Costs and Benefits of Reputation Management Systems

Roberto G. Cascella
University of Trento
Dipartimento di Ingegneria e Scienza dell'Informazione
Via Sommarive 14, I-38100 Povo (TN), Italy
cascella@disi.unitn.it

Abstract

The specific objective of a reputation management system is to facilitate nodes to decide whom to trust before providing resources and the expected quality of service. In self-organized systems, the use of reputation-based mechanisms has a cost in terms of messages required to disseminate feedbacks and to synchronize reputation values.

In this paper we discuss the implementation of a reputation management system and consider two approaches to collect and disseminate data: proactive and reactive schemes. We analyze the cost of these approaches in terms of extra signaling and evaluate the benefit of using reputation management systems, measured as the number of avoided transactions with malicious nodes.

1 Introduction

Self-organized systems consist of heterogeneous components that interact to disseminate and collect data in order to accomplish complex tasks, such as the definition of new services that adapt to the context of the communication. In this scenario, cooperation is required to achieve scalability and maintain the survivability of the system which evolves continuously. Components, or nodes, might act selfishly, i.e., they do not share the data they own, or they can be malicious by injecting false content. In this potentially non-cooperative environment, trust and reputation management schemes have proven to be an effective countermeasure against selfish [2] and malicious users in peer-to-peer [3, 5, 9] and ad hoc networks [8].

Reputation information can work as the fitness criteria for the evolution and adaptation of the systems by sustaining their stability. The use of this information introduces an extra signaling due to the operations to manage trust and to disseminate and collect feedbacks from the nodes. While the collection of feedbacks can be assumed to be part of the process required for the loopback control of context-aware systems, the communication overhead of reputation management schemes must be considered to evaluate the effectiveness of using reputation information.

On one hand, it is important to minimize the communication cost and, on the other hand, cooperation and protection against behavioral attacks should be achieved. For instance, mobile networks are formed primarily by battery-powered components, thus, the number of messages has high impact on the energy-efficiency of the system and must be reduced to increase its survivability. Moreover, in peer-to-peer networks malicious nodes can provide poor quality services or corrupted content to reduce the performance of the application. Reputation can be used to decide whether to interact with a node to reduce the number of malicious transactions and to increase network resilience [2, 5].

In this paper, we analyze the available architectural solutions to implement reputation management systems and we define the messages required by these schemes to function properly. We quantify the cost, in terms of number of messages to handle reputation information and feedbacks, and the benefit, measured as the number of avoided transactions with malicious entities and their correct identification.

The rest of the paper is organized as follows. Section 2 presents the architectural solutions and discusses data dissemination strategies. Section 3 defines the available types of reputation information and the format of messages. Section 4 details the implementation of the reputation system used for this study and Section 5 presents experimental results. Section 6 concludes the paper.

2 Networking of reputation systems

The definition of reputation management systems comes from the need in online communities to create a communication framework similar to traditional social relationships. The real application of these systems has many facets. In fact, the way they must be deployed and their efficacy depend on several factors, such as the application context, type of nodes and risk of the communication. In all cases, it is

important to assess what is the *benefit* of using reputation management schemes to improve the quality of online interactions and, what is their *cost*.

Reputation management systems monitor the behavior of the nodes during transactions and assist them in selecting the interacting peers. These systems consist of three functions to 1) collect feedbacks after transactions, 2) aggregate them to form a useful measure of the trustworthiness of a node and 3) disseminate the reputation value of a particular node to requesting peers. In this section, we discuss the available architectural solutions and analyze the impact of the collection and dissemination methods. These methods depend on the choices of the system designer, the application, and the underlying networking properties. We do not discuss the aggregation function as it does not induce any signaling cost.

2.1 Networking topologies

Networking systems can be categorized in two types: centralized or decentralized. In this paper, we only analyze decentralized systems as we assume that the communication overhead increases linearly with the number of nodes in the other case. Distributed systems do not have a centralized component and the load of this server must be shared among the system components.

