

Practical Strategic Reasoning and Adaptation in Rational Argument-Based Negotiation

Michael Rovatsos¹, Iyad Rahwan², Felix Fischer³, and Gerhard Weiss³

¹ School of Informatics, The University of Edinburgh, Edinburgh EH8 9LE, UK
mrovatso@inf.ed.ac.uk

² Institute of Informatics, The British University in Dubai, P.O. Box 502216, Dubai, UAE
iyad.rahwan@buid.ac.ae

³ Department of Informatics, Technical University of Munich, 85748 Garching, Germany
{fischerf, weissg}@cs.tum.edu

Abstract. Recent years have seen an increasing interest of multiagent system research in employing the theory of *argumentation* for the development of communication protocols. While significant progress has been made in formalising argument-based communication, (possibly *adaptive*) agent-level *argumentation strategies* as a *practical* integration of rational agent reasoning and inter-agent argumentation dialogues have received fairly little attention. In this paper we propose the use of the InFFrA framework in argument-based negotiation. This framework allows for a strategic and adaptive communication to achieve private goals within the limits of bounded rationality in open argumentation communities. The feasibility of the approach is illustrated in an agent-based web linkage scenario, showing that its performance is comparable to that of simple proposal-based negotiation while accommodating much stricter constraints regarding “what can be said” like those used in argumentation.

1 Introduction

Communication between intelligent agents is one of the cornerstones of multiagent system (MAS) technology. Most of the time, this communication is realised in terms of (1) agent communication languages (ACLs) defining the structure of messages (usually in a speech-act like format) and (2) interaction protocols specifying admissible sequences of messages and imposing constraints on the contents of these.

Recent years have seen an increasing interest in employing the theory of *argumentation* for the development of communication protocols. This interest rests on the view that *rational agents* reason and make decisions by constructing and comparing arguments for and against particular conclusions [12, 10]. Hence, it is natural to view *rational interaction* as a disciplined process of argument exchange. As a consequence, significant progress has been made in formalising argument-based communication, founded in various formal theories of argumentation (e.g. [3, 2, 13]). One area of particular interest is *argumentation based negotiation* (ABN) [16], in which agents exchange arguments in order to reach beneficial agreements.

So far, however, fairly little attention has been paid to *argumentation strategies* as a *practical* integration of intra-agent rational reasoning and (hence, rational) inter-agent

dialogue via argument exchange.¹ Even less work has been done on argumentation strategy adaptation. This is paralleled in general ACL research (see e.g. [11, 6] for recent overviews of the field), where the problem of coming up with an optimal communication strategy that *obeys* a given semantics,² but still ensures beneficial interaction outcomes for the agents themselves, is largely unresolved.

One reason why argumentation strategies and their adaptation have received fairly little attention may be the relatively high expressiveness of argumentation protocols, which makes them both difficult to implement and hard to control during execution. In the light of the latter problem, however, the need for adaptive argumentation strategies becomes even more pressing.

In this paper, we take a first step towards the use of adaptive strategies in argument-based dialogues. For this, we conceptually follow a generic agent-centric model of strategic interaction that makes a clear distinction between agents' social behaviour on the one hand and their internal rational reasoning on the other. Rational interaction then means that agents may only engage in communication in a way that respects their alleged mental state (as suggested by their social behaviour). We practically implement a simplified form of argument-based negotiation using a particular instance of the InF-FrA social reasoning framework [19]. We illustrate the feasibility of the approach in a agent-based web linkage scenario, and show that its performance is comparable to that of simple proposal-based negotiation while imposing much stricter and much more realistic constraints.

The paper contributes to the state of the art in argument-based communication in two main ways. First, it presents the first attempt to produce a highly expressive and flexible approach to adaptive communication strategies in argument-based communication in general, and argument-based negotiation in particular. Second, our practical implementation contributes to bridging the gap between global (argumentation) protocol design and rational agent design.

The remainder of this paper is structured as follows. We first introduce a generic model of strategic interaction in a prescriptive social context. In Section 3, we then lay out our approach for adaptive reasoning about communication patterns. Section 4 introduces *interest-based negotiation*, a form of argumentation-based negotiation, which will serve as an underlying argumentation model in an application scenario. Experimental results obtained in this scenario will be given in section 5. Section 6 rounds up with some conclusions.

2 A Generic Model of Strategic Interaction

As suggested in the introductory section, the perspective we adopt towards argumentation (and towards communication in general) is entirely agent-centric. The problem we are trying to solve can be stated as follows:

¹ Preliminary work has begun to investigate the outcomes of different simple strategies in argument-based communication [1, 14].

² We mean *ostensible* adherence to a given semantics here, e.g. by acting *as if* being in a particular mental state (for mentalistic semantics [22, 5]) or *as if* entering into a social commitment (for commitment-based semantics [23, 9]).

