

Bringing Coherence to Agent Conversations

Roberto A. Flores
Computer Science Department
University of Calgary
2500 University Dr., NW,
Calgary, Canada, T2N 1N4
1-403-2207299

robertof@cpsc.ucalgary.ca

Robert C. Kremer
Computer Science Department
University of Calgary
2500 University Dr., NW,
Calgary, Canada, T2N 1N4
1-403-2205112

kremer@cpsc.ucalgary.ca

ABSTRACT

In this paper, we present a social model for software agent conversations for action based on social commitments and their negotiation. We depart from the premises that conversations are cornerstone to support autonomous and heterogeneous agents, and that conversational coherence can be supported through public definitions of speech act and compositional semantics.

We specify a unified social model for conversations in which speech act semantics is an emergent product of identity, conversational use, and expected accomplishments, and where conversational composition is guided by rules of conversational use and their application to the state of conversation instances.

Lastly, we show the effectiveness of this novel approach by formally describing the evolution of a simple conversation for action.

General Terms

Standardization, Languages, Theory.

Keywords

Conversations, social commitments, speech act semantics, compositional semantics.

1. INTRODUCTION

Internet-enabled software is evolving from being mere communication enablers to become active participants in automating the interactions of their users. More and more, these software programs (often called *software agents*) are being invested with the authority and intelligence to autonomously interact with other programs on behalf of their owners.

In open environments such as the Internet agents can be drawn from heterogeneous sources. The common denominator of these agents is not how they are built, but how they *converse*. Cornerstone to the notion of conversation is the issue of their *coherence*, i.e., the shared understanding on the meaning of speech acts (*speech act semantics*), and the connectedness between utterances (*compositional semantics*) [8].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Autonomous Agents '01, May 28-June 1, 2001, Montreal, Canada.

Copyright 2001 ACM 1-58113-000-0/00/0000...\$5.00.

Agent communication languages (ACL) and conversation protocols (CP) are the traditional approaches to support meaning and composition of conversations (respectively) for software agents. The semantics of speech acts in ACL are customarily defined in terms of mental attitudes, such as beliefs, intentions and goals. This approach has been criticized as inadequate for open environments, since agents cannot verify whether the private beliefs of other agents comply with speech act definitions without pre-established constraints on how agents are implemented. On the other hand, CP are static structures that define the sequence of utterances making a coherent conversation. This approach has been criticized for its lack of flexibility, i.e., the lack of compositional rules governing how protocols are extended or merged. An additional shortcoming is that ACL and CP are defined independently from each other. There have been attempts to combine their semantics into a comprehensive approach, but these efforts have failed to recognize the dependencies between speech act and compositional semantics.

Our aim is to support coherency in conversations by integrating speech act semantics and connectedness of utterances in a unified approach.

For that we propose a model for conversations where the meaning of a speech act is the emergent product of their *identity*, *conversational use*, and potential *accomplishments*. The connectedness of utterances is guided by the rules of *conversational use* and their application to the state of conversation instances.

Central to this model is the notion of *social commitments to action*, i.e., engagements to a course of action taken by an agent relative to another agent on whose behalf actions are done. We formally define a small set of fundamental rules to support the negotiation of shared commitments to action, which we see as instrumental for autonomous agents to advance the state of their joint activities.

The structure of this paper is as follows: first, we give a brief account of current approaches to speech act semantics and conversations. This is followed by the description of our unifying social model for agent conversations, and a formal example of its application. We conclude with a brief description of this work.

2. SPEECH ACT SEMANTICS

Linguistic pragmatics is the area of linguistics that studies the use of language as action. One of the dominant theories in this area is that of speech act theory [1][17], which has been traditionally used in computer science to motivate speech act semantics for ACL.

Researchers have determined that ACL semantics should yield definitions that are publicly accessible, such that any utterance of a speech act can be verified to comply with the state of the world [11][13][18][21].

Current ACL approaches that define speech acts in terms of mental attitudes (e.g., Cohen & Levesque’s Joint Intention Theory [3][4][5], KQML [10][15][16], and FIPA ACL [11]) have been criticized as inadequate to comply with this requirement. It is argued that, since mental attitudes are inherently private to agents, verifying the compliance of an utterance given in mental terms would require that agents be constructed in pre-defined ways to allow the inspection of their mental states [18]. We find this restriction unacceptable if a speech act semantics is to be applied to heterogeneous and autonomous agents.

As such, our *first goal* is:

To support the communication of autonomous and heterogeneous agents in open environments we require a publicly verifiable speech act semantics.

To support this goal, we drew inspiration from existing social ACL approaches (e.g., Singh’s [19], and Colombetti’s [6] models) and adopted social commitments as elements to support speech act meaning. Nevertheless, our model is built independently.

3. COMPOSITIONAL SEMANTICS

There are two approaches to derive the sequencing of speech acts in conversations: the first one is to define *compositional rules*, i.e., rules for deriving which utterances can follow any given utterance under specific states; the second one is to define *conversation protocols*, i.e., structures that specify which utterances can follow any given utterance independently of state.

