Dynamic Semantics for Agent Communication Languages

Michael Rovatsos

Abstract: We propose dynamic semantics for agent communication languages (ACLs) as a method for tackling some of the fundamental problems associated with agent communication in open multiagent systems. Based on the idea of providing alternative semantic "variants" for speech acts and transition rules between them that are contingent on previous agent behaviour, our framework provides an improved notion of grounding semantics in ongoing interaction, a simple mechanism for distinguishing between *compliant* and *expected* behaviour, and a way to specify sanction and reward mechanisms as part of the ACL itself.

Introduction

The field of agent communication language (ACL) research has long been plagued by problems of verifiability and grounding:

- mentalistic semantics are not verifiable in open systems (and thus unreliable);
- commitment-based semantics are verifiable but lack grounding (at the time of the utterance, the semantics says nothing about what will happen).

Also, none of the existing approaches allows the ACL to specify how to *respond* to a violation of its semantics by individual agents. Thus, they fail to exploit the possibilities of sanctioning and rewarding certain behaviours in a *communication-inherent* way by modifying the future meaning of messages uttered by agents.

We propose *dynamic semantics* (DSs) for ACLs as a solution to these problems.

Commitments and grounding

Notion of *commitments* based on variation of the framework proposed by Fornara and Colombetti (2002):

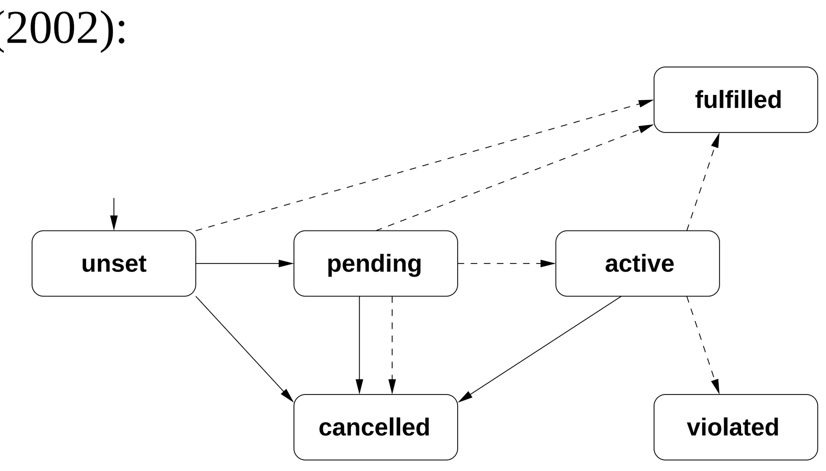

Our definition of commitment:

$$\langle \iota, \mathbf{s} : \chi \oplus \varphi \ominus \psi \rangle_t^{i \to j}$$

where

- ι is a unique *commitment identifier*,
- $\mathbf{s}$ denotes the commitment *state* (any of **u**nset, **p**ending, **a**ctive, **v**iolated, **f**ulfilled, or **c**ancelled, abbreviated by the respective initial),
- i is the *debtor*, j is the *creditor*,
- χ is the *debitum* (i.e. the proposition that i commits to making true toward j),
- φ, ψ are the *activation*/*deactivation conditions*,
- and t is the instant (in a run) at which this commitment entered its current state $\mathbf{s}$.

Example: $\langle x, \mathbf{v} : \mathit{received}(5, \$500) \oplus \mathit{received}(3, \mathit{toys}) \ominus \mathit{returned}(3, \mathit{toys}) \rangle_{12}^{3 \to 5}$

Transition rules manage the contents of commitment stores:

