BLUE: A REPUTATION-BASED MULTI-AGENT SYSTEM TO SUPPORT C2C IN P2P BLUETOOTH NETWORKS

Gianluca Lax
University Mediterranea of Reggio Calabria
via Graziella, loc. Feo di Vito, 89060 Reggio Calabria - Italy

Giuseppe M. L. Sarnè
University Mediterranea of Reggio Calabria
via Graziella, loc. Feo di Vito, 89060 Reggio Calabria - Italy

Keywords: Mobile Commerce, Reputation-based Systems, Multi Agent Systems, Bluetooth Networks.

Abstract: In this paper we propose a multi-agent system, called Blue, providing a reputation mechanism to promote mobile C2C commerce in large and highly dynamic environments based on the Bluetooth wireless technology. In such a scenarios, detecting malicious users assumes a great importance but traditional approaches to obtaining reputation of a user are unfeasible or meaningless. For this purpose, we propose a centralized reputation mechanism, used in a asynchronous way, to manage users' reputation based on trade feedbacks (given by users) suitably weighted. Agents employ such a reputation mechanism to choose their trade counterparts trying to avoid malicious users. A number of experiments show the effectiveness of our proposal.

1 INTRODUCTION

The emerging wireless technologies and devices, as mobile phones and personal digital assistants, will allow mobile-commerce to become the dominant form of trading (Dholakia, 2004). Currently, a user can exploit his mobile phone (or another wireless device) to connect in Internet using centralized services, offered mainly by mobile communication companies, Wi-Fi and Wi-Max providers, or to make an hotel reservation from a train or to purchase stocks from a bank while he walks or downloads music onto an MP3 player and so on. Vice versa, wireless technologies not relying on a centralized server, as Bluetooth (see Section 2), allow different users to realize information exchanges in a peer-to-peer fashion.

According to the wireless technology improvement in terms of operating range, hardware costs and power consumption, more and more highly dynamic[1] wireless networks, where users can exchange a large amount of data, will be realized. In these networks there are three main issues, well-known from other research contexts (Ramchurn et al., 2004; Resnick et al., 2000), that have to be dealt with. Such issues regard how a user may: (*i*) found interesting information; (*ii*) protect himself against malicious users; (*iii*) guarantee his privacy. Solutions of such problems can be provided both by centralized services and by cooperative and distribute techniques.

In particular, centralized services, as Napster (http://www.napster.com, 2006), require that all users inform a given server about the resources (files) they offer. A user needing a resource, has to query the server that provides him with a list of users able to satisfy his request, along with their reputation judgment provided by other users (a famous example of this approach is eBay (http://www.ebay.com, 2006)). A centralized solution, that works well in a wired environment, fails in our context, mainly for the following reasons: *i)* service availability is not assured, usually for connection problems; *ii)* the access time to the services can be significant in a highly dynamic context; *iii)* the server can become a bottleneck because of the high number of updates it has to manage due to the dynamism of the network; *iv)* the users reported by the server can be unavailable since they are too far w.r.t. the operative range of the mobile device.

A different philosophy is employed in a distribute approach where resources and reputation information must be propagated by the users during their interactions (Mui et al., 2002). Also this approach can be unsuitable in our context for the following reasons: *i)* it presumes that the networks are stable enough both in composition and in time living; *ii)* in wide and dy-

[1]From the point of view of a user, each time another user U enters in (resp., exits from) the operative range of his device, it is viewed as the U's joining to (resp., leaving from) the network.

Lax G. and M. L. Sarnè G. (2006).
BLUE: A REPUTATION-BASED MULTI-AGENT SYSTEM TO SUPPORT C2C IN P2P BLUETOOTH NETWORKS.
In *Proceedings of the International Conference on e-Business*, pages 97-104
DOI: 10.5220/0001426300970104

namic communities the acquired knowledge of information and reputation could be broadly partial and consequently ineffective.

In this paper we propose a framework to promote Bluetooth-based trading activities (file for file or money for file, though they could occur also for free) in a profitable and safe way, independently of telephonic or networking providers. It exploits the opportunities offered by Bluetooth networks, smartphones, agent and e-payment technologies. More in detail, Bluetooth is used to realize P2P communications in highly dynamic networks. To realize fast and sure transactions, a Multi-Agent System, called Blue, is implemented. Moreover, a centralized agency supports agents by managing a reputation system.

