A previous version of this article has been published in ECAI-06 proceedings and contained an oversight. (P. 3, column 2: a formula in IE condition of Inform act was missed and has been corrected in this new version of the paper.)

A New Semantics for the FIPA Agent Communication Language based on Social Attitudes

Benoit Gaudou [1] and **Andreas Herzig** [1] and **Dominique Longin** [1] and **Matthias Nickles** [2]

Abstract. One of the most important aspects of the research on agent interaction is the definition of *agent communication languages* (ACLs), and the specification of a proper formal semantics of such languages is a crucial prerequisite for the usefulness and acceptance of artificial agency. Nevertheless, those ACLs which are still mostly used, especially the standard FIPA-ACL, have a communication act semantics in terms of the participating agents' *mental attitudes* (viz. beliefs and intentions), which are in general undeterminable from an external point of view due to agent autonomy. In contrast, semantics of ACLs based on *commitments* are fully verifiable, but not sufficiently formalized and understood yet. In order to overcome this situation, we propose a FIPA-ACL semantics which is fully verifiable, fully formalized, lean and easily applicable. It is based on *social attitudes* represented using a *logic of grounding* in straightforward extension of the BDI agent model.

1 Introduction

The design of agent communication languages (ACLs) has attracted a lot of attention during the last years. Such languages are mostly based on Searle and Vanderveken's speech act theory [9], and are not only relevant for applications involving real software agents or robots, but also for other software entities which need to communicate, like web services.

Among the different existing ACLs, FIPA-ACL is still the most important standard, subsets of which are widely used in agent interaction protocols. FIPA-ACL is semantically rich, and the concepts involved are quite intuitive.

Nevertheless, FIPA-ACL has a feature that has often been criticized in the literature, viz. that the semantics of communication acts (CAs) is defined in terms of the agents' mental states. For example, when agent i informs agent j that φ, then the (rational) effect is that agent j starts to believe φ. In order for such an effect to obtain some hypotheses have to be made; but even in such contexts j is autonomous and might not adopt φ, and in any case i or other agents and the system designer can never verify whether this is the case or not. This is especially felt as being too strong in open environments with black- or gray-box agents where we don't even want to ascribe mental attitudes to other agents.

In contrast, those semantics based on the concept of social commitments [10] is verifiable because they are only based on what has been communicated and the commitments the agents have made by doing that (instead of the beliefs and intentions that are "behind" these commitments and that have caused them). The drawback here is that the existing approaches are only semi-formal, in particular because there is no consensus on what "being committed" actually means. As a consequence, they are rarely used in practice up to now.

The aim of this paper is to resolve the problems of FIPA's CA semantics without loosing its benefits. We propose a novel semantics avoiding the strong hypotheses of the original semantics by "lifting" the BDI-based FIPA semantics to the social level. We do so by replacing the usual private mental attitudes of BDI logics by *public* mental attitudes, i.e. attitudes that have been made public through communication (*social attitudes*). More precisely, our semantics is based on an unified and extended approach to the concept of *communication attitudes (ostensible beliefs and intentions)* [7] and the more or less equivalent concept of *grounding*[3] [5]. For example, the effect of an informing-that-p act is that it is public that the sender believes that p. This does not mean that the sender really believes that p, but only hinders him to subsequently inform that $\neg p$, or to inform that he ignores whether p.

The major benefits of our new semantics are the following:

- It is verifiable, and suitable even for truly autonomous, possibly malevolent agents.
- It is fully formalized.
- It is based on a straightforward extension of BDI, and therefore relatively lightweight.
- It can easily be adapted to similar ACLs, e.g. the widely used KQML/KIF.
- It generalizes the single-addressee FIPA acts to groups of agents, and it distinguishes the group of addressees from the group of bystanders (overhearers), and thus refines FIPA's acts.

All in all, we aim at an agent communication semantics that eliminates the major shortcomings of the still predominant mentalist approaches to ACL semantics while being "upward compatible" to the standard FIPA semantics and similar approaches, lean, and formally well founded.

The remainder of this paper is organized as follows: The next section provides a short account of the logical framework that we have chosen as a formal foundation of our approach. Section 3 presents the new semantics, and Section 4 illustrates our approach by means of a case study. Section 5 concludes.

2 A logic of grounding

In this section we briefly present the logic of belief, intention, action, and grounding defined in [5], that is based on Cohen and

[1] Université Paul Sabatier, IRIT, Toulouse, email: gaudou,herzig,longin@irit.fr
[2] Technical University of Munich, email: nickles@informatik.tu-muenchen.de

[3] We use the term *grounding* as Traum [11], *i.e.* it refers to "the process of adding to the common ground between conversational agent".

