Creating and Using Reputation-based Agreements in Organisational Environments*

Roberto Centeno[1], Ramón Hermoso[1], and Viviane Torres da Silva[2]
{roberto.centeno, ramon.hermoso}@urjc.es, viviane.silva@ic.uff.br

[1] Centre for Intelligent Information Technologies (CETINIA)
University Rey Juan Carlos Madrid (URJC) - Spain
[2] Universidade Federal Fluminense (UFF) - Brazil

Abstract. Reputation mechanisms have been developed during last few years as valid methods to allow agents to better select partners in organisational environments. In most of works presented in the literature, reputation is summarised as a value, typically a number, that represents an opinion sent by an agent to another about a certain third party. In this work, we put forward a novel concept of *reputation-based agreement* in order to support the reputation definition, as well as, some desirable properties about it. We define a reputation service that collects opinions from agents, so creating *agreements* over *situations*. This service will also be in charge of presenting the information by using different *informative mechanisms*. Finally, a case study is presented in order to exemplify our work.

1 Introduction

Reputation mechanisms have been proved to be successful methods to build multi-agent systems where agents' decision-making processes to select partners are crucial for the system functioning [4][5][8]. In models such as in [5][8] the authors focus on let the agent the duty of requesting opinions, aggregating replies and inferring conclusions from the gathered information. In this work we do not intend to leave those duties apart, but complementary endow organisations with a reputation service that may help agents to make decisions when they do not have enough relevant stored information.

In this paper we introduce the concept of *reputation-based agreement* as the cornerstone of the reputation service in an organisational multi-agent system. An *agreement* is usually defined as a meeting of minds between two or more parties, about their relative duties and rights regarding current or future performance. Around this concept new paradigms have been emerged [1] oriented to increase the reliability and performance of agents in organisations by introducing in such

* The primary author of the paper is a student. The present work has been partially funded by the Spanish Ministry of Education and Science under project TIN2006-14630-C03-02 (FPI grants program) and by the Spanish project "Agreement Technologies" (CONSOLIDER CSD2007-0022, INGENIO 2010)

communities these well-known human social mechanisms. With this in mind, we propose a novel approach for the meaning of reputation. From a global point of view, a *reputation-based agreement* is a meeting point on the behaviour of an agent, participating within an organisation, with regard to its reputation. Agreements are evaluated by aggregating opinions sent by participants about the behaviour of agents. We also define some interesting properties that describe different types of agreements. Besides, information about reached agreements will be provided to agents by using the concept of informative mechanism [2].

The paper is organised as follows: Section 2 formalises the reputation service, supported by the idea of reputation-based agreements. In Section 3 we illustrate all concepts introduced by means of a case study. Section 4 discusses some related work and, finally, summarises the paper and presents the future work.

2 Reputation Service based on Reputation-based Agreements

As we have previously pointed out, the current work faces with the task of formalising a reputation service working on organisational multi-agent systems. We adhere the definition of organisation given in [3]. Summarising, an organisation is defined as a tuple $\langle \mathcal{A}g, \mathcal{A}, \mathcal{X}, \phi, x_0, \varphi, \{\mathcal{ON}^{om}, \mathcal{R}^{om}\}\rangle$ where $\mathcal{A}g$ represents the set of agents participating within the organisation; $\mathcal{A}$ is the set of actions agents can perform; $\mathcal{X}$ stands for the environmental states space; ϕ is a function describing how the system evolves as a result of agents actions; x_0 represents the initial state of the system; φ is the agents' capability function describing the actions agents are able to perform in a given state of the environment; $\mathcal{ON}^{om}$ is an organisational mechanism based on organisational norms; and $\mathcal{R}^{om}$ is an organisational mechanism based on roles that defines the positions agents may enact in the organisation (see [3] for more details).

