An Adaptive Reputation Model for VOs

Arturo Avila-Rosas

Instituto Mexicano del Petróleo,
Eje Central 152, México DF, CP 07730, México
aavilar@imp.mx

Abstract. Because Virtual Organisations (VOs) essentially involve cooperating two or more organisations or agents to pursue a common objective, satisfactory cooperation is vital to their success. However, before an agent made the decision to go ahead with the VO, it needs to be *confident* that the rest of the potential partners will be act cooperatively. We show that reputation is a basic ingredient in the formation of VOs. Reputation is computed using an adaptive algorithm, so agents can learn and adapt their reputation models of their partners according to their recent behaviour. Our approach is especially powerful if the agent participates in a VO in which the members can change their behaviour to exploit their partners. The reputation model presented in this paper deals with the questions of deception and fraud that have been ignored in current models of VO formation.

1 Introduction

Recently, a large number of new collaborative, networked organisations have emerged, having as motivation the explosive progress in computer networks and communication systems, but also as a reaction to market pressures that demand customised, high quality products and services at lower costs and, at the same time, shorter production and marketing times. Promising greater flexibility, resource optimisation and responsiveness in *competitive open environments*, VOs are an example of this trend that has pervaded not only business domains but other areas such as e-science. The concept of a VO has been used to describe the aggregation of autonomous and independent organisations connected through a network and brought together to deliver a product or service in response to a customer need [3]. What distinguishes VOs from other forms of organisation is the full mutual dependence of their members to achieve their goal and therefore the need for cooperation. However, open environments in which VOs are embedded involve organisations and individuals that do not necessarily share the same objectives and interests that they might not know in advance, and where they might not trust each other, but should work together and help each other to achieve a common goal. One of the key omissions in the computational representation of VOs relates to the need to take into account more subjective facets like the *reputation* of the individuals, which helps to cope with heterogeneity, autonomy and diversity of interests among members. We observe that current

T. Eymann et al. (Eds.): MATES 2005, LNAI 3550, pp. 204–209, 2005.

solutions underestimate the possibility of swindle in VOs. A common flaw is assuming that the partners selected are fully competent and honest. Since partners represent organisations or individuals who want to maximise their utilities by joining a VO, they have a strong incentive to misrepresent the value of their contributions and enjoy more benefits of cooperative associations [1]. Further, partners are selected in relation to the abilities they claim to have, but it is possible that they do not have such abilities. However, due to lack of information about past interactions, it is difficult to detect and control these situations. This paper considers the introduction of reputation into VOs, by providing a reputation model based on the adaptive evaluation of direct experiences to identify trustworthy individuals to join VO.

The remainder of this paper is organised as follows. The requirement for reputation systems for VOs are discussed in Section 2. In Section 3 we present our reputation model for VOs which is based on reinforcement learning techniques. Section 4 reviews related work, and Section 5 present our conclusions.

2 Requirements

The objective of this section is to delineate the requirements for building a reputation system in order to serve as a decision-making variable in the selection of partners, promote cooperation, produce trust and induce *good* behaviour in the members of a VO.

1. *Distributed reputation management.* Because VOs do not depend on the presence of any centrally trusted authority, there is a need for distributed mechanisms that enable the partners in a VO to collect, store, manage and disseminate reputation in a personalised fashion.
2. *Dynamism.* Due to limitations in time and intense task pressures, partners in a VO should be able to quickly use a reduced number of interactions to estimate the reputation of a partner and; at the same time, take partner selection decisions without having a significant impact, in terms of time consumption, on the formation of a VO.
3. *Adaptability.* Reputation must be updated dynamically to adapt the values of reputation towards true quality of service. This suggests that the updating process should be a *learning* process about another's true abilities, that captures the observed performance through the reputation of the partner. This leads to discarding updating methods that diminish the impact of strategic changes in partner behaviour that intend to milk *high* values of reputation by intentionally deteriorating the provision of the service.
4. *Predictability.* Reputation must provide information to predict the future performance of a partner and eventually the risk involved of interacting with it. That is, based on a partner's previous performance, reputation must provide an indication of its future performance and willingness to accomplish a task.

