

Analysis of a Robust Reputation System
for Self-Organized Networks

Analysis of Reputation System

Jochen Mundinger	Jean-Yves Le Boudec
Statistical Laboratory	EPFL-IC-LCA
Centre for Mathematical Sciences	CH-1015 Lausanne
University of Cambridge	Switzerland
Wilberforce Road	
Cambridge CB3 0WB, UK	

Contact details for proof checking:

J.Mundinger@statslab.cam.ac.uk

Tel: +44/(0)1223/337952

Fax: +44/(0)1223/337956

January 15, 2004

Abstract

Self-organized networks require some mechanism to ensure cooperation and fairness. A promising approach is the use of decentralized reputation systems. However, their vulnerability to liars has not yet been analyzed in detail.

In this paper, we provide a first step to the robustness analysis of a reputation system based on a deviation test. Users accept second hand information only if this does not differ too much from their reputation values. We simplify the original system in order to obtain a one-dimensional formulation and show that it exhibits a phase transition. In the subcritical regime, the reputation system is robust. In the supercritical regime, lying has an impact. We obtain the critical values via a mean-field approach and verify the results by explicit computation. Thus, we provide conditions for the deviation test to make the reputation system robust as well as quantitative results on what goes wrong in the supercritical regime.

1 Introduction

Decentralized systems such as Peer-to-Peer Resource Sharing Networks have recently become more popular, both in practice and research. Novel communication systems that are being considered, such as Mobile Ad-Hoc Networks, are designed to be self-organized so as to achieve minimal administrative and operational costs.

In most applications users are individuals that are primarily interested in their own benefit. However, for a decentralized network to function, its users need to contribute in some form or other to the services of the system without getting any immediate reward. Thus, there is a natural incentive for users to only consume, but not contribute. Cooperation and fairness cannot be guaranteed. This behaviour is called free-riding and is a well-known phenomenon in economics. In the context of Peer-To-Peer Networks, for example, it has been demonstrated in a number of measurement studies [1, 2, 3]. Network performance can be seriously deteriorated.

One possible approach to the free-rider problem is to introduce pricing schemes [4] into the system in order to create the right incentives for the users. For example, see [5] and [6]. Alternatively, contribution rules [7] and, more recently, artificial immune systems are being considered [8].

Another idea is that of a reputation system. Here, users keep track of their peers' behaviour and exchange this information with others. Each user merges his own first hand information with the second hand information he receives in order to compute a reputation value about each of his peers. Users with a good reputation are then favoured. A popular example of a reputation mechanism is the rating used in EBAY [9]. However, this is a centralized mechanism as the ratings are handled by the EBAY server. By contrast, the applications we have in mind are fully decentralized and self-organized. No trusted third party can be assumed. Reputation systems do not distinguish between different causes for misbehaviour. By contrast, pricing or rules mainly address selfish users, but not malicious or faulty behaviour.

The advantage of a reputation system over merely using first hand information is two-fold. Firstly, an accurate estimate of some subject's behaviour can be obtained faster. Secondly, a user 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. Some user might have an interest in spreading false information, so naively believing all second hand information is problematic. Reputation values must be accurate at least to some degree.

A simple idea to address this problem was suggested originally in the context of Mobile Ad-Hoc Networks by Buchegger and Le Boudec [10]. Here, a user believes second hand information only if it does not differ too much from the user's reputation value. This is called the deviation test. In fact, the system considered in [10] is more complex. It also allows for using second hand information from trusted peers. To this end, each user maintains both a reputation and a trust value about each of his peers. Both are updated using a modified Bayesian approach. As opposed to reputation, trust values are based on compatibility and thus indicate agreement.

The system appears to work well. So far, however, performance has only been evaluated through simulations of a Mobile Ad-Hoc Network with a particular set of assumptions (e.g. on the routing protocol). In personal communication the authors reported further simulations suggesting that the deviation test on its own without the trust component nearly performs as well. 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 this paper, we provide the first step by analyzing a simplified model for the case of 2 users, one honest and the other a liar. The model simplifies the reputation counters in order to obtain a one-dimensional system which is easier to analyze. The precise modeling assumptions are listed in Section 3.1 and the model is formulated in Section 3.2. We provide mean-field results in Section 4 and verify them by means of direct computation for certain parameter sets in Section 5. We thus show that the system exhibits a phase transition. That is, there is a threshold rate of lying below which the reputation value of the honest user remains unaffected. Above it, the lying will have an impact and corrupt the reputation system.

