REPUTATION MANAGEMENT SERVICE FOR PEER-TO-PEER ENTERPRISE ARCHITECTURES

M. Amoretti, M. Bisi, M. C. Laghi, F. Zanichelli, G. Conte
Distributed Systems Group - Dip. di Ingegneria dell'Informazione
Università degli Studi di Parma
amoretti, laghi, zanichelli, conte@ce.unipr.it
matteo.bisi@sinfopragma.it

Keywords: Service Oriented Architectures, Peer-to-Peer, Security, Reputation Management.

Abstract: The high potential of P2P infrastructures for enterprise services and applications both on the intranet (*e.g.* project workgroups) and on the Internet (*e.g.* B2B exchange) can be fully achieved provided that robust trust and security management systems are made available.
This paper presents the reputation system we have devised for our P2P framework which supports secure role-based peergroups and service interactions. Our system includes decentralized trust and security management able to cope with several threats. In the paper the underlying analytical model is presented together with a simulation-based evaluation of the robustness against malicious negative feedbacks.

1 INTRODUCTION

Ubiquitous access to networks is deeply changing the ways enterprises organize and perform their business both internally and externally. Intranets and the global Internet allow for seamless and almost instantaneous information and knowledge sharing within organizations thus enabling more efficient processes and activities and giving rise to novel forms of interaction and supporting applications.

Peer-to-peer (P2P) technologies have gained world-wide popularity due to the success of file-sharing applications (and the reactions of several copyright holders) and their decentralized nature appears promising also to the purposes and applications of enterprises. P2P-based instant messaging and file-sharing can be effectively exploited on an intranet to support for example projects workgroups, distributed offices and distributions chains for documents and archives. Internet-enabled inter-firm collaboration can benefit from a P2P approach as well. Business-to-business exchanges are becoming increasingly important and many B2B communities organize themselves to be more competitive in specialized industry sectors and to increase the efficiency of their procurement and supply chains. By leveraging upon P2P technologies, the common tasks of searching for new business partners and exchanging transaction information (*e.g.* quotations) can be improved in terms of instant information, control over shared data (mantained at each P2P node) and reduced infrastructure costs.

The vision of unmediated, instantaneous trading as well as more realistic P2P-based B2B communities can be approached only if enterprise-level solutions are made available to cope with the fundamental trust and security issues. While *identity trust*, namely the belief that an entity is what it claims to be, can be assessed by means of an authentication scheme such as X.509 digital identity certificates, *provision trust*, that is the relying party's trust in a service or resource provider, appears more critical as users require protection from malicious or unreliable service providers. Unlike B2B exchanges based on centralized, third party UDDI directories which offer trustworthy data of potential trading partners (*i.e.* service providers), P2P decentralized interaction lends itself to trust and reputation systems mainly based on first hand experience and second-hand referrals. This information can be combined by a peer into an overall rating or reputation value for a service provider and should influence further interaction with it.

In this paper we present the reputation system we have devised for our P2P framework which supports secure role-based peergroups and service interactions. Our system includes decentralized trust and security management able to cope with several threats, starting from *impersonification*, which refers to the threat caused by a malicious peer posing as another in order

to misuse that peer's privileges and reputation. Digital signatures and message authentication are typical solutions for this kind of attack. As malicious peers can engage in *fraudolent actions*, such as advertising false resources or services and not fulfilling commitments, a consistent reputation management system has been introdued in our P2P framework which also forbids *trust misrepresentation* attempts. In a peer-to-peer system, the most difficult threat to discover and neutralize is *collusion*, which refers to a group of malicious peers working in concert to actively subvert the system. To face this danger, the default policy provided by our security framework is *role-based group membership based on secure credentials (SC policy)*.

The paper is organized as follows. Next section 2 outlines the issues of reputation management systems and the choices available for centralized and decentralized implementation. The analytical model underlying our reputation management is described in section 3. A simulation scenario is then presented, first describing a four roles configuration example (section 4) and then discussing the obtained results (section 5). Section 6 reports on some relevant work in the area of reputations systems for P2P systems. Finally, a few conclusive remarks and an indication of further work conclude the paper.