Blue adopts temporary agent reputation credentials that are required by each agent to the agency and issued on the basis of suitable criteria. Before each trade activity, agents exchange their credentials to evaluate the counterpart reputation and, then, they decide on the opportunity of continuing (or not) the transaction. Note that such credentials can not be verified at the same time of the transaction by means of the issuer or third parties. A more detailed description of this aspect will be provided in Section 4.

Furthermore, the proposed solution solves many questions arisen from centralized and distribute approaches; more in detail: *i)* centralized services and business transactions happen at different time over different communication channels bypassing problems of services availability; *ii)* the Blue mechanism is quick and independent of composition and life of the network; *iii)* users maintain privacy about transaction contents w.r.t. third parties; *iv)* a reciprocal reputation knowledge between two agents is realized; *v)* all agent communications are inexpensive, supporting also micro-trading activities (in fact, the connection costs are usually incompatible with the effective transaction value or constitute themselves a significant part; in any case agent-server communications must happen over an Internet communication channel); the usage of cryptography and digital certificates improves the security in Bluetooth that currently is an unsafe environment (Shaked and Wool, 2005). On the contrary, the main open question is that asynchronism does not allow users to exploit certificates really updated. For this reason Blue is riskier than other P2P systems. However, uncorrect agents are detected and isolated in a short time.

The rest of the paper is organized as follows: in the next section some preliminary notions are introduced. Sections 3 presents the architecture of Blue, consisting of a centralized agency and a number of agents. The adopted reputation model is described in Section 4. Some experiments and results obtained by simulating our framework are reported in Section 5. The comparison with other approaches and techniques are presented in Section 6, and finally, in Section 7 we draw our conclusions.

2 PRELIMINARY NOTIONS

A general overview is presented now about two issues widely exploited in this paper, which are the Bluetooth technology and the cryptography techniques.

Bluetooth[2] (http://www.bluetooth.com, 2006; Dursch and Yen, 2004) realizes radio-frequency communications for short-range connectivity among devices (as personal digital assistants, mobile phones, laptops, printers and so on) in an inexpensive and low power way; it provides fast transmissions of voice and data, also in noisy radio environments[3], and implements data error correction, cryptography and authentication methods. Bluetooth allows us to provide both point-to-point and point-to-multipoint connectivity. In the first modality, the communication channel is shared only between two Bluetooth peers, while in the other modality a small group of Bluetooth units, called *piconets*, can exploit the communication channel in a given time. In a piconet a device acts as master and the others are slaves synchronized to the master's clock. More piconets over a same physical area can form a larger network, called *scatternet*, where various piconet-to-piconet communications at a time can be realized by employing the respective master units. The current Bluetooth standard (ver. 2) permits a transfer rate of 2.1 Mbps, but the Bluetooth SIG members are examining some technologies, as Ultrawide band (http://www.uwbforum.org, 2006), for an integration with Bluetooth to further improve both the transfer rate and the operative range.

The second issue presented in this section regards the "digital sign" that guarantees authenticity and integrity of a message. It employs a *Public Key Cryptosystem* (PKC) based on a cryptographic technique as RSA (Rivest et al., 1978). A PKC requires two complementary cryptographic keys, usually named "secret" and "public". What it is encrypted with one of the two keys can be decrypted only with the other one and vice versa. A trusted authority assures the public key validity and the identity of its owner by means of a certificate. The public key (K_P) should

[2]The original Bluetooth Promoter Group is constituted by the five companies (Ericsson, IBM, Intel, Nokia and Toshiba) that in 1998 had formed a Special Interest Group (SIG) on this technology and by Microsoft, 3Com, Lucent and Motorola.

[3]The adopted unlicensed working frequency band (2.4 GHz) assures an interaction among Bluetooth units also if this band is shared with other devices signals, as garage door openers. Besides, Bluetooth is compliant with global emissions rules and airline regulations.

be accessible to each user while the secret key (K_S) must be not shared with anyone. Moreover, a hash function, as MD5 (Rivest, 1992) or the Secure Hash Algorithm (SHA) (NIST/NSA, 2002), computing a one way message *digest* is exploited. In order to digitally sign a message M, the user U signs the digest of M ($H(M)$) with his K_S^U obtaining the digital signature $DS = K_S^U(H(M))$. Then the message and the signature are sent to V, who can check both the integrity and the authenticity of M by verifying that $H(M)$ equals $K_P^U(DS)$.