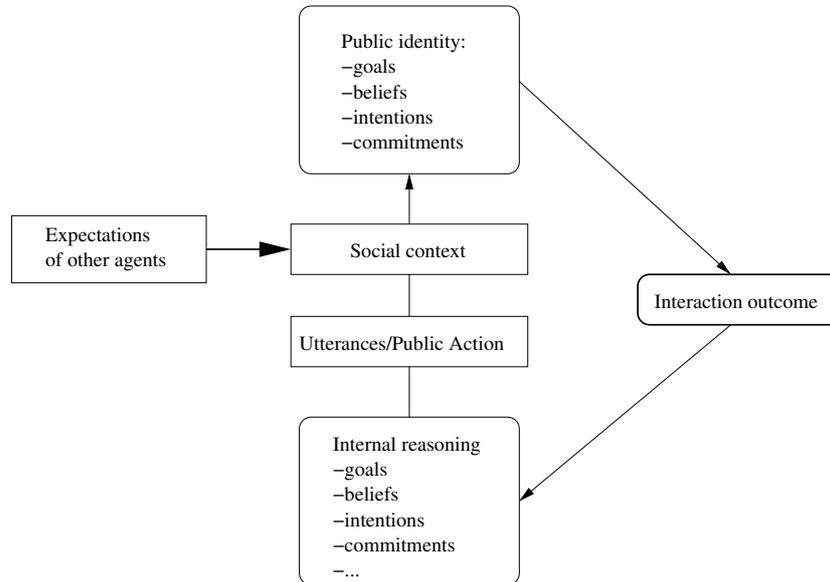


Fig. 1. Agent-centric model of strategic interaction

Given a set of dialogue patterns tied to constraints regarding (among other things) participants' internal structure, how can we design an agent that is capable of employing these patterns – in compliance with the constraints at hand – in order to optimise her own long-term profit?

We should take a minute to describe what is meant by this. In the most general sense, any communication mechanism in a MAS ties the use of certain communicative patterns (protocols, single utterances or publicly observable “non-linguistic” or “physical”, i.e. environment-manipulating, actions) to specific constraints, which may also concern mental states (such as beliefs, desires, and intentions [4]) particularly of those agents taking part in the communication. Since these prescriptive constraints are usually assumed to be common knowledge for all agents situated in the same *social context*, they (together with the actual communication) raise mutual *expectations* regarding these agents' behaviours. As figure 1 suggests, anything that is uttered by an agent is interpreted on the grounds of the social context and leads to the construction of a *public identity* of the communicating agent. This identity reflects (1) what the agent has publicly claimed about her internal state and (2) how this is seen to relate to her actual behaviour by other agents. However, although her utterances are generated by the agent's internal reasoning mechanism, the actual internal state may differ significantly from her public identity, and this is where the strategic aspect of communication comes into play: In contrast to the internal state of an agent, her public identity is subject to inspection by a peer and will evoke certain reactions on her peers' side and hence affect the outcome of any communicative interaction. Thus, the agent herself has an incentive

to manipulate her public identity, albeit the social context confines the range of possible manipulations.

ABN is an instance of this model that allows agents to influence their public identity directly, namely by conveying information about their (alleged) internal state in order to support their proposals or claims. While an agent might put forth arguments that are not in line with her internal state, the public identity itself must be consistent so as not to reveal this difference. Otherwise, other agents might simply refuse to interact with the agent at all due to the latter's inability to interact coherently.

As an example, consider an insincere agent A who is deliberately deceiving agent B when claiming that she is trying to achieve goal G . In a framework of rational communication, A would be expected to act (and speak) in accordance with this commitment to G , e.g. by not claiming to pursue contradictory goals at the same time, by dropping the goal if it is achieved or if it becomes unachievable, etc. It is reasonable to assume that unless A says anything that contradicts these principles of rationality (which can be seen as a social context in the context of communication) or indicates through her physical actions not to be pursuing G , she can maintain her public identity towards B and keep B thinking she is in fact trying to achieve G . This is exactly what is meant by strategic compliance to a social context: to ensure others have certain expectations about our own future behaviour, we must succeed in maintaining a communicative stance that is in concordance with the social context, even if our internal state and reasoning contradicts the construed public identity.

3 Reasoning with Adaptive Communication Patterns

To develop agents that are capable of strategically dealing with the complex dialogue patterns required for argumentation, we make use of the abstract social reasoning framework InFFrA proposed in [21], and, more specifically, the formal model m^2 InFFrA suggested as a concrete instance of InFFrA in [7].

The Interaction Frames and Framing Architecture InFFrA is an abstract framework for reasoning about and learning different classes of interactions in the form of so-called *interaction frames* (henceforth called frames). Each of these frames characterises a category of interaction situations in terms of (1) *roles* held by the interacting parties and *relationships* between them, (2) *trajectories* that describe the observable surface structure of the interaction, and (3) *context* and *belief* conditions that need to hold for the respective frame to be enacted. Further, InFFrA defines *framing* as the activity of constructing, adapting and strategically applying a set of interaction frames from the point of view (and in accordance with the private goals) of a single agent. Roughly speaking, framing consists of four phases:

1. Interpreting the current interaction in terms of a *perceived frame* and matching it against the normative model of the *active frame* which determines what the interaction should look like.
2. Assessing the active frame (based on whether its conditions are currently met, whether its surface structure resembles the perceived interaction sequence, and whether it serves the agent's own goals).

3. Deciding on whether to retain the current active frame or whether to *re-frame* (i.e. to retrieve a more suitable frame from one’s frame repository or to adjust an existing frame model to match the current interaction situation and the agent’s current needs) on the grounds of the previous assessment phase.
4. Using the active frame to determine one’s next (communicative, social) action, i.e. apply the active frame as a prescriptive model of social behaviour in the current interaction.

