Semantic Alignment in the Context of Agent Interactions

Marco Schorlemmer and Manuel Atencia

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 constrained than the actual interaction that is 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 [8, 10]. 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 the cost of being precise about semantics for complex domains is prohibitively high and the cost of ensuring an individual, absolute semantics for agent communication rises rapidly as more participants take part. 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 [9]. 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 [17].

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, 18]. 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. WordNet [5], 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 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.

The approach taken in this paper is very close in spirit to that of Besana and Robertson [3] where meanings of terms may have prior probabilities determined by earlier, similar interactions. Besana and Robertson 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 aiming at reducing the search space of possible a priori mappings between ontological entities, namely by assessing those ones with highest probability in the context of an interaction. We approach semantic heterogeneity from a different angle and attempt to use the interaction itself to do the semantic alignment.

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], 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, 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 forsee:

- a scenario in which agents take hold only of the fragment of the interaction models that concern to them, e.g., when they hold only those specifications that describe the message-passing behaviour of the roles they are capable of playing;
- a scenario in which various versions of an original interaction model are followed by interacting agents, i.e., agents might have downloaded an interaction model in the past, which has subsequently been refined in some way (i.e., the interaction-model may have evolved).

In both cases above, the original messages may not mean 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 tackle 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. In Section 5 we situate our work in the broader picture that sees semantic alignment as a particular case of information-channel refinement. Section 6 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 announces 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 ID 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 sent an offer. 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 machines.

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], the executable interaction-model specification language that is currently used as the core interaction-model language in the OpenKnowledge STREP [14] (see Figure 2 for the definition of LCC's syntax). It specifies the message-passing behaviour of an agent in role of an auctioneer and 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.

A detailed description of LCC lies outside the scope of this paper, 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 machines, which will be the formalism that we will mainly use in this paper. This is the way, for instance, in which particular *scenes* (which are bounded scopes of interaction) are

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

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

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

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

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

Fig. 2. Syntax of LCC interaction models

specified for electronic institutions [4]. Figures 3 and 4 illustrate the message-passing behaviour of an agent in the role of an auctioneer and in the role of a bidder, respectively. 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 label transitions also with timeouts (see i_3 in Figure 3), or with a λ (see j_3 in Figure 4) denoting a state transition not caused by message passing. Variables in messages 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 (!).

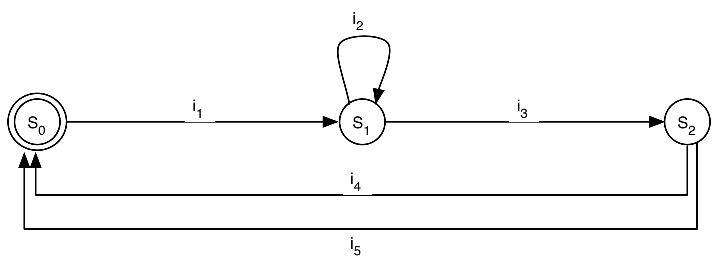

$$i_1 = \langle inform, \langle auctioneer, ?A\rangle, \langle bidder, ?B\rangle, start\_round(?GID, ?BT, ?RP)\rangle$$
$$i_2 = \langle commit, \langle bidder, !B\rangle, \langle auctioneer, !A\rangle, bid(!GID, ?O)\rangle$$
$$i_3 = timeout(!BT)$$
$$i_4 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, sold(!GID, ?P, ?W)\rangle$$
$$i_5 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, withdrawn(!GID)\rangle$$

Fig. 3. Interaction model for the autioneer 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. 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

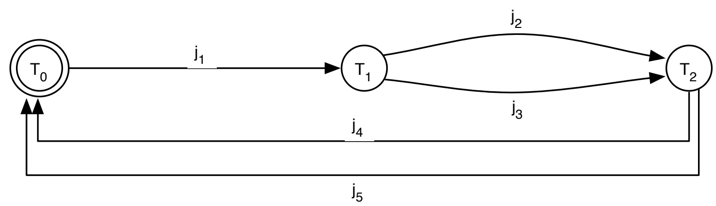

$$j_1 = \langle inform, \langle auctioneer, ?A\rangle, \langle bidder, ?B\rangle, start\_round(?GID, ?BT, ?RP)\rangle$$
$$j_2 = \langle commit, \langle bidder, !B\rangle, \langle auctioneer, !A\rangle, bid(!GID, ?O)\rangle$$
$$j_3 = \lambda$$
$$j_4 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, sold(!GID, ?P, ?W)\rangle$$
$$j_5 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, withdrawn(!GID)\rangle$$

Fig. 4. Interaction model for the bidder role

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.

