current circumstances) and thus as not negotiable. More specifically, because options define world states, a constraint defines world states which an agent does consider as *not* feasible. N3S assumes that each agent a has his own set of constraints $C^{(a)} \subseteq \mathbb{C}$ where \mathbb{C} denotes the set of all possible constraints. The union of all individual constraint sets $C = \bigcup_{a \in A} C^{(a)}$ rules

out the options that are out of question for mutual acceptance. The constraint set C is not necessarily congruent with the set of constraints C_{real} that reality imposes. In the case that the agents settle on an option $o^* \in (C_{real} \setminus C)$, the outcome of the negotiation is void and from that it follows that $o^* = \epsilon$.

Example 4. The constraints for the sales agent a_s from Example 2 may include that she does not sell BMWs going faster than 250 km/h: $(\neg((brand = BMW) \wedge (topSpeed > 250))) \in C^{(a_s)}$. The constraints for the buyer a_b may include that he cannot pay more than 60K: $(\neg(price > 60000)) \in C^{(a_b)}$.

3.3 Interests and Utilities

It is up to an agent to decide whether to accept or reject an option proposed by another agent. To make this decision in a rational way, N3S suggests to take into account the utility of an option (hence of the world state it defines). Thereby, N3S assumes that interests, especially conflicting ones, are central to calculate the utility of options and, with that, to decide on acceptance and rejection: the better the interests are addressed by an option, the higher is the utility of the option. In particular, N3S exploits interests to calculate the utility of an option in two steps, namely, in a first step from option to interest and in a second step from interest to utility. This crucial separation splits the debatable part of an agent's reasoning from the undebatable part. Clearly the mapping from an option to an interest follows rational principles (in the case of rational agents), while the mapping from interest to utility does not. For instance, while it is possible to argue rationally whether air conditioning (part of the option) adds to comfort (interest), it is *not* possible to argue rationally how much comfort it adds to an agent's personal utility. In other words, an agent's utility function is private, but an agent's mapping from options to interests is not. (An agent, as a rational entity, may not only keep track of his own mind state, but may also try to model the mind states – including the interests – of his negotiation partners.)

Formally, there is a mapping from the option $o \in \mathbb{O}$ to the degree of the interest $i \in I$ in o : $\mathbb{O} \times \mathbb{I} \mapsto \mathbb{R}$. This mapping is called *interest-satisfaction-mapping* (*ISM*). Let I_{a_1} be the set of interests of agent a_1 . As this set is finite, *ISM* can be formulated as follows: $\mathbb{O} \mapsto \mathbb{R}^{|I_{a_1}|}$. The result of this formulation is a vector called *interest-satisfaction-vector* *is*. With that, the utility function can be written as $u : \mathbb{R}^{|I_{a_1}|} \mapsto \mathbb{R}$.

Example 5. Consider the car dealer scenario from Example 2 and assume that a_b is interested in some comfort $i_{comfort} \in I_{a_b}$. The availability of air conditioning (AC) fulfills this interest, while not having AC leaves $i_{comfort}$ unaddressed. $hasAC \Rightarrow i_{comfort}$ would then be part of the *ISM*. A utility function could, for example, calculate the utility of an option as the weighted average of the degrees of interest fulfillments: $u(o) = \sum_{i \in I} \omega_i ISM(o, i)$. If the existence of an AC makes up 50 percent of a car's utility for the agent, $\omega_{i_{comfort}}$ would be 0.5 in that case.




Fig. 1. The taxonomy of arguments (see the text for formal definitions).

3.4 Arguments and Argument Taxonomy

A rational agent always aims to maximize the utility he gets from negotiation. This means he will try to get the best of his options to which the other agent’s will agree. According to N3S, argumentation is used to set the stage for a final agreement (negotiation stages distinguished by N3S are described in 4.3). Specifically, N3S allows to distinguish different types of arguments, depending on the way in which they aim to influence the opposing agents’ “mind”. These types, which are summarized in Figure 1, are described below.

Constraint-Arguments Constraint-Arguments, referred to as *c-arguments*, aim at influencing the other agents’ constraint sets. The idea is to widen $O_{feasible}^{(a)}$ where a is the opponent or to narrow $O_{feasible}^{(a)}$ where a is the agent itself. Note that for an effective negotiation it should be in the interest of all parties to achieve a view of the constraint set which is as coherent as possible.

