Semantic Alignment of Agent Interactions through the Communication Product

Manuel Atencia and Marco Schorlemmer

IIIA, Artificial Intelligence Research Institute
CSIC, Spanish National Research Council
Bellaterra (Barcelona), Catalonia, Spain

Abstract. We provide the formal foundation of a novel approach to tackle semantic heterogeneity in multi-agent communication by looking at semantics related to interaction in order to avoid dependency on a priori semantic agreements. We do not assume existence of any ontologies, neither local to interacting agents nor external to them, and we rely only on interactions themselves to resolve terminological mismatches. In the approach taken in this paper we look at the semantics of messages that are exchanged during an interaction entirely from an interaction-specific point of view: messages are deemed semantically related if they trigger compatible interaction state transitions—where compatibility means that the interaction progresses in the same direction for each agent, albeit their partial view of the interaction (their interaction model) may be more simple than the interaction that is actually happening. Our underlying claim is that semantic alignment is often relative to the particular interaction in which agents are engaged in, and, that in such cases the interaction state should be taken into account and brought into the alignment mechanism.

1 Introduction

In multi-agent communication one usually assumes that agents use a shared terminology with the same meaning for message passing. If agents, however, are engineered separately one has to foresee that, when they interact, they will most likely make use of different terminology in their respective messages, and that, if some terms coincide, they may not have the same meaning for all agents participating in an interaction. This is the problem of semantic heterogeneity.

Over the last years various kinds of solutions have been proposed to achieve interoperability at the semantic level, which are applicable to multi-agent communication as well as to database integration, peer-to-peer systems, and the semantic web. One early solution spanning back to the early 1990s goes with agreeing upon a common ontology for the particular domain in which interoperability has to take place [9]. Each agent will have to define its own local terminology in terms of the shared ontology. In this approach, the shared ontology acts as "interlingua", which ultimately means to fall back to the single-ontology view of agent communication.

Common ontologies may be useful for stable domains and closed communities of agents, but being precise about semantics for complex domains is very expensive, and the cost of guaranteeing a global semantics for agent communication increases quickly when the number of participants grows. Current state-of-the-art approaches tackling semantic heterogeneity no more seek to agree on one shared global ontology, but instead attempt to establish correspondences between varying terminologies [10]. There exist many implemented systems, which combine several mature techniques: syntactic-based techniques such as *edit distance* or *n-gram*, structural techniques that exploit the graph structure of ontologies, or semantic-based techniques that consult external source such as upper-level ontologies, dictionaries, and thesauri [5].

In these systems, matching is generally performed at design-time, prior to integration, which means, in our case, prior to agents entering an interaction. This obviously still limits the dynamism and openness of agent communication. Also, matching is done outside the context of the interaction. Furthermore, most current ontology matching techniques follow a classical functional approach, taking two (or more) ontologies as input and producing a semantic alignment of ontological entities as output.

Recent approaches look at applying ontology matching at run-time and only between those fragments of the ontologies that are deemed relevant to the task at hand or to current interaction [11, 17]. This allows for openness and dynamism, and has the additional advantage that we do not need to access the entire ontologies (this is desirable, e.g., when ontologies constitute commercially confidential information). Despite these advantages, dynamic ontology matching techniques still follow a functional approach: when a mismatch occurs, semantic heterogeneity is solved applying current state-of-the-art ontology matching techniques, albeit only for a fragment and at run-time. Furthermore, although done in run-time and more focused on relevant bits of the ontologies, matching is still done separately from the interaction: semantic similarity continues to be established in an interaction-independent fashion, using, e.g., external sources such as WordNet [6], where synonymy between terms was determined prior to interaction and independently from it.

In this paper we provide the formal foundation for a very parsimonious approach to the problem of semantic heterogeneity in multi-agent communication with the aim of complementing the previous solutions applied so far. We claim that semantic alignment is often also relative to the particular interaction in which agents are engaged in, and, more specifically, to the particular state of the interaction. In such cases the interaction state should be taken into account and brought into the alignment mechanism. The meaning of certain terms are often very interaction-specific. For instance, the semantic similarity that exists, in the context of an auction, between the Spanish term “remate” and the English expression “winning bid” is difficult to establish if we are left to rely solely on syntactic or structural matching techniques, or on external sources such as dictionaries and thesauri. The term “remate” may have many different senses, and none of them may hint at its meaning as “winning bid.” But it actually has this very precise meaning when uttered at a particular moment of the interaction happening during an auction.

Our approach shares with that of Besana and Robertson [3] the insight that semantics is often interaction-specific. Besana and Robertson attach probabilities to meanings of terms that are determined by earlier, similar interactions. They use these a priori probabilities to predict the set of possible meanings of a message. As with our approach, meaning is defined relative to a particular interaction, but Besana and Robertson aim at reducing the search space of possible a priori mappings between ontological entities (in a classical sense), namely by assessing those ones with highest probability in the context of an interaction. We approach the semantic heterogeneity problem from a different angle and attempt to use the interaction itself to determine the semantic relationships.

Interaction-Situated Semantic Alignment

We shall address the case in which agents need to establish the semantic relationships with terminologies of other agents on the grounds of their communication within a specific interaction. We call this approach *interaction-situated semantic alignment*. This work is part of a larger research endeavour, carried out in the *OpenKnowledge* Specific Targeted Research Project (STREP) [14] and sponsored by the European Commission under its 6th Framework Program. The project aims at lowering the cost of participation in semantic-intensive distributed systems by focusing on semantics related to interaction (which are acquired at low cost during participation) and using this to avoid dependency on a priori semantic agreements.

