

SOARI: A service-oriented architecture to enable interoperability of agent reputation models

Luis Gustavo Nardin and Jaime Simão Sichman

Programa de Pós-Graduação em Engenharia Elétrica
Escola Politécnica da Universidade de São Paulo
Av. Prof. Luciano Gualberto, 158 – trav. 3, CEP: 05508-970 – São Paulo – SP
{luis.nardin, jaime.sichman}@poli.usp.br

Abstract. Trust and reputation have proved to help protect societies against harmful individuals. Inspired by these principles, many computational models have been and new ones continue to be proposed to protect multi-agent systems. In an open system, it is possible that different agents use different trust and reputation models. Since agents have to exchange information to make their trust and reputation models more robust, and the models use different internal concepts, interoperability among these models becomes a problem. This paper proposes a service-oriented architecture to deal with this problem and illustrates its usefulness by analyzing the effects caused on the agents as a consequence of a more expressive and heterogeneous communication about reputation.

Extended Information

This work was concluded to the partial fulfillment of the requirements for the Degree of Master of Science in Electrical Engineering on March 19th, 2009. The student's advisor was Prof. Dr. Jaime Simão Sichman and the examining board members were:

- Prof. Dr. Jaime Simão Sichman (EP/USP)
- Prof. Dr. Renata Wassermann (IME/USP)
- Prof. Dr. Virginia Dignum (University of Utrecht)

The publications on-line derived from this research work are [11], [12], [10] and [13], whose can be found at <http://www.lti.pcs.usp.br/~gnardin/publications.html>.

The full dissertation can be found at <http://www.teses.usp.br/teses/disponiveis/3/3141/tde-13072009-154149/>.

1 Introduction

Agents present the capabilities of both acting autonomously and engaging in social activities [18]. In open environments, where agents can enter or leave the environment at any time, taking part in such social activities may expose them to risks, for instance, when taking decisions based on information provided by malevolent agents. Some solutions to this problem are based on trust models which serve as a decision

criterion for an agent to act and engage in societies. Some examples of computational trust models are [19, 5, 15, 14, 9], which are mostly based on the concept of reputation borrowed from the social sciences.

Reputation is a social property or a social process. It is a social property when considered as an agent's mental representation about other agents (i.e. an evaluative belief) and a social process when considered as the result of the belief's transmission (i.e. a meta-belief) [5].

In order to accelerate the reputation evaluation and to improve the robustness of their reputation models, the agents generally exchange information about the reputation of third parties. However, since there is no consensus about a single unifying reputation definition, the semantics associated with reputation concepts differ from one model to another, which raises an interoperability problem as depicted on Fig. 1.