2 REPUTATION MANAGEMENT IN ROLE-BASED PEERGROUPS

In our view, a role-based peergroup can achieve stability only if each participant (which is supposed to be authenticated and authorized) bases its actions on previous experience and/or recommendations, *i.e.* which define the reputation of the other participants. Reputation and trust are orthogonal concepts, which require, in a peer-to-peer context, complex management mechanisms such as node identification and digitally signed certificates exchange (Ye et al., 2004).

Our model considers both *peer reputation* and *service reputation*. Peer reputation is important for the reputation manager, which could store service evaluations weighted by the advisor's reputation. Service reputation is important because a peer which has to choose between two apparently identical services, selects the most reputed one. A peer can provide more than one service, each of them with its own reputation which contributes to the overall reputation of the peer.

Peer reputation values are expressed by $(+,-)$ couples, *e.g.* $(12,4)$ which means that peer services have globally received 12 appreciations and 4 negative feedbacks. Peer reputations could have non-zero initial value, fixed by a trusted third party. After each remote service interaction, consumers evaluate the service, assigning $+1$ or -1 to the provider. The reputation manager stores the global reputation of each peer, but also the reputation of each service a peer provides.

When a peer enters the group as newbie, or is promoted to an higher rank, it receives an initial reputation value; based on peer's behaviour, the reputation value changes with time, and represents the trustworthiness of peers on the basis of their transaction with other nodes. Each peer, at the end of an interaction with another member of the group, can provide its feedback about each consumed service; the feedback is used to update the reputation of the provider peer. The reputation manager should weight the received feedback, considering the reputation value of the sender peer.

Two problems arise: where to store the reputation information, and how to guarantee its integrity. Several solutions can be adopted:

1. a stable and recognized peer stores and manages the reputation information of all group members (*centralized* solution);
2. each peer stores its experience against other peers, and when others ask for reputation information of a particular peer, it answers them based on its stored information (*local* solution);
3. the reputation storage is partitioned into several small parts, which are stored in all peers; that is, every peer equally manages some part of the whole reputation information (*global* solution);
4. only stable, recognized and highly-reputed peers are reputation collectors (*mediated* solution).

Not all these solutions are equally scalable, efficient and robust.

Solution 1 is easy to implement as a Centralized Reputation Management Service (CRMS), but it does not scale, *i.e.* it could work only for small peergroups. Solutions 2, 3 and 4 require a Distributed Reputation Management Service (DRMS).

The local solution is slightly efficient and lacks of robustness. If a peer wants more objective reputation about another peer, it should ask many peers. This would generate a lot of messages in the peer-to-peer network. Moreover, if the reputation information is concentrated in few very active customer peers, when these are not online the reputation system is broken.

The global solution is very attractive, in particular if the reputation management system is implemented as a Distributed Hash Table (DHT). In that case, the peer which is responsible for a specific reputation information is determined with a hash function within $O(1)$ time, and its location is found within $O(\log N)$ time.

The mediated solution should be good for unstructured networks, in particular if they have few highly connected and stable nodes, *e.g.* scale-free topologies (Barabási and Albert, 1999). In our role-based

scheme, *admin* peers are the candidates for the realization of the DRMS.

3 ANALYTICAL MODEL

In this section we illustrate an analytical model which describes the fundamental parameters which are involved in the evolution of the reputation, for each role which can be taken in a peergroup based on our SC policy.

The overall reputation value is $Rep = n_+ - n_-$, *i.e.* the difference between the number of positive feedbacks and the number of negative feedbacks. When a peer joins the secure peergroup, and each time it is promoted or degraded, it receives an initial reputation value which is stored by the reputatation management service. The temporal evolution of the reputation value depends on the following parameters:

- *Total votes* $\mathrm{T} = n_+ + n_-$, which represents the sum of all received feedbacks;
- *good Ratio* $\mathrm{R} = \frac{n_+}{n_+ + n_-}$, which is the number of positive feedbacks, versus T; R is a fundamental parameter for the analytical model, because its instantaneous value allows to define the *dependability degree* of the peer.

The dominium of R and T can be easily obtained from their definitions:

$$0 \leq R \leq 1, T \geq 0$$

Depending on values of R and T, we can consider four different conditions for each role a peer can take, which are listed below.