In OpenKnowledge, the specifications of interactions—called *interaction models*—are, along with data, first-class citizens that can be shared between agents. Currently all agents participating in an OpenKnowledge interaction have to follow the same interaction model, but it is realistic to foresee the following scenarios:

- a scenario in which agents take hold only of those fragments of interaction models that concern them, e.g., when they hold only those specifications that describe the message-passing behaviour of the roles they are capable of playing in an interaction;
- a scenario in which agents interact according to slightly different versions of an original interaction model, e.g., when various agents have downloaded an interaction model in the past, which was subsequently refined by one of them (i.e., the interaction-model may have locally evolved).

In both cases above, the original messages may not mean exactly the same to interacting agents, and they all have only a partial view of the actual interaction that is happening. We see these as scenarios in which an interaction-situated semantic alignment approach as the one described in this paper may prove valuable. The second scenario is reminiscent to the ontology refinement scenario of [12]. There, McNeill tackled the problem of terminological mismatch when agents were executing plans based on slightly different ontologies. Here we deal with the problem of terminological mismatch when agents are following interactions based on slightly different interaction models.

The structure of the paper is as follows. In the next section we introduce the basic intuitions of our interaction-situated semantic alignment approach through a concrete interaction model, namely a sealed-bid auction taken from [4]. In Section 3 we formalise the concepts introduced intuitively in order to define, in Section 4, the notion of semantic equivalence as it arises in an interaction such as the one of Section 2. Section 5 concludes the paper discussing our work in progress.

2 An Example: Interaction in a Sealed-bid Auction

In a sealed-bid auction, after the auctioneer has announced the start of a round for auctioning a particular good, bidders are given a period of time to submit their bids (without other bidders knowing it). After that period, the auctioneer announces the winner, namely the bidder that submitted the highest bid. In certain cases the auctioneer may decide to withdraw a good instead (for example if no bids where submitted). Hence the interaction that unfolds is as follows: In the initial state of the interaction, bidders wait for the auctioneer to send a message announcing the *start of round* for a particular good GID at a reserve price RP with bidding time BT. This message passing causes a state transition in the interaction to a state in which bidders are allowed to send their *bids* O for good GID. From the point of view of the auctioneer, the interaction remains in this state until the bidding time BT has elapsed, in which case the interaction moves to a state in which bidding messages are no more expected and in which the auctioneer is supposed to either send a message informing the bidders that the good GID has been *sold* to bidder W for the price P, or to send a message informing that good GID has been *withdrawn*. Either of these messages makes the interaction state change to the initial state, which is also the final state in this case.

From the point of view of the bidders, however, if they have submitted a *bid* O, they consider the interaction to have changed to a state in which they cannot send bids any more, but in which they wait for a message from the auctioneer informing about the outcome of the round. Alternatively they may also assume this state transition without themselves having submitted a bid. This distinction of viewpoints of the auctioneer and the bidders is important to our approach: actual interactions, if modelled as state transition due to message passing, have in general more detail than those specified for each individual roles participating in the interaction. The actual interaction, for instance, is very dependent on the number of agents participating in it. We shall come back to this issue below when we represent interaction models by means of finite state automata.

The above interaction model for a sealed-bid auction can be formally specified in numerous ways. In Figure 1 we shows one such specification in the Lightweight Communication Calculus (LCC) [13], an executable specification language that is used to constrain interactions between agents, and which is currently used as the core interaction modelling language in the OpenKnowledge STREP [14].

An interaction model in LCC is a set of clauses, each of which defines how a role in the interaction must be performed. Roles are described by the type of role and an identifier for the individual agent undertaking that role. The definition of performance of a role is constructed using combinations of the sequence operator ('*then*') or choice operator ('*or*') to connect messages and changes of role. Messages are either outgoing to another agent in a given role ('$\Rightarrow$') or incoming from another agent in a given role ('$\Leftarrow$'). Message input/output or change of role can be governed by a constraint defined using the normal logical operators for conjunction, disjunction and negation. The LCC interaction model of Figure 1 specifies the message-passing behaviour of an agent in the role of an auctioneer and in the role of a bidder. Loops in the interaction are specified via recursive calls to subroles. Here bid_collector is such a subrole of auctioneer.

Figure 2 shows the main definitions of LCC's syntax. The details of the syntax, though, and the operational semantics of LCC lie outside the scope of this paper and are

```
a(auctioneer,A) ::
    inform(start_round(GID,BT,RP)) => a(bidder,_) <-
        good(GID), bidding_time(BT), reserve_price(RP) then
    a(bid_collector(GID,BT),A) then
    ( inform(sold(GID,P,W)) => a(bidder,_) <- winner(W,P) or
      inform(withdrawn(GID)) => a(bidder,_) <- not winner(_,_) ) then
    a(autioneer,A)

a(bid_collector(GID,BT),A) ::
    timeout(BT) or
    ( record_bid(O,B) <- commit(bid(GID,O)) <= a(bidder,B) then
      a(bid_collector(GID,BT),A) )

a(bidder,B) ::
    inform(start_round(GID,BT,RP)) <= a(autioneer,A) then
    ( commit(bid(GID,O)) => a(bid_collector(GID,BT),A) <- make_bid(GID,O,RP) or
      null <- not make_bid(GID,O,RP) ) then
    ( i_won(GID,P) <- inform(sold(GID,P,B)) <= a(auctioneer,A) or
      i_lost(GID,P,W) <- inform(sold(GID,P,W)) <= a(auctioneer,A) or
      no_winner(GID) <- inform(withdrawn(GID)) <= a(auctioneer,A) ) then
    a(bidder,B)
```