We favor the definition of compositional rules over CP since the former allows conversations to dynamically evolve according to their context of use, whether the latter only allows pre-specified patterns of interaction that cannot be manipulated or merged with other protocols [12].

As such, our *second goal* is:

To support flexible conversations, we require a set of rules to specify conversational composition.

To support this goal, we define rules for speech acts’ *conversational use* to govern their sequencing according to the state of social commitments shared among conversing agents.

4. INTEGRATING SPEECH ACTS AND COMPOSITION

Interestingly, in all the current ACL approaches we have studied (whether based on mental attitudes or social commitments) the sequencing of speech acts is defined through conversation protocols.¹ In fact, these approaches treat speech act semantics and CP (as a palliative for conversational rules) as independent but pragmatically complementing definitions to support conversational coherence.

¹ In the case of Cohen and Levesque’s JIT, Smith *et al.* [20] proposed an extension to accommodate for the sequencing of speech acts in the KAOs framework [2].

We do not share this view. Instead of dealing with speech act and compositional semantics as separate attributes for conversations, we aspire to integrate them under a unifying model.

As such, our *third goal* is:

To support a unifying model for speech acts and compositional semantics, we require a common set of principles from which these can be derived.

To support this goal, we define the meaning of speech acts as the emergent product of their *identity*, *conversational use*, and potential *accomplishments*. The connectedness of utterances is guided by the rules of *conversational use* and their application to the state of conversation instances.

5. MODELING CONVERSATIONS FOR ACTION

In this section, we specify a model for conversations based on social commitments to action, and introduce our definition of agents, speech acts, social commitments and various conversation policies for the negotiation of shared social commitments. We formalize these definitions using the Z formal notation [9].

5.1 Speech Acts

We use illocutionary points, i.e., the publicly intended perlocutionary effects, as the basic compositional elements of speech acts. This view allows us to describe the meaning of a speech act as the emergent property of its enclosed illocutionary points.

As shown below, we define speech acts as structures composed of an illocutionary force and a set of illocutionary points.

```
SPEECH_ACT
force: ILLOCUTIONARY_FORCE;
points: P ILLOCUTIONARY_POINT;
```

We also specify that a speech act is a kind of action (physical acts are other type of actions that could be included in this definition).

ACTION ::= SpeechAct «SPEECH_ACT»

To complement this definition, we specify an utterance to be an event occurring at a certain moment in time involving a speech act between a speaker and an addressee.

EVENT ≡ [time:TIME; action: ACTION]

UTTERANCE ≡ [EVENT; speaker:AGENT; addressee:AGENT; speechAct:SPEECH_ACT | (action ∈ ran SpeechAct) ∧ (speechAct = SpeechAct~ action)]

5.2 Social Commitments

We define a social commitment as a structure indicating that there is a debtor committed to an action relative to a creditor on whose behalf the action is done. Based on this definition, we specify that a shared social commitment is a structure with a commitment that is shared among agents.

That speech acts have been declared as actions allows us to have social commitments entailing a speech act.

```
SOCIAL_COMMITMENT
debtor, creditor: AGENT;
action: ACTION;
```

SHARED_SOCIAL_COMMITMENT _____ SOCIAL_COMMITMENT; among: P AGENT; among $\neq \emptyset$
--

5.2.1 Operations on Social Commitments

To denote that social commitments can be adopted or discharged we define the type OPERATION, which is defined in terms of a social commitment.

OPERATION ::= Add «SOCIAL_COMMITMENT» |
Delete «SOCIAL_COMMITMENT»

5.3 Agents

Succinctly, we conceptualize an agent in our model as an entity that maintains a set of shared social commitments and a history of the utterances it has witnessed. Agents autonomously decide whether other agents can affect their set of shared social commitments. This is supported by a negotiation process based on the utterance and sequencing of speech acts. We have called this process the Protocol for Proposals.

AGENT _____ commitments: P SHARED_SOCIAL_COMMITMENT; utterances: P UTTERANCE;
--

5.4 The Protocol for Proposals




Figure 1. Example Interaction Diagram for the Protocol for Proposals.

We define a basic protocol for the negotiation of social commitments. This protocol, which we call the *Protocol for Proposals* (PFP), is informally described as follows: The protocol starts with a proposal from a sender to a receiver to concurrently adopt or discharge a specified social commitment. Either the receiver replies with an acceptance, rejection, or counteroffer, or the sender issues a withdrawal or counteroffer. All utterances except a counteroffer terminate that instance of the protocol. A counteroffer is deemed as a proposal, in the sense that its utterance is followed by any of the reply speech acts (but with speaker-addressee roles inverted if the original addressee is the speaker of the utterance). When an acceptance is issued both speaker and addressee simultaneously apply the proposed commitments to their record of shared social commitments.

