

Extending Virtual Organizations to improve trust mechanisms *

(Short Paper)

Ramón Hermoso, Roberto Centeno, Holger Billhardt and Sascha Ossowski

University Rey Juan Carlos

Tulipán s/n, 28933

Móstoles (Madrid), Spain

{ramon.hermoso, roberto.centeno, holger.billhardt, sascha.ossowski}@urjc.es

ABSTRACT

Virtual Organizations (VOs) are becoming an increasingly important research topic in the field of Multi-Agent Systems (MAS). The problem of selecting suitable counterparts to interact with is of particular relevance for agents belonging to a VO. This issue has been extensively investigated, applying probability or cognitive approaches, but very few focus has been given to the use of internal organizational structures and the improvement they can provide. In this paper we analyze how organizational structures can support the agent selection process based on trust mechanisms. Furthermore, we present a way to extend VOs automatically (e.g., their role taxonomies) by detecting and identifying new roles. We show that such extensions lead to an improvement of the agents' decisions when employing trust mechanisms that take advantage of organizational structures.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence - Multiagent Systems

General Terms

Design

Keywords

Social and organizational structure, trust and reputation

1. INTRODUCTION

The concept of *organization* has become central in Multi-Agent Systems (MAS) research. It is extensively accepted that organizational structures and properties can provide significant advantages when developing agent-based software,

*The present work has been partially funded by the Spanish Ministry of Education and Science under projects TIC2003-08763-C02, TIN2006-14630-C03-02 (FPI grants program) and by the Spanish project "Agreement Technologies" (CONSOLIDER CSD2007-0022, INGENIO 2010)

Cite as: Extending Virtual Organizations to improve trust mechanisms (Short Paper), Ramón Hermoso, Roberto Centeno, Holger Billhardt and Sascha Ossowski, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. 1489-1472.

Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

since they allow building up more complex MAS designs using a reduced set of simple abstractions [5]. Such organizational abstractions often endow the system with some sort of structure that shapes the agents' behavior.

Agents that join an organization usually have to deal with some constraints, such as the need to play particular *roles* in order to participate in certain allowed *interactions*. These notions can be complemented by higher-level abstraction such as sets of *norms* that intend to keep agents from unexpected or undesirable behavior. MAS that have been built based on those types of organizational structures are frequently called *Virtual Organizations* (VOs) [8].

VOs impose limits on the actions the agents can perform, but still allows for a certain autonomy that endows agents with the freedom (within the framework of the organization) *i) to choose the actions to do next* and *ii) to select with whom to perform them*. Regarding the latter, trust and reputation mechanisms are a common way for agents to decide with whom to interact [4, 9, 10, 1].

In this paper, we argue that the organizational structure itself can be used by agents as an additional means to guide their trust-related decision making processes. In particular, virtual organizations usually define roles that agents can play and interactions that agents can participate in, and often describe those concepts in taxonomies. Such taxonomies allow categorizing roles and interactions and, thus, provide a notion of similarity among them, which can be used to infer the trustworthiness of other agents in performing particular interactions. Furthermore, the concept of role can be seen as behavioral patterns that classify agents with respect to their capacities regarding certain interactions. Agents can use the information about the roles another agent plays to determine whether or not it is a good counterpart for a given interaction. Based on this idea, we propose a mechanism, using clustering methods, to extend a given organization by detecting and introducing new roles that can be seen as specializations of existing roles and, thus, extend the role taxonomy. The advantage of more fine grained role taxonomies lies in an improvement of the trust related decisions.

The remainder of this paper is organized as follows: Section 2 outlines a trust model that takes advantage of role and interaction taxonomies. Section 3 presents our approach to extend a given VO by identifying new, more specialized roles. In Section 4 we present an experimental evaluation of the proposed approaches, and Section 5 summarizes our conclusions and points to future lines of work.

2. A CONFIDENCE-BASED TRUST MODEL FOR VOs

In this section we introduce a trust model that takes advantage of the organizational structures of a VO. In particular, we set out from:

- A set \mathcal{I} of *interactions* available in the VO.
- A set \mathcal{R} of *roles* are involved in these interactions.
- A partial order \leq that structures both roles and interactions in taxonomies.

