2009 International Conference on Computational Science and Engineering

Reputation-Based Ontology Alignment for Autonomy and Interoperability in Distributed Access Control

Daniel Trivellato[1] Fred Spiessens[1] Nicola Zannone[1] Sandro Etalle[1,2]
[1]Eindhoven University of Technology
[2]University of Twente
{d.trivellato,a.o.d.spiessens,n.zannone,s.etalle}@tue.nl

***Abstract*—Vocabulary alignment is a main challenge in distributed access control as peers should understand each other's policies unambiguously. Ontologies enable mutual understanding among peers by providing a precise semantics to concepts and relationships in a domain. However, due to the distributed nature of ontology development, ontology alignment is required to allow peers to make informed access control decisions. The alignment should be flexible and accurate to not undermine the autonomy and reliability of peers. This paper addresses the problem of ontology alignment in distributed access control by combining ontology-based trust management with a reputation system.**

I. Introduction

Trust management (TM) [1], [2], [3], [4] is an approach to access control in distributed systems in which access control decisions are based on credentials issued by multiple authorities. One of the main problems of existing TM systems is that they implicitly assume a complete agreement among authorities on the vocabulary for the specification of credentials. This assumption is too restrictive in dynamic coalitions in which parties do not know each other a priori. Moreover, existing TM systems only provide support for a syntactical consensus on the vocabulary. Reaching a semantic agreement is particularly challenging in coalitions of heterogeneous systems.

A partial solution to these problems is offered by the use of ontologies. Ontologies are increasingly seen as a means for enabling interoperability across heterogeneous systems [5]: augmenting trust management with ontologies allows mutual understanding among peers. An example of ontology-based TM framework is our previous work POLIPO (Policies and OntoLogies for Interoperability, Portability, and autOnomy) [6]. POLIPO is a TM system that enables interoperability in dynamic coalitions of heterogeneous systems by requiring every party to specify credentials in terms of concepts from a global shared ontology, understandable by all parties.

However, there are many situations in which this solution is still not satisfactory. Consider, for instance, a dynamic coalition in the Maritime Safety and Security (MSS) domain. Here, different parties may join the coalition on the fly, and a TM system is needed to regulate access to data exchanged among them. In such a setting, it is unrealistic to assume that all peers reach a complete and precise semantic alignment before becoming operative; more likely, they will agree on a common vocabulary (the global ontology), and then they will attempt to realize a mapping from local concepts to concepts in the global ontology.

Such a mapping, though, may be imprecise (e.g., the meaning given to concepts such as "senior officer", "vessel", etc, could differ from peer to peer). If the global ontology has no concepts that exactly describe the attribute of the subject intended by a peer, the peer is forced to issue approximate credentials. Moreover, to disclose sensitive information, a peer receiving a credential expressed in terms of the global ontology has to trust the approximation made by the issuer.

Requiring peers to specify credentials only in terms of the global ontology ensures interoperability among them but may affect credential accuracy because of possible misalignment between local and global concepts. The problem of misalignment occurs every time TM is used in heterogeneous systems or in short- and mid-term cooperation (which is often the case in B2B); in practice, it is faced when reaching a complete and precise alignment would be too costly or take too much time.

In this paper, we propose a solution to the problem of semantic alignment in TM by extending ontology-based TM with a reputation system. In particular, we enhance POLIPO, making the following contributions:

- we increase the autonomy of peers in a distributed system by allowing local vocabulary in credentials;
- we preserve interoperability by providing a technique for ontology alignment;
- we propose and evaluate a metric to compute the ontology alignment accuracy of a peer;
- we propose and evaluate a metric to assess the semantic relatedness of two concepts.

To enable parties to "speak" their own language while preserving global understanding, we adopt the notion of similarity [7], [8], [9]. Similarity represents the degree of semantic resemblance between two concepts. In our approach, similarity is asserted in the form of credentials. Each entity can specify the similarity between two concepts independently from the ontology in which those concepts have been defined.

In this setting, similarity credentials represent subjective appraisals, and each peer may have a different opinion on the similarity between two concepts. For this reason, we combine different similarity statements to assess the similarity between

978-0-7695-3823-5/09 $26.00 © 2009 IEEE
DOI 10.1109/CSE.2009.123

White Star Navy	Admiral ≻ Commodore ≻ Captain-1stRank ≻ Captain-2ndRank ≻ Lieutenant ≻ Midshipman
Grey Cross Navy	GeneralAdmiral ≻ Admiral ≻ Captain ≻ Commander ≻ Lieutenant-Commander ≻ Lieutenant
POSEIDON	OF-6 ≻ OF-5 ≻ OF-4 ≻ OF-3 ≻ OF-2 ≻ OF-1

Fig. 1. Chains of Command

two concepts. Due to the different alignment capabilities of peers, however, the accuracy of similarity assertions may vary from peer to peer. Therefore, when combining different statements, we need to weigh them according to the trustworthiness of their issuer, which reflects their accuracy. Reputation is the metric that is usually adopted to assess the trustworthiness of entities in distributed systems [10], [11], [12], [13], [14]. We thus complement the combination of trust management and ontologies with a reputation system, and propose a metric for computing similarity between two concepts based on reputation. We believe that the solution proposed in this paper can be applied to other TM systems.