$$\begin{aligned}
D:\ & CS \leftarrow CS \cup \{\langle \iota, \mathbf{c} : \chi \oplus \varphi \ominus \psi \rangle_t | \langle \iota, \mathbf{s} : \chi \oplus \varphi \ominus \psi \rangle \in CS, r \models \psi, \mathbf{s} \in \{u, p, a\}, \\
& \qquad \langle \iota, \mathbf{c} : \chi \oplus \varphi \ominus \psi \rangle \notin CS\} \\
A:\ & CS \leftarrow CS \cup \{\langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle_t | \langle \iota, \mathbf{p} : \chi \oplus \varphi \ominus \psi \rangle \in CS, r \models \varphi, \langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle \notin CS\} \\
S:\ & CS \leftarrow CS \cup \{\langle \iota, \mathbf{f} : \chi \oplus \varphi \ominus \psi \rangle_t | \langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle \in CS, r \models \chi, \langle \iota, \mathbf{f} : \chi \oplus \varphi \ominus \psi \rangle \notin CS\} \\
F:\ & CS \leftarrow CS \cup \{\langle \iota, \mathbf{f} : \chi \oplus \varphi \ominus \psi \rangle_t^{i \to j} | \langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle_{t-1}^{i \to j} \in CS, r \models \mathit{Done}(i, a), \mathit{causes}(a, \chi)\} \\
V:\ & CS \leftarrow CS \cup \{\langle \iota, \mathbf{v} : \chi \oplus \varphi \ominus \psi \rangle_t^{i \to j} | \langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle_{t-1}^{i \to j} \in CS, r \models \mathit{Done}(i, a), \neg \mathit{causes}(a, \chi)\}
\end{aligned}$$

Compliance and deviance

Compliant behaviour requires that the following condition be fulfilled:

$$\forall k \leq t \left(\langle \iota, \mathbf{a} : \Gamma \rangle_k^{i \to j} \in CS \Rightarrow \langle \iota, \mathbf{f} : \Gamma \rangle_k^{i \to j} \in CS \right)$$

This can be applied to actual agent functions, thus providing a grounding for action:

$$\begin{aligned}
\mathit{compliant}(CS) := \{ g_i \in G_i(\mathit{Env}, \mathcal{A}) | \\
\forall r \sim g_i . \langle \iota, \mathbf{p} : \chi \oplus \varphi \ominus \psi \rangle^{i \to j} \in CS(r) = CS. \\
\forall r' \sqsupseteq r . \langle \iota, \mathbf{a} : \chi \oplus \varphi \ominus \psi \rangle_{|r'|}^{i \to j} \in CS(r') \Rightarrow \\
(\exists a \in Ac_i . \mathit{causes}(a, \chi) \wedge g_i(r') = a) \}
\end{aligned}$$

Distinguish **expectations** $(\iota, \mathbf{s} : \Gamma)_t^{i \to j}$ from **normative commitments** and introduce the following constructs:

$$\begin{aligned}
\lceil CS \rceil & := \{ \langle \iota, \mathbf{s} : \Gamma \rangle \in CS | \mathbf{s} \in \{\mathbf{u}, \mathbf{p}, \mathbf{a}, \mathbf{f}, \mathbf{v}\}\} \\
\lfloor CS \rfloor & := \{ (\iota, \mathbf{s} : \Gamma) \in CS | (\iota, \mathbf{s} : \Gamma) \in CS, \langle \iota, \mathbf{s}' : \Gamma \rangle \in CS, \mathbf{s}, \mathbf{s}' \in \{\mathbf{u}, \mathbf{p}, \mathbf{a}, \mathbf{f}, \mathbf{v}\}\}
\end{aligned}$$

With this, we can define **expected behaviour**:

$$\mathit{expected}(CS) := \mathit{compliant}(\lfloor CS \rfloor)$$

Dynamic semantics definitions

Static ACL semantics fragment with two alternatives for "accept":

$$RQ: \frac{\mathit{time}(t), \mathit{new}(\iota)}{\boxed{\texttt{request}(i, j, \iota : \Gamma)}} \quad CS \leftarrow CS \cup \{\langle \iota, \mathbf{u} : \Gamma \rangle_t^{i \to j}\}$$

$$RJ: \frac{\langle \iota, \mathbf{u} : \Gamma \rangle_{t'}^{j \to i} \in CS, \mathit{time}(t)}{\boxed{\texttt{reject}(i, j, \iota : \Gamma)}} \quad CS \leftarrow CS \cup \{\langle \iota, \mathbf{c} : \Gamma \rangle_t^{i \to j}\}$$