Example 6. In the car dealer scenario from Example 2, a c-argument could be “I’m sorry Sir, but we don’t sell BMWs that go faster than 250”:
 $(\neg(\text{brand} = \text{BMW} \wedge \text{speed} > 250)) \in C.$

Thread-Arguments Threads target the agents’ non-negotiated options repertoire. Arguments of this type that aim to decrease the utility of the other agent’s $O_{nn}^{(a)}$ are called *active threads*. Moreover, active threads are called *agressive* if the agent announces personal involvement in downgrading the other agent’s $O_{nn}^{(a)}$. Non-negotiated-option arguments, shortly referred to as *n-arguments*, that aim to increase the utility of the agent’s own $O_{nn}^{(a)}$ are *passive threads* from the perspective of the opponents. A rational agent a should be aware of the fact that passive threads only make sense if his $O_{nn}^{(a)}$ is better than the other agents expect it to be. (From an ethical point of view, threads are legitimate as long as they are not aggressive.)

Example 7. In the car dealer scenario from Example 2, a passive thread is to tell the sales agent that one got an offer for a Mercedes with AC and a top speed of 300 km/h for 51,000 € from another car dealer: $(51000, \text{Mercedes}, \text{true}, 300) \in O_{nn}^{(ac)}$.

Disclosure-Argument These arguments, referred to as *d-arguments*, do not change the option spaces. They rather support the other agents in their joint problem solving by disclosing own interests i and by mapping them to a utility measure.

Example 8. A customer could reveal to the sales agent that she is interested in comfort in the car: $(i_{comfort} \in I_a, \text{hasAC} \Rightarrow i_{comfort})$. The first part states that she is interested in comfort, and the second part states that AC needs to be true to fulfill the interest.

Utility-Arguments Utility-Arguments, or *u-arguments* for short, aim at influencing the other agent’s utility function. The idea is to widen $O_{agreeable}^{(a)}$ where a is the opponent (*active* utility-argument) or to narrow $O_{agreeable}^{(a)}$ where a is the agent itself (*passive* utility-argument).

Example 9. Consider Example 8. The sales agent could argue that an AC does not necessarily lead to higher comfort. If the customer intends to use the car in Greenland, she won’t need an AC to keep the car cold: $(\neg AC \rightarrow (\text{Temperature}(car) = \text{Temperature}(country)), \text{Temperature}(country) < 23, \text{Temperature}(car) < 23 \rightarrow \text{comfort})$. So, the sales agent is saying that without an AC the car will have the same temperature as the country and this temperature is below 23 degrees.

4 External Elements of N3S

4.1 Communication

N3S suggests that negotiation essentially is the communication of suggestions (including both options and arguments). Communication starts by suggesting to negotiate, and is continued by the agents through rejecting and accepting suggestions. Communication ends if all agents accept the same suggestion or if one or several agents decide to not longer participate in the negotiation process. Figure 2 depicts the general scheme.

Starting and Terminating a Negotiation Negotiation is a conversation that has a well defined beginning and ending. Before the real negotiation takes place, it has to be clear which agents participate and what the option space is. Thus one agent will suggest a negotiation by sending a message with the following information:




Fig. 2. The logic of suggesting, rejecting and accepting

1. a specification of the option space (*what* is negotiated?);
2. a list of known participating agents (*who* negotiates?);
3. a unique identifier of the (intended) negotiation.

The option space may be defined by an ontology accessible to each participating agent in order to ensure an efficient negotiation. (Note that N3S does not prescribe the provision of an ontology by the user. Rather, it would be possible as well that agents themselves choose or “meta-negotiate” the ontology they want to use.)

Negotiation starts if all participating agents accept the offer to run a negotiation. An ongoing negotiation ends if one or several partners refuse to continue the negotiation process or if all partners agree on (accept) the same suggestion.

Making Proposals An option is always suggested by an agent. N3S requires that each of the participating agents has the choice between sending a reject or an accept message once an option has been suggested. In the case of the acceptance of an option by all agents, the negotiation ends.

Providing Arguments N3S assumes that arguments, just like options, are subject to the suggest-accept-reject scheme. This makes sense because the participating agents can now track whether their negotiation partners accept an argument. In particular, it allows them to adjust their own reasoning and to model the reasoning of the other agents. As another means for increasing negotiation efficiency, a negotiation protocol could be applied that allows to propose only arguments that are (logically) consistent with previously accepted arguments.

