2009 Ninth International Conference on Intelligent Systems Design and Applications

A JADE-based ART-inspired ontology and protocols for handling Trust and Reputation

Javier Carbo
Group of Applied Artificial Intelligence
Computer Science Department
University Carlos III of Madrid
Leganes, Madrid, Spain
Email: javier.carbo@uc3m.es

Jose M. Molina
Group of Applied Artificial Intelligence
Computer Science Department
University Carlos III of Madrid
Colmenarejo, Madrid, Spain
Email: josemanuel.molina@uc3m.es

***Abstract*—Trust and Reputation management play an important role in agent-based Recommender Systems. Although several protocols and ontologies of agents using trust and reputation has been proposed, none of them has been so extensively used and implicitly accepted by research community as those from Agent Reputation and Trust (ART in advane) testbed. The motivation of this adaptation is to facilitate the use of ART principles in real distributed applications instead of a centralized testbed for experimentation. This paper presents an adaptation of the protocols proposed by ART testbed to a codification for the most popular Agent platform: JADE. This implementation follows a coherent API with the FIPA protocols included in JADE distribution for an easy use. We also complement the behaviours of corresponding initiators and responders of the protocols with an ontology formed by a collection of concepts, predicates and agent actions that may represent as the ART application domain as any other service-oriented domain. The proposal has been designed to be applied in domains where multi-agent e-commerce solutions are needed. Future work includes the integration of this ontology and protocols in context-aware scenarios such as an airport.**

***Keywords*-Trust and Reputation Management; Recommending Systems; Agents;**

I. Introduction

The way agents cooperate managing recommendations is a key design factor in Agent based e-commerce solutions. Since often Agent Systems intend to be open, agents have to establish by themselves their own way to apply recommendations into a trust decision. This trust-based selection of partner agents may be defined as a designed or emergent social control [1]. The former, implemented through Electronic Institutions, such as Certification Authorities, takes place when trust is concluded from recommendations of these institutions which are computed by the observation of universal and objective norms. But in many real-world interactions trust is emergent, depends on local and subjective evaluations shared between partner agents (reputation). This last approach is, therefore, our focus of interest: how trust may emerge spontaneously in a distributed way from the interactions between agents. In recent years, research community that works on this issue has grown a lot, and from their members many trust/reputation models have been proposed [2]. But since it was very difficult to compare the respective performances of the different models (as many ad-hoc implementations and metrics have been applied), a testbed platform for agent trust/reputation models was neccesay and it was recently developed: the Agent Reputation and Trust (ART) Testbed [3] [1]. Using such testbed, three international competitions were successfully carried out jointly with the last Autonomous Agents and MultiAgent Systems (AAMAS) International Conferences. During these years the ART testbed has been used by dozens of researchers, and the ART-testbed members have discussed, patched and updated the platform using the feedback from the Competitions (see discussion notes on ART web page) and from the agent trust research community (through the discussion board of ART). These criticism produced some changes in protocols [4], and outlined new directions of work [5]. This wide use of ART just proves that many models (at least as many as participants in competitions) can be adapted to the ontology that ART defined. However this is not a final proof of general acceptance since the goal of ART was to provide a framework to run fair comparisons between different models much more than defining the most extensive and adaptive ontology to allow interoperatibility that could be possible. Therefore, the definition of of an ontology which can handle all all present and future reputation and trust models is still an open issue [6] that ART do not solve, and posibly no static ontology would ever do it. (in fact, some remarkable proposals of ontologies of trust and reputation such as Repage [7] and [8] can not be mapped easily into ART terms). But adapting the ART terms and protocols, that have shown in practice their ability to represent and use dozens of trust and reputation models, is enough to state that they will be a reasonably good ontology to be included in JADE agent applications, even despite the limitations that any ontology would have to map every possible trust model.

