Argumentation over Ontology Correspondences in MAS

Loredana Laera
University of Liverpool, UK
lori@csc.liv.ac.uk

Ian Blacoe
University of Liverpool, UK
blacoe@csc.liv.ac.uk

Valentina Tamma
University of Liverpool, UK
valli@csc.liv.ac.uk

Terry Payne
University of Southampton, UK
trp@ecs.soton.ac.uk

Jerôme Euzenat
INRIA Rhône-Alpes, France
Jerome.Euzenat@inrialpes.fr

Trevor Bench-Capon
University of Liverpool, UK
tbc@csc.liv.ac.uk

ABSTRACT

In order to support semantic interoperation in open environments, where agents can dynamically join or leave and no prior assumption can be made on the ontologies to align, the different agents involved need to agree on the semantics of the terms used during the interoperation. Reaching this agreement can only come through some sort of negotiation process. Indeed, agents will differ in the domain ontologies they commit to; and their perception of the world, and hence the choice of vocabulary used to represent concepts.

We propose an approach for supporting the creation and exchange of different arguments, that support or reject possible correspondences. Each agent can decide, according to its preferences, whether to accept or refuse a candidate correspondence. The proposed framework considers arguments and propositions that are specific to the matching task and are based on the ontology semantics. This argumentation framework relies on a formal argument manipulation schema and on an encoding of the agents' preferences between particular kinds of arguments.

1. INTRODUCTION

Multi-Agent Systems have proved very valuable for solving problems spanning more than one organisation, and in domains where many factors may be dynamic, and thus cannot be readily considered and anticipated at design-time. Open and dynamic environments, such as the Web and its proposed extension, the Semantic Web [4], are by nature distributed and heterogeneous. In these environments *ontologies* [18] are expected to complement agreed communication protocols in order to facilitate mutual understanding and interactive behaviour between such agents. Thus, agents may differ in the domain ontologies they commit to [12]; and in their perception of the world (and hence the choice of vocabulary used to represent concepts). Imposing a single, universally shared ontology on agents is not only impractical because it would result in assuming a standard communication vocabulary (and thus violate the dynamics of open environments); but it also does not take into account the conceptual requirements of agents that could appear in future. Therefore, if their ontologies are different, these need to be reconciled in order to support interaction.

The reconciliation of heterogeneous ontologies in open environments, where no prior assumption can be made on the ontologies used, depends on the different agents' ability to find an agreement on the semantics of the terms used during the interoperation. The agreement usually takes the form of an *alignment* between the ontologies, that is a set of correspondences (or *mappings*) between the concepts, properties, and relationships in the agent ontologies; and on their use to interpret or translate messages exchanged. Reaching this agreement can only come through some sort of negotiation process [1].

The motivation for our work lies in open environments, where one agent wishes to communicate with another in order to request a service. The agents are aware of each other's existence, but have no knowledge about the services that the other may offer. This situation arises when the agents are able to utilise a registry service, such as the FIPA Directory Facilitator, but are not able to understand the services offered due to the fact that they utilise different ontologies. Therefore, the agents need to agree on an ontology alignment to determine whether the services offered match the requested ones. In order to generate a suitable alignment the agents are able to access a mapping repository, OMR, that stores possible mappings between ontology components. These mappings originate from independent mapping engines that employ differing algorithms to calculate potential correspondences between different ontologies.

In this paper we present an approach that makes use of argumentation theory in order to dynamically reach an agreement over heterogeneous ontologies. The set of potential arguments are clearly identified and grounded on the underlying ontology language, and the types of correspondences supported by any such argument are clearly specified. We base our approach on [13] and we tailor it to the specific needs of agent communication, and we evaluate our framework against existing tools for ontology mapping. The novel contribution of the approach presented here lies in the preferences that agents can express on the types of correspondences to use when aligning the ontologies. The argumentation framework allows the agents to reach an agreement on those correspondences that are *mutually acceptable*, because they are determined by those mappings that cannot be refuted by other agents. This approach better reflects the characteristics of autonomy and rationality that are typical

AAMAS'07, May 14–18, 2007, Honolulu, Hawai'i, USA.

978-81-904262-7-5 (RPS) ©2007 IFAAMAS

of agents [21]. A value-based argumentation framework [3] is used to express agents' preferences between the categories of arguments that are built on the types of correspondences.

2. MEANING-BASED ARGUMENTATION

Meaning-based Argumentation is a process that dynamically and automatically enables agents, with different terminologies and interests, to reach consensus on the terminology they use to interact - by *arguing*. This gives them the ability to understand each other enough to carry out their objectives, and in our specific context to request a service. The terminology of each agent is represented according to its own conceptualisation, which we assume is explicitly specified according to its own ontology. An ontology O can be defined as a tuple $\langle C, \leq_C, R, \leq_R, I\rangle$, where C and R are two disjoint sets and their elements are called, respectively, *concepts* and *properties*. $\leq_C$ and $\leq_R$ are partial orders on C and R respectively and are called *concept hierarchy* and *relation hierarchy*. The elements of I may be instances of concepts in C and may be interconnected with other elements in I by a relation in R. Such elements are known as *individuals*. We assume that all agents' ontologes are encoded in the same language, the web standard, OWL[1].