Fig. 1. LCC clauses specifying the interaction models of roles auctioneer (including its subrole bid_collector) and bidder

$$
\begin{aligned}
\mathit{Interaction\_Model} &:= \{\mathit{Clause}, \ldots\} \\
\mathit{Clause} &:= \mathit{Agent}\,\texttt{::}\,\mathit{Ev} \\
\mathit{Agent} &:= \texttt{a(}\mathit{Role}\texttt{,}\ \mathit{Id}\texttt{)} \\
\mathit{Ev} &:= \mathit{Agent} \mid \mathit{Message} \mid \mathit{Ev}\ \texttt{then}\ \mathit{Ev} \mid \mathit{Ev}\ \texttt{or}\ \mathit{Ev} \mid \mathit{Ev}\ \texttt{par}\ \mathit{Ev} \mid \texttt{null} \leftarrow C \mid \texttt{timeout}(N) \\
\mathit{Message} &:= M\texttt{=>}\mathit{Agent} \mid M\texttt{=>}\mathit{Agent}\ \texttt{<-}C \mid M\texttt{<=}\mathit{Agent} \mid C\texttt{<-}\ M\texttt{<=}\mathit{Agent} \\
C &:= \mathit{Term} \mid C\ \texttt{and}\ C \mid C\ \texttt{or}\ C \\
\mathit{Role} &:= \mathit{Term} \\
M &:= \mathit{Term}
\end{aligned}
$$

Where `null` denotes an event which does not involve message passing; Term is a structured term (e.g., a Prolog term); N is a variable; and Id is either a variable or a unique identifier for an agent.

Fig. 2. Syntax of LCC interaction models

given elsewhere [13]. But in order to help in the broad understanding of the semantics of LCC, we have introduced in the above intuitive description of the interaction all relevant variables occurring in the specification, and we have also emphasised those words that constitute the messages.

2.1 Interaction State Transitions

An alternative way to specify interaction models is by means of finite state automata, which is the formalism that we will be using in this paper. This is the way, for instance, in which particular *scenes* (which are bounded scopes of interaction) are specified for electronic institutions [4]. Figure 3 illustrates the message-passing behaviour of an agent in the role of an auctioneer, and corresponds to the first two clauses of Figure 1; Figure 4 illustrates the message-passing behaviour of an agent in the role of a bidder, and corresponds to the third clause of Figure 1. Transitions between states are labelled by means of illocutions, which are tuples consisting of an illocutionary particle, the identifier of the sender together with the role it is playing, the identifier of the receiver together with the role it is playing, the content of the message uttered, and a time stamp. In this paper we shall ignore this last component for the ease of presentation. We may label transitions also with a *timeout* (see, e.g., Figure 3), or with a λ (see, e.g., Figure 4) denoting state transitions not caused by message passing. An arc labelled with $v \mid w$ replaces two arcs. Variables in messages are written in uppercase letters and get their values in those illocutions in which they occur preceded by a question mark (?), and these values are subsequently used in those illocutions in which the corresponding variable occurs preceded by an exclamation mark (!).

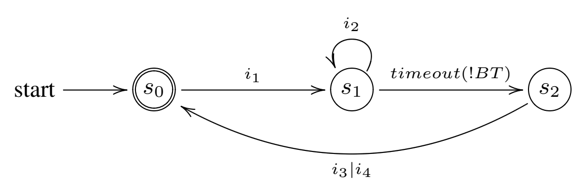

$$i_1 = \langle inform, (?A : auctioneer), (?B : bidder), start\_round(?GID, ?BT, ?RP)\rangle$$
$$i_2 = \langle commit, (!B : bidder), (!A : auctioneer), bid(!GID, ?O)\rangle$$
$$i_3 = \langle inform, (!A : auctioneer), (!B : bidder), sold(!GID, ?P, ?W)\rangle$$
$$i_4 = \langle inform, (!A : auctioneer), (!B : bidder), withdrawn(!GID)\rangle$$

Fig. 3. Interaction model for the auctioneer role

As hinted before, when auctioneers and bidders interact by message passing, an interaction unfolds which contains more detail than the ones specified in Figures 3 or 4. These interaction models capture namely only a partial view of the actual *global* interaction, the view from the perspective of an auctioneer and of a bidder, respectively.