The dynamics of the reputation service is threefold: *i)* agents within an organisation have to send their opinions about *situations* in which they have been involved; *ii)* the reputation service aggregates all opinions received from agents, creating *reputation-based agreements*; and *iii)* information about the agreements reached within the organisation is provided to agents by using different informative mechanisms [2]. In following sections we explained each task in detail.

2.1 How Agents Send Their Opinions

Along the lifetime of an agent within an organisation, it is involved in several different *situations*. A situation is defined as a tuple $\langle \mathcal{A}g, \mathcal{R}, \mathcal{A}, T\rangle$, that represents an agent $\mathcal{A}g$, playing the role $\mathcal{R}$, while performing the action $\mathcal{A}$, through a time period T. Agents usually evaluate those situations in order to compile reliable information that allows them to predict the result of future situations. The rationale of the current work is that if agents share their knowledge about the situations they are involved in, this information might be useful when other

agents have not enough information to select partners to interact with. Situations are evaluated from an agent's individual point of view and, at any time, the agent can send its opinion about a particular situation to the reputation service. We call this information *reputation information message* and it is formalised as follows:

Definition 1. *A reputation information message* $\mathcal{R}^{info}_{ag_i \in \mathcal{A}g}$ *is a tuple, representing an opinion sent by the agent* ag_i *to the reputation service containing an evaluation about a particular situation,* $\mathcal{R}^{info}_{ag_i} = \langle Sit, RepVal \rangle$,

where ag_i stands for the agent, which sends the opinion; Sit is the situation being evaluated; and $RepVal$ represents the evaluation the agent is sending about the situation (typically a number).Therefore, an agent, by using this kind of messages, somehow is making public its opinions – evaluations – about different situations: agents, roles, etc.

2.2 Creating Reputation-based Agreements

In this section we intend to face the task of giving a novel approach for the meaning of reputation, from a centralised point of view, tackling this concept as a partial agreement about a certain situation. When the reputation service receives reputation information messages from agents, it aggregates them creating what we have called *reputation-based agreements*. That is, the aggregation of all the opinions regarding a particular situation is '*per se*' what set of agents – as a whole – actually think about the aforesaid situation. Thus, a reputation-based agreement represents the consensus reached in the reputation opinions space sent by a set of agents about a particular situation. Formally:

Definition 2. *A reputation-based agreement* π *for a particular situation, is a tuple* $\langle Sit, \mathcal{A}g, RepVal, t \rangle$

where Sit stands for the situation about which the agreement is reached; $\mathcal{A}g$ is the set of agents that contributed to the agreement; $RepVal$ represents the reputation value – whatever its representation is (qualitative, quantitative, etc.) – reached as consequences of all opinions sent about the situation; and t stands for the time when the agreement was reached. Therefore, an agreement means a global opinion that a set of agents have on a certain situation. This agreement, as we put forward in next section, can be used as a generalist expectation for a situation in which agents have no (or little) previous information about.

As we have claimed, a reputation-based agreement is reached as consequence of the aggregation of all opinions sent about a particular situation. Thus, the reputation service requires a function that is able to aggregates information reputation messages sent by agents. The aim of such a function is to create agreements from reputation opinions that agents send to the service by means of reputation information messages. We formally define the function as follows:

Definition 3. *Let* f_π *be a function that given all the reputation information messages sent by agents and a particular situation creates a reputation-based agreement for that situation:* $f_\pi : |\mathcal{R}^{info}_{ag_i \in \mathcal{A}g}| \times Sit \rightarrow \Pi$

where $|\mathcal{R}^{info}_{ag_i \in \mathcal{A}g}|$ stands for the set of reputation information messages received by the reputation service; Sit is the set of situations; and Π represents the set of reputation-based agreements.

As aggregation function the module might use any function that is able to aggregate values without any modification. For instance, it is possible to use a simple function to calculate the average of all opinions or a more elaborated function that aggregates the opinions by means of complex calculation, i.e. weighting the assessment by taking into account who is giving the information (the implementation of this function is out of the scope of this paper).