The paper is organized as follows. Next section presents an overview of POLIPO and discusses its limitations. Section III introduces the notion of similarity and shows how it can be used in POLIPO. Section IV defines a metric for computing reputation from similarity statements. Section V defines a reputation-based metric for assessing the similarity between two attributes. Section VI describes a set of experiments for evaluating the proposed metrics. Finally, Section VII discusses related work and Section VIII draws final remarks.

II. The POLIPO framework

POLIPO (Policies & OntoLogies for Interoperability, Portability, and autOnomy) [6] is a security framework that combines TM with ontologies to enable interoperability, portability, and autonomy in dynamic coalitions of heterogeneous systems. In this section, we give an overview of POLIPO and discuss its weaknesses. To make the discussion more concrete, we introduce a scenario in the MSS domain that is used as a running example through the paper.

A. Motivating Example

This section introduces a scenario in the MSS domain which is analyzed in the POSEIDON project[1] to investigate interoperability and dynamism in Systems of Systems (SoS) [15]. As many other domains involving SoS (e.g., air traffic control and Internet), the MSS domain is characterized by the interaction and collaboration between autonomous and heterogeneous systems sharing data, processes, and resources.

In the POSEIDON project, it is assumed that navies (e.g., White Star Navy, Blue Star Navy, Grey Cross Navy, etc.) use a local ontology that formally represents the domain in terms of concepts and relationships, each with a precise semantics. For example, the ontology of each party might include a chain of command (Fig. 1). Different navies can adopt different chains of command. This affects the interoperability between parties because a navy may not know the ranks of other parties or, even worse, different navies may use the same term to indicate different ranks. This is, for instance, the case of rank "lieutenant" that has a different position in the White Star (WS) Navy and Grey Cross (GC) Navy's chains.

In addition to the local ontology, navies joining a coalition refer to a shared ontology (hereafter POSEIDON ontology) to communicate with the other parties. To assist collaborating navies in the comparison of rankings, the POSEIDON ontology defines rank codes (Fig. 1). These codes are used to determine the seniority of an officer. POSEIDON ranks go from OF-1 up to OF-6, where OF-6 is the highest rank and OF-1 the least rank.

When navies form a coalition, an initial assignment of responsibilities is necessary before they become operative. This initial process includes the identification of an entity at each party who is responsible for the fulfillment of assigned duties. In practice, this entity is the authority with the highest rank within the navy. Therefore, the assignment of responsibilities requires a prior alignment between the high level ranks of navies. Conversely, aligning lower ranks has less priority and can be realized in a more flexible way.

B. POLIPO Rules and Policies

POLIPO rules are specified using four types of constructs:

- *ontology atoms*: are used to query ontology concepts and relationships. Every ontology concept is identified by a *conceptURI*, and every relationship by a *relationshipURI*.
- *credential atoms*: represent digitally signed statements made by an issuer about an attribute of a subject. Issuer and subject are identified by a unique name (e.g., a public key) and the attribute is a *conceptURI*. Additional restrictions on the credential atom and on the attribute can be specified as a list of $(property, value)$ pairs, where $property$ is a *relationshipURI* and $value$ is a value of the appropriate type.
- *authorization atoms*: denote the permission of a subject to perform an action on an object. The subject is identified by its unique name, action and object are *conceptURI*s.
- *constraints*: are specified using Constraint Logic Programming (CLP) constraints (e.g., =, $>$, $<$, etc.) or user-defined predicates.

POLIPO rules are Horn clauses of the form $h \leftarrow b_1, \ldots, b_n$, where h, called head, is an atom, and $b_1, \ldots, b_n$, called body, are literals (i.e., positive or negative atoms, where the use of negative atoms is subject to restrictions), with $n \geq 0$. We distinguish three types of rules: *credential release rules*, *authorization rules*, and *predicate definition rules*. Intuitively, the type of rule depends on the type of atom occurring in the head of the rule. We refer to [6] for a complete description of POLIPO rules and present examples of those rules in Fig. 2.

POLIPO policies consist of sets of POLIPO rules. In particular, we distinguish two types of policies: *credential release*

[1] http://www.esi.nl/poseidon/

Credential Release Rule
$\mathsf{cred}(WS, psd{:}OF\text{-}2, X, []) \leftarrow ws{:}Lieutenant(X)$

Authorization Rule
$\mathsf{perm}(psd{:}read, X, Y) \leftarrow aboutSurveillance(Y),$
$\mathsf{cred}(WS, psd{:}OF\text{-}2, X, [[(\text{`}psd{:}ValidUntil\text{'}, Z)]]),$
$Z \geq CurrentTime()$

Predicate Definition Rule
$aboutSurveillance(X) \leftarrow ws{:}aboutMission(X, \text{`}Surveillance\text{'}),$
$ws{:}sensitivityLevel(X, Y), Y < 3$

Fig. 2. Examples of POLIPO Rules

1 $\mathsf{cred}(WS, ws{:}Lieutenant, X, []) \leftarrow ws{:}Lieutenant(X)$
2 $\mathsf{perm}(psd{:}read, X, Y) \leftarrow aboutSurveillance(Y),$
$\mathsf{cred}(WS, ws{:}Lieutenant, X, [])$

Fig. 3. POLIPO Policies Using Local Attributes

policies and *authorization policies*. A credential release policy is a set of credential release rules, and an authorization policy is a set of authorization rules. Predicate definition rules complement credential release and authorization policies with local knowledge.