We can distinguish two topologies: unstructured and structured. In unstructured systems, a node has limited view of the network and it primarily interacts with its neighbors or other entities that it has met in the past. Structured overlays are by definition organized topologies where nodes are *virtually* connected. The reputation management scheme must consider the properties of the overlay network and might follow the same structure of communication.

2.2 Information dissemination

The reputation management system does not require specific communication protocols to achieve its functionality. Hence, the dissemination of the information depends on the network infrastructure available and the overlay network on top of which the reputation management system is built. To reduce the communication overhead several mechanisms can be used, such as piggybacking the information to application data or aggregating multiple feedbacks in one message. Information dissemination can be based on proactive, reactive or hybrid approaches. The hybrid scheme is defined as a combination of the reactive and the proactive.

The use of reputation management schemes is based on feedbacks. They are formed locally after every transaction to judge the amount of satisfaction a node perceives. They are aggregated at several location and form the reputation values [7]. To generalize the results derived in this study, we analyze a reputation management system where there exist designated agents to aggregate and store reputation values. Nodes are assigned to designated agents, which are normal peers that participate in the system and have monitoring capabilities. The nomination of the designated agents depends on the networking topology and the overlay structure.

In our study, the reputation value of a node is computed and stored by several designated agents as many as the number of replica. This choice mitigates the effect of malicious entities that falsify their reports, and guarantees resilience and robustness in the network. For instance, it is not feasible in mobile network to assume all designated agents to be always present during the whole system lifetime.

2.2.1 Collection of feedbacks

Designated agents can collect feedbacks in two distinct ways, proactively or reactively. The proactive approach requires a node to send a feedback on an other peer after every transaction, which is defined as an interaction between two nodes within the context of the application. This enables fast detection of misbehaving nodes but it may incur additional overhead, as nodes must notify all designated agents that are in charge of the peer reputation.

If the information is collected reactively, a node only sends its feedbacks when it is requested to do so. The reputation collecting service at designated agents periodically polls all nodes to gather new feedbacks of all their past transactions, or the aggregated values, since the previous polling round; nodes can also send feedbacks at regular intervals. The effectiveness of this approach is based on the frequency of the queries. If the frequency is too high, nodes might have few interactions to report and the communication overhead introduced by the reputation management system is not compensated by a more accurate measure of the reputation value. On the contrary, if the frequency is too low, reputation values can be outdated.

The reactive approach can use two different strategies for sending feedbacks: 1) upon receiving a request, each node broadcasts all available information and the designated agents process the data for the nodes they are responsible for or 2) each node groups the feedbacks, relevant for every single designated agent, and sends distinct messages.

The first strategy has the advantage of reducing the total number of messages in the system, but it requires every designated agent to process each message even if it does not contain any feedbacks on the nodes which it is responsible for. Since every node in the system can be a designated agent, the total computational cost might be large, as all nodes might process the messages. On the other hand, the second strategy increases the total number of messages, but designated agents only receive the information that they have asked for.

2.2.2 Dissemination of reputation values

The reputation value of a node is aggregated and stored at designated agents. As defined for feedbacks collection, these agents can disseminate this information following two approaches: 1) every time a node participates in a transaction, by following a proactive approach, or 2) at regular intervals by pushing all the reputation values they store.

The most efficient method for reaching all the nodes is *flooding*, i.e., messages are forwarded to all nodes in the community even if they are not interested. In this case, the overhead might be large and the cost of the extra signaling might not be compensated for by the benefit of flooding. Better strategies consist of forwarding messages by inserting in the system *K-copies* of the same data, with $K < N - 1$ and N is the number of nodes, to reduce the dissemination time as well as the number of messages.

Flooding or *lighter* versions are effective in highly connected topologies, since few messages are sufficient to reach all nodes in the system, or in case of disconnected and dynamic networks to enhance the robustness of the system. However, flooding should be avoided in mobile networks, which are energy-aware systems, since processing the messages can consume the batteries of devices fast.

If a hybrid approach is implemented, trust values are pushed periodically to nodes and in case of need they can be retrieved on-demand from designated agents. In the reactive approach, the overhead depends on the frequency of the messages and the number of nodes. In this case the polling rate must be chosen carefully to not overload the system with reputation messages. An adaptive scheme can be used: if the number of nodes is high, the rate is reduced and nodes use a proactive approach to update reputation values.