The $m^2\text{InFFrA}$ framework turns these abstract concepts of interaction frames and framing into a concrete computational model for discrete-time, two-party, turn-taking interactions. A frame in $m^2\text{InFFrA}$ consists of (1) a *trajectory*, i.e. a linear sequence of message or “physical” action patterns (possibly containing variables), (2) condition/substitution pairs that represent past enactments of the frame in terms of variable values and conditions that held at the time of the enactment, and (3) counters that keep track of (i) the frequency with which an *encounter prefix* (i.e. an initial sub-sequence of a perceived conversation) matched the frame trajectory and (ii) the frequency with which certain condition/substitution pairs appeared as instances of the frame.

3.1 Interaction Frames

To define interaction frames formally, we assume a language of speech-act like message and action patterns of the form $\text{perf}(A, B, X)$ or $\text{do}(A, Ac)$ which may contain variables or concrete values in the sender, receiver and content slots. In the case of messages (i.e., exchanged textual signals), perf is a performative symbol (e.g. *request*, *inform*), A and B are agent identifiers or agent variables and X is the content of the message taken from a first-order language \mathcal{L} .

In the case of physical actions (i.e., actions that manipulate the physical environment) with the pseudo-performative do , Ac is the action executed by A (a physical action has no recipient as it is assumed to be observable by any agent in the system). Both X and Ac may contain non-logical “substitution” variables used for generalisation purposes (as opposed to logical “content” variables used by agents to indicate quantification or to ask for a valid binding).

Interaction frames are then defined as tuples $F = (T, \Theta, C, h, h_\Theta)$, where

- $T = \langle p_1, p_2, \dots, p_n \rangle$ is a sequence of message and action patterns p_i , the *trajectory*,
- $\Theta = \langle \vartheta_1, \dots, \vartheta_m \rangle$ is an ordered list of *variable substitutions*,
- $C = \langle c_1, \dots, c_m \rangle$ is an ordered list of *condition sets*, such that $c_j \in 2^{\mathcal{L}}$ is the condition set relevant under substitution ϑ_j ,
- $h \in \mathbb{N}^{|T|}$ is a *trajectory occurrence counter* list counting the occurrence of each prefix of the trajectory T in previous conversations, and
- $h_\Theta \in \mathbb{N}^{|\Theta|}$ is a *substitution occurrence counter* list counting the occurrence of each member of the substitution list Θ in previous conversations.

While the trajectory $T(F)$ models the surface structure of message sequences that are admissible according to frame F , each element of $\Theta(F)$ resembles a past binding of the variables in $T(F)$, and the corresponding element of $C(F)$ lists the conditions required

for or precipitated by the execution of F in this particular case. $h(F)$ finally indicates how often F has been executed completely or just in part, $h_{\Theta}(F)$ is used to avoid duplicates in $\Theta(F)$ and $C(F)$. What hence distinguishes interaction frames from the methods commonly used for the specification of ACL and protocol semantics is that they allow for an explicit representation of *experience* regarding their practical use.

To illustrate these definitions, consider the following example, in which we use an abbreviated notation to capture all elements of the definition more concisely:

$$\begin{aligned}
 F = & \left\langle \overbrace{\left\langle \xrightarrow{5} \text{request}(A, B, X) \xrightarrow{3} \text{do}(B, X) \right\rangle}^T, \right. \\
 & \left. \left\langle \overbrace{\{\text{can}(B, X)\}}^{\Theta_1}, \overbrace{\{\text{can}(B, \text{pay}(S))\}}^{\Theta_2} \right\rangle} \right. \\
 & \left. \left\langle \overbrace{\xrightarrow{2} \langle [A/a], [B/b], [X/\text{pay}(\$100)] \rangle}^{C_1}, \right. \right. \\
 & \left. \left. \overbrace{\xrightarrow{1} \langle [A/b], [B/a], [X/\text{pay}(S)] \rangle}^{C_2} \right\rangle \right\rangle
 \end{aligned}$$

As for the individual elements of F , the *trajectory* T captures the following interaction experience: A has asked B five times to perform (physical) action X , B actually did so in three of these instances.³ Knowledge about the remaining two cases would typically be stored in a different frame. The substitutions Θ and conditions C summarise the following observations: In two of the successful instances (Θ_1/C_1), it was a who asked and b who heeded the request, and the action was to pay \$100. In both cases, $\text{can}(b, \text{pay}(\$100))$ held true (C_i always corresponds to Θ_i in a frame). In the third case, roles were swapped between a and b and the amount S remains unspecified (which does not mean that it did not have a concrete value, but simply that this was abstracted away in the frame).

Thus, $m^2\text{InFFrA}$ frames facilitate the description of (observations about) dialogue sequences by means of generalised message (and action) patterns together with past variable values and context conditions. At the same time, they can be used as a concrete representation for the abstract social context mentioned above, combining behavioural expectations and context conditions in the most general way.

3.2 Frame Semantics

As for the semantics of frames, these are defined in terms of a probabilistic model over the possible *continuations* of a dialogue, i.e. current frame knowledge induces a probability distribution over possible conclusions to a dialogue given a prefix sequence of what has been observed in a dialogue at a certain point in time. This is done using a domain-dependent real-valued similarity measure σ on message patterns (and sequences thereof). σ is defined using a distance metric between messages⁴ and

³ do is used as a special performative to indicate execution of physical actions.