C. Global vs Local Credential Attributes

In [6] we require attributes and properties occurring in credential atoms to be defined in the global ontology. However, this solution has a number of disadvantages. First, concepts and relationships in local ontologies may have no exact correspondent in the global ontology. Therefore, peers may have to issue credentials which are approximations of the intended attribute of the subject. This may not only discourage a peer (for accountability reasons) to release credentials about attributes he is not confident of, but may also lead peers to grant unwanted authorizations. In addition, the certificates issued by peers before joining a coalition cannot be used when interacting with other parties. A solution to these issues is to allow peers to specify policies in terms of their local ontology.

Example 1: The rules in Fig. 3 express the rules in Fig. 2 in terms of the White Star Navy's ontology. Those policies, however, cannot be understood by navies that have no knowledge of the White Star Navy's ontology.

The example above shows that the use of local concepts in policies can affect the interoperability among peers. We tackle this problem by employing the concept of similarity for ontology alignment as discussed in the next section.

III. Similarity for Ontology Alignment

The concept of similarity has been proposed to address the problem of ontology alignment. It expresses the degree of semantic resemblance between two ontology concepts [7], [8], [9]. In coalitions where parties are not assumed to have full knowledge about the local ontology of other parties, similarity between two concepts corresponds to an opinion rather than to a precise measure.

We assume that similarity is symmetric: if a peer states that att_1 is similar to att_2 to a certain degree, it implies that also att_2 is similar to att_1 to the same degree. On the other hand, we make no assumptions on transitivity: if att_1 is similar to att_2 to a certain degree and att_2 is similar to att_3 to a certain degree, we cannot infer any information about the similarity between att_1 and att_3. Adopting a transitive notion of similarity would require knowledge about the ontology of other peers together with the criteria used by peers to express a certain judgment, otherwise it can lead to erroneous similarity values.

To ensure the integrity of similarity statements, they are asserted in form of credentials and are represented by *similarity credential atoms*.

Definition 1 (Similarity Credential Atom): A similarity credential atom is an atom of the form

$$\mathsf{sim}(issuer, att_1, att_2, [p_1 \dots p_n])$$

where $issuer$ is the unique name of the entity signing the similarity credential; att_1 and att_2 are $conceptURI$s; and $[p_1 \dots p_n]$ is a list of properties. Each p_i is a pair $(property, value)$, where $property$ is a $relationshipURI$ related to the similarity credential concept[2] and $value$ bounds the range of the relationship.

Intuitively, a similarity credential atom specifies the similarity between two concepts according to the interpretation of its issuer. Similarity credentials include property $psd{:}timestamp$ and possibly property $psd{:}degree$. Property $psd{:}timestamp$ specifies the time at which the similarity credential has been issued. Property $psd{:}degree$ represents the degree of similarity between the two concepts and its value is in the range $[0 \dots 1]$. If this property is not specified, it assumes default value 1. Examples of similarity credential atoms are presented in Fig. 4.

Example 2: In the first similarity credential of Fig. 4 the White Star Navy states that the degree of similarity between concept $Lieutenant$ defined in its ontology and concept $OF\text{-}2$ defined in the POSEIDON ontology is 0.9. This similarity credential has been issued on 2009/01/15.

Notice that peers can give an opinion about the similarity between two concepts independently from the ontology in which those concepts have been defined (e.g., see the fifth credential in Fig. 4 where GS denotes the Green Star Navy).

Every peer stores the similarity credentials issued by itself and acquired during the interaction with other peers in a *similarity credential repository* (SCR). A peer's repository contains only one credential issued by a certain peer for each pair of concepts. Upon receiving a new similarity credential, the peer checks in its SCR if there is already a credential about the same pair of concepts issued by the same peer. If this is the case, it only keeps the certificate with the most recent timestamp. Peers can enlarge their repository by exchanging similarity credentials with their acquaintances using any gossip protocol.

Example 3: Let the third credential of Fig. 4, be in the SCR of the White Star Navy. During the interaction with some party, the navy receives the fourth credential of Fig. 4. Both credentials are issued by the Grey Cross Navy (GC) and specify the similarity between the same pair of concepts. The

[2] sim is formally defined in the POSEIDON ontology.