In order to make ontologies interoperable, so that the terms in different ontologies are brought into correspondence, we need to provide mappings. These mappings can be provided by a variety of different matching algorithms, such as the ones participating to the Ontology Alignment Evaluation Initiative (OAEI)[2]. An alignment consists of a set of correspondences between the two ontologies. A correspondence (or a mapping) is described as a tuple: $m = \langle e, e', n, R\rangle$, where e and e' are the entities (concepts, relations or individuals) between which a relation is asserted by the correspondence; n is a degree of confidence in that correspondence; and R is the relation (e.g., equivalence, more general, etc.) holding between e and e' asserted by the correspondence [15]. A correspondence over which no agreement has yet been reached by the agents is called a *candidate mapping*. Moreover, we assume that for each correspondence m, the repository is able to provide a set of justifications G, that explain why it has generated the candidate mapping. Such information forms the basis which an agent can dynamically generate arguments and supply the reasons for their mapping choices. In addition, each agent has a (partial or total) pre-ordering of preferences over different types of ontology mismatches ($Pref$), and a private threshold value ε which is compared to the degree of confidence associated with each mapping. These preferences are based on the motivations of the agent, and determine whether a mapping is accepted or rejected. Currently, only a few approaches for ontology alignment provide such justifications [16]. However, tools such as [9] combine different similarity metrics, and these measures can be used to build the required justifications. Moreover, for obvious efficiency reasons, the argumentation approach is applied only over the ontological terms the agents need to understand each other, rather than on the whole ontology.

3. THE ARGUMENTATION APPROACH

The argumentation framework for autonomous agents proposed in this section is based on Value-based Argument Frameworks ($VAFs$) [3]. We start with the presentation of Dung's work [8], upon which the $VAFs$ rely.

Definition An Argumentation Framework (AF) is a pair $AF = \langle AR, A\rangle$, where AR is a set of arguments and $A \subset AR \times AR$ is the *attack* relationship for AF, and is comprised of ordered pairs of distinct arguments in AR. A pair $\langle x, y\rangle$ is referred to as 'x *attacks* y', while a set of arguments S is said to attack an argument y if y is attacked by an argument in S.

An argumentation framework can be simply represented as a directed graph whose vertices are the arguments and whose edges correspond to the elements of A.

In Dung's framework, attacks always succeed. This is reasonable when dealing with deductive arguments, but in many domains, including the one under consideration, arguments lack this coercive force: they provide reasons which may be more or less persuasive, and their persuasiveness may vary according to their audience. Therefore, we distinguish *attacks* from *successful attacks*, i.e., those which defeat the attacked argument. We use a Value-based Argumentation Framework, which prescribes different strengths to arguments on the basis of the values they promote and the ranking given to these values by the audience for the argument. This allows us to systematically relate strengths of arguments to their motivations, and to accommodate different audiences with different interests and preferences.

Definition A *Value-Based Argumentation Framework* (VAF) is defined as $\langle AR, A, \mathcal{V}, \eta\rangle$, where (AR, A) is an argumentation framework, $\mathcal{V}$ is a set of k *values* which represent the types of arguments and $\eta\colon AR \rightarrow \mathcal{V}$ is a mapping that associates a value $\eta(x) \in \mathcal{V}$ with each argument $x \in AR$.

In section 3.1, the set of values $\mathcal{V}$ is defined as the different types of ontology mismatch, which we use to define the categories of arguments and to assign to each argument one category.

Definition An *audience* for a VAF is a binary relation $\mathcal{R} \subseteq \mathcal{V} \times \mathcal{V}$ whose (irreflexive) transitive closure, $\mathcal{R}^*$, is asymmetric, i.e. at most one of (v, v'), (v', v) are members of $\mathcal{R}^*$ for any distinct v, $v' \in \mathcal{V}$. We say that v_i *is preferred to* v_j *in the audience* $\mathcal{R}$, denoted $v_i \succ_{\mathcal{R}} v_j$, if $(v_i, v_j) \in \mathcal{R}^*$. Let $\mathcal{R}$ be an audience, α is a *specific audience* (compatible with $\mathcal{R}$) if α is a *total ordering* of $\mathcal{V}$ *and* $\forall\ v,\ v' \in \mathcal{V},\ (v, v') \in \alpha \Rightarrow (v', v) \notin \mathcal{R}^*$.

In this way, we take into account that different agents (represented by different audiences) can have different perspectives on the same candidate mapping. Given a set of arguments and counter arguments, it is necessary for the agents to consider which of them they should accept. Acceptability of an argument is defined in the following way:[3]

Definition Let $\langle AR, A, \mathcal{V}, \eta\rangle$ be a VAF and $\mathcal{R}$ an audience. For arguments x, y in AR, x is a *successful attack* on y *with respect to the audience* $\mathcal{R}$ if: $(x, y) \in \mathcal{A}$ *and* it is *not* the case that $\eta(y) \succ_{\mathcal{R}} \eta(x)$. An argument x is *acceptable to the*

---

[1] http://www.w3.org/OWL/

[2] http://oaei.ontologymatching.org/

[3] Note that all these notions are now relative to some audience.

subset S with respect to an audience $\mathcal{R}$ if: for every $y \in AR$ that *successfully attacks* x with respect to $\mathcal{R}$, there is some $z \in S$ that successfully attacks y with respect to $\mathcal{R}$. A subset S of AR is *conflict-free with respect to the audience* $\mathcal{R}$ if: for each $(x, y) \in S \times S$, either $(x, y) \notin \mathcal{A}$ or $\eta(y) \succ_{\mathcal{R}} \eta(x)$. A subset S of AR is *admissible* with respect to the audience $\mathcal{R}$ if: S is conflict free with respect to $\mathcal{R}$ and every $x \in S$ is acceptable to S with respect to $\mathcal{R}$. A subset S is a *preferred extension* for the audience $\mathcal{R}$ if it is a maximal admissible set with respect to $\mathcal{R}$.

The key notion here is the *preferred extension* which represents a consistent position within AF, which is defensible against all attacks and which cannot be further extended without becoming inconsistent or open to attack. In order to determine whether the dispute is resolvable, and if it is, to determine the preferred extension with respect to value orderings promoted by distinct audiences, [3] introduces the notion of objective and subjective acceptance as follows:

Definition Given a VAF, $\langle AR, A, \mathcal{V}, \eta \rangle$, an argument $x \in AR$ is *subjectively acceptable* if and only if, x appears in the preferred extension for some specific audiences but not all. An argument $x \in AR$ is *objectively acceptable* if and only if, x appears in the preferred extension for *every* specific audience. An argument which is neither objectively nor subjectively acceptable is said to be *indefensible*.

