

Gossip-based aggregation of trust in decentralized reputation systems

Yoram Bachrach · Ariel Parnes · Ariel D. Procaccia · Jeffrey S. Rosenschein

Published online: 9 December 2008
Springer Science+Business Media, LLC 2008

Abstract Decentralized Reputation Systems have recently emerged as a prominent method of establishing trust among self-interested agents in online environments. A key issue is the efficient aggregation of data in the system; several approaches have been proposed, but they are plagued by major shortcomings. We put forward a novel, decentralized data management scheme grounded in gossip-based algorithms. Rumor mongering is known to possess algorithmic advantages, and indeed, our framework inherits many of their salient features: scalability, robustness, a global perspective, and simplicity. We demonstrate that our scheme motivates agents to maintain a very high reputation, by showing that the higher an agent's reputation is above the threshold set by its peers, the more transactions it would be able to complete within a certain time unit. We analyze the relation between the amount by which an agent's average reputation exceeds the threshold and the time required to close a deal. This analysis is carried out both theoretically, and empirically through a simulation system called *GossipTrustSim*. Finally, we show that our approach is inherently impervious to certain kinds of attacks.

Keywords Reputation systems · Trust · Gossip · Manipulation · Game theory

A preliminary version of this article appeared in the proceedings of IJCAI 2007.

Y. Bachrach · A. Parnes · J. S. Rosenschein
School of Engineering and Computer Science, The Hebrew University of Jerusalem, Jerusalem, Israel
e-mail: yori@cs.huji.ac.il

A. Parnes
e-mail: arielp02@cs.huji.ac.il

J. S. Rosenschein
e-mail: jeff@cs.huji.ac.il

A. D. Procaccia (✉)
Microsoft Israel R&D Center, Herzeliya, Israel
e-mail: arielpro@gmail.com

1 Introduction

In open multiagent environments, self-interested agents are often tempted to employ deceit as they interact with others. Fortunately, dishonest agents can expect their victims to retaliate in future encounters. This “shadow of the future”, as Axelrod termed it [7], motivates cooperation and trustworthiness.

However, as the size of the system grows, agents have a progressively smaller chance of dealing with another agent that they already know; as a consequence, building trust in domains populated by numerous agents becomes much harder. One prominent example of such an environment is the internet, which has the important benefit of allowing interaction among entities that are not familiar with one another. Since the value of a transaction depends heavily on the reliability of the parties involved in it, it is very important in such environments to have mechanisms that allow for the obtaining of information regarding parties involved in a possible transaction.

Reputation systems address this problem by collecting, maintaining, and spreading reports among agents, so that agents may learn from others’ experience. To put it differently, agents are intimidated by the “shadow of the future” today, even though tomorrow they are most likely to meet total strangers. Also, such information may allow an agent to *choose* the most suitable party with whom to carry out a transaction. Such systems have been used in a wide variety of applications: electronic commerce, auctions, and many peer-to-peer systems.

Reputation systems can be decomposed into two major components: (1) the trust model, which describes whether an agent is trustworthy, and (2) the data management scheme. The latter component poses some interesting questions, since it is imperative to efficiently aggregate trust-related information in the system. A simple solution is to maintain a central database that contains the feedback gathered from past transactions.

Unfortunately, the centralized solution is inappropriate in distributed environments where scalability is a major concern, as the database soon becomes a bottleneck of the system. Moreover, this approach is not robust to failures. Previous work on *decentralized* reputation schemes suffered from other major problems: agents have to maintain complex data structures, evaluation of trust is based only on local information, or there are restrictive assumptions on the trust model.¹

We approach this difficult area by designing a novel method of trust aggregation (i.e., a reputation system’s data management scheme). The method is demonstrated in this article for a simple trust model, but it can be extended to more complex models.

The roots of our *gossip-based* approach can be traced to a seminal paper by Frieze and Grimmett [19]: a rumor starts with one agent; at each stage, each agent that knows the rumor spreads it to another agent chosen uniformly at random. The authors show that the rumor reaches all agents quickly (a result that coincides with real life).

We directly rely on more recent results, surveyed in the next section. It has been shown that aggregate information, such as averages and sums of agents’ inputs, can be calculated using similar methods of uniform gossip in a way that scales gracefully as the number of agents increases. Furthermore, the approach is robust to failures, and the results hold even when one cannot assume a point-to-point connection between any two agents (as one cannot assume in peer-to-peer [P2P] networks).

In our setting, each agent merely keeps its private evaluation of the trustworthiness of other agents, based on its own interactions.² When an agent wishes to perform a transaction

¹ The “or” is not exclusive.

² The question of how agents set this valuation is outside the scope of this article.

with another, it obtains the *average* evaluation of the other's reputation from all agents in the system, using a gossip-based technique.

Although the algorithms we present estimate the *average* reputation, they can be easily adapted to estimating whether a certain agent has a high reputation in the eyes of the *majority* of the agents, or certain other similar metrics. Thus, the framework we advocate for aggregating reputation information accommodates more sophisticated trust models.

Some advantages are immediately self-evident. Each agent stores very little information, which can be simply and efficiently organized, and evaluation of trust is based on global information. Additionally, this framework inherits the advantages of gossip-based algorithms: scalability, robustness to failure, decentralization, and as a consequence, applicability in peer-to-peer networks.

We have implemented a system called *GossipTrustSim*, for simulating transactions among agents who use our suggested procedure to decide whether or not to perform a transaction. This simulation system is used to investigate some of the properties of our suggested method.

An important desideratum that one would like a reputation system to satisfy is motivating agents to maintain an untarnished reputation, i.e., to be *absolutely trustworthy* (as opposed to, say, being generally trustworthy but occasionally cheating). We claim that our data management scheme, together with an extremely simple trust model, satisfies this property, by showing that the higher an agent's reputation is above the threshold set by the peers with whom it wishes to interact, the more transactions it will be able to complete within a certain time unit. This claim is supported by both theoretical results, as well as by empirical results obtained using *GossipTrustSim*.

In order to achieve their goals, reputation systems must be hard to manipulate. In such systems, agents may attempt to deviate from the protocol of interaction in order to artificially inflate a certain agent's reputation. We demonstrate that our scheme is *inherently resistant to some attacks* (with no assumptions on the trust model). This is a positive side effect of the exponential convergence rates of the algorithms we use.

In this article we do *not* address the problem of designing a trust model. Rather, we suggest an approach for agents to aggregate distributed trust information so as to decide with whom to carry out transactions.

The article proceeds as follows. Section 2 discusses gossip based methods for aggregating information, which lie at the core of the method we suggest. Section 3 introduces our method of gossip-based aggregation of trust information for decentralized reputation systems. In Sect. 4 we theoretically analyze how keeping a very high reputation allows an agent to shorten the time it requires to achieve successful transactions. Section 5 discusses simulation results that also support this claim. In Sect. 6 we consider certain kinds of attacks that can be used by agents who attempt to manipulate reputation systems, and show that our method is resistant to such attacks. Section 7 discusses important related work. We conclude in Sect. 8.

2 Gossip-based information aggregation

In this section, we survey the relevant results of Kempe et al. [24]. These algorithms allow us to estimate the average of values held at network nodes (in our case, these values will be the reputation values concerning a particular agent). Kempe et al. [24] also shows how to calculate other functions over these values, such as the majority function and sum. Thus our algorithms can be adapted for other, more sophisticated models of trust.

2.1 Push-Sum

We begin by describing a simple algorithm, PUSH-SUM, to compute the average of values at nodes in a network. There are n nodes in the system, and each node i holds an input $x_i \geq 0$. At time t , each node i maintains a *sum* $s_{t,i}$ and a *weight* $w_{t,i}$. The values are initialized as follows: $s_{0,i} = x_i, w_{0,i} = 1$. At time 0, each node i sends the pair $s_{0,i}, w_{0,i}$ to itself; at every time $t > 0$, the nodes follow the protocol given as Algorithm 1.