2.3 Reputation-based Agreements: Properties

From previous definitions – definitions 2 and 3 – it is possible to define some desirable properties about reputation-based agreements. These properties should be taken into account when agreements are created and may also provide useful extra information when informing about different issues.

Property 1 *A reputation-based agreement π is **complete** iff. all agents participating in an organisation, at time t, contribute to reach that agreement:*

$$\pi^* \Leftrightarrow \begin{cases} \mathcal{O} = \langle \mathcal{A}g, \mathcal{A}, \mathcal{X}, \phi, x_0, \varphi, \{\mathcal{ON}^{om}, \mathcal{R}^{om}\}\rangle \ \wedge \\ \pi = \langle Sit, \mathcal{A}g', RepVal, t\rangle \ \wedge \\ (\mathcal{A}g = \mathcal{A}g') \end{cases}$$

That is, given a time t every participant $ag \in \mathcal{A}g$ in the organisation $\mathcal{O}$ has necessarily sent a reputation information message indicating its opinion about the situation concerning the agreement ($\mathcal{A}g = \mathcal{A}g'$). The more complete agreements are in the system, the more reliability the information will be supposed to offer.

Property 2 *A reputation-based agreement π is α-**consistent** iff. the reputation value of π differs, at most, $1-\alpha$ from the reputation value sent by every agent that contributed to reached that agreement:*

$$\pi^\alpha \Leftrightarrow \begin{cases} \pi = \langle Sit, \mathcal{A}g, RepVal, t\rangle \ \wedge \\ \forall ag \in \mathcal{A}g \ \ [\forall r \in Rep^{info}_{ag}[(r = \langle Sit_i, RepVal_i\rangle) \wedge \\ (Sit_i = Sit) \ \wedge \ (|RepVal_i - RepVal| \leq 1 - \alpha)]] \end{cases}$$

This property represents how agents sending their opinions about a situation agree in a certain extent. Therefore, the higher α is, the more similar the opinions are.

Property 3 *A reputation-based agreement π is **full** iff. it is complete and 1-consistent:* $\pi^\phi \Leftrightarrow (\pi^* \ \wedge \ \pi^\alpha \ \wedge \ \alpha = 1)$

In the case α is 1 means that all agents have the same opinion about a given situation. This property is very desirable when seeking reputation-based agreements, because the more agents contribute to the agreement, the stronger validity the latter gets. Thus, the likelihood of capturing what is actually happening in the organisation tends to be higher.

Property 4 *A reputation-based agreement π is $\mathcal{R}$-**consistent** iff. all the agents contributing the agreement play the same role in the system:*

$$\pi_{\mathcal{R}} = \langle Sit, \mathcal{A}g, RepVal, t\rangle \Leftrightarrow \forall ag \in \mathcal{A}g \;\; play(ag, \mathcal{R})$$

where $\mathcal{R}$ stands for the role the consistency is based on, $\mathcal{A}g$ is the set of agents that contribute to reach the agreement, and $play : \mathcal{A}g \times \mathcal{R} \rightarrow [true, false]$ is a function that returns *true* if the agent $\mathcal{A}g$ plays the role $\mathcal{R}$.

This property is useful in cases in which a new agent, joining an organisation, wants to know what other agents – that are executing in the organisation and playing the same role to it is, think about a given situation. For instance, someone who is thinking of *buying* something would like to know which are the opinions of those who have previously played the role *buyer*.