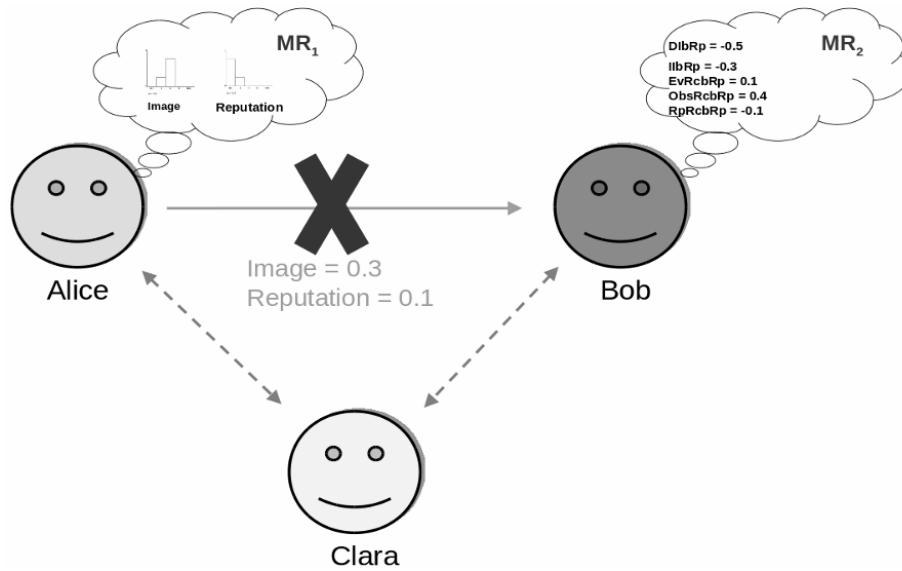


Fig. 1. Reputation Interaction Scenario

Fig. 1 is composed of three agents (Alice, Bob and Clara), where agents Alice and Bob have represented internally the reputation evaluation about Clara, each one using different reputation models, respectively MR₁ and MR₂. Their reputation evaluations are initially based on a direct interaction of those agents with Clara. However, in order to propagate her reputation evaluation about Clara, the agent Alice wants to transmit it to agent Bob, but since they use different reputation models, they are unable to communicate using its internal reputation model concepts. In order to overcome this limitation a mechanism or architecture is required to support the agent reputation interoperability. In this paper, we propose a service-oriented architecture (SOA), named SOARI, to

support agent reputation interoperability. Moreover, we perform some experiments and present their results where SOARI is used to enable interoperability of two reputation models: Repage [14] and L.I.A.R. [9]. These experiments evaluate the impact that the reputation models interoperability may cause on agents evaluation accuracy. The rest of the document is organized as follows. Section 2 presents an overview of the SOARI architecture and its components, followed by a presentation of its usage in section 3. In section 4, some experiments using SOARI, as well as their results and analysis are shown. Finally, our conclusions and future work are presented in section 5.

2 SOARI Overview

SOARI is a service-oriented architecture to support the semantic interoperability among agents that implement heterogeneous reputation models. The main underlying idea of SOARI is that the mapping among different reputation models, represented as ontologies, may be realized off-line and be available on-line as a service for other agents that use the same reputation model.

Fig. 2 shows the SOARI architecture and highlights its components integrated into an agent architecture.

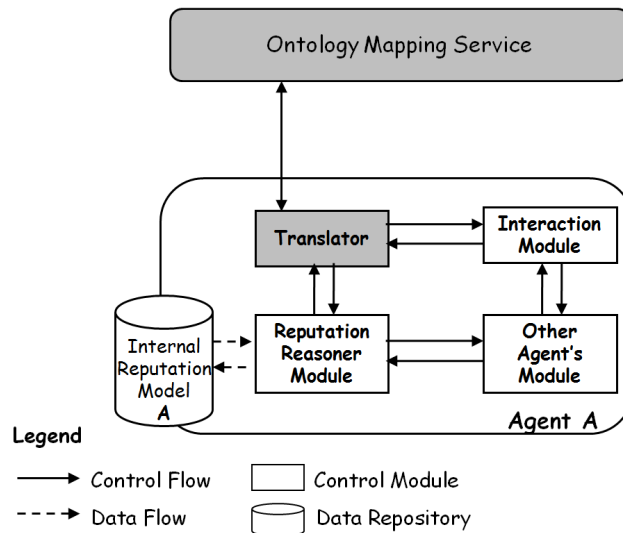


Fig. 2. Service-Oriented Architecture for Reputation Interaction

In this architecture, it is considered that different agents may have heterogeneous reputation models and it uses an hybrid approach proposed by Visser et al. [17] to allow the interoperation among them based on a common vocabulary. In this work, a common functional ontology of reputation that subsumes most agent reputation models, named

FORE [4], is used as the interlingua.

Interoperation requires agents to perform two distinct but interrelated ontology functions: *Mapping* and *Translation*. *Mapping* is a collection of functions assigning the concepts and relations in one ontology to the concepts and relations in another ontology. *Translation* is the application of the mapping functions to translate sentences from a source to a target ontology [8].