4.2 Option Spaces

N3S structures the option space from the perspective of rationally acting agents. An obvious option class being relevant to rational agents is the class of options deemed feasible by all agents:

$$O_{feasible} = \bigcap_{a \in A} O_{feasible}^{(a)} \quad (1)$$

(Sub-)Space	Description
\mathbb{O}	includes every possible option
$O_{feasible}^{(a)} \subseteq \mathbb{O}$	includes all options that agent a judges to be possible in respect to $C^{(a)}$
$O_{feasible} = \bigcap_{a \in A} O_{feasible}^{(a)}$	includes all options that are perceived as possible by all agents in A
$O_{agreeable}^{(a)} \subseteq O_{feasible}^{(a)}$	all options that are agreeable for agent a
$O_{agreeable} = \bigcap_{a \in A} O_{agreeable}^{(a)}$	all options agreeable to all agents in A

Table 1. Option hierarchy induced by N3S.

To be considered as feasible by all agents, however, is not sufficient for an option to be accepted. A rational agent can not be expected to agree upon an option if he fares better in not agreeing to anything. If no agreement is reached, each agent has to fall back to his best non-negotiated option which is given by:

$$\text{argmax}(\text{util}^{(a)}(o_{nn} \in O_{nn}^{(a)})) \quad (2)$$

So, every option o which (i) has a higher utility than the best of his non-negotiated options and (ii) is feasible is in the set of agreeable options:

$$\forall_o [(\text{util}(o) \geq \max(\text{util}^{(a)}(o_{nn} \in O_{nn}^{(a)})) \wedge o \in O_{feasible}^{(a)}) \Rightarrow o \in O_{agreeable}^{(a)}] \quad (3)$$

The set of options agreeable to all agents in A is defined as the intersection of all agents' individual agreeable option sets:

$$O_{agreeable} = \bigcap_{a \in A} O_{agreeable}^{(a)} \quad (4)$$

From (1), (3) and (4) it follows that

$$O_{agreeable} \subseteq O_{feasible} \quad (5)$$

Table 1 summarizes the above considerations.

Example 10. Consider the care dealer situation introduced in Example (2). As we already saw in Example (4), the constraints for the customer a_b include that he cannot pay more than 60K: $\neg(\text{price} > 60000) \in C^{(a_b)}$. Thus an option (61000, BMW, true, 200) would be in \mathbb{O} but not in $O_{feasible}^{(a)}$. Against that, (59000, BMW, true, 200) would be feasible. However, if a_b also got the offer (58000, BMW, true, 200) from another sales agent, (59000, BMW, true, 200) might not be in $O_{agreeable}^{(a)}$. Depending on his *ISM* and his utility function (59000, BMW, true, 300) could again be in $O_{agreeable}^{(a)}$, if the difference in speed (300 km/h instead of 200 km/h) is worth the additional 1,000 € to a_b .




Fig. 3. The transition diagram for the N3S negotiation stages.

4.3 Stages of the Negotiation Process

Based on the option space hierarchy, N3S distinguishes different stages of the negotiation process and allows to assign specific negotiation strategies to these stages. These stages are described below. The stage transition diagram is shown in Figure 3.

Pre or post Negotiation (0) Before or after a negotiation process the option space of an agent does only consist of the non-negotiated options. These non-negotiated options may well include options that were made possible by preceding or still ongoing negotiations. An obvious negotiation strategy for a rational agent in this stage is to broaden his option portfolio. Thus an agent should check whether there is a prospect of making new options available through negotiation. A way to achieve this is to start a negotiation process in order get information about the options of other agents. Initiation of a negotiation process will lead to one of the tree following stages. The goal is to reach stage III which is the only stage at which each participating agent has an incentive to agree.

Negotiation without available options (I) In this stage, at least two agents are engaged in negotiation, but $O_{feasible} = \emptyset$. A strategy to identify options in this stage is to argue about the other agents' constraints $C^{(a)}$. Thus c-arguments will dominate the conversation.

Negotiation without incentive to agree (II) In this stage, at least two agents are engaged in negotiation and there are feasible options (i.e., $O_{feasible} \neq \emptyset$), but for at least one agent none of these options seems to have a higher utility than his best non-negotiated option (i.e., $O_{agreeable} = \emptyset$). There are two strategies to open up $O_{agreeable}$: to influence and change the utility functions of the other agents through u-arguments; and to attack the other agents' $O_{nn}^{(a)}$ through n-arguments.