2.2.3 Local operations

Nodes record all transactions, that they have participated in, to form the feedbacks, i.e., opinions on the amount of satisfaction they perceive in a transaction. Locally nodes can cache temporarily these opinions as well as the reputation values for a subset of nodes with which they have already interacted. This information is useful when interactions between the same nodes are frequent and when the communication with designated agents is not possible, such as in congested or disconnected networks like in mobility scenarios.

Since storage capabilities of nodes are limited, outdated reputation values or transactions' feedbacks can be safely filtered and discarded. Indeed, a timestamp can be used to indicate the freshness of the data. This has a twofold meaning: 1) to reduce the amount of information stored at designated agents or local peers and 2) to age both the importance of a feedback, to compute the reputation value, and the reputation value itself.

3 Type of reputations

In reputation management systems, the burden of the communication is mainly on designated agents which have the function of collecting, aggregating and disseminating reputation values. The communication is based on packets or on the notion of messages. A messages consists of a payload and metadata that carry the necessary information for a node to process the payload. In this section we define the reputation messages and we discuss different types of *reputation* that can be associated to a node.

Reputation information can be grouped in four different classes: 1) first-hand information on nodes (or *opinion* (O_n), 2) second-order information on nodes (or *credibility* C_n) 3) reputation information on nodes (R_n), and 4) the subjective trust information on a node (T_n). A node's trustworthiness (T_n) is computed by aggregating locally the reputation values, weighted by the credibility of the reporting nodes, and the personal experience of the nodes.

3.1 Opinion

Each transaction between two entities results in the formation of an opinion that consists of the amount of satisfaction measured during the interaction. This information is part of the feedback that nodes send to *designated agents* responsible for aggregating the peer reputation value. The feedback also includes the confidence that the reporting node has in its opinion, i.e., the opinion quality.

Feedbacks are associated with a timestamp, that indicates when the transaction takes place. The size of the feedback depends on the number and type of fields. In general, one feedback can be encoded on 28 bytes as defined below.

Node identifier Id is encoded on a fix number of bytes and it is used not for forwarding decisions but to record correctly a transaction. The hash value of the identifier can be used to provide anonymity and avoid possible collusion attacks [1, 3, 6]. A MD5 hash function produces a 16 bytes identifier.

Timestamp TS is used to age the information and it can be encoded on 4 bytes. This number can be further reduced by defining different formats.

Opinion value O is the subjective view of the node. This value is between 0 and 1. A value of 0 means that a node is not trustworthy and a value of 1 means completely trustworthy. Opinion values can be encoded on 4 bytes.

Quality of the opinion $Q(O)$ is the *confidence* associated to the opinion and it can be encoded on 4 bytes. This value might be used for reputation aggregation at designated agents as it quantifies the confidence a node has on the reported opinion. For instance, a low quality value is associated to an opinion when it is based on few observations or these observations are not consistent [3, 4].

3.2 Reputation

Reputation is the global system view of the behavior of a node. It is maintained locally by designated agents for those nodes they are responsible for. In some case, other nodes can also decide to store a copy. The computation of the reputation value depends on the aggregating function. The payload of the message carrying the reputation value might have a length of 24 bytes.

Node identifier Id is encoded on 16 bytes as defined earlier.

Reputation value R is the reputation value aggregated at designated agents. Similarly to opinion, reputation is defined between 0 and 1 and can be encoded on 4 bytes.

Quality of the reputation value $Q(R)$ is the *confidence* of designated agents associated to reputation values. It can be encoded on 4 bytes.

3.3 Credibility and trust information

Credibility, or second-order trust information, is computed by designated agents, on all nodes reporting opinions, and by nodes, on the reporting designated agents. It is not shared with other nodes and it is evaluated on a direct-experience basis only. The second-order trust information is defined as a measure of the judging capability of reporting peers and it is used to weigh feedbacks or reputation values received from nodes and designated agents respectively.