Figure 1 shows a limited example of how the interaction in the PFP could occur. In this figure, the interaction starts with an utterance from agent A containing a PROPOSE illocutionary point. This utterance can be replied by agent B with a speech act containing an ACCEPT, REJECT, or COUNTER illocutionary points. An acceptance or a rejection terminates the interaction; a counteroffer continues it by requiring a reply from agent A in which he accepts, rejects or counteroffers the just counteroffered proposal (not shown in the figure). A more complex example showing how the PFP is applied to conversations is given in a later section.

5.4.1 Requirements for the Protocol for Proposals

To support this protocol, we define the following requirements (and corresponding effects):

- **Requirement 1.** Agents shall be able to propose other agents to consider the shared uptake of social commitments.
 - Effect a)* A proposal commits proposed agents to reply.
- **Requirement 2.** Proposed agents shall be able to accept a proposal.
 - Effect a)* An acceptance releases proposed agents of the commitment to reply.
 - Effect b)* An acceptance (if uttered within the *window of interaction*² specified in the proposal) realizes the shared uptake of proposed social commitments.
- **Requirement 3.** Agents shall be able to reject a proposal.
 - Effect a)* A rejection releases proposed agents of the commitment to reply.
- **Requirement 4.** Agents shall be able to counteroffer a proposal.
 - Effect a)* A counteroffer releases proposed agents of the commitment to reply to the proposal.
 - Effect b)* A counteroffer commits counteroffered agents to reply within the window of interaction.
- **Requirement 5.** Counteroffered agents shall be able to accept a counteroffer.
 - Effect a)* An acceptance releases counteroffered agents of the commitment to reply.
 - Effect b)* An acceptance (if uttered within the *window of interaction* specified in the counteroffer) realizes the speaker and addressee shared uptake of counteroffered commitments.
- **Requirement 6.** Agents shall be able to reject a counteroffer.
 - Effect a)* A rejection releases counteroffered agents of the commitment to reply.
- **Requirement 7.** Agents shall be able to counteroffer a counteroffer.
 - Effect a)* A counteroffer releases previously counteroffered agents of the commitment to reply to that counteroffer within the window of interaction.
 - Effect b)* A counteroffer commits counteroffered agents to reply.

These requirements (and embedded effects) are supported by the definitions in the sections below.

5.4.2 Illocutionary Points

To model the PFP, we define four basic illocutionary points: PROPOSE, ACCEPT, REJECT, and COUNTER. We define them as illocutionary points as follows:

ILLOCUTIONARY_POINT ::= Propose «PROPOSE» | Accept «ACCEPT» | Reject «REJECT» | Counter «COUNTER»

² The *window of interaction* is a time interval in which a reply is expected.

5.4.2.1 Proposing

We define the illocutionary point PROPOSE to support *requirement 1* (agents shall be able to propose to other agents that they consider the shared uptake of social commitments). These are supported by the specification of the variables proposing as the operation on a social commitment being proposed, and replyBy as the time indicating the window of interaction.

PROPOSE

```
proposing: OPERATION;
replyBy: TIME;
```

5.4.2.2 Accepting

We define the illocutionary point ACCEPT to support *requirement 2* (proposed agents shall be able to accept a proposal) and *requirement 5* (counteroffered agents shall be able to accept a counteroffer). These requirements are supported by the specification of accepting as the social commitment that is being agreed to uptake.

ACCEPT

```
accepting: OPERATION;
```

5.4.2.3 Rejecting and Withdrawing

We define the illocutionary point REJECT to support *requirement 3* (agents shall be able to reject a proposal) and *requirement 6* (agents shall be able to reject a counteroffer). These requirements are supported by the specification of rejecting as the social commitment being rejected. Note that—from the standpoint of a speech act—there is no difference between rejecting and withdrawing, since both discard a previously made proposal. Rather, that one is rejecting or withdrawing is defined in the context of previously uttered speech acts. This point is addressed in a later section on conversational commitments.

REJECT

```
rejecting: OPERATION;
```

5.4.2.4 Counter-Offering

Lastly, we define the illocutionary point COUNTER to support *requirement 4* (agents shall be able to counteroffer a proposal), *requirement 7* (agents shall be able to counteroffer a counteroffer). These requirements are supported by the specification of COUNTER in terms of REJECT and PROPOSE, where the former indicates the commitment previously proposed and now being rejected, and the latter presents the new proposed commitment along with a time indicating the window of interaction.

COUNTER

```
REJECT;
PROPOSE;
```

5.4.3 Conversational Commitments

We identify conversational commitments as those commitments that entail the utterance of a speech act. In this section, we present conversation policies for the adoption and discharge of commitments when uttering the illocutionary points for the PFP. These policies can be taken as norms of conversational behavior that agents are expected to follow in a society that mandates the use of the protocol.

5.4.3.1 Adopting Conversational Commitments

Policy 1A and Policy 1B support the expectation that agents being addressed by an utterance with a PROPOSE or a COUNTER illocutionary point (respectively) will commit to reply.

