Conversational Semantics with Social Commitments

Roberto A. Flores, Philippe Pasquier, and Brahim Chaib-draa

Université Laval, Département d'Informatique et de Génie Logiciel,
Sainte-Foy, Québec, Canada G1K 7P4
{flores, pasquier, chaib}@damas.ift.ulaval.ca

Abstract. Message semantics are traditionally defined in terms of mental states, which is a trend that is criticized for assuming the sincerity and cooperativeness of agents. To circumvent these limitations, several proposals have been put forth to define the semantics of messages using social commitments. We follow this trend and present a conversational model where the meaning of messages is based on their use as coordinating devices advancing conversations that advance the state of social commitments and the state of the activities in which agents participate.

1 Introduction

Agent communication languages (ACLs) mandate the common elements upon which *coherent conversations* [4] can take place. The most influential ACLs in the agent community are KQML [5] and its *de facto* successor FIPA-ACL [8], which define the semantics of messages using mental states, and the sequencing of messages through conversation protocols. The main reasons to challenge the viability of these approaches in open environments lie on the practical impossibility of agents to verify that uttered messages comply with their semantic definitions without assuming the goodwill of interlocutors to abide by them (*sincerity condition*) [18], and the disassociation between message definitions and their use in conversations [15]. As an alternative to circumvent these limitations, a second trend has emerged that makes use of the notion of social commitments to define coherent conversations. Notably within this trend is the work advanced in [7][10][11][22], which support various aspects of conversational coherence.

We borrow from these experiences and propose an approach to define the conversational semantics of messages, i.e., the meaning that messages could have according to their use in conversations. To this end, we drew inspiration from the study of language use [3], which highlights two complementary types of meaning: *speaker's meaning*, which is based on the use of messages for the communication of intent, and *signal meaning*, which is based on the use of messages as coordinating devices incrementing the common ground of interacting agents.

We advocate this latter type of meaning and conceptualize messages as coordinating devices in conversations that advance the state of social commitments in joint activities, where the states of conversations, commitments and activities are part of the common ground of interacting agents.

Following this view, we propose a model where the meaning of messages is incrementally defined based on the following levels: a *structural level*, where the meaning of messages is given according to their constituents; an *agreement level*, where the meaning of messages is given based on their occurrence as part of a conversation in which agents concur to advance the state of commitments; a *commitment state level*, where the meaning of messages is given according to the state of the commitments these messages manipulate; and a *joint activity level*, where the meaning of messages is given according to their use in joint activities.

In Section 2, we describe our notion of social commitment and its life cycle, i.e., the states in which a commitment could be, as well as the transitions between these states. In particular, we consider the message exchanges that accomplish conversational transitions. This is followed (in Section 3) by our view of agents as image holders, where an image is an agent representation that stores both the messages that the agent represented has exchanged with the agent holding the image, and the commitments these agents have established through these communications. It is in this context that we derive our definitions of shared utterance and shared commitment. Based on these notions, we present (in Section 4) our four-level model upon which messages could be incrementally defined, and illustrate its application by defining a *Call for Proposals* message in a Contract Net Protocol [20] activity. We then conclude with brief remarks on related work.

Throughout this paper, we use the Object-Z specification language [19] to formalize definitions. We chose this language mainly due to the straightforwardness it affords to translate definitions into object-oriented implementations.

2 Social Commitments

The notion of *social commitments* [2][17][21] has been advanced as a way to raise expectations about other agents' performances. Specifically, a social commitment can be defined as an engagement in which an agent (the debtor) is responsible relative to another agent (the creditor) for the performance of an action.[1].

We share with others (e.g., [1][10]) the view that social commitments have a life cycle made of states and transitions between states. As shown in Figure 1, a commitment could be either *accepted* or *rejected* according to whether or not agents are engaged in it. If accepted, a commitment is either *active*, *violated* or *fulfilled*; if rejected, it is either *inactive* or *cancelled*. Commitments can move between states through four transition types: *adoption*, where an inactive commitment becomes accepted; *violation* and *fulfilment*, where an active commitment becomes violated or fulfilled, respectively; and, *discharge*, where an accepted commitment becomes cancelled. Initially, all commitments are inactive, but can become accepted upon adoption. Adopted commitments are classified as either active, violated or fulfilled according to the state of achievement of their conditions of satisfaction (i.e., whether these conditions could be met, cannot be met, or have been met, respectively), and can become cancelled upon discharge.

[1] In this paper, we do not explicitly consider propositional content

It is worth noticing that violation and fulfillment depend on commitments' conditions of satisfaction and that adoption and discharge are accomplished through (conversational) agreement. In this paper, we model only conversational transitions and assume that transitions based on the conditions of satisfaction are carried out automatically.

