

# Analysis of a Reputation System for Mobile Ad-Hoc Networks with Liars

Jochen Mundinger  
Statistical Laboratory  
University of Cambridge  
Cambridge CB3 0WB, UK  
j.mundinger@statslab.cam.ac.uk

Jean-Yves Le Boudec  
EPFL-IC-LCA  
CH-1015 Lausanne, Switzerland  
jean-yves.leboudec@epfl.ch

## Abstract

*Using decentralized reputation systems is a promising approach to ensuring cooperation and fairness in Mobile Ad-Hoc Networks. However, they are vulnerable to liars and robustness has not been analyzed in detail. With our work, we provide a first step to the analysis of a reputation system based on a deviation test. Nodes accept second hand information only if this does not differ too much from their reputation values. Whereas our earlier paper [13] dealt with a simplified one-dimensional model, we now consider the original two-dimensional system. We show that the system exhibits a phase transition: In the subcritical regime, it is robust and lying has no effect. In the supercritical regime, lying does have an impact. We compute the critical values via a mean-field approach and use simulations to verify our results. Thus, we obtain conditions for the deviation test to make the reputation system robust and provide guidelines for a good choice of parameters.*

## 1. Introduction

Performance of Mobile Ad-Hoc Networks is well-known to suffer from free-riding as there is a natural incentive for nodes to only consume, but not contribute to the services of the system. We would also like to protect the network from malicious attacks. One of the ideas that have been considered to this end is that of a reputation system. Nodes keep track of their peers' behaviour and exchange this information with others in order to compute a reputation value about each peer. Nodes with a good reputation are then favoured.

By using second hand information, an accurate estimate of some subject's behaviour can be obtained faster. Moreover, a node can have a reputation value about a subject without ever having interacted with it himself. However, an inherent problem with any such mechanism is the vulnerability to liars.

A simple idea to protect the system from liars was suggested by Buchegger and Le Boudec [3], [5]. A node believes second hand information only if it does not differ too much from the node's reputation value. This is called the deviation test. In fact, the system considered in [3], [5] is more complex. It also allows for using second hand information from trusted peers. Therefore each node maintains both a reputation and a trust value about each of his peers.

The system appears to work well. Performance has only been evaluated through simulations of a network with a particular set of assumptions, though (e.g. on the routing protocol). Further simulations suggested that the deviation test on its own without the trust component performs nearly as well [4]. It seems surprising that such a simple idea works so well and we consider it worth analyzing in more detail and in a more general context. This is the aim of our research. In our earlier paper [13] we dealt with a simplified one-dimensional model. We now consider the original two-dimensional system.

In this paper, we analyze the case of 2 nodes, one honest and the other a liar. The precise modeling assumptions are listed in Section 2 and the model is formulated in Section 3. We provide mean-field results in Section 4 and verify them by means of simulation in Section 5. We thus show that the system exhibits a phase transition. That is, there is a threshold proportion of lying below which the reputation value of the honest node remains unaffected. Above it, the lying will have an impact and corrupt the reputation system. Thus, we provide guidelines for a good choice of parameters.

A number of reputation mechanisms have been suggested and studied. Michiardi and Molva propose the collaborative reputation mechanism CORE [11]. The CONFIDANT Protocol was introduced by Buchegger and Le Boudec [2]. Aberer and Despotovic [1] suggest a mechanism for P-Grid, a Peer-To-Peer system, that spreads negative information only. Collaboration enforcement in Peer-To-Peer systems has also been considered by Moreton and Twigg [12]. Carbone et al. [6] introduce a formal model

for trust in dynamic networks. Jurca and Faltings [8] and Fernandes et al. [7] consider incentives for truthful reporting itself. The reader is referred to [9] for the EigenTrust algorithm, a method to compute global trust values in the presence of pre-trusted peers. However, to the best of our knowledge, our work is the first analytical approach to evaluate a reputation system.

## 2. Assumptions

### 2.1. Subject behaviour

We consider the case when there is a single subject whose reputation is considered. This subject might be one of the  $N$  nodes themselves, but it might also be the provider of some external service such as Internet access. At each observation, its actual behaviour is assumed to be either positive or negative with probabilities  $\theta$  and  $1 - \theta$  respectively, independently of all other observations. The more practical case when there are  $M$  subjects of interest can be decomposed into  $M$  instances of our model, as the  $M$  different sets of reputation values do not interfere with each other.