⁴ See [8] for details on this metric.

extended into a similarity $\sigma(\vartheta, F)$ between a substitution ϑ and an entire frame F with trajectory $T(F)$ and substitutions $\Theta(F)$ by summing over individual similarities $\sigma(T(F)\vartheta, T\Theta_i(F))$ between the message pattern sequence induced by ϑ and the past cases stored in $\Theta(F)$. Moreover, $\Theta_i(F)$ is only considered if the corresponding condition set $C_i(F)$ is currently satisfied.⁵ This means that certain valuations for variables in the message “templates” of the trajectory that have been observed in the past will only be considered if the conditions under which they were observed hold in the current state.

Given a frame repository $\mathcal{F} = \{F_1, \dots, F_n\}$ representing the agent’s interaction experience, and a (possibly empty) sequence w of messages perceived in the current conversation, the probability of encounter prefix w being concluded with w' computes as

$$P(w'|w) = \sum_{F \in \mathcal{F}, ww' = T(F)\vartheta} P(\vartheta|F, w)P(F|w),$$

i.e. the probability that some F is enacted under a specific substitution ϑ such that ww' equals the trajectory of F under ϑ . To compute the probabilities on the right-hand side of this equation, we assume that $P(\vartheta|F, w) \propto \sigma(\vartheta, F)$ in the sense that the likelihood of any substitution is proportional to its similarity to a frame as compared to that of any other substitution still possible. The probability $P(F|w)$ is computed by looking at the occurrence counter value corresponding to the last element of $T(F)$ (i.e. to $T(F)$ as a whole).

3.3 Decision Making and Frame Adaptation

Based on this probabilistic semantics, [19] defines a two-layer decision-making and learning process: at the (lower) action level, agents use utility estimates $u(w, KB)$ (which are obtained, for example, by computing the utility of physical (do) actions that occur along a dialogue sequence w under current knowledge KB and assigning a small communication cost to each “non-physical” message) to maximise the expected utility *within* the activated frame (i.e. among all substitutions that this frame still permits). This involves adversarial search in the space of variable substitutions that the agent and her peer may apply in their respective part(s) of the conversation (since values for some of these variables can be chosen by the agent and for others this is done by her peer).

At the (upper) *framing* level, which is concerned with choosing a frame to activate from a given frame repository \mathcal{F} , agents use a variant of hierarchical Q-learning (based on the *options* framework proposed in [15]) to learn optimal re-framing strategies for changing frames during an encounter if (1) the current frame trajectory no longer matches the perceived encounter message sequence, (2) frame conditions no longer apply, or (3) the frame no longer seems to offer positive utility under the optimal substitution.

The intuition behind this layered approach is that frames provide decision-making blocks for communication behaviour that help the agent distinguish between different

⁵ This assumes the agent maintains some kind of knowledge base KB and can verify $KB \models \varphi$ for any φ entailed by KB .

communication situations, e.g. different phases in a dialogue (e.g. “preference elicitation” and “proposal exchange” in automated negotiation contexts). Once an appropriate frame has been identified for the current situation (which also depends on the context conditions contained in that frame) the alternatives provided by other frames can safely be discarded, at least until the frame has been successfully executed or some problem arises and *re-framing* occurs. While no re-framing occurs, the agent activates the lower decision-making layer to make optimal decisions regarding the degrees of freedom that are provided by the currently active frame (in the form of unbound variables). Since this search space is manageable, we conduct exhaustive (adversarial) search based on maximum expected utility of all ground continuation sequences the frame caters for. This allows for communicative decision making under *bounded rationality* conditions.

The full m²InFFrA architecture contains several additional components:

- *A mechanism for frame adaptation generalisation from experience.* Since agents may deviate from pre-defined frames, frame adaptation cannot be restricted to a mere update/extension of counter values and condition/substitution sets. To allow for the adaptation of frames from actual interaction experience, [8] extends the aforementioned distance metric on message sequences to frames and interprets frames as clusters in the space of possible conversations. *Cluster validation* techniques are then used to decide whether (and how) new observations should be merged into existing frames or whether they should be used to create a new frame.
- *Heuristics for making decisions about deviating from existing frames.* Often, agents would prefer to deviate from the currently executed frame because it does not seem desirable under changed environmental circumstances (e.g. by not executing costly physical actions that are part of the trajectory). However, this would jeopardize the long-term stability of the whole system of frames and trust in their use and therefore we have developed heuristics to facilitate an explicit trade-off between local desirability and global predictability of communication processes [20].
- *Methods for deriving encounter state abstractions.* At the “upper” frame selection decision-making level, agents must base their choices on the current “communication state”, which has to be modelled in a different way from the general state of affairs (which is essentially described by all facts in the agent’s knowledge base *KB*). In [19], we describe methods for deriving such abstractions of the general state of the world by focusing on those aspects of it that are relevant to the current dialogue, such as the role of the agent in it, the subject of the dialogue and its projected effect(s).

For lack of space, we will omit these details of the architecture here. For our purposes, it shall suffice to note that m²InFFrA provides a framework for decision-theoretic (boundedly) rational selection and long-term adaptation of dialogue patterns in the form of simple interaction frames. This is achieved by providing agents with an initial set of admissible patterns, to which they will add their experiences over time. Based on similarity considerations and long-term accumulation of feedback regarding the usefulness of different frames in different interaction situations, they can optimise their frame and action choices. In the following, we will use this architecture for learning and decision-making in a complex social context, namely that of interest-based negotiation.