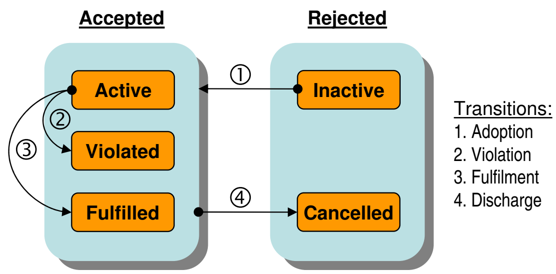

Fig. 1. Social Commitment states.

2.1 Social Commitment Messages

We define message interactions as communicative actions where a speaker communicates to an addressee a (non-empty) set of conversational tokens. We identify four tokens for the negotiation of social commitments: *propose*, which indicates a social commitment operation (that could be either to adopt or to discharge a commitment), and a time interval by which a reply to this proposal is expected; *accept* and *reject*, which are replies indicating either an acceptance or rejection to modify a social commitment state; and *counter*, which simultaneously rejects a modification and proposes a different one to be considered instead. We further specify utterances as events marking the occurrence of communicative actions at a certain moment in time.

2.2 Achieving Conversational Transitions

It is one thing to define communicative acts and quite another to describe how they are used and what they can accomplish in conversations. To that end, we use an interaction protocol called *protocol for proposals* (*pfp*) [6] as the fundamental vehicle to adopt and discharge commitments. As shown in Figure 2, the protocol starts with a proposal (i.e., a communicative act containing a *propose* token) from agent a to agent b. This message can be followed (before the expiration of a reply deadline) by the interaction patterns α or β. Interaction pattern α indicates that either agent b sends an accepting message to agent a, or that the interaction continues with pattern β (but with agents a and b's participatory roles inverted, that is, the role of the agent that in pattern α was agent a will be agent b in pattern β, and likewise for agent b). Interaction pattern β indicates that agent a sends a rejection or counterproposal message to agent b, in which case the interaction follows (before the expiration of a reply deadline) by either pattern α or pattern β. All replies except a counterproposal terminate the protocol; and when an acceptance is issued, both a and b simultaneously apply the proposed and accepted social commitment operation to their record of social commitments.

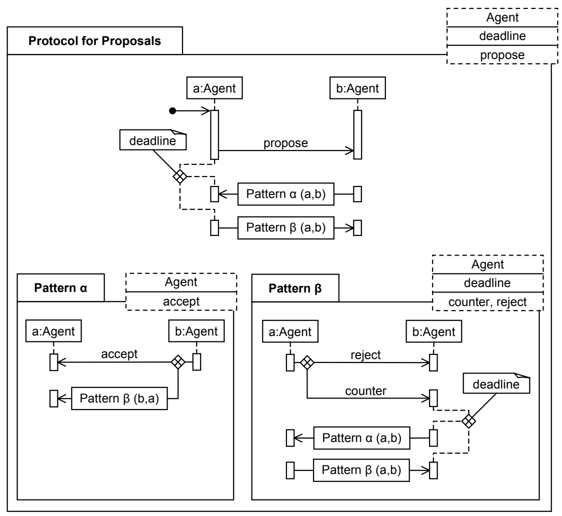

Fig. 2. The *protocol for proposals*.

3 Images and Agents

We conceptualize agents as image holders, and images as agent representations listing utterances and commitments. To restrict complex constructs (e.g., agents holding images that hold images, *ad infinitum*, of other agents), we limit these definitions through the following two properties: first, agents only capture utterances in which they are involved as either the speaker or the addressee (thus circumventing intricate ascriptions, such as agents being thought to have witnessed other agents' communications); and, second, communications are reliable (i.e., issuing an utterance implies that the speaker and the addressee are aware that it occurred). These properties help us capture the shared state of witnessed utterances, in the same spirit as that of shared-basis common ground [3].

Images are specified as repositories of utterances and social commitments, where each social commitment is associated with a unique state. As such, we define (as shown below) a mapping between a social commitment and a state.

$$SocialCommitmentState == \downarrow SocialCommitment \times \downarrow State$$

This definition does not preclude (as a function would do) that agents establish duplicated commitments, i.e., commitments that have the same debtor, creditor and action but different states. This feature requires common constructs and policies to unambiguously manipulate commitments and keep their shared state consistent as they undergo transitions. Rather than providing a unique identifier field as part of the structure of commitments, we kept the minimal debtor-creditor-action structure and specified that all identical commitments could only undergo conversational transitions through independent utterances, which guarantees that each duplicated commitment maps to a unique state.

As shown below, an image is specified as an *UtteranceHolder* and a *SocialCommitmentHolder*, where the former holds a set of utterances occurring at different times. On the other hand, a *SocialCommitmentHolder* (not shown) holds a set of social commitment states in which identical adopted (or discharged) commitments were proposed and accepted using different utterances.