3.1 Arguments about mappings

In this paper we focus on arguments about mappings. We define these arguments as follows:

Definition An argument $x \in AF$ is a triple $x = \langle G, m, \sigma \rangle$ where m is a correspondence $\langle e, e', n, R \rangle$; G is the grounds justifying a prima facie belief that the correspondence does, or does not hold; and σ is one of $\{+, -\}$ depending on whether the argument is that m does or does not hold.

A key issue is that the interaction between arguments is based on a notion of attack; an argument x is attacked by the assertion of its negation $\neg x$, namely the *counter-argument*, defined as follows:

Definition An argument $y \in AF$ *rebuts* an argument $x \in AF$ if x and y are arguments for the same mapping but with different signs, e.g. if x and y are in the form $x = \langle G_1, m, + \rangle$ and $y = \langle G_2, m, - \rangle$, x counter-argues y and vice-versa.

Moreover, if an argument x supports an argument y, they form the argument $(x \rightarrow y)$ that attacks an argument $\neg y$ and is attacked by argument $\neg x$. These arguments are clearly identified and grounded on the underlying ontology language OWL. Therefore, the grounds justifying correspondences can be extracted from the knowledge in ontologies. This knowledge includes both the extensional and intensional OWL ontology definitions. Our classification of the grounds justifying correspondences is the following:

semantic (M): the sets of models of two entities do or do not compare;
internal structural (IS): two entities share more or less internal structure (e.g., the value range or cardinality of their attributes);
external structural (ES): the set of relations, each of two entities have, with other entities do or do not compare;
terminological (T): the names of two entities share more or less lexical features;
extensional (E): the known extensions of two entities do or do not compare.
These categories correspond to the type of categorizations underlying ontology matching algorithms [20].
In our framework, we use the types of arguments described above as types for the VAF; hence $\mathcal{V} = \{M, IS, ES, T, E\}$. Therefore, for example, an audience may specify that terminological arguments are preferred to semantic arguments, or vice versa. Note that this may vary according to the nature of the ontologies being aligned. Semantic arguments are given more weight in a fully axiomatised ontology, compared to that in a lightweight ontology where there is very little reliable semantic information on which to base such arguments. Table 1 presents a sample of reasons for the justification of candidate OWL ontology alignments. The table represents an (extensible) set of argument schemes, the instantiations of which include AR. Attacks between these arguments arise when we have arguments for the same mapping but with conflicting values of σ, thus yielding attacks that can be considered symmetric. Moreover, the relations in the mappings can also give rise to attacks: if relations are not deemed exclusive, an argument against inclusion is a fortiori an argument against equivalence (which is more general).
For example, given a candidate mapping $m = \langle c, c', \_, \equiv \rangle$ between two OWL ontologies O_1 and O_2, with concepts c and c' respectively, an argument for accepting the mapping m may be that the labels of c and c' are synonymous. An argument against may be that no mapping is defined for some of their sub-concepts.

3.2 Argument generation process

In this section we examine the way arguments are generated and the main characteristics of agent architecture used in order to exchange arguments on correspondences. Each agent Ag_i has access to its individual ontology:

Definition An agent Ag_i is characterised by a 4-tuple $\langle O_i$, VAF_i, $Pref_i$, $\varepsilon_i \rangle$ where O_i is the OWL ontology; $VAF_i = \langle AR_i, A_i, \mathcal{V}, \eta_i \rangle$ is the Valued-based Argumentation Framework of the agent Ag_i; $Pref_i$ is the private pre-ordering of preferences over $\mathcal{V}$ and ε_i is the private threshold value.

A set of agents $A = \{Ag_1, \ldots, Ag_n\}$ forms a multi-agent system (MAS). The set of arguments shared by all agents are not necessarily disjoint, while the values $\mathcal{V}$ are common and shared by all agents.
As mentioned, the preferences and threshold selected by an agent depend on its context and situation. A major feature of this context is the agent's ontology, and its structural features, such as the depth of the subclass hierarchy and branching factor, ratio of properties to concepts, etc. The analysis of the structural components of the ontology is inline with the approach to ontology evaluation, demonstrated in [6], and can be formalized in terms of feature metrics. An agent can then determine its preferences and threshold based on the characteristics of its ontology. For example, selecting a preference for terminological mapping if the ontology is lacking in structure, or preferring extensional mapping if its ontology is rich in instances.
In our framework, the arguments and counter-arguments are generated by an agent using these preferences and thresholds. Specifically, we assume n agents, committing to different ontologies, and a mapping repository, OMR, storing

Table 1: Argument scheme for OWL ontological alignments

Mapping	σ	Grounds	Comment
$\langle e, e', n, \equiv\rangle$	+	$\exists m_i = \langle ES(e), ES(e'), n', \equiv\rangle$	e and e' have mapped neighbours (e.g., super-entities, etc.) of e are mapped in those of e'
$\langle e, e', n, \sqsubseteq\rangle$	+	$\exists m_i = \langle ES(e), ES(e'), n', \equiv\rangle$	(some or all) neighbours (e.g., super-entities, etc.) of e are mapped in those of e'
$\langle e, e', n, \equiv\rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv\rangle$	(some or all) Instances of e and e' are mapped
$\langle e, e', n, \sqsubseteq\rangle$	+	$\exists m_i = \langle E(e), E(e'), n', \equiv\rangle$	(some or all) Instances of e are mapped to those of e'
$\langle e, e', n, \equiv\rangle$	-	$\nexists m_i = \langle E(e), E(e'), n', \equiv\rangle$	No instances of e and e' are mapped
$\langle e, e', n, \equiv\rangle$	+	$label(e) \approx_T label(e')$	Entities's labels share lexical features (e.g., synonyms, and lexical variants)

the several sets of correspondences between the agent ontologies, generated by a variety of different mapping services. Different sets of correspondences can be generated by mapping services depending on the methods used and their configurations[4]. In order to enable the agents to come to agreement on a suitable alignment for the service requested, without requiring a complete alignment between the ontologies, the requesting agent specifies which of the components from its own ontology are involved in the service request, and only seeks to generate an alignment with the ontology of the other agent with regard to these entities.