- $T \geq T_{th}$, where $T_{th} \geq 0$ is the *confidence threshold*, evaluated on all received feedbacks. The reason of considering such a threshold is that the dependability of the reputation value of a peer depends on the total number of performed transactions (the more they are, the more the value is dependable).
- $T < T_{th}$ means that the peer has recently joined the group, thus the reputation value must be weighted to consider a potentially less dependability.
- $R \geq R_{th}$, where $0 \leq R_{th} \leq 1$ is the *trust threshold*. In this case, the peer can be trusted.
- $R < R_{th}$, on the other side, means that the peer cannot be trusted and should be degraded.

From the definitions of T and R, we obtain

$$Rep = (2R - 1)T \tag{1}$$

from which we can derive the *role preservation condition*

$$Rep \geq (2R_{th} - 1)T \tag{2}$$

Moreover, by retrieving T and R values of a peer, it is possible to compute the probability that next feedback will be positive (P_{good}) for that peer. Figure 1 illustrates P_{good} versus the good ratio value R. We can observe that if R is greater or equal than the average value R_{avg}, which is an assigned parameter, P_{good}'s value is constant (P^*). If the good ratio is included between R_{avg} and R_{th}, a mechanism named *recovery window* allows the peer to increase its good ratio. In details, a probability bonus P_b is conceded, defined by:

$$P_b = \frac{R_{avg} - R}{R_{avg} - R_{th}}[P_{good}(R_{th}) - P^*]$$

The value of $P_{good}(R_{th})$ is an assigned paramater.

This mechanism is general, but the recovery window size should be different for each role, decreasing with the importance of the peer.

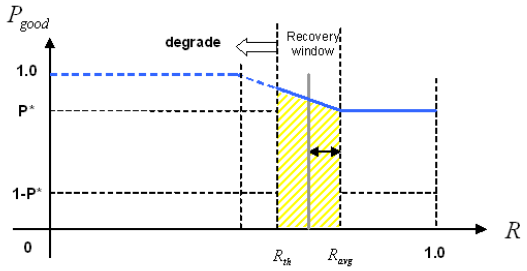

Figure 1: Relation between the probability for next feedback to be positive and the good ratio.

It is possible to compute, from previous parameters, the maximum rate $max\{f_{bad}\}$ with which one or more malicious peers (maybe cooperating) can provide negative feedbacks without affecting the peer's role. Over this rate value, the peer is degraded for insufficient good ratio, according to the rules which have been illustrated above.

We first consider the case of null recovery window, *i.e.* $P_{good} = P^*$ for all $R \geq R_{th}$. In the time unit Δt, the average number of received feedbacks is defined as

$$\Delta T = \Delta n_+ + \Delta n_- + \Delta n_{bad}$$

where

$$\Delta n_+ = P^* \Delta n$$

and

$$\Delta n_- = (1 - P^*)\Delta n$$

are justified feedbacks, while Δn_{bad} represents the average number of unjustified negative feedbacks, sent by malicious peers. The reputation changes by

$$\Delta Rep = \Delta n_+ - \Delta n_- - \Delta n_{bad} \tag{3}$$

Suppose the peer is at the trust threshold, thus (from eq. 2)

$$Rep = (2R_{th} - 1)T$$

Next Δt provides

$$\begin{aligned} \Delta Rep &= (2R_{th} - 1)\Delta T \\ &= (2R_{th} - 1)(\Delta n_+ + \Delta n_- + \Delta n_{bad}) \end{aligned}$$

With the latter and with eq. 3, we can obtain the maximum value of Δn_{bad} over which the peer is degraded, and then

$$max\{f_{bad}\} = \frac{\Delta n_{bad}}{\Delta T} \quad (4)$$

If there is the recovery window, the average number of positive feedbacks per time unit is

$$\Delta n_+ = (P^* + P_b)\Delta n$$

thus the maximum value of Δn_{bad} over which the peer is degraded is not constant but depends on R.

In the following section, we show how this model can be applied to a real system.

4 FOUR-ROLE CONFIGURATION EXAMPLE

In this section we consider a four-role system, and for each role we set the initial values for the parameters we illustrated in section 3. In details, the list of roles is:

- admin - the peer is highly-reputed, and trusted by the group founder, or it is the group founder itself; the actions it is allowed to perform are: service sharing/discovery, group monitoring, voting for changing member ranks, store reputation information (if the mediated solution is adopted);
- newbie - the peer is a new member; it only can search for an admin peer, to ask for a promotion;
- searcher - the peer is allowed to search for services and to interact with them;
- publisher - the peer can search for services but also publish its own services in the peergroup.