Levesque's [2]. AGT is a finite set of agents, ACT is a set of actions, $ATM = \{p, q, ...\}$ is the set of atomic formulas. Complex formulas are denoted by $\varphi, \psi \ldots$ A model $\mathcal{M}$ is a 5-tuple that is made up of: a set of possible worlds W; a mapping

$$\mathcal{V} : W \longrightarrow (ATM \longrightarrow \{0,1\})$$

associating a valuation $\mathcal{V}_w$ to every $w \in W$; a mapping

$$\mathcal{A} : ACT \longrightarrow (W \longrightarrow 2^W)$$

associating actions $\alpha \in ACT$ and worlds $w \in W$ with the set of worlds resulting from the execution of α in w; a mapping

$$\mathcal{G} : (2^{AGT} \setminus \emptyset) \longrightarrow (W \longrightarrow 2^W)$$

associating sets of agents $I \subseteq AGT$ and worlds $w \in W$ with the set of worlds that are publicly possible for the group I at w (the worlds that are compatible with what has been uttered in I's presence); and finally the mapping

$$\mathcal{I} : AGT \longrightarrow (W \longrightarrow 2^{2^W})$$

associating every $i \in AGT$ and world w with the set of propositions (alias sets of worlds) that are intended by i. (The $\mathcal{I}_i$ are neighborhood functions in Chellas' sense [1].)

The logical language contains modal operators of action $After_\alpha$ and $Before_\alpha$, for every $\alpha \in ACT$, modal operators of groundedness G_I for every group I, and modal operators of intention Int_i for every agent $i \in AGT$.

The formula $After_\alpha \varphi$ reads "φ is true after every execution of the action α", and $Before_\alpha \varphi$ reads "φ is true before every execution of the action α". Semantically, $w \Vdash After_\alpha \varphi$ iff $w' \Vdash \varphi$ for each $w' \in \mathcal{A}_\alpha(w)$, and $w \Vdash Before_\alpha \varphi$ iff $w' \Vdash \varphi$ for each w' such that $w \in \mathcal{A}_\alpha(w')$. The logic of $After_\alpha$ and $Before_\alpha$ is the tense logic K^t, i.e. standard normal modal logic K plus the conversion axioms $\varphi \to Before_\alpha \neg After_\alpha \neg\varphi$ and $\varphi \to After_\alpha \neg Before_\alpha \neg\varphi$. The abbreviation $Done_\alpha \varphi \stackrel{def}{=} \neg Before_\alpha \neg\varphi$ reads "α has just been done before which φ was true". We note $Done(\alpha) \stackrel{def}{=} Done_\alpha \top$ for convenience. Moreover, $Before_{\alpha\cup\alpha'} \varphi$ abbreviates $Before_\alpha \varphi \wedge Before_{\alpha'} \varphi$. (Hence $Done(\alpha\cup\alpha')$ stands for $Done(\alpha) \vee Done(\alpha')$.)

$G_I \varphi$ reads "it is grounded for group I that φ is true", or for short: "φ is grounded for I". When I is a singleton, $G_{\{i\}} \varphi$ means that for agent i, φ is grounded. In this (and only in this) degenerated case 'public' grounding is the same as private belief. We write $G_i \varphi$ for $G_{\{i\}} \varphi$. The accessibility relations of grounding operators must satisfy the constraints for the standard normal modal logic KD (seriality), plus the following, for groups I, I' such that $I' \subseteq I$:

(i) if $u\mathcal{G}_{I'}v$ and $v\mathcal{G}_I w$ then $u\mathcal{G}_I w$
(ii) if $u\mathcal{G}_{I'}v$ and $u\mathcal{G}_I w$ then $v G_I w$
(iii) if $u\mathcal{G}_I v$ and $v\mathcal{G}_{I'} w_1$ then there is w_2 such that $u\mathcal{G}_I w_2$ and $V(w_1) = V(w_2)$
(iv) $\mathcal{G}_I \subseteq \bigcup_{i \in I} \mathcal{G}_I \circ \mathcal{G}_i$

Constraint (i) stipulates that subgroups are aware of what is grounded in the group: whenever w is a world for which it is grounded for I' that all I-grounded propositions hold in w, then all I-grounded propositions indeed hold in w. This is a kind of *attention* property: each subgroup participating in a conversation is aware of what is grounded in the group. Similarly (ii) expresses that subgroups are aware of what is ungrounded in the group, too. (i) and (ii) correspond to the axioms of strong rationality (SR_+) and (SR_-):

$$G_I \varphi \to G_{I'} G_I \varphi \qquad (SR_+)$$

$$\neg G_I \varphi \to G_{I'} \neg G_I \varphi \qquad (SR_-)$$

which express that if a proposition φ is grounded (resp. ungrounded) for a group I then it is grounded for each subgroup that φ is grounded (resp. ungrounded) for I[4].