3 THE BLUE FRAMEWORK

In this section we describe the Blue framework, that consists of a centralized *agency* and a number of agents providing same features. Each agent, associated to both a specific user and a SIM, stores and handles both system and user information on one hand and monitors and supports user's activities on the other hand. All agents are affiliated to the agency. This latter provides them with some tools and information and takes care to promote a trust atmosphere. For this purpose, the agency exploits the reputation model described in Section 4. Agents exploit Bluetooth features to realize P2P connections among them in such a way agents can support their users during searching, buying or selling activities.

3.1 The Blue Agency

The *Blue* Agency, denoted by Ag, is a centralized service provider which supports agent activities in order to guarantee the correct behaviour of all affiliated agents and their trading activities. To this goal the agency keeps the following information: *i)* its own secret and public cryptographic keys (resp., K_S^{Ag} and K_P^{Ag}); *ii)* an ontology that consists of one or more XML-schema employed by the agents to describe the resources to sell or to buy; *iii)* the list of affiliated agents; *iv)* for each agent, the value of its reputation (whose computing is described in Section 4), as well as a specific digital certificate, called *(reputation) credential*, used during transactions are stored. The credential C of an agent is a tuple $\langle ID, R, K_P, Ex\rangle$, where ID is an agent identifer[4], R is the agent reputation rating, K_P is the public cryptographic keys of the agent and Ex is the expiration dates of C, respectively. Observe that the temporal validity of C, (depending on Ex) is tightly correlated to R, so credentials of trustworthy agents have longer validity. Moreover, the agency stores also user's financial account information. The agency provides three main activities, more specifically:

- ***affiliate managing*** - When Ag receives a new agent affiliation, it provides the agent with K_P^{Ag}, an initial reputation value (see Section 4) and the current ontology. Then the agency adds the new agent to the list of affiliated agents.
- ***credential providing*** - On demand, the agency supplies an agent with a valid and updated credential, signed by the agency with K_S^{Ag}.
- ***reputation managing*** - Ag manages and updates the reputation rate of agents (see Section 4).

3.2 The Blue Agent

Each agent is associated to a user and supports his activities by managing (in terms of insertion, deletion and updating) the resources he wants to sell or to buy. The agent is automatically activated (resp., deactivated) when the user's device is on (resp., off "per se" or for an explicit user's choice). In order to support user activities, an agent keeps the following information: *i)* the user's secret and public keys; *ii)* the ID of the agency (e.g., the Internet address) and its public key; *iii)* user's financial account information; *iv)* the Blue ontology; *v)* the resources to sell (resp., to buy), described using such an ontology and including also the price of selling (resp., buying).

Observe that each agent has to be affiliated to the agency. The first time an agent is activated, it receives the public key K_P^{Ag} of the agency, the current ontology and an initial credential C (signed by Ag). Moreover, the user (associated to the agent) has to provide some information exploited by the agency during the user identification task, as well as some initial parameters, like an individual hazard threshold (Falcone and Castelfranchi, 2001a; Tan and Thoen, 2000), used to select the counterpart agent for a transaction, and a list of resources to sell and to buy.

Now we describe the activities performed by the agents during trading. Without lost of generality, we consider two agents, named a_b and a_s, the former interested in buying a resource, the latter in selling. The trading activities are:

1. ***research*** - Each agent searches another agent close enough to communicate by a Bluetooth network.
2. ***presentation*** - Then each agent verifies the identity of the other agent, as follows:

[4]To avoid malicious identity change of a compromised reputation with a spotless one, easy in virtual communities but potentially easier in Bluetooth networks, in Blue the SIM of smart-phone is employed as agent ID (i.e., a mobile phone number used to identify the agent) (Pfitzmann et al., 1997) and agent data are stored in a persistent way by the agency. As a consequence each identity change, even though possible, requires necessarily another SIM.

(a) a_b sends to a_s its credential C_b along with its current time now_b;

(b) a_s verifies the integrity of C_b (that is signed by the agency) from which extracts ID_b and the public key of a_b, that is K_P^b. Then a_s sends to a_b its credential C_s, its current time now_s and the signature (by a_s) of a message M_s containing ID_b and now_b.

(c) a_b in its turn, verifies the integrity of C_s, and extracts K_P^s to verify M_s authenticity. Then a_b sends to a_s a signed (by a_s) message M_b containing ID_b and now_b. M_b (resp. M_s) is the proof for a_s (resp. a_b), of the encounter with a_b (resp. a_s), and will be used if the transaction will be performed;

(d) finally, a_s completes the presentation task by verifying the authenticity of M_b.