On the one hand the global interaction joins together all illocutions, timeouts and λ-transitions occurring in role interaction models, since it has to capture all potential state transitions. On the other hand 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 the global interaction model in Figure 5. Observe, for instance, transitions k_2 and k_3 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 (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 j_3). Consequently, this transition is reflected in the global interaction (λ-transition k_3), although there is no corresponding transition in the auctioneer's interaction model. The auctioneer does not distinguish any state change. Hence, each local transition has a corresponding transition in the global interaction, while each global state has a corresponding state in the local interaction.

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

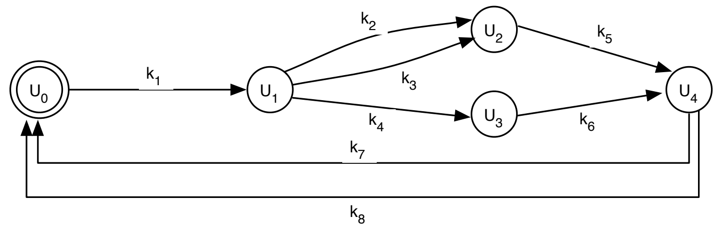

$$k_1 = \langle inform, \langle auctioneer, a\rangle, \langle bidder, b\rangle, start\_round(?GID, ?BT, ?RP)\rangle$$
$$k_2 = \langle commit, \langle bidder, b\rangle, \langle auctioneer, a\rangle, bid(!GID, ?O)\rangle$$
$$k_3 = \lambda$$
$$k_4 = timeout(t)$$
$$k_5 = timeout(t)$$
$$k_6 = \lambda$$
$$k_7 = \langle inform, \langle auctioneer, a\rangle, \langle bidder, b\rangle, sold(!GID, ?P, ?W)\rangle$$
$$k_8 = \langle inform, \langle auctioneer, a\rangle, \langle bidder, b\rangle, withdrawn(!GID)\rangle$$

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

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:

$$j_1 = \langle inform, \langle auctioneer, ?A\rangle, \langle bidder, ?B\rangle, nueva\_ronda(?GID, ?BT, ?RP)\rangle \quad (1)$$
$$j_2 = \langle commit, \langle bidder, !B\rangle, \langle auctioneer, !A\rangle, postura(!GID, ?O)\rangle$$
$$j_3 = \lambda$$
$$j_4 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, remate(!GID, ?P, ?W)\rangle$$
$$j_5 = \langle inform, \langle auctioneer, !A\rangle, \langle bidder, !B\rangle, sin\_ganador(!GID)\rangle$$

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, although we shall treat messages as propositions, 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.

Agents are able to communicate by sending messages. We assume that a set $\mathcal{IP}$ of *illocutionary particles* is shared by all the agents (e.g., those of KQML [6] or FIPA ACL [7]).

We model an interaction model as a non-deterministic finite-state machine whose transitions are either illocutions, or special transitions such as timeouts or λ-transitions (null transitions):

Definition 1. *An* interaction model *is a tuple* $IM = \langle Q, q^0, F, M, \Sigma, \delta \rangle$ *where:*

- Q *is a finite set of* states
- q^0 *is a distinguished element of* Q *named the* initial state
- F *is a subset of* Q *which elements are called* final states
- M *is a set of* messages
- $\Sigma = \mathcal{I} \cup \mathcal{C} \cup \{\lambda\}$*, where*
  - $\mathcal{I}$ *is the set of all tuples* $\langle \iota, (id : r), (id' : r'), m \rangle$ *with* $\iota \in \mathcal{IP}$*,* $m \in M$*, and* $(id : r), (id' : r')$ *agents with* $id \neq id'$*. Elements of* $\mathcal{I}$ *are called* illocutions*. If* $\sigma = \langle \iota, (id : r), (id' : r'), m \rangle$ *then* $(id : r)$ *and* $(id' : r')$ *are the* sender *and* receiver *of* σ *respectively, and we write* $\mathcal{S}(\sigma) = (id : r)$ *and* $\mathcal{R}(\sigma) = (id' : r')$*. We also write* $\mathcal{IP}(\sigma)$ *instead of* ι
  - $\mathcal{C}$ *is a set of* special transitions
  - $\mathcal{I} \cap \mathcal{C} = \emptyset$
- δ *is the* transition function*, mapping* $Q \times \Sigma$ *to* 2^{Σ}