(iii) stipulates that for every objective proposition grounded for I it is publicly established for I that each subgroup of I is grounded on it (which does not imply that it is grounded for the latter): whenever w is a world for which all propositions grounded for I' are grounded for I, then all those propositions are indeed grounded for I in w. It validates the axiom (WR)

$$G_I \varphi \to G_I G_{I'} \varphi, \text{ for } \varphi \text{ objective} \qquad \text{(WR)}$$

which says that if the objective formula φ is grounded for a group K then it is necessarily grounded for K that for each subgroup K' the formula is grounded.[5] Note that this does not imply that for every subgroup φ is actually grounded, *i.e.* (WR) does not entail $G_K \varphi \to G_{K'} \varphi$. In particular, the fact that φ is grounded for group K does not imply that the members of K believe that φ.

(iv) expresses that if it is grounded for a set I that a proposition is established for every agent then it is grounded for I, too. This corresponds to axiom (CG)

$$(\bigwedge_{i \in I} G_I G_i \varphi) \to G_I \varphi \qquad \text{(CG)}$$

which says that if a proposition is established for every agent in I, then it is established for the whole group I. Together, (WR) and (CG) stipulate that for objective φ we have $(\bigwedge_{k \in K} G_K G_k \varphi) \leftrightarrow G_K \varphi$. Note that $G_K \varphi$ does NOT imply $G_k \varphi$ where $k \in K$. Indeed, a proposition can be grounded in a group independently of the private belief of each agent of the group about this proposition: there is thus no sincerity hypothesis.

$Int_i \varphi$ reads "agent i intends that φ be true". The Int_i are non-normal modal operators which only validate the rule of equivalence: $\frac{\varphi \leftrightarrow \psi}{Int_i \varphi \leftrightarrow Int_i \psi}$. They neither validate $Int_i (\varphi \wedge \varphi') \to (Int_i \varphi \wedge Int_i \varphi')$ nor $(Int_i \varphi \wedge Int_i \varphi') \to Int_i (\varphi \wedge \varphi')$.

Intentions and actions are related by the principle of intentional action saying that if α has just been performed by agent i then i had the intention to do so immediately before.

$$Before_{i:\alpha} Int_i Done(i{:}\alpha) \qquad \text{(IA)}$$

where $i{:}\alpha$ denotes that action α is performed by agent i.

To highlight our proposal for the semantics of grounding consider the following example. There are three agents $AGT = \{0, 1, 2\}$. Let agent 0 (privately) believe that 2 sells high-quality products, formally written $G_0 q_2$. Now suppose that in private conversation agent 0 tells 1 that the contrary is the case (for example to trigger some attitude of 1 that benefits 0). The (illocutionary) effect is $G_{\{0,1\}} G_0 \neg q_2$. Then agent 2 joins in the conversation, and later on 0 informs 1 and 2 that q_2: The illocutionary effect is $G_{\{0,1,2\}} G_0 q_2$. This illustrates that even for nested groups $\{0\} \subset \{0,1\} \subset \{0,1,2\}$, mutually inconsistent states of public group belief might hold simultaneously.

3 Communication act semantics

Following and extending the ACL syntax used in [4], a single communication act (CA) is denoted as $\langle i, \mathsf{ActName}(J, \varphi), K\rangle$, where i

[4] In particular, we have the modal axioms (4) and (5) for G_I operators as theorems of our logic.

[5] (WR) concerns only objective formulas, i.e. formula that does not contain any modality. If we applied (WR) to some mental states, we would restrict the agents' autonomy. For example, when an agent performs the speech act $\langle i, \mathsf{Inform}(J, p), K\rangle$, he expresses publicly that he believes p. Thus if agent i expresses: $\langle i, \mathsf{Inform}(J, G_J p), K\rangle$ the formula $G_K G_i G_J p$ holds, and the agents $j \in J$ cannot afterwards express that they believe $\neg p$. If he made this speech acts, the formulae $G_K G_J \neg p$ and, thanks to (WR), $G_K G_i G_J \neg p$ would hold, which is inconsistent with the above formula $G_K G_i G_J p$.

is the performing agent, J is a group of recipients (whereas FIPA only allows one addressee). *ActName* is the name of the act (in our model not necessarily corresponding to exactly one speech act type, see below). φ is the propositional content of the act. K, which is missing in FIPA, denotes a group of attending agents who overhear the respective utterance, with $i \in K$, $J \subseteq K \setminus \{i\}$ and $J \neq \emptyset$. For a dialogue of only two agents i and j we have $J = \{j\}$ and $K = \{i, j\}$.