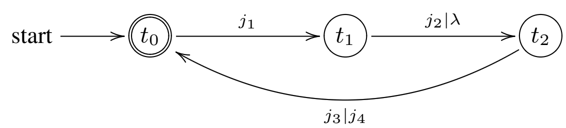

$$j_1 = \langle inform, (?A : auctioneer), (?B : bidder), start\_round(?GID, ?BT, ?RP)\rangle$$
$$j_2 = \langle commit, (!B : bidder), (!A : auctioneer), bid(!GID, ?O)\rangle$$
$$j_3 = \langle inform, (!A : auctioneer), (!B : bidder), sold(!GID, ?P, ?W)\rangle$$
$$j_4 = \langle inform, (!A : auctioneer), (!B : bidder), withdrawn(!GID)\rangle$$

Fig. 4. Interaction model for the bidder role

Actually, neither needs to be aware of the model followed by the other for the interaction to unfold correctly in its totality. In general, two (or more agents) are capable of interacting following separate interaction models if their states are assumed to be projections of states of a global interaction—which in general is not known to each of the agents—and each state transition that separate agents follow when an illocution is uttered, has a corresponding state transition in the global interaction. Figure 5 shows the global interaction model for a scenario with one auctioneer a and one bidder b.

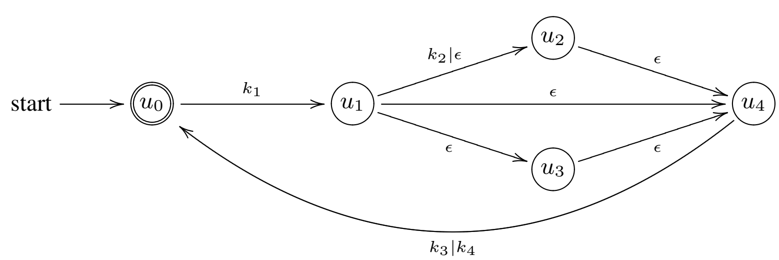

$$k_1 = \langle inform, (a : auctioneer), (b : bidder), start\_round(?GID, ?BT, ?RP)\rangle$$
$$k_2 = \langle commit, (b : bidder), (a : auctioneer), bid(!GID, ?O)\rangle$$
$$k_3 = \langle inform, (a : auctioneer), (b : bidder), sold(!GID, ?P, ?W)\rangle$$
$$k_4 = \langle inform, (a : auctioneer), (b : bidder), withdrawn(!GID)\rangle$$

Fig. 5. Global interaction model for one agent in the auctioneer and one agent in the bidder role

Ideally, a global interaction model matches together all messages occurring in compatible illocutions of role interaction models, i.e., illocutions with the same illocutionary particle, sender, and receiver, and that trigger the same state transition. In addition, each actual state of the global interaction should have a corresponding state in each of the role interaction models. This means that the states of the interaction models in Figures 3 and 4 are projections of states of a global interaction model such as the one shown in Figure 5 (actually the global interaction model is more complex, as we shall see later in Section 3; we have simplified it here for ease of explanation). Observe, for instance, transition k_2 from state u_1 to u_2 in the global interaction. The bidder considers that the interaction changes its state from t_1 to t_2 when it utters the illocution j_2 (which corresponds to k_2 in the global interaction model), while the auctioneer does not perceive this as a state change (illocution i_2 in the auctioneer's interaction model) and considers that the interaction remains in state s_1. Therefore, global interaction states u_1 and u_2 both project onto s_1 for the auctioneer, while they project onto t_1 and t_2, respectively, for the bidder. The bidder may also consider the interaction state to change without message passing (λ-transition). Consequently, this transition is reflected in the global interaction as an ϵ-transition from u_1 to u_2, although there is no corresponding arc in the auctioneer's interaction model. (We provide a precise account on ϵ-transitions in Section 3). The auctioneer does not distinguish any state change.

2.2 Aligning while Interacting

This fact is what we shall exploit for solving mismatch and semantic heterogeneity when agents use different vocabularies in message-passing. A Spanish-speaking bidder, for instance, with its interaction model labelled using Spanish auction terminology and participating in an auction managed by an English-speaking auctioneer could infer the semantic alignment existing between its Spanish terminology and the English one by the fact that interaction states followed by an auctioneer and a bidder are projections of an actual interaction generally unknown by participants in the interaction, but in which auctioneer and bidder participate and move between states together by message passing.

Imaging now the interaction model of Figure 4, but for a Spanish-speaking bidder, with illocutions given below:

$$j_1 = \langle inform, (?A : auctioneer), (?B : bidder), nueva\_ronda(?GID, ?BT, ?RP)\rangle$$
$$j_2 = \langle commit, (!B : bidder), (!A : auctioneer), postura(!GID, ?O)\rangle$$
$$j_3 = \langle inform, (!A : auctioneer), (!B : bidder), remate(!GID, ?P, ?W)\rangle$$
$$j_4 = \langle inform, (!A : auctioneer), (!B : bidder), sin\_ganador(!GID)\rangle$$

The Spanish-speaking bidder initially expects a "nueva_ronda" message from the auctioneer. The English-speaking auctioneer initially is supposed to broadcast a "start_round" message to bidders. When this illocution is uttered the Spanish-speaking bidder may safely assume that "start_round" means "nueva_ronda", which makes the interaction change to the state in which the English-speaking auctioneer expects "bid" messages from buyers and the Spanish-speaking bidder is supposed to either send a "postura" or change state without sending or receiving any message. Consequently, if "postura" is uttered the English-speaking auctioneer can safely assume that "postura" means "bid".

Notice that these equivalences stem from the assumption that auctioneer and bidder are always in the same state of the global interaction and follow the same state transition when a illocution is uttered (see Figure 5). Or, more precisely, their local states in each of their own interaction models are projections from the same state of the actual global interaction. In the next two sections we formalise this approach, providing a definition of "global interaction model" through the idea of a product of interaction models, what we call the *communication product*. This product represents all compatible state transitions and, from this product we define the notion of semantic equivalence that arises from compatible interactions. We shall treat messages as propositions, however, i.e., as grounded atomic sentences, leaving the generalisation to first-order sentences for future work.

