Applied Ontology X (2007) Y–Z
IOS Press

Let’s *talk* about our ”being”: a linguistic-based ontology framework for coordinating agents

Maria Teresa Pazienza[a], Savino Sguera[b] and Armando Stellato[a]

DISP, University of Rome, Tor Vergata,
Via del Politecnico 1 - 00133 Rome (Italy)
[a] *E-mail: {pazienza, stellato}@info.uniroma2.it*
[b] *E-mail: s.sguera@ieee.org*

Abstract. In open scenarios, agents willing to cooperate must impact the communication barrier between them and their unknown partners. If agents are not relying on any agreement about the meaning they ascribe to the symbols used in the conversation, semantic misalignments will arise on the discourse domain as well, thus making the communication impossible. In wide and heterogeneous scenarios offered by the Web (and its newborn semantic incarnation), traditional and meaning-safe methodologies for communication need to keep abreast of new dynamic communicative interaction modalities. The uncertainties arising from communication between actors which base their behavior on different and heterogeneous knowledge models have to be taken into account, wisely balancing the extent of “reachable knowledge” with the trustworthiness of its information. Main objective of this work is to augment the FIPA Ontology Service Specification with linguistically-aware methodologies for communication, describing a wide-scope framework for multi-agent systems design, semantic integration and coordination.

Keywords: FIPA-compliant multi-agent systems, ontological primitives, semantic integration, linguistic resources

1. Introduction

In the latest years the approach to multi-agent systems design has increasingly moved towards an interdisciplinary *territory*, showing a growing overlap between artificial intelligence (AI) and software engineering (SE) scopes: intelligent agents and agencies are not an AI prerogative, whilst it turns out to be quite crucial to consider some SE *ground* concepts throughout the specification and design process, especially as the system grows in size or complexity. The behavior-based artificial intelligence (BBAI) model, borrowed from robotics and industrial automation scenarios and underlying most popular FIPA[1]-compliant agents development platforms and frameworks, has proven to be effective, addressing agents design in an incremental and iterative fashion, enforcing modularity and code reuse.

Anyway, if on one hand BBAI provides a smart, tiny, modular approach to agent development, on the other hand it is very well known that the design burden moves on the interaction side, as agency’s intelligence has to be coded *by hand*, and inter-agent communication assumes a pivotal role in multi-agent systems design: most of the complexity usually resides in systems communicational layer, and easily grows up along with the number of actors populating the agency and with the particular workflow instance.

Enabling communication between heterogeneous semantic peers (agents, semantic web services etc…) is a fundamental issue for the developing of a Web where machines can effectively *understand* the

[1] Foundation for Intelligent Physical Agents, http://www.fipa.org

meaning ascribed to knowledge representations. This is especially true in relation to the high levels of heterogeneity, evolution, distribution and autonomy of information which already characterize the Web as we know it now.

Moreover, if we look at the Web (r)evolution towards a semantic interoperability together with the proliferation of wireless devices – PDAs and mobile phones – it becomes quite clear what the web is likely to become in the next years: a *dense* and highly dynamic network, populated by a multitude of *nomadic* elaboration nodes able to *understand* information sense, communicating in a more effective and *intelligent* way. We believe it is natural – to some extent – to look at these *nodes* as autonomous entities communicating in a peer-to-peer fashion, *living* in an heterogeneous environment, providing and asking services one another; that is, actually, as agents.

Agents willing to cooperate in wide and open environments must however impact the *communication barrier* between themselves: they will have to make their services available to the community, and recognize those of other *actors* in the society (to whichever extent it is considered, up to the Web as a whole); recognize, interpret and respond properly to communicative acts initiated by other agents and *understand* messages content.

Agents shall have to be *self-describing*, putting other agents in the condition of identifying the ones they need to cooperate with, in order to successfully communicate and exchange information towards a shared goal – to whatever extent it is considered. To put it in Burstein's and Uschold's words (Burstein & Uschold, 2005), "*every agent must, in effect, wear its description 'on its sleeve'*".

The need for a *meaningful* conversational capability of intelligent agents clashes however with the limitations imposed by open environments. How can agents' ability to describe themselves be useful in contexts where even no minimal knowledge is guaranteed to be shared between the communicating partners?

Following past works by (Pazienza & Vindigni, 2002) which explored the possibility of adopting Natural Language as the minimal agreed resource of shareable knowledge, in this essay we revise the FIPA proposal on Agents Ontology Service (FIPA XC00086D, 2001) by introducing new solutions exploiting the linguistic properties of ontologies (the language(s) in which they are expressed, linguistic coverage of pure semantic data, expected pertinence of the adopted terminology etc…) owned by the partners of communication as a basis for achieving semantic consensus.

To provide the core elements for exposing this so called "linguistic expressivity" of ontologies, we defined two metadata ontologies (that is, ontologies containing mostly properties which can be used to describe other ontologies): the *Linguistic Watermark* and the *Ontological Linguistic Watermark* ontologies, whose content can be analyzed by the same partners of communication to coordinate themselves at the best of their possibility, or even by dedicated agents whose existence is centred on supporting communication between third parties.

In Section 3 we will describe the motivations which lead us to propose such methodologies and solutions, while we will add more details on the approach followed, in Section 4.

Section 5 describes the different elements characterizing our framework, reconsidering the ontology development process from a linguistic perspective, which is essential for boosting the expressiveness of ontologies in open scenarios where no semantic consensus is established across different, heterogeneous knowledge resources. In this section the *Linguistic Watermark* ontology and the concept of *Linguistic Enrichment* will be introduced.

Section 6 will extend the linguistically-flavoured ontology design approach, introducing the *Ontological Linguistic Watermark*, a set of metadata describing ontologies in terms of their *linguistic expressivity*.

In Section 7 the FIPA Ontology Service Specification application scenarios will be revised in light of our proposed framework, showing *why* and *how* natural language should represent a crucial element in agents communication.

2. Related works

Expressions such as "Semantic Coordination" and "Meaning Negotiation" have been adopted in literature to address different aspects of the general process of resolving communication barriers between interacting peers. In their research work, Reed, Norman and Jennings (Reed et al. 2002) investigate on possible methodologies for making agents agree on the meaning of the very basic communication primitives which constitute their range of performatives. Their objective is to make agents relying on different communication protocols be able to communicate in open environments Their solution takes the form of a *semantic space* framework, a sort of meta-layer (which, in the end, must be known a-priori by agents willing to communicate), where communication primitives can be described and classified along multiple relevant dimensions. Agents can thus converge, during an initial "semantic boostrap" stage, on the mutually agreeable semantics which throw the basis for their communication.

Opposite to this "assumptions-free" approach, is the research on architectures supporting information exchange in multi-agent systems (Bayardo et al, 1997; Tamma & Bench-Capon, 2001; Li et al., 2005) and in modern web services architectures for the Semantic Web (Paurobally et al., 2006). In these investigations, agents fit on communication protocols which are completely agreed (even considering open environments, as these protocols may be expressed through standards which are supposed to be diffused world-wide), and in some cases they may even agree upon specific implementations of their performatives (e.g. which kind of requests can be submitted inside the agent system). The "semantic coordination" here refers to the agreement which agents (through dedicated figures called *mediators*) are trying to establish on the knowledge content carried by their messages; this is typically represented by ground/schema data explicitly defined in formal ontologies, which may or may not be shared by the communicating partners. The focus of the "coordination" thus moves from the communication protocol to the content which is carried by the protocol.

Our research is closer (though not identical) to these second range of approaches, being applicable to all those cases where the partners of communication share at least a common communication protocol (or, equally, when this issue has already been solved in some way) and share those performatives which make them able to express their knowledge in terms of ontological objects. This scenario is not a too optimistic one, in that communication protocols are converging towards very few and largely accepted standards (FIPA, KQML for agents, SOAP for Web Services), while ontologies have been consecrated by the W3C as a new standard for sharing and reusing data on the web. Still, even considering possible mismatches at the sole content level, communication in open environments is not a trivial issue, since many ontologies are going to exist and be adopted to represent same or overlapping domains and may offer incompatible vocabularies for addressing similar or identical concepts and objects. With this in mind, we tried to move the problem to the very last form of shareable knowledge which may be found along different conceptualizations: Natural Language.

Interesting research has been conducted on the acquisition of shared meaning in linguistic communication, by using techniques and results borrowed from game theory (Jager, 2004; van Rooij, 2004; var Rooij and Sculz, 2004). On the one hand, their results are quite distant from our scenario: software agents utter their knowledge which is defined according to well defined ontologies, have roughly no understanding of the linguistic characterization of what they are communicating (which they merely see as identification labels for their conceptualizations), and are willing to make their knowledge as explicit as they can (namespacing, concept references to widely assessed upper ontologies etc…). On the other side, some of their practical results (see, as a trivial example, the KNIFE and CUTTER example in Benz et al., 2005) could be applied if agents were able to get a richer "linguistic awareness" of the knowledge they need to express.

The importance of a deeper integration between conceptual and linguistic knowledge is well known in literature (Basili et al., 2003; Benjamins et al., 2004; Buitelaar et al., 2006; Oltramari et al, 2006; Pazienza et al., 2003; Pazienza and Stellato, 2006a;2006b;2006c;2006d); moreover, many approaches to ontology matching are today based on exploitation of linguistic resources (Maedche et al., 2002; Bouquet et al., 2005; van Hage et al., 2005) and many systems have been realized to demonstrate the quality of these approaches, but their limitation to practical use in real open scenarios is that agents embedding these systems do not know which ***strategy*** is best to follow in a given situation.

A mediator agent has no explicit knowledge of what he is trying to mediate: are the ontologies written in the same (natural) language? Is at least one of the ontologies expressed in different languages? (i.e. is it a multilingual ontology?), is it possible to combine lexical anchors in different languages to improve the ontology matching process? May the exploitation of a given linguistic resource be of any help in improving quality of retrieval of conceptual matches from the ontologies to be aligned, or is it only a redundancy because the ontologies have already been enriched with the same resource? This kind of investigation represents, in our perspective, a fundamental aspect which could really contribute to the feasibility of a dynamic coordination between heterogeneous semantic peers. Rather than promoting the nth mapping technique which shows surprising results in a confined laboratory experiment, we want to stress the focus here on reducing the ***entropy*** which characterizes every first contact between unknown agents: we will adopt the expression ***Linguistic Coordination*** to address this kind of "awareness" which should be acquainted by agents before they even try to start a meaning negotiation activity.

3. Motivations and Approach

In our premised scenario characterized by a wide, open and uncontrolled environment (such as the web in its entirety), agents involved in a communication will not show, on average, any agreement about the ***meaning*** they ascribe to the symbols used in their speech acts. The communication will thus collapse under the lack of the same foundational elements which are necessary for a proper conveyance of meaning to exist.

Indeed, the model of agent communication in FIPA is based upon the assumption that two agents willing to communicate, share a common ontology for the domain of discourse; this entails that the realization of agent communication needs an agreement over the ***language*** and the ***model*** to refer to. Although an important effort has been conducted by FIPA towards a standardization of agent platforms and agent message transport layer, too many ***dimensions of variation*** still remain to be taken into account. While adherence to the FIPA message protocol grants agents with the ability to establish a communication, distinguish between requests and responses, and bind them to a given subject of conversation, nothing is said about the real content of these messages, thus making impossible for agents belonging to different communities to effectively communicate. We still have to cope with the heterogeneities of agent societies and, mostly, with the diversities in agent frameworks design which are often developed for ad-hoc environments and applications. All these aspects represent a threat against ***large-scale*** interaction among different multi-agent systems.

If peers, even belonging to different agencies, were able to autonomously discover each other's services and negotiate the issues characterizing the particular communicative act instance, it would be much easier to design open systems, having a solid base to address the ***impedance mismatch*** among agents' different knowledge representations from a shared starting point, and making a further abstraction layer available to designers. Thus for the communication to be fruitfully carried out, as stated in the introduction, we must allow agents to have a certain degree of shared ***meta-knowledge*** about communication itself.