Property 5 *A reputation-based agreement π is $\mathcal{R}$-**complete** iff. is $\mathcal{R}$-consistent and is complete for all the agents that play the role $\mathcal{R}$ at time t:*

$$\pi^{*}_{\mathcal{R}} = \langle Sit, \mathcal{A}g', RepVal, t\rangle \Leftrightarrow \begin{cases} \pi_{\mathcal{R}} \wedge \mathcal{O} = \langle \mathcal{A}g, \mathcal{A}, \mathcal{X}, \phi, x_0, \varphi, \{\mathcal{ON}^{om}, \mathcal{R}^{om}\}\rangle \wedge \\ \forall ag \in \mathcal{A}g \;\; (play(ag, \mathcal{R}) \rightarrow ag \in \mathcal{A}g') \end{cases}$$

Property 6 *A reputation-based agreement π is $\mathcal{R}$-**full** iff. is $\mathcal{R}$-complete and is 1-consistent: $\pi^{\phi}_{\mathcal{R}} \Leftrightarrow (\pi^{*}_{\mathcal{R}} \;\wedge\; \pi^{\alpha} \;\wedge\; \alpha = 1)$*

2.4 Providing Information about Reputation-based Agreements

Once we have defined an agreement as a distributed consensus-based expectation for a set of agents on a certain situation, we now present how the reputation service can present the relevant information on the reached agreements to the agents participating in the organisation. Reputation-based agreements somehow capture the general thinking about a particular situation – when the more α-consistent the agreement is the more the reality is captured. Thus, information about the agreements reached until that moment may be very useful for agents. In particular, when agents have recently joined the organisation, they do not have any hint about situations in which they might be involved in, so if the reputation service provides information about agreements, agents may improve their utility from the very beginning.

With this in mind, we deal with the problem of how the reputation service may provide such information. To that end, we part from the notion of *informative mechanism* [2]. Those types of mechanisms are in charge of providing some kind of information to agents in order to regulate a multi-agent system. Thus, an *informative mechanism* $\Gamma : \mathcal{S}' \times \mathcal{X}' \rightarrow \mathcal{I}$ is a function that given a partial description of an internal state of an agent (S') and, taking into account the partial view that the service has of the current environmental state ($\mathcal{X}'$), provides certain information ($\mathcal{I}$). We adhere this definition to create mechanisms over the agreements for different situations, creating information valuable for participants in the organisation. We formally define them as follows:

Definition 4. *An informative mechanism providing information about reputation-based agreements is: $\Gamma_{\Pi} : Sit \times \mathcal{X}' \rightarrow \mathcal{I}_{\Pi}$*

where Sit,$\mathcal{X}'$ are already defined and $\mathcal{I}_{\Pi}$ stands for the information provided by the mechanism by using the set of agreements Π reached over the situation Sit.

We have chosen a very general definition about information in order to cover all possible types of information the reputation service could offer taking into account the reputation-based agreements reached. The information provided may consist of a ranking sorting the best agents for a particular situation, such as $\langle \_, \mathcal{R}, \mathcal{A}, \_ \rangle$, created from the agreements reached for that situation, a value representing the reputation value for a situation, reached as consequence of the agreement for that situation, information about the properties of the agreement reached for a particular situation, if it is full, complete, etc.

3 Case Study: Pubs Area

In this section, we illustrate the proposed model by means of a simple case study. The scenario we use involves four different agents: *Anna*, *John*, *Jessica* and *Albert*. In this organisation agents can *order* and *delivery* drinks. That organisation is created with the aim of getting in touch pubs' owners and providers of drinks. Thus, agents join the organisation playing the roles of *pub* and *provider*, representing a pub's owner and a company provider of drinks, respectively. In our particular example, agents are playing the following roles: *Anna* - *pub*, *John* - *pub*, *Jessica* - *pub* and *Albert* - *provider*.

In this scenario, agents representing pubs' owners are interested in collaborating sharing information about providers, because the pubs are situated in the same area and they collaborate each other so as to foster the attraction customers to that area. That is, although they try to maximize their own benefit, one of their goals is to foster the pubs area where they are, even if that entails to exchange information about drink providers.