In the standard semantics of FIPA CAs [4] (henceforth called FIPA-S), semantics is specified by providing the *feasibility preconditions* (FPs) and the *rational effects* (REs) of single CAs. The former denote which logical conditions need to be fulfilled in order to execute the respective act, and the latter specify which conditions hold after the successful performance of that act. FPs characterize both the ability of the speaker to perform the act and the context-dependent relevance of the act (i.e., that performing the act is relevant given a certain dialogue context). In contrast, REs specify the desired and rationally-expectable direct perlocutionary effect of the utterance, i.e. what becomes true in case the perlocutionary act succeeds.

We think there are at least three reasons not to qualify a CA by its rational effect. Firstly, it is possible to desire and expect different kinds of RE of the same CA; secondly, Searle shows in [9, Sec. 2.1] that the effect of a speech act cannot be a rational (or perlocutionary) effect simply because a lot of speech acts just do not have any perlocutionary effect. He also shows that even if a speech act can have a perlocutionary effect, we can always exhibit a context where the speaker does not intend this perlocutionary effect. Thirdly, strong hypotheses (such as sincerity, competence, credibility…) must be made about the agents to enable the inference of the expected RE, which is too restrictive in our context of open multi-agent systems, possibly with conflicts and malevolent, egocentric agents...

In contrast to FIPA-S, the FPs and IEs (for illocutionary effects) in our model do not make any statement about mental attitudes, but specify the preconditions and effects in terms of groundings of group K (the public, so to say). They are chosen such that the respective communication act is both executable given all realistic preconditions, and succeeds reliably with a publicly verifiable effect. The only (self-evident) exception follows from the bridge axioms (SR_+) and (SR_-) given in the previous section, stating that an agent or subgroup of a certain group knows about the existence of the respective grounded beliefs or intentions of their group — this means merely that the agents keep track of the ongoing course of communication in terms of FPs and IEs.

In the sequel we use the term *Social Attitudes Based Semantics* (SABS) for our modelling, and will define the SABS semantics of the four *primitive CAs* of FIPA-ACL: *Inform*, *Request*, *Confirm* and *Disconfirm*, and we will also present the respective FIPA-S specifications for comparison. All other FIPA-CAs are macros composed of these primitives in a more or less straightforward manner.

3.1 Inform: Asserting information

We start with the FIPA-S version of the semantics:

$\langle i, \mathsf{Inform}_{\mathsf{FIPA}}(j, \varphi), K\rangle$

FP: $Bel_i\, \varphi \wedge \neg Bel_i\, (Bel_j\, \varphi \vee Bel_j\, \neg\varphi \vee U_j\varphi \vee U_j\neg\varphi)$

RE: $Bel_j\, \varphi$

At this, $U_j\varphi$ denotes that agent j is uncertain about φ, but thinks that φ is more likely than $\neg\varphi$. The terms "uncertain" and "more likely" are not precisely defined in [4]. The essential preconditions of *Inform* in the FIPA-S are thus that agent i truthfully believes what he asserts, and that the receiver does not have any definite opinion about the asserted proposition. Whereas the former condition is obviously unrealistic given a truly autonomous agent i, the latter disallows (problematically) the use of *Inform* to convince the addressee. We consider the latter usage as crucial e.g. in the context of computational argumentation and argumentation-based negotiation. We could introduce an additional conviction-act extending the syntax, or try to emulate it with a construct like $Request(Inform(\varphi))$, but this would not only be unnecessary and inelegant, but would also blur the fact that there exists a continual transition from "pure information" to conviction. It is also not clear why the absence of an opinion shall be a realistic precondition for the success of an information act, or, conversely, why the existence of an opinion (which could be very weak, or "by default" only) shall hinder the receiver to adopt the asserted information (e.g., consider that the addressee might trust the sender more than herself).

The rational effect of $Inform$ in FIPA-S is simply that the addressed agent believes what she has been told (in case the act succeeds). Of course, this effect cannot be verified with autonomous agents. Even if it could be verified, it would be too strong and unlikely. Moreover it is not verifiable that the informing agent (truthfully) intends the adoption of a certain belief.