Given an agent $Ag_i = \langle O_i, VAF_i, Pref_i, \varepsilon_i\rangle$, and a candidate mapping $m \in OMR$ with a set of justifications G, the agent Ag_i first evaluates the acceptability of m. A mapping m is accepted by an agent Ag_i if there exist justifications G for m that correspond to the highest preference $Pref_i$ (with respect to the pre-ordering), assuming n is greater than its private threshold ε_i. Consequently, the agent generates a set of arguments $x = \langle G, m, +\rangle$, by instantiating the argumentation schema. In the other case, the agent Ag_i rejects the mapping m and it generates arguments against: $x = \langle G, m, -\rangle$ (see algorithm 1). Figure 1 depicts the architecture of our approach for achieving dynamic negotiation of ontology alignments between agents requesting and providing services. The pattern of communication between the requesting agent, providing agent and the mapping repository is shown in table 2.

Algorithm 1 Generation of Arguments

```
Require: a set of agents ${Ag_1, ..., Ag_n}$, an ontology alignment repository $OMR$ , a set of candidate mappings $m_j$
Ensure: a set of arguments and counter-arguments $x_k$
for all agent $Ag_i = \langle O_i, VAF_i, Pref_i, \varepsilon_i\rangle$ do
    for all mapping $m_j = \langle e_j, e'_j, n_j, R_j\rangle$ do
        if $n \geq \epsilon$ and $\exists\ \mathcal{G}_j \in \mathcal{G}$ such that $\mathcal{G}_j = Pref$ then
            Generate arguments for $m : x_k = \langle G_k, m, +\rangle$
        else
            Generate arguments against $m : x_k = \langle G_k, m, -\rangle$
        end if
    end for
end for
```

Table 2: Communication between Agents

1	$Agent_1$ wants a service X and knows which components of its own ontology O_1 are involved in requesting X
2	$Agent_1$ knows $Agent_2$ exists, but nothing else.
3	$Agent_1$ requests $Agent_2$'s ontology O_2.
4	$Agent_1$ sends $Agent_2$ the terms from O_1 involved in X.
5	Both agents get the candidate mappings from OMR.
6	Both agents exchange their arguments sets.
7	Both agents instantiate their argumentation frameworks.
8	Both agents then calculate their preferred extensions.
9	Both agents then determine the agreed alignment, by exchanging these preferred extensions.
10	$Agent_1$ then sends its service request X to $Agent_2$.
11	$Agent_2$ compares X with its service descriptions - $S2_1, S2_2$, etc. using the agreed alignment, and sends any matching service descriptions to $Agent_1$.
12	$Agent_1$ examines these service descriptions using the alignment to confirm that they match the required service X.
13	If a matching service is confirmed, $Agent_1$ sends the service request to $Agent_2$.

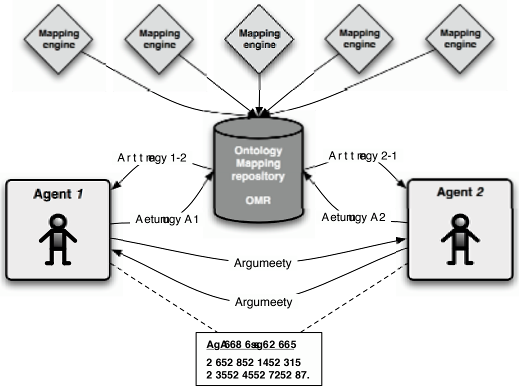

Figure 1: Architecture

4. AGREED & AGREEABLE ALIGNMENTS

Although in $VAFs$ there is always a unique non-empty preferred extension with respect to a specific audience, provided the AF does not contain any cycles in a single argument type [3], an agent may have multiple preferred extensions either because no preference between two values in a cycle has been expressed, or because a cycle in a single value exists. The first may be eliminated by committing to a specific audience, but the second cannot be eliminated in this way. In our domain, where many attacks are symmetric, two cycles are frequent and in general an audience may have multiple preferred extensions. Thus, given a set of arguments justifying mappings organised into an argumentation framework, an agent is able to determine which mappings are acceptable by computing the preferred extensions with respect to its preferences. If there are multiple preferred extensions, the agent must commit to the arguments present in all preferred extensions, but has some freedom of choice with respect to those in some but not all of them. This partitions arguments into three sets: *desired arguments*, present in all preferred extensions, *optional arguments*, present in some but not all, and *rejected arguments*, present in none. If we have two agents belonging to different audiences, these

---

[4] Note that the mapping repository does not store every possible mapping between entities, but applies a threshold to the submitted mappings in order to store only those plausibly supported by the ontological knowledge. The assumption of using mapping services does not imply any loss of generality, since these services are now becoming increasingly available. See: http://oaei.ontologymatching.org/

sets may differ. Doutre et al. [7] describe a means by which agents may negotiate a joint preferred extension on the basis of their partitioned arguments so as to maximise the number of desired arguments included, whilst identifying which optional arguments need to be included to support them.

Based on the above considerations, we thus define an *agreed alignment* and an *agreeable alignment* as follows. An *agreed alignment* is the set of correspondences supported[5] by those arguments which are in every preferred extension of every agent. An *agreeable alignment* extends the agreed alignment with those correspondences supported by arguments which are in some preferred extension of every agent. Whilst the mappings included in the agreed alignments can be considered valid and consensual for all agents, the agreeable alignments have a uncertain background, due to the different alternative positions that each agent can take. However, given our context of agent communication, we seek to accept as many candidate mappings as possible. We therefore take into consideration both set of alignments - agreed and agreeable.