The FIPA Ontology Services Specification (FOSS, from now on) described in document XC00086D, (FIPA XC00086D, 2001), deals with technologies enabling agents to manage explicit ontologies, and specifies an ontology service for a community of agents featuring a dedicated Ontology Agent whose role in the community is to provide (at least) a subset of the following tasks:

– discovery of public ontologies in order to access them,
– maintain a set of public ontologies,
– translate expressions between different ontologies and/or different content languages,
– respond to query for relationships between terms or ontologies, and
– facilitate the identification of a shared ontology for communication between two agents.

The above mentioned FIPA specification (which is still marked as "Experimental") presents many problems which may occur in wide and heterogeneous scenarios of communication, where use of different (natural) languages, or simply ***different uses*** of the (same) language, blurred by potential sources of ambiguities like synonymy and polysemy, may oppose a solid barrier to a proper reconciliation of conveyed meanings. This is a fundamental aspect which must be necessarily taken into proper consideration in real wide and open contexts like, one for all, those offered by the Semantic Web, which is "Semantic", in that it will be based on shared vocabularies for communication, but indeed still a "Web" of heterogeneous sources of information. The Ontology Service Specification though, remains a bit underspecified on identifying proper solutions for semantic coordination. As a first consideration, it is perfectly reasonable that agents should at least share a minimal meta-vocabulary for being able to ***talk*** about their inner knowledge and, in this sense, the Open Knowledge Base Connectivity, OKBC (Chaudhri et al., 1998) is a possible choice: OKBC shares a lot in common with RDF(S)/OWL (today standards for representing ontologies) since both these languages expose an object-oriented perspective over knowledge modeling, characterized by similar concepts like classes, properties (***slots*** in OKBC) and instances (***individuals*** in OWL). At the same time, OKBC primitives may seem inappropriate to cover the inferential layer and the other features characterizing OWL and RDFS, so the FOSS would need some rethinking about this aspect: maybe that, in presence of a well-acclaimed standard, it is no more convenient to adopt a meta-ontological layer, like OKBC, which sacrifices completeness of information in exchange for a wider compatibility towards other unused formalisms.

The second consideration, which opens up the needs pursued in this research work, is related to the impossibility, for agents, to have an effective (intended as: "with productive outcomes") communication when their respective knowledge is expressed through different vocabularies. In those cases, it is no more sufficient for agents to be able to exchange ontology terms through an agreed meta-ontology: the exchanged terms would appear, in such a babel of different conceptualizations, as meaningless strings thrown from agent to agent. Agents should thus rely on some basic form of knowledge which possesses the following characteristics:

1. it is some kind of meta-knowledge, not directly related to any aspect of the specialized information owned by agents
2. it is expected to be possessed by any agent in the community
3. it is not an ***interlingua*** limited to the capacity of declaring and exchanging the knowledge owned by communicating agents, it should instead be exploitable to understand similarities and differences arising from agents' different vocabularies, and to reconcile their content.

Natural language seems to us the least common denominator between different, heterogeneous, conceptualizations (Pazienza & Vindigni, 2002), embodying all of the three above characteristics. It is independent (point 1) from the specific knowledge owned by agents, but is necessary to represent it, in every circumstance. It is the sole form of shared knowledge (point 2) which is inherently adopted when expressing (through label descriptors, concept documentation etc...) the content of knowledge resources, just because it is the natural form adopted by humans to express their knowledge. Finally (point 3), because of its "popularity" and its firm use by humans, it is the only source of information which can be expected to retain some "invariants" even in completely different conceptualizations, developed in complete independence from each other: this is a key feature towards the process of reconciliation of incompatible resources. Under these premises, ontology content should always be accompanied by rich and evocative textual information, possibly extracted from well-assessed and qualitatively appreciated lexical resources (vocabularies, terminologies, etc..), to maximally qualify its scope and intentions.

Obviously, the sole "awareness" about the importance of natural language does not suffice: the concept of natural language comprises a lot of different, incompatible, idioms (Benjamins et al., 2004), while ambiguity pitfalls due to linguistic phenomena like polysemy and synonymy need to be taken into account when relying on this natural form of knowledge. Agents thus need a dedicated, shared, meta-ontological layer for expressing the way they are able to "speak" about themselves, and the ability to inspect this kind of knowledge, to identify the context of every communication.

The main objective of this essay is thus to revise the FIPA Ontology Service Specification, rediscussing the scenarios presented in that document under the rationale of a wider awareness of the linguistic properties which characterize the agent ontologies, and to describe a wide-scope framework supporting multi-agent systems design, semantic integration and coordination, hence assuming natural language as ***the*** unifying element between heterogeneous semantic peers.

We hereby propose an ontology framework supporting coordination between heterogeneous semantic peers (being them Agents, Semantic Web Services, or whatever entity may be exposed to some form of "meaning negotiation") acting in an open domain. This framework offers a shared meta-data vocabulary for ontologies, which can be used to synthesize the "linguistic expressivity" of the knowledge exposed by communicating agents. This form of "linguistic metadata", bringing information about the ability of agents to "talk" about their own knowledge, could thus be exploited to make mediator agents able to set up, under the best possible conditions, a semantic coordination activity.

Our proposed framework is characterized by two ontologies: the first one, the *Linguistic Watermark* Ontology (LW), provides a common ground for describing linguistic resources (from now on, LRs) and their characteristics (represented language, availability of syntactic/semantic relations, structure etc…). The second ontology, the *Ontological Linguistic Watermark* Ontology (OLW) is the one in charge of providing the proper meta-data for describing the linguistic expressivity of ontologies, by addressing general linguistic properties of the described ontologies (availability of *rdfs:label*(s) for different languages, coverage of ontology concepts for each language) and providing precise metrics for describing the adoption of a given linguistic resource to linguistically enrich the ontology content.

4. Boosting ontology expressiveness: linguistically motivated ontology design

Following our premise, it is clear how ontologies should not "stand on the shoulders" of their namespaces and expect to become the ultimate knowledge resource about the domain they describe, being referenced by any existing knowledge based application. The Web is not a centralized and easy to manage knowledge repository, in which each world object is associated to a single entity within a universally accepted namespace. As evidenced in (de Bruijn and Polleres, 2004), "in the context of an open environment such as the Web it is very unlikely that there will be very few ontologies shared by many parties; we will have to deal with many different heterogeneous ontologies with overlapping domains". In a scenario like this, it is thus important for these ontologies to make their content easy understandable and identifiable by machines as well as by humans.

The same best practices which apply to software documentation ("do not use strange acronyms or personally conceived nicks for variables and functions, but *evocative* names which can universally lead to a

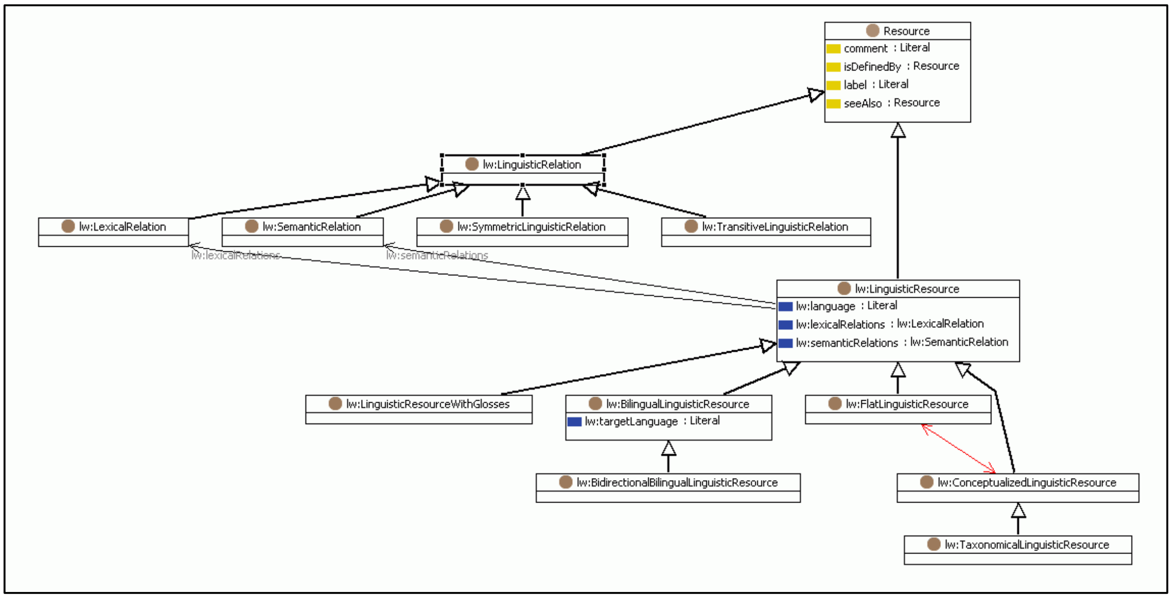

Figure 1: some aspects of the Linguistic Watermark Ontology

deep understanding of their role in their context; add *rich textual descriptions* commenting each piece of software up to its elementary elements, etc…) could, or better, *should* be applied in documenting knowledge data as well. This is even more important than in code programming, since the documented data is conceived to be shared on a much wider scale (World Wide!) and thought to be as self-explanatory as possible, since it could not be accompanied by any further external documentation and should be reusable for really different needs. A *linguistically motivated approach* to ontology design (Pazienza and Stellato, 2006a;2006b) could thus contribute to the degree of "shareability" of ontological resources upon open scenarios characterized by several layers of heterogeneities such as: different idioms (multilinguism), different application domains (involving different jargons, terminologies etc…) as well as to the exchange of ground knowledge data defined upon different logical spaces.

4.1. Linguistic resources: the Linguistic Watermark Ontology

Reuse of available linguistic resources, lexical resources in particular, such as, among the others, terminologies, dictionaries, thesauri and wordnets, is a well expected step in the development of these linguistically enriched ontologies. In (Pazienza and Stellato, 2006a;2006b;2006c) a framework for performing enrichment of ontologies with information borrowed from several lexical resources has been introduced.

Unfortunately, while ontologies have undergone a process of standardization which culminated, in 2004, with the promotion of OWL (Dean, M. & Schreiber, G., 2004), as the official ontology language for the Semantic Web, linguistic resources still maintain heterogeneous formats and follow different models, which make tricky the proper identification of their role in the linguistic enrichment process.

To cope with all of these heterogeneities, we introduced (see Pazienza&Stellato2005a; Pazienza&Stellato2005b for preliminary work on this subject) the Linguistic Watermark Ontology (Figure 1): a classification of Linguistic Resources (at the current stage of specification, it is actually limited to *Lexical Resources*), describing the different characteristics which distinguish a particular resource inside our framework.

Core element of the ontology is the concept of LinguisticResource, a general descriptor for every kind of LR. A linguistic resource can then be defined according to a more specific category. Here we provide a few details for the most relevant ones:

- a Linguistic Resource can be classified as a FlatLinguisticResource, meaning that it is mainly arranged around simple entries of the form:

  IndexWord: WORDLIST

  where IndexWord is a word (or complex linguistic expression) belonging to the "domain" of the resource, and WORDLIST is a flat collection of words/linguistic expressions "pertaining" to the given IndexWord. The semantics (i.e., what "pertaining" means) of this class of entries depends strictly on the nature of the resource; for instance, in bilingual dictionaries (particular instances of class BilingualLinguisticResource) WORDLIST will contain translations of IndexWord, while in a dictionary of synonyms, WORDLIST will expose synonymic expressions for IndexWord, expressed in its same idiom. In this class of resources, WORDLIST is an heterogeneous container for words pertaining to IndexWord and no further clusterization is defined inside it (e.g. according to their different senses).
- Opposed to Flat Linguistic Resources, there are those resources which can be categorized upon the class ConceptualLinguisticResource. In this class there fall all the LRs where translations and/or synonymical expressions of words can be clustered according to a given kind of semantic aspect (typically, but not limited to, words' senses). A notable instance of this class is WordNet (Miller et al., 1990; Fellbaum, 1998), a lexical database which is centered around semantic anchors called synsets (*syn*onymy *sets*), i.e. identifiers for collection of linguistic expressions considered synonymical upon the leibnizian notion of syntactyc substitution. Since these resources are, by definition, opposed to Flat Resources, their related class has been described as owl:disjointWith from the previous one.
- A TaxonomicalLinguisticResource is a Conceptualized LR whose semantic aggregators are arranged in a hierarchy of broader/narrower meanings.
- A LinguisticResourceWithGlosses is a Conceptualized LR whose semantic aggregators are accompanied by a rich natural language description (a *gloss*) of their meaning.
- A BilingualLinguisticResource is a LR providing translations for expressions from one language to another. The specific structure of the resource depends on other facets expressed through its watermark (for example, if it is Conceptualized or Flat).
- A BidirectionalBilingualLinguisticResource is a kind of bilingual LR offering navigation of translations in both ways, from one language to the other one and vice versa.