This paper proposes an adaptation of the generally ac-

[1] http://www.art-testbed.net/

978-0-7695-3872-3/09 $26.00 © 2009 IEEE
DOI 10.1109/ISDA.2009.48

Figure 1. ART domain outline. Source [3]

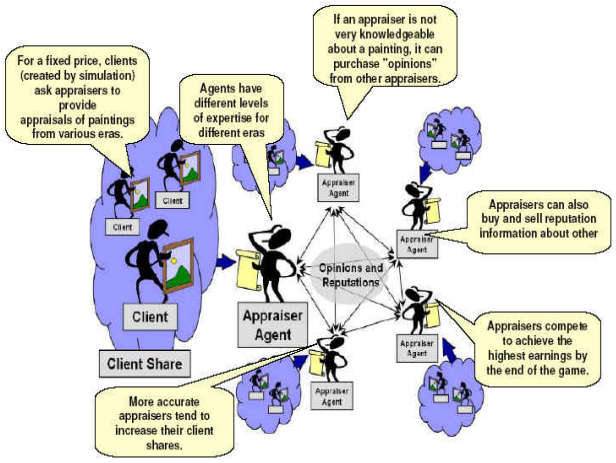

Figure 2. The role of simulation engine in the ART testbed. Source: [11]

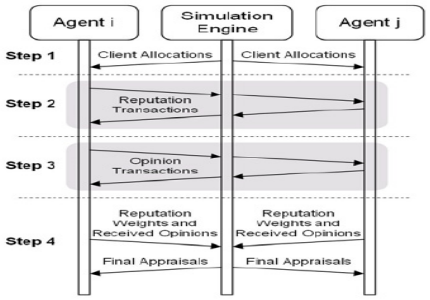

cepted protocols and concepts of ART to an implementation that would facilitate the integration of trust/reputation models in real-world agent applications such as a distributed Recommender System, not just local simulations for experimentation. Although there are already previous ontologies on trust/reputation concepts [7] [8], that could be approximately translated into ART terms and functionality, we state that they are more extensive and theoretical. The intention of moving trust and reputation models from theory to real uses, allowing heterogeneous agents to interoperate about reputation is not new, but previous approaches such as [9] are proposals of ad hoc architecture of agents, while our proposal intends to be neutral in terms of the internal architecture and reasoning about reputation to facilitate a generalized adoption of the ontology and protocols. In order to achieve this goal the corresponding protocols, messages, concepts, actions and predicates of ART testbed were implemented using one of the most extensively used agent platforms: JADE [10], where the centralized approach of the ART testbed is replaced by a truly distributed alternative, which seems to be more appropriate as part of real applications of agents that provide and recommend services. In section 2 we explain the terms and protocols used in the ART testbed. Section 3 presents our adaptation of ART ontology and protocols into a JADE agent system. Finally section 4 shows the conclusions of our work.

II. ART TERMINOLOGY AND PROTOCOLS

A. ART outline

The ART testbed compares different trust/reputation models and strategies in the art appraisal domain (see figure 1). In this domain, the agents are the players/competitors that earn money/points appraising paintings (with the assistance of certain trust strategies). They do not communicate directly (to avoid synchronization problems, collusions and cheats in competitions). While clients that own the paintings to be appraised, instead of being also agents, their existence is emulated by the simulation engine. The simulation engine also rules the execution of the corresponding appraisal ability of the agents (called expertise in ART terms) [3]. In each timestep, the simulator engine presents each appraiser agent with paintings (generated by the simulation engine, see figure 2) to be appraised, paying in advance a fixed fee f for each appraisal request. All the agents start with the same number of paintings to appraise, and therefore with the same initial money/points. Very close valuations of paintings to the real value would lead to more future clients, and therefore to more earnings to win the competition. Each painting belongs to an era among a finite set of possible artistic eras while agents have different levels of expertise (ability to appraise) in each artistic eras. An agent can appraise its own paintings and may request opinions (at a fixed cost) from other appraisers to get its valuation of the painting close to the real value (specially useful in the eras where the agent has low expertise). An agent can act also as provider of appraisals in response to opinion (about paintings) requests from other agents. Additionally, an agent can similarly request reputation information about other appraisers (at a fixed and much lower cost than opinions).

B. Opinion Transaction Protocol

The opinion transaction protocol suffered changes in 2008 competition [4] from the original specification of the testbed [3] (figure 3). The final opinion transaction protocol consists of two independent distinct protocols.