These concerns lead to the following SABS semantics:

$\langle i, \mathsf{Inform}(J, \varphi), K\rangle$

FP: $\neg G_K\, G_J\, \varphi \wedge \neg G_K\, Int_i\, G_J\, \varphi \wedge \neg G_K\, \neg G_i\, \varphi$

IE: $G_K\, G_i\, \varphi \wedge G_K\, Int_i\, G_J\, \varphi$

In the FP, $\neg G_K\, G_J\, \varphi$ specifies that the addressed agent has not expressed φ before (with group K attending), corresponding to the $\neg Bel_i\, Bel_j\, \varphi$ part of the FIPA-S FP (the *relevance precondition*). It simply expresses that asserting an information would be unnecessary if the receiver has already expressed its belief in it. However, our new FP does not demand that group J has no opinion at all about φ, allowing to use $Inform$ also to *convince* J in case this group has already expressed its disbelief in φ earlier. $\neg G_K\, Int_i\, G_J\, \varphi$ in FP effectively demands that agent i did not assert this information using an assertive communication act before, which is also an aspect of the relevance precondition. $\neg G_K\, \neg G_i\, \varphi$ ensures that the asserted opinions of agent i are mutually consistent. (The latter is a precondition of *rationality*).

In the IE, $G_K\, G_i\, \varphi$ denotes that with asserting φ, it becomes grounded that agent i believes that φ, regardless if she does so privately (i.e., mentally) or not.

As usual we define $\langle i, \mathsf{InformIf}(J, p), K\rangle$ as an abbreviation of $\langle i, \mathsf{Inform}(J, p), K\rangle \cup \langle i, \mathsf{Inform}(J, \neg p), K\rangle$. Hence $Done(\langle i, \mathsf{InformIf}(J, p), K\rangle) \equiv Done(\langle i, \mathsf{Inform}(J, p), K\rangle) \vee Done(\langle i, \mathsf{Inform}(J, \neg p), K\rangle)$.

However, in many cases, we can safely assume that group J immediately starts to publicly believe the asserted information, namely when this group apparently trusts the uttering agent in regard to this information. (A notorious exception are exam situations.) An important particular case is expressed by the following axiom, for $J \subseteq K$ and $\alpha = \langle i, \mathsf{InformIf}(J, \varphi), K\rangle$:[6]

$$(G_K\, Done_\alpha \bigwedge_{j \in J} Int_j\, Done(\alpha)) \rightarrow G_K\, G_J\, \varphi \qquad (1)$$

This specifies that if an agent has requested a certain information before from agent i in form of a closed question (like with "Is it raining

[6] We here consider that the group J have asked i to publicly declare that φ. (And not each j, as it would be the case if α was $\langle i, \mathsf{InformIf}(\{j\}, \varphi), K\rangle$ in (1).)

outside?"), it becomes grounded that she believes the answer.[7]

3.2 Request: Requesting an action to be done

Again, we state the FIPA version of the semantics first:

$\langle i, \mathsf{Request_{FIPA}}(j, \alpha), K\rangle$

FP: $FP(\alpha)[i\backslash j] \wedge Bel_i\, Agent(j, \alpha) \wedge Bel_i\, \neg PG_j\, Done(\alpha)$

RE: $Done(\alpha)$

Here, α is an action expression, $FP(\alpha)[i\backslash j]$ denotes the part of the feasibility preconditions of action α where the mental attitudes are those of agent i. $Agent(j, \alpha)$ states that j is the only agent that ever performs, has performed or will perform α, and $PG_j\, Done(\alpha)$ denotes that $Done(\alpha)$ (i.e., action α has just been performed successfully) is a *persistent goal* [8] of agent j. The RE just specifies that the intended perlocutionary effect of this communication act is to get α done.

Obviously, these specifications again require strong assumptions about mental properties, which are equally problematic as in the case of $Inform$. In addition, $Agent(j, \alpha)$ reduces the scope of this communication act unnecessarily, disallowing concurrent intention of j to perform the same action herself.