Apart from the cases where it is explicitly specified, these classes are not disjoint from each other. As an example, the original English WordNet can be both categorized as a TaxonomicalLinguisticResource (being its synsets arranged in a hierarchical graph) as well as a LinguisticResourceWithGlosses (and thus as a ConceptualizedLinguisticResource, which is a superclass for both these categories).

A minimal number of properties provides further descriptions for characterizing linguistic resources represented according to the Linguistic Watermark Ontology. Each resource is associated to a given natural language through the language property. Bilingual Linguistic Resources expose one more property – targetLanguage – for identifying the language in which translations are expressed. Two more properties, lexicalRelations and semanticRelations, inform about which kind of linguistic relations are expressed inside a given resource. Linguistic Relations are described inside the same ontology and can be further specified as being SymmetricLinguisticRelation and/or TransitiveLinguisticRelation; while LexicalRelation is disjointWith SemanticRelation, instances of both of them can indifferently be categorized as symmetric and/or transitive. While new individuals representing linguistic relations could be added, upon necessity, in the watermark of each linguistic resource, it should be important to define an exhaustive list of linguistic relations which could be referred by any instance of the LinguisticResource class. This would provide a shared vocabulary for unifying, harmonizing or simply coordinating information extracted from different LRs. Obviously, the ontology depicted in fig. 1, even if completed with the instance data related to the specific linguistic relations, and other schema-level details which have not been reported for clarity of the picture, is still not intended to provide a complete model for representing lexical resources. Its scope is quite distant from notable past efforts like the Text Encoding Initiative (www.tei-c.org) and the LREEAGLES (Expert Advisory Group on Linguistic Engineering Standards) project (Calzolari et al., 1996) or, to cite a more recent one, like LMF (Lexical Markup Language, Francopoulo et al., 2006). The Linguistic Watermark is instead a necessary high-level description of the characteristics of lexical resources which come into play in our framework with three well defined objectives:

1. Being able to identify, among available linguistic resources, those which are best suited for enriching a given ontology (depending on needed natural language and desired linguistic information)
2. Being able to select those resources which can be useful to assist a given semantic coordination
3. Being able to trace back the quality and the nature of the enrichments which have been carried on a given ontology

Should some of the efforts described above reach some level of consensus or, even better, become accepted by some standardization body, we will develop a mapping between our ontology and the new standard or, even better, if represented in some RDF based model, revise our framework to adopt that standard as a model.

5. An ontology for supporting Linguistic Coordination: the Ontological Linguistic Watermark

In our perspective, the Ontology Agents foreseen by the FIPA Ontology Service Specification should be able to characterize the linguistic characteristics of the knowledge resources they host to improve retrieval of knowledge from their content as well as to identify similar resources hosted by other Agents.

It is thus important, should an ontology had been already linguistically enriched, to know in advance to which extent and through which modalities this enrichment process has been conducted; if agents knew the history of the process which contributed to the development of the knowledge resources they adopt, they would be more capable of sharing its contents with other agents. In our framework, this information is represented (Figure 2) by the *Ontological Linguistic Watermark* (OLW, from now on), a collection of meta-data descriptors carrying information about the (natural) language(s) adopted in describing ontology contents, and (in case) about the linguistic resources which have been exploited to do that. These metadata play a central role in the Linguistic Coordination phase, as agents are able to inspect each other’s watermarks at the start of their first contact and plan the best strategy for establishing a communication: this may also imply the exploitation of external resources as well as invoking the help of other agents specialized in supporting communication and meaning negotiation. The key element of the Ontological Linguistic Watermark ontology is the rdf property: linguisticEnrichment, which has owl:Ontology as domain (it is thus a metadata descriptor usable for any ontology). This property points to a class LinguisticEnrichment, describing a generic linguistic enrichment process which may have been

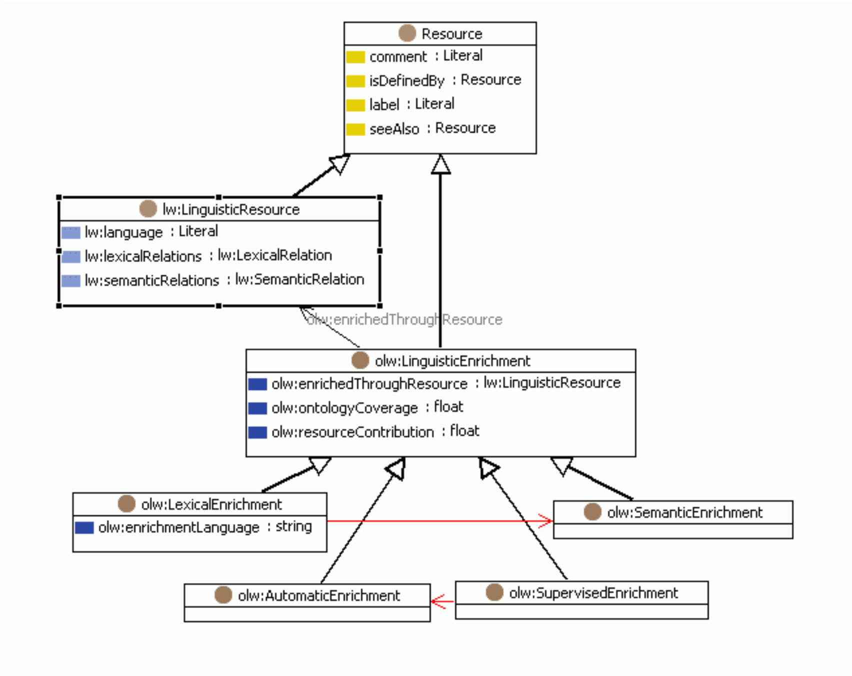

Figure 2: an overview of the Ontological Linguistic Watermark

applied to the characterized ontology. The linguisticEnrichment property has no maxcard, in that a single ontology may have been subject to different enrichment processes (for example, to have a wide coverage of a given language or to produce a multilingual ontology). The LinguisticEnrichment class may be further specialized according to two different perspectives. First, it may be completely partitioned into two disjoint classes:

- LexicalEnrichment: in which the information extracted from the LR consists in terms and natural language descriptions, successively added as labels and comments to the objects of the enriched ontology
- SemanticEnrichment: in which the enrichment has been carried out by extracting the semantic descriptors of a ConceptualizedLinguisticResource. This practice is rather frequent (as proposed, for example, in Jarrar, 2006), as it may provide more precise semantic anchors in scenarios where a given LR is well known. If one of the communicating agents is however not aware of the LR adopted to semantically enrich the ontology of the other one, this kind of enrichment may offer little help in the communication, since the semantic identifiers extracted from the LR appear as meaningless strings to the unaware agent.

Second, the Linguistic Enrichment may be carried on either manually or automatically, through dedicated tools, like those described in (Pazienza & Stellato, 2006a); then the need to represent two other disjoint classes have been added: AutomaticEnrichment and SupervisedEnrichment. For the latter class, it is possible to specify to which extent the automatic process has been supervised by a human.

If the enrichment is a lexical enrichment, a dedicated property (enrichmentLanguage) expresses the natural language for which the particular enrichment has been carried on.

The enrichedThroughResource property provides information (through reference to instances of lw:LinguisticResource class of the *LinguisticWatermark* Ontology) on the LRs which are exploited for linguistic enrichment (instances of the LinguisticEnrichment class). A single instance of a Linguistic Enrichment is thus intended to have been conducted on a single LR, so this property is *owl:functional* (maxcard = 1).

Two other properties (ontologyCoverage and resourceContribution) add quantitative information for knowing how much a linguistic enrichment has covered the knowledge data expressed by the ontology and how much this contributed with respect to other resources. The ontologyCoverage property expresses (as reported in its *rdfs:label*) the "percentage of ontology concepts which are represented by at least a term from the linguistic resource" (the LR adopted for the described linguistic enrichment). The resourceContribution expresses the "Percentage of terms common to both ontology and linguistic resource wrt total number of terms (for the same language) in the ontology. This information is useful for knowing how much the specific resource contributed to the linguistic enrichment of the ontology" (with respect to other resources for the same – natural – language which may have been used to enrich the ontology).

Last, there are two properties in charge of providing conclusive information on the range of (natural) languages available for accessing the knowledge content of the ontology. The languagesForNames property tells the language(s) used to name the resources of the described ontology. With "name" it is intended the actual true ontology name of the ontology resource. Note that this property only informs on the languages used to describe resources in its own namespace. If other ontologies imported by the one described by this property use different languages, this should be reported in their own languagesForNames description. An agent willing to communicate the languages in which its whole knowledge is expressed, should explicitly report the set of languages of all its imported ontologies; in this way it is however also able to reply to specific queries on the languages adopted for a subset of its imported ontologies. The languagesForLabels property provides much the same information, with the sole difference that it informs on the use of the annotation property *rdfs:label* to add natural language descriptors for the ontological resources.

The OLW thus provides a complete overview of the linguistic expressivity of ontologies, which can be exposed by the agents willing to cooperate in open scenarios. The metadata provided by the OLW can be analyzed by the same agents involved in a communication or by specific agents dedicated to supporting communication between heterogeneous semantic peers, much in the spirit of (Pazienza & Vindigni, 2002).

6. Applications and use cases

We report in this section some sample application scenarios and use cases for our framework, revisiting some of those exposed in the FIPA Ontology Service Specification from a linguistic perspective, with the aim of practically showing how our proposed ontologies may support several communication needs in open domains.

All of the examples in this paper assume agents are communicating through the FIPA Agent Communication Language (though, without loss of generality, any agreed communication protocol would fit as well). The scope of investigation here is in fact centred on finding agreement on the content of the messages which are exchanged by agents, while the semantics of the performatives they use are assumed to be shared in their communicative context.