In order to describe the relationship between local and global interaction models we need to define morphisms between interaction models (or IM-morphisms, for short) in a way that they capture the intuitions described in Section 2, namely that while illocutions are mapped in one direction, interaction states are mapped in the opposite direction:

Definition 2. *Let* IM_1 *and* IM_2 *be interaction models,* $IM_i = \langle Q_i, q_i^0, F_i, M_i, \Sigma_i, \delta_i \rangle$*. An* IM-morphism *from* IM_1 *to* IM_2 *is a contravariant*[1] *pair of functions* $f = \langle g, h \rangle$*,* $g : \Sigma_1 \to \Sigma_2$ *and* $h : Q_2 \to Q_1$*, such that:*

[1] *Contravariant* means going in opposite directions.

1. $h(q_2^0) = q_1^0$
2. $q \in F_2$ *if and only if* $h(q) \in F_1$*, for all* $q \in Q_2$
3. $\mathcal{IP}_1(\sigma) = \mathcal{IP}_2(g(\sigma)), \mathcal{S}_1(\sigma) = \mathcal{S}_2(g(\sigma))$ *and* $\mathcal{R}_1(\sigma) = \mathcal{R}_2(g(\sigma))$*, for all* $\sigma \in \mathcal{I}_1$
4. $\sigma \in \mathcal{C}_1$ *if and only if* $g(\sigma) \in \mathcal{C}_2$*, for all* $\sigma \in \Sigma_1$
5. $q' \in \delta_2(q, g(\sigma))$ *if and only if* $h(q') \in \delta_1(h(q), \sigma)$*, for all* $q, q' \in Q_2$ *and* $\sigma \in \Sigma_1$

In the remainder of this paper, we shall write $f(\sigma)$ and $f(q)$ instead of $g(\sigma)$ and $h(q)$, respectively.

Definition 3. *Given two IM-morphisms* $f_1 = \langle g_1, h_1 \rangle$ *from* IM_1 *to* IM_2 *and* $f_2 = \langle g_2, h_2 \rangle$ *from* IM_2 *to* IM_3*, the composite IM-morphism is* $f_2 \circ f_1 = \langle g_2 \circ g_1, h_1 \circ h_2 \rangle$.

It is easy to prove that $f_2 \circ f_1$ is also an IM-morphism. The combination of interaction models into a global interaction model is naturally given by the coproduct of interaction models, since it makes the disjoint union of illocutions while projecting global states to local states. State transitions in the coproduct are given either by special transitions, or transitions by messages of the same illocution and from the same sender to the same receiver:

Definition 4. *Let* IM_1 *and* IM_2 *be two interaction models. The* coproduct *of* IM_1 *and* IM_2*,* $IM_1 + IM_2 = \langle Q, q^0, F, M, \Sigma, \delta \rangle$*, is defined as follows:*

- $Q = ((Q_1 \setminus (\{q_1^0\} \cup F_1)) \times (Q_2 \setminus (\{q_2^0\} \cup F_2))) \cup \{\langle q_1^0, q_2^0 \rangle\} \cup (F_1 \times F_2)$
- $q^0 = \langle q_1^0, q_2^0 \rangle$
- $F = F_1 \times F_2$
- $M = M_1 \uplus M_2 = \{\langle 1, m \rangle : m \in M_1\} \cup \{\langle 2, m \rangle : m \in M_2\}$
- $\Sigma = \mathcal{I} \cup \mathcal{C}$*, where* $\mathcal{I}$ *is the set of illocutions over* M *as defined above, and* $\mathcal{C}$ *is the disjoint union of* $\mathcal{C}_1$ *and* $\mathcal{C}_2$*, i.e.,* $\mathcal{C} = \mathcal{C}_1 \uplus \mathcal{C}_2 = \{\langle 1, c \rangle : c \in \mathcal{C}_1\} \cup \{\langle 2, c \rangle : c \in \mathcal{C}_2\}$
- $\langle q_1', q_2' \rangle \in \delta(\langle q_1, q_2 \rangle, \sigma)$ *if*
  - $\sigma = \langle \iota, (id : r), (id' : r'), \langle i, m \rangle \rangle$ *and* $q_i' \in \delta_i(q_i, \langle \iota, (id : r), (id' : r'), m \rangle)$
  - $\sigma = \langle i, c \rangle \in \mathcal{C}$ *and* $q_i' \in \delta_i(q_i, c)$