- The first one is dedicated to ask for a certainty value of an agent in a given artistic era (where certainty refers to an assessment of agent expertise in that era). This certainty protocol begins when a requester agent sends a 'certainty request' message to another appraiser agent (noted as provider in advance), identifying the era of the painting to be appraised. Upon receiving a certainty request, if the potential provider is willing to provide the requested information, it responds by sending a 'certainty reply' message with the corresponding cer-

Figure 3. ART Opinion and Reputation Transaction Protocols. (Source: [11][4])

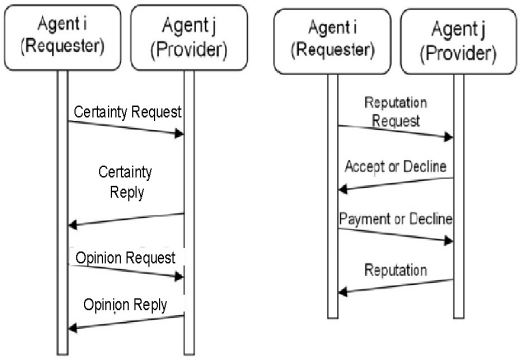

tainty assessment about the opinion it could provide about paintings of that era, defined as a real number between zero and one. If the provider agent did not wish to participate in the requested transaction, it may not decline the certainty request although it might not send any certainty reply at all. The cost of each certainty transaction is a fixed non-returnable amount that is much lower than the cost of asking for an opinion and it is paid when certainty is requested with independence of the trustworthiness of the reply, and even of the existence of a reply.

- The second protocol corresponds to the opinion transaction itself where the opinion requester sends a payment to the provider if it chooses to ask for an opinion from the provider. Upon receipt of payment, in an equivalent way to certainty protocol, the provider is not required to send a truthful opinion about the painting, neither has the provider to send an 'opinion reply' at all. And finally, the provider can not explicitly 'decline' the opinion request.

Since direct communications between agents are not allowed, the simulation engine is in charge of forwarding such certainty and opinion messages, implementing both payments, and producing the opinions according to the corresponding expertise of the agents and to the will to provide a truthful opinion (expressed in a numerical value). However certainty assessment is generated by the agents rather than produced by the simulation engine.

C. Reputation Transaction Protocol

Additionally, appraiser agents can exchange reputations, acting as reputation providers and requesters, following the protocol of figure 3. ART designers consider reputation as the information about the personal view of provider about the expected expertise of other appraiser agents. A reputation transaction begins when a requester sends a 'reputation request' message to a reputation provider, including the agent about whom it is requesting reputation information and furthermore including the era of the expected expertise it is asking for. Then, after the provider agent has received a reputation request, it may send an accept or decline message depending on whether it is willing to provide the requested reputation (or even because it has no information at all about the trustworthiness of the given agent and era). If the provider accepts the transaction, then the requester will send the corresponding payment to the provider in order to receive the requested reputation information. The cost of each reputation transaction is again a fixed amount (and much lower than the opinion cost). Finally the provider is not required to send its actual reputation value, neither has the provider to send any reputation value at all.

III. A JADE-based ART-inspired ontology and protocols

The intention in this contribution is to facilitate the integration of reputation services into wider agent-based application domains. Instead of defining new ontology and protocols, we think that it would be more accepted by the research community reusing a generalized version of ART terms and protocols, since competitions, testing and research carried out with this platform have in some way validated them. Therefore the proposal here is to adapt, implement and to freely distribute a reputation ontology and protocols in a widely used agent platform such as JADE [10] that is compliant with FIPA ACL standard [12] for the communication process.

A. A JADE-based Ontology of ART terms

An ontology allows domain representation in order to be exchanged/shared and re-used by agents. The JADE platform that we chose, classify the elements of an ontology into Predicates, Concepts and AgentActions. So first of all, we have to identify them in ART reputation and opinion transactions. We can see Opinion Transaction as two independent and similar dialogs. In both cases the requester is asking the provider to inform of an Certainty/Opinion assessment according to the descriptors Era and Painting. On its part, Reputation Transaction consists of two sequential and different dialogs. The second one is similar to those of opinion transaction: the requester is asking the provider to inform of a Reputation assessment according to some descriptors: Era, AppraiserAgent. But the first dialog is semantically different. It consists of a requester agent asking the provider about its will to perform some future action (sharing a future reputation assessment). Therefore, we can conclude the next elements of an ART-inspired ontology:

- ART Concepts: Certainty, Painting, Era, Reputation, Appraiser and Opinion.
- ART Predicates: HasCertainty(Who: Appraiser, On: Era, Value: Certainty), HasOpinion(Who: Appraiser,

Figure 4. Extract from the owl codification of ART-JADE ontology

About: Painting, On: Era, Value: Opinion), HasReputation (Who: Appraiser On: Era, About: Appraiser, Value: Reputation)

- ART AgentActions: ShareReputation(Who: Appraiser, On: Era, About: Appraiser)

But we have to generalize them in order to be valid for a set of heterogeneous domains. The resulting ontological high level concepts has to be meta-concepts of the ART domain-dependent terms such as Painting, Era, Appraiser and Certainty. We propose the following generalization: Service instead of Painting, Type of Service instead of Era, jade.core.AID instead of Appraiser and Ability instead of Certainty.