3. ***evaluation*** - At this point, each agent extracts the reputation rating of the other agent from the credential previously received and will continue the trading only if the counterparts agent is considered trustworthy (i.e., such a reputation rating must be higher than the hazard threshold), otherwise it terminates the communication.

4. ***buy offer*** - a_b sends to a_s a list of resources it wants to buy. These resources are described using the common agent ontology and in order to preserve a_b privacy, only the hash values[5] of each item are reported (Han et al., 2004). Then, a_b receives the list of resources of a_s matching the request of a_b, along with their prices.

5. ***sell offer*** - When a_s receives the buy offer from a_b it compares the hash values of the resources (to sell) it owns with the hash values received from a_s in order to find a matching. Finally, the list containing all the matching resources, along with their price, is sent to a_b. In the case of an empty result, a message for closing the protocol is sent to a_b.

6. ***transaction*** - In this step, the trade can be carried out if there is the matching between offer and demand. At the end of the transaction, a_b signs again the proof of the encounter (M_b) generating $\overline{M_b}$ that is sent to a_s and is the proof that the transaction has occurred. In its turn, a_s performs the same operation.

Observe that in our system the transaction is completely performed *locally*. This is one of the main advantages given by our proposal. On the other hand, for the same reason some fraud can happen since no check can be performed at the same time of the transaction, but it is guaranteed that fraud is detectable, traceable and unprofitable.

[5]For an efficient implementation we may assume that hash values are pre-computed and stored.

4 THE REPUTATION MODEL

In this section we describe the reputation model used in our framework. In the following, according to (Abdul-Rahman and Hailes, 2000), we consider that the reputation is "*an expectation about an agent's behavior based on information about or observations of its past behavior*".

When economical interests are involved, the actors should be trustworthy. To this aim, trade agent activities need of reputation information that can be be based on a direct or, more usually, indirect agent knowledge exploiting in this case some reliable propagation mechanisms where the number of (independent) information sources (Falcone and Castelfranchi, 2001b) and their credibility will be relevant.

In the previous sections we have seen that the agency manages the *reputation rating* R of each agent exploiting the feedbacks provided confidentially by the past agent's trade counterparts. To prevent malicious behaviours, before a trade activity, each agent presents, as a visit card, its temporary *Reputation Credential* described in Section 3.2.

The Blue reputation mechanism should satisfy some properties (Xiong and Liu, 2004; Ramchurn et al., 2004) and more specifically: *i)* taking into account the trade history of each agent; *ii)* differentiating dishonest from honest feedbacks, to avoid malicious reputation manipulations; *iii)* recognizing different transaction contexts, to avoid reputation gain in small transaction value for cheating in high one; *iv)* identifying agents provided of dynamic personality that alternate their behaviours between honesty and cheat, independently of the transaction context.

Observe that, assuming a large agent population and a dynamic environment, the probability of reencounters among agents is very low. As a consequence, agents' reputation based on subjective agents' impressions (direct dimension of reputation) is not significative; in fact, in Blue agent reputation is based only on indirect information aggregated in a centralized way (indirect dimension of reputation). Another consequence is that the presence of collusive agent coalitions, that could maliciously influence reputation evaluations, can be neglected, since collusive agents should encounter effectively their victims (see the observation above).

Thus, the proposed reputation metrics is based only on the following factors: *i)* The value of feedbacks obtained; *ii)* The number of feedbacks obtained; *iii)* The credibility of the feedbacks sources; *iv)* The transaction contexts. To describe the reputation model proposed, we consider a framework consisting of the centralized agency Ag and a set of agents $\{a_1, \ldots, a_n\}$ that, suitably supported by Ag, interact to carry out trading activities focused on the research/sell/buy/exchange of resources in a prof-

itable way over Bluetooth networks.

Each transaction T_i performed by a_i is a tuple $\langle ID_j, val, \overline{M_j}, r_j \rangle$ where: ID_j is the identifier of the trading agent counterpart, val is the monetary value of the transaction describing the transaction context (i.e., its relevance), $\overline{M_j}$ is the proof that an encounter with a_j has occurred and it is carried out during the agent presentation (see Section 3.2); r_j is the appreciation of a_i about the transaction. The value of such an appreciation ranges from 0 (for unsatisfying transactions – for example, when a fraud occurs) to 1 (for full satisfaction). Such a value is explicitly provided by the user at the end of the transaction.