6.1. Scenario 1: Definition of terms querying

This scenario (Figure 3) allows the usage of an Ontology Agent OA to access definition of terms when large linguistic ontologies need to be used (FIPA XC00086D, 2001, section 2.3.1). The standard

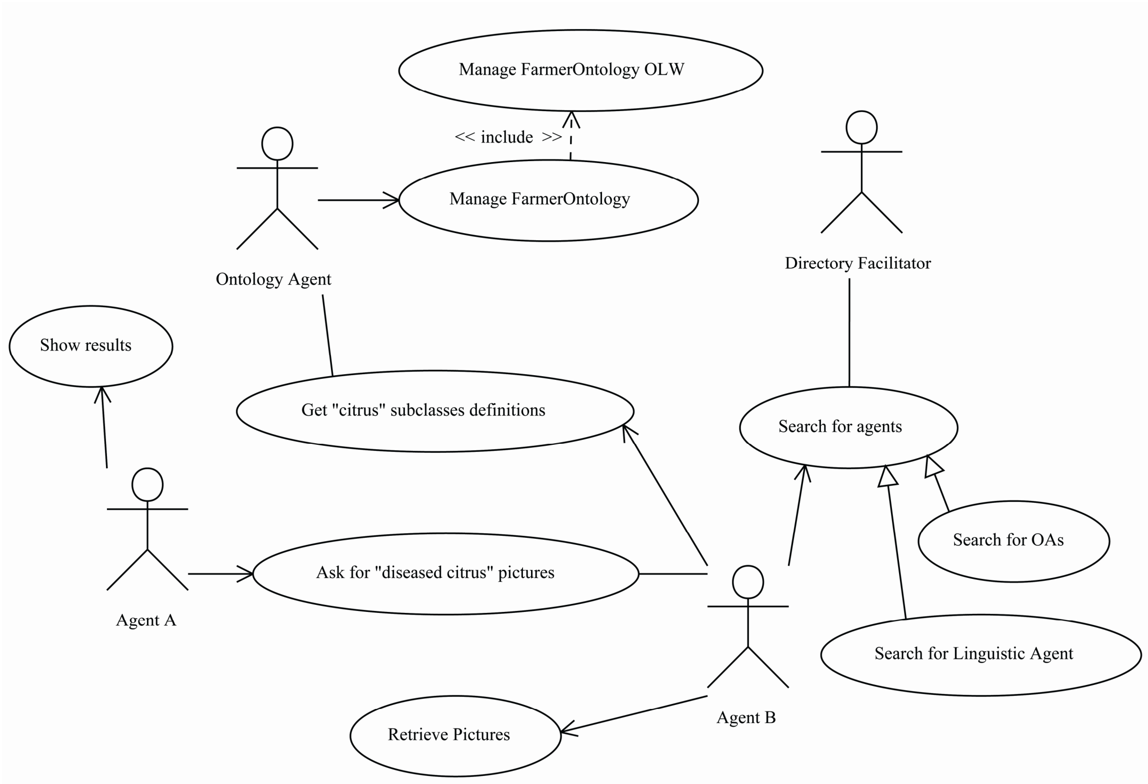

Figure 3: An UML Use Case sketch for Scenario 1

scenario proposed by the FOSS includes an Agent B with caption-based picture indexing capabilities, retrieving pictures matching other agents' queries. The scenario evolves as follows:

– An user interface Agent A is retrieving pictures of *diseased citrus* for its user, a farmer needing to discover which disease his trees are suffering from. Agent A requests Agent B to send pictures of *diseased citrus*, a concept defined in a given domain ontology, for example, the farmer ontology.
– Agent B attempts to retrieve pictures but no results matching the *citrus* caption are available; thus it searches for the appropriate OA, which is managing the farmer ontology, and contacts it in order to get *citrus* subclasses' definitions.
– The OA informs Agent B that *oranges* and *lemons* are more specific concepts for class *citrus*.
– Agent B now can search pictures of *diseased orange* and *diseased lemon*, and send them to the Agent A, which will in turn show the results to its user.

The one above is the original scenario as presented in the FOSS; we can make the situation a bit more complex and suppose *citrus* has no subclasses in the farmer ontology, or that they are just too few (lemons and oranges are not the totality of citrus fruits). Furthermore, we could now add some details which would reasonably be available in our framework:

– Agent B knows several agricultural image repositories whose images are written in *English*, and another repository with captions written in *Italian*.
– Both repositories are arranged upon folders which are indexed too by the agent, so it would be best to get an exact match on the name of the folder for the citrus concept, rather then exploring

subconcepts and accept only partial results.

– The farmer ontology is enriched through the Ontological Linguistic Watermark Ontology, and the OA agent can provide information about these ontological metadata. In the specific case, OA is able to tell that the farmer ontology’s concepts are expressed in the English language, and that it has been subject of two linguistic enrichments: one with English labels and comments extracted from WordNet 2.1 and another one, manually, for the Italian language.

This more complex scenario will thus evolve in this way:

– Agent B queries the OA for English synonyms (expressed through *rdfs:labels* with the *xml:lang* attribute set to *en*) of the concept *citrus*. The OA replies with no label, so, Agent B (trusting WordNet as a complete lexical resource), infers there are no real synonyms of the word citrus. Its research on English repositories ends (at least, for what can be inferred from the farmer ontology).
– Agent B knows from Agent A (which in turn has reported info from the OLW of the farmer ontology) that the ontology is written in English and that there are translations for the Italian language, so it queries OA for *rdfs:label*(s) (with language set to *it*) of the term *citrus*, getting back the words *citro* and *agrume*, which help successfully in the retrieval of figures from the Italian repository.

Though this scenario (even in its original form, taken from FOSS) recalls use of ontologies more as reference dictionaries for terms to be adopted in text-based searches (in the original scenario from FOSS, Agent B was getting subconcepts of *citrus* to be used in its search, while in this case it asks for synonyms and translations), the advantage of using these formalized conceptualizations again emerges. Had *citrus* been an ambiguous word, Agent B was however accessing, through OA, to the meaning of citrus it was interested in, and through its labels, only to the synonyms and the translations which pertain to the proper sense of the word.

There is an apparently fruitless information which has been obtained from the OLW in the first part of Agent B’s search: the *citrus* had no other labels in the English language. This is indeed a negative answer, though, by knowing from the OLW that an enrichment for the English language had been done, this information could actively help in getting the dimension of the completeness of the available results. Agent B knows that, if “citrus” is an English word (and, at the same time, if it exists in WordNet), the absence of English labels for the specific ontology, coincides with the absence of English synonyms for the word (at least for the meaning of the word it is interested in), thus it explicitly *knows* that it cannot do any better than move to other degraded information (as in the original scenario, by asking subconcepts of *citrus*). Another information which must not be underestimated is the explicit awareness about the language adopted for the ids of the ontology concepts. In the second part of the search, if the ontology had no labels for the Italian language, Agent B could have made use of a bilingual dictionary for retrieving Italian translation of the term citrus (e.g. through a request to its Directory Facilitator (DF), invoking the help of a *linguistic agent* offering bilingual translation services). To query the proper resource (invoke the proper linguistic agent/service), it is thus important to know in advance which is the source language.

6.2. Scenario 2: Ontological equivalence testing

In this use case (Figure 4) an agent has to check the *logical* equivalence of two given ontologies.

– An American ontology designer declares a car-product ontology associated to the ontology agent OA2. The ontology is referred by the OA2 as the car-product ontology.

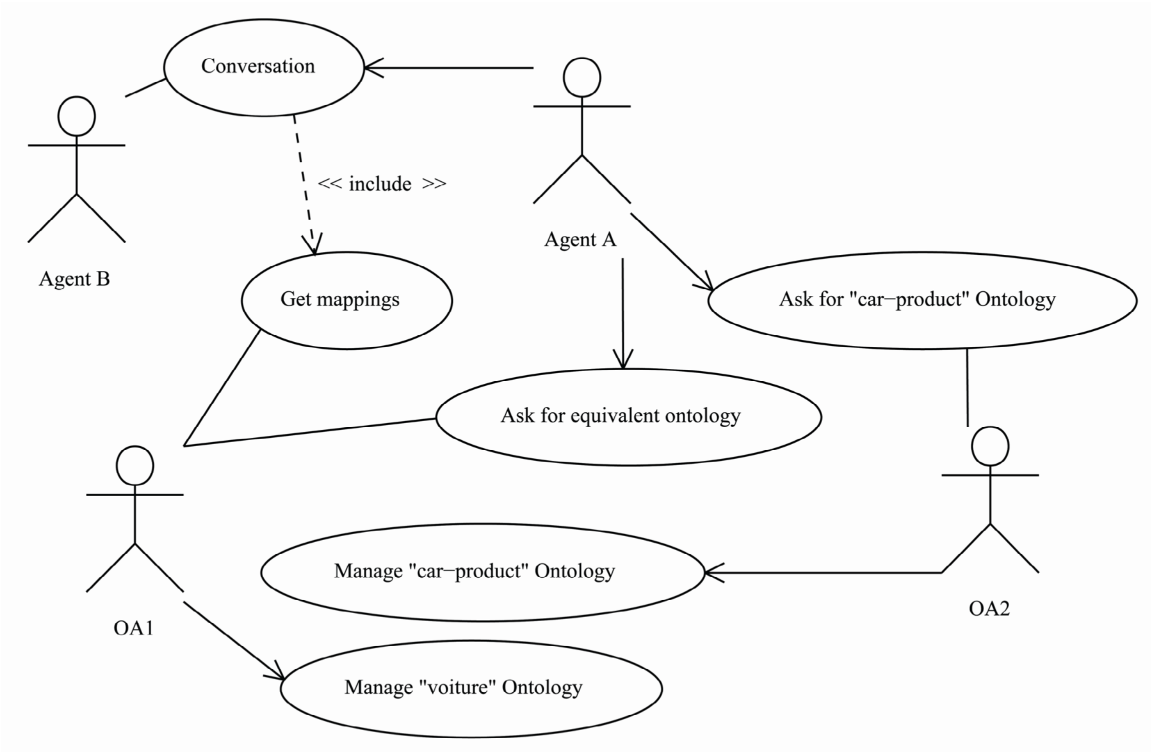

Figure 4: An UML Use Case sketch for Scenario 2

- The designer declares then a complete French translation of his car-product ontology as the voiture ontology, and associates it to the ontology agent OA1. The two ontologies are declared equivalent to OA1, which is providing *exact* mappings between the two ontologies.
- An Agent A (in the US) requests OA2 to provide an ontology in the *cars* domain, getting back the ontology name car-product.
- Agent A needs now to communicate with Agent B (in France) about *cars* using the car-product ontology. Note that Agent A does not know this ontology.
- Ontology agent OA1 is queried by Agent A which needs a car-product-equivalent ontology, if any. Then OA1 returns voiture to Agent A.
- Agent A informs Agent B that these two ontologies are equivalent and that OA1 can be used as a formal knowledge mediator, since explicit mappings between voiture and car-product exist.
- Agent A and Agent B may now start their conversation about *cars*.

Again, the one above is the original scenario from the FOSS; in our framework, probably the ontology designer would have not published two different ontologies by translating the ids from one original resource: agents expecting to rely on the presence of labels in different languages would expect as well to be able to access ontologies whose ids are written upon languages which are not familiar to them, but featuring proper translations through labels for their native language (or the native language of their human user).

A single ontology agent would thus be necessary, providing multilingual access to the same resource: this simplifies things a lot, because the awareness about the equivalence of the expressions in different