Image

UtteranceHolder
SocialCommitmentHolder

UtteranceHolder

$witnessed : \mathbb{P}\ Utterance$

$\forall u_1, u_2 : witnessed \mid u_1 \neq u_2 \bullet u_1.time \neq u_2.time$

As shown below, we define agents as entities that hold images. To this end, we first define the class *ImageHolder*, which specifies a function mapping agents and images, and where 1) an image only holds utterances whose speaker or addressee is the agent this image represents, 2) an image only holds commitments where its agent is the creditor or debtor, and 3) an image has records of all the utterances that have changed the state of all the adopted and discharged commitments it holds. Lastly, we specify agents as image holders, where each held image records utterances in which the agent is either the speaker or addressee, and commitments where the agent is either the creditor or debtor.

Agent

ImageHolder

$$
\begin{array}{l}
\forall\, agent : \downarrow Agent;\ image : Image \mid \\
\quad image = awareof(agent) \\
\bullet\ (\forall\, utterance : image.witnessed \\
\quad \bullet\ self \in utterance.speechact.performers) \wedge \\
\quad (\forall\, time : Time \\
\quad\quad \bullet\ \forall\, sc : \mathrm{dom}(image.commitments(time)) \\
\quad\quad\quad \bullet\ self \in \{sc.creditor, sc.debtor\})
\end{array}
$$

$$
\begin{array}{l}
\textit{ImageHolder} \\
\hline
\textit{awareof} : \downarrow \textit{Agent} \rightarrow \textit{Image} \\
\hline
\forall \textit{agent} : \downarrow \textit{Agent};\ \textit{image} : \textit{Image} \mid \\
\quad \textit{image} = \textit{awareof}(\textit{agent}) \\
\bullet (\forall \textit{utterance} : \textit{image.witnessed} \\
\quad \bullet \textit{agent} \in \textit{utterance.speechact.performers}) \wedge \\
\quad (\forall \textit{time} : \textit{Time} \\
\quad \bullet (\forall \textit{sc} : \mathrm{dom}(\textit{image.commitments}(\textit{time})) \\
\qquad \bullet \textit{agent} \in \{\textit{sc.debtor}, \textit{sc.creditor}\}) \wedge \\
\qquad (\forall \textit{add} : \downarrow \textit{Adopted} \mid \\
\qquad\quad \textit{add} \in \mathrm{ran}(\textit{image.commitments}(\textit{time})) \\
\qquad \bullet \textit{add.adopt.proposal} \in \textit{image.witnessed} \wedge \\
\qquad\quad \textit{add.adopt.acceptance} \in \textit{image.witnessed}) \wedge \\
\qquad (\forall \textit{del} : \downarrow \textit{Discharged} \mid \\
\qquad\quad \textit{del} \in \mathrm{ran}(\textit{image.commitments}(\textit{time})) \\
\qquad \bullet \textit{del.discharge.proposal} \in \textit{image.witnessed} \wedge \\
\qquad\quad \textit{del.discharge.acceptance} \in \textit{image.witnessed}))
\end{array}
$$

3.1 Sharing Utterances and Commitments

An utterance is shared between its speaker and addressee if they are aware that the utterance has been witnessed by both of them—which holds true given our assumption of reliable communications. Thus, an utterance is shared if its speaker and addressee hold images in which they have witnessed its occurrence.

$$
\begin{array}{l}
\textit{SharedUtterance} : \textit{Utterance} \rightarrow \mathbb{B} \\
\hline
\forall u : \textit{Utterance} \\
\bullet \textit{SharedUtterance}(u) \Leftrightarrow \\
\quad (\forall \textit{agent} : u.\textit{speechact.performers} \\
\quad \bullet \exists \textit{speaker}, \textit{addressee} : \textit{Image} \mid \\
\qquad \textit{speaker} = \textit{agent.awareof}(u.\textit{speechact.speaker}) \wedge \\
\qquad \textit{addressee} = \textit{agent.awareof}(u.\textit{speechact.addressee}) \\
\quad\ \bullet u \in \textit{speaker.witnessed} \wedge \\
\qquad u \in \textit{addressee.witnessed})
\end{array}
$$

Likewise, a social commitment is shared between two agents if these agents are the creditor and debtor of the commitment and if they have images in which this commitment has the same state. Accordingly, that a shared commitment is in an adopted or discharged state implies that agents also share the proposing and accepting utterances that brought the commitment to its current state.

$$
\begin{array}{|l}
\mathit{SharedSocialCommitments} : \mathit{Time} \times \downarrow\!\mathit{Agent} \times \downarrow\!\mathit{Agent} \rightarrow \mathbb{P}\,\mathit{SocialCommitmentState} \\
\hline
\forall\, t : \mathit{Time};\ a_1, a_2 : \downarrow\!\mathit{Agent} \\
\bullet\ \mathit{SharedSocialCommitments}(t, a_1, a_2) = \\
\quad \{sc : \mathit{SocialCommitmentState} \mid \\
\quad\ \forall\, agent : \{a_1, a_2\} \\
\quad\ \bullet\ \exists\, i_1, i_2 : \mathit{Image} \mid \\
\qquad\quad i_1 = agent.awareof\,(a_1) \wedge \\
\qquad\quad i_2 = agent.awareof\,(a_2) \\
\qquad \bullet\ sc \in i_1.commitments(t) \wedge \\
\qquad\quad sc \in i_2.commitments(t)\}
\end{array}
$$

4 Message Definitions

Based on the above specifications, we incrementally built our message definitions through four levels: at the *structural level*, where messages can be classified according to the order and type of their constituents; at the *agreement level*, where the significance of messages is given based on their occurrence as part of conversations seeking agreement to advance the state of commitments; at the *commitment state level*, where meaning is given according to the state of manipulated commitments; and at the *joint activity level*, where meaning is given according to the use of messages as part of joint activities.