As in our formalism the propositional content of a CA is a formula, a request to do action α is defined as a request that $Done(\alpha)$ be true. Furthermore, in our case the addressee of a speech act is a group of agents. Thus a request is addressed to each agent of the group in the aim that either at least one agent of the group do the requested action ("open that door"), or each agent of the group do it ("clean that room", addressed to a group of children). $i : \alpha$: denotes that i is the author of action α (making superfluous the FIPA *Agent* predicate). We thus have two kinds of request (whereas there is only one in FIPA):

$$\langle i, \mathsf{RequestSome}(J, J{:}\alpha), K\rangle \overset{def}{=} \langle i, \mathsf{RequestSome}(J, \bigvee_{j\in J} Done(j{:}\alpha)), K\rangle$$

$$\text{FP: } \left(\neg G_K \bigvee_{j\in J} Int_j\, Done(j{:}\alpha)\right) \wedge \neg G_K\, \neg Int_i \left(\bigvee_{j\in J} Done(j{:}\alpha)\right)$$

$$\text{IE: } G_K\, Int_i \left(\bigvee_{j\in J} Done(j{:}\alpha)\right) \wedge G_K\, \neg G_i \left(\bigvee_{j\in J} Int_j\, Done(j{:}\alpha)\right)$$

So our FP specifies that is not grounded that at least one of the agents in J intends to achieve α already (relevance precondition), and that it is not grounded that agent i does not intend $Done(\alpha)$ (rationality precondition). The IE is also straightforward: the act results in the grounding that agent i intends that at least one agent in J intends $Done(\alpha)$ become true, and that i does not believe that one agent in J intends $Done(\alpha)$.

Second, we define:

$$\langle i, \mathsf{RequestEach}(J, J{:}\alpha), K\rangle \overset{def}{=} \langle i, \mathsf{RequestEach}(J, \bigwedge_{j\in J} Done(j{:}\alpha)), K\rangle$$

$$\text{FP: } \left(\neg G_K \bigwedge_{j\in J} Int_j\, Done(j{:}\alpha)\right) \wedge \neg G_K\, \neg Int_i \left(\bigwedge_{j\in J} Done(j{:}\alpha)\right)$$

$$\text{IE: } G_K\, Int_i \left(\bigwedge_{j\in J} Done(j{:}\alpha)\right) \wedge G_K\, \neg G_i \left(\bigwedge_{j\in J} Int_j\, Done(j{:}\alpha)\right)$$

which specifies that i intends that each agent of J perform the requested action α. For compatibility reasons, we also define

$$\langle i, \mathsf{Request}(J, \alpha), K\rangle \overset{def}{=} \langle i, \mathsf{RequestSome}(J, J{:}\alpha), K\rangle$$

FIPA also defines the acts $Confirm$ (for the confirmation of an uncertain information) and its pendant $Disconfirm$ as primitives. But since our $Inform$ semantics has an adequately weakened FP that does not require that the asserted information is not uncertain, $Confirm$ and $Disconfirm$ simply map to $Inform$ in our semantics.

4 Case study

In order to demonstrate the properties and the application of our approach, this section presents a brief case study in form of an *agent purchase negotiation* scenario. In particular, we aim to demonstrate the following crucial features of SABS, all not being present in FIPA-S or, by principle, any other BDI-based ACL semantics:

- Pre- and post-conditions of communication acts being only dependent from publicly observable agent behavior, thus being fully verifiable;
- Communication acts with contents being inconsistent with the beliefs and intentions of the participating agents;
- Communication acts addressing groups of agents;
- Multiple communication acts uttered by the same sender, but with mutually inconsistent contents (even towards nested groups);
- Persuasive Inform-acts.

In addition, the example shows how the logging of the grounding state of the negotiation dialogue can replace *commitment stores*, which are usually used to keep track of the various commitments arising during the course of an interaction (like to sell or buy a product). In contrast, by the use of our semantics we obtain the publicly available information about the state of commitment of the participating agents directly in terms of logical post-conditions of communication acts, namely publicly expressed intentions. As explained in Section 1, we consider this to be simpler and formally clear compared to the use of social commitments in the sense of [10].

The interaction roughly follows protocols for *purchase negotiation dialogue games* as known from, e.g., [6], but omitting several details of such protocols which are not relevant for our demonstrative purposes (like the specification of selling options in detail). Also, such protocols often make use of proprietary negotiation locutions, whereas we get along with FIPA-ACL constructs, since in our context, no acts not contained in FIPA-ACL (like the "Promise" and "Threaten" acts in protocols for argumentation-based negotiation) are required. Nevertheless, our scenario is clearly beyond FIPA's *contract net* specification [3].

Our scenario consists of four agents $MAS = \{s_1, s_2, b_1, b_2\}$, representing potential car sellers and customers. In the discourse universe exists two instances θ_1 and θ_2 of some car type θ (e.g., specimen of the Alfa Romeo 159).

We present now the interaction course, consisting of sequential steps in the following form. Note that the interaction course consists of multiple interlaced conversations among different sender/receiver pairs and different overhearers (i.e., different "publics" so to say). In particular, agent b_2 is involved in two selling dialogues at the same time.