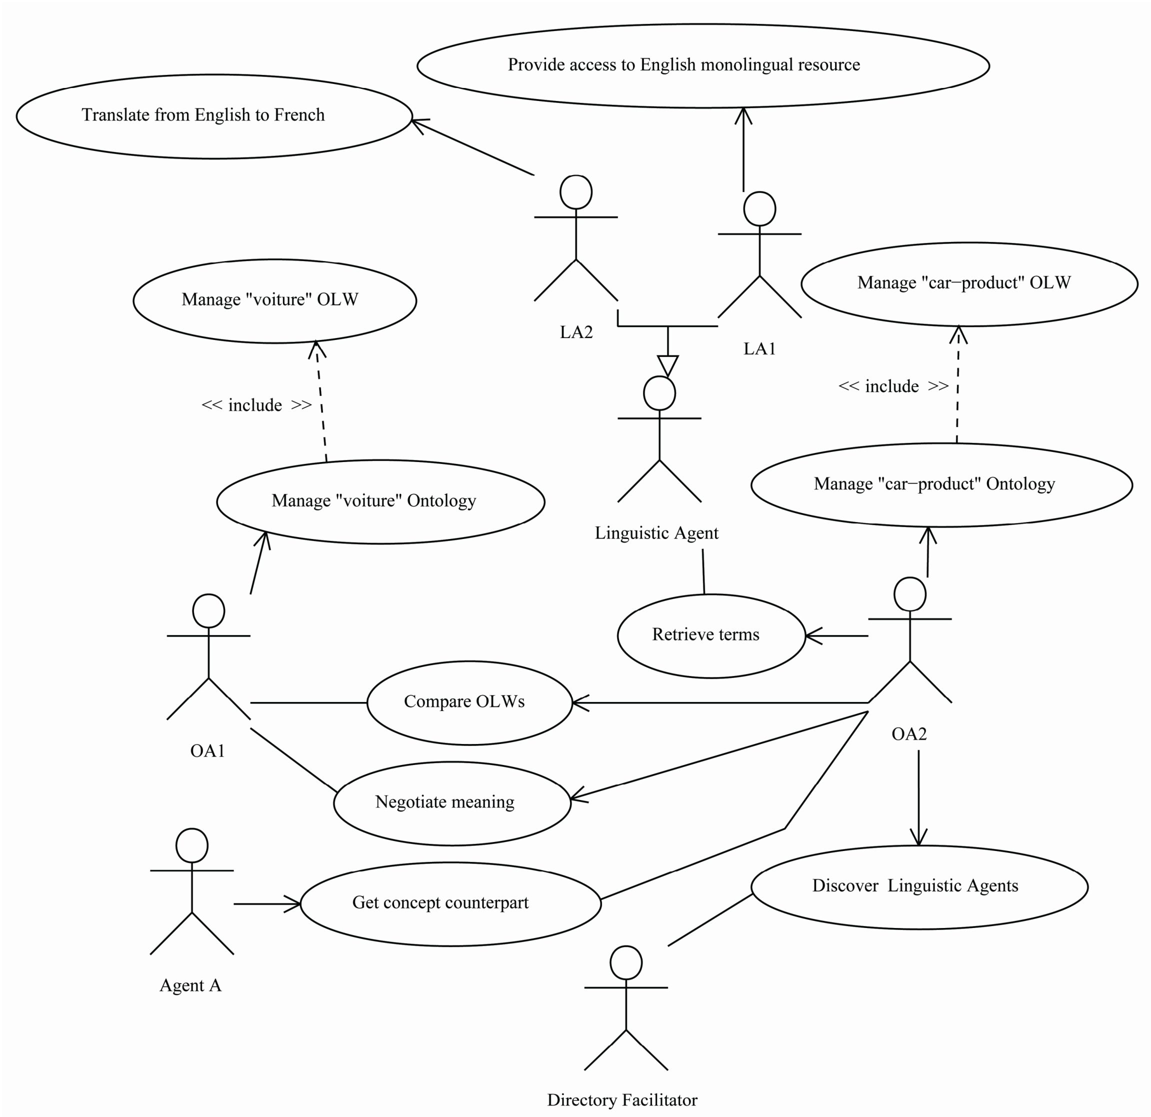

Figure 5: UML Use Case for Scenario 3 with Linguistic Agents supporting Semantic Coordination

languages is kept in the ontology itself and is of public domain through the OLW, while in the original scenario it was confined to Agent A. Also, the OLW helps in formalizing the availability (and the extent) of translations for a given ontology in different languages, and the source language(s) used for representing its ids.

6.3. Scenario 3: Term translation

As a further possible alternative to the previous scenario, we must consider the case in which an explicit equivalence between two given ontologies does not exist. In open contexts it is likely to happen that agents with different, independent conceptualizations of the same application domain are willing to communicate. Section 2.3.5 of FOSS presents a similar scenario, but, in that case, agents rely on an explicit mapping provided by an ontology agent.

What can happen in reality is that the two ontologies have been discovered in complete independence (or, equivalently, that two agents based on two different ontologies need to communicate) and that there is no available mapping between them. In those cases, automatic services for mapping ontologies *on-the-fly*, should be provided by agents focused on this task (we could see them as specialized roles of the already known Ontology Agent). Several approaches to ontology mapping have been extensively discussed in literature (see Rahm and Bernstein, 2001; de Bruijn et al., 2003; Kalfoglou and Schorlemmer, 2003; Shvaiko and Euzenat, 2005; to get an overall view of ontology matching approaches), many of which base their algorithms on the exploitation of existing linguistic resources: as already stated in the introduction, we do not enter into the details of the adopted techniques, while we want to stress that all of these approaches lack a real plausible framework in which they could be applied.

Many ontology mapping tools available from the ontology research communities need to be tuned and configured according to the specific situation in which they are put on work (see Sure et al., 2004, for some application scenarios which have been considered inside the EON workshop on evaluation of Ontology Mapping Tools), though they claim to be usable in scenarios where mappings are to be computed on demand, and expected to be initialized by their hosting agents, without human intervention. It is easy to verify that, through our framework, a complete description of the characteristics of the two ontologies is made available, thus helping communicating agents in automatically setting the ideal environment for starting a semantic coordination activity.

Let us now provide a few details of such a scenario (Figure 5): suppose that the car-product ontology and the voiture ontology are not formally equivalent, but are products of two distinct conceptualization processes (being the former developed in the US, and the latter in France) over the same application domain (i.e. *cars* domain).

Obviously, differences in the conceptualizations involve not only a problem of translation, but also the possible presence of strong heterogeneities both on a macroscopic perspective (the whole structure of the ontology) as well as on a detailed investigation (use of different ontology patterns and/or even different datatypes). Finally, the two ontologies may as well present different linguistic characteristics.

The scenario is now much less formally defined; we have no explicit correspondences in voiture for concepts defined in the car-product ontology, and vice-versa: languages are different, concept hierarchies are different, and the possible linguistic resources used for lexical and semantic annotation are different too. In the following, we provide a ground example for such a scenario:

- OA1 is managing the voiture ontology, and its Ontological Linguistic Watermark too. OA2 is doing the same for the car-product ontology.
- OA2 declares that it has semantic coordination skills, so it proposes to handle the mediation. OA1 accepts.
- Agent OA2 compares the ontologies' watermarks to know which information sources could be useful to reduce the mismatch between them. The voiture ontology reveals a single manual linguistic enrichment, providing *rdfs:label*s and *rdfs:comment*s for the French language. The car-product ontology presents no linguistic enrichment.
- OA2 decides to invoke the help of two linguistic agents: it needs a good monolingual linguistic resource for the English language, to extend the vocabulary of English words (by finding synonyms of the terms in the car-product ontology) which can then be used in conjunction with a bilingual linguistic resource (from English to French) necessary for finding matches across the two ontologies.
- OA2 contacts the Directory Facilitator to check availability of two linguistic agents: the first of-

fering access to a LinguisticResource with language=en, while the second providing query facilities for a BilingualLinguisticResource with language=en and targetLanguage=fr

- The DF replies with the addresses of two linguistic agents, the first one (LA1) providing access to a WordNet 2.0 resource, the second one (LA2) offering a bilingual translation service from English to French based on a FreeDict dictionary[2]
- OA2 is requested by Agent A to find a counterpart for the DrivingWheel concept from the car-product ontology. Linguistic Agent LA1 is queried by OA2 for synonyms of this concept and, though the entry "driving_wheel" is present in WordNet, no synonyms are available for this term. OA2 then queries LA2 for obtaining translations of "driving wheel", obtaining the following ones:
  - "roue de transmission"
  - "roue d'entraînement"
  - "roue motrice"
- OA1 is queried by OA2 for the presence of the retrieved French terms in its voiture ontology. *RoueDeTransmission* is a concept in the voiture ontology, so it is a good candidate for matching the original *DrivingWheel* concept requested by Agent A.
- On-The-Fly meaning negotiations continue to arise between OA1 and OA2, following requests from Agent A. These may go even beyond simple translation of terms, requiring complex instance transformations and composition of objects from the two ontologies, though the initial linguistic investigation still remains the key to extract the evidences needed to reconcile the two different ontologies.

6.4. Scenario 4: Ontology location

In this scenario an Agent A wants to know the list of ontologies referring to the term *car*. The agent believes that such an ontology exists because it has received a natural language request from a user including this term. However, it has no idea of the kind of concepts underlying this symbol, and it would like to access its definition without any human intervention.

- Agent A wants to know the list of ontologies dealing with a given term.
- Agent A queries the DF for the list of the OAs available in the agency.
- Agent A queries in turn each OA for a list of ontologies including the given term.
- The OA queries each of the ontologies for which it is a provider, seeking concepts, properties or instances labelled with the given term.

This use case though can simply be extended (Figure 6): suppose an Agent A wants to know the list of ontologies modelling the *cars* domain, no matter the language they are expressed into. This example has a wider, more general scope than the previous one:

- Agent A wants to know the list of ontologies dealing with the term *car*.
- Agent A (which is probably a user agent), *knows* that the expression *car* requested by the user is in the English language (this information has been provided explicitly by the user, has been derived by the profile of the user, has been extracted by a log of his interaction with Agent A etc…).
- Agent A queries the DF for the list of the OAs available in the agency.
- Agent A queries in turn each OA for a list of ontologies including the given term.
- Each OA has already inspected the OLW of all the ontologies that it is able to access, and knows,

[2] http://www.dict.org/bin/Dict.

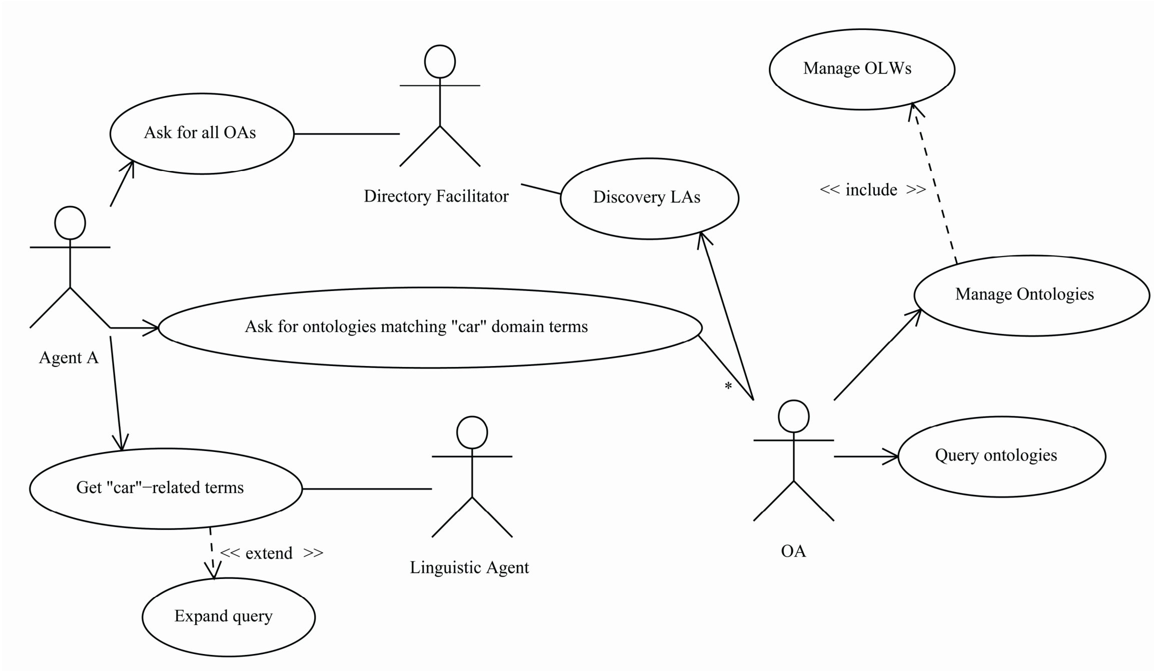

Figure 6: UML Use Case for Scenario 4

for each one, the set of natural languages available for accessing their content. Each OA has thus a complete list of natural languages that are required to search its ontologies (notice that even the expression “query” used in the FOSS implies not a true logical query to the ontology, but a free-textual retrieval of knowledge, which could be supported by indexing modules, search and string matching techniques etc…) and is able to compute the minimal list of languages (basing on those which are common to diverse OLWs) needed to support its search.