1	$\mathsf{sim}(WS, ws{:}Lieutenant, psd{:}OF\text{-}2, [(psd{:}degree, 0.9), (psd{:}timestamp, 2009/01/15)])$
2	$\mathsf{sim}(BS, ws{:}Lieutenant, psd{:}OF\text{-}2, [(psd{:}degree, 0.7), (psd{:}timestamp, 2009/01/20)])$
3	$\mathsf{sim}(GC, gc{:}Lieutenant, psd{:}OF\text{-}1, [(psd{:}timestamp, 2008/01/01)])$
4	$\mathsf{sim}(GC, gc{:}Lieutenant, psd{:}OF\text{-}1, [(psd{:}degree, 0.8), (psd{:}timestamp, 2009/02/10)])$
5	$\mathsf{sim}(GS, ws{:}Lieutenant, gc{:}Lieutenant, [(psd{:}degree, 0.9), (psd{:}timestamp, 2009/03/01)])$
6	$\mathsf{sim}(BS, gc{:}Lieutenant, ws{:}Lieutenant, [(psd{:}degree, 0.5), (psd{:}timestamp, 2009/03/10)])$
7	$\mathsf{sim}(WS, gc{:}Lieutenant, psd{:}OF\text{-}1, [(psd{:}degree, 0.8), (psd{:}timestamp, 2009/03/20)])$
8	$\mathsf{sim}(GC, gc{:}Lieutenant, ws{:}Lieutenant, [(psd{:}degree, 0.6), (psd{:}timestamp, 2009/04/10)])$

Fig. 4. Similarity Credentials

$$\mathsf{perm}(psd{:}read, X, MI) \leftarrow \mathsf{cred}(WS, psd{:}Ally, Y, []), \mathsf{cred}(Y, Z, X, []), \mathsf{similar}(Z, ws{:}Lieutenant) \geq 0.6$$

Fig. 5. POLIPO Policies Using Similarity Constraints

White Star Navy replaces the old similarity credentials with the new one because of a more fresh timestamp.

As several peers can express an opinion about the similarity between two concepts, these opinions should be combined to estimate the actual similarity between those concepts. An intuitive and simple way in which a peer can estimate it is to calculate an average similarity degree. For the sake of compactness, hereafter we use the notation $\mathsf{sim}(i, att_1, att_2, d)$ to indicate the credential issued by i stating that att_1 and att_2 have degree of similarity d.

Definition 2 (Simple Attribute Similarity): Let att_1 and att_2 be two $conceptURI$s, and $S = \{\mathsf{sim}(i, att_1, att_2, d) \mid \mathsf{sim}(i, att_1, att_2, d) \in SCR\}$. The simple attribute similarity between att_1 and att_2, denoted by $\mathsf{similar}(att_1, att_2)$, is defined as

$$\mathsf{similar}(att_1, att_2) = \frac{\sum_S d}{|S|} \quad (1)$$

where $|S|$ denotes the cardinality of S.

Example 4: Let the similarity credentials in Fig. 4 (except for the third one) be part of the *SCR* of the White Star Navy. The White Star Navy can assess that the similarity between $ws{:}Lieutenant$ and $gc{:}Lieutenant$ is 0.67.

Peers can use similarity to increase their autonomy while maintaining interoperability among them. A peer might accept a credential whose attribute is similar to a given attribute for at least a certain degree. For this purpose, we define the similarity constraint

$$\mathsf{similar}(att_1, att_2) \geq threshold$$

where $threshold$ is the minimum degree of similarity that there must be between attributes att_1 and att_2. This constraint can be used as any other constraint in POLIPO rules.

Example 5: The White Star Navy authorizes access to mission instruction (MI) to the officers of allied navies who provide a certificate whose attribute is at least similar for 0.6 to $ws{:}Lieutenant$ (Fig. 5). According to Example 4, the White Star Navy allows users presenting a certificate of lieutenant as defined in the Grey Cross Navy's ontology to read MI.

IV. A Reputation Metric Based on Agreement

In the previous section, we have introduced similarity credentials to represent the opinion of a peer about the similarity between two concepts. Not all peers, however, are equally trustworthy. Therefore, similarity statements should be discriminated according to the trustworthiness of their issuer, giving more weight to the opinion of more trustworthy peers. We address this concern by complementing POLIPO with a reputation system. In this section, we present a metric for computing the reputation of peers.

The reputation of a peer reflects the accuracy of its similarity statements. It is defined on the basis of its *agreement* with the other peers, which depends on the *affinity* between their statements. Intuitively, the higher is the affinity among the statements, the higher the agreement between their issuers.

Affinity measures the level of correspondence between two non-contradicting statements. We introduce a local *similarity threshold* st to allow peers to decide when two similarity statements are contradicting. The value of st is in the range $[0, 1)$.

Definition 3 (Affinity): Let $s_1 = \mathsf{sim}(i, a, b, d_1)$ and $s_2 = \mathsf{sim}(j, a, b, d_2)$ be two similarity credentials issued by peers i and j and specifying the degree of similarity between attributes a and b. Let st be the similarity threshold. The affinity between s_1 and s_2, denoted by $aff(s_1, s_2)$, is defined as

$$aff(s_1, s_2) = \begin{cases} \dfrac{1 - |d_1 - d_2| - st}{1 - st} & \text{if } 1 - |d_1 - d_2| \geq st \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The agreement between two peers depends on the affinity between the similarity credentials they issued.