4.1 Structural Level

This level sets the foundations to build our message classification. Definitions at this level identify messages based on the type and ordering of their components. These definitions are independent of the occurrence of messages as utterances, and allow their analysis outside the scope of conversations.

In the context of the *pfp*, agents agree on the conversational transition of social commitment states. In this view, a message is a well formed proposal (defined through the *ToPropose* function below) if it contains a *propose* conversational token whose speaker and addressee are the creditor and debtor of the proposed commitment. Similarly, a message is a well formed reply (as defined in *ToReply*) if there is a *reply* token (either an *accept*, *reject* or *counter*) whose speaker and addressee are the creditor and debtor of the commitment, and if there is no other reply token referring to this commitment (thus avoiding ambiguity on the termination of a *pfp* instance within a single message).

$$
\begin{array}{|l}
\mathit{ToPropose} : \mathit{ToSpeak} \rightarrow \mathbb{P}\downarrow\!\mathit{Propose} \\
\hline
\forall\, s : \mathit{ToSpeak} \\
\bullet\ \mathit{ToPropose}(s) = \\
\quad \{p : \downarrow\!\mathit{Propose} \mid \\
\quad\ (p \in s.tokens) \wedge \\
\quad\ (s.performers = \{p.proposing.commitment.creditor, \\
\qquad\qquad\qquad\qquad p.proposing.commitment.debtor\})\}
\end{array}
$$

$$
\begin{array}{|l}
ToReply : ToSpeak \to \mathbb{P} \downarrow Reply \\
\hline
\forall\, s : ToSpeak \\
\bullet\ ToReply(s) = \\
\quad \{ r : \downarrow Reply \mid \\
\qquad (r \in s.tokens) \land \\
\qquad (s.performers = \{ r.replying.commitment.creditor, \\
\qquad\qquad\qquad r.replying.commitment.debtor \}) \land \\
\qquad (\nexists r_1 : \downarrow Reply \mid \\
\qquad\quad r_1 \in s.tokens \land \\
\qquad\quad r_1 \neq r \\
\qquad \bullet\ r_1.replying.commitment = r.replying.commitment) \}
\end{array}
$$

These definitions can be used to specify other messages with more refined meanings. For example, acceptances and rejections could be defined as messages containing an accept and reject token (respectively) that are well formed replies; and counterproposals could be defined as messages containing a counter token that is a well formed proposal and rejection. Other feasible definitions are an offer, which could be a proposal where the speaker is the debtor of the commitment, and a request, where the hearer is the debtor of the commitment.

4.2 Agreement Level

This level builds upon the structural level, and indicates the significance of messages once they are uttered. Definitions take into account the time when an utterance was issued, the previous utterances that are shared between its speaker and addressee, and its occurrence as part of a *pfp* instance.

To support definitions at this level, we specify *SharedProposals* (below) to refer to all proposals shared by two agents, within a certain time interval, that contain a propose token matching a given commitment operation. Likewise, *SharedReplies* (not shown) refers to all shared replies that occurred in a time interval.

$$
\begin{array}{|l}
SharedProposals : Interval \times \downarrow Agent \times \downarrow Agent \times \downarrow Operation \to \mathbb{P}\ Utterance \\
\hline
\forall\, i : Interval;\ a_1, a_2 : \downarrow Agent;\ op : \downarrow Operation \\
\bullet\ SharedProposals(i, a_1, a_2, op) = \\
\quad \{ u : SharedUtterances(a_1, a_2) \mid \\
\qquad i.from \leq u.time \leq i.until \land \\
\qquad (\exists\, p : ToPropose(u.speechact) \\
\qquad\ \bullet\ p.proposing = op) \}
\end{array}
$$

Based on these definitions, we specify that a proposal between two agents at a given time is a sound attempt to reach agreement (as shown in *SoundProposal* below) if 1) there exists (at the given time) a shared utterance between these agents that proposes the given commitment operation, and 2) this proposal can be replied, which we specify simply as having the reply time in the proposal start after the utterance of the proposal.