– Each OA queries the DF for accessing a list of linguistic agents which could support their natural language searches. This support is not limited to translations of the searched term for the required languages, but even to a wider understanding of the linguistic expressions which are available for the same language in which the original term is expressed. For example, it would be important to know that “auto”, “automobile”, “machine” and “motorcar” are all valid synonyms for the term *car*, because some ontologies could adopt one of those alternative expressions to refer to the same concept. Notice that, in those cases where the OLW reveals that a complete linguistic enrichment has been conducted on a given ontology through the adoption of a known lexical resource, the OA would decide not to use any linguistic agent because the ontology is expected to be already fully represented in the language adopted for that enrichment. For instance: if a complete linguistic enrichment for the English language through the WordNet resource has been conducted on a certain ontology, the OA would be sure that, if present, the concept *car* would surely contain the term “car” (together with its synonyms) among its available rdfs:labels, so there would be no need to ask for linguistic support to any linguistic agent.
– As an additional possibility (specific case of Figure 6), Agent A could have preventively queried the DF for linguistic agents providing monolingual support in the native language of the query (*English*). This way Agent A could interact with the user asking him to disambiguate the term wrt its different possible meanings (this would *expand* the query by maintaining at the same time pre-

cise results).

- The OA queries each of the ontologies it is able to access, seeking concepts, properties or instances labelled with the given term. Thanks to the OLW, the OA knows, for ***each*** ontology, which languages to use and, in case, which linguistic resources it can exploit to expand its query. The OA also knows from the OLW if its search can be performed only on the ontology names or it can be extended to the *rdfs:label*(s).

7. An inside look of the agent framework

A lot of work has been done by researchers towards flexible architectures to support semantic coordination in open multi-agent systems (Bayardo et al, 1997; Tamma & Bench-Capon, 2001; Li et al., 2005) or in modern web services architectures for the Semantic Web (Paurobally et al., 2006). The scenarios described in the previous phase (which were revisions of those published in the FOSS) merely anticipate the presence of an agent architecture supporting linguistic and semantic coordination. In this section we will provide more details about the agent framework which we think could add a decisive edge in supporting communication between heterogeneous semantic peers.

We will first introduce a *general purpose* agent taxonomy, describing each agent's peculiarities and prerogatives; we will then focus upon the relationships between the proposed agent roles, the former FIPA proposal for Ontology Services and the presented Linguistic Watermark ontologies, by describing a typical behavioral pattern between agents involved in a Semantic Coordination activity. Finally some details about a prototypical software implementation of the surveyed framework and references to past works will be given.

7.1. Linguistic Agents

Linguistic agents encapsulate one or more linguistic resources, and provide an interface to access their content, through different kind of services. They register themselves at a *Directory Facilitator* (DF) living in the agent platform, and publish services for getting resources' URIs and accessing their contents.

In presence of an Ontology Agent, the Linguistic Watermark Ontology describing the provided resource(s) can be explicitly sent to the OA itself which takes the charge of inspecting its content, thus making easier the discovery of the needed LR and their use in the previously exposed scenarios.

7.2. Domain/Resource/User Agents

A Resource Agent is, in a wide sense, an autonomous software entity living in an agent platform, having its own goals, knowledge and beliefs, likely to fit a well defined application domain and thus formally defined in a domain ontology. Whether it is a user agent offering an interface to the knowledge of a MAS, a Resource/Domain agent providing information from any kind of knowledge repository, we simply assume that this figure may "speak" in terms of a given conceptualization, which can be part of its inner knowledge (implicit ontology) or referred through explicit namespacing (explicit ontology).

It is – together with the mediator agent – one of the fundamental actors of the meaning negotiation process, as it represents the (virtual) end user of the communication: usually it is such an agent to start the semantic coordination procedure and to request a conceptual mediation against another agent's ontology.

For the semantic coordination to take part in an effective manner, each of the domain agents willing to communicate publish an explicit fingerprint of the conceptualizations they own (the Ontological Linguistic Watermark Ontology, OLW), containing information about the linguistic enrichment processes (as described in Section 5) which contributed to describe their knowledge content. As an ontology may have not been subject to any linguistic enrichment process at all though, the corresponding OLW may even be reduced to the sole information about the idiom used to represent its concept identifiers. FIPA Ontology Agents may as well play the same above role, in those cases where they provide access to an ontology which is explicitly referenced by a communicating domain agent.

By comparing ontologies’ watermarks, the linguistic coordinator (see next paragraph) agent will then be able to easily decide which supporting agents are likely to be contacted, and negotiate the (natural) language(s) to be used throughout the whole mapping process.

7.3. Agents for Linguistic Coordination (Linguistic Coordinators)

This would typically be identifiable as an agent role more than as an agent per se. Linguistic Coordinators are able to inspect the OLWs of the agents involved in a communication (being these their specific OLWs or the OLWs of ontologies provided by Ontology Agents and explicitly referenced by them), and decide which supporting agents (linguistic agents and mediator agents) need to be invoked to harmonize the conversation between the dialoguing agents.

They act more as conversation managers and do not embody real mediation skills for solving communication problems. Diverse agents of this framework could embody the role of Linguistic Coordinators: it is quite natural for Ontology Agents, as they need to be, for their specific role, aware of the linguistic characteristics of the ontologies they host.

Domain/User agents could do the same, so that they would be able to handle the situation by themselves, invoking those agents which can help their communication needs.

Finally, Mediator Agents (see next paragraph) could act both as mediators and coordinators, thus concentrating into one single entity all the skills needed to semantically coordinate a communication between two agents. This latter case would save much of the communication overhead between these two roles, and would also simplify and improve the whole process, as the coordinator role of the agent would more precisely consider those available linguistic agents whose content and services best fit its mediation capabilities.

7.4. Mediator Agents

While the Coordinator Agent can be considered as the “Manager” of the linguistic coordination, Mediator Agents play the role of the technicians. Mediator Agents are in fact in charge of negotiating the meaning of concepts exchanged during a communication between semantically incompatible peers.

We first need to distinguish between two different kinds of mediators: those which provide semantic coordination facilities by inspecting available ontology mapping documents, that is, resources available on the web which state conceptual correspondences between ontology resources, and those which offer the same service by evaluating correspondences on-the-fly.

The former will mostly carry on basic look-up over available explicit mapping resources (even more than one, as domain ontologies may have been built compositionally from several available smaller ones), supported by inferential abilities to entail new mappings other than those which are explicitly declared.

When no a-priori mapping is available (or it is available for only a part of the conceptualizations which need to be harmonized), these agents need to come to a solution by calculating mappings on demand, following the requests of the user/domain agents. This second kind of mediator agent usually encapsulates ontology matching techniques, instance migration and transformation services and other kind of solutions for reconciling differences between heterogeneous sources of information. They are informed by the Linguistic Coordinator agent about which linguistic resources can be exploited for performing their task (e.g. which agents have been invoked and are available for the current coordination activity), so that they can configure an optimal application of their mapping/matching skills.

7.5. Language-aware Ontology Agents

The Ontology Agent (OA) as it is defined in FIPA XC00086D, (2001), maintains ontologies by defining, modifying or removing terms and definitions contained in ontologies. It responds to queries about the terms in an ontology or relationships between ontologies.

The OA can provide the translation service of expressions between different ontologies or different content languages by itself, possibly as a wrapper to an ontology server.

An OA though is an agent that provides access to one or more ontology servers and which provides ontology services to an agent community. As well as all the other agents, the OA registers its service with the DF and it also registers the list of maintained ontologies and their available translations to other knowledge resources in order to allow agents to query the DF for the proper OA that manages a given ontology. Every agent can then request the services of the OA.

In our framework, OAs would be expected to have a good linguistic knowledge of the ontologies they host, so that they could provide the best search services according to the linguistic characteristics of each of its ontologies and of the agents querying their content. This linguistic awareness could be made concrete through inspection of ontologies' meta-data provided by their OLW. For this reason, OAs are often expected to behave as Linguistic Coordinators, though this role could as well be assumed by an external agent.

7.6. Linguistic and Semantic Coordination: Interaction between the introduced Agent Roles

The agent framework introduced so far opens up a wide range of possible interactions aiming at consolidating the required knowledge for handling under the best possible circumstances a Semantic Coordination activity. This wide variety of situations mainly suggests that the involved agents may be considered under every aspect as Intelligent Agents, as they have not only to reply to requests but they must understand and judge the situation in which they are operating and take decisions which can then be accepted or neglected by other agents, according to their availability and intentions. In Figure 8 a sequence diagram representing the typical interaction in a Semantic Coordination activity is represented: notice that, for easiness of reading, some details have been voluntarily masked (like the presence of a DF providing the addresses of the needed agents).

As shown in figure 8, following a first request for communication, the generic semantic coordination process usually has to start with a first hand-shake, in which information about the peers' respective knowledge is exchanged (usually in the form of ontology namespaces). Should the two agents discover that they are not able to communicate on the basis of a shared ontology vocabulary (as this is the case of the example in figure), they will first try to coordinate in order to invoke the figures which can support their communication. Agent A queries the DF for an OA providing static translation services between its

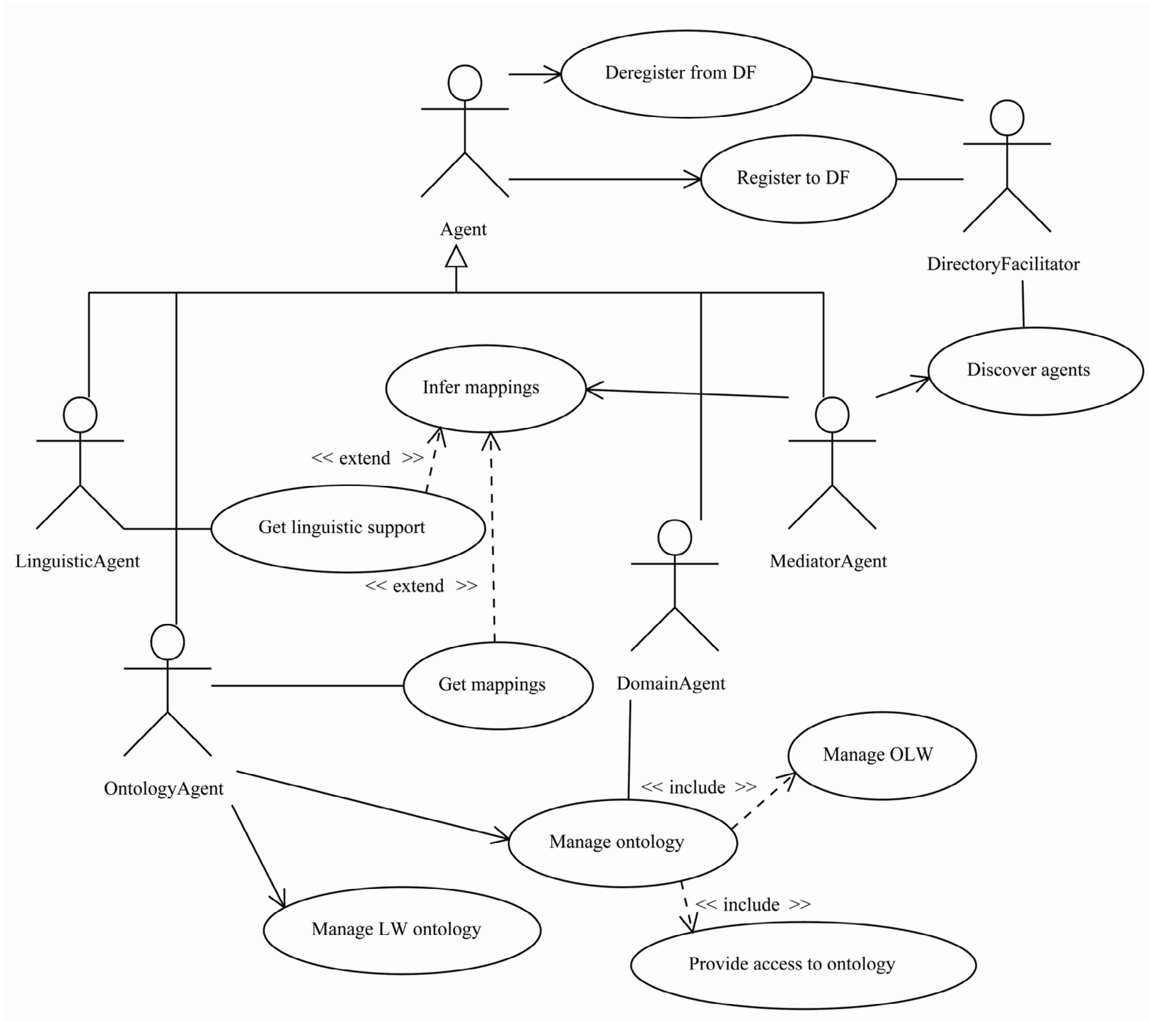

Figure 7: General use case sketch for the surveyed framework

ontologies and those of Agent B. Had the DF found such an OA, Agent A would have simply requested to it translations for its ontological terms. In the depicted situation, Agent A is instead returned with the address of an OA able to setup a dynamic semantic coordination: this case only requires a further step, with AgentA sending its OLW and the one of Agent B to the OA, so that it can setup the best environment for coordinating the two peers. This part of the semantic coordination is what we referred to as the *Linguistic Coordination* phase, in which the OA contacts a LinguisticCoordinator agent which inspects the OLWs of the two agents willing to communicate and invokes the resource and mediation agents needed to support the ontology mediation process.