Therefore the final java classes that belong to our generalized ART-inspired ontology are:

- Concepts: Ability, Service, TypeOfService, Reputation, AID and Opinion.
- Predicates: HasAbility(Who: AID, On: TypeOfService, Value: Ability), HasOpinion(Who: AID, ABOUT: Service, ON: TypeOfService, VALUE: Opinion), HasReputation (WHO: AID, ON: TypeOfService, ABOUT: AID, VALUE: Reputation)
- AgentActions: ShareReputation(WHO: AID, ON: TypeOfService, ABOUT: AID)

The OWL codification of this ontology (figure 4) and the implementation in java can be freely downloaded from www.giaa.inf.uc3m.es/miembros/jcarbo/ARTJADE.html

Figure 5. FIPA-compliant ART-inspired Opinion and Reputation Protocols

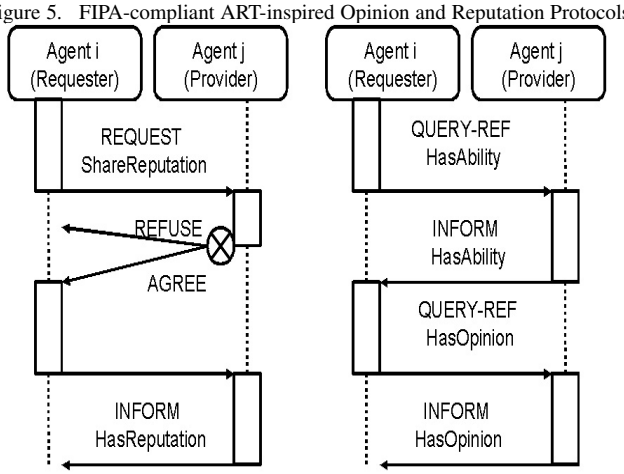

B. a JADE-based version of ART protocols

FIPA also requires the content of each message to have a proper semantics according to a given set of ACL performatives. Therefore we assigned the FIPA ACL performative 'query-ref' to the 'ReputationRequest', 'AbilityRequest' and 'OpinionRequest' messages where the VALUE attribute of predicates 'HasReputation', 'HasAbility' and 'HasOpinion' is a variable which is instantiated by the corresponding responses. These 'ReputationReply', AbilityReply' and 'OpinionReply' messages take therefore the performative 'inform' assigned. Finally the message that ask an agent to perform a 'ShareReputation' agent action has the performative 'request' according to the semantics of such message, while the corresponding message of response can be in the form of an 'AGREE' or 'REFUSE' performatives. Figure 5 shows the FIPA ACL performatives of these messages: These protocols have been implemented as FSM Finite State Behaviours of JADE that combine in a sequential way the corresponding behaviours for tackling the received messages and for sending the corresponding responses.

C. The Payment Problem: A Completely Distributed ART-inspired Agent System

We have then adapted reputation and opinion transactions to a JADE implementation of a FIPA-compliant ontology and protocols, but we have not yet reflected the complete interactions that ART simulates. The communications with the user (possibly also a JADE agent), and the payments (ability, opinion and reputation assessments have a cost, and client allocation of a painting has an income) are lacking in our JADE-based approach. The ART platform solves this problem with the central control of the simulation engine over all the communications, because it emulates the existence of clients and updates the bank balance of all appraiser agents according to the observation of all the messages exchanged (all of them has to pass through the simulation engine). Since