Negotiation with incentive to agree (III) In this stage, at least two agents are engaged in negotiation and $O_{agreeable} \neq \emptyset$, and thus it is rational for both parties to successfully finish the negotiation. A negotiation protocol like the monotonic concession protocol presented in [31] to clarify specific terms of agreement and to find a fair option for all parties. Besides this, there is still the possibility to strengthen one’s position with n-arguments and u-arguments. Note that it is well possible to fall back to stage II by using (e.g.) passive threads.

5 Putting the Pieces Together – A Sample Flow of Arguments

The examples provided in the preceding sections illustrate specific aspects of N3S. In this section, an example is given that illustrates the interdigitation of agent-internal and external elements of N3S. The emphasis of this example is on the overall flow of argumentation. (For the sake of clarity, the used formalism is kept as simple as possible and utilities – functions and u-arguments – are not considered. Note that there is no commonly accepted benchmark negotiation problem in the field; we think the problem introduced below may be a good benchmark candidate for both theoretical and practical analysis.)

5.1 The “Polygon Negotiation Problem”

Basic Setting Assume there are two agents (a_1, a_2) and a polygon P of order n ($y = c_0x^0 + c_1x^1 + \dots + c_nx^n$). Coefficients are in the range $[-1, 1]$ and initially set to zero: $c_0 = c_1 = \dots = c_n = 0$. If the agents do not find an agreement, the initial values are kept. Each agent wants to maximize P at certain positions x ; thus, these positions represent their interests. If a_1 is interested in maximizing P at $x = 0.5$, this is written as “ $\max(P(0.5)) \in I_{a_1}$ ”. Moreover, assume each agent has only one interest (hence, reasoning about utilities is not necessary) and possesses the capability to manipulate at least one coefficient. If a_1 can set the coefficients c_0 and c_3 , this is written as “ $\text{canSet}(a_1, \{c_0, c_3\})$ ”.

Instantiation The following information determines the initial conditions:

$$\begin{aligned}
 &P \text{ of order } n = 3 \\
 &\max(P(1.0)) \in I_{a_1} \Rightarrow u_{a_1} = c_0 + c_1 + c_2 + c_3 \\
 &\max(P(-1.0)) \in I_{a_2} \Rightarrow u_{a_2} = c_0 - c_1 + c_2 - c_3 \\
 &\text{canSet}(a_1, \{c_3\}) \\
 &\text{canSet}(a_2, \{c_0\})
 \end{aligned}$$

5.2 Sample Flow of Argumentation

Table 2 shows a possible sequence of arguments exchanged by the two agents involved in the Polygon Negotiation Problem. According to this sequence, agent a_2 starts right away with a proposal (1). As a_1 can not change c_1 and c_2 and

1	a_2 propose	$\text{set}(a_1, \{c_1 = -1.0, c_2 = 1.0, c_3 = -1.0\}),$ $\text{set}(a_2, \{c_0 = 1.0\})$
2a	a_1 reject	$\text{set}(a_1, \{c_1 = -1.0, c_2 = 1.0, c_3 = -1.0\}),$ $\text{set}(a_2, \{c_0 = 1.0\})$
2b	c-argument	$\neg\text{canSet}(a_1, \{c_2, c_1\})$
3	a_2 d-argument	$\text{canSet}(a_2, \{c_0\})$
4	a_1 d-argument	$\text{canSet}(a_1, \{c_3\})$
5	a_2 propose	$\text{set}(a_1, \{c_3 = -1.0\}), \text{set}(a_2, \{c_0 = 1.0\})$
6a	a_1 reject	$\text{set}(a_1, \{c_3 = -1.0\}), \text{set}(a_2, \{c_0 = 1.0\})$
6b	d-argument	$\max(P(1.0)) \in I_{a_1}$
6c	t-argument	$\{c_0 = 0.0, c_1 = 0.0, c_2 = 0.0, c_3 = 0.0\} \in O_{n_n}^{(a_1)}$
7	a_2 propose	$\text{set}(a_1, \{c_3 = -0.5\}), \text{set}(a_2, \{c_0 = 1.0\})$
8a	a_1 reject	$\text{set}(a_1, \{c_3 = -0.5\}), \text{set}(a_2, \{c_0 = 1.0\})$
8b	propose	$\text{set}(a_2, \{c_0 = 1.0\})$
9	a_2 accept	$\text{set}(a_2, \{c_0 = 1.0\})$