SOARI, which is an extension of the general agent architecture for reputation interoperability proposed by [16], extends it in two ways in order to perform those two functions: (i) it subdivides the Reputation Mapping Module (RMM) in two distinct and specialized modules: the Ontology Mapping Service (OMS), which performs the ontology mapping function, and the TRANSLATOR module, which performs the ontology translation function. (ii) It performs the ontology mapping function as a service outside the agent architecture.

- The OMS module is a service outside the agent that implements the ontology mapping function and presents two main functionalities: (i) to map concepts from the target’s reputation model ontology to the concepts of the common ontology; and (ii) to answer concept translation requests from the TRANSLATOR module.
- The TRANSLATOR module resides inside the agent and it translates reputation messages. It has four main activities: (i) to translate the reputation messages from the common ontology to the target agent’s reputation model ontology whenever the message comes from the Interaction Module (IM); (ii) to translate the reputation messages from the agent’s reputation model ontology to the common ontology whenever the message is sent to IM; (iii) to trigger some function in the Reputation Reasoner Module (RRM) based on the interpretation of messages written using the reputation model ontology; and (iv) to create a message using the reputation model ontology whenever requested by RRM.

The advantage of using a service-oriented architecture, from a design/programming perspective, is that the agents become simpler since they do not need to perform the mapping function internally. In the other hand, the advantage of using the hybrid approach comes from the fact that the agents do not know the other agents internal reputation model, and hence cheating is avoided. By defining such extension, we intended to alleviate the agent dynamic workload, since the agent does not need to perform the mapping function on-line. Moreover, the results of such mapping is stored in the service and it may be reused by new agents that enter the system and have an internal reputation model that was already mapped and stored in the service.

Detailed information about this architecture can be found in [12].

3 SOARI Usage

In order to consider the use of the SOARI architecture proposed, at least one OMS must exist in the system to support the mapping and translation of reputation model ontology concepts to a common reputation ontology concepts and vice-versa. The following two steps are required prior to use OMS [11]:

1. to design the reputation model ontologies of the reputation models, if the reputation models are not described in ontological terms, since the OMS only maps ontologies.
2. to align the reputation models ontology to the common reputation ontology, since the OMS processes only ontologies that are already described in terms of a common reputation ontology. In this work, the common reputation ontology used is the FORE.

3.1 Designing the Reputation Model Ontologies

We built the reputation model ontologies manually by using Protégé as the editor and OWL [1] as the ontology language, which is the most recent standard ontology language from the World Wide Web Consortium (W3C¹). In the sequence, the terminologies identified as concepts in two reputation models, L.I.A.R. [9] and Repage [14], were described. Those concepts compose the design of the reputation model ontologies.

L.I.A.R. (Liar Identification for Agent Reputation) is a model for the implementation of social control of agent interaction. The idea is to provide tools that allow agents to (1) reason about other agent's interaction; (2) detect any interaction rules violation; and (3) maintain a reputation model of other agents. Its reputation model distinguishes reputation in five different types, which are based on seven roles involved in the reputation-related processes. Each of the seven roles is defined by the source and kind of information used to calculate the reputation value. The seven roles are: *Target* role, *Participant* role, *Observer* role, *Evaluator* role, *Punisher* role, *Beneficiary* role and *Propagator* role. The five different types of reputation are: *Direct Interaction-based Reputation* (DIbRp), *Indirect Interaction-based Reputation* (IIBRp), *Observation Recommendation-based Reputation* (ObsRcbRp), *Evaluation Recommendation-based Reputation* (EvRcbRp) and *Reputation Recommendation-based Reputation* (RpRcbRp). Each reputation is associated to a facet, which is the subject the evaluation is about.

Repage. The Repage (Reputation and ImAGE) system is a computational module based on a reputation model proposed by [5]. *Image* and *reputation* are the two main concepts in this model and they represent social evaluations. *Image* is an evaluative belief, which is formed using information, acquired by agent experience or propagated third-party images. *Reputation* is a meta-belief, which is formed, based on anonymous reputation value transmitted on the social network about the target agent. The social evaluations are context-based which means that the agent may hold different social evaluation for the same target (*AgentImage* and *AgentReputation*). The model distinguishes the types of agents involved in the image or reputation formation in four different types: *Target* agents, *Evaluator* agents, *Propagator* agents and *Beneficiary* agents.