Algorithm 1

- 1: **procedure** PUSH-SUM
 - 2: Let $\{(\hat{s}_l, \hat{w}_l)\}_l$ be all the pairs sent to i at time $t - 1$
 - 3: $s_{t,i} \leftarrow \sum_l \hat{s}_l$
 - 4: $w_{t,i} \leftarrow \sum_l \hat{w}_l$
 - 5: Choose a target $f_t(i)$ uniformly at random
 - 6: Send the pair $(\frac{1}{2}s_{t,i}, \frac{1}{2}w_{t,i})$ to i and to $f_t(i)$
 - 7: $\frac{s_{t,i}}{w_{t,i}}$ is the estimate of the average at time t
 - 8: **end procedure**
-

2.2 Convergence and the diffusion speed of uniform gossip

Let $U(n, \delta, \epsilon)$ (the *diffusion speed* of uniform gossip) be an upper bound on the number of turns PUSH-SUM requires so that for all $t \geq U(n, \delta, \epsilon)$ and all nodes i ,

$$\frac{1}{\sum_k x_k} \cdot \left| \frac{s_{t,i}}{w_{t,i}} - \frac{1}{n} \sum_k x_k \right| \leq \epsilon$$

(the relative error is at most ϵ) with probability at least $1 - \delta$.

Theorem 1 ([24])

1. $U(n, \delta, \epsilon) = O(\log n + \log \frac{1}{\delta} + \log \frac{1}{\epsilon})$.
2. The size of all messages sent at time t by PUSH-SUM is $O(t + \max_i \text{bits}(x_i))$, where $\text{bits}(x_i)$ is the number of bits in the binary representation of x_i .

2.3 Advantages of Push-Sum

A major advantage of gossip-based algorithms is their robustness to failures: the aggregation persists in the face of failed nodes, permanent communication failures, and other unfortunate events. Further, no recovery action is required. The assumption is that nodes can detect whether their message has reached its destination; PUSH-SUM is modified so that if a node detects its target failed, it sends its message to itself.

Theorem 2 ([24]) Let $\mu < 1$ be an upper bound on the probability of message loss at each time step, and let U' be the diffusion speed of uniform gossip with faults. Then:

$$U'(n, \delta, \epsilon) = \frac{2}{(1 - \mu)^2} U(n, \delta, \epsilon).$$

In several types of decentralized networks, such as P2P networks, point-to-point communication may not be possible. In these networks, it is assumed that at each stage nodes

Algorithm 2

```

1: procedure EVAL-TRUST( $i, j, \delta, \epsilon$ ) ▷  $i$  evaluates  $\bar{r}^j$  with accuracy  $\epsilon$ , confidence  $1 - \delta$ 
2:   for all  $k \in N$  do
3:      $x_k \leftarrow r_k^j$  ▷ Inputs to PUSH-SUM are  $j$ 's reputation w.r.t. agents
4:   end for
5:   run PUSH-SUM for  $U = U(n, \delta, \epsilon)$  stages
6:   return  $\frac{s_{U,i}}{w_{U,i}}$ 
7: end procedure

```

send messages to all their neighbors (*flooding*). When the underlying graph is an expander, or at least expected to have good expansion, results similar to the above can be obtained. Fortunately, it is known that several peer-to-peer topologies induce expander graphs [25, 29].

2.4 Push-Sum notes

In the rest of the article, we have $x_i \leq 1$, and in particular $\sum_i x_i \leq n$. Therefore, it is possible to redefine U to be an upper bound on the number of turns required so that for all $t \geq U$ and all nodes i , the *absolute error* $\left| \frac{s_{t,i}}{w_{t,i}} - \frac{1}{n} \sum_k x_k \right|$ is at most ϵ with confidence $1 - \delta$, and it still holds that $U(n, \delta, \epsilon) = O(\log n + \log \frac{1}{\delta} + \log \frac{1}{\epsilon})$. Hereinafter, when we refer to U we have this definition in mind.

Remark 1 The protocol PUSH-SUM is presented in terms of a synchronized starting point, but this assumption is not necessary. A node that poses the query may use the underlying communication mechanism to inform all other nodes of the query; convergence times are asymptotically identical.

3 Our framework

Let the set of agents be $N = \{1, \dots, n\}$. Each agent $i \in N$ holds a number $r_i^j \in [0, 1]$ for each agent $j \in N$ (including itself). This number represents j 's reputation with respect to i , or to put it differently, the degree to which i is willing to trust j . As agents interact, these assessments are repeatedly updated. We do not in general concern ourselves with how agents set these values.

When an agent i is deliberating whether to deal with another agent j , i wishes to make an informed evaluation of the other's reputation. Let $\bar{r}^j = \frac{\sum_k r_k^j}{n}$ be the average of j 's reputation with respect to all agents. Knowledge of \bar{r}^j would give i a good idea of how trustworthy j is (this is, of course, a simple model of trust).

3.1 Gossip-based trust

We show that in the above scenario, agents can use gossip-based algorithms to decide with whom to carry out transactions. Also, in such a setting, agents are encouraged to keep a completely untarnished reputation. Similar results can be obtained for more complex trust models.

A simple way to compute the average trust is via PUSH-SUM.

The protocol EVAL-TRUST is given as Algorithm 2. PUSH-SUM is executed for $U = U(n, \delta, \epsilon)$ stages. At time U , it holds for all $k \in N$, and in particular for agent i , that $\left| \frac{s_{t,i}}{w_{t,i}} - \bar{r}^j \right| \leq \epsilon$, with probability $1 - \delta$. In other words, the algorithm returns a very good approximation of j 's average reputation.

3.2 Use in practice

In practice, when two agents i and j interact, i may evaluate j 's reputation (and vice versa) by calling EVAL-TRUST. The protocol quickly returns the approximation of \bar{r}^j , based on the values r_k^j at the time EVAL-TRUST was called. Each agent i keeps different values $s_{t,i}$ and $w_{t,i}$ for every different query that was issued by some other agent in the system, and updates these values repeatedly according to PUSH-SUM. Thus, at any stage every agent participates in many parallel executions of PUSH-SUM.³

A possible cause for concern is the amount of communication each agent has to handle at every turn. However, the quick convergence of PUSH-SUM implies that often the burden would not be too great. Indeed, assume that the number of new interactions at each turn is bounded by a constant c (or at worst is very small compared to n). Each such new interaction results in at most two new executions of EVAL-TRUST, but the execution lasts at most U turns. To conclude the point, under the foregoing assumption, each agent sends at most $c \cdot U = O(\log n)$ messages per turn.

Remark 2 The size of messages depends on how the r_i^j are calculated, and as mentioned above, this issue is outside the scope of this article. Nevertheless, there would usually be a constant number of reputation levels (say, for instance, that r_j^i can assume the values $r_j^i \in \{0, 0.1, 0.2, \dots, 1\}$), so the message size would normally be constant.

As the above method of aggregating an agent's average reputation relies on the gossip-based algorithm PUSH-SUM, it inherits all the latter's benefits, in particular robustness to failure and applicability in peer-to-peer networks.

4 The benefit of an unstained reputation

It is very desirable (indeed, crucial) that a reputation system be able to induce truthfulness in agents. Naturally, an agent with a stained reputation would be shunned by its peers, while an agent with a good reputation would easily solicit deals and transactions. A further step in this direction is motivating agents *never* to cheat.

4.1 Reputation and speed of transaction

An agent with a *generally* good reputation, one that is slightly above the required threshold of its peers, that only occasionally cheats, would probably be able to win the confidence of its peers; there is seemingly no reason why an agent should not play false now and again. Nevertheless, we consider in this section an extremely simple and general trust model, and

³ Throughout the article, we consider the case where each agent has a trust value for every other agent. In many domains, some agents are simply unknown to some other agents (who thus hold no reputation information regarding those agents). One possible way to overcome this is to use a strategy similar to handling node failures, where the agent simply sends the Push-Sum message to itself in the case where the target of the message is a failed node.

show that with the data management scheme that we have presented, there is a social benefit to having a very high reputation: the higher the agent’s reputation, the shorter the time required to close deals. Consider for example a computational grid, where agents provide a certain computational service at a cost, and where each such computation itself is not extremely time-consuming. Computational resources are wasted while an agent is not involved in a transaction, so in order to maximize its utility, an agent should be able to perform as many transactions as possible in a given time unit. Thus, the speed at which transactions are agreed upon is very important. Another example is file sharing in peer-to-peer networks, where information is provided at a cost. Again, any time spent while resources are idle (in this case, time spent when network resources are not used to transmit information) is wasted, so agents are incentivized to spend as little time as possible closing deals.