3 Direct Reputation Model

In this section we introduce our model of reputation, which is based upon SPORAS. SPORAS was chosen because some of its design guidelines are consistent with our specific requirements. Particularly, since SPORAS is an adaptive reputation model that updates the reputation values after each interaction, removing the effects of *obsolete* data in some manner, it is ideal for environments where the behaviour of VO members is changing through the time because the relationships among them are themselves changing as a function of their interests and goals. We start by defining mathematically the concepts of reputation and impressions. Next we describe the methods used in our model for updating reputation.

3.1 Reputation

We define the reputation of an agent as *a perception regarding its intention and competences, which is held by other agents through the formation and dissemination of subjective evaluations based on experiences and observations of past actions.* Here, these evaluations are called *impressions*. From the definition, the observed behaviour of others is collected through: (i) direct experiences, with interaction histories serving as a strong evidence for estimating someone reputation or (ii) via the testimony of others, known as *recommenders*. On the basis of the source of reputation, two concepts of reputation may be derived, namely *direct reputation* and *social reputation*. The concept of social reputation is beyond the scope of this paper.

3.2 Direct Reputation

Direct reputation (DR) is defined as the weighted average evaluation that an agent makes of another's competence, and gives the extent to which the target is *good* or *bad* with respect to a given behaviour or action. Direct reputation is context-dependent so that an agent is reputed according to the service provided. In our algorithm we adopt the ideas of Shapiro [5], then direct reputation is computed as the average of *impressions* received within the most recent time window,

$$W = [t - \epsilon, t]\,, \tag{1}$$

where ϵ defines a time interval that limits the set of interactions and in which impressions are used to compute a direct reputation value. Impressions are weighted from 0 to 1 to indicate the notion of importance of an impression in relation to others for calculating reputation. The direct reputation values vary in the range of [0,1] and are used only to represent comparative values in this continuous space from bad reputation (values near 0) to good reputation (values near 1). The direct reputation of i in the perspective of j in context k is represented as:

$$DR^k_{ij} \in [0, 1].$$

3.3 Impression

We define an impression as an evaluative opinion that is formed by any entity (individual, organisation, etc.) based on a discrete experience with another partner, coupled with the partner's performance. An impression is related with a dimension that describes just one of the qualities of the service as required by agent j. Mathematically, the impression appear as follows,

$$imp_{ij}^{d} \in [0, 1],$$

$$Q_{ij} = \{d \in k | k \; is \; a \; context\}, \tag{2}$$

where i is the service provider whose interaction with the service consumer j left in it the strong impression *imp* in relation to dimension d, and Q_{ij} is the set of dimensions for evaluating a service provider in context k. The numbers used for impressions are merely reference values for making comparisons, each consumer establishes a personal threshold of *acceptable* values for the dimension d evaluated.

3.4 Updating Direct Reputation

Each agent updates its reputation value of a service provider every time it receives impressions from either direct (immediate or observed interactions) or indirect experiences. In order to update the reputation values (after receiving t rated experiences or impressions) we use the following reinforcement learning based action update rules:

$$DR_t = DR_{t-1} + \alpha \cdot [imp_t - DR_{t-1}]. \tag{3}$$

Reputation, in Eq.(3), can be interpreted as the aggregation of the previous value of reputation plus a factor that strengthens or weakens that value. This factor indicates the proximity of the recent impression to the past reputation, and shows of how well the previous reputation predicts the latest given impression. The update rule in Eq.(3) is a linear function which is required in an open environment where the number of prior interactions may be reduced, and reputation cannot be updated in the long term through a non-linear function because an agent could cheat on many occasions before the reputation is updated. Now, if α is near 1 then all the previous history will be forgotten, otherwise, if α is near 0 then the previous history will be preserved.

The factor α is also known as a learning rate, and is an indicator of how long past experiences will last in the memory of the system. For our purposes, we consider α as a function $\alpha(DR_{t-1}, imp_t)$ with the following properties that are based on the ideas of Carbo et al. [2]:

– The function $\alpha(DR_{t-1}, imp_t)$ determines how fast the reputation value changes after an experience and how this affects the memory of the system. This depends on the accuracy of the predictions suggested by the *impressions* received; that is, how much similarity exists between the expectation formed by the previous reputation values and the last rating.