The trust model is based on the assumption that *agents tend to behave similarly when enacting similar roles in similar interactions*. Using this assumption, an agent is able to assess (to a certain extent) the future behavior of another agent in a certain situation by considering its past behavior in “similar situations”. That is, an agent can infer trustworthiness, even if it has, first, no direct past experience about a specific situation and, second, it can not collect opinions from other agents either because the opinions from others are unreliable (“liars”) or none of the agents has enough proper experience. So, even though we adhere to the standard notion of *reputation* and *confidence* used in literature [4, 6], in this paper we will exploit organizational information with regards to the latter.

In our model, a confidence value $c_{A \rightarrow \langle B, R, I \rangle}$ is built up from A ’s past interactions with agent B playing role R and performing interactions of type I . We call LIT – *Local Interaction Table* – the agent’s data structure dedicated to store confidence values for past interactions with any counterpart the agent has interacted with. Each entry corresponds to a *situation*: an *agent* playing a specific *role* in a particular *interaction*. LIT_A denotes agent A ’s LIT. Each entry in a LIT consists of:

- the Agent/Role/Interaction identifier $\langle X, Y, Z \rangle$,
- the confidence value for the situation ($c_{A \rightarrow \langle X, Y, Z \rangle} \in [0..1]$),
- a reliability value ($r_{A \rightarrow \langle X, Y, Z \rangle}$) that measures how certain an agent is about its own confidence in situation $\langle X, Y, Z \rangle$.

If an agent participates in an interaction I with agent B playing role R , the corresponding entry in its LIT will be updated as follows:

- **confidence value:** let $g_{\langle X, Y, Z \rangle} \in [0..1]$ denote the evaluation value an agent A calculates for a particular experience with the agent X playing role Y in the interaction of type Z . In our work, we use the following equation to update confidence:

$$c_{A \rightarrow \langle X, Y, Z \rangle} = \epsilon \cdot c'_{A \rightarrow \langle X, Y, Z \rangle} + (1 - \epsilon) \cdot g_{\langle X, Y, Z \rangle}, \quad (1)$$

where $c'_{A \rightarrow \langle X, Y, Z \rangle}$ is the confidence value in A ’s LIT before the interaction is performed and $\epsilon \in [0..1]$ is a parameter specifying the importance given to A ’s past confidence value.

- **reliability value:** we calculate reliability by using the approach proposed by Huynh, Jennings and Shadbolt [4]. This approach takes into account the number of interactions a confidence value is based on, and the variability of the individual values across past experiences.

The counterpart selection of an agent A in an interaction I that requires a role R is guided by the trustworthiness value $t_{A \rightarrow \langle B, R, I \rangle} \in [0..1]$ for each known agent B . It is calculated from A ’s LIT as follows:

$$t_{A \rightarrow \langle B, R, I \rangle} = \frac{\sum_{\langle X, Y, Z \rangle \in LIT_A} c_{A \rightarrow \langle X, Y, Z \rangle} \cdot w_{A \rightarrow \langle X, Y, Z \rangle}}{\sum_{\langle X, Y, Z \rangle \in LIT_A} w_{A \rightarrow \langle X, Y, Z \rangle}} \quad (2)$$

$w_{A \rightarrow \langle X, Y, Z \rangle}$ is the weight given to agent A ’s confidence on situation $\langle X, Y, Z \rangle$. The weights combine the confidence reliability with the similarity of the situation $\langle X, Y, Z \rangle$ to the target situation $\langle B, R, I \rangle$ in the following way:

$$w_{A \rightarrow \langle X, Y, Z \rangle} = r_{A \rightarrow \langle X, Y, Z \rangle} \cdot sim(\langle X, Y, Z \rangle, \langle B, R, I \rangle) \quad (3)$$