### 2.2. Reputation

Each node  $i$  maintains a reputation value  $R^i = R_n^i$  about the subject that reflects its belief about  $\theta$  at time  $n$ . This opinion might change with new observations, arising either from interactions with the subject itself or with a peer.

A direct (first hand) observation is an observation of the subject's behaviour. Direct observations are always accepted and the reputation values updated accordingly. An indirect (second hand) observation arises from interactions with peers who report about their own direct observations. Indirect observations are only accepted if the reported observation does not deviate too far from the current reputation  $R_n^i$ . This deviation test is controlled by the parameter  $d$ . Moreover, if accepted, the impact of an indirect observation is scaled by a weighting parameter  $\omega$ .

We shall restrict attention to the case of 2 nodes, one honest with reputation value  $R_n$  about the subject and the other a liar. This is closely related to the general case, because several liars can be thought of as a single liar by aggregating their influence and we can focus on one out of the honest nodes by symmetry. Moreover, it looks like ignoring the other honest ones could be accounted for by increasing the proportion of (necessarily truthful) direct interactions, but this will have to be investigated in more detail (cf. Section 6).

### 2.3. Interaction model

We shall assume that each node  $i$  makes observations at certain points in time, giving rise to the discrete-time indices  $n = 1, 2, \dots$ . With fixed probabilities these are either direct observations or indirect ones. For the two node case, an observation by the honest node is assumed direct with probability  $p$  and indirect (i.e. from the liar) with probability  $\bar{p} = 1 - p$ . All these events are assumed to be independent.

### 2.4. Adversary model

One needs to make precise assumptions on the adversaries' abilities in order to give performance guarantees. We shall assume that liars follow the plain strategy to always lie maximally, i.e. they will always report either extremely negative or extremely positive behaviour about the subject when interacting with their peers in an attempt to achieve maximal impact. It suffices to focus on the extremely negative part, as the other one is similar by symmetry.

### 2.5. Performance

A reputation system works well if good nodes in the network benefit from it and bad nodes do not, or at least not as much. We claim that this can be achieved by suitable reaction mechanisms based on the reputation values, provided that these values are accurate. Notice that liars can easily change their own reputation values to anything they want. So the question is whether they can influence the reputation values of the honest nodes. The faster the nodes can obtain accurate estimates, the better the system will work, but there is a fundamental trade-off between robustness and speed. We shall assess robustness in detail. It will then be possible to choose parameters such that the system will be as fast as possible subject to satisfying a given accuracy condition.

## 3. Model

A natural scheme, suggested by the reputation system in [5] and other proposals, is to keep a history of previous events. Two counters,  $\alpha_n$  and  $\beta_n$ , are updated whenever there is a new observation, either direct or indirect.  $\alpha_n$  keeps track of positive observations,  $\beta_n$  keeps track of negative observations. Thus we are led to consider the following two-dimensional process  $(\alpha_n, \beta_n)$  for  $n \geq 0$ .

$$(\alpha_{n+1}, \beta_{n+1}) = u(\alpha_n, \beta_n) + \begin{cases} (1, 0) \\ (0, 1) \\ (0, \omega)1_{\{\frac{\alpha_n}{\alpha_n + \beta_n} \leq d\}} \end{cases} \quad (1)$$

with probabilities  $p\theta$ ,  $p(1-\theta)$  and  $\bar{p}$  respectively. The three possibilities correspond to a positive direct observation, a negative direct observation and an indirect observation respectively (cf. Sections 2.1 and 2.3).

Direct observations are always accepted and counted with 1. Indirect observations have to pass the deviation test with parameter  $0 < d < 1$ , modeled by the indicator function, and are weighted by  $\omega > 0$  as described in Section 2.2. Indirect observations are negative, because they are obtained from the liar who is assumed to report extremely negative behaviour (cf. Section 2.4).

Moreover, we discount both components individually with a discount factor  $0 < u < 1$ , typically very close to 1. This allows the system to gradually forget about old observations. We want discounting in the model to be able to track changing behaviour.

The initial conditions are  $(\alpha_0, \beta_0)$ . Note that if  $\omega = 1$  then starting with  $\alpha_0 + \beta_0 = 1/(1-u)$  will leave  $\alpha_n + \beta_n$  unchanged, starting elsewhere the sum will converge to  $1/(1-u)$ . Starting with such a converged value is in fact reasonable, because we would like to allow for tracking changing behaviour. However, there can be no a priori knowledge of a change, so we cannot simply reset the system to an arbitrary starting value.