4.1 Instantiating argumentation frameworks

In order to reach agent consensus about ontology alignments, first we have to build the argumentation frameworks and evaluate them to find which arguments are agreed and agreeble. There are four main steps in applying our argumentation approach:

Each agent individually constructs an argumentation framework for each candidate mapping, by considering the repertoire of argument schemes available to it, and instantiating these schemes with respect to its interests. Each argument either supports or rejects the conclusion that the mapping is valid. Having established the set of arguments, the agent then determines the attacks between them by considering their mappings and signs, and the other factors discussed above. This step produces several VAFs for reasoning about the candidate correspondences.

Given a MAS, each agent considers its individual frameworks with all the argument sets of all the other agents and then extends the attack relations by computing the attacks between the arguments present in its framework with the other arguments. Then, for each VAF, the agents determine which arguments are undefeated by attacks from other arguments. The algorithm in [3] can be employed for computing the preferred extensions of a VAF given a value ordering. The global view is considered by taking the union of these preferred extensions for each audience.

Finally, each agent considers which arguments are in every preferred extension of every audience. The mappings that have only favourable arguments are included in the agreed alignments, and the mappings that have only arguments against are rejected. For mappings whose acceptability cannot be established, agents extend the search space to consider those arguments which are in some preferred extension of every audience. The mappings supported by those arguments are part of the set of agreeable alignments.

The dialogue between agents can thus consist of the exchange of arguments sets, from which agents can individually compute acceptable mappings. If necessary and desirable, these can then be reconciled into a mutually acceptable position through a process of negotiation, as suggested in [7] which defines a dialogue process for evaluating the status of arguments in a VAF, and shows how this process can be used to identify mutually acceptable arguments. In constructing a position, an ordering of values best able to satisfy the joint interests of the agents concerned is determined. However, such issues are the subject of ongoing research.

The above technique considers sets of mappings and complete argumentation frameworks. If instead the problem is to determine the acceptability of a single mapping it may be more efficient to proceed by means of a dialectical exchange, in which a mapping is proposed, challenged and defended.

4.2 Argumentation protocol

In this section we briefly introduce an argumentation protocol which can be used to evaluate the acceptability of a single mapping. The idea behind the protocol is to allow n agents ($n \geq 2$) to argue about the acceptability of a potential mapping, arriving at a joint solution that is based on their preferences and the information they exchange during argumentation.

Formally, our argumentation protocol is a tuple $\langle$ Mapping, Agents, Acts, Replies, Move, Dialogue, Result$\rangle$ such that:

Mapping: is a candidate mapping $m \in OMR$, subject to evaluation between the agents.

Agents: is the set of agents taking part in the dialogue, $Ag_1, \ldots, Ag_n$

Acts: is the set of possible argumentation speech acts: $Acts = \{Support, Contest, Withdraw\}$

Replies:$Acts \rightarrow Acts$ is a mapping that associates to each speech act its possible replies: *Replies(Support)* = { *Support, Contest, Withdraw* } ; *Replies(Contest)* = { *Support, Contest, Withdraw* } ; *Replies(Withdraw)* = $\emptyset$.

Move: is a tuple $M_i = \langle S_i, H_i, Move_i \rangle$, where $S_i \in Agents$ is the agent which makes the move; $H_i \subseteq Agents$ is the set of agents to which the move is addressed and $Move_i = Act_i(c)$ is the uttered move, with $Act_i(c)$ as the speech act applied to a content c.

Dialogue: is a finite non-empty sequence of moves $Dialogue = \{M_0, \cdots, M_p\}$

Result: $Mapping \rightarrow \{agreed, agreeable, rejected\}$ is a mapping which returns the result of the acceptability of the mapping m.

Initially, each agent Ag_i evaluates the acceptability of the mapping m, and then generates a set of arguments $\{x_1^+ = \langle G_1, m, + \rangle, \ldots, x_r^+ = \langle G_r, m, + \rangle\}$ or counter-arguments $\{x_1^- = \langle G_1, m, - \rangle, \ldots, x_r^- = \langle G_r, m, - \rangle\}$. Each agent Ag_i sends $support(m, \{x_1^+, \ldots, x_r^+\})$ or $contest(m, \{x_1^-, \ldots, x_r^-\})$ to the other agents. If the agent does not have any arguments or counter-arguments to propose, then it can withdraw by $Withdraw(m)$ and the dialogue terminates. Each agent Ag_i checks whether it agrees with the set of mappings $m_1, \ldots, m_k$ used in the arguments that are exchanged. If they have arguments or counter-arguments to present, they start a new dialogue to evaluate each mapping in $m_1, \cdots, m_k$. The dialogue terminates when each mapping involved in the initial dialogue on m has been evaluated or when an agent withdraws[6]. When the dialogue terminates, each agent builds the argument framework, and checks the acceptability of the mapping by computing its preferred extensions.

---

[5] A mapping m is *supported* by an argument $x = \langle G, m, + \rangle$.

[6] Termination of the dialogue is ensured by having a finite number of candidate mappings, plus a time-out to handle undecidable arguments.

5. AN ILLUSTRATIVE EXAMPLE

We illustrate the ideas presented in this paper with the aid of a simple example, where agent Ag_1 interacts with another agent Ag_2 to request a service. We assume that the only terms used for such a request are $\mathit{Paper\_Author}$ and $\mathit{Research\_Topic}$. The first agent, Ag_1, uses the EKWA ontology[7]; whereas the second agent, Ag_2 uses the OpenConf Ontology[8]. For space reasons, we only consider a subset of these ontologies, shown in Figure 2, where the first and second ontologies are represented by O_1 and O_2 respectively. Ag_2 does not know if it is able to satisfy the request until the ontologies are aligned. Thus, the agents access the mapping repository, OMR, that stores the following set of possible correspondences, originating from different mapping engines that may employ differing algorithms:

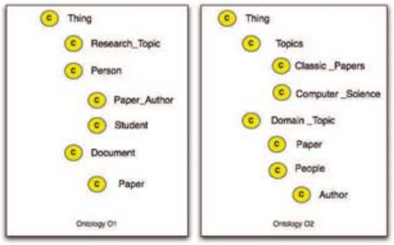

Figure 2: O_1 and O_2 ontologies

$m_1{=}\langle O_1\colon \mathit{Paper\_Author}, O_2\colon \mathit{Author}, 0.45, =\rangle$; [9]
$m_2{=}\langle O_1\colon \mathit{Paper\_Author}, O_2\colon \mathit{Paper}, 0.54, =\rangle$;
$m_3{=}\langle O_1\colon \mathit{Research\_Topic}, O_2\colon \mathit{Topics}, 0.44, =\rangle$;
$m_4{=}\langle O_1\colon \mathit{Research\_Topic}, O_2\colon \mathit{Domain\_Topic}, 0.44, =\rangle$;
$m_5{=}\langle O_1\colon \mathit{Person}, O_2\colon \mathit{People}, 0.9, =\rangle$;

Note that the agents also argue about the mapping m_5 since it is part in the argumentation. Assume now that Ag_1 selects the audience $\mathcal{R}_1$, which prefers terminology to external structure, $(T \succ_{\mathcal{R}_1} ES)$. Ag_2 selects the audience $\mathcal{R}_2$, which prefers external structure to terminology $(ES \succ_{\mathcal{R}_2} T)$. The pre-ordering of preference *Pref* corresponds to each agent's audience. Moreover, all the above candidate mappings have a degree of confidence n that is above the threshold of each agent, and so this does not influence their acceptability.

Agent Ag_1 accepts all of the correspondences, whilst Ag_2 accepts only the mapping m_1 and rejects the mappings m_2, m_3 and m_4. Both agents accept m_5. The arguments and counter-arguments generated are shown in Table 3 overleaf, that shows each argument, labeled with an identifier Id, its type $\mathcal{V}$, the attacks A that can be made on it by opposing arguments, and the agent that proposes the argument. Based upon these arguments and the attacks, each agent can construct the argumentation framework which brings the arguments together so that they can be evaluated, shown in Figure 3. The nodes represent arguments (labelled with their Id) with the respective type value $\mathcal{V}$. The arcs represent the attacks A, and the direction of the arcs represents the direction of the attack.

We have two arguments supporting the correspondence m_1.

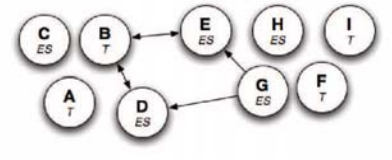

Figure 3: Argumentation framework

Argument A states that $O_1\colon \mathit{Paper\_Author}$ and $O_2\colon \mathit{Author}$ do have lexical similarity. Argument C supports the mapping m_1 since their ancestors, $O_1\colon \mathit{Person}$ and $O_2\colon \mathit{People}$, are mapped as established by the mapping m_5. Moreover, we have two arguments against m_2, and one for it: D is against the correspondence m_2, since none of the superconcepts of $O_1\colon \mathit{Paper\_Author}$ are mapped to any superconcept of $O_2\colon \mathit{Paper}$; E is against the correspondence m_2, since none of the sibling-concepts of $O_1\colon \mathit{Paper\_Author}$ are mapped to any sibling-concepts of $O_2\colon \mathit{Paper}$; and C argues for m_2, because $O_1\colon \mathit{Paper\_Author}$ and $O_2\colon \mathit{Paper}$ do have lexical similarity. The arguments D and E attack symmetrically the argument B. The arguments F and G justify the mapping m_5, since, respectively, the labels of $O_1\colon \mathit{Person}$ and $O_2\colon \mathit{People}$ are synonymous and their subconcepts are mapped by m_1 and m_2. Clearly, argument G attacks the arguments D and E. Finally, the arguments H and I state that the concept $O_1\colon \mathit{Research\_Topic}$ is mapped with $O_2\colon \mathit{Domain\_Topic}$ and $O_2\colon \mathit{Topic}$ since both have lexical similarity. The agent Ag_2 does not have any counter-arguments to rebut H and I.

Given the two audiences, $\mathcal{R}_1$ and $\mathcal{R}_2$, the preferred extensions achieved are shown in Table 4.

The arguments that are accepted by both audiences are $\{A, C, G, F, H, I\}$. Thus the *agreed alignment* is m_1, m_3, m_4 and m_5, while the mapping m_2 is rejected (since B is unacceptable to $\mathcal{R}_2$).

Table 4: Preferred Extensions

Preferred Extensions	Audience
$\{I,H,B,C,G,A,F\}$, $\{I,H,B,G,C,E,A,F\}$, $\{I,D,H,B,C,G,E,A,F\}$, $\{I,H,D,B,G,C,A,F,G\}$	$\mathcal{R}_1$
$\{I,D,H,B,C,G,E,A,F\}$, $\{I,H,D,G,C,E,A,F\}$	$\mathcal{R}_2$

6. EVALUATION

In order to gauge the effectiveness of the proposed approach, we have evaluated it empirically. The experiments aim to measure how the argumentation affects the accuracy of the possible alignments generated from the mappings in the *OMR*. For this purpose, we use the OAEI benchmark test suite. This test set consists of one reference ontology for a bibliographic domain, to be compared with other ontologies. Most of these ontologies originate from the reference ontology by making arbitrary changes. The experiments involved only two agents, Ag_1 and Ag_2 - with Ag_1 having the reference ontology in each test and Ag_2 having the respective test ontology. Their threshold has been set to zero, and so will not influence the process. The evaluation sets we used are the following:

simple tests: The reference ontology is compared with itself, with another irrelevant ontology or the same ontology restricted to OWL-Lite. Tests 101, 102, 103, 104.

systematic tests: The reference ontology is compared with modified ones. These modifications involved discarding some features, such as names, comments, hierarchy, instances, re-

---

[7] http://nb.vse.cz/∼svabo/oaei2006/data/ekaw.owl

[8] http://nb.vse.cz/∼svabo/oaei2006/data/OpenConf.owl

[9] m_1 states an equivalence correspondence with confidence 0.45 between the concept $\mathit{Paper\_Author}$ in the ontology O_1 and the concept Author in the ontology O_2.