the final intention of this contribution is to facilitate the integration of an ART-inspired reputation service into more complex and fully distributed applications, this problem should not be ignored. The solution that we propose consists of integrating clients as agents, and performing all the payments in a fully distributed mode. This last proposal may be elegantly implemented using a micropayments scheme if we consider that reputation, ability and opinion costs are low enough. For instance MicroMint [13] payment scheme coud be used. The basic idea of Micromint is that using hash functions a broker produces unforgeable coins, then sell these coins to users. A user can use these coins to pay vendors. In return, the vendors will get real money from the broker by redeeming these coins. This restricts the possible expenses of an agent to provide a service to the amount it received in payment for the service. Micromint also imposes that the prefixed and constant opinion ability and reputation costs should be a multiple of the value of a coin. Both limitations are acceptable in this context. In our context the client is who produces 'coins' that sell to appraisal agents, then appraisal agents may use these 'coins' to pay other appraisal agents the corresponding amount of ability, opinion and reputation costs. Afterwards appraisal agents who received these 'coins' in payment of ability, opinion and reputation transactions will get real money from the client redeeming these 'coins' (see the last message of figure 6). Therefore, we have added a new concept called 'Coins' that will be the value of attribute 'BY' in the predicates 'HasOpinion', 'HasAbility' and 'HasReputation'. Additionally providers of services, opinions and reputations of a given painting, should ask for redeeming the Coins they received in payment. There will be then a message of type 'request' with a new AgentAction Pay(WHO: AID, TO: AID, WHAT: Coins). So, at last, we have to include the dialogue with clients in a FIPA-compliant way. Since in ART domain, a client request an agent to perform a particular service (an estimation on the economic value of a painting called 'appraisal'), this message has assigned the FIPA ACL performative 'request' and the content of this message will be the AgentAction Provide(WHO: AID, WHAT: ServiceDescription, ON:Type, BY: Coins). Afterwards, the inform-done message in response to this request will content a predicate 'HasProvided(WHO: AID, WHAT: ServiceDescription, ON: Type, VALUE: ServiceImplementation). And finally, in order to complete the simulation of ART execution cycle, clients will let know agents the quality of provided services (level of success of the appraisals, pointing out the final real value of the paintings). So there will be an additional message where Clients 'inform' agents about the predicate HasQuality(WHAT: ServiceDescription, ON: TypeOfService, VALUE: ServiceEvaluation). Where ServiceDescription, ServiceImplementation and ServiceEvaluation are new concepts that we have aggregated to the ontology instead of the previous Service concept. In figure 6 we can observe this interaction, which includes further communications of type 'HasOpinion' and 'HasReputation' where coins are given in payment of corresponding received Opinions and Reputations. Although MicroMint is just an example (it has some limitations and practical problems), we have shown how a micropayment scheme can successfully match with a reputation-dedicated ontology and protocols to provide a fully distributed reputation service that could implement any reputation model of those who participated in the ART competitions.

Figure 6. FIPA-compliant ART-inspired Appraisal Protocol with Payments

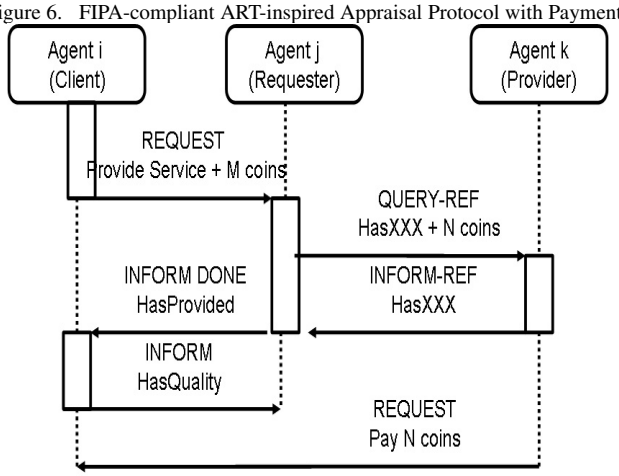

IV. Conclusions

From the point of view of agent-based Recommender Systems, the distributed way trust/reputation emerges from the interaction of agents is a key issue. In this research line, several trust/reputation models were ad hoc tested that leaded to the development of an Agent Reputation and Trust (ART) testbed. Due to the relative success of the ART testbed, its protocols and concepts have a significant acceptation, possibly greater than any of the proposed trust/reputation ontologies. Therefore, this definition effort can also be used for other purposes, such as a good design foundation to spread the inclusion of reputation and trust communications into more general service-oriented systems that would be truly distributed. According to that intention, we have generalized and implemented ART-inspired ontology and protocols using JADE for re-using them in different domains. Specifically, we intend to integrate this ontology and protocols into an airport context aware system of [14]. Furthermore, the inclusion of distributed personalized trust and reputation management into recommendation services is of public and general interest. Our work can also be considered a support and an extension of the ART initiative, since this contribution shows that the particular approach followed by ART designers can be straight-ahead applied and directly integrated in more complex realistic scenarios