The quantity we are interested in is  $R_n = \alpha_n/(\alpha_n + \beta_n)$ . We examine how well this compares to the true value of  $\theta$ . By the initial state  $R_0$  we shall mean  $\frac{1}{1-u}(R_0, 1 - R_0)$ . Notation is summarized in Table 1.

Note that, although we have defined the process in order to estimate  $\theta$ , it does not converge to a constant (in probability). For all  $n$ , there is positive probability of the next state taking either one of two values which differ by a constant. This is due to the discounting. So, we assess convergence (in distribution) to some limiting distribution from which we infer  $\theta$ .

symbol	meaning
$\theta$	prob. of positive subject behaviour (cf. 2.1)
$p$	prob. of an obs. being direct (cf. 2.3)
$\bar{p}$	prob. of an obs. being indirect (cf. 2.3)
$\alpha_n$	positive reputation component (cf. 3)
$\beta_n$	negative reputation component (cf. 3)
$R_n$	inferred reputation value (cf. 3)
$d$	deviation test parameter (cf. 2.2)
$\omega$	weighting factor for indirect obs. (cf. 2.2)
$u$	discount factor (cf. 3)

**Table 1. Notation**

## 4. Mean-field approach

In this section we will derive and solve the ‘mean’ ordinary differential equation (ODE) associated with our process. This is a standard approach in stochastic approximation theory. For a comprehensive reference see Kushner and Yin [10]. Usually, one would then show that noise effects average out asymptotically so that the actual behaviour of the process is determined by that of the mean ODE. Instead, in the next section, we will use simulations to verify the mean-field predictions from our ODE approach directly.

From (1), averaging the dynamics, we obtain the mean ODE as

$$\begin{aligned}\dot{\alpha}(t) &= (u-1)\alpha(t) + p\theta \\ \dot{\beta}(t) &= (u-1)\beta(t) + p(1-\theta) + \omega\bar{p}1_{\{\frac{\alpha}{\alpha+\beta} \leq d\}}\end{aligned}\quad (2)$$

This system is discontinuous, but linear above and below the line of discontinuity  $\alpha/(\alpha + \beta) = d$ . Solving (2), we find that the system has either one or two solutions, depending on the parameters of the model.

$$(\alpha, \beta) = \frac{p}{1-u}(\theta, 1-\theta) \quad (3)$$

is a solution on  $\alpha/(\alpha + \beta) > d$  and


$$(\alpha, \beta) = \left( \frac{p\theta}{1-u}, \frac{p(1-\theta)}{1-u} + \frac{\omega\bar{p}}{1-u} \right) \quad (4)$$

on  $\alpha/(\alpha + \beta) \leq d$ . (3) is a solution if  $\theta > d$  and (4) is a solution if  $\theta \leq d$  or  $\bar{p} \geq (\theta - d)/(\theta - d + \omega d)$ . Both solutions are globally asymptotically stable on their respective region. Moreover, if only one exists, trajectories in the other region converge to it also. In this sense it is globally asymptotically stable for the whole system. Otherwise, both are stable on their respective region only and, in this sense, both are locally stable overall.

**Theorem 1** *If  $\theta > d$ ,  $\frac{p}{1-u}(\theta, 1-\theta)$  is a solution of the mean ODE (2). For  $\bar{p} < \bar{p}_c = (\theta - d)/(\theta - d + \omega d)$  it is globally asymptotically stable. Otherwise, there exists a second, false solution  $\frac{1}{1-u}(p\theta, p(1-\theta) + \omega\bar{p})$  and both are locally stable. If  $\theta \leq d$  then the latter, false one is globally asymptotically stable.*

Thus, the reputation system exhibits a phase transition behaviour. Assuming  $\theta > d$ , we find a bifurcation in terms of the parameter  $\bar{p}$ . In the subcritical regime, that is, for  $\bar{p} < \bar{p}_c$ , the solution corresponding to the true  $\theta$  is unique. In the supercritical regime where  $\bar{p} \geq \bar{p}_c$  there is a second, false solution.

Alternatively, this can be phrased in terms of the system parameter  $d$ . For small  $d$  there is only one solution, the true



**Figure 1. Bifurcation plot in terms of  $d$ : As  $d$  increases from 0 to 1 the number of solutions changes from 1 (the true one) to 2 and back to 1 (the false one). Here,  $\pi$  is short hand for  $(p\theta)/(p + \omega\bar{p})$ .**

one. For intermediate  $d$  there are two and for large  $d$  there is only the false solution. This is illustrated in Figure 1.