Table 3: Arguments for and against the correspondences m_1, m_2, m_3, m_4, m_5

Id	Argument	A	$\mathcal{V}$	Agent
A	$\langle Label(Paper\_Author) \approx_T Label(Author), m_1, +\rangle$		T	Ag_1, Ag_2
B	$\langle Label(Paper\_Author) \approx_T Label(Paper), m_2, +\rangle$	D,E	T	Ag_1
C	$\langle \exists m = \langle superclass(Paper\_Author), superclass(Author), 0.67, \equiv, \rangle, m_1, +\rangle$		ES	Ag_1, Ag_2
D	$\langle \nexists m = \langle superclass(Paper\_Author), superclass(Paper), 0, \equiv, \rangle, m_2, -\rangle$	B	ES	Ag_2
E	$\langle \nexists m = \langle sibling(Paper\_Author), sibling(Paper), 0, \equiv, \rangle, m_2, -\rangle$	B	ES	Ag_2
F	$\langle Label(Person) \approx_T Label(People), m_5, +\rangle$		T	Ag_1, Ag_2
G	$\langle \exists m = \langle subclass(Person), subclass(People), \equiv, \rangle, m_5, +\rangle$	D,E	ES	Ag_1, Ag_2
H	$\langle Label(Research\_Topic) \approx_T Label(Domain\_Topic), m_3, +\rangle$		T	Ag_1
I	$\langle Label(Research\_Topic) \approx_T Label(Topic), m_4, +\rangle$		T	Ag_1

lations, restrictions, etc. Tests 201, 202, 204, 205, 206, 221, 222, 223, 224, 225, 228, 230.

complex tests: The reference ontology is compared with four real-life ontologies for bibliographic references found on the web and left unchanged. Tests 301, 302, 303, 304.

We used standard information retrieval metrics to assess the results of our tests: *Recall*, *Precision* and *F-measure*. *Precision* measures the ratio between the number of correct mappings and the number of all mappings found. *Recall* measures the ratio between the number of correct mappings and the total number of correct mappings that should be found. *F-measure* combines the measures of precision and recall as single measure. In order to evaluate how the argumentation influences the matching accuracy, we calculated the *F-measure* before and after the argumentation. The preferences $Pref_1$ and $Pref_2$, chosen on the basis of the ontological information, for Ag_1 and Ag_2 are presented in Table 5. The experimental results are shown in Figure 6.a. The pre-argumentation F-measure represents the matching accuracy over all mappings present in the *OMR*, without considering the agents' preferences. The post-argumentation F-measure represents the matching accuracy achieved by the agents using our approach, based on the same initial set of candidate mappings. The results demonstrate that our approach not only allows the selection of those mappings that best suit the interests of each agent, but, in most cases, improves retrieval in terms of *F-measure*. We have also compared our argumentation approach with a number of other ontology alignment tools representative of the state-of-the-art. These tools[10] are: Foam [10], Falcon [14] and OWL-Lite Alignment (OLA) [11]. The experiments conducted on each tool are the same as those described above, and the results of this comparison can be seen in Figure 6.b. The results demonstrate that, in the majority of the test cases, our argumentation approach produces an information recall F-measure that is in a similar range to those produced by the other approaches. Examination of the results in the different test groups shows that for the *simple tests* all the approaches produce comparable results, with, for example, test 102 showing a 0 F-measure because this test ontology has a non-overlapping domain and so no alignment would be expected. In the *systematic tests* the results for argumentation are only slightly below those for the other approaches, with two main exceptions. For tests 205 and 206 the argumentation produces a 0 F-measure because the information in the two ontologies causes the agents to select directly opposing preferences, which leads to an inability to reach agreement on many of the mappings. For tests 201 and 202, argumentation performs less well than the other approaches due to a lack of candidate mappings in the repository to argue over (because in these test ontologies the concept labels have been replaced by random strings). In the *complex tests* our approach can be seen to perform well in comparison to the other tools (out-performing OLA and performing similarly to Foam) which is largely due to the fact that the four test ontologies involved are real, independently engineered ontologies over the same domain, that provide richer ontological information than many of the other tests.

Table 5: Ag_1 and Ag_2's preferences for each Test

Test	$Pref_1$	$Pref_2$
101,102,103,104,204,224,225,228,230	$T \succ ES$	$T \succ ES$
201,202	$T \succ ES$	$ES \succ IS \succ T$
205,206,223	$T \succ ES$	$ES \succ T$
221,301,302	$T \succ ES$	$T \succ IS$
222,304	$T \succ ES$	$T \succ IS \succ ES$
303	$T \succ ES$	$T \succ ES \succ IS$

a) Pre and Post Argumentation F-measure

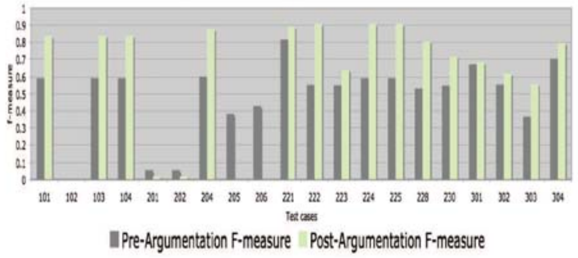

b) Comparison with Falcon, Foam and OLA

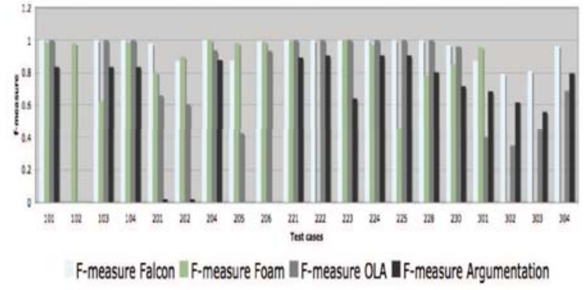

Figure 4: F-measure tests and comparisons

7. RELATED WORK

Few approaches have addressed the use of argumentation or negotiation between agents w.r.t. ontology alignments. An ontology mapping negotiation [17] has been proposed to establish a consensus between different agents using the MAFRA alignment framework. It is based on the utility and meta-utility functions used by the agents to establish if