The Linguistic Coordinator thus receives a request for its services by the OA and then, upon positive acknowledgement, receives a copy of the OLWs of Agent A and Agent B. After having analyzed the two watermarks, the Coordinator first decides which linguistic resources should be adopted (distinguishing between *helpful* and *necessary* resources) and then checks the existence of available linguistic agents providing access to these resources. If this first phase is successfully concluded, it then tries to contact a mediator agent able to conduct an ontology mediation activity by exploiting the available resources.

When the linguistic coordinator has checked that every needed resource and agent are available, the *linguistic coordination* phase ends and the coordinator can acknowledge the request from the OA by forwarding to it the addresses of the identified supporting agents.

Notice that, though in figure the agent roles have been kept separate, we could suppose the OA be endowed with Linguistic Coordination abilities¸ or even with mediation skills. Another possibility is that one of the two communicating agents could declare itself to be a coordinator, and thus assume control of the specific situation in which they are put. However, in the example, to clearly identify agent competences, we show the interaction between roles as they were different, separated, physical agents.

Also, for sake of simplicity, and without loss of generality, we limited to one the number of involved OAs and Linguistic Agents. In a realistic scenario, two OAs (one for Agent A and the other for Agent B) would plausibly be involved, while there is no theoretical limit to the number of Linguistic Resource (and their hosting Linguistic Agents) which could be summoned to provide support to the mediation activity; consider LingAgent as a collector for all the Linguistic Agents involved in the process.

When the linguistic coordination phase is concluded and Agent A is being acknowledge positively of its request for a Semantic Coordination, it can requests to the OA translation services from its ontology vocabulary to the one of Agent B. The OA then forwards the received requests to the chosen mediator agent, which carries over its assigned mediation tasks alternating reasoning and lexical analysis with submission of queries to the linguistic agents and to the OA (as in), to get new useful data.

While we do not enter into the details of the meaning negotiation activity which immediately follows the linguistic coordination (as they are out of the scope of this specific research and are, as stated in the introduction, widely discussed in literature), we stress that such an approach does not violate the dynamics of the process and the heterogeneous nature of actors and resources which are typically associated to the idea of Semantic Coordination. On the mere perspective of domain agents (which could be considered as the end users of Semantic Coordination), they just need to be aware of the existence of Ontology Agents (as in the original FIPA proposal) which act both as knowledge repositories as well as mediators (in the broader sense of the expression) in the community of agents.

7.7. Framework implementation

A former study about this agent framework had already been presented in (Pazienza & Vindigni, 2002), while a prototypical implementation – ALINAs (Architecture for LINguistic Agents) – has been provided in (Pazienza et al., 2002). In (Sguera et al., 2006) a newer version (J-ALINAs) of this architecture has been presented, reflecting recent research work about the linguistic enrichment of ontologies, and migrating the ALINAs architecture to the JADE agent platform (Bellifemine et al., 1999), which actually simplifies the development of agent platforms through a middleware layer complying to FIPA standards, implementing some of the required agents for the platform to be FIPA-compliant, and allowing agents to be movable from one machine to another.

The latest framework implementation, J-ALINAs, maintains much of the same classification of agent roles which has been presented above. Yet this approach is not to be considered as a rigid one: indeed the boundary between agents will sometimes be not so clear, and will be possible for an agent to behave on more than a side.

J-ALINAs encompasses a set of JADE behaviors, realizing the system's reactive layers hierarchy. In the following paragraphs we will describe the most significant ones, *moving* our perspective from one agent to another. Further details may be found in (Sguera et al., 2006).

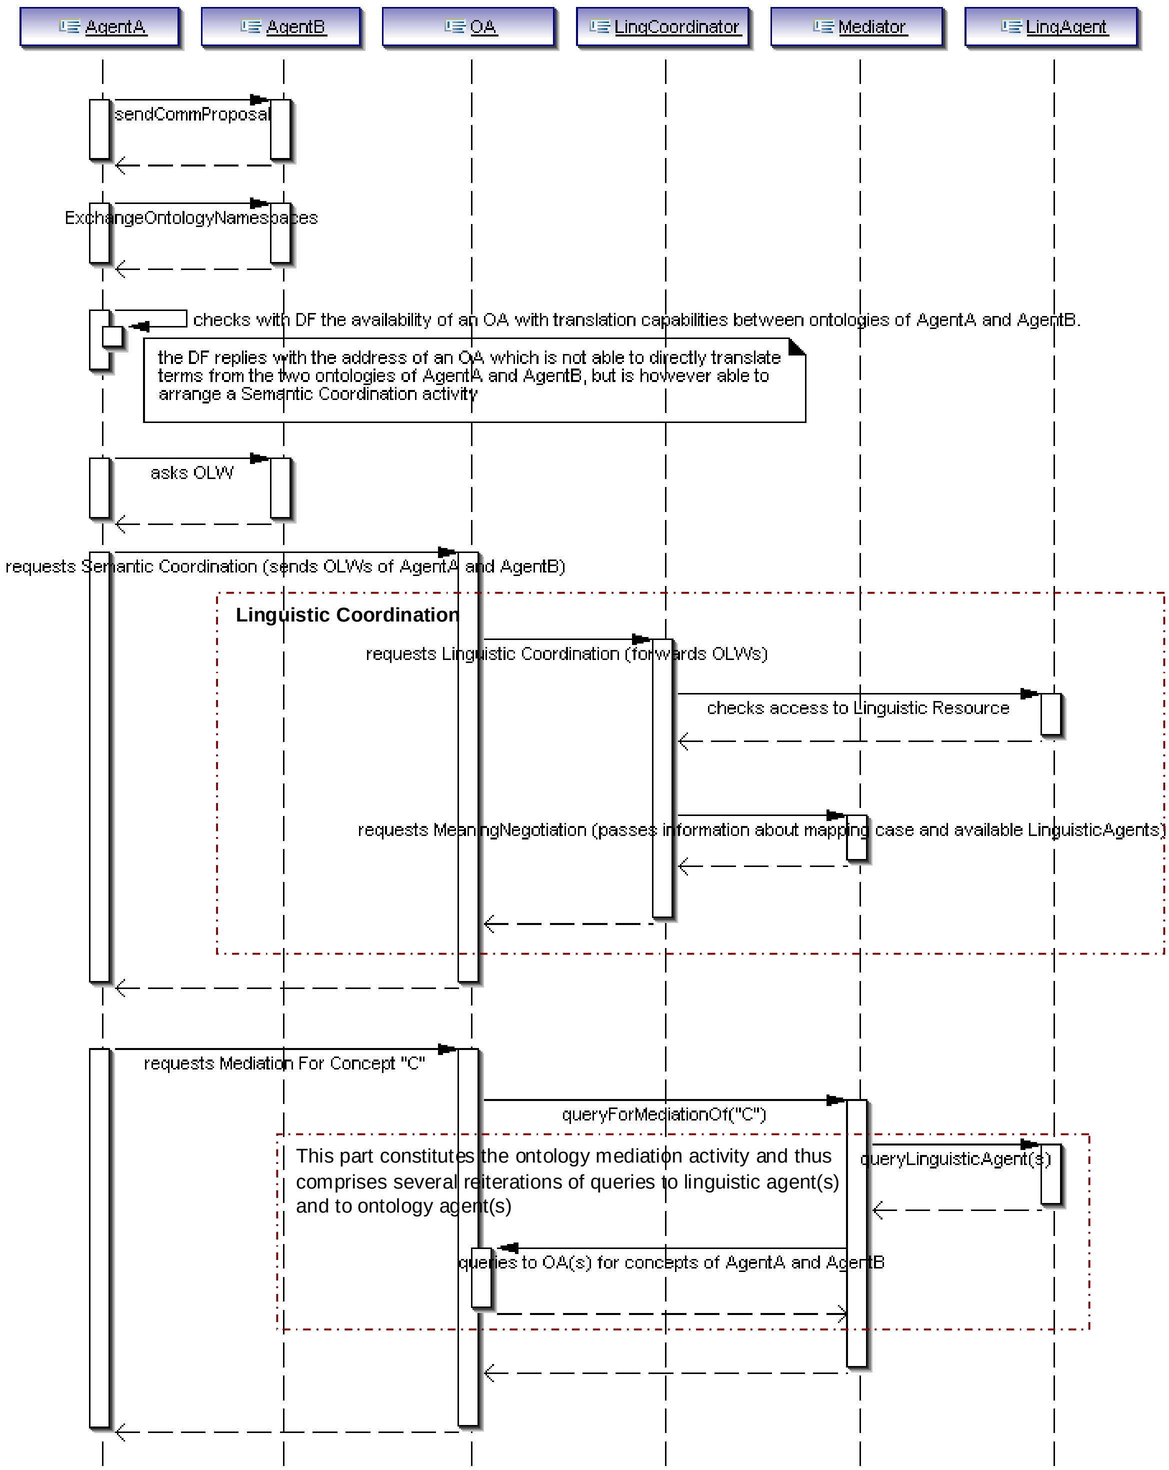

Figure 8: A general sequence diagram for the typical Semantic Coordination scenario

6.7.1 Linguistic support to meaning negotiation

The behavior encapsulating access to linguistic resources and providing linguistic support to the meaning negotiation process is a cyclic one, enabling the linguistic agent to react to mediator agent's stimuli. In other words it models a simple reactive behavior, implementing a reactive agent paradigm.

A linguistic agent in the ***setup*** procedure will add an instance of such behavior to its execution queue; it will eventually reply to queries originated by the mediator agent with semantic anchors for the particular term, or with a *NOT_UNDERSTOOD* FIPA message performative in such cases where the query is malformed (or, more simply, out of its comprehension).

Queries are matched by the agent against an application-defined message template, providing developers with high freedom of choice with respect to the type and number of linguistic services the agent publishes in the system.

The ***linguistic-querying*** task is delegated to a specific behavior residing in the mediator agent, which is responsible for persisting anchors upon the agent itself; however the concept of ***persisting*** anchors is absolutely abstract for the framework, and it is responsibility of the application developer to implement the ***persist*** and ***retrieval*** functions.

6.7.2 The twofold domain agent behavior

The provided domain agent paradigm has been designed to act indifferently as ***requester*** or ***responder*** in the mapping process. We actually reckon this has positive implications both towards bidirectional knowledge mediation between agents, and the integration of the application-specific ontological agent into the surrounding environment, reducing the number of different ***actors*** for the developers to take into account.

If the agent is required to act as ***requester*** in the process it will start discovering the surrounding environment, contacting agents and ***handshaking*** them.

Otherwise, if the stimulus is a *PROPOSE* FIPA message performative, the agent will eventually accept to participate in a mediation and wait for a mediator agent to ask for its ontological linguistic watermark or information about its ontology concepts and relations among them.