Utterance no. sender→receiver: Descriptive act title
Message [8]

[7] The intention Int_j can be triggered with FIPA's *QueryIf* act. The schema would work analogously for $\langle i, \mathsf{InformIf}(\{j\}, \neg\varphi), K\rangle$.

[8] Using syntactical macros according to [4]. Only in case the message primitives are semantically relevant in our context, the respective macros are expanded.

Effect (optionally) gives the effect of the act in terms of grounded formulas, according to SABS and the axioms in Section 2 (so this may go beyond the direct IE).

In contrast, *Private information (PI)* optionally unveils relevant mental attitudes before or after an act has been uttered and understood by the respective agents. The PIs are not determined by preceding communication acts, due to agent autonomy. They are also of course usually not available to observers, and just given for explanatory purposes.

U1 $s_1 \rightarrow \{b_1, b_2\}$**: Initialize dialogue**
$\langle s_1, \mathsf{RequestEach}(\{b_1, b_2\}, enterDialogue(\theta_1)), \{s_1, b_1, b_2\}\rangle$
U2 $b_1 \rightarrow \{s_1\}$**: Enter dialogue**
$\langle b_1, \mathsf{Agree}(\{s_1\}, enterDialogue(\theta_1)), \{s_1, b_1, b_2\}\rangle$
U3 $b_2 \rightarrow \{s_1\}$**: Enter dialogue**
$\langle b_2, \mathsf{Agree}(\{s_1\}, enterDialogue(\theta_1)), \{s_1, b_1, b_2\}\rangle$
U4 $s_2 \rightarrow \{b_2\}$**: Initialize dialogue**
$\langle s_2, \mathsf{Request}(\{b_2\}, enterDialogue(\theta_2)), \{s_2, b_2\}\rangle$
U5 $b_2 \rightarrow \{s_2\}$**: Enter dialogue**
$\langle b_2, \mathsf{Agree}(\{s_2\}, enterDialogue(\theta_2)), \{s_2, b_2\}\rangle$
PI_{s_1}**:** $Bel_{s_1}\ discounts$
U6 $s_1 \rightarrow \{b_1, b_2\}$**: Information about discount**
$\langle s_1, \mathsf{Inform}(\{b_1, b_2\}, \neg discounts), \{s_1, b_1, b_2\}\rangle$
Effect:
$G_{\{s_1,b_1,b_2\}} G_{s_1} \neg discounts$
$\wedge G_{\{s_1,b_1,b_2\}} Int_{s_1} G_{\{b_1,b_2\}} \neg discount$
Seller s_1 asserts that no discounts can be given while believing (PI_{s_1}: $Bel_{s_1}\ discount$) that the opposite is true (there might be the company policy that discounts should be given, but that might reduce the seller's individual profit).
U7 $s_1 \rightarrow \{b_2\}$**: Information about discount**
$\langle s_1, \mathsf{Inform}(\{b_2\}, discounts), \{s_1, b_2\}\rangle$
Effect:
$G_{\{s_1,b_2\}} G_{s_1} discounts$
$\wedge G_{\{s_1,b_2\}} Int_{s_1} G_{b_2} discount$
While seller s_1 informed group $\{b_1, b_2\}$ that there would be no price discounts, he informs customer b_2 that this is not true (likely because s_1 thinks that b_2 is a valued customer whereas b_1 is not).
U8 $b_2 \rightarrow \{s_1\}$**: Query if car type has high accident rate**
$\langle b_2, \mathsf{Request}(\{s_1\}, InformIfAccidentRateHigh), \{s_1, b_2\}\rangle$
Effect:
$G_{\{s_1,b_2\}} Int_{b_2} Done(s_1 : InformIfAccidentRateHigh) \wedge$
$G_{\{s_1,b_2\}} \neg G_{b_2} Int_{s_1} Done(s_1 : InformIfAccidentRateHigh)$, with
$InformIfAccidentRateHigh \stackrel{def}{=}$
$\langle s_1, \mathsf{InformIf}(\{b_2\}, accidentRateHigh(\theta)), \{s_1, b_2\}\rangle$
$PI_{s_1} : Bel_{s_1}\ accidentRateHigh(\theta_1)$
U9 $s_1 \rightarrow \{b_2\}$**: Information about accident rate**
$\langle s_1, \mathsf{Inform}(\{b_2\}, \neg accidentRateHigh(\theta)), \{s_1, b_2\}\rangle$
Effect:
$G_{\{s_1,b_2\}} G_{s_1} \neg accidentRateHigh(\theta)$
$\wedge G_{\{s_1,b_2\}} G_{b_2} \neg accidentRateHigh(\theta)$
Note that due to her closed question before and axiom 1 it becomes immediately grounded that b_2 believes the asserted information. In addition, b_2 privately believes this information also (see PI_{b_2} below), but revises this later.
Seller s_1 asserted $\neg accidentRateHigh(\theta_1)$ though thinking the opposite.
$PI_{b_2} : Bel_{b_2} \neg accidentRateHigh(\theta)$
U10 $b_2 \rightarrow \{s_2\}$**: Query if car type has high accident rate**
$\langle b_2, \mathsf{Request}(\{s_2\}, InformIfAccidentRateHigh), \{s_2, b_2\}\rangle$
U11 $s_2 \rightarrow \{b_2\}$**: Information about accident-damage**
$\langle s_2, \mathsf{Inform}(\{b_2\}, accidentRateHigh(\theta)), \{s_2, b_2\}\rangle$
Again, b_2 publicly believes the information, and trusts it for some reason privately more than the information given by seller s_1 earlier. Nevertheless, $G_{\{s_1,b_2\}} G_{b_2} \neg accidentRateHigh(\theta_1)$ remains true.
$PI_{b_2} : Bel_{b_2}\ accidentRateHigh(\theta)$
U12 $b_2 \rightarrow \{s_2\}$**: Propose to buy at a certain price**
$\langle b_2, \mathsf{Propose}(\{s_2\}, buy(\theta_2, 10000\pounds)), \{s_2, b_2\}\rangle$
U13 $s_2 \rightarrow \{b_2\}$**: Accept proposal**
$\langle s_2, \mathsf{AcceptProposal}(\{b_2\}, buy(\theta_2, 10000\pounds)), \{s_2, b_2\}\rangle$
Effect (with the previous act):
$G_{\{s_2,b_2\}} Int_{b_2} buy(\theta_2, 10000\pounds)$ (i.e., b_2 is publicly committed to buy θ_2 at the price of $10000\pounds$ now).