(including even payment mechanisms) using the JADE platform and with neutrality respect to the internal architecture of agents. In order to encourage other participants in ART competitons, and any other researcher into adapting their ART models to JADE implementations of agents, we will facilitate the general use of these ART-inspired protocols and ontology in further applications publishing the owl definition and the java code of the ontology and these protocols in our website[2]. Finally our agents that previously participated in the ART competitions are being adapted to this JADE implementation of ART-inspired ontology and protocols. These agents implemented AFRAS reputation model, and Alpha-Beta, Kalman and IMM estimation filters [15].

ACKNOWLEDGEMENTS

This work was supported in part by Projects CICYT TIN2008-06742-C02-02/TSI, CICYT TEC2008-06732-C02-02/TEC, SINPROB, CAM MADRINET S-0505/TIC/0255 and DPS2008-07029-C02-02

REFERENCES

[1] R. Conte and M. Paolucci, *Reputation in Artificial Societies*. Kluwer Academic Publishers, 2002.

[2] S. D. Ramchurn, D. Huynh, and N. R. Jennings, "Trust in multi-agent systems," *Knowl. Eng. Rev.*, vol. 19, no. 1, pp. 1–25, 2004.

[3] K. Fullam, T. Klos, G. Muller, J. Sabater, A. Schlosser, Z. Topol, K. S. Barber, J. 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 *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2005)*, 2005, pp. 512–518.

[4] J. Sabater, M. Gomez, G. Muller, and J. Carbo, "Changes for the 2008 competition," 2008. [Online]. Available: http:megatron.iiia.csic.es/art-testbed/changes_2008.htm

[5] M. Gomez, J. Sabater-Mir, J. Carbo, and G. Muller, "Improving the art testbed, thoughts and reflections," in *Procs. of 12th CAEPIA Conference*, 2007, pp. 1–15.

[6] I. Pinyol, J. Sabater-Mir, and G. Cuní, "How to talk about reputation using a common ontology: From definition to implementation," 2007, pp. 90—101.

[7] J. Sabater-Mir, M. Paolucci, and R. Conte, "Repage: Reputation and image among limited autonomous partners," *Journal of Artificial Societies and Social Simulation*, vol. 9, no. 2, Mar. 2006.

[8] S. Casare and J. Sichman, "Towards a functional ontology of reputation," in *The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2005)*, 2005, pp. 505–511.

[2]www.giaa.inf.uc3m.es/miembros/jcarbo/ARTJADE.html

[9] L. Nardin, A. Brandao, J. Sichman, and L. Vercouter, "A service oriented architecture to support agent reputation models interoperability," in *Trust in Agent Societies. Lecture Notes in Computer Science 5396*, 2008, pp. 292–307.

[10] F. Bellifemine, A. Poggi, and G. Rimassa, "Developing multi-agent systems with a fipa-compliant agent framework," *Software - Practice And Experience*, vol. 31, pp. 103–128, 2001.

[11] K. Fullam, T. Klos, G. Muller, J. Sabater, Z. Topol, K. S. Barber, J. Rosenschein, and L. Vercouter, "The agent reputation and trust (art) testbed architecture," in *Workshop on Trust in Agent Societies at The Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2005)*, 2005, pp. 50–62.

[12] FIPA, "Foundation for intelligent phisical agents 97 specification." [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.55.9502

[13] R. Rivest and A. Shamir, "Payword and micromint: Two simple micropayment schemes," in *Proceedings of the International Workshop on Security Protocols. Lecture Notes In Computer Science 1189*. London, UK: Springer-Verlag, 1997, pp. 69–87.

[14] N. Sanchez, J. Carbo, and J. M. Molina, "Jade/leap agents in an aml domain," in *Hybrid Artificial Intelligence Systems. Lecture Notes in Computer Science 5271*, 2008, pp. 62–69.

[15] J. Carbo, J. Garcia, and J. M. Molina, "Noise detection in agent reputation models using imm filtering," in *Trust in Agent Societies. Lecture Notes in Computer Science 5396*, 2008, pp. 25–42.