The similarity function $sim(\langle I_1 \rangle, \langle I_2 \rangle)$ ($\langle I_1 \rangle \equiv \langle X, Y, Z \rangle$ and $\langle I_2 \rangle \equiv \langle B, R, I \rangle$) is computed as the weighted sum of the similarities of the individual elements (agent, role and interaction) as it is shown in the following equation:

$$sim(\langle I_1 \rangle, \langle I_2 \rangle) = \begin{cases} \beta \cdot sim_R(R, Y) + \gamma \cdot sim_I(I, Z), & \text{if } B = X \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $sim_R(R, Y)$, $sim_I(I, Z) \in [0..1]$ measure the similarity between roles and interactions, respectively, and β and γ , with $\beta + \gamma = 1$, are parameters specifying the sensibility regarding the individual similarities. $sim_R(R, R')$ and $sim_I(I, I')$ can be determined by some measure of the *distance* of concepts in the corresponding taxonomies. We use they simple measure

$$sim_R(x, y) = sim_I(x, y) = 1 - \frac{h}{h_{MAX}} \quad (5)$$

where x, y are either roles or interactions, h is the number of hops between x and y in the corresponding taxonomy, and h_{MAX} is the longest possible path between any pair of elements in the hierarchy tree. Other functions have been described in [2].

3. CONFIDENCE WITH EVOLVING ORGANIZATIONAL STRUCTURES

So far, we have assumed organizational structures to be static. In this section we present a mechanism, based on clustering, that extends the a priori structures in order to improve performance.

During the execution of a VO, there are two elements that can evolve:

- *Agents.* Agents belonging to a VO may change their behavior at their will.
- *Organizational structures.* Role and interaction taxonomies may evolve over time due to agents’ behavioral changes.

We propose to make the changes on the aforementioned elements explicit by using cluster analysis. Clustering techniques can support the learning of these changes that affect the VO, and can prompt the whole organization to evolve its taxonomies.

As outlined previously, the trust values that agent A holds for others in its LIT are based on confidence and reliability values ($c_{A \rightarrow \langle X, Y, Z \rangle}$, $r_{A \rightarrow \langle X, Y, Z \rangle}$) calculated from past interactions. These trust tuples can be conceived as a bi-dimensional trust space. For each pair $\langle R_i, I_j \rangle$, belonging to

an agent A , a trust space TS such that

$$TS_A = \{(c_k, r_k); c_k \in \mathcal{C}_{A \rightarrow \langle R_i, I_j \rangle}, r_k \in \mathcal{R}_{A \rightarrow \langle R_i, I_j \rangle}, 0 \leq c_k \leq 1, 0 \leq r_k \leq 1\}$$

can be constructed, where $\mathcal{C}_{A \rightarrow \langle R_i, I_j \rangle}$ is a set containing confidence values stored by agent A related to interactions of type I_j with counterparts playing role R_i . The same applies to $\mathcal{R}_{A \rightarrow \langle R_i, I_j \rangle}$ for reliability values. Thus, organizational information is subdivided into groups of *role specializations* for specific interactions, giving extra information to agents when using their trust models.

3.1 Role Clustering Algorithm

To specialize roles we use a *K-means* algorithm, where k represents the number of clusters to be made in each execution. The algorithm will take as input a set of points $(x, y) \in TS_i$ for each tuple $\langle R, I \rangle$ and each agent i that may enact R in I with available information, and will return a group of clusters. A cluster mean point represents mean behaviour for all the agents belonging to it, and the whole cluster represents a pattern of behavior for all the agents included. Hence, groups created by clustering can be used to refine role and interaction taxonomies.

An example of how a clustering algorithm groups data for a role in a particular interaction is shown in Figure 1. Let \mathcal{O} be an organization in the domain of medical services. Furthermore, suppose that agents can join \mathcal{O} with two different types of roles: *Patient* and *Physician*. Once the clustering algorithm has been applied, role *Physician* is specialized, for example, in “*Surgeon*” and “*Psychiatrist*”, since there are two groups of agents that perform better than the others for interactions related to *operations* and *mental problems*, e.g., *Cardiovascular Surgery* and *Low Self-Steem Treatment*. New clusters represent new roles (specializations) in the taxonomy, and new roles in the taxonomy may change the participant roles in interactions: if we join \mathcal{O} enacting role *Patient* and we need to look for a Surgeon to perform interaction *Cardiovascular Surgery*, we will use our trust models just with agents that can play that role, and not with those that “only” are *Physician*, since the agents in the former group are supposed to be more trusted than the others.