Definition 4 (Agreement): Let i and j be two peers, S_i and S_j be the sets of similarity credentials issued by i and j respectively, and $S = \{\langle \mathsf{sim}(i, a, b, d_i), \mathsf{sim}(j, a, b, d_j)\rangle | \mathsf{sim}(i, a, b, d_i) \in S_i \wedge \mathsf{sim}(j, a, b, d_j) \in S_j\}$. The agreement between i and j, denoted by $agree(i, j)$, is defined as

$$agree(i, j) = \begin{cases} \dfrac{\sum_{\langle s_1, s_2\rangle \in S} aff(s_1, s_2)}{|S|} & \text{if } |S| \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Every party uses the similarity credentials stored in its *SCR* to determine the agreement among all parties in the coalition. Agreement values are represented as a matrix A, called *agreement matrix*, whose elements are calculated using Equation 3 (i.e., $a_{ij} = agree(i, j)$) and are updated every time a new similarity credential is acquired.

Example 6: Let the similarity credentials in Fig. 4 be the similarity credentials in the White Star Navy's repository and $st = 0.6$. The agreement matrix A of the White Star Navy is

$$A = \begin{bmatrix} 1 & 0.5 & 1 & 0 \\ 0.5 & 1 & 0.75 & 0 \\ 1 & 0.75 & 1 & 0.25 \\ 0 & 0 & 0.25 & 1 \end{bmatrix}$$

where the order of navies in A is WS, BS, GC, and GS.

It is worth noting that the agreement matrix is symmetric, $a_{ij} \in [0,1]$, and the elements in the diagonal are equal to 1 (i.e., a peer always agrees with itself).

The agreement among peers is used to determine the reputation that peers have about each other.

Definition 5 (Reputation): The reputation of a peer j in the view of a peer i, denoted by r_{ij}, is the measure of accuracy of the similarity statements of j according to the agreement matrix of i.

A simple way for a party i to compute the reputation of another party k would be to consider the agreement between itself and k (i.e., a_{ik}). This "simple" approach is, however, imprecise, because it only considers the opinions of i, which might be inaccurate; moreover, it works only if i and k issued a significant number of credentials about the same pairs of concepts. Therefore, i also uses information concerning the agreement between k and the other parties. The agreement between k and other peers is weighed by their reputation in the view of i.

$$r_{ik} = \sum_j \frac{r_{ij}}{\sum_l r_{il}} a_{jk} \tag{4}$$

where r_{ik} represents the reputation of party k in the view of party i, based on the agreement between k and every acquaintance j of i. Notice that the reputation value r_{ij} may be different from the agreement value a_{ij}. This allows i to give more weight to the statements of some parties (e.g., vocabulary experts). The initial value of r_{ij} can be any value in the range $[0,1]$. Reputation values are normalized with respect to the sum of all reputation values of i (i.e., $\sum_l r_{il}$) to guarantee r_{ik} to be in the interval $[0,1]$. If we define $\vec{r}_i$ to be the vector containing the values r_{ik} for every party k, Equation 4 can be rewritten as follows:

$$\vec{r}_i = A \frac{\vec{r}_i}{\|\vec{r}_i\|_1} \tag{5}$$

where $\|\vec{r}_i\|_1$ denotes the norm 1 of vector $\vec{r}_i$. Notice that Equation 5 should contain the transpose of A, which is, however, equal to A, because of its symmetry.

To find the reputation of all peers in the view of party i, we iterate Equation 5. The reputation vector $\vec{r}_i$ will converge after t iterations (under the assumption that A is irreducible and aperiodic).

$$\vec{r}_i^{(t)} = A \frac{r_i^{(t-1)}}{\|r_i^{(t-1)}\|_1} \tag{6}$$

A peer may want to strongly bias its computation on the initial reputation values, because, for instance, it firmly believes in vocabulary experts' judgment. For this purpose, we generalize Equation 6 as follows:

$$\vec{r}_i^{(t)} = (1-\alpha) \frac{A\vec{r}_i^{(t-1)}}{\|\vec{r}_i^{(t-1)}\|_1} + \alpha \vec{p}_i \tag{7}$$

where α is a constant in the range $[0,1]$ and $\vec{p}_i$ is the initial reputation vector of party i. Besides influencing reputation values, this equation also guarantees the convergence of the method by making matrix A irreducible and aperiodic [11]. We discuss the convergence of the method in Section VI-A.

$$\mathsf{perm}(psd{:}read, X, MI) \leftarrow \mathsf{cred}(WS, psd{:}Ally, Y, []), \mathsf{cred}(Y, Z, X, []), \mathsf{similar}(0.5, Z, ws{:}Lieutenant) \geq 0.6$$

Fig. 6. POLIPO Rule Using Reputation-Based Similarity Constraints

Example 7: Let the matrix in Example 6 be the agreement matrix of the White Star Navy, $\vec{p} = [\ 1 \quad 0 \quad 0 \quad 0\ ]^T$ be its initial reputation vector (where ^{T} indicate the transpose of a vector), and $\alpha = 0$. The reputation vector $\vec{r}$ of the navy converges to $[\ 0.81 \quad 0.70 \quad 0.89 \quad 0.14\ ]^T$.