Each **admin** peer has the following initial configuration:

$$\begin{cases} n_{+,init} = 50 \\ n_{-,init} = 10 \\ T_{init} = n_{+,init} + n_{-,init} = 60 \\ R_{init} = \dfrac{n_{+,init}}{n_{+,init} + n_{-,init}} = 0.8\overline{3} \\ Rep_{init} = n_{+,init} - n_{-,init} = 40 \\ R_{avg} = 0.8 \\ R_{th} = 0.6 \end{cases}$$

Figure 2 illustrates the probability P_{good} for next feedback to be positive, versus the good ratio value R of an admin peer. Assuming that $\Delta n = 12$, we have $\Delta n_+ = 9.6$ and $\Delta n_- = 2.4$. Without recovery window, the maximum value of malicious negative feedbacks per time unit, over which the peer is degraded, is $\Delta n_{bad} = 4$. Thus, the maximum malicious feedbacks rate which can be accepted is $max\{f_{bad}\} = 25\%$. If there is the recovery window, Δn_{bad} depends on R. In particular, considering the worst case $R = R_{th} = 0.6$, we obtain $max\{f_{bad}\} = 40\%$.

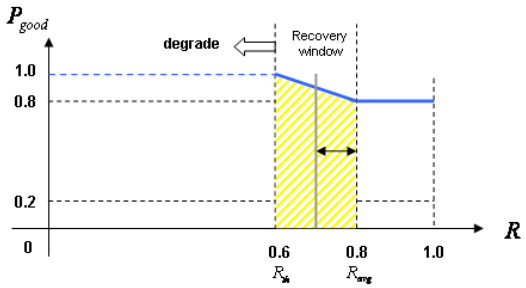

Figure 2: Relation between the probability for next feedback to be positive and the good ratio.

Each **publisher** peer has the following initial configuration:

$$\begin{cases} n_{+,init} = 35 \\ n_{-,init} = 10 \\ T_{init} = n_{+,init} + n_{-,init} = 45 \\ R_{init} = \dfrac{n_{+,init}}{n_{+,init} + n_{-,init}} = 0.\overline{7} \\ Rep_{init} = n_{+,init} - n_{-,init} = 25 \\ R_{avg} = 0.75 \\ R_{th} = 0.6 \end{cases}$$

Figure 3 illustrates that, compared with previous case, the recovery window for a publisher is smaller: 0.15 versus 0.2. With this window, the maximum tolerable rate of malicious negative feedbacks is $max\{f_{bad}\} = 33.3\%$. Without the recovery windows, it would be 20%.

Each **searcher** peer has the following initial configuration:

$$\begin{cases} n_{+,init} = 25 \\ n_{-,init} = 10 \\ T_{init} = n_{+,init} + n_{-,init} = 35 \\ R_{init} = \dfrac{n_{+,init}}{n_{+,init} + n_{-,init}} = 0.714 \\ Rep_{init} = n_{+,init} - n_{-,init} = 15 \\ R_{avg} = 0.7 \\ R_{th} = 0.6 \end{cases}$$

For a searcher, whose recovery window is 0.1 large (see figure 4), the maximum tolerable rate of malicious negative feedbacks is $max\{f_{bad}\} = 25\%$. It would be 14.3%, without the recovery window.

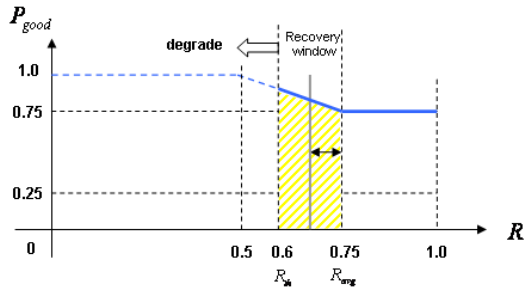

Figure 3: Relation between the probability for next feedback to be positive and the good ratio, in the case of a publisher peer.

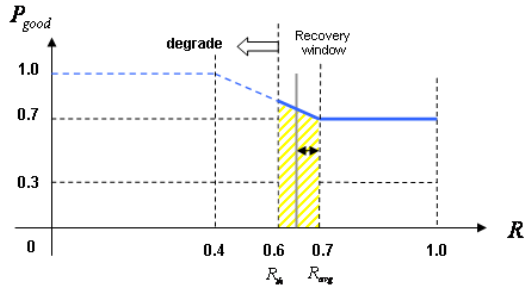

Figure 4: Relation between the probability for next feedback to be positive and the good ratio, in the case of a searcher peer.