We consider a model in which each agent i has a reputation threshold r_i^{thr} (similar to [33]) and a confidence level δ_i : agent i is willing to deal with an agent j if and only if i knows that j ’s average reputation is at least r_i^{thr} , with confidence $1 - \delta_i$. Agent i evaluates j ’s reputation as above, using EVAL- TRUST. Recall that when the algorithm terminates, agent i only has an ϵ -close approximation of \bar{r}^j . If $\frac{s_{t,i}}{w_{t,i}}$ is very close to r_i^{thr} , i would have to increase the accuracy.

Remark 3 We still do not commit to the way the values r_i^j are determined and updated, so the above trust model is quite general.

Algorithm 3

```

1: procedure DECIDE- TRUST( $i, j$ )                                ▷  $i$  decides if it wants to deal with  $j$ 
2:    $\epsilon \leftarrow 1/2$                                              ▷ Initialization
3:    $k_1 \leftarrow 0$ 
4:   loop
5:      $k_2 \leftarrow U(n, \delta_i, \epsilon)$ 
6:     run EVAL- TRUST( $j$ ) for another  $k_2 - k_1$  stages           ▷ A total of  $k_2$  stages
7:     if  $s_{t,i}/w_{t,i} < r_i^{thr} - \epsilon$  then
8:       return false
9:     else if  $s_{t,i}/w_{t,i} > r_i^{thr} + \epsilon$  then
10:      return true
11:    end if
12:     $k_1 \leftarrow k_2$ 
13:     $\epsilon \leftarrow \epsilon/2$ 
14:  end loop
15: end procedure

```

The procedure DECIDE- TRUST, given as Algorithm 3, is a straightforward method of determining whether $\bar{r}^j \geq r_i^{thr}$. Agent i increases the accuracy of the evaluation by repeatedly halving ϵ , until it is certain of the result. In this context, a stage of EVAL- TRUST corresponds to a stage of PUSH- SUM.

4.2 Theoretical analysis of speed of transaction

Proposition 3 *Let $i, j \in N$, and $\Delta_{ij} = |\bar{r}^j - r_i^{thr}|$. With probability at least $1 - \delta_i$, DECIDE- TRUST correctly decides whether agent j ’s reputation is at least r_i^{thr} after $O(\log n + \log \frac{1}{\delta_i} + \log \frac{1}{\Delta_{ij}})$ stages of EVAL- TRUST (Fig. 1).⁴*

⁴ The probability is the chance that the algorithm will answer incorrectly; the bound on the number of stages is always true.

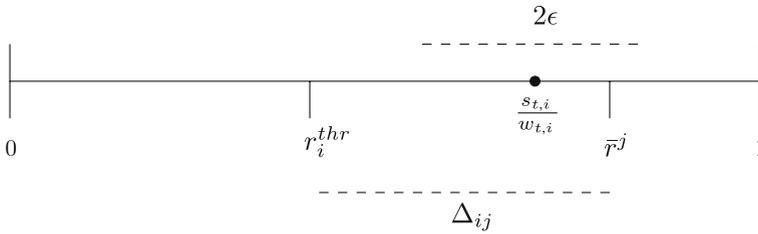


Fig. 1 Illustration for the proof of Proposition 3. Once $\epsilon < \Delta_{ij}/2$ and $s_{t,i}/w_{t,i} > r_i^{thr} + \epsilon$, it holds with probability $1 - \delta_i$ that $\bar{r}^j \geq r_i^{thr}$

Proof Assume w.l.o.g. that $r_i^{thr} < \bar{r}^j$, and that the algorithm reached a stage t_0 where $\epsilon < \Delta_{ij}/2$. At this point, it holds that $|\frac{s_{t,i}}{w_{t,i}} - \bar{r}^j| \leq \epsilon$ (with probability $1 - \delta_i$), and therefore:

$$\begin{aligned} \frac{s_{t,i}}{w_{t,i}} &\geq \bar{r}^j - \epsilon \\ &= r_i^{thr} + \Delta_{ij} - \epsilon \\ &> r_i^{thr} + \epsilon. \end{aligned}$$

Hence, the algorithm surely terminates when $\epsilon < \Delta_{ij}/2$. Now the proposition follows directly from the fact that $U(n, \delta_i, \Delta_{ij}) = O(\log n + \log \frac{1}{\delta_i} + \log \frac{1}{\Delta_{ij}})$, so we give the correct answer with the required probability. \square

To conclude, Proposition 3 implies that there is a benefit for agent j in maintaining a high reputation: for any agent i with a reasonable threshold, Δ_{ij} is significant, and this directly affects the running time of DECIDE-TRUST. As explained at the beginning of this section, this is important in many real world domains, such as computations in grids or file sharing in peer-to-peer networks. Due to these theoretical results, we can see that in such settings it is worthwhile for agents not only to have a reputation high enough to perform transactions, but to maintain a reputation significantly above this required threshold, in order to increase the rate at which these transactions are performed.

Remark 4 The result is limited, though, when the number of agents n is large, as the time to evaluate an agent’s reputation is also proportional to $\log n$.

5 Reputation and speed of transactions: simulation results using *GossipTrustSim*

In Sect. 4 we discussed how a highly positive reputation has a social benefit for agents, as it shortens the time an agent needs to close deals. In that model, agent i was willing to interact with agent j if and only if i believed j ’s average reputation was above a certain threshold value r_i^{thr} with confidence level of δ_i . We introduced a simple algorithm, *Decide-Trust*, that agents can use in this model. Using the algorithm, an agent can bound its probability of accepting a transaction by mistake (i.e., the probability of the agent deciding that $\bar{r}^j > r_i^{thr}$ when in fact it is not) by a desired confidence level, δ_i . The algorithm automatically chooses the number of Push-Sum steps to perform, based on the Push-Sum results obtained along the way.

We denoted the amount by which an agent j ’s average reputation truly is above the desired transaction threshold required by agent i as $\Delta_{ij} = |\bar{r}^j - r_i^{thr}|$. In Sect. 4 we have shown that

the higher an agent's reputation is above the required threshold (the higher Δ_{ij} is), the smaller the number of Push-Sum steps that are required. More precisely, the number of *Eval-Trust* calls is $O(\log n + \log \frac{1}{\delta_i} + \log \frac{1}{\Delta_{ij}})$. We now provide the results of several simulations, that precisely characterize how the time required to close a deal depends on Δ_{ij} .

Although our empirical results only concern certain models, they indicate that agents who are significantly above the required reputation threshold can indeed significantly shorten the time they need to close a single deal, and are thus able to complete more deals per given time unit. These simulation results complement our theoretical analysis in the previous section.

We built a system called *GossipTrustSim*, for simulating transactions among agents that use *Decide-Trust* (Algorithm 3) to decide whether or not to perform a transaction. *Gossip-TrustSim* has been implemented in Java, and includes an implementation of *Push-Sum* and of our suggested *Eval-Trust* and *Decide-Trust* (Algorithms 2 and 3). It supported various simulations of the models we suggest in this article.

5.1 Empirical analysis of speed of transaction

The simulations we present here deal with a given agent, x , who is interested in a transaction with agent j . In each such simulation, for each agent i the system randomly chooses r_i^j , j 's reputation in the eyes of i (i.e., the degree to which i is willing to trust j), according to a certain model; in all simulations presented in this article we use a simple model, where r_i^j is drawn from the same distribution for any agent i . Once the r_i^j values are set for all agents i , we then simulate what happens when agent x attempts to decide whether to interact with agent j , using our suggested *Decide-Trust* algorithm. Each such simulation runs *Decide-Trust* until agent x decides whether to trust agent j or not, and we count the total number of Push-Sum steps performed.