V. Reputation-Based Attribute Similarity

In this section, we combine the results of Sections III and IV by defining a metric for computing attribute similarity based on reputation. The similarity metric in Equation 1 gives the same weight to all similarity credentials independently from the trustworthiness of their issuer. To alleviate this problem, we revise Equation 1 by taking into account the reputation of issuers. In addition, we introduce a local *reputation threshold* rt that allows parties to filter out credentials issued by parties with low reputation. In this way, we can be more confident that the computed similarity is accurate. The value of rt is in the range $(0,1]$.

Definition 6 (Weighted Attribute Similarity): Let i be the party in the coalition assessing similarity between concepts a and b, rt be its reputation threshold, and $S = \{\mathsf{sim}(j,a,b,d) \,|\mathsf{sim}(j,a,b,d) \in SCR_i \wedge r_{ij} \geq rt\}$ with r_{ij} being the reputation of party j in the view of i. The similarity between a and b according to i's perspective is

$$\mathsf{similar}(rt,a,b) = \sum_S \frac{r_{ij}}{\sum_S r_{ij}} d \tag{8}$$

Example 8: Fig. 6 shows how to specify rules using Equation 8 to compute the similarity of an attribute Z with attribute $ws{:}Lieutenant$. The reputation threshold is set to 0.5. This implies that the similarity credentials about attribute Z and $ws{:}Lieutenant$ issued by parties with reputation lower than 0.5 are not considered in the computation of attribute similarity. If we consider the credentials in Fig. 4, the similarity between $ws{:}Lieutenant$ and $gc{:}Lieutenant$ is 0.56. In contrast to Example 5, a certificate stating that an officer is a lieutenant with respect to the Black Star Navy's chain of command is now not valid to access MI.

It is worth noting that Equation 8 is equivalent to Equation 1 if all parties have the same reputation (and rt is lower or equal to this value).

VI. Experiments

This section presents two experiments: the first investigates the convergence of the reputation metric; the second shows the correlation between reputation and similarity credential accuracy.

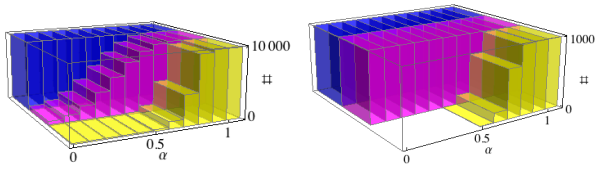

(a) Converging 20×20 matrices (b) Converging100×100 matrices

Fig. 7. Influence of α on convergence

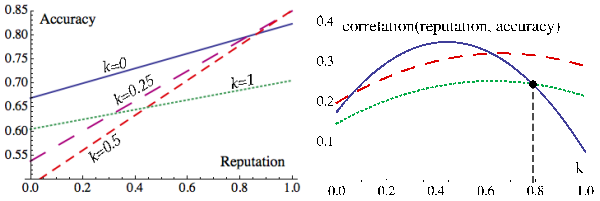

(a) Regression Lines (b) Correlation Coefficients

Fig. 8. Influence of Reputation

A. Influence of α on the Speed of Convergence

We evaluated the influence of α on the number of iterations Equation 6 takes to converge to a stable reputation vector, given a randomly generated agreement matrix. We made 10000 experiments on 20×20 matrices and 1000 on 100×100 matrices. To account for the relatively low number of matching similarity credentials we expect in real situations, we generated matrices with 80% sparseness. We set a limit to the number of iterations first to 10, then to 20 and 100. We have convergence in t steps when $\|\vec{r}^{(t)} - \vec{r}^{(t+1)}\|_2 < 0.001$. We started with $\alpha = 0$, incrementing it at each step by 0.1.

Fig. 7 shows two combined histograms in which the number of successful convergences is on the y-axis, and the values of α on the x-axis. Fig. 7(a) shows the results for the 20×20 matrices; the front row corresponds to the lowest iteration limit (10) and the back row to highest one (100). Reputation values start converging in 10 steps with $\alpha = 0.6$, in 20 step with $\alpha = 0$. Most experiments converge in 10 steps with $\alpha = 0.8$, in 20 steps with $\alpha = 0.7$, and in 100 steps with $\alpha = 0$. Fig. 7(b) shows the results for the 100×100 matrices. As before, converging in 10 steps starts with $\alpha = 0.6$, but here all experiments converge within 20 steps, even with $\alpha = 0$.