Policy 1A (Figure 2) supports *effect 1a* (a proposal commits the proposed agents to reply), and Policy 1B (not shown) supports *effects 4b* and *7b* (a counteroffer commits addressees to reply).

POLICY_1A

```
Δ AGENT
```

```
utterance?: UTTERANCE;
```

```
set: P SHARED_SOCIAL_COMMITMENT;
```

```
utterance?.time = now;
```

```
#set = #(getPROPOSEpoints utterance?.speechAct);
```

```
∀ propose:PROPOSE
```

```
| propose ∈ getPROPOSEpoints utterance?.speechAct
```

```
• ∃ sc:SHARED_SOCIAL_COMMITMENT
```

```
| sc.among = {utterance?.speaker, utterance?.addressee} ∧
```

```
sc.debtor = utterance?.addressee ∧
```

```
sc.creditor = utterance?.speaker ∧
```

```
(∃ action:ACTION
```

```
| action ∈ ran SpeechAct ∧
```

```
(∃ s:SPEECH_ACT
```

```
| ∃ ip:ILLOCUTIONARY_POINT
```

```
| ip ∈ s.points
```

```
• (ip ∈ ran Accept ∧ (∃ a:ACCEPT | a = Accept~ ip
```

```
• a.accepting = propose.proposing)) ∨
```

```
(ip ∈ ran Reject ∧ (∃ r:REJECT | r = Reject~ ip
```

```
• r.rejecting = propose.proposing)) ∨
```

```
(ip ∈ ran Counter ∧ (∃ c:COUNTER | c = Counter~ ip
```

```
• c.rejecting = propose.proposing))
```

```
• s = SpeechAct~ action)
```

```
• action = sc.action)
```

```
• sc ∈ set;
```

```
commitments' = commitments ∪ set;
```

Figure 2. A proposal commits addressee agents to reply.

This policy can be read as “for each PROPOSE illocutionary point in a just uttered speech act, agents add to their commitments database that speaker and addressee have as a shared social commitment that the addressee (debtor) will utter to the speaker (creditor) a reply speech act (action) with either an ACCEPT, REJECT or COUNTER with the same operation on commitment as the one found in the just uttered PROPOSE.”

Policy 1B reads the same as this previous policy but with a COUNTER instead of the PROPOSE illocutionary point.

5.4.3.2 Discharging Conversational Commitments

Once that agents have adopted conversational commitments, they can expect their discharge when any of the following three conditions hold:

- *Condition 1.* A proposed (or counteroffered) agent utters a speech act containing an ACCEPT with the same operation on commitment as that of the previously uttered PROPOSE (or COUNTER). This condition is defined in Policy 2A (Figure 3), which supports *effect 2a* (an acceptance releases proposed agents of the commitment to reply) and *effect 5a* (an acceptance

releases counteroffered agents of the commitment to reply).

- *Condition 2.* The proposing or the proposed agent utters a speech act containing a REJECT illocutionary point with the same operation on commitment as that of the previously uttered PROPOSE. This condition is defined in Policy 2B (not shown), which supports *effect 3a* (a rejection releases proposed agents of the commitment to reply) and *effect 6a* (a rejection releases counteroffered agents of the commitment to reply).
- *Condition 3.* The proposing or proposed (or counteroffering or counteroffered) agent utters a speech act containing a COUNTER illocutionary point rejecting the same operation on commitment as that of the previously uttered PROPOSE (or COUNTER). This condition is defined in Policy 2C (not shown), which supports *effect 4a* (a counteroffer releases proposed agents of the commitment to reply) and *effect 7a* (a counteroffer releases counteroffered agents of the commitment to reply).

Policy 2A can be read as “for each ACCEPT illocutionary point in a just uttered speech act, agents delete from their commitments database that the speaker and addressee have as a shared social commitment that the addressee (creditor) will utter to the speaker (debtor) a reply speech act (action) with either an ACCEPT, REJECT or COUNTER with the same operation on commitment as

POLICY_2A
Δ AGENT utterance?: UTTERANCE; set: P SHARED_SOCIAL_COMMITMENT;
utterance?.time = now; #set = #(getallPROPOSEforAllACCEPT(utterances, utterance?)) + #(getallCOUNTERforAllACCEPT(utterances, utterance?));
\forall accept:ACCEPT accept \in getACCEPTpoints utterance?.speechAct \wedge (isthereaPROPOSEforthisACCEPT(utterances, utterance?.speaker, utterance?.addressee, accept) \vee isthereaCOUNTERforthisACCEPT(utterances, utterance?.speaker, utterance?.addressee, accept))
• \exists sc:SHARED_SOCIAL_COMMITMENT sc.among = {utterance?.speaker, utterance?.addressee} \wedge sc.debtor = utterance?.speaker \wedge sc.creditor = utterance?.addressee \wedge (\exists action:ACTION action \in ran SpeechAct \wedge (\exists s:SPEECH_ACT \exists ip:ILLOCUTIONARY_POINT ip \in s.points • (ip \in ran Accept \wedge (\exists a:ACCEPT a = Accept~ ip • a.accepting = accept.accepting)) \vee (ip \in ran Reject \wedge (\exists r:REJECT r = Reject~ ip • r.rejecting = accept.accepting)) \vee (ip \in ran Counter \wedge (\exists c:COUNTER c = Counter~ ip • c.rejecting = accept.accepting)) • s = SpeechAct~ action) • action = sc.action) • sc \in set;
commitments' = commitments \ set