$$
\begin{array}{|l}
\mathit{SoundProposal} : \mathit{Time} \times \downarrow\mathit{Agent} \times \downarrow\mathit{Agent} \times \downarrow\mathit{Propose} \rightarrow \mathbb{B} \\
\hline
\forall\, \mathit{time} : \mathit{Time};\ \mathit{agent}_1, \mathit{agent}_2 : \downarrow\mathit{Agent};\ \mathit{propose} : \downarrow\mathit{Propose} \\
\bullet\ \mathit{SoundProposal}(\mathit{time}, \mathit{agent}_1, \mathit{agent}_2, \mathit{propose}) \Leftrightarrow \\
\quad (\mathit{SharedProposals}(\mathit{at}(\mathit{time}), \mathit{agent}_1, \mathit{agent}_2, \mathit{propose.proposing}) \neq \varnothing) \wedge \\
\quad (\mathit{time} < \mathit{propose.reply.from})
\end{array}
$$

Likewise, the function *SoundReply* (below) specifies that a reply is sound if, at the time it occurs, there is a proposal that could be answered and has not been answered yet. This outcome is achieved by the partial functions *proposed*, which maps each shared proposal that could be replied at the given time to a set of replies that could answer it; and *replied*, which maps a subset of the proposals in *proposed* with one of its corresponding replies, where each reply replies to only one proposal. Thus, a reply would be sound if there are unanswered proposals, i.e., if the proposals in *replied* is a proper subset of proposals in *proposed*.

$$
\begin{array}{|l}
\mathit{SoundReply} : \mathit{Time} \times \downarrow\mathit{Agent} \times \downarrow\mathit{Agent} \times \downarrow\mathit{Reply} \rightarrow \mathbb{B} \\
\hline
\forall\, t : \mathit{Time};\ s, a : \downarrow\mathit{Agent};\ r : \downarrow\mathit{Reply} \\
\bullet\ \mathit{SoundReply}(t, s, a, r) \Leftrightarrow \\
\quad (\forall\, \mathit{proposed} : \mathit{Utterance} \nrightarrow \mathbb{P}\, \mathit{Utterance} \mid \\
\qquad \mathrm{dom}\, \mathit{proposed} = \\
\qquad\qquad \{u : \mathit{SharedProposals}(\mathit{before}(t), s, a, r.\mathit{replying}) \mid \\
\qquad\qquad\ \exists\, p : \mathit{ToPropose}(u.\mathit{speechact}) \mid \\
\qquad\qquad\quad p.\mathit{reply.from} \leq t \leq p.\mathit{reply.until} \wedge \\
\qquad\qquad\quad p.\mathit{proposing} = r.\mathit{replying} \\
\qquad\qquad\ \bullet\ \mathit{proposed}(u) = \mathit{SharedReplies}(\mathit{within}(p.\mathit{reply.from}, t), s, a, r.\mathit{replying})\} \\
\quad\ \bullet\ \forall\, \mathit{replied} : \mathit{Utterance} \nrightarrow \mathit{Utterance} \mid \\
\qquad\ \forall\, u : \mathrm{dom}\, \mathit{replied} \\
\qquad\ \bullet\ \mathit{replied}(u) \in \mathit{proposed}(u) \wedge \\
\qquad\quad (\nexists u_1 : \mathrm{dom}\, \mathit{replied} \mid \\
\qquad\qquad u_1 \neq u \\
\qquad\quad\ \bullet\ \mathit{replied}(u) = \mathit{replied}(u_1)) \\
\qquad \bullet\ \mathrm{dom}\, \mathit{replied} \subset \mathrm{dom}\, \mathit{proposed})
\end{array}
$$

Based on the above, we define that an utterance would be a proposal if it is a well formed, sound proposal (as shown in *Proposing* below).

$$
\begin{array}{|l}
\mathit{Proposing} : \mathit{Utterance} \rightarrow \mathbb{P}\downarrow\mathit{Propose} \\
\hline
\forall\, u : \mathit{Utterance} \\
\bullet\ \mathit{Proposing}(u) = \\
\quad \{p : \mathit{ToPropose}(u.\mathit{speechact}) \mid \\
\quad\ \mathit{SoundProposal}(u.\mathit{time}, u.\mathit{speechact.speaker}, u.\mathit{speechact.addressee}, p)\}
\end{array}
$$

Likewise, an utterance would be a reply (as defined in *Replying* below) if it is a well formed, sound reply.

$$
\begin{array}{|l}
\mathit{Replying} : \mathit{Utterance} \rightarrow \mathbb{P} \downarrow \mathit{Reply} \\
\hline
\forall\, u : \mathit{Utterance} \\
\bullet\ \mathit{Replying}(u) = \\
\quad \{ r : \mathit{ToReply}(u.\mathit{speechact}) \mid \\
\quad\ \ \mathit{SoundReply}(u.\mathit{time}, u.\mathit{speechact}.\mathit{speaker}, u.\mathit{speechact}.\mathit{addressee}, r) \}
\end{array}
$$

In the same manner as it was explained in the structural level, definitions at this level could also be specialized to create other more refined definitions, such as accepting, rejecting, offering and requesting, among others.