6.7.3 The domain agent's introspection

Concurrently, the domain agent runs an instance of a ***simple reflex*** behavior aimed to answer specific queries about ontology's instances, properties or relations among concepts. This behavior can be exploited by developers to fit messages exchange among agents to support the specific mapping algorithms, and does not affect the handshaking procedure at all.

6.7.4 The generic mediation process' instance

Modules providing ontology mediation based upon external linguistic resources have been designed to give the handshaking procedure a ***well-known*** form, and at the same time to give developers using the framework the chance to easily ***extend*** the system with custom behaviors encapsulating application-specific mapping algorithms.

Indeed, after the ***handshaking*** phase took place the mediator agent shall own all the elements needed to conduct a ***generic*** ontology mediation, and to effectively decide which other agents in the society are likely to be involved in the process; in other words we provide developers embracing J-ALINAs architecture a solid and shared base to design their application on, building a further abstraction level from the underlying environment.

The abstraction itself comes to evidence even from a strictly ***sequential*** perspective: because every mediation process starts with a communication among agents and the system's facilitator whereas the agents discover the surrounding environment, we reckon it is convenient to let developers the chance not to care – to some extent – about the *coordination* and *synchronization* layer, focusing on their system's peculiar aspects.

8. Conclusions

In this essay we have presented an agent framework, and a set of ontologies, formalizing different aspects of agents communication, starting from and extending the specifications expressed in the FIPA XC00086D document on Ontology Services. We believe that, though the FIPA proposal needs not to be reformulated in its original intentions, it requires further specifications on the way agents can coordinate to improve their chances of communication. We showed how, while the diverse techniques and algorithms for ontology mediation can be hidden as technical details in the scope of a semantic coordination activity, the premises which allow for this process to be conducted need to be formalized, exposed and thus shared on a common basis, by agents willing to communicate.

To this end, the proposed Linguistic Watermark ontology and Ontological Linguistic Watermark metadata endow agents embracing our described framework, with the necessary "knowledge about themselves" (in particular, pertaining to mere linguistic aspects of their knowledge) which can be exhibited to other agents to better understand the concrete possibilities of establishing a communication, and exploited by supporting agents (directory facilitator, ontology agents and those specialized figures we introduced in our framework) to make these possibilities become ground truth.

References

Basili, R., Vindigni, M. and Zanzotto, F.M. Integrating Ontological and Linguistic Knowledge for Conceptual Information Extraction. IEEE Computer Society, *WI '03: Proceedings of the IEEE/WIC International Conference on Web Intelligence*, Washington, DC, USA, pages 175, 2003

Bayardo, R. J. J., Bohrer, W., Brice, R., Cichocki, A., Fowler, J., Helal, A., Kashyap, V., Ksiezyk, T., Martin, G., Nodine, M., Rashid, M., Rusinkiewicz, M., Shea, R., Unnikrishnan, C., Unruh, A. & Woelk, D. (1997) *InfoSleuth: Agent-Based Semantic Integration of Information in Open and Dynamic Environments* ACM Press, Proceedings of the ACM SIGMOD International Conference on Management of Data, New York, 26,2, pages 195-206, 13--15~ 1997

Bellifemine, F., Poggi , A. and Rimassa, G., JADE - A. FIPA-compliant agent framework, in Proc. of PAAM'99. (London, April 1999), pp.97-108.

Benjamins V. R., Contreras, J., Corcho O. & Gómez-Pérez, A. (2004). Six Challenges for the Semantic Web. *SIGSEMIS Bulletin, April 2004*

Benz, A., G. Jäger & R. van Rooij (2005), An introduction to game theory for linguists, in Benz et al., (eds.), *Game Theory and Pragmatics*, Palgrave MacMillan, 1-82.

Bouquet, P.; Serafini, L.; Zanobini, S.: Peer-to-Peer Semantic Coordination. Journal of Web Semantics, 2(1), 2005

Buitelaar, P., Declerck, T., Frank, A.,Racioppa, S., Kiesel, M., Sintek, M., Engel, R., Romanelli, M., Sonntag, D., Loos, B., Micelli, V., Porzel, R., & Cimiano, P. (2006). LingInfo: Design and Applications of a Model for the Integration of Linguistic Information in Ontologies. *Proceedings of the OntoLex Workshop at LREC, pp. 28-32. May 2006.*

Burstein, M.H. & Uschold, M. (2005). Infrastructure for Semantic Interoperability and Integration: Breakout Discussion Summary. 2005

Chaudhri, V. K., Farquhar, A., Fikes, R., Karp, P. D., & Rice, J. P. (1998). OKBC: A Programmatic Foundation for Knowledge Base Interoperability. *Proc. AAAI'98 Conference, Madison, WI, July 1998*

Cole, R. A., Mariani, J., Uszkoreit, H., Zaenen, A. & Zue, V. (1997) Eds. Survey of the State of the Art in Human Language Technology, Cambridge University Press, Cambridge, UK, 1997.

Dean, M., & Schreiber, G. (2004). Eds. OWL Web Ontology Language Guide. *W3C Recommendation (10 February 2004).*

de Bruijn, J., Martín-Recuerda, F., Manov, D. and Ehrig, M.: State-of-the-art survey on Ontology Merging and Aligning. Project Deliverable d4.2.1, SEKT project IST-2003-506826

de Bruijn, J. and Polleres, A.: *Towards and Ontology Mapping Specification Language for the Semantic Web.* DERI Technical

Report DERI-2004-06-30, 2004

Fellbaum, Christine. Ed. (1998). *WordNet: An Electronic Lexical Database.* Cambridge, MA: MIT Press

FIPA XC00086D, (2001). FIPA Ontology Service Specification. *http://www.fipa.org/specs/fipa00086/XC00086D.html*

Francopoulo,G., George, M., Calzolari, N., Monachini, M., Bel, N., Pet, M., Soria., C. *LEXICAL MARKUP FRAMEWORK (LMF), LREC2006, Geneva, Italy, 2006.*

van Hage, W. R., Katrenko, S. and Schreiber, G. A Method to Combine Linguistic Ontology-Mapping Techniques. *International Semantic Web Conference*, pages 732-744, 2005

Jäger, G. (2004). Evolutionary Game Theory and typology: a case study. Manuscript, University of Potsdam and Stanford University

Jarrar, M.: Towards the notion of gloss, and the adoption of linguistic resources in formal ontology engineering. Proceeding of the 15th International World Wide Web Conference, WWW2006. Edinburgh, Scotland. May 2006. ACM, 2006

Kalfoglou, Y. and Schorlemmer, M. *Ontology mapping: the state of the art.* In The Knowledge Engineering Review Journal, number 18(1), pages 1–31, 2003

Li, L., Wu, B. & Yang, Y. (2005). Semantic mapping with multi-agent systems, *2005 IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE'05)*.

Maedche, A., Motik, B., Silva, N. and Volz, R. Mafra: a mapping framework for distributed ontologies. *In Proceedings of the 13th European Conference on Knowledge Engineering and Knowledge Management EKAW-2002*, Madrid, Spain, 2002

Miller, G. A., Beckwidth, R., Fellbaum, C., Gross, D., Miller, K.J., Introduction to WordNet: An On-Line Lexical Database, *International Journal of Lexicography*, 3, no. 4, 1990, 235-244

Oltramari, A., Huang, C., Lenci, A., Buitelaar, P., Fellbaum, C. (2006) *OntoLex2006: Second Workshop on Interfacing Ontologies and Lexical Resources for Semantic Web Technologies*, held jointly with LREC2006, Magazzini del Cotone Conference Center, Genoa, Italy

Paurobally, S., Tamma, V. & Wooldridge, M. (2005). Cooperation and agreement between Semantic Web Services. *W3C Workshop on Frameworks for Semantics in Web Services, 2005.*

Pazienza, M.T., Stellato, A. and Vindigni, M. ALINAs:un'architettura multilayer ad agenti per il supporto alla comunicazione linguistica Dagli oggetti agli agenti (WOA2002), Milano (Italy), 18 November 2002

Pazienza, M. T. and Stellato, A. (2006a). An Environment for Semi-automatic Annotation of Ontological Knowledge with Linguistic Content. *3rd European Semantic Web Conference (ESWC 2006) Budva, Montenegro, June 11-14, 2006.*

Pazienza, M. T. and Stellato, A. (2006b) Exploiting Linguistic Resources for building linguistically motivated ontologies in the Semantic Web. *Second Workshop on Interfacing Ontologies and Lexical Resources for Semantic Web Technologies (OntoLex2006), held jointly with LREC2006, Magazzini del Cotone Conference Center, Genoa, Italy, 24-26 May 2006.*

Pazienza, M. T. & Stellato, A. (2006c). Linguistic Enrichment of Ontologies: a methodological framework. *Second Workshop on Interfacing Ontologies and Lexical Resources for Semantic Web Technologies (OntoLex2006), held jointly with LREC2006 ,Magazzini del Cotone Conference Center, Genoa, Italy, 24-26 May 2006*

Pazienza, M.T., Stellato, A. (2006d). An open and scalable framework for enriching ontologies with natural language content. *The 19th International Conference on Industrial, Engineering & Other Applications of Applied Intelligent Systems (IEA/AIE'06), special session on Ontology & Text. Annecy, France, June 27-30, 2006*

Pazienza, M. T. & Stellato, A. (2005a). Linguistically motivated Ontology Mapping for the Semantic Web. *SWAP 2005, the 2nd Italian Semantic Web Workshop, Trento, Italy, December 14-16, 2005.*

Pazienza, M. T. & Stellato, A. (2005b). The Protégé Ontoling Plugin - Linguistic Enrichment of Ontologies in the Semantic Web. *4th International Semantic Web Conference (ISWC-2005) Galway, Ireland, November, 2005*

Pazienza, M.T., Stellato, A., Henriksen, L., Paggio, P., & Zanzotto, F. M. (2005). Ontology Mapping to support ontology-based question answering. *Proceedings of the second MEANING workshop. Trento, Italy, February 2005*

Pazienza, M.T., Stellato, A., Vindigni, M., Valarakos, A., & Karkaletsis, V. (2003). Ontology integration in a multilingual e-retail system. *HCI International 2003, Crete, Greece, 2003*

Pazienza, M. T. & Vindigni, M. (2002). Agent based ontological mediation in IE systems. *LNAI 2700, Summer Convention on Information Extraction 2002, Information Extraction in the Web Era*, 92-128.

Rahm, E. & Bernstein, P. (2001). A survey of approaches to automatic schema matching. *In Very Large Databases Journal, number 10(4),* 334-350.

Reed, C., Norman, T.J., Jennings, N.R. Negotiating the semantics of agent communication languages. *Computational Intelligence* 18 (2002) 229—252

van Rooij, R. (2004). Signalling games select Horn strategies. *Linguistics and Philosophy*, 27, 493–527.

van Rooij, R. and K. Schulz (2004). Exhaustive interpretation of complex sentences. *Journal of Logic, Language and Information*, (13), 491–519.

Sguera, S., Stellato, A., Griesi, D., & Pazienza, M. T. (2006). J-ALINAs: A JADE-based Architecture for Linguistic Agents. *Proceedings of WOA 2006. Catania, Italy, September, 26-27, 2006.*

Shvaiko, P. and Euzenat, J. *A Survey of Schema-Based Matching Approaches*. J. Data Semantics IV, pages 146-171, 2005

Sure, Y., Corcho, O., Euzenat J., and Hughes, T. eds.: Proceedings of the 3rd Evaluation of Ontology-based tools (EON), located at the 3nd International Semantic Web Conference ISWC 2004, Hiroshima, Japan, November 2004

Tamma, V. & Bench-Capon, T.J.M. (2001). A conceptual model to facilitate knowledge sharing in multi-agent systems. *In Proceedings of the Autonomous Agents 2001 Workshop on Ontologies in Agent Systems (OAS 2001). Montreal, May, 29th, 2001,* 69-76.

Uschold, M. & Gruninger, M. (2005). Architectures for semantic integration. *Dagstuhl Seminar Proceedings: Semantic Interoperability and Integration 2005.*