3.2 Aligning the Reputation Model Ontologies to FORE

Alignment is the establishment of binary relations between the concepts of two ontologies [8]. The binary relations used to perform this operation in our case are defined in

¹ <http://www.w3c.org>

FORe.

We manually defined each of the reputation model concepts identified previously in terms of FORe. For example, the L.I.A.R. instances of the *Direct Interaction-based Reputation* concept have at least one association through the *hasInformationSource* relation to instances of the *DirectExperience*, formally defined:

$$\exists \text{hasInformationSource}(\text{DirectExperience})$$

In addition, the Repage instances of the Image concept have at least one association through the *hasInformationSource* relation to instances of the *DirectExperience* or *Observation* or *SecondHandInformation*, formally defined:

$$\exists \text{hasInformationSource}(\text{DirectExperience or Observation or SecondHandInformation})$$

Having the alignment, it is manually stored in the Ontology Repository. Therefore, the Ontology Mapping Service detects this new ontology and executes the Classifier Module to process it. As a result, the output of such process is stored in the Translation Repository. This is made for each reputation model ontology. The detailed mapping results generated by the OMS can be seen in [11].

Besides the existence of one OMS, the TRANSLATOR module must exist in each one of the agents using SOARI. The integration of the TRANSLATOR module to the agent has to be performed by the agent developer and requires adaptations in the IM and RRM modules.

4 Experiments

In this section, we provide an analysis of the effects caused on the agents' reputation evaluation accuracy as a consequence of a more expressive and heterogeneous communication about reputation using SOARI. Basically, this section intends to answer two questions: (1) is there any improvement in the accuracy of the agents' reputation evaluation when enabling more expressive communication about reputation? (2) how does the heterogeneity influence the accuracy of dishonest agent's reputation evaluation?

In order to answer those questions, some experiments were performed using two testbeds: ART [7] and FOREART [3].

The ART testbed (Agent Reputation and Trust testbed) is currently the unique platform freely available to perform benchmarks with heterogeneous reputation models. However, this platform does not allow the agents to communicate about reputation using their distinct semantically reputation model concepts, thus losing expressiveness.

In the other hand, the FOREART testbed, which is an extension of the ART testbed, allows a more expressive communication among the agents. In order to reach this goal, this latter platform allows the transmission of symbolic messages composed of concepts of FORe.

Both testbeds simulate an art appraisal game, where agents evaluate paintings for clients and gather opinions and reputations from other agents to produce accurate appraisals. However, they differ because the agents in ART testbed interoperate about reputation

using numeric values, while in the FOREART testbed they use symbolic values. In those experiments, one agent deliberately lies about the other agents' reputation and about paintings evaluation. The analysis was performed to determine how accurate the other agents are in the evaluation of the reputation of the liar agent. The experiments include two types of agents: *Honest* and *Dishonest*. The *Honest* agents answer to the requests only when they have expertise about the requested painting era and with information coherent to their internal state. The *Dishonest* agents answer to all the requests, even when they do not have expertise about that painting era and they never answer the requests with information coherent to their internal state.

4.1 Experiments Description

The main objective of these experiments was to identify the mean value of the reputation assigned by the *Honest* agents to the *Dishonest* agent. In order to enable comparison between the experiments, the initial painting era knowledge and clients distribution were identical in all the experiments. Moreover, all the agents used the same configuration parameters in all the simulations.

In order to reach this goal, we executed 10 simulations ($p = 10$) for each experiment with 100 cycles each. Each simulation was composed of 21 agents ($n = 21$), where 20 agents were *Honest* and 1 agent was *Dishonest* ($i = [1, 20]$ and $j = 21$). The mean value of the reputation assigned to the *Dishonest* agent by each *Honest* agent (r_j) considered only the value obtained in the last simulation cycle ($l = 100$ and $m = 100$). The value of the last simulation cycle was used because it was considered the most accurate evaluation.