Table 2. An argumentation sequence in the Polygon Negotiation Problem.

thus he rejects the proposal (2a) and supplies a reason with a c-argument (2b). a_2 replies with a d-argument (3) informing a_1 about her capabilities. a_1 does the same (4). These two disclosures lead from stage I, in which no feasible options are known, to stage II. In this stage the agents can make feasible proposals (which, however, may be not acceptable for the respective other agent). a_2 brings forth a new and feasible proposal (5), but a_1 rejects for a good reason (6a), reveals his interests (6b), and tells a_2 that rejection is a non negotiated option for him (6c). At that point it is clear to a_2 that, from a_1 's point of view, rejecting yields a utility of 0.0. In (7) stage III is entered, as the proposal of a_2 is acceptable (for a_1 it yields a utility greater than 0.0). Thus there is an incentive for a_1 to agree. However, instead of agreeing immediately, a_1 tries to get more: he rejects again (8) and makes a counter proposal (8b). At this point a_2 accepts (9) and thus the negotiation ended successfully with a joint agreement.

6 Discussion and Conclusions

N3S has the following main characteristics:

- it is generic negotiation framework that is problem- and domain-independent, covers both agent-internal and external elements of negotiation, and is applicable to a broad range of negotiation scenarios;
- it introduces the concept of option and option hierarchy as a key ingredient of rational negotiation;
- it provides a practical taxonomy of arguments that supports an agent in arguing in a systematic way; and
- it effectively structures the negotiation process into distinct stages.

Together these characteristics make N3S unique and distinct from other ABN frameworks (see Section 2). Based on these characteristics, a particular strength of N3S is its potential to serve both as a “technical ” guideline for implementing automated negotiation processes and a “conceptual” guideline for unifying and integrating available approaches addressing selected aspects of negotiation. Examples of such aspects are society protocols for agent interaction [7], nested argumentation [12], argumentation based on strategic reasoning [2], commitment based argumentation [8], trust-based negotiation [22], and learning of argumentation strategies [4]. The integration of all these approaches, or of components of them, into a coherent whole is, in our opinion, *the* research and application challenge in the field.

The key contribution of N3S to the state of the art lies in the identification of both arguments types and negotiation stages. Specifically, to our knowledge no other currently available ABN framework, including those mentioned in Section 2, *relates* argument types and negotiation stages in a structured way as N3S does. By relating argument types and negotiation stages, N3S opens new possibilities in enabling rational agents to argue *efficiently*. We already addressed facets of this issue in our approaches to practical and social reasoning (e.g., [24,25]), and our current research aims to explore the relationships and possible synergies between N3S and these approaches as well as related ones (e.g., [2,18,23]).

An important question raised by N3S is how its argument taxonomy, which essentially is “option-centered”, compares to the argument taxonomies proposed within the related frameworks. Another relevant research topic raised by N3S is how the negotiation stages of N3S are related to the types of conflicts among agents identified in other frameworks (e.g., [17]). Last but not least, it is currently an open issue how the N3S negotiation strategies (see 4.3) are related to negotiation strategies, tactics and heuristics proposed elsewhere (see, e.g., the various papers on strategic aspects of negotiation in [11]).

Acknowledgements. The authors thank Achim Rettinger, Ana Balevic and Florian Echter for many useful and inspiring discussions automated negotiation in general and N3S in particular. Special thanks goes to the Chair of Systems Design of ETH Zurich, headed by Frank Schweitzer, for generous financial support of our joint research efforts on N3S.

References

1. M.R. Andersson and T.W. Sandholm. Leveled commitment contracts with myopic and strategic agents. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98)*, pages 38–45, 1998.
2. J. Bentahar, M. Mbarki, and B. Moulin. Strategic and tactic reasoning for communicating agents. In Maudet et al. [11]. Available at <http://homepages.inf.ed.ac.uk/irahwan/argmas/argmas06/>.
3. E.H. Durfee and V.R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-21(5):1167–1183, 1991.