Figure 3. An acceptance releases proposed agents of the commitment to reply.

the one found in the just uttered PROPOSE (or COUNTER).”

Conversational commitments, as defined in these policies, support the continuity of conversations. As a conversation evolves, conversational commitments are added and others are deleted. In this view, conversations end when there are no outstanding conversational commitments for that conversation instance.

5.4.4 Agreeing to Uptake Social Commitments

Nevertheless, one thing is to uptake conversational commitments and another one to uptake proposed social commitments, i.e., the social commitments being negotiated. We consider that conversing agents are only justified to affect their set of shared commitments if there is the awareness that a speech act containing a PROPOSE or COUNTER illocutionary point is followed (within specified times) by a speech act containing an ACCEPT with the same operation on commitment as that of the PROPOSE or COUNTER previously uttered.

We define Policy 3 (Figure 4) to denote the uptake of social commitments upon the correct sequencing of a pair of utterances containing a PROPOSE or a COUNTER and an ACCEPT illocutionary points. These policies support *effects 2b* and *5b* (an acceptance realizes the shared uptake of proposed/counteroffered commitments), respectively.

This policy can be read as “for each ACCEPT illocutionary point in a speech act uttered by a speaker to an addressee that has an identical operation on commitment as that of a PROPOSE or a COUNTER previously uttered by the addressee to the speaker (and before the expiration time indicated in the PROPOSE or COUNTER

POLICY_3
Δ AGENT utterance?: UTTERANCE; setToAdd: P SHARED_SOCIAL_COMMITMENT; setToDelete: P SHARED_SOCIAL_COMMITMENT;
utterance?.time = now; #setToAdd + #setToDelete = #(getallPROPOSEforAllACCEPT(utterances, utterance?)) + #(getallCOUNTERforAllACCEPT(utterances, utterance?));
\forall accept:ACCEPT accept \in getACCEPTpoints utterance?.speechAct \wedge (isthereaPROPOSEforthisACCEPT(utterances, utterance?.speaker, utterance?.addressee, accept) \vee isthereaCOUNTERforthisACCEPT(utterances, utterance?.speaker, utterance?.addressee, accept))
• \exists c:SOCIAL_COMMITMENT c = Add~ accept.accepting \vee c = Delete~ accept.accepting • \exists sc:SHARED_SOCIAL_COMMITMENT sc.among = {utterance?.speaker, utterance?.addressee} \wedge sc.debtor = c.debtor \wedge sc.creditor = c.creditor \wedge sc.action = c.action • (accept.accepting \in ran Add \wedge sc \in setToAdd) \vee (accept.accepting \in ran Delete \wedge sc \in setToDelete);
commitments' = (commitments \cup setToAdd) \ setToDelete

Figure 4. A proposal followed by an acceptance realizes the proposed operation on commitment.

has elapsed), agents apply to their commitments database as a shared commitment between speaker and addressee the negotiated operation on commitment.”

6. EXAMPLE: ASKING THE TIME

In the dynamics of our model, agents will join a society where the description of activities is specified in terms of roles, sequencing of communicational actions, and the description of actions and their results.

Currently our model only accounts for the expected sequencing of communicative actions in an activity. We acknowledge though the importance of action definitions, but such study will not be pursued here since they are not significant to define our rules of conversational composition.

In this section, we present a small example of the application of our model to guide the evolution of conversations. This example is intended for agents that want to know about the time of the day. As such we specify an activity in terms of an action *getTime*, which we informally define as “to get the current local time,” along with two policies to indicate the adoption and discharge of the commitment to communicate the obtained time.

The first policy, which we identify here as Policy 4 (Figure 5), indicates that the acceptance of a proposal for adopting the action *getTime* causes the automatic adoption of the shared commitment that the accepting agent will utter a speech act in which he proposes to discharge the commitment that he performs the action. This is, that an agent is accepting to do the action *getTime* causes the adoption as a shared commitment that he will propose its discharge. Although this commitment could have also being included in the proposing speech act, its definition as part of the activity description allows agents to know before hand how the interactions in the activity are expected to evolve.

Policy 5 (not shown) indicates that once the creditor agent has accepted the proposal to discharge the action *getTime*, there is the automatic shared commitment that the debtor has to propose discharging the action. That is, that the action has being discharged causes the release of the commitment that its discharge has to be proposed. At this point in a conversation, agents can expect that there will be no outstanding conversational commitments, and thus that the conversation is over.