5 Conclusion

We've proposed a novel approach to the semantics of agent communication, based on verifiable social attitudes which are triggered by observable communication acts. We believe that this approach is more adequate for open systems in comparison both to traditional mentalistic and commitment-based semantics, as it allows to analyze the meaning of messages on the social level without the need to know about mental agent properties or architectural details, while being easily comprehensible, downward compatible to BDI, and fully formalized. A subject of future work in this respect will be the practical application of our approach in the field of interaction protocols, and argumentation and negotiation frameworks.

REFERENCES

[1] B. F. Chellas, *Modal Logic: an introduction*, Camb. Univ. Press, 1980.
[2] Philip R. Cohen and Hector J. Levesque, 'Intention is choice with commitment', *Artificial Intelligence*, **42**(2–3), 213–261, (1990).
[3] Foundation for Intelligent Physical Agents. FIPA Interaction Protocol Library Specification, 2000. URL: http://www.fipa.org/specs/fipa00025.
[4] Foundation for Intelligent Physical Agents. FIPA Communicative Act Library Specification, 2002. URL: http://www.fipa.org/repository/aclspecs.html.
[5] B. Gaudou, A. Herzig, and D. Longin, 'Grounding and the expression of belief', in *Proc. 10th Int. Conf. on Princ. of Knowledge Repr. and Reasoning (KR 2006)*, (2006).
[6] P. McBurney, R. M. van Eijk, S. Parsons, and L. Amgoud, 'A dialogue-game protocol for agent purchase negotiations', *Journal of Autonomous Agents and Multi-Agent Systems*, **7**(3), 235–273, (2003).
[7] M. Nickles, F. Fischer, and G. Weiss, 'Communication attitudes: A formal approach to ostensible intentions, and individual and group opinions', in *Proceedings of the Third International Workshop on Logic and Communication in Multiagent Systems (LCMAS 2005)*, eds., W. van der Hoek and M. Wooldridge, (2005).
[8] M. D. Sadek, 'A study in the logic of intention', in *Proc. Third Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'92)*, eds., Bernhard Nebel, Charles Rich, and William Swartout, pp. 462–473. Morgan Kaufmann Publishers, (1992).
[9] J. R. Searle, *Speech acts: An essay in the philosophy of language*, Cambridge University Press, New York, 1969.
[10] M. P. Singh, 'A social semantics for agent communication languages', in *Proceedings of the IJCAI Workshop on Agent Communication Languages*, (2000).
[11] D. R. Traum, *Computational theory of grounding in natural language conversation*, Ph.D. dissertation, Computer Science Departement, University of Rochester, December 1994.