For instance, a designated agent, misreporting reputation information, can be detected by comparing the received reputation scores from other designated agents. Then, its credibility, locally computed, is adjusted in accordance [3].

Similar to credibility, nodes compute the trust information locally. It defines how a node judges trustworthy another peer for providing a specific service. This information is used ultimately to decide upon a transaction. It can be based on both local opinions and global reputation values, or on one of the two scores.

4 Implementation and design considerations

The reputation management scheme can function either in unstructured and structured overlay networks. In our model we assume that the system is organized in a Distributed Hash Table (DHT) to simplify the nomination of designated agents. The hash of the node identifiers gives the position in the DHT, and the node responsible for that portion of the table is elected as designated agent. Multiple designated agents can be nominated by using different seeds to compute the hash. In this study, we assume that nodes have only one identifier, which cannot be spoofed.

Designated agents collect feedbacks encapsulated in messages either in reactive or proactive mode, as defined in Section 2. In a DHT, the cost of this operation is proportional to two factors: 1) the number of lookups, as we assume multiple designated agents, and 2) the cost of each lookup. The latter depends on the number of messages in the DHT scheme to reach the designated agents. The lookup cost can be the order of $O(logN)$, where N is the number of nodes in the DHT. The dissemination is done in a way similar to collection.

Table 1. Parameters' setting

Network Topology	random
Number of iterations	45,000
Experiments run	5
Malicious Nodes	30%
Type of maliciousness	Transaction and feedback
Feedbacks collection	Proactive or reactive
Reputation dissemination	Proactive or reactive
Polling frequency	2,500 or 5,000 iterations
Trustworthiness threshold	0.5
Designated agents	5 (if otherwise specified)

Our model targets malicious nodes that misbehave when interacting with other peers, both when reporting feedbacks to designated agents and when reporting reputation values to nodes in the role of designated agents. Reputation values of the nodes are aggregated at designated agents and they are computed by using the feedback received and the credibility of the reporting node. Initially, the credibility of all nodes is set at 0.5 by default and designated agents or nodes can increase or decrease their credibility by reporting accurate or inaccurate information. Multiple agents, 5 if not otherwise specified, are also used to mitigate the effect of malicious nodes in the system.

In our experiments we use a simplified model of the ROCQ reputation scheme [3, 4] and the simulation parameters listed in Table 1. We run a number of initial transactions equal to the number of nodes to bootstrap the reputation management scheme. We simulate a population with 30% of malicious nodes both in transactions and in reporting scores; they report the inverse of the estimated opinion, or reputation value for the case of designated agents.

At each iteration (or transaction), two nodes, which are selected randomly, compute the trust value of each other. If both values are above the deterministic threshold 0.5, the nodes interact. The trust value is a function of the opinions, that a node forms after each transaction, and of the reputation value reported by designated agents.

The *benefit* of the reputation management system is evaluated in terms of success rate, defined as the number of correct decisions made (i.e., interactions with good peers plus the number of avoided interactions with malicious peers) as proportion of the total number of decisions made. Only de-

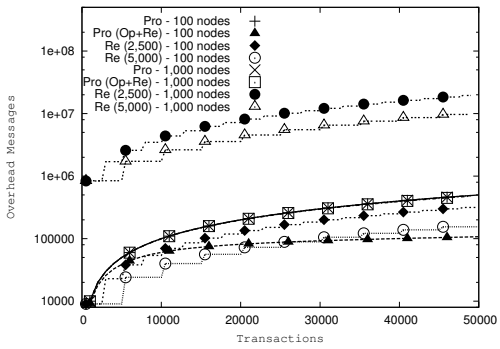

Figure 1. Overhead messages in a system composed of 100 or $1,000$ nodes using the Proactive (*Pro*) or the Reactive (*Re* - polling every $2,500$ or $5,000$ transactions) approach.