Figure 6 shows one possible conversation for this activity³. In this case, the conversation starts with the utterance of a speech act from agent A to agent B that contains a PROPOSE illocutionary point proposing the addition of a social commitment in which agent B is to perform the action *getTime* on behalf of agent A. As shown in the figure, this utterance triggers Policy 1A (uttering a proposal causes the shared conversational commitment that this proposal will be replied) which results in the addition of shared commitment number 1 to the state of shared commitments⁴.

Next is an utterance from agent B to agent A in which the former accepts committing to do the action. This acceptance triggers the following policies: Policy 2A (the reply to a proposal discharges

³ For the clarity of the diagram, we omitted some of the parameters in the structures used (e.g., the window of interaction in the PROPOSE illocutionary points is not shown).

⁴ In this figure, we use the symbol \otimes to denote XOR (exclusive-or).

POLICY_4

```

Δ AGENT
utterance?: UTTERANCE;
set: P SHARED_SOCIAL_COMMITMENT;
getTime: ACTION;

utterance?.time = now;
#set = #(getAllRequestForTimeforallAcceptToGiveTime(utterances,
                                                    utterance? ));

getTime ∈ ran getTime;

∀ accept:ACCEPT
| accept ∈ getACCEPTpoints utterance?.speechAct ∧
(∃ c:SOCIAL_COMMITMENT
 | c = Add~ accept.accepting
 • c.debtor = utterance?.speaker ∧
   c.creditor = utterance?.addressee ∧
   c.action = getTime) ∧
(#({propose:PROPOSE |
      (∀ u:UTTERANCE
 | u ∈ utterances ∧
   u.time < now ∧
   u.speaker = utterance?.addressee ∧
   u.addressee = utterance?.speaker
 • propose ∈ getPROPOSEpoints u.speechAct ∧
   propose.proposing = accept.accepting ∧
   propose.replyBy > now)) > 0)
• ∃ sc:SHARED_SOCIAL_COMMITMENT
 | sc.among = {utterance?.speaker, utterance?.addressee} ∧
   sc.debtor = utterance?.speaker ∧
   sc.creditor = utterance?.addressee ∧
(∃ action:ACTION
 | action ∈ ran SpeechAct ∧
(∃ s:SPEECH_ACT
 | ∃ propose:PROPOSE
 | propose ∈ getPROPOSEpoints s
 • ∃ c:SOCIAL_COMMITMENT
 | c = Delete~ propose.proposing
 • c.debtor = sc.debtor ∧
   c.creditor = sc.creditor ∧
   c.action = getTime
 • s = SpeechAct~ action)
 • action = sc.action)
 • sc ∈ set;

commitments' = commitments ∪ set

```

Figure 5. The acceptance to commit to the action *getTime* causes the commitment to propose its discharge.

the commitment to reply), which deletes commitment number 1; Policy 3 (the acceptance of a proposal causes the shared uptake of the proposed commitment), which adds commitment number 2; and, Policy 4 (accepting to perform the action *getTime* causes the shared commitment that the agent doing the action will propose the discharge of the commitment to do the action—presumably because he has done it, or because he is polite enough to communicate that he will not be able to do it), which causes the adoption of shared commitment number 3.

Next is an utterance from agent B to agent A in which he proposes to discharge that he is to do the action. This proposal triggers

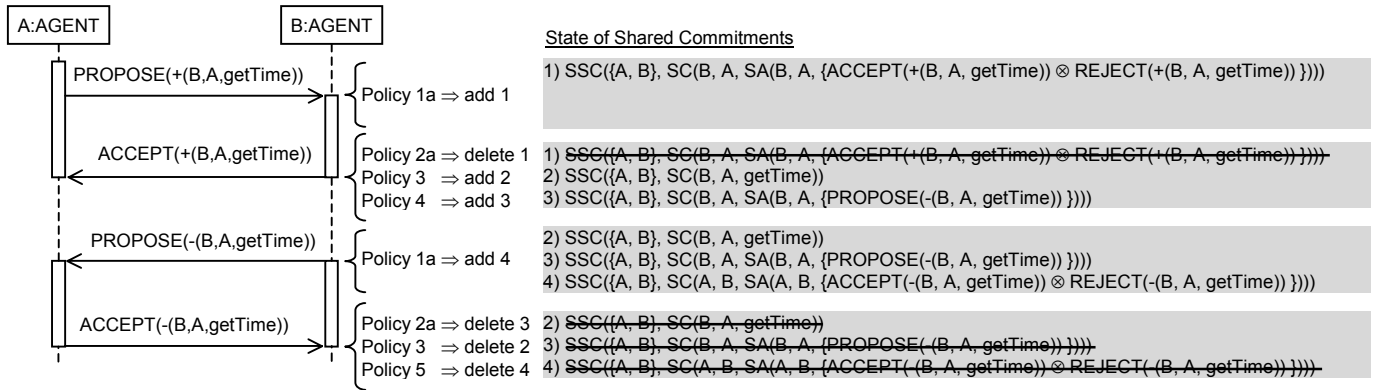


Figure 6. Interaction diagram for a conversation asking the time.