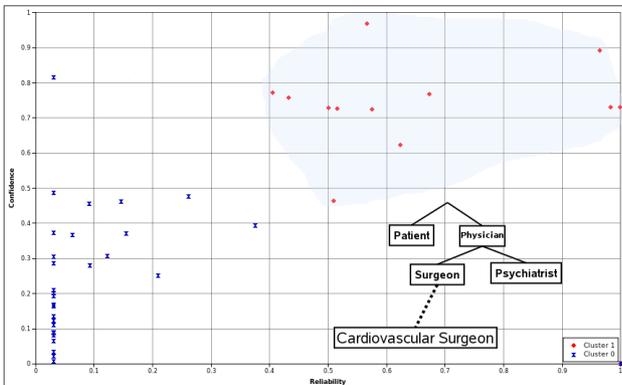


Figure 1: An example of agents behavior pattern recognition for role *Surgeon* and interaction *Cardiovascular operation* using *K-means* with $k = 2$. (Dashed line represents a new specialized role).

3.2 Filtering techniques for Roles

Clustering algorithms divide information in groups, depending on the similarity of their members. Thus, it is reasonable to think that in some cases some clusters can be unnecessary, for instance:

- when the groups that have been formed are clearly different but there is not much trust on them (low confidence and reliability).
- when a cluster is trusted enough but only contains very few agents.

To solve those problems, our approach proposes two different and compatible filtering mechanisms:

1. There cannot exist a cluster whose mean point (mean agent) is outside the 2nd quadrant; that is, in order to specialize a role, mean agent must have a confidence value $c_k > 0.5$ and a reliability value $r_k > 0.5$, since 0.5 represents a point with no information, whilst a lower value of c_k or r_k means a non-efficient cluster that does not entail a necessity to specialize a new role.
2. A minimum number of agents in a cluster has to be imposed to avoid that groups with only a few individuals (compared to the population of the VO) can contribute to form a new specialized role.

4. EVALUATION

In this section we report on a case of study aimed at evaluating our models. Suppose *Travel Services* domain, where the roles that can be played at the beginning are *Customer* or/and *Provider*, and the services are those shown in Figure 2. In our experiments, we used the TOAST testbed developed to evaluate trust models in VOs [3]. Our simulation targeted a customer/provider environment where a number of providers offer different services and a number of customers want to use those services.

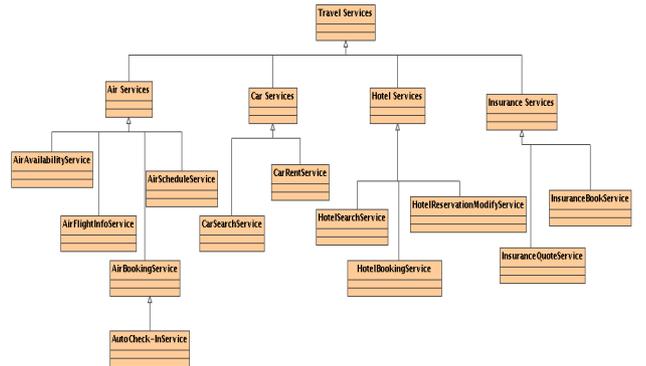


Figure 2: Services taxonomy for tourism scenario

We ran the same experiment for the model presented in Section 2, which we will term *inference model*, and its extension with clustering mechanism with the same scenario, using the same number of customers, goals, services and providers. In particular, we used a collection of 20 clients and 20 providers. Our results have been obtained from a collection of 15000 generated goals (service requests). Furthermore, we repeated each experimental run several times

with different random seeds. The same *Inference model* but using clustering classification (*Inference model evolving*) is used to train and make taxonomies evolve according to role and interaction confidence-based evolution. The clustering algorithm was executed once each 500 interactions. *Inference with new taxonomies* is essentially an *Inference model*, but using learnt taxonomies resulting from executing the *Inference model evolving*. Finally, in the *Basic model* agents evaluate the expected behaviour of the potential candidates for a situation exclusively by using the corresponding confidence value stored in their own LITs. If no entry exists about a situation, e.g., no previous experiences are available, the counterpart is selected randomly.