3 Formalising Interaction Models and their Relations

We model a multi-agent system as a set MAS of agents. Each agent in MAS has a unique identifier and may take one (or more) roles in the context of an interaction. Let $Role$ be the set of roles and Id the set of agent identifiers. We write $(id : r)$, with $r \in Role$ and $id \in Id$, for the agent in MAS with identifier id playing role r.

Each agent is able to communicate by sending messages from a set M, which is local to the agent. We assume that a set $\mathfrak{I}_{\mathrm{P}}$ of *illocutionary particles* is shared by all agents (e.g., those of KQML [7] or FIPA ACL [8]).

Definition 1. *Given a non-empty set M of messages, the set of* illocutions generated by *M, denoted by $\mathfrak{I}(M)$, is the set of all tuples $\langle \iota, (id : r), (id' : r'), m \rangle$ with $\iota \in \mathfrak{I}_{\mathrm{P}}$, $m \in M$, and $(id : r), (id' : r')$ agents such that $id \neq id'$.*

If $i = \langle \iota, (id : r), (id' : r'), m \rangle$ is an illocution then $(id : r)$ is the *sender* of i and $(id' : r')$ is the *receiver* of i.

3.1 Interaction Models

We model an interaction model as a (partial) deterministic finite-state automaton whose transitions are labelled either with illocutions, or with special transitions such as, for example, timeouts or null transitions (λ-transitions):

Definition 2. *An* interaction model *is a tuple* $\mathrm{IM} = \langle Q, q^0, F, M, C, \delta \rangle$ *where:*

- *Q is a finite set of* states*,*
- *q^0 is a distinguished element of Q named the* initial state*,*
- *F is a non-empty subset of Q which elements are called* final states*,*
- *M is a finite non-empty set of* messages*,*
- *C is a finite set of* special transitions*, and*
- *δ is a partial function from $Q \times (\mathfrak{I}(M) \cup C)$ to Q called the* transition function*.*

Remark 1. If $\mathrm{IM} = \langle Q, q^0, F, M, C, \delta \rangle$ is an interaction model, IM is associated with an automaton, $Aut(\mathrm{IM}) = \langle Q, q^0, F, \Sigma, \delta \rangle$, where $\Sigma = \mathfrak{I}(M) \cup C$. In the remainder of the paper, we refer indifferently to either IM or $Aut(\mathrm{IM})$ if no confusion is likely.

Example 1. Figures 6 and 7 illustrate interaction models for the English-speaking auctioneer role and the Spanish-speaking bidder role, respectively.[1] They are equivalent to Figures 3 and 4, except for the propositional messages, and also the use of Spanish terms for the bidder.

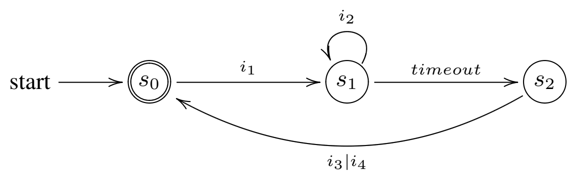

$$i_1 = \langle inform, (a : auctioneer), (b : bidder), start\_round\rangle$$
$$i_2 = \langle commit, (b : bidder), (a : auctioneer), bid\rangle$$
$$i_3 = \langle inform, (a : auctioneer), (b : bidder), sold\rangle$$
$$i_4 = \langle inform, (a : auctioneer), (b : bidder), withdrawn\rangle$$

Fig. 6. Interaction model for the auctioneer role

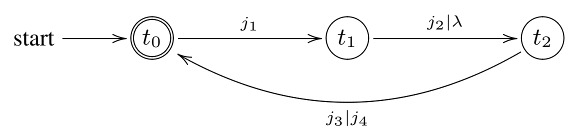

$$j_1 = \langle inform, (a : auctioneer), (b : bidder), nueva\_ronda\rangle$$
$$j_2 = \langle commit, (b : bidder), (a : auctioneer), postura\rangle$$
$$j_3 = \langle inform, (a : auctioneer), (b : bidder), remate\rangle$$
$$j_4 = \langle inform, (a : auctioneer), (b : bidder), sin\_ganador\rangle$$

Fig. 7. Interaction model for the bidder role

3.2 The Communication Product

As hinted in Section 2, we shall use the algebraic product of two interaction models to capture all possible interactions between agents. In general, a product of two objects is

[1] An arc from p to q labelled with $v \mid w$ replaces two arcs, so it means $p = \delta(q, v)$ and also $p = \delta(q, w)$.

the natural algebraic construction that represents all possible behaviours of the combination of those two objects. The *communication product* defined below, thus, captures the global interaction with respect to the message-passing behaviour of agents of two interaction models. It is not an unconstrained product, since it takes into account the compatibility of illocutions and special transitions in terms of illocutionary particles, senders, and receivers.