Policy 1A, which causes the addition of shared commitment number 4 (that this proposal for discharging will be replied).

As previously mentioned, that this proposal is accepted or not is for agent A to decide (e.g., the proposal may include the time obtained, or may provide a reason to forfeit performing the action). In this case, agent A graciously accepts the proposal, which triggers Policy 2A (which deletes commitment 3), Policy 3 (which deletes commitment 2), and Policy 5 (which deletes commitment 4).

At this point, all shared conversational commitments have been deleted, signaling the termination of the conversation.

7. REVISITING GOALS

The major goals of our model (which were described in sections 2, 3, and 4) were:

1. To support autonomous and heterogeneous agents in open environments we require publicly verifiable speech act semantics.
2. To support flexible conversations, we require a set of rules to specify conversational composition.
3. To integrate speech act meaning and conversational composition, we require a common set of principles from which these can be derived.

The remainder of this section elaborates on how these goals are achieved by this theory.

7.1 Message Semantics

In our model, speech act meaning is the emergent product of the following descriptive elements:

- *identity*: where a speech act is identified by their illocutionary points;
- *conversational use*: where the utterance of a speech act enables (and obligates) agents to utter other speech acts; and
- *accomplishments*: where a speech act specifies the actions it can bring about through negotiation. Agents can use speech acts in conjunction with the PFP to request the performance of actions.

To illustrate this point we describe the meaning of a proposing speech act; where its identity is that of a speech act containing the illocutionary point PROPOSE; its conversational use is that of a speech act that once uttered is expected to be replied (Policy 1A); and, its accomplishment is that of a speech act that (if replied within a specific time frame) can cause the mutual uptake in

speaker and addressee of the previously proposed operation on a social commitment (Policy 3).

Furthermore, the definitions and policies specified in our model are defined in terms of the observable characteristics and effects of communications. As such we deemed these elements to be public and thus verifiable in open environments, thus fulfilling our first goal.

7.2 Composing Conversations

Flexible conversations require the specification of rules governing their sequencing. In our model, the PFP supports a simple yet powerful method to advance the state of conversations involving social commitments to action.

As previously specified, the PFP starts with a proposal that can be accepted, rejected, or counteroffered by the addressee, or withdrawn or counteroffered by the speaker. An acceptance, rejection, or withdrawal terminates the interaction, while a counteroffer requires a reply in the form of a withdrawal, acceptance, rejection, or another counteroffer. In theory, a counteroffer can follow another counteroffer *ad infinitum*; in practice, the number of successive counteroffers might be limited by the reasoning, competence, or endurance of interacting agents. In any event, complex conversations for action are dynamically composed to suite the current requirements of the agents, thus fulfilling our second goal.

7.3 Integrating Semantics and Composition

Our model implements a novel approach to speech act semantics in which conversational composition is an integral part of speech act meaning. As such we deem our model to support our last goal, i.e., to integrate message meaning and conversational composition under a common set of principles.

8. CONCLUSION

In this paper we set out to address the problem of the lack of conversational coherence supported by current ACL semantics and CP. CP are relatively inflexible, and ACL (typically based on mental attributes) require unrealistic knowledge of the inner workings of other agents. CP and ACL have also not been combined in a way that recognizes the dependencies between the two.

We have taken the approach of using only observable behavior and the concept of shared social commitments to drive the

coherence of agent conversations for action. The choices an agent can make at any juncture in a conversation are guided by a small set of fundamental policies formally defined in our model. The atomic element of all conversations is a simple proposal followed by one of five responses (plus the possibility of no response at all): the receiver's acceptance, rejection, or counter-offer, and the sender's rejection or counteroffer. Acceptance messages imply the shared uptake of conversational and social commitments, which are the driving force behind the larger conversation modules.

In this paper we presented a first step toward a unifying model of rational agent conversations. On the one hand, the basic conversational rules in our model can be applied for the off-line design of conversation protocols that simple agents (i.e., without complex rational mechanisms) could follow. These protocols would probably not be as rich and flexible as those of rational agents, but their interactions will nevertheless be consistent. Additionally, in the case of rational agents, we foresee that occurrences of communicative acts could be dynamically qualified (e.g., ascribed with mental attitudes) according to context, thus, helping agents to select engagements that do not lead to (potentially) conflicting goals, or undesirable communicational patterns (e.g. deadlocks, livelocks), for example.

We plan to further investigate these issues. Currently we are working on the formal specification of more complex interactions (such as the Contract Net protocol), and on the implementation of an experimental engine that agents can use to support our model in an environment where agents could be engaged in multiple simultaneous conversations.