– Similarity will be estimated through a similarity function $\beta(DR_{t-1}, imp_t) \in (0, 1)$:

$$\beta(DR_{t-1}, imp_t) = 1 - e^{-10 \cdot ABS(E-imp)}, \tag{4}$$

where E is the estimated rating based on the past reputation and rating:

$$E = \frac{DR_{t-1} + imp_{t-1}}{2}. \tag{5}$$

– Finally, the function $\alpha(DR_{t-1}, imp)$ is updated as follows:

$$\alpha(DR_t, imp) = \frac{\alpha(DR_{t-1}, imp) + \beta(DR_{t-1}, imp)}{2}. \tag{6}$$

4 Related Work

Zacharia and Maes in [6] present SPORAS, which is a *centralised* reputation system that establishes reputation for users in an on-line community (for example chat rooms, auctions or newsletters groups), based on the aggregation of *rates* given by users after each transaction. Reputation in SPORAS aims to predict future performance of the users. In order to make accurate predictions using a small computational space, a recursive and adaptive algorithm for updating reputation is used. This aggregation method then allows newer rates to count more than older ones. Because SPORAS is a centralised reputation system, it is not viable for VOs where partners need personalised reputation values calculated from assembled rates of those they trust already rather than those they do not know. Moreover, mediators are designed and operated by parties whose interests may sometimes diverge from those of the electronic community. Although the assumption made in SPORAS to make reputation values dependent on the reputation of the entity who is providing a feedback is correct, it mixes two different dimensions of reputation. While a user can be reputed as completely unable to cheat on deals, nonetheless that same user may be a bad evaluator of other users. That is, being an excellent service provider does not mean being an honest evaluator.

REGRET is a reputation system developed by Sabater and Sierra [4] that adopts a sociological approach for computing reputation in societies of agents trading well defined products inside an e-commerce environment. Although REGRET provides a very simple method for aggregating rates (or *impressions* that are the result of evaluating direct interactions) based on the weighted sum of the impressions (more relevance is given to the recent ones), its major contribution is the vision of reputation through of three dimensions. These dimensions are called the *individual dimension*, *social dimension* and *ontological dimension.* As discussed earlier, VOs require to a certain extent that the reputation of a partner is assessed in a *reactive* form to detect possible opportunistic behaviour. However, REGRET's main idea consists of emphasising the freshness of information. Computations in REGRET give a *fixed* high relevance to recent rates over older ones according to a time dependent function, and, moreover the rates are aggregated in a way that can be sensitive to noise since they are simply summed.

5 Conclusions and Future Work

We have provided a critical overview of the state of the art in the field of VOs and reputation. We argue that subjective aspects of partners such as their *competences* and *trustworthiness* should be taken into account in partner selection decisions, since these aspects ultimately influence cooperation between partners. Moreover, we assert that reputation plays an important role in VOs when members decide who to interact with and when to interact, by providing information about the past behaviour of potential partners, their abilities and reliability. In particular, we assert the importance of reputation not only in the formation process of VO but in the operation process.

Additionally, we discussed the requirements for building reputation systems that pursue three basic objectives in the formation and operation of VOs: (1) they provide useful information about potential partners for selecting the most *appropriate*, and eventually enable the formation of VOs; (2) they foster trust among the partners of the VO by revealing each partner's capabilities and predicting its future behaviour; and (3) they offer a means for enhancing cooperation by detecting and deterring deceptive behaviour through imposing *collective sanctions* on defectors.

Although this paper has answered how reputation is relevant to recognise cooperative partners through direct interactions, it opens up more research opportunities and questions that are unanswered. Moreover, there are other issues that were not faced in this paper, due to the bounds imposed on the research, and still need to be addressed.

References

1. S. Braynov and T. Sandholm. Trust revelation in multiagent interaction. In *Proceedings of CHI'02 Workshop on Philosophy and design of Socially Adept Technologies*, pages 57–60, Minneapolis, USA, 2002.
2. J. Carbo, J. Molina, and J. Davila. Trust management through fuzzy reputation. *International Journal of Cooperative Information Systems*, 12(1):135–155, 2003.
3. E. Oliveira and A. Rocha. Agents advanced features for negotiation in electronic commerce and virtual organisations formation processes. In *Agent Mediated Electronic Commerce, the European AgentLink Perspective*, volume 1991 of *Lectures Notes in Artificial Intelligence*, pages 77–96, 2000.
4. J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First International Joint Conference on AAMAS*, pages 475–482, Bologna, Italy, 2002.
5. Carl Shapiro. Consumer information, product quality, and seller reputation. *The Bell Journal of Economics*, 13:20–35, 1982.
6. G. Zacharia and P. Maes. Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(8):881–907, 2000.