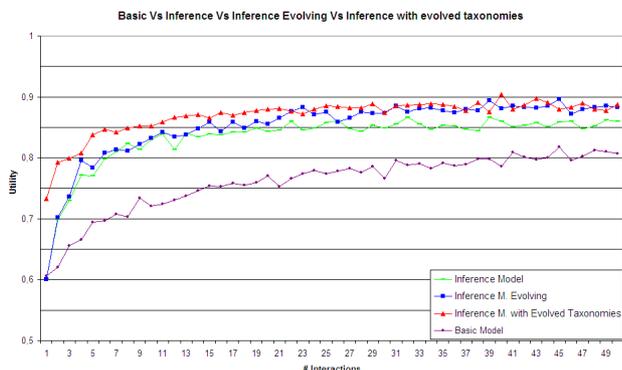


Figure 3: Overall utility performance of different trust models.

Figure 3 shows the evolution of the overall system utility over the number of interactions that have taken place. The overall system utility is calculated as the average of the utilities of all individual customers. As utility values we use the quality values a client obtains after using a service [3]. As it can be observed, the utility improves in the model that uses learnt taxonomies, above all when initiating the execution, since it can make use of more refined and specified taxonomies.

5. CONCLUSIONS

In this paper we have presented a trust model for VOs which takes advantage of organizational information, and use structural similarities to evolve and create new specialized roles and interactions when needed. Our experimental evaluation has provided evidence that this approach improves the previous trust models in a VO, as well as the overall efficiency of the latter.

In contrast to other approaches to trust systems (most of them based on reputation distribution – reputation values exchange about third parties) [4, 10, 7], we have presented a way of evaluating trust at a local level that focuses on the experience of agents obtained in past interactions. The FIRE model proposed by Huynh, Jennings and Shadbolt [4] is also related to *interaction trust* and *role-based trust*. As in our approach, the former is built from direct experience of an agent, while the latter is the rating that results from role-based relationships between agents. However, the FIRE model neither consider inference on VO structures nor learn new taxonomies from the organizational evolution.

As future work we plan to test our trust model with other more complex clustering algorithms, as for example some

based on fuzzy logic for dealing with imprecise information, as well as to take into account not only best cluster as valuable information, but also others clusters that can bring other kind of information to agents. We also intend to introduce untrustworthy agents in clusters and study their behaviors among the others. Finally, we are interested in studying how agents reach a consensus to institutionalize a new specialized role.

6. REFERENCES

- [1] K. K. Fullam and K. S. Barber. Dynamically learning sources of trust information: Experience vs. reputation. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*, pages 1055 – 1062, 2007.
- [2] P. Ganesan, H. Garcia-Molina, and J. Widom. Exploiting hierarchical domain structure to compute similarity. *ACM Trans. Inf. Syst.*, 21(1):64–93, 2003.
- [3] R. Hermoso, H. Billhardt, R. Centeno and S. Ossowski. Effective Use of Organisational Abstractions for Confidence Models. In G. O’Hare, M. O’Grady, O. Dikenelli and A. Ricci, Editors, *Engineering Societies in the Agents World*, volume 4457 of *LNAI*, pages 368–384. Springer-Verlag, 2007.
- [4] T. D. Huynh, N. R. Jennings, and N. R. Shadbolt. FIRE: An integrated trust and reputation model for open multi-agent systems. In *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI)*, 2004.
- [5] A. Omicini and S. Ossowski. Objective versus subjective coordination in the engineering of agent systems. In M. Klusch, S. Bergamaschi, P. Edwards, and P. Petta, editors, *Intelligent Information Agents: An AgentLink Perspective*, volume 2586 of *LNAI: State-of-the-Art Survey*, pages 179–202. Springer-Verlag, Mar. 2003.
- [6] S. D. Ramchurn, C. Sierra, L. Godó, and N. R. Jennings. A computational trust model for multi-agent interactions based on confidence and reputation. In *Proceedings of 6th International Workshop of Deception, Fraud and Trust in Agent Societies*, pages 69–75, 2003.
- [7] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 475–482, New York, NY, USA, 2002. ACM Press.
- [8] M. Schumacher and S. Ossowski. The governing environment. In Weyns, Parunak, and Michel, editors, *Environments for Multiagent Systems II*, volume 3830, pages 88–104. Springer-Verlag, 2006.
- [9] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. Travos: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.
- [10] B. Yu and M. P. Singh. An evidential model of distributed reputation management. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 294–301, New York, NY, USA, 2002. ACM Press.