9. ACKNOWLEDGMENTS

Many thanks to Douglas Norrie, Michal Pechoucek, Martyn Fletcher, Frank Maurer, Jörg Denzinger, Brian Gaines, and the researchers at the Intelligent Systems Group at the University of Calgary for their thoughtful comments and suggestions to improve the material presented in this paper. We are also grateful to the anonymous reviewers for their suggestions to improve this paper.

This work was supported by the Natural Sciences and Engineering Council of Canada (NSERC), Smart Technologies, Inc., the Alberta Software Engineering Research Consortium (ASERC), and the Alberta Informatics Circle of Research Excellence (iCORE).

10. REFERENCES

- [1] Austin, J.L. *How to Do Things with Words*. Harvard University Press, 1962.
- [2] Bradshaw, J.M., Duffield, S., Benoit, P. & Woolley, J.D. KAoS: Toward An Industrial-Strength Open Agent Architecture. *Software Agents*, J.M. Bradshaw (ed.), AAAI Press, 1997, pp. 375-418.
- [3] Cohen, P.R. & Levesque, H.J. Rational Interaction as the Basis for communication. *Intentions in Communication*, P.R. Cohen, J. Morgan & M.E. Pollack (eds.), MIT Press, 1990, pp. 221-255.
- [4] Cohen, P.R. & Levesque, H.J. Intention is Choice with Commitment. *Artificial Intelligence*, Number 42, Elsevier Science Publishers, 1990, pp. 213-261.
- [5] Cohen, P.R. & Levesque, H.J. Teamwork. *Noûs*, Special Issue on Cognitive Science and Artificial Intelligence, Volume 25, Number 4, 1991, pp. 487-512.
- [6] Colombetti, M. A Commitment-based Approach to Agent Speech Acts and Conversations. *Fourth International Conference on Autonomous Agents*, Workshop on Agent Languages and Conversation Policies, M. Greaves, F. Dignum, J. Bradshaw, and B. Chaib-draa (eds.), Barcelona, Spain, 2000, pp. 21-29.
- [7] Conte, R. & Castelfranchi, C. *Cognitive and Social Action*. University College London Press, 1995.
- [8] Craig, R.T. & Tracy K. (eds.) *Conversational Coherence: Form, Structure, and Strategy*, Sage Publications, 1983.
- [9] Diller, A. Z. *An Introduction to Formal Methods*. John Willey & Sons, 1990.
- [10] Finin, T., Labrou, Y. & Mayfield, J. KQML as an Agent Communication Language. *Software Agents*, J.M. Bradshaw (ed.), AAAI Press, 1997, pp. 291-316.
- [11] FIPA Agent Communication Language Specifications, *Foundation for Intelligent Physical Agents*, 1997 <http://www.fipa.org/>
- [12] Greaves, M., Holmback, H. & Bradshaw, J. What is a Conversation Policy? *Third International Conference in Autonomous Agents*, Workshop on Specifying and Implementing Conversation Policies, M. Greaves & J. Bradshaw (eds.), Seattle, WA, 1999, pp. 1-9.
- [13] Huhns, M.N. & Singh, M.P. Agents and Multiagent Systems: Themes, Approaches, and Challenges. *Readings in Agents*, M.N. Huhns & M.P. Singh (eds.), Morgan Kaufmann Publishers, 1999, pp. 1-23.
- [14] Karnow, C.E.A. Liability for distributed artificial intelligences. *Berkeley Technology Law Journal*, Volume 11, Number 1, 1996, pp.147-204.
- [15] Labrou, Y. *Semantics for an Agent Communication Language*. Ph.D. Thesis, Computer Science and Electrical Engineering Department, University of Maryland, 1997.
- [16] Labrou, Y. & Finin, T. Semantics and Conversations for an Agent Communication Language. *Readings in Agents*, M.N. Huhns & M.P. Singh (eds.), Morgan Kaufmann Publishers, 1999.
- [17] Searle, J. *Speech Acts*. Cambridge University Press, 1969.
- [18] Singh, M.P. Agent Communicational Languages: Rethinking the Principles. *IEEE Computer*, Volume 31, Number 12, 1998, pp. 40-47.
- [19] Singh, M.P. A Social Semantics for Agent Communication Languages. *International Joint Conference in Artificial Intelligence*, Workshop on Agent Communication Languages, F. Dignum, B. Chaib-draa & H. Weigand (eds.), Springer-Verlag, Stockholm, Sweden, 1999.
- [20] Smith, I.A., Cohen, P.R., Bradshaw, J.M., Greaves, M. & Holmback, H. Designing conversation policies using joint intention theory. *Third International Conference on Multi-Agent Systems*, Paris, France, 1998, pp. 269-276.
- [21] Wooldridge, M. Verifiable Semantic for Agent Communication Languages. *Third International Conference on Multi-Agent Systems*, Y. Demazeau (ed.), IEEE Press, 1998.