In practical terms, this suggests that the reputation system works and that the liar cannot achieve anything if  $\theta > d$  and  $\bar{p} < \bar{p}_c$ . However, the liar does have an impact otherwise. As for the latter condition, it is intuitively clear that the deviation test can filter out extreme lies only if they do not occur too often. As for the first condition, it is clear that if the true  $\theta$  is too close to the extreme 0 behaviour, the deviation test will not filter the lies and the liar will have an impact. The deviation test cannot protect a ‘very bad’ subject behaviour to be pushed by the liar to an ‘extremely bad’ perception by the honest node. However, there is a range of parameters for which it does protect the reputation system.

As mentioned in Section 2, we can repeat the analysis to show that the reputation system similarly protects against extremely positive reports. Combining the two, we obtain the following necessary and sufficient conditions for the true solution to be unique when both, positive and negative lying are permitted:  $\min\{\theta, 1 - \theta\} > d$  and  $\bar{p} < (\min\{\theta, 1 - \theta\} - d)/(\min\{\theta, 1 - \theta\} - d + \omega d)$ .

## 5. Simulations

In this section we report our simulation results. We used formulation (1) to compute  $10^5$  steps and then plot  $R_n = \alpha_n/(\alpha_n + \beta_n)$  against  $n$ . 100 independent runs were carried out for each parameter set.

In Figure 2 we show a typical sample path for  $\theta = 0.8$  with  $d = 0.4$ ,  $u = 0.99$ ,  $\omega = 1$  and  $R_0 = 0$  when  $\bar{p} = 0.2$ , i.e.  $p = 0.8$ . The lower and upper boundaries in the plot correspond to  $R_n = 0$  and 1 respectively. The lower and upper intermediate lines correspond to  $(p\theta)/(p + \omega\bar{p})$  and  $\theta$  respectively. We note that  $R_n$  approaches  $\theta$  and then remains within its neighbourhood until the end of the simula-

tion. All 100 independent runs showed the same qualitative behaviour.

Note that for the parameter set above the predicted critical value is  $\bar{p}_c = 0.5$  and the subcritical behaviour is as expected. We obtained the same qualitative behaviour when  $\bar{p} = 0.2$  is replaced by  $\bar{p} = 0.4$  and  $\bar{p} = 0.45 < \bar{p}_c$ . Here, starting with  $R_0 = 0$  can be viewed as a worst case. For other starting values, too, including the other extreme  $R_0 = 1$ , we obtained the same qualitative behaviour.

In Figure 3 we illustrate the effect of the discount parameter  $u$ . A typical sample path is shown for the same parameters as in Figure 3 except  $u = 0.999$ . The variability around  $\theta$  is smaller and it takes longer for the process to approach  $\theta$ .

In Figure 4 we consider the supercritical case with now  $\bar{p} = 0.8$ . Also, as opposed to the earlier parameter set (Figure 2) we take  $u = 0.95$  for clearer illustration. As a side effect, variability is increased. We note that the process settles down for some time in the neighbourhoods of  $(p\theta)/(p + \omega\bar{p}) = p\theta = 0.16$  and  $\theta = 0.8$  in an alternating fashion. This is in agreement with the mean-field prediction that both  $(p\theta)/(p + \omega\bar{p})$  and  $\theta$  are solutions. It is due to the discounting that we do not have convergence to a constant, but there is always positive probability of moving from one solution to the other. This can be viewed as another advantage of discounting, because the process cannot get stuck at the false solution forever.

The same qualitative behaviour is observed in all 100 independent runs and also for  $\bar{p} = 0.6$  and  $\bar{p} = 0.55 > \bar{p}_c$ , only the proportion of time spent near  $\theta$  is higher. Note also that the second solution depends on  $\bar{p}$  and  $\omega$ : it is at 0.32 and 0.36 respectively. In Figure 5 we demonstrate that the long-term behaviour does not depend on the starting value  $R_0$ . Whereas in the mean-field approach each (locally stable) solution has a basin of attraction, we know that there is always positive probability for the process to move from one to the other, so the independence of the starting value is as expected.

With a different choice of parameters the prediction of only one solution at  $(p\theta)/(p + \omega\bar{p})$  for the case  $\theta \leq d$  can also be verified. Finally, we carried out simulations with the same parameters as in Figure 2 except  $\omega = 2$  to verify, in particular, the critical value  $\bar{p}_c = 1/3$ .

Thus, we have verified the results of the mean-field approach. In addition, Figures 2 – 5 give us an idea of the proportion of time near the false  $(p\theta)/(p + \omega\bar{p})$ . This time increases with  $\bar{p}$ .