The simulation contain 2500 agents. The initial value for ϵ is set to $\epsilon_1 = 0.5$. Since this is a very large value, EVAL- TRUST is typically called more than once. The threshold value in our simulation is $r_x^{thr} = 0.5$, so agent x will accept the transaction only if j 's average reputation is $\bar{r}^j > 0.5$. Our confidence level is $\delta_x = 0.1$, so agent x will accept the transaction only if $\bar{r}^j > r_x^{thr}$ with a probability of at least $1 - \delta_x = 0.9$.

In our first simulation, each r_i^j is distributed uniformly between a minimal value a and a maximal value $b = 1$. Thus, its expectation is $\mathbb{E}[r_i^j] = \frac{1+a}{2}$. Since there are many agents, the average reputation of agent j is an approximately normal distribution. We simulate the *Decide-Trust* run when in fact $\bar{r}^j > r_x^{thr}$ (and a is chosen so that this is indeed the case), so agent x should eventually accept the transaction. It is easy to see that the higher a is, the higher $\mathbb{E}[r_i^j]$ for all agents i , and the higher $\mathbb{E}[\Delta_{xj}]$ is.

We used *GossipTrustSim* to find the required number of Push-Sum steps for agent x to approve the transaction with agent j . We did this for several values of a , and for each value we repeated the simulation 100 times, then calculated the average number of required Push-Sum steps. The results are presented in Fig. 2, which shows the relationship found between Δ_{xj} , the amount by which agent j 's reputation exceeds the threshold required by agent x , and the number of Push-Sum steps agent x has to perform before approving the transaction.

It is easy to see from Fig. 2 that agent j 's situation is indeed better when its average reputation is well above agent x 's threshold. The higher its reputation is, the fewer Push-Sum steps are required for agent x to accept the transaction. This means that although an agent can successfully complete transactions with another agent when its reputation is only slightly above the other agent's required threshold, it would be able to complete more transactions per

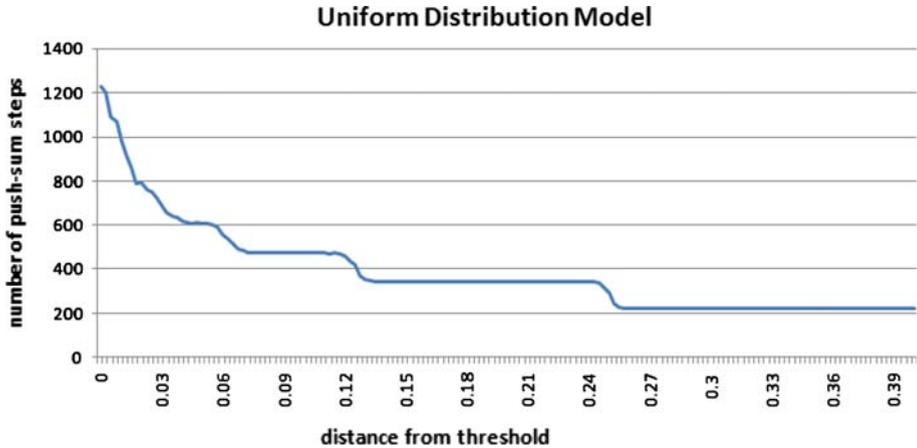


Fig. 2 Uniform distribution model: The relation between the distance from the threshold, Δ_{xj} , and the required number of Push-Sum steps

time unit by making sure its reputation is much higher than that threshold value. However, the speedup achieved by increasing the reputation diminishes when the reputation is well above the threshold.

Another thing to note is that there are “bumps” in the graph, where suddenly a small increase in the distance from the threshold, Δ_{xj} , results in a large decrease in the required number of Push-Sum steps. This occurs since Algorithm 3, Decide-Trust, calls Algorithm 2, Eval-Trust, several times, and halves the value of ϵ each time. Since each such call to Eval-Trust requires several Push-Sum steps, typically when one less call to Eval-Trust is required, there is a significant decrease in the required number of Push-Sum steps.

In our second simulation, each r_i^j has a normal distribution, with expectation μ and variance $\sigma^2 = 0.0001$. As before, since there are many agents, the average reputation of agent j is also an approximately normal distribution, but with a much smaller variance. Again, μ is chosen such that in fact $\bar{r}^j > r_x^{thr}$ (so we only simulate for values $\mu > r_x^{thr}$), so agent x should eventually accept the transaction. The higher μ is, the higher $\mathbb{E}[\Delta_{xj}]$ is. As before, we tested the number of Push-Sum steps required for agent x to approve the transaction with agent j for several values of μ . For each value, we repeated the simulation 100 times, and calculated the average number of required Push-Sum steps. The results are presented in Fig. 3, which shows the relationship found between Δ_{xj} and the number of required Push-Sum steps.

Similarly to the uniform distribution model, Fig. 3 shows that agent j 's situation is better when its average reputation is well above agent x 's threshold. In this model as well, the higher its reputation is, the fewer Push-Sum steps are required for agent x to accept the transaction, so it can complete more transactions per time unit.

To conclude, the simulations support our claim that when using the gossip-based approach, agents are incentivized to ensure that their reputation is well above the threshold value needed to complete transactions. Agents that occasionally cheat may still be able to complete transactions, as long as their reputation does not drop below the threshold chosen by the agents with whom they wish to interact, but they would be able to perform fewer transactions per time unit.

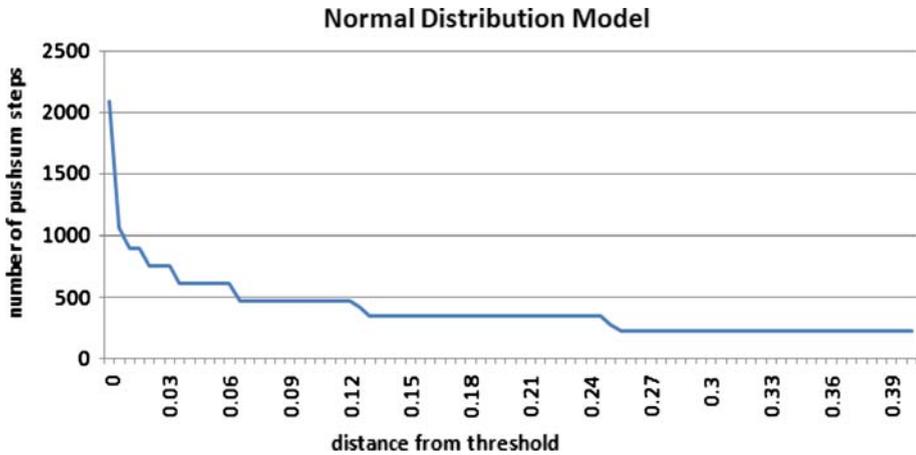


Fig. 3 Normal distribution model

6 Resistance to attacks

We have seen that information about an agent’s reputation can be efficiently propagated, as long as all agents consistently follow EVAL- TRUST. However, with reputation systems we are usually dealing with self-interested agents. In our context, a manipulative agent may artificially increase or decrease the overall evaluation of some agent’s reputation by deviating from the protocol. Several such manipulations are mentioned in Sect. 7. We now examine some possible manipulations that can be carried out when agents are using a reputation system such as the one we suggest here, and show that our methods are resistant to such attacks.

In the framework we have presented, trust is evaluated on the basis of global knowledge, i.e., the average of all reputation values in the system. Therefore, any small coalition cannot significantly change the average reputation of some agent j by setting their own valuations r_i^j to legal values in $[0, 1]$, and then following the protocol EVAL- TRUST.⁵

6.1 Arbitrarily setting reputation

An interesting setting which should be considered arises when a manipulator is allowed to set its reputation value arbitrarily. As a simple motivating example, consider a setting where agents propagate agent j ’s average reputation ($x_i = r_i^j$ for all i), and a manipulator i^m wants to ensure that for all i , $\frac{s_{t,i}}{w_{t,i}}$ converges to a high value as the time t increases. At some stage t_0 , the manipulator updates s_{t_0,i^m} to be n , but except for this harsh deviation follows the protocol to the letter. In particular, the manipulator might initially set $r_{i^m}^j = x_{i^m} = n$. We refer to this strategy as *Strategy 1*. Clearly, for all i , $\frac{s_{t,i}}{w_{t,i}}$ eventually converges to a value that is at least 1.