Formally, consider a set of n agents, where $i = \{1, 2, \dots, n-1\}$ are *Honest* agents and $j = n$ is a *Dishonest* agent. Moreover, consider that r_{ij}^{sk} is the reputation value assigned by agent i to agent j in cycle k on simulation s . Typically, the reputation value assigned by agent i to agent j on simulation s corresponds to the mean reputation value of a set

of cycles. Thus, $r_{ij}^s = \frac{\sum_{k=l}^m r_{ij}^{sk}}{m-l+1}$, where l and m represents, respectively, the lower and upper cycle limits. The mean reputation value assigned by the *Honest* agents to the

Dishonest agent on simulation s is $r_j^s = \frac{\sum_{i=1}^{n-1} r_{ij}^s}{n-1}$. Finally, given a set of simulations $s = 1, \dots, p$ that compose an experiment, the mean value of the *Dishonest* agent is

$$r_j = \frac{\sum_{s=1}^p r_j^s}{p}.$$

The experiments performed were classified based on two dimensions: (1) reputation models used by the agents in the experiment (Repage, L.I.A.R. or both), and (2) reputation communication method (numeric or symbolic) (Table 1). Moreover, the mixed experiments are split in two others based on the reputation model of the *Dishonest* agent. This distinction is indicated by the *D/L.I.A.R.* and *D/Repage* suffix in the experiment's name. In the other experiments, the *Dishonest* agent uses the same reputation model than the *Honest* agents.

Table 1. Summary of experiments

ID	Experiment Name	Reputation Model	Reputation Communication
exp1	ART/L.I.A.R.	L.I.A.R.	Numeric
exp2	ART/Repage	Repage	Numeric
exp3.1	ART/Mixed-D/L.I.A.R.	L.I.A.R. and Repage	Numeric
exp3.2	ART/Mixed-D/Repage	L.I.A.R. and Repage	Numeric
exp4	FOReART/L.I.A.R.	L.I.A.R.	Symbolic
exp5	FOReART/Repage	Repage	Symbolic
exp6.1	FOReART/Mixed-D/L.I.A.R.	L.I.A.R. and Repage	Symbolic
exp6.2	FOReART/Mixed-D/Repage	L.I.A.R. and Repage	Symbolic

4.2 Experiments Results and Analysis

Here, we present an analysis of the results obtained from the experiments in order to answer the two questions posed at the beginning of this section. The analysis methodology used to answer the questions raised on this section is based on nonparametric hypothesis test called Wilcoxon’s Rank Sum Test [2]. This hypothesis test was selected since not all data to be analyzed were normally distributed.

The analysis performed on this section was based on the L.I.A.R. and Repage reputation models attributes. For reputation model attribute, we mean the different concepts of reputation defined in each reputation model presented in section 3.1.

The L.I.A.R. reputation model defines five different types of reputation: *Direct Interaction-based Reputation* (DIbRp); *Indirect Interaction-based Reputation* (IIbRp); *Observation Recommendation-based Reputation* (ObsRcbRp); *Evaluation Recommendation-based Reputation* (EvRcbRp); and *Reputation Recommendation-based Reputation* (RpRcbRp). The Repage reputation model defines two reputation concepts: *Image* and *Reputation*.

Effect of the expressiveness of communication. In order to analyze the effects of the more expressive communication, it was verified if the mean value of the *Dishonest* agent’s attributes (r_j) obtained using numerical reputation values (ART experiments) were higher than the similar ones obtained on the symbolic testbed (FOReART experiments). The underlying idea was that if these results were statistically different, then it would mean that the *Dishonest* agent is better identified if reputation is expressed and exchanged in a more richer level. Thus, using Wilcoxon’s Rank Sum Test, a set of hypotheses was required to demonstrate it. The general form of the hypothesis is:

The mean value of the reputation model attribute from ART experiments is higher than the same attribute’s mean value from the FOReART experiments, where higher reputation mean value means a worse detection of the *Dishonest* agent. This hypothesis, from the point of view of the reputation model attribute is expressed mathematically as $Q_{ART}^X > Q_{FOReART}^X$, where X is a L.I.A.R. or Repage reputation model attribute.