Since each agent credential has an expiration date, periodically an agent has to require a new C to the agency. On this occasion, the agent transfers to the agency the list of transactions done. The agency on the basis of the transaction received, updates the reputation of each agent involved in the transaction. In particular, for each transaction T performed by a_i, assuming a_j be the counterpart of the transaction, then the reputation of a_j (that is R_j) is updated as follows:

$$R_j = (1 - \alpha) \cdot \widetilde{R_j} + \alpha \cdot r_j$$

with $\alpha = R_i \cdot \left(\frac{val}{Val_{max}} \right)$, where the new value of the reputation of a_j is computed weighting in a complementary way, by means of α, two contributions, (1) its previous reputation value (denoted by $\widetilde{R_j}$) and (2) the appreciation about the trade event done by counterpart agent a_i (that is r_j). α allows us to tune the two components. In particular, the weight of the second contribution is considerable whenever the value of the transaction (val) is high or the counterpart agent is authoritative (it has a high reputation R_i).

Observe that we assume the maximum value of a transaction is Val_{max}. It is computed by the agency that stores the value of the transactions performed. Moreover, the initial reputation rating is fixed to 0.5. This choice takes into account two opposite needed: the former is of not penalizing new agents (Ramchurn et al., 2004) while the latter is of penalizing agents having a bad reputation that want reenter into the system (Zacharia and Maes, 2000).

5 EXPERIMENTS

In this section we describe a number of experiments performed in order to test the efficacy of our proposal. For this purpose we implemented a C++ prototype simulating the interactions among (buyer and seller) agents during trading. Now we describe the parameters and the evaluation metrics considered in our experiments.

Number of agents. We considered a population of 1,000 agents.

Number of malicious agents. This parameter varies from 0.1% to 25% of the overall number of agents.

Malicious behaviours. In order to considerate different behaviours of malicious agents, we assumed that a malicious agent behaves incorrectly with a probability MB. In particular, $MB = 1$ means that the agent always misbehaves, whereas $MB = 0.5$ means that each 100 transactions, it behaves well during 50 (randomly chosen) transactions and badly in the remaining ones. We varied MB between 0.1 and 1.

Number of transactions. The number of transactions performed by each agent ranges varies from 10 to 100. Thus, the overall number of transactions varies from 500 to 50000.

Value of transactions. We represented the value of each transaction by a random integer between 1 and 10, thus assuming that the maximum value is 10.

Initialization. We initially set the agent rate to 0.5 and the hazard threshold of each agent to a random value between 0.3 and 0.7.

Evaluation metrics. We have computed several evaluation measures (Rijsbergen, 1979; Srivihok and Sukonmanee, 2005; McLaughlin and Herlocker, 2004), that are *Precision*, *Recall* and *Accuracy*. Before defining such measures, we introduce some notations. Let TP (true positive) be the overall number of *good* transactions taken place, where the word good means that the transaction ends positively. In words, TP represents the number of time the reputation model forecasts correctly that the counterpart agent is not malicious. Analogously, let TN (true negative) be the number of *bad* transactions taken place, let FP (false positive) be the number of good transactions that did not occur (because the counterpart agent reputation was not enough), and finally, let FN (false negative) be the number of bad transactions that did not occur. Observe that TP and FN represent cases in which the reputation model worked well, conversely TN and FP represent incorrect predictions of the reputation model. We are ready to define Precision, Recall and Accuracy. Precision $Pre = \frac{TP}{TP+FP}$ is the percent of positive predictions that are correct and is indicative of the model correctness. Recall $Rec = \frac{TP}{TP+TN}$ is the percent of the positive cases that are caught by the model and denotes the model completeness. Accuracy $Acc = \frac{TP+FN}{TP+TN+FP+FN}$ is the percent of predictions that are correct. Obviously, the higher the value of such measures, the better the accuracy of the system is. Moreover, as a metrics we consider also a system property, named *Malicious Rate*, that is the average reputation rate of malicious agents. We expect it to decrease as the number of transactions increases.

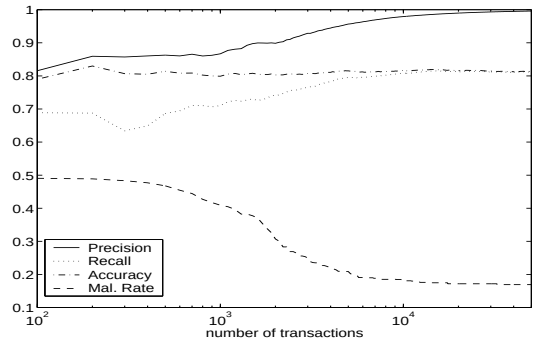

Figure 1: Results vs. number of transactions.