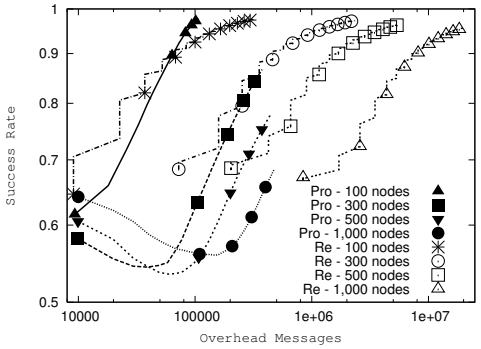

Figure 2. Impact of the number of nodes on the proportion of correct decisions and on the number of overhead messages. (Collection and dissemination every $2,500$ - $Op+Re$).

cisions made by good peers are counted. We also calculate the *cost*, i.e., the number of messages exchanged due to the use of a reputation management scheme, and evaluate the impact of this overhead on the success rate. To count the messages we define the reactive and proactive approaches as follows.

Proactive: nodes request to designated agents the reputation value of the interacting node before each interaction. The feedback is sent immediately after.

Reactive: each node aggregates the reports, which it must send to each designated agent, and these feedbacks are collected at regular intervals. The reputation values of all nodes in the system are stored locally and they are updated when designated agents disseminate new values. We simulate two frequencies for collection and dissemination: $2,500$ or $5,000$ transactions between two subsequent polls.

5 Experimental results

In this section we evaluate the extra signaling introduced by the reactive and proactive approaches and we analyze the impact of different parameters on the performance of the reputation management scheme.

5.1 Overhead messages

Fig. 1 shows the number of overhead messages. We only count the messages required by the reputation management scheme to function properly, i.e., for the collection and dissemination of information. We simulate two different strategies for the node to decide if carrying on a specific transaction with another peer: the decisions are based on the peer reputation only or on a combination of the peer reputation and the local opinion. In this last case, reputation is used instead of opinion, if the two nodes have interacted less than 5 times, i.e, there are few samples to form a correct personal opinion on the behavior of the node. A node is considered trustworthy if the score, either reputation or opinion, is above the threshold 0.5.

The combination of opinion and reputation, indicated by *Op+Re* in Fig. 1, is only used for the proactive approach, as reputation is always disseminated in the reactive one. As expected, a higher polling frequency, $2,500$ compared to $5,000$, increases the amount of overhead messages. In this case, the extra signaling is also highly dependent on the size of the system. For the proactive approach, the overhead is the same when the number of nodes differs unless decisions are based on opinion first. In fact, this reduces the number of queries to the designated agents if the same nodes interact often and have sufficient information to form a correct opinion. This is evident for the case of 100 nodes, *Pro (Op+Re)* in Fig. 1.

5.2 Reactive and proactive approaches

In Fig. 2 we plot the success rate as a function of the number of messages transmitted in the system to analyze the trade-off between the benefit of using reputation information for taking decisions on transactions and the cost measured by overhead messages. The curves show that the number of nodes has a great impact on the performance. In particular, for the reactive approach the amount of overhead messages is proportional to the number of nodes. Indeed, when the system is composed of few nodes, a smaller number of messages is sent and the success rate increases.

The results for the proactive approach stress that an accurate estimation of the nodes' behavior is a function of the

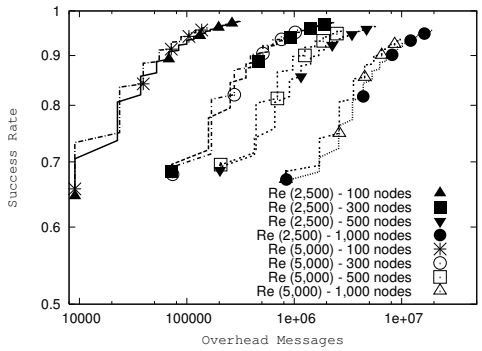

Figure 3. Reactive approach: impact of system size and polling frequency.