Figure 2. Typical sample path for  $\theta = 0.8$  with  $d = 0.4$ ,  $u = 0.99$ ,  $\omega = 1$  and  $R_0 = 0$  when  $\bar{p} = 0.2$ . We plot  $R_n$  against  $n$ , the upper line corresponding to  $\theta$  and the lower line to  $(p\theta)/(p + \omega\bar{p})$ .  $R_n$  approaches and then remains close to  $\theta$ .




Figure 3. Typical sample path for the same parameters as in Figure 2 except  $u = 0.999$ . The variability is smaller and it takes longer for the process to approach  $\theta$ .




Figure 4. Typical sample path for the same parameters as in Figure 2 except  $\bar{p} = 0.8$  and also  $u = 0.95$ . The process settles down for some time in a neighbourhood of  $(p\theta)/(p + \omega\bar{p})$  and  $\theta$  in an alternating fashion.




Figure 5. Typical sample path for the same parameters as in Figure 4 except  $R_0 = 1$ . Qualitatively, the long-term behaviour remains the same.

## 6. Conclusions and further work

The reputation system will be most robust against lying if  $d$  is chosen small. We have quantified the effect on the robustness due to a change in  $d$ . This is important for the fundamental trade-off, because smaller  $d$  means less use of second hand information. For a system that is as fast as possible subject to being accurate on a given range we would choose the largest  $d$  that still satisfies the required accuracy. In practical terms, we have seen that there is a reasonable range of parameters for which the deviation test will protect the reputation systems from liars. Given any cost function with arbitrary weights on accuracy and speed, we can compute the optimal choice of the system parameter  $d$ . One might also want to think about individually controlled  $d^i$  based on the nodes' current information.

The scenario of two peers that we have considered thus far can also be viewed as an extreme case. Even if all other nodes are malicious so that all second hand information is manipulated, the reputation system protects against the lying if the aggregate proportion of lying is below a threshold. In a real-world scenario one would typically be able to assume that at least some if not most nodes are honest. To examine this in more detail, the next step is to consider the case of three peers: one honest node making direct observations with probability  $p$ , indirect ones originating from the liar with probability  $q$  and indirect ones originating from the other honest peer with probability  $1 - p - q$ .

The next extension is then to consider strategic lying, that is adversaries attempting something more subtle than simply lie maximally. For example, they could lie in some proportion of reports only or they could always report intermediate behaviour in an attempt to conceal their lies. It would also be interesting to consider random noise instead of fake reports.

**Acknowledgment** The authors would like to thank Sonja Buchegger for valuable discussions.

## References

- [1] K. Aberer and Z. Despotovic. Managing trust in a peer-to-peer information system. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKC 2001)*, 2001.
- [2] S. Buchegger and J.-Y. Le Boudec. Performance analysis of the confidant protocol: Cooperation of nodes — fairness in dynamic ad-hoc networks. In *Proceedings of IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*, Lausanne, CH, June 2002.
- [3] S. Buchegger and J.-Y. Le Boudec. The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In *Proceedings of WiOpt '03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Sophia-Antipolis, France, March 2003.
- [4] S. Buchegger and J.-Y. Le Boudec. Personal communication, February 2004.
- [5] S. Buchegger and J.-Y. Le Boudec. A robust reputation system for peer-to-peer and mobile ad-hoc networks. In *Proceedings of P2PEcon 2004*, Harvard University, Cambridge MA, USA, June 2004.
- [6] M. Carbone, M. Nielsen, and V. Sassone. A formal model for trust in dynamic networks. Technical report, BRICS Report RS-03-4, 2003.
- [7] A. Fernandes, E. Kotsovinos, S. String, and B. Dragovic. Incentives for honest participation in distributed trust management. In *Proceedings of iTrust 2004*, Oxford, UK, March 2004.
- [8] R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, June 2003.
- [9] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p network. In *Proceedings of the Twelfth International World Wide Web Conference 2003*, May 2003.
- [10] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, second edition, 2003.
- [11] P. Michiardi and R. Molva. CORE: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. Sixth IFIP conference on security communications, and multimedia (CMS 2002), Portoroz, Slovenia., 2002.
- [12] T. Moreton and A. Twigg. Enforcing collaboration in peer-to-peer routing services. In *Proceedings of the First International Conference on Trust Management*, Heraklion, Crete, May 2003.
- [13] J. Munding and J.-Y. Le Boudec. Analysis of a robust reputation system for self-organized networks. *To appear in European Transactions on Telecommunications*, 2005.