Despite the apparent effectiveness of Strategy 1, it is easily detected. Indeed, unless for all $i \neq i^m$ it holds that $s_{t_0,i} = 0$ at the time t_0 when the manipulator deviated by assigning $s_{t_0,i^m} = n$, the expressions $\frac{s_{t,i}}{w_{t,i}}$ would eventually converge to a value that is strictly greater than 1; this would clearly unmask the deceit. It is of course possible to update s_{t_0,i^m} to be less than n , but it is difficult to determine *a priori* which value to set without pushing the average reputation above 1.

⁵ In fact, this holds for every coalition that does not constitute a sizable portion of the entire set of agents.

6.2 Empirical analysis of manipulations

We now provide empirical results regarding manipulations that arbitrarily set an agent's reputation. Specifically, we investigate how the magnitude of the manipulation (which also makes it more easily detectable) is related to its effect, as measured by the inflation in the target agent's reputation, and how this relates to the accuracy required of our procedure.

In our setting, a certain manipulator attempts to artificially inflate the reputation of a target agent j using *Strategy 1*. As explained in Sect. 6.1, the manipulator i^m sets s_{t_0, i^m} to be extremely high. We measure the magnitude of the manipulation as $M = \frac{s_{t_0, i^m}}{n}$. When many Push-Sum steps are performed, the estimated average trust $\frac{s_{r,i}}{w_{r,i}}$ (for any i) would eventually converge to a value of at least M . However, our method is designed to run as few Push-Sum steps as possible, so the manipulation is typically not that effective. For a certain run regarding agent j , we can measure the effectiveness of a manipulation as the amount by which the target agent's reputation was artificially inflated, i.e., the difference between the true average trust towards the target agent, $\bar{r}^j = \frac{\sum_k r_k^j}{n}$, and the estimate returned by the procedure, $E = \frac{s_{r,i}}{w_{r,i}} - \bar{r}^j$.

Consider a manipulator i^m that attempts to inflate the reputation of a target agent j who has a very low average trust, \bar{r}^j . When many Push-Sum steps are performed, $\frac{s_{r,i}}{w_{r,i}}$ would converge to a value of at least M , so E should be very close to M (as \bar{r}^j is very small). However, when very few Push-Sum steps are performed, much of the weight remains "concentrated" on few agents, and is not distributed among all the agents. In such a case, E is likely to be less than M . We have used GossipTrustSim to empirically analyze this effect.

The simulation was carried out in the following way, in a setting similar to that of Sect. 5.1. We considered a system with 100 agents, and a manipulator i^m that attempts to inflate the reputation of the target agent j . The reputations of each agent i towards j was chosen using the normal distribution, with expectation $\mu = 0.3$ and variance $\sigma^2 = 0.01$. The manipulator used *Strategy 1* of Sect. 6.1, and has set s_{t_0, i^m} to a very high value. We tested the relationship between the magnitude of the manipulation $M = \frac{s_{t_0, i^m}}{n}$ and its effectiveness $E = \frac{s_{r,i}}{w_{r,i}} - \bar{r}^j$. For each measured manipulation magnitude M , we performed 5000 repetitions, and in each measured the manipulation effectiveness E . Each of the graphs in Fig. 4 presents the relationship between M and E . The effectiveness of the manipulation, E , is affected by the number of Push-Sum steps performed. When we require the procedure to be very accurate, more Push-Sum steps have to be performed, as explained in Sect. 2.2, so the effectiveness of the procedure increases. All of the graphs below show the relationship between M and E given our procedure, and all of them have used a confidence $\delta = 0.1$, but they differ in the accuracy required for the procedure. We used values of 0.3, 0.03, 0.003, 0.0003 for the accuracy ϵ .

Figure 4 indicates that although higher magnitudes of manipulation are more effective, the inflation in the target agent's reputation E is smaller than the magnitude of manipulation M . The reason for this is that the Push-Sum procedure does not manage to fully distribute the excess weight set by the manipulator, s_{t_0, i^m} , across all agents. The simulations also show that the higher the accuracy required of our procedure, the more effective the manipulation is, as more Push-Sum steps are performed, so the excess weight is better distributed.

In Sect. 6.1 we have noted that *Strategy 1* is a manipulation that is very easy to detect. In fact, the higher the magnitude of the manipulation, the more likely it is that it would be detected. Since the above simulations indicate that such a manipulation can indeed artificially inflate an agent's reputation, and since such manipulations are easily detectable, we believe any reputation system implementing gossip mechanisms similar to those we present here

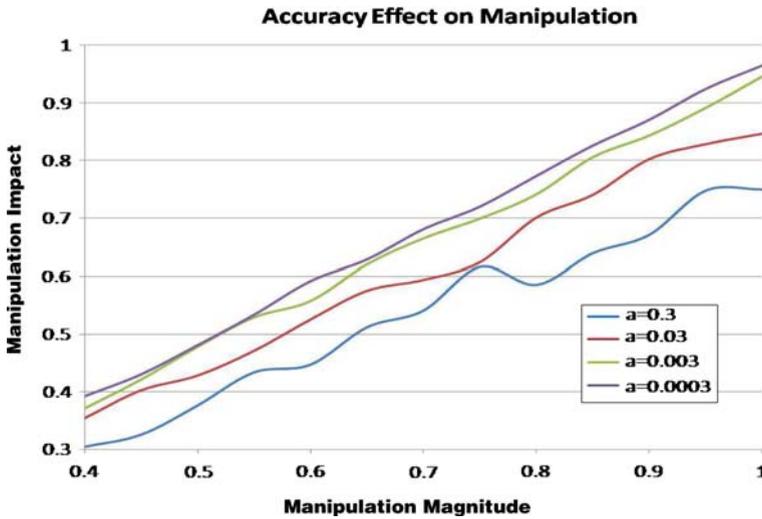


Fig. 4 The relation between the magnitude of the manipulation and its effectiveness, for various values of required accuracy of the procedure

should also include simple safeguards against such a manipulation. In the following section, we investigate another manipulation strategy that is harder to detect.

6.3 A more subtle approach to manipulating reputation

We now consider a more subtle way to increase the values $\frac{s_{t,i}}{w_{t,i}}$, a deceit that is indeed difficult to detect; we call this strategy *Strategy 2*. For the first T stages of the algorithm, the manipulator i^m follows PUSH-SUM as usual, with the exception of the updates of s_{t,i^m} : after updating $w_{t,i^m} = \sum_l \hat{w}_l$ (as usual), i^m updates: $s_{t,i^m} = w_{t,i^m}$. In other words, the manipulator sets its personal evaluation of the average $\frac{s_{t,i^m}}{w_{t,i^m}}$ to be 1 at every stage $t = 1, \dots, T$. For time $t > T$, the manipulator abides by the protocol.

Using this strategy, it always holds that $\frac{s_{t,i}}{w_{t,i}} \leq 1$ for all i . In addition, for all t , it still holds that $\sum_i w_{t,i} = n$. Therefore, without augmenting the system with additional security measures, this manipulation is difficult to detect. We shall presently demonstrate formally that the manipulation is effective in the long run: $\frac{s_{t,i}}{w_{t,i}}$ converges to 1 for all i .

Proposition 4 Under Strategy 2, for all $i \in N$, $\frac{s_{2T,i}}{w_{2T,i}} \xrightarrow{T \rightarrow \infty} 1$ in probability.

Proof We first notice that $\sum_i s_{t,i}$ is monotonically increasing at stage t . Moreover, as noted above, it holds that at every stage, $\sum_i w_{t,i} = n$, as for all $i \in N$: $\frac{s_{t,i}}{w_{t,i}} \leq 1$, and thus:

$$\sum_i s_{t,i} \leq \sum_i w_{t,i} = n.$$

Let $\epsilon, \delta > 0$. We must show that it is possible to choose T large enough such that for all $t \geq 2T$ and all $i \in N$, $\Pr[\frac{s_{t,i}}{w_{t,i}} \geq 1 - \epsilon] \geq 1 - \delta$.