In order to validate this hypothesis using the Wilcoxon's Rank Sum Test, the following test is performed:

$$H0 : Q_{ART}^X \leq Q_{FOReART}^X$$

$$H1 : Q_{ART}^X > Q_{FOReART}^X$$

The complete set of hypotheses to demonstrate the effects of the more expressive communication is composed of hypotheses from A to E, where each one is related to a reputation model attribute, respectively, *Dlbrp*, *Iibrp*, *RpRcbRp*, *Image* and *Reputation*.

Table 2. Expressiveness hypotheses result

Pair	Hypotheses				
	A	B	C	D	E
(exp1, exp4)	✓	✓	✗	-	-
(exp2, exp5)	-	-	-	✗	✗
(exp3.1, exp6.1)	✓	✓	✗	✗	✗
(exp3.1, exp6.2)	✓	✓	✗	✗	✗
(exp3.2, exp6.1)	✓	✓	✗	✗	✗
(exp3.2, exp6.2)	✓	✓	✗	✗	✗

When applied to the results of the following pairs of experiments: (exp1, exp4), (exp2, exp5), (exp3.1, exp6.1), (exp3.1, exp6.2), (exp3.2, exp6.1) and (exp3.2, exp6.2), considering the risk level (α) of 0.05 and the degree of freedom of 18, those hypotheses generate the results presented in Table 2 (✗ means that $H0$ was rejected, which confirms the hypothesis; ✓ means that $H0$ was not rejected, thus the hypothesis cannot be confirmed; and – (dash) means that the hypothesis is not applicable for the pair of experiments).

Analyzing the information in Table 2, we can verify that the hypotheses C, D and E reject the $H0$ (indicated by ✗) confirming those hypotheses, while the hypotheses A and B do not (indicated by ✓). From the reputation model point of view, the hypotheses D and E are associated to the Repage reputation model (*Image* and *Reputation* attributes), while the hypotheses A, B and C are associated to the L.I.A.R. reputation model (*Dlbrp*, *Iibrp* and *RpRcbRp* attributes).

Considering that the Repage agents update *Image* and *Reputation* attributes with received information, while the L.I.A.R. agents update only the *RpRcbRp* attribute, we can conclude that a more expressive communication about reputation has a significant statistical impact on agents using both reputation models, since the detection of the *Dishonest* agent is improved on the hypotheses that consider the attributes impacted by communication.

Effect of the reputation model heterogeneity. The analysis of the effect of reputation model heterogeneity was performed by testing if the mean value of the *Dishonest* agent's reputation model attributes (r_j) obtained on experiments with homogeneous

reputation model were higher than the similar ones obtained on mixed experiments. The underlying idea was that if these results were statistically different, then it would mean that heterogeneous environments, composed of agents with different reputations models, would better identify the *Dishonest* agent, since different aspects of the behavior of this latter can be better captured in different reputation model. Thus, to demonstrate it using Wilcoxon's Rank Sum Test a set of hypotheses was required. The general form of the hypotheses is:

The mean value of the reputation model attribute from experiments with homogeneous reputation model is higher than the same attribute's mean value from mixed experiments, where higher reputation mean value means a worse detection of the *Dishonest* agent. This hypothesis, from the point of view of the reputation model attribute is expressed mathematically as $Q_{P/M}^X > Q_{P/Mixed}^X$, where M is the reputation model (L.I.A.R. or Repage), X is its attribute and P is the testbed platform (ART or FORE-ART).