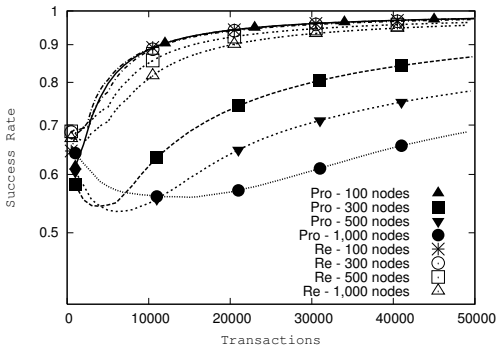

Figure 4. Performance of reactive (Re - polling frequency $2,500$) and proactive (Pro) approaches in terms of success rate.

number of the interactions. When there are few nodes in the system, nodes interact more frequently, thus, there are more samples to estimate the nodes' reputation value, as shown also in Fig. 1. Fig. 2 confirms our hypothesis, as the curves for the proactive approach show that the success rate decreases initially and then increases as more information is available. This is more evident when the number of nodes becomes larger.

Fig. 3 shows the performance of the reactive approach with different values of the frequency for feedbacks collection and dissemination of reputation values, such as $2,500$ and $5,000$ transactions between two subsequent polls.

With a frequency of $5,000$ transactions, the amount of new information is higher, thus, the reputation system is able to have a better approximation of the behavior of the nodes by using the same amount of number of messages compared to a higher frequency. With few nodes in the system, the frequent collection of feedbacks has the advantage of updating the reputation value faster and the evaluation of the reputation value is more accurate. As a guideline, the polling frequency must consider the system size and the frequency of transactions.

In Fig. 4, we plot the success rate as a function of the number of transactions for the proactive and reactive approaches. We can conclude that the reactive approach introduces higher overhead, but it gives better performance as all the nodes have global knowledge in the system. Fig. 4 shows that the success rate of the proactive approach is comparable to the reactive one in the case of 100 nodes. As observed earlier, having few nodes in the system means more samples to estimate locally their behavior.

5.3 Number of designated agents

In this section, we analyze the impact of the number of designated agents to the overall computation of the communication overhead. Fig. 5 plots the success rate as a function of the number of messages transmitted to disseminate information in proactive or reactive (every $2,500$ transactions) mode. It is worth noticing that, when only reputation is used as a metric to judge the trustworthiness of the nodes, a high number of designated agents is important to predict better the behavior of the nodes, plot (a) in Fig. 5.

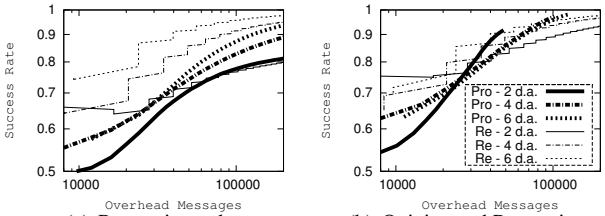

(a) Reputation only (b) Opinion and Reputation

Figure 5. Impact of the number of designated agents (2, 4, or 6 d.a.) on the success rate with 100 nodes and a frequency of $2,500$ transactions for the reactive approach (Re).

In our simulation, we consider the case of designated agents that can also be malicious. Thus, multiple designated agents for the same node means more estimations of the node's reputation. This guarantees that the malicious behavior of a small fraction of designated agent in reporting reputation values does not influence the evaluation of the trustworthiness of the node. However, the cost of multiple designated agents is paid in terms of number of messages used to collect and disseminate the data.

Similar conclusions are derived from the plot (b) in Fig. 5, when nodes use also opinions to take decisions on a transaction. This reduces both the need for reputation values in the proactive approach and the impact of malicious

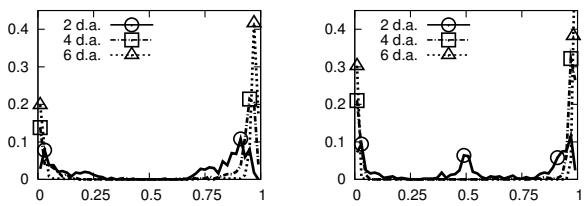

(a) Proactive (b) Reactive $2,500$

Figure 6. Probability Distribution Function (PDF) of the reputation value for the proactive (a) and the reactive (b) approach with a polling frequency of $2,500$ interactions. (100 nodes - after $22,000$ transactions)

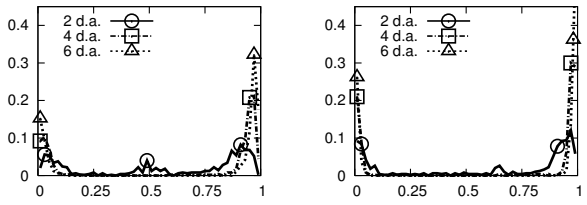

(a) After $22,000$ transactions (b) After $45,000$ transactions

Figure 7. Probability Distribution Function (PDF) of the reputation value when a reactive approach is used. (100 nodes - collection and dissemination every $5,000$ interactions)

designated agents in general. Thus, we can conclude that the success rate increases if there are more designated agent and opinion is used first.