Assume that at time t it holds that:

$$\frac{\sum_i s_{t,i}}{n} < 1 - \epsilon/2. \tag{1}$$

Let $I_t = \{i \in N : \frac{s_{t,i}}{w_{t,i}} \geq 1 - \epsilon/4\}$, $w(I_t) = \sum_{i \in I_t} w_{t,i}$. It holds that:

$$\begin{aligned} n(1 - \epsilon/2) &\geq \sum_{i \in N} s_{t,i} \\ &\geq \sum_{i \in I_t} s_{t,i} \\ &\geq \sum_{i \in I_t} w_{t,i} \cdot (1 - \epsilon/4) \\ &= w(I_t)(1 - \epsilon/4). \end{aligned}$$

It follows that $w(I_t) \leq n \cdot \frac{1-\epsilon/2}{1-\epsilon/4}$. The total weight of agents in $N \setminus I_t$ is at least $n - w(I_t)$. There must be an agent $i_t \in N \setminus I_t$ with at least a $1/n$ -fraction of this weight:

$$w_{t,i_t} \geq \frac{n - w(I_t)}{n} \geq \frac{\epsilon}{4 - \epsilon}. \tag{2}$$

In order for the choice of i_t to be well-defined, assume i_t is the minimal index that satisfies Eq. 2.

Now, let s'_{t,i^m} be the manipulator’s sum had it updated the sum according to the protocol, i.e., $s'_{t,i^m} = \sum_l \hat{s}_l$ for all messages l sent to i^m . With probability $1/n$ (and independently of other stages), $f_t(i_t) = i^m$; if this happens, it holds that:

$$\begin{aligned} s'_{t+1,i^m} &\leq (w_{t+1,i^m} - 1/2 \cdot w_{t,i_t}) + 1/2 \cdot s_{t,i_t} \\ &\leq (w_{t+1,i^m} - 1/2 \cdot w_{t,i_t}) \\ &\quad + 1/2 \cdot w_{t,i_t} \cdot (1 - \epsilon/4). \end{aligned} \tag{3}$$

For all stages t it holds that $\sum_i s_{t+1,i} - \sum_i s_{t,i} = s_{t+1,i^m} - s'_{t+1,i^m}$, as the manipulator is the only agent that might change $\sum_i s_{t,i}$. Therefore, in the conditions of Eq. 3,

$$\begin{aligned} \sum_i s_{t+1,i} - \sum_i s_{t,i} &= s_{t+1,i^m} - s'_{t+1,i^m} \\ &= w_{t+1,i^m} - s_{t+1,i^m} \\ &\geq 1/2 \cdot w_{t,i_t} \cdot \frac{\epsilon}{4} \\ &\geq \frac{\epsilon^2}{32 - 8\epsilon} \\ &= \Delta(w). \end{aligned}$$

So far, we have shown that for each stage t where Eq. 1 holds and $f_t(i_t) = i^m$, it is the case that $\sum_i s_{t+1,i} - \sum_i s_{t,i} \geq \Delta(w)$. This can happen at most $\frac{n(1-\epsilon/2)}{\Delta(w)}$ times before Eq. 1 no longer holds, or to put it differently, before $\frac{\sum_i s_{t,i}}{n} \geq 1 - \epsilon/2$.

Let X_t be i.i.d. binary random variables, that are 1 if and only if $f_t(i_t) = i^m$. It holds that for all t where Eq. 1 is true, $\mathbb{E}[X_t] = 1/n$. By Chernoff’s inequality, it holds that:

$$\Pr \left[\frac{1}{T_1} \sum_{t=1}^{T_1} X_t \leq \frac{1}{2n} \right] \leq e^{-\frac{T_1}{2n^2}}.$$

It is possible to choose T_1 to be large enough such that this expression is at most $\delta/2$, and in addition $\frac{1}{2n} \cdot T_1 \geq \frac{n(1-\epsilon/2)}{\Delta(w)}$. Therefore, at time T_1 , the average $\frac{\sum_i s_{T_1,i}}{n} \geq 1 - \epsilon/2$ with probability $1 - \delta/2$.

Recall that after T stages (where i^m deviated from the protocol), it still holds that $\sum_i w_{T,i} = n$. Assume that indeed $\frac{\sum_i S_{T_1,i}}{n} \geq 1 - \epsilon/2$. By modifying the proof of Theorem 3.1 from [24], it is possible to show that after another $T_2 = T_2(n, \delta, \epsilon)$ stages where all agents observe the protocol, it holds with probability $1 - \delta/2$ that for all i , $\left| \frac{S_{T_1+T_2,i}}{w_{T_1+T_2,i}} - \frac{\sum_i S_{T_1,i}}{n} \right| < \epsilon/2$, and thus for all i and $t \geq T_1 + T_2$, $\frac{S_{t,i}}{w_{t,i}} > 1 - \epsilon$ with probability $1 - \delta$.

The proof is completed by simply choosing $T = \max\{T_1, T_2\}$. □

Proposition 4 indicates that even this subtle manipulation can increase an agent’s reputation to the maximal value when it is used for long enough periods of time. In the next section we examine how the short convergence times of Push-Sum make our approach resistant to such manipulations.

6.4 Resistance to subtle reputation manipulation

Proposition 4 implies that Strategy 2 poses a provably acute problem, when PUSH-SUM is run a large number of turns. Fortunately, PUSH-SUM converges exponentially fast, and thus it is usually the case that the manipulator is not able to significantly affect the average reputation, as the following proposition demonstrates.

Proposition 5 *Let $T_1 \leq T$. Under Strategy 2 it holds that $\mathbb{E} \left[\frac{\sum_i S_{T_1,i}}{n} - \bar{r}^j \right] \leq \frac{T_1}{2n}$.*

Proof Let $\{\hat{s}_l, \hat{w}_l\}$ be the messages that the manipulator received at time $t + 1$. The manipulator sets $s_{t+1,i^m} = w_{t+1,i^m} = \sum_l \hat{w}_l$. Essentially, this is equivalent to setting for all $l \hat{s}_l = \hat{w}_l$, or in other words, raising each \hat{s}_l by $\hat{w}_l - \hat{s}_l$. At turn t it was already true that $s_{t,i^m} = w_{t,i^m}$ (w.l.o.g. this is also true for $t = 0$), so it is enough to consider messages at time t from all $i \neq i^m$.

Therefore, for all stages t , it holds that:

$$\begin{aligned} \mathbb{E} \left[\sum_i s_{t+1,i} - \sum_i s_{t,i} \right] &= \sum_{i \neq i^m} \left(\Pr[f_t(i) = i^m] \cdot \left(\frac{1}{2} w_{t,i} - \frac{1}{2} s_{t,i} \right) \right) \\ &= \frac{1}{2n} \sum_{i \neq i^m} (w_{t,i} - s_{t,i}) \\ &\leq \frac{1}{2n} \sum_{i \neq i^m} w_{t,i} \\ &\leq \frac{1}{2n} \sum_{i \in N} w_{t,i} \\ &= \frac{1}{2}. \end{aligned}$$

The last equality follows from the fact that for all t , $\sum_i w_{t,i} = n$.

As $\bar{r}^j = \frac{\sum_i s_{0,i}}{n}$, and from the linearity of expectation, we obtain that

$$\mathbb{E} \left[\frac{\sum_i S_{T_1,i}}{n} - \bar{r}^j \right] = \frac{1}{n} \mathbb{E} \left[\sum_{t=0}^{T_1-1} \left(\sum_i s_{t+1,i} - \sum_i s_{t,i} \right) \right]$$

$$\begin{aligned}
&= \frac{1}{n} \sum_{t=0}^{T_1-1} \mathbb{E} \left[\sum_i s_{t+1,i} - \sum_i s_{t,i} \right] \\
&\leq \frac{1}{n} T_1 \cdot \frac{1}{2}.
\end{aligned}$$

□

In particular, since $U(n, \delta, \epsilon) = O(\log n + \log \frac{1}{\delta} + \log \frac{1}{\epsilon})$, PUSH-SUM is executed $O(\log n)$ stages, and thus the difference in the average is at most $O(\frac{\log n}{n})$, which is quite insubstantial.

Remark 5 It is not guaranteed at time T_1 that each $\frac{s_{t,i}}{w_{t,i}}$ is close to \bar{r}^j , because the inputs were dynamically changed during the execution of PUSH-SUM.