Definition 3. *Given two interaction models* IM_1 *and* IM_2*,* $\mathrm{IM}_k = \langle Q_k, q_k^0, F_k, M_k, C_k, \delta_k\rangle$ $(k = 1, 2)$*, the* communication product *of* IM_1 *and* IM_2*, denoted by* $\mathrm{IM}_1 \otimes \mathrm{IM}_2$*, is the interaction model* $\langle Q, q^0, F, M, C, \delta\rangle$ *where:*

- Q *is the Cartesian product of* Q_1 *and* Q_2*; specifically,* Q *states are all possible ordered pairs* $\langle q_1, q_2\rangle$ *with* $q_1 \in Q_1$ *and* $q_2 \in Q_2$*,*
- *the initial state* q^0 *is the pair* $\langle q_1^0, q_2^0\rangle$*,*
- F *is the Cartesian product of* F_1 *and* F_2*,*
- M *the Cartesian product of* M_1 *and* M_2*,*
- C *is the singleton set* $\{\epsilon\}$*; and finally*
- *the transition function* δ *is defined as follows:* $\langle q_1', q_2'\rangle = \delta(\langle q_1, q_2\rangle, \sigma)$ *if*
    - σ *is an illocution* $\langle \iota, (id : r), (id' : r'), \langle m_1, m_2\rangle\rangle$ *and for every* $k \in \{1, 2\}$*,* $q_k' = \delta_k(q_k, \langle \iota, (id : r), (id' : r'), m_k\rangle)$*,*
    - $\sigma = \epsilon$ *and there exist* $c_1 \in C_1$ *and* $c_2 \in C_2$ *such that* $q_k' = \delta_k(q_k, c_k)$ *for every* $k \in \{1, 2\}$*,*
    - $\sigma = \epsilon$ *and for some* $k \in \{1, 2\}$ *there exists* $c \in C_k$ *such that* $q_k' = \delta_k(q_k, c)$ *and* $q_l' = q_l$ *with* $l \in \{1, 2\}$ *and* $l \neq k$*.*

Let IM_1 and IM_2 be two interaction models. The communication product $\mathrm{IM}_1 \otimes \mathrm{IM}_2$ is associated with a finite automaton with ϵ-moves in a natural way. The language generated by $\mathrm{IM}_1 \otimes \mathrm{IM}_2$ is the language generated by this automaton.

Example 2. The communication product of interaction models for auctioneer role and bidder role is depicted in Figure 8. Discontinuous lines are arcs that are not involved in the language generated by the communication product. Notice that this diagram without discontinuous lines fits with the diagram depicted in Figure 5.

4 Semantic Alignment in Interaction Models

Being a model of all compatible interactions of varying interaction models, the communication product is the place to look at the semantic relations between messages. From a theoretical point of view, in order to establish these relations, we look at the language generated by the communication product. Messages of different interaction models are semantically related if they are paired in illocutions whose utterance make the interaction reach a final state (i.e., make the interaction succeed) according to the global interaction determined by the communication product. This is formally given below. We use '$\sqsubseteq$' to denote semantic subsumption of messages, and use '$\sqcup$' to denote disjunction. Semantic equivalence between messages, denoted with '$\equiv$', arises when they subsume each other. We also pair messages with natural numbers to keep syntactically equivalent messages separate, as they may not be semantically equivalent.

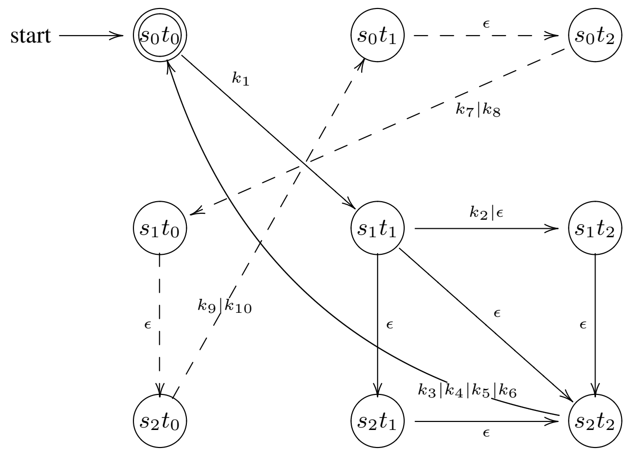

$$
\begin{aligned}
k_1 &= \langle inform, (a : auctioneer), (b : bidder), \langle start\_round, nueva\_ronda\rangle\rangle \\
k_2 &= \langle commit, (b : bidder), (a : auctioneer), \langle bid, postura\rangle\rangle \\
k_3 &= \langle inform, (a : auctioneer), (b : bidder), \langle sold, remate\rangle\rangle \\
k_4 &= \langle inform, (a : auctioneer), (b : bidder), \langle sold, sin\_ganador\rangle\rangle \\
k_5 &= \langle inform, (a : auctioneer), (b : bidder), \langle withdrawn, remate\rangle\rangle \\
k_6 &= \langle inform, (a : auctioneer), (b : bidder), \langle withdrawn, sin\_ganador\rangle\rangle
\end{aligned}
$$

$$
\begin{aligned}
k_7 &= \langle inform, (a : auctioneer), (b : bidder), \langle start\_round, remate\rangle\rangle \\
k_8 &= \langle inform, (a : auctioneer), (b : bidder), \langle start\_round, sin\_ganador\rangle\rangle \\
k_9 &= \langle inform, (a : auctioneer), (b : bidder), \langle sold, nueva\_ronda\rangle\rangle \\
k_{10} &= \langle inform, (a : auctioneer), (b : bidder), \langle withdrawn, nueva\_ronda\rangle\rangle
\end{aligned}
$$