In order to validate this hypothesis using the Wilcoxon's Rank Sum Test, the following test is performed:

$$H0 : Q_{P/M}^X \leq Q_{P/Mixed}^X$$

$$H1 : Q_{P/M}^X > Q_{P/Mixed}^X$$

The complete set of hypotheses to demonstrate the effects of heterogeneous reputation models is composed of hypotheses from F to O. The hypotheses from F to J are related to the ART platform, where each one is related to a reputation model attribute, respectively, *Dlbrp*, *Iibrp*, *RprcbRp*, *Image* and *Reputation*. The hypotheses from K to O are related to the FOREART platform, where each one is related to a reputation model attribute, respectively, *Dlbrp*, *Iibrp*, *RprcbRp*, *Image* and *Reputation*.

When applied to the results of the following pairs of experiments: (exp1, exp3.1), (exp1, exp3.2), (exp2, exp3.1), (exp2, exp3.2), (exp4, exp6.1), (exp4, exp6.2), (exp5, exp6.1) and (exp5, exp6.2), considering the risk level (α) of 0.05 and the degree of freedom of 18, those hypotheses generate the results presented in Tables 3 and 4 (✘ means that $H0$ was rejected, which confirms the hypothesis; ✓ means that $H0$ was not rejected, thus the hypothesis cannot be confirmed; and – (dash) means that the hypothesis is not applicable for the pair of experiments).

Table 3: ART hypotheses result

Pair	Hypotheses				
	F	G	H	I	J
(exp1, exp3.1)	✓	✓	✘	-	-
(exp1, exp3.2)	✓	✓	✘	-	-
(exp2, exp3.1)	-	-	-	✓	✓
(exp2, exp3.2)	-	-	-	✓	✓

Table 4: FOREART hypotheses result

Pair	Hypotheses				
	K	L	M	N	O
(exp4, exp6.1)	✓	✓	✘	-	-
(exp4, exp6.2)	✓	✓	✘	-	-
(exp5, exp6.1)	-	-	-	✓	✓
(exp5, exp6.2)	-	-	-	✓	✓

Analyzing the Tables 3 and 4, we can see that the only hypotheses that rejected $H0$ (indicated by ✘) are related to the *RprcbRp* attribute. Since this attribute is the only one from the L.I.A.R. reputation model affected by the communication, we can conclude

that the reputation model heterogeneity has a significant statistical impact on agents using L.I.A.R. reputation model. Thus, we can presume the existence of some intrinsic characteristics that enable the L.I.A.R. reputation model the use of a more expressive communication on an heterogeneous environment.

5 Conclusions

In this paper we presented the SOARI architecture that extends the general architecture for reputation interoperation and enables the interoperability of reputation among heterogeneous agents. Then, its usage considering RePage and L.I.A.R. reputation models was shown. Moreover, some experiments were performed to answer two questions: (1) is there any improvement in the reputation evaluation accuracy when enabling a more expressive communication? and (2) how does the heterogeneity influence the evaluation accuracy of the dishonest agents' reputation?

The results obtained about a more expressive communication about reputation presented an improvement on the accuracy of the reputation evaluation of other agents. In the other hand, the results of reputation model heterogeneity did not allow us to conclude such improvement. Therefore, the results have shown that the L.I.A.R. reputation model always benefits, which leads us to think that there are some intrinsic or implementation model's characteristics that provided it.

In the future, we intend to perform experiments using more and different reputation models, thus expanding the analysis related to the effects of heterogeneity on the accuracy of reputation evaluation. Based on those results, we expect to have enough information to perform a detailed analysis to identify the relationship between the reputation models characteristics and the benefits of using the SOARI architecture. We also intend to evaluate the possible application of the general approach, i.e. providing agent interoperability using a service-oriented architecture based on ontology translation in different aspects of multi-agent systems other than reputation, for instance, organizational models [6].

Acknowledgements

This project is partially supported by FAPESP/Brazil Process 2008/06356-3. Jaime S. Sichman is partially supported by CNPq/Brazil.