There exist two natural IM-morphisms, $s_i : IM_i \to IM_1 + IM_2$ $(i = 1, 2)$*, called* injections*, defined as follows:*

- *On transitions:*
  - $s_i(\langle \iota, (id : r), (id' : r'), m \rangle) = \langle \iota, (id : r), (id' : r'), \langle i, m \rangle \rangle$
  - $s_i(c) = \langle i, c \rangle$
- *On states:* $s_i(\langle q_1, q_2 \rangle) = q_i$

Lemma 1. $IM_1 + IM_2$ *is an interaction model and* s_1 *and* s_2 *are IM-morphisms.*

Proof. We just prove that:

$$\langle q_1', q_2' \rangle \in \delta(\langle q_1, q_2 \rangle, s_i(\sigma)) \text{ iff } s_i(\langle q_1', q_2' \rangle) \in \delta_i(s_i(\langle q_1, q_2 \rangle), \sigma)$$

for each $i = 1, 2$, $\sigma \in \Sigma_i$ and $\langle q_1, q_2 \rangle, \langle q_1', q_2' \rangle \in Q$. If $\sigma = c \in \mathcal{C}_i$:

$$\begin{aligned} \langle q_1', q_2' \rangle \in \delta(\langle q_1, q_2 \rangle, s_i(c)) & \text{ iff } \langle q_1', q_2' \rangle \in \delta(\langle q_1, q_2 \rangle, \langle i, c \rangle) \\ & \text{ iff } q_i' \in \delta_i(q_i, c) \\ & \text{ iff } s_i(\langle q_1', q_2' \rangle) \in \delta_i(s_i(\langle q_1, q_2 \rangle), c) \end{aligned}$$

The proof is analogous when σ is an illocution. □

The following theorem states that the interaction model defined in Definition 4 is indeed a coproduct because its universal property is satisfied:

Theorem 1. *Given two IM-morphisms f_1 from IM to IM_1 and f_2 from IM to IM_2, there exists exactly one IM-morphism f defined from $IM_1 + IM_2$ to IM that makes the following diagram commute:*

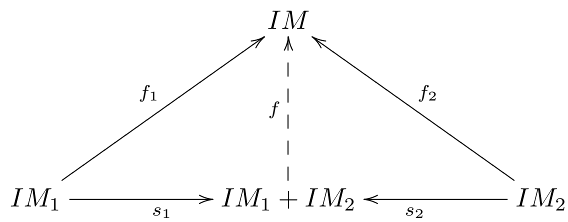

Proof. For each $\sigma \in \Sigma$, there exists $i \in \{1, 2\}$ and $\sigma_i \in \Sigma_i$ such that $s_i(\sigma_i) = \sigma$; we define $f(\sigma) = f_i(\sigma_i)$. On states, $f(q) = \langle f_1(q), f_2(q)\rangle$. It is easy to prove that f is an IM-morphism and makes the above diagram commutes. The uniqueness is guaranteed by definition.

4 Semantic Alignment in Interaction Models

Recall that our approach to tackle semantic heterogeneity in multi-agent communication was stemming from our claim that semantic alignment is often relative to the particular state of the interaction in which agents are engaged in. In this view, the meaning of a message in an uttered illocution is given by the state transitions it brings forward in the scope of an interaction model. Consequently, from the interactional point of view taken in this paper, two messages are semantically equivalent if, and only if, they yield the same state transitions (when occurring in equivalent illocutions, i.e., illocutions with the same illocutionary particle, and the same sender and receiver).