Fig. 8. The communication product

Definition 4. *Let* IM_1 *and* IM_2 *be two interaction models,* $\mathrm{IM}_k = \langle Q_k, q_k^0, F_k, M_k, C_k, \delta_k \rangle$ $(k = 1, 2)$*. Let* $m \in M_1$ *and* $m^1, \ldots, m^n \in M_2$*. We write:*

$$\langle 1, m \rangle \sqsubseteq \langle 2, m^1 \rangle \sqcup \cdots \sqcup \langle 2, m^n \rangle$$

if for all strings x *accepted by the communication product* $\mathrm{IM}_1 \otimes \mathrm{IM}_2$*, if the illocution* $\langle \iota, (id : r), (id' : r'), \langle m, m' \rangle \rangle$ *appears in* x *then* $m' = m^i$ *for some* $i \in \{1, \ldots, n\}$.

Analogously, it is defined:

$$\langle 2, m \rangle \sqsubseteq \langle 1, m^1 \rangle \sqcup \cdots \sqcup \langle 1, m^n \rangle$$

We can also establish relationships among messages with regard to a specific illocution particle.

Definition 5. *Let* IM_1 *and* IM_2 *be two interaction models,* $\mathrm{IM}_k = \langle Q_k, q_k^0, F_k, M_k, C_k, \delta_k \rangle$ $(k = 1, 2)$*. Let* $m \in M_1$ *and* $m^1, \ldots, m^n \in M_2$*. Let* ι *be an illocution particle. We write:*

$$\langle 1, m \rangle \sqsubseteq_\iota \langle 2, m^1 \rangle \sqcup \cdots \sqcup \langle 2, m^n \rangle$$

if for all strings x *accepted by the communication product* $\mathrm{IM}_1 \otimes \mathrm{IM}_2$*, if the illocution* $\langle \iota, (id : r), (id : r'), \langle m, m' \rangle \rangle$ *appears in* x *then* $m' = m^i$ *for some* $i \in \{1, \ldots, n\}$.

Analogously, it is defined:

$$\langle 2, m \rangle \sqsubseteq_\iota \langle 1, m^1 \rangle \sqcup \cdots \sqcup \langle 1, m^n \rangle$$

Example 3. In our example, we have the following relationships among messages (without pairing them with natural numbers because messages of the auctioneer and the bidder are disjoint):

$$\begin{aligned}
start_round &\equiv nueva_ronda \\
bid &\equiv postura \\
sold &\sqsubseteq remate \sqcup sin_ganador \\
withdrawn &\sqsubseteq remate \sqcup sin_ganador \\
remate &\sqsubseteq sold \sqcup withdrawn \\
sin_ganador &\sqsubseteq sold \sqcup withdrawn
\end{aligned}$$

4.1 Converging to a Semantic Alignment

As said before, interaction models specify the space of interactions that are allowed, and its communication product captures the entire space of actual interactions when combining particular ones. The above semantic relationships are, thus, those justified by the entire space of actual interactions. This product, however, is obviously not accessible to agents in general, which may only be aware of their local interaction model. It is therefore necessary to provide agents with the mechanism to somehow discover the above semantic relationship while interactions unfold—in the sort of manner as intuitively described for our example in Section 2.2—assuming that for all agents participating in the interaction, the state of the interaction they perceive stems from the actual global state (i.e., their locally managed states are projections of the actual global state), and this throughout the entire interaction.

In [1] we described an alignment process by which two agents establish the semantic relationship between terms of their respective vocabularies based on the assumption that mismatching terms describe a partial perspective of a shared physical environment state, a state that is not accessible (i.e., completely and faithfully perceived) to any of the two agents. As agents go through more and more states of the environment, the semantic alignment between their vocabularies is further and further refined. In the scenario described in this paper agents do not share a physical environment such as in [1], but they share the same interaction. Hence their "environment" is captured by the communication product that captures the entire space of actual interactions, but which is not accessible to agents in general. An uttered illocution, though, provides a "description" of the interaction state, because its utterance "means" that the illocution was allowed in the current interaction state according to the partial perspective of the uttering agent. An agent receiving the illocution can now establish a semantic alignment based on the assumption that both agents where sharing the same interaction state.

Providing a detailed computational mechanism by which agents gradually approximate the set of semantic relationships that arise during an interaction is subject of our current work in progress. The aim of this paper was to first provide the required theoretical foundation to be able to specify such mechanism in a sound manner. We are certain, however, that such gradual approximation is theoretically feasible, because first, the communication product can be seen as defining as information channel in channel theory, the mathematical theory of information flow put forward by Barwise and Seligman [2]; and second, in [15, 16, 1] it is shown how semantic alignment can be seen as a process of information-channel refinement.

5 Conclusion

In this paper we have laid the formal foundations for a novel approach to tackle the problem of semantic heterogeneity in the context of multi-agent communication. We look at the semantics of messages from an interaction-based point of view, as it arises in the context of interaction models. Messages are deemed semantically related if they trigger compatible interaction state transitions—where compatibility here means that the interaction progresses in the same direction for each agent, albeit their view of the interaction (their interaction model) may be more constrained than the interaction that is actually happening.

One advantage of this approach is that it takes into account meaning that is very interaction-specific and which cannot be derived from sources that are external to the interaction. In this sense we see it as a complementary approach to current state-of-the-art semantic alignment techniques as it may provide valuable information for pruning the search space or disambiguating the results of candidate semantic alignments computed with toady's ontology-matching technology.