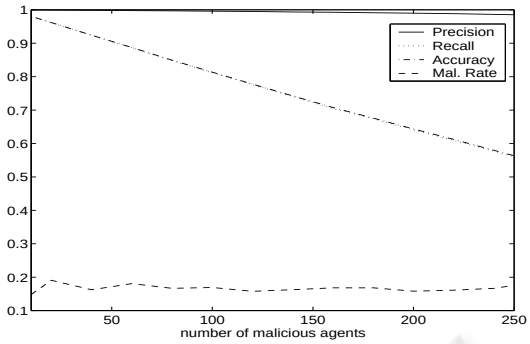

Figure 2: Results vs. number of malicious agents.

5.1 Selected Results of Experiments

In the first experiment the behavior of the system as the number of transactions increases is analyzed. Figure 1 reports the value of Precision, Recall, Accuracy and Malicious Rate obtained by a simulation consisting of 50,000 transactions. We have fixed the number of malicious agents to 10% of the agent population, and we have set the probability of malicious action (MB) to 1. We may observe that the considered metrics assume a stable value after about 10,000 transactions, that is when each agent has performed about 5 transactions. Precision is always quite high (at least 0.8) and converges to 1. It means that when the reputation model suggests performing a trading with an agent, the probability that such a transaction ends positively is very high. The value of Recall ranges from 0.6 to 0.8 in the first 1,000 transactions, then it remains around 0.8. This result shows that a number (about 20%) of true negative cases may occur, and this is the price to pay for obtaining a high Precision. Indeed, the reputation model suggests performing a transaction only if it is reputed *sure* and such a *protective* policy increases the number of good transactions that are not suggested (true negative cases) by the model. Also Accuracy is always quite high (around 0.8). As observed above, we remark that a lot of wrong predictions regard true negative cases that are not too penalizing for users.

In Figure 2 we show the value of the adopted metrics after a simulation of 50,000 transactions, conducted varying the number of malicious agents. We note that Malicious Rate is always very low (less than 0.2) and is independent of the tested parameter. Thus our reputation model works well in the case of both small and large number of malicious agents, as shown in all experiments conducted. Precision decreases slightly as the number of malicious agents increases, anyway maintaining a high value (always more than 0.95). Finally we observe that the measures of Recall and Accuracy are always very similar and decrease as the considered parameter increases. Indeed, as the number of malicious increases, the average of the reputation rate of all agents decreases. As a consequence, the number of true negative cases increases and this reduces Recall and consequently Accuracy.

In the last experiment we studied the performance of the reputation model as the probability of malicious action varies. For space limitation, we report only the results obtained for Malicious Rate. However, similar considerations may be done also for Precision, Recall and Accuracy. In Figure 3 the value of Malicious Rate versus the probability of malicious action is reported. We recall that Malicious Rate is a system property computed as the average reputation rate of malicious agents, thus it ranges from 0 to 1. We have considered three different times of the simulation: after 500 transaction (short time), after 5,000 transactions (medium time) and after 50,000 transactions (long time). The first result we may observe is that for long time, Malicious Rate is always very low (about 0.2), whereas for short time the reputation model does not produce good results. This may be explained considering that for short time each agent has performed (on the average) 1 transaction that is not enough for evaluation its behaviour. Moreover, malicious agents are more difficult to detect when their probability of malicious behaviour is low, as shown in Figure 3. For medium time, we note that the capability of the reputation model to detect malicious agents is almost linearly depending on the probability of malicious behaviour. However, the more interesting result is that our reputation model is able to detect (after a adequate number of transactions) malicious agents even though they alternate between good and bad transactions.

6 RELATED WORK

The relevance of reputation issues is witness by the rich literature produced in these latter years. In partic-

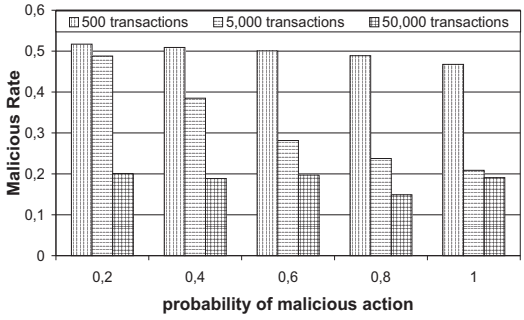

Figure 3: Malicious rate vs. probability of malicious action.

ular, different notions of reputation belonging to different disciplines are presented in (Mui et al., 2002), while different reputation properties and models are described in (Ramchurn et al., 2004). An interesting approach is proposed by (Dellarocas, 2003), where reputation is dealt with game and economic theories.