In our semantic heterogeneity example of Section 2 where an English-speaking auctioneer interacts with a Spanish-speaking bidder, the messages "bid" and "postura" turn out to be semantically related, since they make the global interaction change in the same way (see Figure 5). And so are the messages "sold", "withdraw", "remate" and "sin_ganador"; they are semantically equivalent from an interactional point of view because they all make the interaction state change in the same way, at least for the simplified interaction models considered in Section 2.[2]

The most general global interaction for the auctioneer and the bidder is precisely given by the coproduct interaction model of those of the auctioneer and the bidder, as defined in Definition 4. Semantic relationships between messages uttered by the auctioneer and messages uttered by the bidder are hence to be sought in this global interaction: a message m in one agent's interaction model is more general than a message n in

[2] By enriching our model of interactions with additional constraints on the illocutions, commitments brought forward by certain illocutions, etc., we will have more items determining the semantic similarity or dissimilarity of messages.

another agent's interaction model (provided they both occur in equivalent illocutions) if the injection of m into the coproduct of interaction models occurs in illocutions that yield at least the same coproduct interaction state transitions as the illocutions in which the injection of n occurs. Or put differently: Within the scope of an interaction model, let us define the meaning of a message in an illocution to be the set of all state transitions it brings forward. Then, a message m in one agent's interaction model is more general than a message n in another agent's interaction model (again, whenever they occur in equivalent illocutions) if the meaning of the injection of m into the coproduct of interaction models subsumes (contains) the meaning of the injection of n. Following is the formal definition of this semantic relationship:

Definition 5. *Let IM_1 and IM_2 be two interaction models whose set of messages are M_1 and M_2, respectively. Let Q and δ be the set of states and the transition function of the coproduct $IM_1 + IM_2$. Let s_1 and s_2 be the injections of the coproduct.*

A message $m_i \in M_i$ is more general *than a message $m_j \in M_j$ ($i, j \in \{1, 2\}, i \neq j$) for illocutionary particle ι, sender s and receiver r if, and only if, for all $q \in Q$ such that $\delta(q, s_j(\langle \iota, s, r, m_j \rangle)) \neq \emptyset$, we have that $\delta(q, s_j(\langle \iota, s, r, m_j \rangle)) \subseteq \delta(q, s_i(\langle \iota, s, r, m_i \rangle))$. We write $m_j \sqsubseteq m_i$.*

Related semantic relationships such as equivalence, overlap and disjointness can be derived in a straightforward way from the definition of $\sqsubseteq$ above. For its importance, though, we provide a direct definition of semantic equivalence below:

Definition 6. *Let IM_1 and IM_2 be two interaction models whose set of messages are M_1 and M_2 respectively. Let Q and δ be the set of states and the transition function of the coproduct $IM_1 + IM_2$. Let s_1 and s_2 be the injections the above coproduct.*

A message $m_i \in M_i$ is equivalent *to a message $m_j \in M_j$ ($i, j \in \{1, 2\}, i \neq j$) for illocutionary particle ι, sender s and receiver r if, and only if, for all $q \in Q$ we have that $\delta(q, s_i(\langle \iota, s, r, m_i \rangle)) = \delta(q, s_j(\langle \iota, s, r, m_j \rangle))$. We write $m_i \equiv m_j$.*

4.1 Semantic Alignment of an English Auctioneer and a Spanish Bidder

In our sealed-bid auction example of Section 2 modelling the interaction of an English-speaking auctioneer with a Spanish-speaking bidder we have:

- interaction model $IM_a = \langle Q_a, S_0, F_a, M_a, \Sigma_a, \delta_a \rangle$, where $Q_a = \{S_0, S_1, S_2\}$, $F_a = \{S_0\}$, $M_a = \{start\_round, bid, sold, withdrawn\}$, and where Σ_a is the set of all illocutions over M_a with illocutionary particles $\iota \in \{inform, commit\}$, including λ and $timeout$, and δ_a is the transition function given in Figure 3;[3]
- interaction model $IM_b = \langle Q_b, T_0, F_b, M_b, \Sigma_b, \delta_b \rangle$, where$Q_b = \{T_0, T_1, T_2\}$, $F_b = \{T_0\}$, $M_b = \{nueva\_ronda, postura, remate, sin\_ganador\}$, and where Σ_b is the set of all illocutions over Ω_b with illocutionary particles $\iota \in \{inform, commit\}$, including λ, and δ_b is the transition function given in Figure 4 with illocutions those listed in (1) of Section 2.2.[3]

[3] But now messages are just predicate names instead of full first-order terms.