References

1. S. Bechhofer, F. Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein. OWL web ontology language reference. W3C recommendation, W3C, February 2004.
2. S. Boslaugh and D. P. A. Watters. *Statistics in a nutshell*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, 2008.
3. A. A. F. Brandão, L. Vercouter, S. J. Casare, and J. S. Sichman. Extending the art testbed to deal with heterogeneous agent reputation models. In C. Castelfranchi, S. Barber, J. Sabater-Mir, and M. S. (eds.), editors, *Proceedings of the 10th International Workshop on Trust in Agent Societies*, Honolulu, USA, 2007.

4. S. Casare and J. S. Sichman. Towards a functional ontology of reputation. In *AAMAS '05: Proceedings of the 4th Joint Conference on Autonomous Agents and Multiagent Systems*, pages 505–511, Utrecht, The Netherlands, 2005. ACM.
5. R. Conte and M. Paolucci. *Reputation in Artificial Societies. Social Beliefs for Social Order*. Kluwer, Boston, 2002.
6. L. Coutinho, A. A. F. Brandão, J. S. Sichman, and O. Boissier. Model-driven integration of organizational models. In *AOSE '08: Proceedings of the 9th International Workshop on Agent Oriented Software Engineering*, Lisbon, Portugal, 2008.
7. K. K. Fullam, T. B. Klos, G. Muller, J. Sabater-Mir, A. Schlosser, Z. Topol, K. S. Barber, J. S. Rosenschein, L. Vercouter, and M. Voss. A specification of the agent reputation and trust (art) testbed: experimentation and competition for trust in agent societies. In *AAMAS '05: Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 512–518, New York, USA, 2005. ACM.
8. Y. Kalfoglou and M. Schorlemmer. Ontology mapping: the state of the art. *The Knowledge Engineering Review*, 18(1):1–31, 2003.
9. G. Muller and L. Vercouter. L.i.a.r. achieving social control in open and decentralised multi-agent systems. Technical report, École Nationale Supérieure des Mines de Saint-Étienne, Saint-Étienne, France, 2008.
10. L. G. Nardin, A. A. F. Brandão, J. S. Sichman, and L. Vercouter. Soari: A service oriented architecture to support agent reputation models interoperability. pages 292–307, 2008.
11. L. G. Nardin, A. A. F. Brandão, J. S. Sichman, and L. Vercouter. An ontology mapping service to support agent reputation models interoperability. In *Proceedings of the 11th International Workshop on Trust in Agent Societies*, pages 140–144, Estoril, Portugal, 2008.
12. L. G. Nardin, A. A. F. Brandão, J. S. Sichman, and L. Vercouter. A service-oriented architecture to support agent reputation models interoperability. In *Proceedings of the 3rd Workshop on Ontologies and Their Applications*, Salvador, Brasil, 2008.
13. L. G. Nardin, G. Muller, A. A. F. Brandão, J. S. Sichman, and L. Vercouter. Effects of expressiveness and heterogeneity of reputation models in the art testbed: Some preliminary experiments using the soari architecture. In *Proceedings of the 12th International Workshop on Trust in Agent Societies*, Budapest, Hungary, 2009.
14. J. Sabater-Mir, M. Paolucci, and R. Conte. Repage: Reputation and image among limited autonomous partners. *Journal of Artificial Societies and Social Simulation*, 9(2), 2006.
15. J. Sabater-Mir and C. Sierra. Social regret, a reputation model based on social relations. *SIGecom Exch.*, 3(1):44–56, 2002.
16. L. Vercouter, S. J. Casare, J. S. Sichman, and A. Brandão. An experience on reputation models interoperability based on a functional ontology. In *IJCAI '07: Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 617–622, Hyderabad, India, 2007.
17. U. Visser, H. Stuckenschmidt, H. Wache, and T. Vögele. Enabling technologies for interoperability. In U. Visser and H. Pundt, editors, *Proceedings of the Workshop on the 14th International Symposium of Computer Science for Environmental Protection*, pages 35–46, Bonn, 2000. TZI.
18. M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley & Sons, Chichester, England, 2002.
19. G. Zacharia and P. Maes. Trust management through reputation mechanisms. *Journal of Applied Artificial Intelligence*, 14(9):881–907, 2000.