A well know reputation system is employed in eBay (http://www.ebay.com, 2006); it has been deeply investigate and doubtless the main advantage is its simplicity. On the contrary, it is sensitive to many malicious behaviours. In the Blue multi-agent context, the reputation is certified (as in (Huynh et al., 2004)) and limits the effects of malicious elements. Differently from (Sen and Sajja, 2002), Blue does not suffer the presence of liars even when a boolean rate is adopted.

In accord with Sporas (Zacharia and Maes, 2000), Blue disincentives the change of identity, but in Sporas a very low initial reputation rate (differently from Blue) is used and in this way it is hardest for a new agent to gain reputation. Moreover, Sporas try to avoid collusive agent alliances for increasing reciprocally their reputation rates, by limiting the number of times an agent may increase the reputation of another agent. A study of the dynamics of honesty and dishonesty behaviours in a semi-competitive multi-agent environment where agents can have incentives to be honest or dishonest is presented in (Lam and Leung, 2005).

Reputation models are more trustworthy if more agents as possible cooperate to provide their evaluations (Birk, 2000). Two approaches exist, in the first the agents are free to not provide their feedbacks (positive reputation system), in the other the agents are penalized (or likely promoted) if they do not provide their feedbacks (negative reputation systems). The second approach is needed in Blue, but in usual C2C contexts the first approach seems to be more effectiveness (Yamamoto et al., 2004) where, using an iterate prisoner's dilemma approach, have been investigate different strategies.

In REGRET (Sabater and Sierra, 2001) the agent reputation is computed aggregating (social dimension) the *impressions* that agents (altruistic and cooperative) obtain by direct interactions (individual dimension) weighting the terms of a common semantic (ontological dimension) in accord with their personal point of views. By knowing these weights, it is possible for an agent to uniform the rates provided by other agents to its point of view. A similar approach could be applied in a future version.

In the field of P2P many works have explored reputation issues. Usually, a great attention is given to preserve the correctness of the reputation rates from the effects of the malicious agents by realizing a mix of different techniques. All the following cited works implement frameworks where agents do not have particular difficulty to contact other agents or centralized services for knowing reputation information, differently from Blue.

In particular, a safe approach is reported by (Kamvar et al., 2003) that uses pretrusted peers to minimize the influence of malicious agents in collusion activities. It implements a mechanism taking into account the entire system history. A well formed reputation models is PeerTrust (Xiong and Liu, 2004), where it is implemented a reputation-based trust framework in an adaptive manner able to solve many issues of a P2P system. Some aspects of this work, also if they are referred to a different scenario, are similar to Blue. Finally, in (Damiani et al., 2002) a reputation-based protocol is proposed to support and preserve anonymous and secure services, for choosing reliable resources in P2P networks.

7 CONCLUSIONS

In this paper we have presented Blue, a framework to support C2C commerce activities over Bluetooth networks. The adopted reputation model allows the users involved in trading to detect possible malicious users. It exploits temporary agent reputation credentials that are managed by a centralized agency. A number of experiments show the effectiveness of our proposal in finding malicious users and avoiding frauds. This work has opened many research directions to explore the different possibilities of the Bluetooth networks in a P2P networks. Currently, a real prototype is in an advanced stage of implementation.

REFERENCES

Abdul-Rahman, A. and Hailes, S. (2000). Supporting trust in virtual communities. In *HICSS '00: Proc. of the 33rd Hawaii Int. Conf. on System Sciences - Vol.*

6, pages 1769–1777, Washington, DC, USA. IEEE Comp. Soc.

Birk, A. (2000). Boosting cooperation by evolving trust. *Applied Artificial Intelligence*, 14(8):769–784.

Damiani, E., di Vimercati, D. C., Paraboschi, S., Samarati, P., and Violante, F. (2002). A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *CCS '02: Proc. of the 9th ACM Conf. on Computer and communications security*, pages 207–216, New York, NY, USA. ACM Press.

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424.

Dholakia, R. R. (2004). Editorial: electronic markets in the post-euphoric phase: Relationships, values and behaviors. *Telematic Information*, 21(2):115–121.

Dursch, A. and Yen, D. C. (2004). Bluetooth technology: An exploratory study of the analysis and implementation frameworks. *Computer Standards and Interfaces*, 26(4):263–277.