5.4 Estimation of reputation values

Another important metric to evaluate the performances of the proactive and reactive approaches is the probability distribution of the reputation values in the system. As discussed in previous sections, varying the number of designated agents and nodes in the system has a great impact on the success rate of the reputation management scheme. This also influences the accuracy of reputation values.

In Fig. 6–8 we plot the distribution of the reputation values for 2, 4 and 6 designated agents (d.a. in the plots) when 30% of the nodes are malicious. Thus, we expect to have a 30% of the nodes with a reputation value below 0.5 and in the best case close to 0. In Fig. 6 and Fig. 7 we consider a system with 100 nodes while in Fig. 8 with $1,000$ nodes.

As discussed in Section 5.3, a high number of designated agents is required to mitigate the effect of a fraction of misbehaving agents. Plot (a) in Fig. 6 shows that when 2 designated agents are assigned to each node, trustworthy nodes cannot be rewarded with a high reputation value for their cooperative behavior. However, with 6 designated agents, malicious nodes are correctly identified already after $22,000$ transactions. We can notice that the population is divided in two groups, malicious nodes (on the left of the plot) and trustworthy nodes (on the right of the plot).

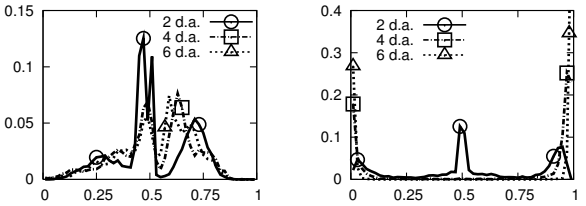

(a) Proactive (b) Reactive $2,500$

Figure 8. Probability Distribution Function (PDF) of the reputation value with a different scale for the proactive (a) and the reactive (b) approach with a polling frequency of $2,500$ interactions. ($1,000$ nodes - after $22,000$ transactions)

Plot (b) in Fig. 6 shows the distribution of the reputation value when information is collected and disseminated reactively. In this case, designated agents poll nodes every $2,500$ iterations. As anticipated in Section 5.2, the reactive approach converges faster for the prediction of the nodes' trustworthiness than a proactive one. In particular, 4 designated agents per node are sufficient to give a good estimation of the reputation values. If we assuming frequent transactions between the same two nodes in a system that consists of 100 entities, feedbacks are accurate and designated agents can predict better reputation values consequently. Hence, a false report from a malicious designated agent is detected timely and its credibility is lowered.

In Fig. 7, we decrease the polling frequency down to every $5,000$ transactions. For this setting, we might expect better accuracy compared to the previous case as nodes form feedbacks on a grater number of transactions. But, this is not completely verified as designated agents query nodes less often. As a result, nodes might use old reputation values that do not map the current state of the system to take decisions. This results in slower convergence time of the reputation management scheme, shown in plots (a) and (b) of Fig. 7. This uncertainty also creates room for malicious designated agents to bias nodes' decisions.

Fig. 8 shows the distribution of reputation values for $1,000$ nodes in the system. The number of nodes influences the accuracy of the reputation values' estimation and the

converging time of the management scheme.

We notice that when a proactive approach is used for collection and dissemination, the reputation scheme is not able to differentiate between trustworthy and malicious nodes, plot (a) in Fig. 8. This is the result of sporadic transactions, which means few samples per node to predict the reputation value. This uncertainty is clearly shown in plot (a) of Fig. 8 as the reputation values are all almost close to 0.5 even if more designated agents are used.