Remark 6 The above discussion focused on a setting where the manipulator attempts to increase the average reputation of an agent. It is likewise possible for a manipulator to decrease an agent's average reputation, or indeed set it eventually to any value it wants.

Remark 7 Jelasy et al. [21] propose a general method to completely prevent malicious agents from deviating in gossip-based algorithms, by augmenting the protocol with exchange of certificates. However, the authors describe their approach in a very general manner, and so this approach was not implemented here.

7 Related work

The main focus of research on trust and reputation systems has been on the semantic aspects of these systems, and their effect on social welfare. Previous work has highlighted the advantages of reputation systems in overcoming social pitfalls in several domains. Akerlof [4], for instance, has considered markets where information asymmetry exists between buyers and sellers, in the sense that buyers can only guess the quality of goods; in such a setting, a reputation system can improve social welfare.

Several papers provide detailed background concerning reputation systems [20, 23, 26]. Several such systems are used in practice, so it is possible to observe their actual real-world impact. A prominent reputation system (and one of the earliest) is that used by eBay; the efficiency of eBay-like reputation systems has been studied in [15]. This is an extremely simplistic reputation system, which has several drawbacks [30]. For example, a user may misbehave in only a fraction of his transactions, and still have an increasing total score. Nevertheless, it has been shown [30] that the system works surprisingly well—reputation profiles are predictive of future behavior.

Other research has analyzed manipulations of general reputation mechanisms. Friedman and Resnick [18] have discussed the effects of *cheap pseudonyms*. When agents can enter the system using pseudonyms, and the cost of recreating an identity is cheap, agents that have a stained reputation may easily shed it. The authors have considered several solutions to this problem: disallowing anonymity, entry fees (which make pseudonyms more expensive), and using a central authority for irreplaceable (“once-in-a-lifetime”) pseudonyms. However, each approach has major drawbacks. Sybil attacks are discussed in [10], where an agent creates many false identities in order to boost the reputation of its primary identity.

Dellarocas [14] has studied a setting where agents manipulate a reputation system by providing unfair ratings to some of their peers, and suggests some solutions. Several other mechanisms can make reputation systems resistant to manipulations; [31] discusses scoring rules that help obtain honest feedback, while [6] proposes a method for learning who the good raters are.

The next few paragraphs survey previous work on distributed reputation systems. An early work is that of Abdul-Rahman and Hailes [1], which relied on the results of Marsh [28] to design a model of trust in online environments. In this framework, each agent must maintain and update large data structures, which contain knowledge about the entire system. Updating this data may be inefficient, and in particular it is not certain that the scheme scales well when the number of agents grows.

P2PRep [11] and Xrep [13] are peer-to-peer reputation systems that can be piggybacked on existing peer-to-peer protocols (such as Gnutella). P2PRep allows peers to estimate trustworthiness of other peers by polling. XRep takes another step forward: each peer keeps trust evaluations both of other peers and of resources. No guarantees are given with respect to computational efficiency and scalability.

Aberer and Despotovic [3] introduce a reputation system that consists of both a semantic model and a data management scheme. The latter relies on P-Grid [2], and uses distributed data structures for storing trust information; the associated algorithms scale gracefully as the number of agents increases. In addition, a limited resistance to manipulation and failure is achieved through replication of data. This approach suffers from several shortcomings compared to ours. Agents in this scheme assess others' reputation only on the basis of complaints filed in the past; the framework is generally limited to such binary trust information. In addition, trust is evaluated only according to referrals from neighbors, whereas in our scheme the evaluation is based on all the information in the system.

Xiong and Liu [33] introduced a sophisticated framework specifically applicable in peer-to-peer networks, where the decision whether to trust a peer is based on five metrics: satisfaction, number of transactions, credibility of feedback, transaction context, and community context. This work was extended in [32]. Both papers concentrate on the trust model, and generally do not elaborate on the data management scheme. Specifically, in [33] a P-Grid [2] is used. Therefore, this work is in a sense orthogonal but complementary to ours. Dewan and Dasgupta [17] propose self-certification and IP-Based safeguards as ways of inducing trust; this work also complements ours.

As mentioned in Sect. 1, the roots of our *gossip-based* approach can be traced to a the seminal paper by Frieze and Grimmett, which studied the telephone call gossip problem; this led to a line of work regarding such gossip methods. The current article follows the approach taken by Kempe et al. [24], using the Push-Sum procedure. There are several similar approaches [21, 22, 27].

Gossip-based algorithms⁶ have many applications in other domains. One important application is load balancing in distributed operating systems, such as the distributed operating system MOSIX [8, 9]. This is a management system for Linux clusters and organizational grids, that provides a single-system image (SSI).⁷ Gossip-based approaches are used in MOSIX to balance load across the computing nodes [8, 5]. Yet another application of gossip algorithms is replicated database maintenance [16].

⁶ Also called *epidemic algorithms*.

⁷ This is roughly the equivalent of an operating system for a cluster as a whole.

8 Conclusions and future research

We have presented a data management scheme based on gossip-based algorithms, and have demonstrated that it possesses the following features:

- Decentralization: no central database, and further, applicability in networks where point-to-point communication cannot be assumed.
- Scalability: the time to evaluate an agent's average reputation with confidence $1 - \delta$ and accuracy ϵ is $O(\log n + \log \frac{1}{\delta} + \log \frac{1}{\epsilon})$.
- Robustness to failure.
- Global perspective: evaluation of trust is based on all relevant information in the system, rather than local information.
- Extremely simple data structures: each agent merely keeps an assessment of the agents with which it personally interacted.
- Motivates absolute truthfulness, as the time to close deals may decrease as reputation increases.
- Resistance to some attacks, such as carefully tampering with the updates performed by PUSH-SUM.

We have focused on the data management scheme, and have largely ignored the trust model (with the exception of Sect. 4). However, we believe that many existing trust models can be integrated with our framework. A very simple example is the binary trust model of [3], where agents can file complaints against other agents. In our framework, each agent i sets its value r_i^j to be 0 if it wishes to file a complaint against j ; otherwise, the value is 1. More sophisticated models may require tweaks in the framework. Consider the trust model presented in [33], where five factors are taken into account. Three of the factors mentioned simply determine the way an agent updates its own values r_i^j , and our framework of course supports any update formula. The “number of transactions” factor is already taken into account, as we compute the average reputation. The “credibility of feedback” factor requires a small change: given credibility ratings c_i for agents, a weighted average can be computed; the initial inputs of agents are $x_i = s_{0,i} = c_i \cdot r_i^j$, and their weights are $w_{0,i} = c_i$.

An interesting opportunity for future research is to examine systems where the agent graph is not fully connected. Since Push-Sum may also be used on graphs that are not fully connected, it would be interesting to examine how well this method functions in such graphs. This would enable the extension of the work to certain peer-to-peer networks.

Another interesting direction for future research is augmenting the framework with an option to efficiently choose among *service providers*:⁸ agent i requires a specific service, and there are m other agents that offer to service i 's request. Agents can be matched to service providers using a *matchmaking* service, but this problem has been dealt with [12]. We are concerned with the following question: once an agent is given a list M of m -service providers, which one should it choose? The obvious answer is, the one with the highest reputation: $\operatorname{argmax}_{j \in M} \bar{r}^j$. However, as the size of M may approach n , it is difficult to estimate the reputation of all agents in M .

One possible solution is to hold an election: the voters are the n agents, and the candidates are the m service providers. It is possible, for instance, to determine the winner using the simple *plurality* rule: each agent votes for one candidate; the winner is the candidate that secured the largest number of votes. It is also possible to resolve this election in our framework, using PUSH-SUM. Denote $M = \{j_1, j_2, \dots, j_m\}$. The input x_i of each agent for

⁸ In a sense similar to [11].

PUSH-SUM is (conceptually) a base $n + 1$ number with m coordinates; the l 'th coordinate of x_i is n if i votes for j_l , and 0 otherwise. The average is calculated by using PUSH-SUM with an *absolute* error $\epsilon = 1/3$, and some confidence $1 - \delta > 0$ (the x_i are translated to base 10). After the average $\frac{\sum_i x_i}{w_i}$ is calculated, the result is rounded to the nearest integer and again translated to base $n + 1$; the agent j_l such that the l 'th coordinate is largest wins the election.