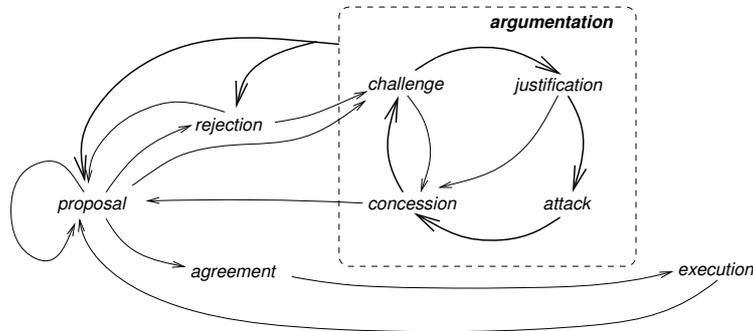


Fig. 2. Basic control flow of interest-based negotiation

4 Interest-Based Negotiation

In contrast to *proposal-based* negotiation, in which agents merely exchange proposals (such as contracts in contract nets, deals in bargaining or goods and prices in auctions), *argumentation-based* negotiation [16] allows agents to exchange information about their internal state in order to convince the other that a particular course of joint action will be mutually beneficial. *Interest-based negotiation* (IBN) [17] is a particular ABN framework which allows agents to argue over each others' beliefs, goals, and the means for achieving these goals. In this paper, we are concerned with a simplified variant of IBN in which an agents' proposal may be (1) challenged by asking for reasons (in terms of the agent's beliefs, goals, etc.) that lead to her negotiation stance and (2) justified by the agent, whereupon the challenging agent may (3) attack this justification until finally the attack, if successful, leads to a (4) concession that brings the agents closer to an agreement. Figure 2 illustrates the basic control flow of the dialogues we model here. The original IBN framework described in [17] enables more flexible dialogues (e.g. involving shifts in focus during the dialogue). As a starting point and to make the simulation viable, however, we consider only a subset of these possible dialogues and explore strategy learning within this subset.⁶

The process of IBN can be seen as traversal of a so-called *goal graph* that facilitates the representation of goal hierarchies, preferences and justifications. Each node in a goal graph represents a fact or a goal, and directed links between goal nodes can be used to represent goal hierarchies. Furthermore, a link (viz set of links) leading from a fact (viz set of facts) to a goal node and labelled with an action identifier denotes that execution of this action requires the respective fact to be true, and contributes to the respective goal. In terms of the model presented in section 2, a goal graph constructed from the arguments put forward by a specific agent can serve as a representation of this agent's public identity. An example of a goal graph for a particular domain will be given in section 5.

What makes IBN attractive for the study of argumentation strategy learning and the reason why IBN lends itself well to an implementation in $m^2InFFrA$ is that it provides

⁶ In the remainder of the paper, we will use the term IBN to refer to this simpler version of interest-based negotiation.

Table 1. Example frames for single-shot IBN

$$\begin{aligned}
& \langle \langle \rightarrow \text{request}(A, B, X) \rightarrow \text{reject}(B, A, X) \rightarrow \text{ask-reason}(A, B, X) \\
& \quad \rightarrow \text{inf-problem}(B, A, P) \rightarrow \text{concede}(A, B, P) \rangle, \\
& \langle \{ \text{problem}(P, X) \} \rangle, \langle \langle \rangle \rangle \rangle \\
& \\
& \langle \langle \rightarrow \text{request}(A, B, X) \rightarrow \text{ask-reason}(B, A, X) \rightarrow \text{inf-goal}(A, B, G) \\
& \quad \rightarrow \text{att-means}(B, A, Y) \rightarrow \text{concede}(A, B, Y) \rightarrow \text{do}(B, Y) \rangle, \\
& \langle \{ \text{goal}(A, G), \text{achieves}(X, G), \text{achieves}(Y, G), X \neq Y \} \rangle, \langle \langle \rangle \rangle \rangle
\end{aligned}$$

a rich set of social rules with which agents have to comply when engaging in rational argumentation. The most prominent of these are:

1. No proposal can be considered viable if it cannot be implemented under current circumstances, i.e. if its environmental preconditions are not met. If an agent is informed (believably) that her proposal rests on false assumptions, she must withdraw it.
2. A proposal has to be dropped if it can be shown that its effects have already been achieved or that they are unachievable.
3. No proposal is acceptable that violates a higher-level goal even though it achieves some lower-level goal. In fact, to make things more difficult we will require that a proposal that jeopardises *any* goal will be considered unacceptable in our experiments.
4. No alternative to a proposal can be rejected once the fact has been accepted that it will achieve the alleged goal (i.e. if it achieves the same thing, there is no reason to reject an alternative).

Note that in this list, “goals” always refer to the agent’s *own* goals, i.e. we do not assume any “collective rationality” that would force the agent to justify her stances with respect to a global set of goals. We rather assume that the public identity of the agent is described by a goal and belief structure that the *agent* is supposed to have, and in communication she has to act in accordance with this purported goal structure.

Quite interestingly, while these rules are based on principles of agent-level rationality (some of them in fact reflect fundamental elements of BDI theory [4]), in an argumentation scenario they constitute society-level rules of communicative behaviour: Any agent who violates them would no longer be treated as rational by others and might be excluded from the society altogether (simply for lacking the ability to participate in reasonable communication).