– coproduct $IM_a + IM_b = \langle Q, \langle S_0, T_0 \rangle, F, M, \Sigma, \delta \rangle$, where $Q = \{\langle S_0, T_0 \rangle, \langle S_1, T_1 \rangle, \langle S_1, T_2 \rangle, \langle S_2, T_1 \rangle, \langle S_2, T_2 \rangle\}$, $F = \{\langle S_0, T_0 \rangle\}$, $M = \{start\_round, bid, sold, withdrawn, nueva\_ronda, postura, remate, sin\_ganador\}$, and where Σ is the set of all illocutions over M with illocutionary particles $\iota \in \{inform, commit\}$, including λ and $timeout$, and δ is the transition function given in Figure 6 (we have only drawn those transitions between states that are special or for which the illocution, the sender, and the receiver coincide, and are thus relevant for establishing semantic relationships).

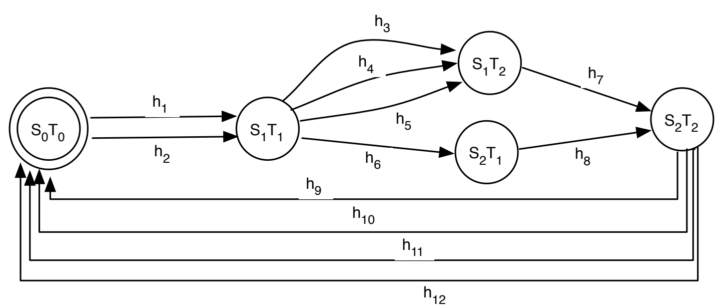

$$h_1 = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, start\_round \rangle$$
$$h_2 = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, nueva\_ronda \rangle$$
$$h_3 = \langle commit, \langle bidder, b \rangle, \langle auctioneer, a \rangle, bid \rangle$$
$$h_4 = \langle commit, \langle bidder, b \rangle, \langle auctioneer, a \rangle, postura \rangle$$
$$h_5 = h_8 = \lambda$$
$$h_6 = h_7 = timeout$$
$$h_9 = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, sold \rangle$$
$$h_{10} = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, withdrawn \rangle$$
$$h_{11} = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, remate \rangle$$
$$h_{12} = \langle inform, \langle auctioneer, a \rangle, \langle bidder, b \rangle, sin\_ganador \rangle$$

Fig. 6. Coproduct interaction model of the English-speaking auctioneer and the Spanish-speaking bidder interaction models

The set of semantic relationships between the auctioneer's message terminology and the bidder's message terminology that arise from this coproduct are as follows:

$$nueva\_ronda \equiv start\_round$$
$$postura \equiv bid$$
$$remate \equiv sin\_ganador \equiv sold \equiv withdrawn$$

4.2 Converging to a Semantic Alignment

Interaction models specify the space of interactions that are allowed, and the coproduct of interaction models 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 coproduct, however, is 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 coproduct interaction model that captures the entire space of actual interactions, but which are 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 compute a semantic alignment based on the assumption that both agents where sharing the same interaction state.

Providing a detailed description of how to gradually approximate the set of semantic relationships that arise during an interaction is subject of our current work in progress. We are certain, however, that such gradual approximation is theoretically feasible, because first, the coproduct of interaction models can be seen as a particular instance of an 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.

In the next section we show in more detail how this link to channel theory is established, assuming that the reader is familiar with this mathematical theory —a description of channel theory lies outside the scope of this paper. The complete understanding of the next section, however, is not necessary to grasp our proposed approach to semantic alignment. It serves, though, to put our work in the broader picture provided by information-flow theory.

5 Information Flow in Interaction-Situated Semantic Alignment

In [1] we studied the problem of aligning differing conceptualizations of two or more agents relative to their respective perception of the environment or domain they are acting in. The syntactic entities that agents use in order to describe their own perceptions of an environment state are aligned precisely because these perceptions come from the same environment state. This environment state acts as a bridge that links agents perceptions, and eventually agents' terminologies describing these perceptions. Channel theory is the framework we chose to formalise these ideas.

In this paper, we have faced a similar situation. Agents' viewpoints refer here to their own interaction models and the coproduct of these interaction models plays the role of the environment. Consequently there exists a natural relationship between the ideas exposed in this paper and channel theory.

To each interaction model one can associate a *classification*, by taking state transitions (pairs of states) as *tokens* and illocutions as *types* and classifying a transition $\langle q, q' \rangle$ to an illocution σ, if and only if, the illocution σ causes the interaction to change from state q to state q'.