Unfortunately, since in this case the inputs x_i are large, obtaining such a small absolute error requires a large number of iterations of PUSH-SUM, and furthermore, the message size is large. Is there a way of holding an election (using some other voting rule, perhaps) in our framework in a way that scales well with respect to both the running time and message size?

A different direction that could be taken is to use our framework to prevent attacks based on cheap pseudonyms [18]. This is made possible due to the fast aggregation of information and its global perspective. If an agent cheats, *all* other agents will soon know. It is possible to restrict newcomers, with an unestablished reputation, to only one transaction in every period of length $O(\log n)$. This way, each identity would be good for only a single deceitful transaction, since after the period is over, the information could have already been obtained by any agent. Granted, it would still be possible to shed stained identities, but the flow of transactions a cheater would be able to complete would be severely diminished.

Acknowledgments This work was partially supported by Israel Science Foundation grant #898/05. The work was done while Ariel Procaccia was at the Hebrew University of Jerusalem, and was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

References

1. Abdul-Rahman, A., & Hailes, S. (2000). Supporting trust in virtual communities. In *HICSS'00: Proceedings of the 33rd Hawaii International Conference on System Sciences* (Vol. 6, p. 6007). Washington, DC: IEEE Computer Society. ISBN: 0-7695-0493-0.
2. Aberer, K. (2001). P-grid: A self-organizing access structure for P2P information systems. In C. Batini, F. Giunchiglia, P. Giorgini, & M. Mecella (Eds.), *Proceedings of the 9th International Conference on Cooperative Information Systems (CoopIS 2001)*, Lecture Notes in Computer Science (Vol. 2172, pp. 179–194). Springer-Verlag, London, UK. <http://citeseer.ist.psu.edu/aberer01grid.html>.
3. Aberer, K., & Despotovic, Z. (2001). Managing trust in a peer-2-peer information system. In *Proceedings of the Tenth International Conference on Information and Knowledge Management* (pp. 310–317).
4. Akerlof, G. A. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84, 488–500.
5. Amar, L., Barak, A., Levy, E., & Okun, M. (2007). An on-line algorithm for fair-share node allocations in a cluster. In *CCGRID'07: Proceedings of the Seventh IEEE International Symposium on Cluster Computing and the Grid* (pp. 83–91). Washington, DC: IEEE Computer Society. <http://dx.doi.org/10.1109/CCGRID.2007.22>. ISBN: 0-7695-2833-3.
6. Awerbuch B., & Kleinberg, R. D. (2005). Competitive collaborative learning. In *COLT: Proceedings of the Workshop on Computational Learning Theory*, Morgan Kaufmann Publishers.
7. Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.
8. Barak, A., Guday, S., & Wheeler, R. G. (1993). *The MOSIX distributed operating system, load balancing for UNIX. Lecture notes in computer science (LNCS)* (Vol. 672). Berlin; New York: Springer.
9. Barak, A., & La'adan, O. (1998). The MOSIX multicomputer operating system for high performance cluster computing. *Future Generation Computer Systems*, 13(4–5), 361–372.
10. Cheng, A., & Friedman, E. (2005). Sybilproof reputation mechanisms. In *P2PECON'05: Proceeding of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-peer Systems* (pp. 128–132). New York, NY, USA: ACM.
11. Cornelli, F., Damiani, E., De Capitani di Vimercati, S., Paraboschi, S., & Samarati, P. (2002). Choosing reputable servants in a P2P network. In *Proceedings of the 11th International Conference on the World Wide Web* (pp. 376–386).
12. Cybenko, G., & Jiang, G. (1999). Matching conflicts: Functional validation of agents. In *Proceedings of the AAI Workshop on Agent Conflicts* (pp. 14–19).

13. Damiani, E., De Capitani di Vimercati, S., Paraboschi, S., Samarati, P., & Violante, F. (2002). A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *CCS'02: Proceedings of the 9th ACM Conference on Computer and Communications Security*, Washington, DC (pp. 207–216). New York: ACM. <http://doi.acm.org/10.1145/586110.586138>. ISBN: 1-58113-612-9.
14. Dellarocas, C. (2000). Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *EC'00: Proceedings of the 2nd ACM Conference on Electronic Commerce* (pp. 150–157). New York: ACM. <http://dx.doi.org/10.1145/352871.352889>. ISBN: 1581132727.
15. Dellarocas, C. (2001). Analyzing the economic efficiency of eBay-like online reputation mechanisms. In *ACM Conference on Electronic Commerce (EC-01)*, Tampa, Florida.
16. Demers, A., Greene, D., Hauser, C., Irish, W., Larson, J., Shenker, S., et al. (1987). Epidemic algorithms for replicated database maintenance. In *PODC'87: Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing*, Vancouver, British Columbia, Canada (pp. 1–12). New York: ACM. <http://doi.acm.org/10.1145/41840.41841>. ISBN: 0-89791-239-4.
17. Dewan, P. (2004). Peer-to-peer reputations. In *Proceedings of the 18th International Parallel and Distributed Processing Symposium*.
18. Friedman, E. J., & Resnick, P. (2001). The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2), 173–199.
19. Frieze, A. M., & Grimmett, G. R. (1985). The shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10, 57–77.
20. Fudenberg, D., & Levine, D. K. (1995). *Reputation and equilibrium selection in games with a patient player*. Levine's Working Paper Archive 103, UCLA Department of Economics, January. <http://ideas.repec.org/p/cla/levarc/103.html>.
21. Jelasty, M., Montesor, A., & Babaoglu, O. (2003). *Towards secure epidemics: Detection and removal of malicious peers in epidemic-style protocols*. Technical Report UBLCSS-2003-14, Department of Computer Science, University of Bologna.
22. Jelasty, M., Montesor, A., & Babaoglu, O. (2005). Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems*, 23(3), 219–252.
23. Kandori, M. (1992). Social norms and community enforcement. *The Review of Economic Studies*, pp. 63–80.
24. Kempe, D., Dobra, A., & Gehrke, J. (2003). Gossip-based computation of aggregate information. In *FOCS'03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science* (p. 482). Washington, DC: IEEE Computer Society. ISBN: 0-7695-2040-5.
25. Keyani, P., Larson, B., & Senthil, M. (2002). Peer pressure: Distributed recovery from attacks in peer-to-peer systems. In *Proceedings of the International Workshop on Peer-to-Peer Computing* (pp. 306–320).
26. Kreps, D. M., & Wilson, R. (1982, August). Reputation and imperfect information. *Journal of Economic Theory*, 305–15.
27. Leita, J., Pereira, J., & Rodrigues, L. (2007). Epidemic broadcast trees. In *Proceedings of the 26th IEEE International Symposium on Reliable Distributed Systems*, 10–12 October 2007 (pp. 301–310). Washington, DC: IEEE Computer Society.
28. Marsh, S. P. (1994). *Formalising trust as a computational concept*. PhD thesis, University of Stirling.
29. Pandurangan, G., Raghavan, P., & Upfal, E. (2001). Building low diameter P2P networks. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science* (pp. 492–499).
30. Resnick, P., & Zeckhauser, R. (2002). Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. *Advances in Applied Microeconomics*, 11.
31. Resnick, P., & Zeckhauser, R. J. (2003). *Eliciting honest feedback in electronic markets*. KSG Working Paper Series RWP02-039.
32. Srivatsa, M., Xiong, L., & Liu, L. (2005). TrustGuard: Countering vulnerabilities in reputation management for decentralized overlay networks. In *WWW'05: Proceedings of the 14th International Conference on World Wide Web*, Chiba, Japan (pp. 422-431). New York: ACM. <http://doi.acm.org/10.1145/1060745.1060808>. ISBN: 1-59593-046-9.
33. Xiong, L., & Liu, L. (2003). A reputation-based trust model for peer-to-peer ecommerce communities [Extended Abstract]. In *EC'03: Proceedings of the 4th ACM Conference on Electronic Commerce*, San Diego, CA (pp. 228–229). New York: ACM. <http://doi.acm.org/10.1145/779928.779972>. ISBN: 1-58113-679-X.