$$AC: \frac{\langle \iota, \mathbf{u} : \Gamma \rangle_{t'}^{j \to i} \in CS, \mathit{time}(t)}{\boxed{\texttt{accept}(i, j, \iota : \Gamma)}} \quad CS \leftarrow CS \cup \{\langle \iota, \mathbf{p} : \Gamma \rangle_t^{i \to j}\}$$

$$AC2: \frac{\langle \iota, \mathbf{u} : \Gamma \rangle_{t'}^{j \to i} \in CS, \mathit{time}(t)}{\boxed{\texttt{accept}(i, j, \iota : \Gamma)}} \quad CS \leftarrow CS \cup \{\langle \iota, \mathbf{p} : \Gamma \rangle_t^{i \to j}\} \cup \{(\iota, \mathbf{c} : \Gamma)_t^{i \to j}\}$$

To define DS for ACLs we now introduce a state transition system in which each state specifies an "ordinary" (static) commitment-based semantics and a "range" of agent pairs for which these semantics are assumed to apply.

A **dynamic semantics** (DS) is a structure $\langle O, S, s_0, \Delta \rangle$ where

- $O = \{o_1, o_2, \ldots, o_n\}$ a set of dialogue operators,
- $S \subseteq \wp(O)$ is a set of *semantic states* specified as subsets of dialogue operators which are valid in this state, $s_0 \in S$ is the initial semantic state,
- and the *transition relation* $\Delta \subseteq S \times \wp(\mathcal{C}) \times \wp(Ag \times Ag) \times S$ defines the transitions over S triggered by conditions expressed as elements of $\wp(\mathcal{C})$ ($\mathcal{C}$ is the set of all possible commitments).

Example:

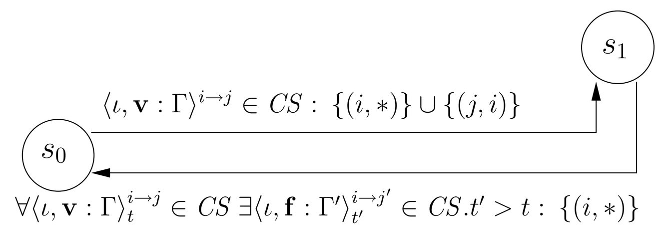

State of a dynamic semantics $\langle O, S, s_0, \Delta \rangle$ after run r with immediate predecessor r' is defined as a mapping act_r as follows:

1. $r = \varepsilon$: $act_\varepsilon(i, j) = s_0$ for all $i, j \in Ag$
2. $r \neq \varepsilon$:

$$act_r(i, j) = \begin{cases} s' & \text{if } \exists \delta = (s, c, A, s') \in \Delta . \\ & \quad (i, j) \in A(\delta, CS(r)) \\ act_{r'}(i, j) & \text{else} \end{cases}$$

Discussion

Our framework allows us to exploit **reciprocity** in responding to deviance/compliance (by using reputation-based adaptation, mutuality of expectations, and recovery mechanisms).

Possible desiderata for dynamic ACL semantics design:

- *Respect for commitment autonomy*: The semantics must not allow an agent to create a pending commitment for another agent or to violate a commitment on behalf of another agent.
- *Avoiding commitment inconsistency:* The ACL must either disallow commitment to contradictory actions or beliefs, or at least provide operators for rectifying such contradictory claims.
- *Unprejudiced judgement:* Expected behaviour prediction must not deviate from compliant behaviour prediction if deviant behaviour has not been observed so far.
- *Convergence:* The semantic state of each of the dialogue operators will remain stable after a finite number of transitions, regardless of any further agent behaviour (debatable).
- *Forgiveness:* After initial deviance, further compliant behaviour of an agent should lead to a semantic state that predicts compliant behaviour for that agent again.
- *Equality:* Unless this is required by domain-specific constraints, the same dynamics of semantics should apply to all parties involved.

Conclusion

Summary of contributions:

- Extension of commitment-based ACL semantics to provide an improved notion of *grounding* commitments in agent interaction
- Simple way of distinguishing between *compliant* and *expected* behaviour with respect to an ACL specification
- Mechanism for specifying how meaning *evolves* with agent behaviour and how this can be used to describe *communication-inherent* sanctioning and rewarding

The University of Edinburgh, School of Informatics. Enquiries to: mrovatso@inf.ed.ac.uk