In the context of $m^2\text{InFFrA}$, these rules can be used directly to define argumentation frames. Consider the two frames for single-shot⁷ IBN quoted from [19] and shown in table 1, which implement rules 1 and 4, respectively. In the first frame, *B* justifies her

⁷ I.e., involving only one iteration of the challenge-justification-attack-concession loop shown in figure 2. For lack of space, the frames developed in [19] for iterative IBN are omitted here.

refusal to perform the proposed action X by pointing to a problem P that inhibits execution of X . In the second frame, B attacks A 's justification for action X (namely a goal G achieved by it) with an alternative action Y that achieves G as well. The logical predicates *problem*, *goal* and *achieves* in these examples (their meaning should be obvious from the context) refer to knowledge states of the individual agents. While it may be fairly easy for both agents to check specific instances of *problem* and *achieves* (e.g. by “inspecting” the environment), this is certainly not the case for the *goal* predicate. However, this statement still has to be consistent with the public identity of agent A as given by her past (and future) statements. As a result, A cannot attack the alternative means Y for achieving G , independent of the fact if she really holds G as a goal.⁸

To allow for the exchange of multiple arguments, [19] further defines six frames for iterative IBN, corresponding to a successful proposal, challenged proposal, and rejected proposal (i.e. edges leading out of the proposal node of figure 2) and to successful challenge, successful justification, and successful attack (i.e. edges leading into the concession node). Using these frames in practice (or, more precisely, using an indefinitely long sequence of frames in a single encounter) requires a more complex control flow than that currently possible in $m^2lnFFrA$ (e.g. storing the state of a particular argument or proposal when a shift in focus occurs). While beyond the scope of this paper, this certainly is on our research agenda in order to increase the expressiveness and flexibility of our implementation.

The overall workings of a society of IBN- $m^2lnFFrA$ agents in the experiments that we report on below are as follows:

- We equip all agents in the society with a set of (identical) interaction frames that enable them to conduct an IBN process as depicted in 2). Initially, all counters in these frames are set to 0 and substitution sets are empty.
- We construct a goal graph which can be inspected by any agent and that reflects the goal structure of a rational agent in the domain. In principle, agents could describe their internal (alleged) goal structures to each other through discussions from scratch, but we assume a commonly known goal structure to simplify things because we are only interested in rational argumentation *given* some publicised goal structure.
- Although this is not required in the general $m^2lnFFrA$ architecture, we force agents to adhere to these frames, i.e. once activated, agents cannot deviate from them unless in ways permitted by the initial . In other words, we require agents to obey the overall communicative regime, which will in this case force them to concede to anything that follows from their assumed goal structure and beliefs. As a consequence, the strategic choices of agents are restricted to *which* of the currently matching frames to use *when* and with which concrete values for unbound variables in that frame in any given step.
- As interaction unfolds, agents will adapt their frame repositories according to observations and attempt to optimise their long-term strategy using the hierarchial

⁸ For reasons of simplicity, we suffice with this, somewhat naive, approach for handling the complex notion of *commitment* in dialogue [25]. An elaborate account of commitment management is beyond the scope of this paper.

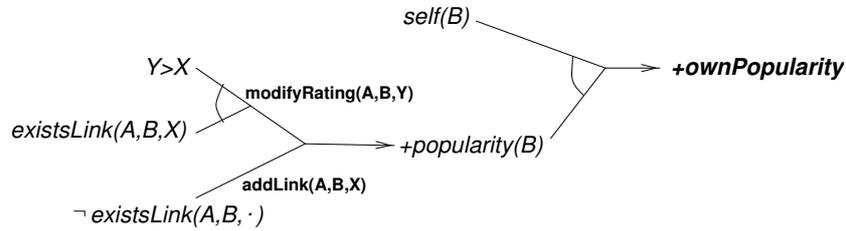


Fig. 3. Part of the LIESON goal graph for the (sub)goal of increasing one’s own popularity (by either obtaining an in-link with arbitrary rating value or increasing an having an existing link’s rating value X be increased to Y), which itself contributes to the (super)goal of increasing one’s score

learning and optimisation process described for $m^2InFFrA$ above. They will be solely judged by the “physical” utilities they obtain from the environment, i.e. no genuine, immediate gain can be derived from communication itself (other than a small communicative (negative) cost).

5 Experimental Results

As a proof of concept, interest-based negotiation using interaction frames has been implemented and tested in the multiagent-based link exchange system LIESON. In this system, agents representing Web sites engage in communication to negotiate over mutual linkage with the end of increasing the popularity of one’s own site and that of other preferred sites.

Available physical actions in this domain are the addition and deletion of numerically rated links originating from one’s own site and the modification of these ratings (where the probability of attracting traffic through a link depends on the rating value). Agent performance is computed based on the flow through the link network as well as on private ratings the agents hold towards each other. It is worth noticing that these private ratings also introduce a form of “social standing”, since linkage decisions by higher rated agents have a greater impact on individual as well as overall performance. Figure 3 shows a part of the goal graph for LIESON.

Technically, LIESON agents consist of a non-social BDI [18] reasoning kernel that projects future link network configurations and prioritises goals according to utility considerations. If these goals involve actions that have to be executed by other agents, the $m^2InFFrA$ component starts a framing process which runs until the goal of communication has been achieved or no adequate frame can be found. Note that the goal-prioritising internal reasoning mechanism need not generate intentions that are in keeping with the goal graphs talked about in conversations with other agents. This is exactly what is meant by strategically exploiting the possibilities of public identity management while pursuing one’s private agenda.