4.3 Commitment State Level

This level builds upon the agreement level, and refines the definitions of messages according to the shared state of the commitment being manipulated. As such, an utterance proposing the discharge of a social commitment (as indicated in *ProposingStateDischarge* below) would be one that contains a propose token attempting to delete an accepted social commitment.

$$
\begin{array}{|l}
\mathit{ProposingStateDischarge} : \mathit{Utterance} \rightarrow \mathbb{P}\, \mathit{SocialCommitmentState} \\
\hline
\forall\, u : \mathit{Utterance} \\
\bullet\ \mathit{ProposingStateDischarge}(u) = \\
\quad \{ \mathit{sc} : \mathit{SocialCommitmentState} \mid \\
\quad\ \ \exists\, p : \mathit{Proposing}(u);\ a : \mathit{Accepted} \mid \\
\quad\quad p.\mathit{proposing} \in \mathit{Delete} \\
\quad\ \ \bullet\ \mathit{sc} = (p.\mathit{proposing}.\mathit{commitment} \mapsto a) \wedge \\
\quad\quad \mathit{sc} \in \mathit{SharedSocialCommitments}(u.\mathit{time}, u.\mathit{speechact}.\mathit{speaker}, \\
\qquad\qquad\qquad\qquad\qquad\qquad u.\mathit{speechact}.\mathit{addressee}) \}
\end{array}
$$

This definition could then be refined as a withdrawal (as indicated in *Withdrawal*, shown below) if the involved commitment is in an active state, and if its discharge is being proposed by the same agent that proposed its adoption. Likewise, this definition could be refined as a *Release* (not shown) if the proposing and proposed agents are the creditor and debtor of the withdrawn commitment.

$$
\begin{array}{|l}
\mathit{Withdrawal} : \mathit{Utterance} \rightarrow \mathbb{P}\, \mathit{SocialCommitmentState} \\
\hline
\forall\, u : \mathit{Utterance} \\
\bullet\ \mathit{Withdrawal}(u) = \\
\quad \{ \mathit{sc} : \mathit{ProposingStateDischarge}(u) \mid \\
\quad\ \ \forall\, a : \mathit{Active} \mid \\
\quad\quad a = \mathit{state}(\mathit{sc}) \\
\quad\ \ \bullet\ a.\mathit{adopt}.\mathit{proposal}.\mathit{speechact}.\mathit{speaker} = u.\mathit{speechact}.\mathit{speaker} \}
\end{array}
$$

4.4 Joint Activity Level

The joint activity level builds upon the commitment state level, and refers to the meaning given to messages when they are used as part of joint activities. In retrospect, the meaning of messages is not only given by their constituents, their use as devices advancing the state of commitments, and the shared state of the commitments they refer to, but also by the type of actions these commitments bring about, and by the roles that interacting agents play in these actions.

To exemplify this point, we refer to a contract net joint activity defined in [7], which specifies a manager and a bidder roles that interact to bring about three interdependent actions: one in which the bidder produces a bid, a second one in which the manager evaluates the bid, and a third one in which (if offered by the manager) the bidder performs the then-bid now-contract. These actions were defined in independent activities with independent roles, and then merged into the contract net activity, where dependencies between roles and actions were defined, e.g., the bid resulting from the bidding action is the bid evaluated in the evaluating action, the bidder is the producer of a bid and the executor of the contract. In this view, a *Call for Proposals* message (below) would be one in which a manager requests to a bidder the adoption of a commitment where the bidder produces a bid (which indicates the requirements she could fulfill).

Message definitions at this level could not only be used to design the roles that agents could be programmed to play in activities but could also be used by deliberative agents to dynamically direct their conversations based on the messages issued and the commitments these messages entail (which is an approach explored in [14]).

$$
\begin{array}{|l}
\mathit{CallforProposals} : \mathit{Utterance} \rightarrow \mathbb{P}\, \mathit{SocialCommitmentState} \\
\hline
\forall\, u : \mathit{Utterance} \mid \\
\quad u.\mathit{speechact}.\mathit{speaker} \in \mathit{Manager} \wedge \\
\quad u.\mathit{speechact}.\mathit{addressee} \in \mathit{Bidder} \\
\bullet\ \mathit{CallforProposals}(u) = \\
\quad \{ \mathit{sc} : \mathit{RequestingAdoption}(u) \mid \\
\quad\ \ \forall\, \mathit{act} : \mathit{ToOfferPerformance} \mid \\
\quad\quad \mathit{act} = (\mathit{commitment}(\mathit{sc})).\mathit{action} \\
\quad\ \ \bullet\ \mathit{act}.\mathit{producer} = u.\mathit{speechact}.\mathit{addressee} \wedge \\
\quad\quad \mathit{act}.\mathit{receiver} = u.\mathit{speechact}.\mathit{speaker} \}
\end{array}
$$

4.5 Resolving Ambiguities in Transitions when Duplicated Commitments Exist

There are a few issues that must be resolved to keep the state of social commitments shared during transitions in cases when duplicated commitments exist, and when commitments can only be identified through a creditor, debtor and action descriptors. In this section we separately explore this issue in the cases of conversational and satisfaction transitions.