Finally, each **newbie** peer has the following initial configuration:

$$\begin{cases} n_{+,init} = 15 \\ n_{-,init} = 10 \\ T_{init} = n_{+,init} + n_{-,init} = 25 \\ R_{init} = \dfrac{n_{+,init}}{n_{+,init} + n_{-,init}} = 0.6 \\ Rep_{init} = n_{+,init} - n_{-,init} = 5 \\ R_{avg} = 0.6 \\ R_{th} = 0.52 \end{cases}$$

Thus a newbie, which is characterized by the smallest recovery window (0.08), has $max\{f_{bad}\} = 23\%$, which would be 13% without the recovery window.

5 SIMULATION RESULTS

By means of SP2A (Amoretti et al., 2005), our middleware for the development and deployment of service-oriented peer-to-peer architectures, we realized a centralized reputation management service, and

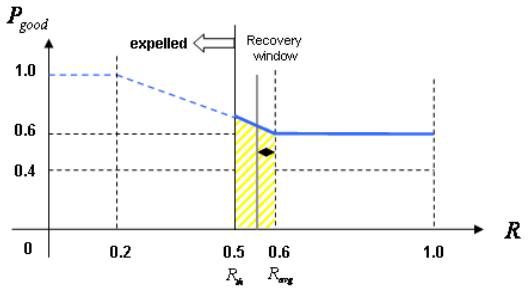

Figure 5: Relation between the probability for next feedback to be positive and the good ratio, in the case of a newbie peer.

we simulated the interaction of that service with an hypothetical network of peers which provide positive and negative feedbacks. We emphasize that our purpose was not to investigate reputation retrieval and maintainance performance, which is obviously different from centralized to distributed services. The deployed testbed allowed us to verify the correctness of the proposed analytical model, and to tune the paramater values for the four-role secure group configuration.

The reputation management service maintains a reputation table, and randomly assigns feedbacks to peers, with $\Delta n = 12$ as assumed in section 4. All simulations started with each peer having $R = R_{avg}$, and lasted the time necessary to observe significant results (we set $\Delta t = 1$ minute). We tracked the evolution of the reputation value Rep for each peer, and we computed the average behaviour for each role. We firstly simulated a peergroup of righteous peers, in which positive and negative feedbacks per unit time are distributed according to R_{avg}. Then we performed several simulations of a system in which some peers provided unjustified negative feedbacks, with increasing rate. For each role, we found the maximum tolerable rate of malicious negative feedbacks, over which the target peer is degraded, or banned from the peergroup if its role is newbie.

In general, to know if the role preservation condition $Rep \geq (2R_{th} - 1)T$ will be fulfilled, in a stable condition with fixed f_{bad} and R_{avg}, we must compare the average derivative of the current reputation curve, with the average derivative of the minimum reputation curve $Rep(R_{th})$ (using, for example, the least squares method). There are two possible situations:

- $\frac{d}{dt}Rep \geq \frac{d}{dt}Rep_{th}$: the curves diverge, *i.e.* the reputation of the peer increases more quickly than the minimum reputation, under which the peer is degraded;

- $\frac{d}{dt}Rep < \frac{d}{dt}Rep_{th}$: the curves converge, *i.e.* in a non-infinite time the peer will be degraded.

Also note that both curves depend on f_{bad}, because $\Delta T = \Delta n_+ + \Delta n_- + \Delta n_{bad}$.

Starting from $R = R_{avg}$, the good ratio decreases if the number of negative feedbacks per time unit is higher than the number of positive feedbacks. In particular, this eventuality can arise if $f_{bad} > 0$. In our simulations, for each role we set a recovery window (according to the parameters illustrated in section 4), which enters the game when $R < R_{avg}$, and contributes to maintain the reputation value over R_{th}, *i.e.* $\frac{d}{dt}Rep \geq \frac{d}{dt}Rep_{th}$.