Definition 7. *Given an interaction model* $IM = \langle Q, q^0, F, M, \Sigma, \delta \rangle$*, the classification* generated *by* IM*, written* $\mathbf{IM}$ *or* $Cl(IM)$*, is the classification with* $tok(\mathbf{IM}) = Q \times Q$*,* $typ(\mathbf{IM}) = \Sigma$ *and such that* $\langle q, q' \rangle \models_{\mathbf{IM}} \sigma$ *if* $q' \in \delta(q, \sigma)$*.*

Equally, to each IM-morphism one can associate an infomorphism by taking the pair of functions mapping illocutions and pairs of states.

Definition 8. *Given an IM-morphism* $f = \langle g, h \rangle : IM_1 \to IM_2$*, the infomorphism* generated by *f is* $Cl(f) = \langle Cl(f)\hat{}, Cl(f)\check{} \rangle : Cl(IM_1) \to Cl(IM_2)$ *with* $Cl(f)\hat{}(\sigma_1) = g(\sigma_1)$ *and* $Cl(f)\check{}(q_2, q_2') = \langle h(q_2), h(q_2') \rangle$*.*

Consequently, a coproduct of interaction models together with its injection IM-morphisms, determines an information channel linking their associated classifications:

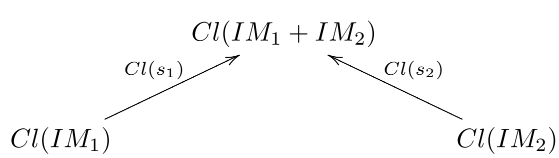

The classification $Cl(IM_1 + IM_2)$ is the core of the channel. An element $\langle \langle q_1, q_2 \rangle, \langle q_1', q_2' \rangle \rangle \in tok(Cl(IM_1) + Cl(IM_2))$ then connects $\langle q_1, q_1' \rangle \in Cl(IM_1)$ and $\langle q_2, q_2' \rangle \in Cl(IM_2)$. Intuitively, a pair of states of the coproduct interaction model links its corresponding projections onto the agents' interaction models. It is by virtue of these links that semantic alignment is established. Actually, the semantic relationships given in

Definitions 5 and Definitions 6 arise from the so call *distributed logic* of the above information channel.

In [1] we formalised semantic alignment as a sequence of information-channel refinements, and in [16] we gave a formal foundation for ontology-alignment interaction models in the context of channel theory. Consequently, it should be theoretically possible to translate the alignment techniques developed in these papers to the semantic alignment approach described here.

6 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 interactional 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 actual interaction 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 disambiguate 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 interaction model morphism and coproduct of interaction models which includes a duality between transitions and interaction states. In this sense the coproduct of interaction models is maximal with respect to transitions, because the global interaction model has to include all possible transitions specified in local interaction models. But the coproduct is minimal with respect to states, because every global interaction state should have a corresponding state in the local interaction models. This work is part of an ongoing effort to characterise semantic alignment as a certain kind of information flow in distributed systems as brought forward by Barwise and Seligman.

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. Definitions 5 and 6 define a semantic relationship not only relative to the interaction in which they are uttered but also relative to the illocution they 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 coproduct 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. C. Fellbaum, editor. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.
6. T. Finin, Y. Labrou, and J. Mayfield. KQML as an agent communication language. In J. M. Bradshaw, editor, *Software Agents*. MIT Press, 1997.
7. Foundation for Intelligent Physical Agents. FIPA 97 specification, part 2: Agent communication language. Technical report, Foundation for Intelligent Physical Agents, 1997.
8. T. Gruber. A translation approach for portable ontology specifications. *Knowledge Engineering*, 5(2):199–220, 1993.
9. Y. Kalfoglou and M. Schorlemmer. Ontology mapping: The state of the art. *The Knowledge Engineering Review*, 18(1):1–31, 2003.
10. D. Lenat. CyC: A large-scale investment in knowledge infrastructure. *Communications of the ACM*, 38(11), 1995.
11. V. López, M. Sabou, and E. Motta. PowerMap: Mapping the 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. P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. In S. Spaccapietra, S. March, and K. Aberer, editors, *Journal on Data Semantics IV*, volume 3730 of *Lecture Notes in Computer Science*, pages 146–171. Springer, 2005.
18. 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.