On the one hand, conversational transitions deal with the issue of identifying the commitment that is being referred to in utterances. Ambiguity could arise if 1) a reply occurs at a time when more than one proposal with identical commitments could be answered, since this may result in agents selecting different proposals as the one being replied to; and 2) a subsequent reply occurs that is regarded by one of the agents as answering the remaining proposal while the other agent does not (e.g., if the reply time of the message that the latter agent retained as unanswered has expired), and this reply is an acceptance changing the state of a social commitment, since it will result in one of these agents changing the state of the referred commitment while the other agent does not. These cases result (or may result, in the case of the former) in discrepancies on the shared state of commitments: in the first case, agents hold commitments with different replied proposals; in the second case, an agent holds an accepted commitment while the other does not, which may eventually lead to a clash of their expectations of each other within the involved joint activity.

These problems can be prevented either by expanding the structure of communicated commitments with a disambiguating feature (e.g., the time when the proposal being answered was issued, a unique commitment identifier), by engaging on a subsequent dialog requesting the explicit identification of the proposal being replied (e.g., asking for the time when the proposal occurred), or by mandating that agents use the same criteria to select the proposal being replied (e.g., a reply answers the proposal whose reply time expires first). Although we do not model these strategies, we intuitively favor that agents engage in dialogues, rather than augmenting the structure of communicated commitments with proprietary information, or attempting to standardize the functionality of agents, which would be impossible to enforce in open environments.

On the other hand, satisfaction transitions deal with the issue of identifying the performances that satisfy accepted commitments. That an agent holds identical active commitments only means that it has recorded those commitments given independent conversational transitions, not necessarily that these commitments will be satisfied as many times as recorded. Since the possibility of optimizing performances (i.e., whether one performance satisfies all identical commitments or if independent performances are required) may be bound by the expectations of involved agents, they are not modelled in our analysis. Rather, we assume that transitions based on conditions of satisfaction are automatically traversed according to the state of these conditions.

5 Related Work

Conversations and social commitments have been the subject of previous studies. On the one hand, some efforts have aimed at the study of social commitments in argumentation [21], in which the evolution of conversations is motivated by the commitments that are implied in communications. Of particular interest for the agent communication languages community are the proposals furthered in [13][16]. On the other hand, other efforts have focused on the mechanics of conversations based on the operations advancing the state of social commitments, which is a view independent of the intentional motives behind their advancement. We share this latter view, and aim at the identification of public elements binding the evolution of conversations. In addition to our proposal, there are other approaches pursuing this goal, such as those advanced by Fornara and Colombetti [10] (who specify a compositional approach where messages are categorized as speech acts whose meaning is given by operations to manipulate the state of commitments), and Yolum and Singh [22] (whose approach indicates the meaning of messages according to operations and reasoning rules applied to commitments). We concur with these approaches in the view that message meaning could be based on the (shared) state of commitments, and intuitively conform with the practical aspects of *pre-commitments* (to indicate the sequencing of messages establishing commitments) and *conditional commitments* (in the restricted sense of their sequencing), which in our approach are afforded through the *pfp*, and the constraints in agent roles and joint activities, respectively.

There are, however, other aspects where our approaches differ: for example, on the assumptions between the dynamics and the pragmatics of commitment state transitions, which in the aforementioned approaches is not distinguished, e.g., a commitment can be created only by its debtor [22](p. 530) (c.f., could a debtor create a commitment in the context of an offer?)[2], a commitment can be cancelled only by its creditor [9](p. 536) (c.f., could a debtor cancel a commitment if sanctions were applied?). One of the strengths of our approach is that it puts forth a simple incremental model where these aspects can be distinguished. To complement this model, we are currently exploring the role of sanctions in conversational transitions.

An additional concern, noted in [12] regarding [22], is the view that commitment operations should not be unilaterally applied but rather must be jointly approved by interacting agents (unless mandated by the context or by meta-commitments). Although we endorse this view, our current analysis is restricted to explicit manipulations approved by consensus, as afforded by the *pfp*, for any conversational transition adopting or discharging social commitments.

6 Conclusions

As noted in [3], utterances are signals with two complementary types of meaning: *speaker's meaning*, which is defined in terms of their use for the communication of intent, and *signal meaning*, which is defined in terms of their use as coordinating devices to advance conversations.

Within the multiagents community, message semantics has traditionally emphasized speaker's meaning, as reflected by FIPA-ACL and KQML's adoption of speech acts and mental states for their message definitions. The main advantage of these approaches is that they straightforwardly communicate intent, i.e., agents can readily know the intended meaning of a message by just observing its definition rather than by inferring it from the context of interaction. However, and as noted earlier, the practical use of these approaches in open environments is handicapped by their assumption that agents are sincere and cooperative.

Signal meaning, on the other hand, has been kept as a low profile component of meaning and is not addressed by these standardizing efforts. We contend that

[2] See [12](p. 369) for a discussion on this issue.

this type of meaning should be taken into account as part of message definitions, and explore this possibility in a restricted context where messages aim at the negotiated manipulation of social commitment states within joint activities. We chose this type of messages due to the fitness of social commitments to coordinate the expectations of agents and advance the state of their interactions.