We report on two different sets of experiments in order to compare the performance of simple proposal-based negotiation (PBN) to that of (single-shot) IBN. PBN has been implemented by supplying agents with a set of frames that allow for requesting action

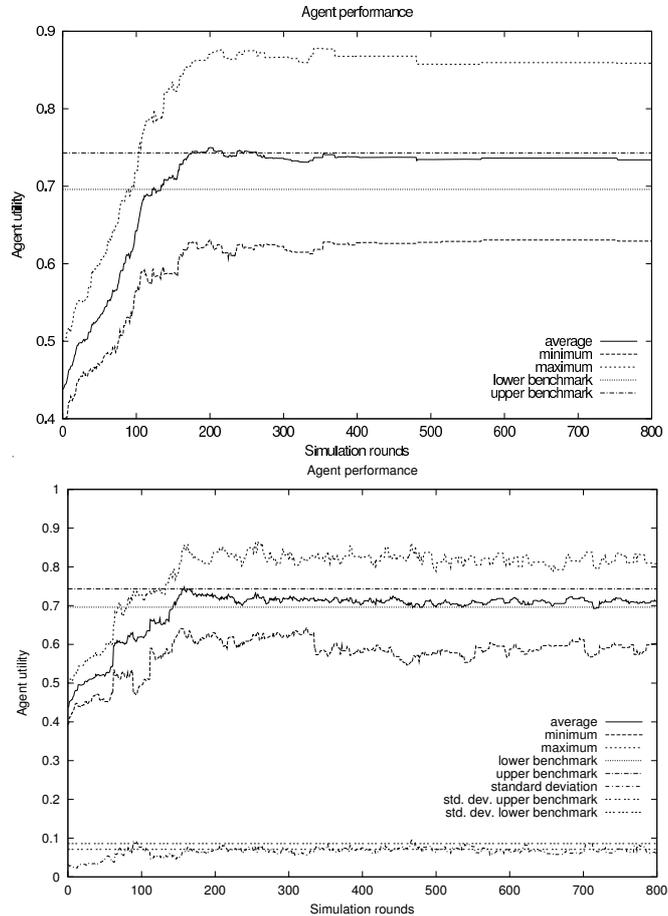


Fig. 4. Performance plots for proposal-based (above) vs. interest-based negotiation (below)

execution from another agent, proposing alternative actions, or proposing actions the other has to perform in return for one's cooperation. In contrast to the IBN case, agents are free to perform or not perform these actions without giving a reason. Single-shot IBN is realised using a set of frames, one for each path in the graph of figure 2 (two examples for these were given in the previous section, for exact definitions of the remaining frames cf. [19]). As compared to proposal-based negotiation, IBN enforces a much stricter *communication regime* by requiring agents to justify their stance, to accept any alternative suggested for the same goal, to abandon any proposal that threatens at least one goal, etc. In particular, this also implies that agents cannot simply reject a proposal because it does not seem desirable in terms of utility.

Figure 4 shows the average agent performance (in a society of ten agents) as well as the individual performances of the best and worst agent for PBN (above) and IBN (below), both of them averaged over 50 independent runs. The constant lines depicted

in the plots correspond to benchmark values that are relevant in the linkage domain: the lower benchmark corresponds to the average utility that can be achieved if every agent “honestly” expresses her opinion towards any other agent by laying a link that is weighted with her actual rating for the target site, while the upper benchmark is reached if every agent lays “politically correct” links to all other agents by not laying any links with negative ratings (i.e. concealing any critical views of other sites). Note that to start off with, agents know nothing about these benchmarks and link configurations that will yield high utility scores to them. In particular, laying more and more links in an honest way is only slightly dominated by the strategically superior, politically correct linkage pattern and it is quite impressive that agents achieve a performance close to the upper benchmark.

The significance of the results shown in figure 4 is that the agents manage to attain (and maintain) a reasonable level of long-term utility even under these stricter – and much more realistic – circumstances (albeit with bigger fluctuations indicating frequent “loss of an argument”). This illustrates nicely that $m^2InFFrA$ is capable of combining decision-theoretic learning with complex knowledge-based reasoning about constraint-governed conversation patterns. The ability to record experiences with certain communication patterns (by extending the pre-specified negotiation frames with new substitutions and situation-dependent conditions) and to reinforce their use depending on the environmental feedback obtained while using them in a particular interaction allows agents to adapt not only to a set of communication patterns but in fact to the (evolving) *communication practice* of a MAS.

The results also suggest that IBN has an *equilibratory* effect on the social outcome since the utility difference between most and least successful agents is somewhat smaller than is the case for proposal-based negotiation. Quite naturally, the requirement to “give reasons” (and hence to act rationally in accordance with public identity) seems to reduce the impact of “having more power”. Indeed, a closer look at individual interactions reveals that agents are capable of “winning an argument” independent of their power and that of their peer.

6 Conclusions

In this paper, we have presented a practical, adaptive approach to argumentation for artificial agents. Starting with a brief discussion of the general issue of strategic interaction, we have argued for a separation between *internal reasoning* and *social behaviour* to allow for the combination of the decision-theoretic design of a rational, self-interested agent with the prescriptive, society-level constraints entailed by typical argumentation protocols.

We have introduced the abstract architecture $InFFrA$ and the notions of *interaction frames* as a representation of a class of interaction in terms of surface structure and contextual conditions and of *framing* as the process of strategically applying a set of interaction frames and adapting them from experience.

We have further presented a simplified yet flexible version of interest based negotiation, which allows for rejecting, challenging, justifying and attacking arguments in the form of agents’ mental states. This version has then been implemented using a par-

ticular instance of InFFrA for two-party, turn-taking conversations. The feasibility of the approach has been illustrated in a agent-based web linkage scenario, and its performance has been shown to be comparable to that of simple proposal-based negotiation while imposing much stricter and much more realistic constraints.