---

[10] The tools selected were the best performers in OAEI'05.

a mapping is accepted, rejected or negotiated, but is highly dependent on the MAFRA framework and cannot be flexibly applied in other environments. van Diggelen et al. [19] present an approach for agents to agree on a common ontology in a decentralised way. Rather than being the goal of any one agent, the ontology mapping is a common goal for every agent in the system. Beun et al. [5] present a computational framework for the detection of ontological discrepancies in multiagent systems by using feedback utterances. Bailin and Truszkowski [2] present an ontology negotiation protocol that enables agents to exchange parts of their ontology, by a process of successive interpretations, clarifications, and explanations. The end result of this process is that each agent will converge on a single, shared ontology, whereas, in our context, agents keep their own ontologies.

Many ontology alignment tools have been proposed in the area of the Semantic Web. QOM [9], and its extension Foam[10], are based on heuristically calculated similarity of the individual ontology entities, and is distinguished by an emphasis on the efficiency of alignment. OLA [11] is dedicated to the alignment of OWL-Lite ontologies, and aims to use all the available information (i.e. lexical, internal and external structure, extensional, and data-types) extracted from two given ontologies. Falcon [14] is an automatic tool for aligning ontologies, which employs three distinct matchers in combination: a string-similarity matcher, a Vector Space Model of domain terms, and an RDF graph matcher. None of these approaches consider the preferences of agents seeking to align ontologies, however, each of them could be used as a mapping engine in our architecture.

8. CONCLUSIONS

In this paper we have outlined a framework that provides a novel way for agents, with different ontologies, to agree upon an ontology alignment. This is achieved using an argumentation process in which candidate correspondences are accepted or rejected, based on the ontological knowledge and the agent's preferences. This will give agents the ability to understand each other enough to carry out their objectives, for example to request a service. We believe that this approach will facilitate reaching more sound and effective mutual understanding and communication in a multi-agent-system. Currently, the preferences are limited to a few general types of ontology mismatch. In future, we will extend the preferences to suit more ontology alignment algorithms and also consider their combination. We are intending to demonstrate the effects of varying the preferences on the quality of the alignments reached, and evaluate the framework for more than two agents. Moreover, in future work we will investigate how to argue about the entire alignment, rather than only the individual candidate mappings. These arguments could occur when a global similarity measure between the whole ontologies is applied.

Acknowledgements This work was partially supported by Knowledge Web (FP6-IST 2004-507482) and PIPS (FP6-IST 2004-507019).

9. REFERENCES

[1] K. Aberer and *et al.* Emergent semantics principles and issues. In *Proc. of DASFAA*, 2004.

[2] S. C. Bailin and W. Truszkowski. Ontology Negotiation: How Agents Can Really Get to Know Each Other. In *Proc. of the WRAC 2002*, 2002.

[3] T. Bench-Capon. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation*, 13(3):429–448, 2003.

[4] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5), 2001.

[5] R. Beun, R. van Eijk, and H. Prüst. Ontological feedback in multiagent systems. In *Proc. of Int. Conference on Autonomous Agents and Multiagent Systems*, 2005.

[6] K. Dellschaft and S. Staab. On how to perform a gold standard based evaluation of ontology learning. In *Proc. of the ISWC*, 2006.

[7] S. Doutre, T. Bench-Capon, and P. E. Dunne. Determining Preferences through Argumentation. In *Proceedings of AI*IA'05*, pages 98–109, 2005.

[8] P. Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-person Games. In *Artificial Intelligence*, volume 77, pages 321–358, 1995.

[9] M. Ehrig and S. Staab. QOM - Quick Ontology Mapping. In *Proc. of the ISWC*, 2004.

[10] M. Ehrig and Y. Sure. Foam - framework for ontology alignment and mapping; results of the ontology alignment initiative. In *Proc. of the Workshop on Integrating Ontologies*, 2005.

[11] J. Euzenat, D. Loup, M. Touzani, and P. Valtchev. Ontology alignment with ola. In *Proceedings of the 3rd EON Workshop, ISWC*, 2004.

[12] T. R. Gruber. A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2):199–220, 1993.

[13] L. Laera and *et al.* Reaching agreements over ontology alignments. In *Proc. of the Fifth International Semantic Web Conference (ISWC'06)*, 2006.

[14] G. C. N. Jian, W. Hu and Y. Qu. Falcon-AO: Aligning Ontologies with Falcon. In *Proc. of the K-CAP Workshop on Integrating Ontologies*, 2005.

[15] P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on Data Semantics*, 4:146–171, 2005.

[16] P. Shvaiko, F. Giunchiglia, P. Pinheiro da Silva, and D. McGuinness. Web explanations for semantic heterogeneity discovery. In *Proceedings of ESWC*, 2005.

[17] N. Silva, P. Maio, and J. Rocha. An Approach to Ontology Mapping Negotiation. In *Proc. of the Workshop on Integrating Ontologies*, 2005.

[18] R. Studer, V. Benjamins, and D. Fensel. Knowledge engineering, principles and methods. *Data and Knowledge Engineering*, 25(1-2):161–197, 1998.

[19] J. van Diggelen, R. Beun, F. Dignum, R. van Eijk, and J.-J. Meyer. A decentralized approach for establishing a shared communication vocabulary. In *Proc. of the AMKN*, 2005.

[20] P. Visser, D. Jones, T. Bench-Capon, and M. Shave. Assessing Heterogeneity by Classifying Ontology Mismatches. In *Proc. of the FOIS'98*, 1998.

[21] M. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *Knowl. Eng. Rev.*, 10(2):115–152, 1995.