Falcone, R. and Castelfranchi, C. (2001a). *Social trust: a cognitive approach*. Kluwer Academic Publishers, Norwell, MA, USA.

Falcone, R. and Castelfranchi, C. (2001b). The socio-cognitive dynamics of trust: Does trust create trust? In *Proc. of the Workshop on Deception, Fraud, and Trust in Agent Societies held during the Autonomous Agents Conf.*, pages 55–72, London, UK. Springer-Verlag.

Han, P., Xie, B., Yang, F., and Shen, R. (2004). A scalable p2p recommender system based on distributed collaborative filtering. *Expert Systems with Application*, 27(2):203–210.

http://www.bluetooth.com (2006).

http://www.ebay.com (2006).

http://www.napster.com (2006).

http://www.uwbforum.org (2006).

Huynh, T. D., Jennings, N. R., and Shadbolt, N. R. (2004). Fire: An integrated trust and reputation model for open multi-agent systems. In *ECAI 2004: Proc. of the 16th Europ. Conf. on Artificial intelligence*, pages 18–22. IOS Press.

Kamvar, S. D., Schlosser, M. T., and Garcia-Molina, H. (2003). The eigentrust algorithm for reputation management in p2p networks. In *WWW '03: Proc. of the 12th Int. Conf. on World Wide Web*, pages 640–651, New York, NY, USA. ACM Press.

Lam, K. M. and Leung, H. F. (2005). A trust/honesty model in multiagent semi-competitive environments. In *PRIMA 2004: Proc. of the 17th Pacific rim Int. Workshop*, volume Lecture Notes in Computer Science Vol. 3371, pages 128–147. Springer.

McLaughlin, M. R. and Herlocker, J. L. (2004). A collaborative filtering algorithm and evaluation metric that accurately model the user experience. In *SIGIR '04: Proc. of the 27th Int. Conf. on Research and development in information retrieval*, pages 329–336. ACM Press.

Mui, L., Mohtashemi, M., and Halberstadt, A. (2002). Notions of reputation in multi-agents systems: a review. In *AAMAS '02: Proc. of the first Int. Joint Conf. on Autonomous agents and multiagent systems*, pages 280–287, New York, NY, USA. ACM Press.

NIST/NSA (2002). Fips 180-2. secure hash standard (shs). NIST/NSA.

Pfitzmann, A., Pfitzmann, B., Schunter, M., and Waidner, M. (1997). Trusting mobile user devices and security modules. *Computer*, 30(2):61–68.

Ramchurn, S. D., Huynh, D., and Jennings, N. R. (2004). Trust in multi-agent systems. *Knowl. Eng. Rev.*, 19(1):1–25.

Resnick, P., Kuwabara, K., Zeckhauser, R., and Friedman, E. (2000). Reputation systems. *Commun. ACM*, 43(12):45–48.

Rijsbergen, C. J. V. (1979). *Information Retrieval*. Butterworth.

Rivest, R. L. (1992). The md5 message-digest algorithm. Network Working Group and RSA Data Security, Inc.

Rivest, R. L., Shamir, A., and Adleman, L. (1978). A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126.

Sabater, J. and Sierra, C. (2001). Regret: reputation in gregarious societies. In *AGENTS '01: Proc. of the 5th Int. Conf. on Autonomous agents*, pages 194–195, New York, NY, USA. ACM Press.

Sen, S. and Sajja, N. (2002). Robustness of reputation-based trust: boolean case. In *AAMAS '02: Proc. of the 1st Int. Joint Conf. on Autonomous agents and multiagent systems*, pages 288–293, New York, NY, USA. ACM Press.

Shaked, Y. and Wool, A. (2005). Cracking the bluetooth pin. In *MobiSys '05: Proc. of the 3rd Int. Conf. on Mobile systems, applications, and services*, pages 39–50, New York, NY, USA. ACM Press.

Srivihok, A. and Sukonmanee, P. (2005). E-commerce intelligent agent: personalization travel support agent using q learning. In *ICEC '05: Proc. of the 7th Int. Conf. on Electronic commerce*, pages 287–292. ACM Press.

Tan, Y.-H. and Thoen, W. (2000). An outline of a trust model for electronic commerce. *Applied Artificial Intelligence*, 14(8):849–862.

Xiong, L. and Liu, L. (2004). Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):843–857.

Yamamoto, H., Ishida, K., and Ohta, T. (2004). Modeling reputation management system on online c2c market. *Comput. Math. Organ. Theory*, 10(2):165–178.

Zacharia, G. and Maes, P. (2000). Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(9):881–907.