Following this perspective, we propose a four-level incremental model that focuses on the characteristics of messages (*structural level*) through which agents agree through conversations (*agreement level*) to advance the state of social commitments (*commitment state level*) within joint activities (*joint activity level*). We explored the feasibility of *pfp* messages to describe signal meaning given their support for building flexible and modular conversation protocols [7].

7 Acknowledgements

We are grateful for the support received from the National Science and Engineering Research Council (NSERC) of Canada during the preparation of this work. We are also thankful to the anonymous reviewers for their invaluable comments.

References

1. J. Bentahar, B. Moulin, and B. Chaib-draa. Commitment and argument network: A new formalism for agent communication. In F. Dignum, editor, *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Artificial Intelligence*, pages 146–165. Springer Verlag, 2004.
2. C. Castelfranchi. Commitments: From individual intentions to groups and organizations. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 41–48, San Francisco, CA, June 1995.
3. H.H. Clark. *Using language.* Cambridge University Press, 1996.
4. R.T. Craig and K. Tracy. *Conversational Coherence: Form, Structure, and Strategy.* Sage Publications, 1983.
5. T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J.M. Bradshaw, editor, *Software Agents*, pages 291–316. MIT Press, 1997.
6. R.A. Flores and R.C. Kremer. To commit or not to commit: Modelling agent conversations for action. *Computational Intelligence*, 18(2):120–173, 2003.
7. R.A. Flores and R.C. Kremer. A principled modular approach to construct flexible conversation protocols. In A.Y. Tawfik and S.D. Goodwin, editors, *Advances in Artificial Intelligence: Proceedings of the $17^t h$ Canadian Conference on Artificial Intelligence*, volume 3060 of *Lecture Notes in Computer Science*, pages 1–15, London, Canada, May 2004. Springer Verlag.
8. Foundation for Intelligent Physical Agents (FIPA). http://www.fipa.org.
9. N. Fornara and M. Colombetti. Operational specification of a commitment-based agent communication language. In C. Castelfranchi and W.L. Johnson, editors, *Proceedings of the 1^{st} International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 535–542, Bologna, Italy, July 2002.
10. N. Fornara and M. Colombetti. Defining interaction protocols using a commitment-based agent communication language. In J.S. Rosenschein, T. Sandholm, M.J. Wooldridge, and M. Yokoo, editors, *Proceedings of the 2^{nd} International Joint*

Conference on Autonomous Agents and Multiagent Systems, pages 520–527, Melbourne, Australia, July 2003. ACM Press.
11. M.A. Labrie, B. Chaib-draa, and N. Maudet. Diagal: A tool for analyzing and modelling commitment-based dialogues between agents. In Y. Xiang and B. Chaib-draa, editors, *Advances in Artificial Intelligence*, volume 2671 of *Lecture Notes in Artificial Intelligence*, pages 353–369. Springer Verlag, 2003.
12. P. McBurney and S. Parsons. The posit spaces protocol for multi-agent negotiation. In F. Dignum, editor, *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Artificial Intelligence*, pages 364–382. Springer Verlag, 2004.
13. S. Parsons, P. McBurney, and M.J. Wooldridge. The mechanics of some formal inter-agent dialogues. In F. Dignum, editor, *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Artificial Intelligence*, pages 329–348. Springer Verlag, 2004.
14. P. Pasquier and B. Chaib-draa. The cognitive coherence approach for agent communication pragmatics. In J.S. Rosenschein, T. Sandholm, M.J. Wooldridge, and M. Yokoo, editors, *Proceedings of the 2^{nd} International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 544–552, Melbourne, Australia, July 2003. ACM Press.
15. J. Pitt and A. Mamdani. Some remarks on the semantics of FIPAs agent communication language. *Autonomous Agents and Multi-Agent Systems*, 2(4):333–356, 1999.
16. C.A. Reed. Dialogue frames in agent communication. In *Proceedings of the 3^{rd} International Conference on Multiagent Systems (ICMAS 98)*, pages 246–253, Paris, France, 1998. IEEE Press.
17. M.P. Singh. Social and psychological commitments in multiagent systems. In *AAAI Fall Symposium on Knowledge and Action at Social and Organizational Levels*, Monterey, California, November 1991.
18. M.P. Singh. Agent communicational languages: Rethinking the principles. *IEEE Computer*, 31(12):40–47, December 1998.
19. G. Smith. *The Object-Z Specification Language*. Kluwer Publishers, 2000.
20. R.G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104–1113, 1980.
21. D.N. Walton and E.C.W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, 1995.
22. P. Yolum and M.P. Singh. Flexible protocol specification and execution: Applying event calculus planning using commitments. In C. Castelfranchi and W.L. Johnson, editors, *Proceedings of the 1^{st} International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 527–534, Bologna, Italy, July 2002.