Figure 6 illustrates the average evolution of an admin peer's reputation over the simulation time interval. If no malicious peers provide unjustified negative feedbacks, the measured reputation of the target peer is represented by the fat continuous curve. Comparing this curve with the graph of the reputation which we obtain if $R = R_{th}$ (the thin continuous curve in the figure), we can observe that they diverge, thus we expect that the peer will not be degraded unless $f_{bad} = 0$. In the same figure, dotted curves refer to the case of $f_{bad} = 20\%$; they still diverge. Finally, dashed curves show the limit over which the role preservation condition is not fulfilled, *i.e.* $f_{bad} \simeq 35\%$. This result is compatible with the analytical model, for which $max\{f_{bad}\}$ is 40% (in the worst case $R = R_{th}$).

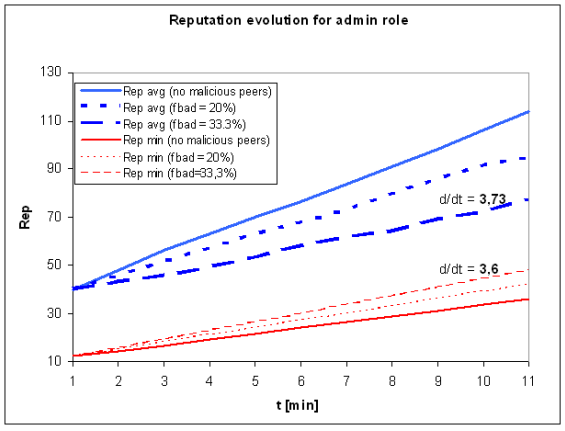

Figure 6: Average and mimunum reputation dynamics, for an admin peer, for different rates of malicious negative feedbacks.

The most interesting simulation results for the publisher role are illustrated in figure 7, which compares the case of no malicious peers (continuous lines) with the case of unjustified negative feedbacks with $f_{bad} = 29.4\%$ rate (dashed line). We can observe that, in the latter case, the derivatives demonstrate that the curves converge. We measured $max\{f_{bad}\} = 29\%$, over which the publisher is degraded in a non-infinite time. Also this result is compatible with the analytical model, for which $max\{f_{bad}\}$ is 33.3% (in the worst case $R = R_{th}$).

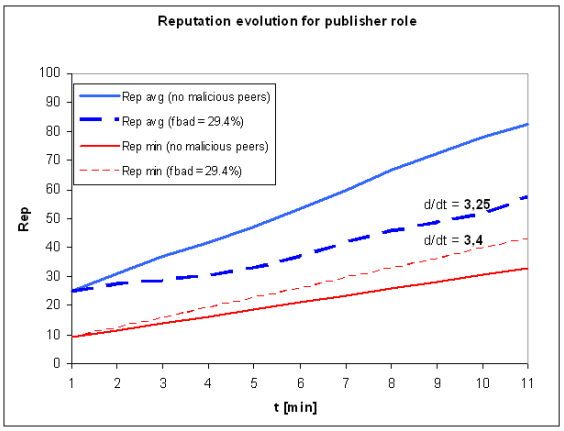

Figure 7: Average and mimunum reputation dynamics, for a publisher peer, for different rates of malicious negative feedbacks.

Figures 8 and 9 illustrates, respectively, simulation results for the searcher and the newbie roles. We measured a maximum malicious feedbacks rate $max\{f_{bad}\} = 25\%$, for a searcher peer. The figure illustrates what happens when this rate is overthrown, *i.e.* the reputation curves (average and minimum) converge. For a newbie peer, the measured maximum rate of unjustified negative feedbacks is $max\{f_{bad}\} = 21\%$. The figure illustrates a less dangerous situation. Both these simulations gave satisfactory results, which respect the numerical constraints obtained with the analytical model.

All results are summarized in table 1. The first and second columns report the analytical results, respectively for a reputation management service without and with recovery window. The third column reports the simulation results, which refer to a reputation management service with recovery window and target peer with righteous behaviour, *i.e.* a peer which would maintain the initially assigned good ratio R_{avg} in abscence of malicious negative feedbacks.