B. Correlation between Reputation and Credential Accuracy

This experiment aims to show the correlation between the reputation of an issuer and the accuracy of its similarity credentials. We set an objective similarity degree $d_0(a, b)$ for all pairs of concepts (a, b) and define the accuracy of a credential $\mathsf{sim}(i, a, b, d_i)$ based on its distance from d_0 (i.e., $1 - |d_0(a, b) - d_i|$). We consider a variable rate k of unreliable peers, with $0 \leq k \leq 1$. For unreliable issuers, we generated d_i randomly with a uniform distribution. For reliable issuers, we generated d_i with a triangular distribution centered in $d_0(a, b)$.

Fig. 8(a) shows the results for $k \in \{0, 0.25, 0.5, 1\}$ in the form of regression lines with issuer's reputation on the x-axis and credential accuracy on the y-axis. With $k = 0$ all issuers are reliable and the accuracy is maximal (full line). When $k = 1$ all issuers are unreliable and the accuracy is much lower (dotted line). For $k \in \{0.25, 0.5\}$ the regression lines (dashed) are steeper, indicating that reputation strongly reflects credential accuracy in situations where both reliable and unreliable peers issue similarity credentials.

In Fig. 8(b) the full line plots the strength of the correlation between reputation and credential accuracy for the different values of k. The figure shows that correlation is lowest for extreme values of k. If all peers are (un-)reliable, all similarity credentials are (in-)accurate and the influence of reputation is less pronounced. The strong correlation for moderate values of k shows that reputation effectively helps to distinguish accurate from inaccurate similarity statements.

The dotted line in Fig. 8(b) plots the correlation between reputation and accuracy when using the simple attribute similarity metric in Equation 1 (which does not consider reputation) to estimate the objective similarity degree d_0. The correlation here stems only from the way in which reputation values are calculated. The dotted line is therefore a lower bound on the relevance of correlation. When k reaches 0.79, the full and the dotted curves cross and the reputation of an issuer no longer reflects the accuracy of its statements.

The dashed line in Fig. 8(b) plots the correlation between reputation and accuracy when using the weighted attribute similarity metric (Equation 8, with $rt = 0$) to estimate the objective similarity degree d_0. The influence of reputation is significant as the correlation is stronger than the previous one for every value of k.

VII. Related Work

Ontologies are increasingly seen as a key factor for enabling interoperability across heterogeneous systems [5]. The distributed nature of ontology development has led to different ontologies for the same or overlapping domains. Therefore, ontology alignment is required for combining those ontologies. Several ontology alignment systems and algorithms have been proposed (e.g., [16], [17], [18], [19]), exploiting various types of information (e.g., element names, structural properties, characteristics of data instances).

Simply identifying the mapping between concepts from different ontologies does not provide sufficient information about their relationship. To achieve this, the mapping has to be complemented by a degree of semantic resemblance. Metrics for measuring the degree of semantics similarity between two concepts are presented in [19], [7], [8], [9]. In [7] the similarity degree of two concepts depends on the number of identical relationships to which they take part. Ngan et al. [8] measure similarity on the basis of a combination of syntactic properties, neighborhood, context, and equivalent similarity. Nguyen et al. [9] introduce a similarity measure based on ontology structure and corpus-based features. The application of presented systems and metrics for ontology alignment to our scenario would require a complete knowledge of the local

ontology of other parties. This, however, conflicts with our assumptions.

Due to our partial knowledge assumption, we consider similarity statements as opinions rather than absolute values. To obtain a reliable approximation of absolute similarity values, we combine different opinions, weighing them by the trustworthiness of their issuer. Reputation systems [12], [14], [10], [13], [11] provide a way to assess the trustworthiness of unknown entities in a distributed system. They are usually based on a transitive notion of trust, i.e. they compute the reputation of an unknown entity through the opinion (reputation) that known peers have of it. Our reputation system does not consider the opinion of other parties directly, but is based on local observations (i.e., similarity credentials in the local SCR). The aim of Eigentrust [11] and PeerTrust [14] is to compute global reputation values for each peer. From this perspective, our work is similar to Credence [13], in which reputation is a personalized estimation rather than a universal value. In PeerTrust [14] and TRUMMAR [10] reputation is time-dependent, i.e. most recent transactions are weighted more than older ones. Also our system is time-dependent as we only consider the most recent similarity credential issued by a peer for a certain pair of concepts. However, differently from those systems, we do not weigh statements with respect to their timestamp. Our reputation metric is similar to the one presented in Eigentrust [11]; however, we do not normalize the agreement matrix to be able to use an absolute reputation threshold when assessing the similarity of two concepts. Similarly to [13], [14], our system uses public key cryptography to guarantee integrity of information. Therefore, it is not vulnerable to many threats typical of P2P systems that other reputation mechanisms have to face [20].

VIII. Conclusions

In this paper, we have addressed the problem of semantic alignment in distributed access control for dynamic coalitions of heterogeneous systems. The novelty of our proposal is represented by a technique for accurate ontology alignment when peers have limited knowledge of the local ontology of other peers. Using credentials to express opinions about attribute similarity and weighing such statements with respect to the trustworthiness of their issuer make it possible to approach ontology alignment in a flexible way, guaranteeing interoperability among peers without undermining their autonomy.