4. C.D. Emele, T. Norman, and P. Edwards. A framework for learning argumentation strategies. In Maudet et al. [11]. Available at <http://homepages.inf.ed.ac.uk/irahwan/argmas/argmas06/>.
5. R. Fisher, W. Ury, and B. Patton. *Getting to Yes: Negotiating Agreement Without Giving In*. Random House Business Books, 2nd edition, January 2004.
6. N.R. Jennings. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250, 1993.
7. A. Kakas, N. Maudet, and P. Moraitis. Layered strategies and protocols for argumentation-based agent interaction. In Rahwan et al. [19], pages 64–77.
8. N. Karunatilake, N. Jennings, I. Rahwan, and T. Norman. Argument-based negotiation in a social context. In Parsons et al. [14].
9. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104(1-2):1–69, 1998.
10. V.R. Lesser, K.R. Decker, T. Wagner, N. Carver, A. Garvey, B. Horling, and D. Neiman (et al.). Evolution of the GPGP/TAEMS domain-independent coordination framework. *Autonomous Agents and Multi-Agent Systems*, 9(1/2):87–144, 2004.
11. N. Maudet, S. Parsons, and I. Rahwan, editors. *Argumentation in Multi-Agent Systems. Proceedings of the Third International Workshop*. 2006. Available at <http://homepages.inf.ed.ac.uk/irahwan/argmas/argmas06/>.
12. S. Modgil. Nested argumentation and its application to decision making over actions. In Parsons et al. [14].
13. H.J. Müller. Negotiation principles. In G.M.P. O’Hare and N.R. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, pages 211–230. Wiley, New York et al., 1996.
14. S. Parsons, N. Maudet, P. Moraitis, and I. Rahwan, editors. *Argumentation in Multi-Agent Systems. Proceedings of the Second International Workshop*. Lecture Notes in Artificial Intelligence, Vol. 4049. Springer-Verlag, Berlin et al., 2006.
15. S. Parsons, C. Sierra, and N. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):267–274, 1998.
16. P. Pasquier, I. Rahwan, F. Dignum, and L. Sonenberg. Argumentation and persuasion in the cognitive coherence theory: preliminary report. In Maudet et al. [11]. Available at <http://homepages.inf.ed.ac.uk/irahwan/argmas/argmas06/>.
17. I. Rahwan. *Interest-based Negotiation in Multi-Agent Systems*. PhD thesis, Department of Information Systems, University of Melbourne, Australia, 2004.
18. I. Rahwan and L. Amgoud. An argumentation-based approach for practical reasoning. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pages 347–354, 2006.
19. I. Rahwan, P. Moraitis, and C. Reed, editors. *Argumentation in Multi-Agent Systems. Proceedings of the First International Workshop*. Lecture Notes in Artificial Intelligence, Vol. 3366. Springer-Verlag, Berlin et al., 2005.
20. I. Rahwan, S. Ramchurn, N. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *Knowledge Engineering Review*, 18(4):343–375, 2004.
21. H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, Mass., 1982.
22. S.D. Ramchurn. *Multiagent negotiation using trust and persuasion*. PhD thesis, Faculty of Engineering and Applied Science, University of Southampton, UK, 2005.

23. B. Roth, R. Riveret, A. Rotolo, and G. Governatori. Strategic argumentation: a game theoretical investigation. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law (ICAIL 2007)*, 2007.
24. M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiß. Practical strategic reasoning and adaptation in rational argument-based negotiation. In S. Parsons, N. Maudet, P. Moraitis, and I. Rahwan, editors, *Argumentation in Multiagent Systems*, Lecture Notes in Artificial Intelligence, Vol. 4049, pages 122–137. Springer-Verlag, 2006.
25. M. Rovatsos, G. Weiß, and M. Wolf. An approach to the analysis and design of multiagent systems based on interaction frames. In *Proceedings of the First International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, 2002.
26. F. Sadri, F. Toni, and P. Torroni. Abductive logic programming architecture for negotiating agents. In *Proceedings of the Eight European Conference on Logics in Artificial Intelligence (JELIA 2002)*, volume 2424 of *Lecture Notes in Computer Science*, pages 419–431. Springer Verlag, 2002.
27. S. Saha and S. Sen. A Bayes net approach to argumentation based negotiation. In Rahwan et al. [19], pages 208–222.
28. C. Sierra, N.R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In M. P. Singh, A. Rao, and M. J. Wooldridge, editors, *Fourth International Workshop on Agent Theories Architectures and Languages (ATAL-97)*, volume 1365 of *Lecture Notes in Computer Science vol. 1365*, pages 177–192. Springer-Verlag, 1998.
29. R.G. Smith. The contract-net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.
30. G. Weiß, editor. *Multiagent Systems. A Modern Approach to Distributed Artificial Intelligence*. The MIT Press, Cambridge, MA, 1999.
31. Michael J. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons Ltd, march 2002.