6 RELATED WORK

There have been several studies about managing reputation in P2P networks, most of them related to content sharing, which has been the killer application for these architectures. A reputation management system in DHT-based structured P2P networks is proposed

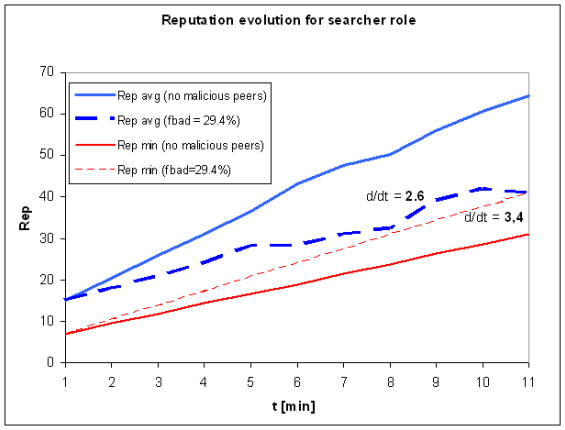

Figure 8: Average and mimunum reputation dynamics, for a searcher peer, for different rates of malicious negative feedbacks.

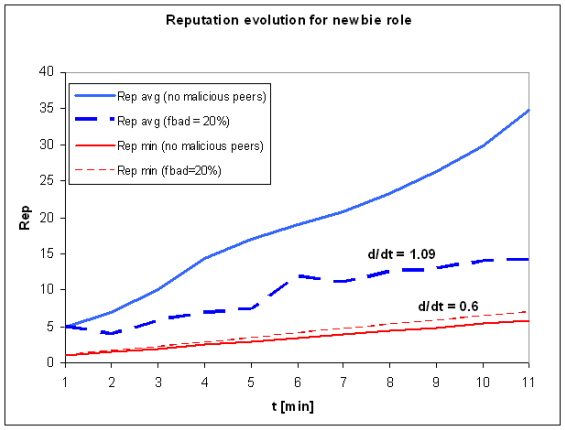

Figure 9: Average and mimunum reputation dynamics, for a newbie peer, for different rates of malicious negative feedbacks.

Role	th_{norec}	th_{rec}	sim_{rec}
A	25%	40%	35%
P	20%	33,3%	28%
S	14,3%	25%	25%
N	13%	23%	21%

Table 1: Maximum tolerable malicious feedbacks rate $max\{f_{bad}\}$: analytical results without and with recovery window, and simulated results.

in (Lee et al., 2005); this model uses file reputation information as well as peer reputation information, and the system uses a global storage for reputation information, that is available when evaluator is not on-line. The reputation information consists, as in our model, of two values representing the number of positive and negative feedbacks.

In (Mekouar et al., 2004b) a reputation management system for partially-decentralized P2P systems is described, in which the reputation information is managed by supernodes. The authors assume that the supernodes are selected from a set of trusted peers, and they share a secret key used to digitally sign the reputation data. Good reputation is obtained by having consistent good behaviour through several transactions. The proposed scheme is based on four values associated to each peer and stored at the supernode level; two of them are used to provide an idea about the satisfaction of users, and the others express the amount of uploads provided by the peer. The reputation information is updated according with the peer transactions of upload and download. In a successive work (Mekouar et al., 2004a) the same authors propose an algorithm to detect malicious peers which are sending inauthentic files or are lying in their feedbacks. They distinguish righteous peers from those which share inauthentic files and provide false feedbacks about other peers. The model introduces the concept of suspicious transaction, that is a transaction whose appreciation depends on the reputation of the sender and the concept of credibility behaviour, as an indicator of the liar behaviour of peers. These schemes are able to detect malicious peers and isolate them from the system, but do not consider the tolerable rate of malicious negative feedbacks, and suppose that supernodes are always trustworthy.

A distributed method to compute global trust values, based on power iteration, is illustrated in (Kamvar et al., 2003). The reputation system aggregates local trust values of all users, by means of an approach based on transitive trust: a peer will have a high opinion of those peers which have provided authentic files and it is likely to trust the opinions of those peers, since peers which are honest about the files they provide are also likely to be honest in reporting their local trust values. The idea of transitive trust leads to a system where global trust values correspond to the left principal eigenvector of a matrix of normalized local trust values. All peers in the network cooperate to compute and store the global trust vector, taking into consideration the system history with each single peer. The scheme is reactive, *i.e.* it requires reputations to be computed on-demand, through the cooperation of a large number of peers. This introduces additional latency and requires a lot of time to collect statistics and compute the global rating.