Therefore, after several interactions among them – performing actions of ordering and delivering different types of drinks – *Anna* decides to make public her opinion about *Albert* as provider. Thus, she uses the reputation information messages to send to the reputation service her opinion, as follows:

$$\mathcal{R}^{info}_{Anna} = \langle \langle Albert, provider, \_, \_ \rangle, 0.8 \rangle$$

This information shows that *Anna* has had bad experiences while she was ordering drinks to *Albert* (0.8) because *Albert* almost never violates contracts and offers low prices. Similarly, *John* and *Jessica* send their opinions about *Albert* as providers, by using the following reputation information messages:

$$\mathcal{R}^{info}_{John} = \langle \langle Albert, provider, \_, \_ \rangle, 0.7 \rangle \ ; \ \mathcal{R}^{info}_{Jessica} = \langle \langle Albert, provider, \_, \_ \rangle, 0.9 \rangle$$

It seems that both *John* and *Jessica* agree that *Albert* is a reliable provider.

When the reputation service receives this information, it is able to create reputation-based agreements by using a function that aggregates the reputation information messages. Let us suppose that it aggregates the messages by cal-

culating the average of reputation values sent by agents over exactly the same situation[3]:

$$f_\pi(Sit) = \frac{\sum_{i=1}^{n} \mathcal{R}_{ag_i}^{info} = \langle Sit, RepVal_i \rangle}{n}$$

Therefore, from the set of messages sent by the agents, so far, the reputation service can create two reputation-based agreements regarding to two different situations:

$$\pi_1 = \langle \langle Albert, provider, \_, \_ \rangle, \{Anna, John, Jessica\}, 0.8, t \rangle$$

π_1 represents that there exists an agreement within the organisation regarding to *Albert* as *provider* – regardless the action he performs – is evaluated as 0.8, and such an agreement is reached by the collaboration of *Anna*, *John* and *Jessica*, at time t.

In order to provide information about agreements the reputation service makes available three different informative mechanisms:

- $\Gamma_{\Pi}^{1}(\langle \mathcal{A}g, \mathcal{R}, \_, \_ \rangle)$ given a situation where an agent and a role are specified, returns meta-information[4] about the reputation-based agreement reached regarding that situation;
- $\Gamma_{\Pi}^{2}(\langle \mathcal{A}g, \mathcal{R}, \_, \_ \rangle)$ given a situation where an agent and a role are specified, returns the reputation-based agreement reached. In particular, it returns the reputation value in the agreement of that situation;
- $\Gamma_{\Pi}^{3}(\langle \_, \mathcal{R}, \_, \_ \rangle)$ given a situation where a role is specified, returns a ranking of agents playing that role, sorted by the reputation value they have as consequence of the reputation-based agreements reached until the current time t.

Let us suppose that a new pub is opened in the same area by *Alice*, so she joins the organisation playing the role *pub*. Since she is looking for a drinker *provider*, a ranking of "best" providers would be a great solution to select the best one. She can thus use Γ_{Π}^{3}: $\Gamma_{\Pi}^{3}(\langle \_, provider, \_, \_ \rangle)$. Lets suppose that by using this information *Alice* knows that there exists an agreement within the organisation showing that *Albert* is among the best providers. But, how good is him?. To answer this question *Alice* queries the informative mechanism Γ_{Π}^{2} as follows:

$$\Gamma_{\Pi}^{2}(\langle Albert, provider, \_, \_ \rangle) \Rightarrow 0.8$$

In addition, she would like to know if the agents that have testified about *Albert*'s behaviour has similar opinions about him. Therefore, she performs the following query:

$$\Gamma_{\Pi}^{1}(\langle Albert, provider, \_, \_ \rangle) \Rightarrow \pi^{0.9}$$

With this information *Alice* knows that all opinions sent about *Albert* are coincident because of the reputation-based agreement reached is 0.9-consistent ($\pi^{0.9}$).

[3] It could be used whatever other function that is able to aggregate the information received from agents

[4] with meta-information we mean the α-consistency of the agreement, if it is full, complete, etc.