For the reactive approach, we can conclude that the number of designated agents is important to evaluate reputation values correctly. Indeed, plot (b) of Fig. 8 shows that 2 designated agents are not sufficient to make the reputation scheme to function properly as their possible malicious behavior increases the uncertainty due to the limited number of samples for the estimation of reputation.

From Fig. 8, we can conclude that increasing the number of nodes causes slower convergence of reputation values to a good approximation of the nodes' behavior.

6 Discussion

In this paper, we perform a detailed experimental evaluation of the communication overhead of reputation management schemes and analyze the conditions under which the use of reputation is beneficial to the system. We use a simplified version of ROCQ to simulate a reputation scheme, and the conclusions can be generalized to other distributed reputation management mechanisms. We simulate different approaches for their implementation with varying characteristics in terms of number of nodes, number of designated agents, frequency of information dissemination and evaluation of nodes' trustworthiness.

We find that the number of overhead messages depends on the size of the system and on the frequency of direct transactions between the same two nodes. A good choice of the frequency for information collection and dissemination depends on how much available information is new and updated. In our simulations, we have defined the polling frequency as the number of interactions between nodes in the system. In this way we are in the position to analyze the impact of the amount of new information rather than its freshness. Our conclusions only hold in the cases we explore, but we can predict that when nodes transact sporadically, a proactive approach results in updated reputation values at smaller cost. When the frequency is high, a reactive approach updates the reputation values periodically without making the nodes ask.

In conclusion, the collection of feedbacks and dissemination of reputation values must not overwhelm the normal functionality of the application as the communication overhead can cause network congestion or in the worst case it can drain all resources. The frequency of these operations can be adjusted to respond to specific needs of the application or to the context of communication. As guideline, in a hostile environment a higher frequency increases the extra signaling but it enables fast detection of malicious attacks.

An interesting area of future research is churn, i.e., nodes that join and leave the system. In our results, we find that the number of nodes decreases the convergence for the estimation of reputation values without considering churn. Churn might cause a slower convergence of the reputation values [3], but it would have slightly affected the number of overhead messages that mainly depends on the system size. In particular, in small systems, the use of local opinions is sufficient to quantify the trustworthiness of a node, as it is shown by the increase of the success rate.

Our main contribution is the analysis of the applicability of reputation management systems. The results show the benefit of reputation systems, applied to any self-organized system, in sustaining the availability of resources at the cost of extra signaling.

7 Acknowledgment

Work partially supported by projects DAMASCO funded by the Italian Ministry of Research and BIONETS (FP6-027748) funded by European FET Program.

References

[1] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *CIKM*, New York, USA, 2001.

[2] R. G. Cascella. The "Value" of Reputation in Peer-to-Peer Networks. In *CCNC 2008*, Las Vegas, USA, Jan 10-12 2008.

[3] A. Garg, R. Battiti, and R. Cascella. Reputation management: Experiments on the Robustness of ROCQ. In *WAGEN Workshop at ISADS*, pp. 725–730, China, Apr. 2005.

[4] A. Garg, R. Battiti, and G. Costanzi. Dynamic Self-management of Autonomic Systems: The Reputation, Quality and Credibility (RQC) scheme. In *WAC 2004 (LNCS 3457)*, pp. 165–176, Berlin, Germany, Oct. 2004.

[5] A. Garg and R. Cascella. Reputation management for collaborative content distribution. In *ACC Workshop at WoWMoM*, pp 547–552, Taormina, Italy, June 13-16 2005.

[6] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in P2P networks. In *Proc. of the 12th international conference on World Wide Web*, pp. 640–651, Budapest, Hungary, 2003..

[7] S. Marti and H. Garcia-Molina. Taxonomy of trust: categorizing p2p reputation systems. *Computer Networks*, 50(4):472–484, 2006.

[8] P. Michiardi and R. Molva. Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In *the IFIP TC6/TC11 6th Joint Working Conference on Communications and Multimedia Security*, 2002.

[9] L. Xiong and L. Liu. PeerTrust: Supporting reputation-based trust in peer-to-peer communities. *IEEE Trans. on Data and Knowledge Engineering*, 16(7):843–857, 2004.