To identify malicious peers and to prevent the spreading of malicious content, a reputation-based architecture is proposed in (Selcuk et al., 2004). The protocol aims to distinguish malicious responses from benign ones, by using the reputation of the peers which provide them. The protocol relies on the P2P infrastructure to obtain the necessary reputation infor-

mation when it is not locally available at the querying peer. The outcomes of past transactions are stored in trust vectors; every peer maintains a trust vector for every other peer it has dealt with in the past. The trust query process is similar to the file query process except that the subject of the query is a peer about whom trust information is inquired. The responses are sorted and weighted by the credibility rating of the responder, derived from the credibility vectors maintained by the local peer, which are similar to the trust vectors.

In (Garg et al., 2005), the use of a scheme named ROCQ (Reputation, Opinion, Credibility and Quality) in a collaborative content-distribution system is analyzed. ROCQ computes global reputation values for peers on the basis of first-hand opinions of transactions provided by participants. Global reputation values are stored in a decentralized fashion using multiple score managers for each individual peer. When a peer wishes to interact with another peer, it retrieves the reputation values for that peer from its score managers. The final average reputation value is formed by two aggregations, first at the score managers and second at the requesting peer; if a peer has had interactions with the prospective partner before, it may wish to prefer its own first-hand experience to the information being provided by the trust management system or to use a combination of the global reputation and its first hand experience.

All these works consider the situation in which a peer with a bad reputation is simply isolated from the system, while the analytical model we are proposing describes different roles for peers, associated with different actions. So a peer with a suspect malicious behaviour can be first degradated, and eventually isolated from the system. It is also possible to compute the probability that the next feedback will be positive for a peer, that allows a peer to increase its good ratio, and the maximum rate with which one or more malicious peer can provide negative feedbacks without affecting the peer's role.

7 CONCLUSIONS

In this work we have illustrated the analytical model of a reputation management service for role-based peergroups. The model defines some parameters and indicators, such as the maximum tolerable rate of malicious negative feedbacks. We applied the reputation model to a four-role security policy, giving a parameter set for each role, and computing the theoretical values for the main indicators. These results have been confirmed by those we obtained from several simulations, which we realized using a centralized reputation management service.

Further work will follow two directions. To complete the analytical model, we must consider also malicious positive feedbacks. For example, we could check for suspiciously rapid increasing of good ratios, and introduce a recovery window not only to prevent unjustified degradations, as in current model, but also to contrast malicious promotion attempts. Once the model is completed, and all parameters are tuned, we can search for the best distributed solution for reputation storage and retrieval.

8 ACKNOWLEDGEMENTS

This work has been partially supported by the "STIL" regional project, and by the "WEB-MINDS" FIRB project of the National Research Ministry.

REFERENCES

Amoretti, M., Zanichelli, F., and Conte, G. (2005). SP2A: a Service-oriented Framework for P2P-based Grids. In *3rd International Workshop on Middleware for Grid Computing, Co-located with Middleware 2005.*

Barabási, A. and Albert, R. (1999). Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512.

Garg, A., Battiti, R., and Cascella, R. (2005). Reputation Management: Experiments on the Robustness of rocq.

Kamvar, S. D., Schlosser, M., and Garcia-Molina, H. (2003). The Eigentrust Algorithm for Reputation Management in peer-to-peer Networks. In *The 12th International World Wide Web Conference*.

Lee, S. Y., Kwon, O.-H., Kim, J., and S.J.Hong (2005). A Reputation Management System in Structured peer-to-peer Networks. In *The 14th International Workshops on Enabling Technologies:Infrastruture for Collaborative Enterprise*.

Mekouar, L., Iraqui, Y., and Boutaba, R. (2004a). Detecting Malicious Peers in a Reputation-based peer-to-peer System.

Mekouar, L., Iraqui, Y., and Boutaba, R. (2004b). A Reputation Management and Selection Advisor Schemes for peer-to-peer Systems. In *The 15th IFIP/IEEE International Workshop on Distributed Systems: Operations and Management.*

Selcuk, A. A., Uzun, E., and Pariente, M. R. (2004). A Reputation-Based Trust Management System for peer-to-peer peer-to-peer Networks. In *IEEE International Symposium on Cluster Computing and the Grid.*

Ye, S., Makedon, F., and Ford, J. (2004). Collaborative Automated Trust Negotiation in Peer-to-Peer Systems. In *Fourth International Conference on Peer-to-Peer Computing (P2P'04).*