From a formal point of view, the formalisation of what semantic alignment of messages means in the context of interaction models yields a notion of communication product of interaction models which captures all possible compatible message-passing behaviours of interacting agents with mismatching message terminology. From a conceptual perspective about what a message means in the context of an interaction, and of what semantic equivalence is, we found that, by developing the approach described in this paper, we encountered new questions to explore, e.g., how interaction-specific, and even illocution-specific semantic relationships might be. Definition 4 defines semantic relationship relative to the interaction in which message are uttered but Definition 5 does the same also relative to the illocution messages are part of. This view would allow a term to be more general than another when uttered together with one kind of illocutionary particle or more specific when uttered together some other kind of illocutionary particle.

Finally, as we have said, we are certain that the very same interaction that unfolds during agent communication may be used to approximate the semantic relationships underlying the interaction, and which we have modelled as a communication product of interaction models. We have already formalised the idea of semantic alignment as information-channel refinement in our previous work, and we are currently looking at how this translate to an interaction-situated semantic alignment approach.

Acknowledgements. This work is supported under the UPIC project, sponsored by Spain's Ministry of Education and Science under grant number TIN2004-07461-C02-02 and under the OpenKnowledge Specific Targeted Research Project (STREP), sponsored by the European Commission under contract number FP6-027253. The OpenKnowledge STREP comprises the Universities of Edinburgh, Southampton, and Trento, the Open University, the Free University of Amsterdam, and the Spanish National Research Council (CSIC). M. Schorlemmer is also supported by a *Ramón y Cajal* Research Fellowship from Spain's Ministry of Education and Science, which is partially funded by the European Social Fund.

References

1. M. Atencia and M. Schorlemmer. A formal model for situated semantic alignment. In E. H. Durfee and M. Yokoo, editors, *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems. May 14–18, 2007. Honolulu, Hawai'i*, pages 1270–1277. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 2007.
2. J. Barwise and J. Seligman. *Information Flow: The Logic of Distributed Systems*, volume 44 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1997.
3. P. Besana and D. Robertson. Probabilistic dialogue models for dynamic ontology mapping. In P. C. G. da Costa, K. B. Laskey, K. J. Laskey, F. Fung, and M. Pool, editors, *URSW '06 Uncertainty Reasoning for the Semantic Web - Volume 2. Proceedings of the Second ISWC Workshop on Uncertainty Reasoning for the Semantic Web*, volume 218 of *CEUR Workshop Proceedings*. CEUR-WS, 2006.
4. M. Esteva. *Electronic Institutions: From Specification to Development*, volume 19 of *Monografies de l'Institut d'Investigació en Itel·ligència Artificial*. CSIC, 2003.

5. J. Euzenat and P. Shvaiko. *Ontology Matching*. Springer, 2007.
6. C. Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.
7. T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. M. Bradshaw, editor, *Software Agents*. MIT Press, 1997.
8. Foundation for Intelligent Physical Agents. FIPA 97 specification, part 2: Agent communication language. Technical report, Foundation for Intelligent Physical Agents, 1997.
9. T. Gruber. A translation approach for portable ontology specifications. *Knowledge Engineering*, 5(2):199–220, 1993.
10. Y. Kalfoglou and M. Schorlemmer. Ontology mapping: The state of the art. *The Knowledge Engineering Review*, 18(1):1–31, 2003.
11. V. López, M. Sabou, and E. Motta. PowerMap: Mapping the real semantic web on the fly. In I. F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, and L. Aroyo, editors, *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, volume 4273 of *Lecture Notes in Computer Science*, pages 414–427. Springer, 2006.
12. F. McNeill. *Dynamic Ontology Refinement*. PhD thesis, School of Informatics, The University of Edinburgh, 2006.
13. D. Robertson. Multi-agent coordination as distributed logic programming. In B. Demoen and V. Lifschitz, editors, *Logic Programming. 20th International Conference, ICLP 2004*, volume 3132 of *Lecture Notes in Computer Science*, pages 416–430. Springer, 2004.
14. D. Robertson, F. Giunchiglia, F. van Harmelen, M. Marchese, M. Sabou, M. Schorlemmer, N. Shadbolt, R. Siebes, C. Sierra, C. Walton, S. Dasmahapatra, D. Dupplaw, P. Lewis, M. Yatskevich, S. Kotoulas, A. P. de Pinninck, and A. Loizou. Open Knowledge: Semantic webs through peer-to-peer interaction. Project Manifesto available at http://www.openk.org, 2006.
15. M. Schorlemmer and Y. Kalfoglou. Progressive ontology alignment for meaning coordination: An information-theoretic foundation. In F. Dignum et al., editors, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, Utrecht, The Netherlands, July 25–29, 2005*, pages 737–744. ACM Press, 2005.
16. M. Schorlemmer, Y. Kalfoglou, and M. Atencia. A formal foundation for ontology-alignment interaction models. *International Journal on Semantic Web and Information Systems*, 3(2):50–68, 2007.
17. J. van Diggelen, R. Beun, F. Dignum, R. van Eijk, and J.-J. Meyer. Ontology negotiation: Goals, requirements and implementation. *International Journal of Agent-Oriented Software Engineering*, 1(1):63–90, 2007.