The introduction of reputation, however, may introduce risks such as the granting of illegitimate access. However, we have good reasons to believe that this potential problem is limited in domains in which access control decisions are based on credentials (e.g., MSS, B2B, etc.). Contrarily to most P2P systems, the peers in such domains are not anonymous and the credentials they issue support accountability and non-repudiation. The strong authentication and identification demanded by the MSS domain, for instance, is an effective measure against Sybil attacks and related threats that plague many P2P reputation systems [20], [11], [13], [14]. Future work has to investigate the threats that may affect the proposed approach and devise strategies that protect against it.

The experiments presented in this paper focus mainly on the accuracy of the proposed metrics. Other experiments will be set up to provide insight into the practical scalability of the similarity calculation and to analyze the robustness of our approach against ontology misalignment attacks.

Acknowledgments This work has been carried out as part of the POSEIDON project under the responsibility of the Embedded Systems Institute (ESI). This project is partially supported by the Dutch Ministry of Economic Affairs under the BSIK03021 program. This work has been also funded by the EU TAS3 project.

References

[1] M. Y. Becker and P. Sewell, "Cassandra: Distributed access control policies with tunable expressiveness," in *Proc. of POLICY'04*. IEEE Computer Society, pp. 159–168.

[2] M. Czenko and S. Etalle, "Core TuLiP Logic Programming for Trust Management," in *Proc. of ICLP'07*, ser. LNCS 4670. Springer, 2007, pp. 380–394.

[3] N. Li, B. N. Grosof, and J. Feigenbaum, "Delegation logic: A logic-based approach to distributed authorization," *ACM Trans. Inf. Syst. Secur.*, vol. 6, no. 1, pp. 128–171, 2003.

[4] N. Li, J. C. Mitchell, and W. H. Winsborough, "Design of a Role-Based Trust-Management Framework," in *Proc. of Symp. on Security and Privacy*. IEEE Computer Society, 2002, pp. 114–130.

[5] N. Choi, I.-Y. Song, and H. Han, "A survey on ontology mapping," *SIGMOD Rec.*, vol. 35, no. 3, pp. 34–41, 2006.

[6] D. Trivellato, F. Spiessens, N. Zannone, and S. Etalle, "POLIPO: Policies & OntoLogies for Interoperability, Portability, and autOnomy," in *Proc. of POLICY'09*. IEEE Computer Society, 2009.

[7] R. Culmone, G. Rossi, and E. Merelli, "An Ontology Similarity Algorithm for Bioagent," in *Proc. of NETTAB'02*, 2002.

[8] L. D. Ngan, T. M. Hang, and A. E. S. Goh, "Semantic Similarity between Concepts from Different OWL Ontologies," in *Proc. of IEEE International Conference on Industrial Informatics*, 2006, pp. 618–623.

[9] H. Nguyen and H. Al-Mubaid, "A Combination-based Semantic Similarity Measure using Multiple Information Sources," in *Proc. of IRI'06*. IEEE Press, 2006, pp. 617–621.

[10] G. Derbas, A. Kayssi, H. Artail, and A. Chehab, "TRUMMAR - A Trust Model for Mobile Agent Systems Based on Reputation," in *Proc. of ICPS'04*. IEEE Press, 2004, pp. 113–120.

[11] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *Proc. of WWW'03*. ACM, 2003, pp. 640–651.

[12] P. Michiardi and R. Molva, "Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks," in *Proc. of the Conf. on Comm. and Multimedia Security*, 2002, pp. 107–121.

[13] K. Walsh and E. G. Sirer, "Experience with an object reputation system for peer-to-peer filesharing," in *Proc. of NSDI'06*. USENIX Association, 2006, pp. 1–1.

[14] L. Xiong and L. Liu, "PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities," *TKDE*, vol. 16, no. 7, pp. 843–857, 2004.

[15] M. W. Maier, "Architecting principles for systems-of-systems," *Systems Engineering*, vol. 1, no. 4, pp. 267–284, 1999.

[16] P. Bouquet, L. Serafini, and S. Zanobini, "Semantic coordination: a new approach and an application," in *Proc. of ISWO'03*, 2003.

[17] A. Doan, J. Madhavan, R. Dhamankar, P. Domingos, and A. Halevy, "Learning to match ontologies on the semantic web," *The VLDB Journal*, vol. 12, no. 4, pp. 303–319, 2003.

[18] P. Mitra and G. Wiederhold, "Resolving terminological heterogeneity in ontologies," in *Proc. of OSI*, 2002.

[19] H.-H. Do and E. Rahm, "COMA: a system for flexible combination of schema matching approaches," in *Proc. of VLDB'02*, 2002, pp. 610–621.

[20] K. Hoffman, D. Zage, and C. Nita-Rotaru, "A survey of attack and defense techniques for reputation systems," Purdue University, Tech. Rep. CSD TR 07-013, 2007.