To our knowledge, the implementation of IBN frames in m^2 InFFrA constitutes both the first practical approach and the first application of machine learning methods to argumentation-based negotiation in MASs. This example illustrates that the combination of logical constraints (that can be used to describe knowledge-level or social-level communication semantics) and probabilistic models of communication processes (which allow for an application of decision-theoretic learning and optimisation methods) make m^2 InFFrA a prime candidate for achieving rational agent behaviour in other, similarly complex communication contexts that are defined by modern ACLs and interaction protocols.

In the future, we would like to extend frame representations to enable more complex communication constraints and capture protocol information beyond simple turn-taking message sequences, in particular by allowing cycles, branching and multi-party dialogues. Also, we are interested in “context mining” for frame conditions, i.e. the automated discovery of those aspects in the context that are responsible for the success (or failure) of a frame. Finally, as frames suggest the combination of symbolic communication and constraints with it would be interesting to integrate interaction frames with existing *relational reinforcement learning* [24] methods.

References

1. L. Amgoud and N. Maudet. Strategic considerations for argumentative agents (preliminary report). In S. Benferhat and E. Giunchiglia, editors, *Proceedings of the 9th International Workshop on Non-Monotonic Reasoning (NMR 2002): Special session on Argument, Dialogue and Decision*, pages 399–407, 2002.
2. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, pages 338–342, Amsterdam, Netherlands, 2000. IOS Press.
3. K. Atkinson, T. Bench-Capon, and P. McBurney. A dialogue game protocol for multi-agent argument over proposals for action. In Iyad Rahwan, Pavlos Moraitis, and Chris Reed, editors, *Proceedings of the First International Workshop on Argumentation in Multi-Agent Systems (ArgMAS'04)*, volume 3366 of *Lecture Notes in Computer Science*. Springer Verlag, Berlin, Germany, 2005.
4. M. E. Bratman. *Intentions, Plans and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.
5. P. R. Cohen and H. J. Levesque. Communicative actions for artificial agents. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS)*, pages 65–72, 1995.
6. V. Dignum. *A model for organizational interaction: based on agents, founded in logic*. PhD thesis, Utrecht University, The Netherlands, 2004.
7. F. Fischer and M. Rovatsos. Reasoning about communication: A practical approach based on empirical semantics. In *Proceedings of the 8th International Workshop on Cooperative Information Agents (CIA)*, volume 3191 of *Lecture Notes in Artificial Intelligence*, Erfurt, Germany, 2004. Springer-Verlag.

8. F. Fischer, M. Rovatsos, and G. Weiß. Acquiring and adapting probabilistic models of agent conversation. In *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Utrecht, The Netherlands, 2005.
9. N. Fornara and M. Colombetti. Operational specification of a commitment-based agent communication language. In Maria Gini, Toru Ishida, Cristiano Castelfranchi, and W. Lewis Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 536–542, Bologna, Italy, 2002. ACM Press.
10. A. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 883–890, Melbourne, Australia, 2003.
11. M. T. Kone, A. Shimazu, and T. Nakajima. The state of the art in agent communication languages. *Knowledge and Information Systems*, 2:259–284, 2000.
12. P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11:113–131, 1995.
13. P. McBurney. *Rational Interaction*. PhD thesis, University of Liverpool, 2002.
14. S. Parsons, M. J. Wooldridge, and L. Amgoud. Properties and complexity of formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.
15. D. Precup. *Temporal Abstraction in Reinforcement Learning*. PhD thesis, Department of Computer Science, University of Massachusetts, Amherst, 2000.
16. I. Rahwan, S. D. Ramchurn, N. R. Jennings, P. McBurney, S. Parsons, and L. Sonenberg. Argumentation-based negotiation. *Knowledge Engineering Review*, 18(4), 2003.
17. I. Rahwan. *Interest-based Negotiation in Multi-Agent Systems*. PhD thesis, Department of Information Systems, University of Melbourne, Melbourne, Australia, 2004.
18. A. S. Rao and M. P. Georgeff. An abstract architecture for rational agents. In W. Swartout C. Rich and B. Nebel, editors, *Proceedings of Knowledge Representation and Reasoning (KR&R)*, pages 439–449, 1992.
19. M. Rovatsos. *Computational Interaction Frames*. PhD thesis, Department of Informatics, Technical University of Munich, 2004.
20. M. Rovatsos, M. Nickles, and G. Weiss. Interaction is meaning: A new model for communication in open systems. In *Proceedings of the 2nd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2003.
21. M. Rovatsos, G. Weiß, and M. Wolf. An Approach to the Analysis and Design of Multiagent Systems based on Interaction Frames. In Maria Gini, Toru Ishida, Cristiano Castelfranchi, and W. Lewis Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Bologna, Italy, 2002. ACM Press.
22. M. P. Singh. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence*, 8(1–2):47–71, 1993.
23. M. P. Singh. A social semantics for agent communication languages. In *Proceedings of the IJCAI Workshop on Agent Communication Languages*, 2000.
24. P. Tadepalli, R. Givan, and K. Driessens. Relational Reinforcement Learning: An Overview. In *Proceedings of the Workshop on Relational Reinforcement Learning*, Banff, Alberta, Canada, 2004.
25. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany NY, USA, 1995.