4 Conclusions

As stated before, one of the main advantages of having a centralised reputation service is that it decreases the cost and makes it feasible for an agent to know a more consistent reputation about another agent based on numeral experiences. In the case of distributed mechanisms (such as [5][8]), the agent itself would need to participate in several interactions with the given agent and also to ask other agents for their experiences with others. In the case of a centralised mechanism, the agent can easily get information about the reputation showing the behaviour of other agents within the system. In [9], Sabater et al. present a centralised reputation mechanism that is incorporated as a service in Electronic Institutions (EIs). From a global perspective, this work has many similarities with ours, since uses also a reputation service in an organisational environment (EIs). However, the authors do not focus on how to exploit the collected information as agreements that can be presented to agents in different ways. In [6] the authors present an approach to create rankings able not only to provide the most trustful agents but also a probabilistic evidence of such reputation values. Those rankings are also computed by a centralised system by aggregating the reputations reported by the agents. This approach and the one presented in our paper could be complementary, since that paper focuses on defining the ranking algorithms and ours focuses on describing the mechanism that allows to receive the reputation information and to provide the already evaluated agreements (for instance by using rankings). Another work that could be also complementary to the approach presented here, is the one presented in [7]. They describe the algorithm *NodeRanking* that creates rankings of reputation ratings. Therefore, our reputation service could use this algorithm so as to provide information about the reputation-based agreements reached within the organisation.

Summarising, this work puts forward a novel approach of reputation-based agreement concept by supporting on a reputation service that creates reputation-based agreements as aggregations of opinions sent by participants within an organisation. Besides, we also define some desirable properties that can be derived and should be taken into account when providing the information they contain. Furthermore, we also propose to use the agreements by utilising the concept of informative mechanisms [2], so providing agents with useful information. Finally, an example has been analysed so that it illustrates how the reputation service works in a collaborative domain where agents are interested in sharing their opinions. In future work we plan to experimentally test our approach by implementing a case study presented here, as well as, running several experiments comparing our approach with similar ones. We also intend to investigate new properties about reputation-based agreements to provide agents participating in an organisation with more useful information. Finally, we plan to extend the concept of reputation-based agreement by creating agreements aggregating "similar" situations, so we must go into the concept of similar situations in depth.

References

1. C. Carrascosa and M. Rebollo. Modelling agreement spaces. In *Proc. of WAT@IBERAMIA*, pages 79–88, 2008.
2. R. Centeno, H. Billhardt, R. Hermoso, and S. Ossowski. Organising mas: A formal model based on organisational mechanisms. In *SAC*, pages 740–746, 2009.
3. R. Centeno, V. Torres da Silva, and R. Hermoso. A reputation model for organisational supply chain formation. In *Proc. of COIN@AAMAS*, pages 33–48, 2009.
4. Chrysanthos Dellarocas. *Handbook on Economics and Information Systems*, chapter Reputation Mechanisms. Elsevier, 2005.
5. T. Huynh, N. Jennings, and N. Shadbolt. Fire: An integrated trust and reputation model for open multi-agent systems. In *Proc. of ECAI*, pages 18–22, 2004.
6. A. Ignjatovic, N. Foo, and C. Lee. An analytic approach to reputation ranking of participants in online transactions. In *Proc. of WI-IAT*, pages 587–590. IEEE Computer Society, 2008.
7. J. Pujol, R. Sangüesa, and J. Delgado. Extracting reputation in multi agent systems by means of social network topology. In *Proc. of AAMAS*, pages 467–474, New York, NY, USA, 2002. ACM.
8. J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *AAMAS*, pages 475–482, 2002.
9. Jordi Sabater-Mir, Isaac Pinyol, Daniel Villatoro, and Guifr Cun. Towards hybrid experiments on reputation mechanisms: Bdi agents and humans in electronic institutions. In *Proc. of CAEPIA*, volume 2, pages 299–308, 2007.