Note that the idea of a deviation test is a very natural one. In a social network of acquaintances humans are likely to disbelieve information that, to them, seem

highly unlikely. At least, if they have no means of verifying it themselves. As such, the results are of interest in the context of social sciences also.

2 Related Work

A number of reputation mechanisms have been suggested and studied. However, to the best of our knowledge, our work is the first analytical approach to evaluate a reputation system.

A selection of references focusing on decentralized reputation systems is given below. Further references can be found therein.

Michiardi and Molva propose the collaborative reputation mechanism CORE [11]. The CONFIDANT Protocol was introduced by Buchegger and Le Boudec [12]. Aberer and Despotovic [13] 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 [14]. Carbone et al. [15] introduce a formal model for trust in dynamic networks. Jurca and Faltings [16] and Fernandes et al. [17] consider incentives for truthful reporting itself. The reader is referred to [18] for the EigenTrust algorithm, a method to compute global trust values in the presence of pre-trusted peers.

3 Model

3.1 Modeling Assumptions

Subject Behaviour

We study the case when there is a single subject whose reputation is considered. Its actual behaviour is assumed to be either positive or negative with probabilities θ and $1 - \theta$ respectively. Thus, when a user interacts with the subject itself it observes positive behaviour with probability θ and negative behaviour otherwise. This is assumed to be independent of all other observations. Hence the actual behaviour is represented by the parameter θ , a real number in $[0, 1]$.

Note that this subject is not necessarily one of the N users themselves. Alternatively, users of the network might be interested in the behaviour of some external subject. In the context of a Mobile Ad-Hoc Network, for example, this might be the provider of an external service such as Internet access. Our model captures this case as well. The more practical case when there are M subjects of interest can be decomposed into M instances of our model. The M sets of reputation values do not interfere with each other and can be considered independently. In particular, if we take N subjects, one for each user, the model corresponds to the scenario described in the introduction.

Reputation

There are N users $1, 2, \dots, N$, with corresponding reputation values $R^i(t)$ about the subject. These are also real numbers in $[0, 1]$ and reflect the belief that user i has about θ at time t . This opinion might change with new observations, arising either from interactions with the subject itself or with a peer.

A direct observation is an observation of the subject's behaviour. The collection of direct observations constitutes a user's first hand information. An indirect observation arises from interactions with peers who report about their own direct observations. The collection of indirect observations is the second hand information available to the user.

Direct observations are always accepted and the reputation values updated accordingly. Indirect observations are only accepted if the reported observation does not deviate too far from the current opinion $R^i(t)$. This deviation test is controlled by the parameter Δ in $(0, 1)$.

Interaction Model

The interaction model describes how users interact with the subject and their peers. We shall assume that each user i makes direct observations at the points of a Poisson process in time, at rate μ^i . Interactions of user i with j (such that i receives second hand information from j) occur according to a Poisson process with rate λ^{ij} . All

processes are assumed to be independent.

In all applications the interaction pattern is influenced by the call model or some model of the users activity. In mobile applications, it is further influenced by the mobility model. Although the interaction pattern might differ between applications, the model above is a natural one to examine.

Adversary Model

One needs to make precise assumptions on the adversary's 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. They do so in attempt to achieve maximal impact. It suffices to focus on the extremely negative part, as the other one is similar by symmetry. The adversary model can be gradually extended to capture more sophisticated attacks on the reputation system.

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. The faster users 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 being accurate.

Further Assumptions

We first consider the case of one honest user and one liar. The honest user makes direct observations at rate $\mu = \mu^1 = 1$ and indirect ones originating from the liar at rate $\lambda = \lambda^{12} > 0$. Define $p = \mu/(\mu + \lambda) = 1/(1 + \lambda)$. Then, at each step, the observation is direct with probability p and indirect with probability $1 - p$. Notation is summarized in Table 1.

Note that there is a close relationship between the case of two peers only and the general case. We can focus on one out of the honest users by symmetry. Several liars can be considered as one by aggregating their influence. This can be accounted for by increasing λ for the one liar. In fact, it looks like ignoring the other honest ones could be accounted for by increasing μ , but this will have to be investigated in more detail (cf. Section 6).

3.2 Model Formulation

A natural scheme, motivated by the reputation system suggested in [10] and other proposals, is to keep a history of prior events, that is a count of positive and negative observations. Thus we are led to consider the following two-dimensional process $z_n = (x_n, y_n)$ for $n \geq 0$.

$$(x_{n+1}, y_{n+1}) = \rho(x_n, y_n) + \begin{cases} (1, 0) & \text{w.p. } p\theta \\ (0, 1) & \text{w.p. } p(1 - \theta) \\ (0, \omega)1_{\{x_n/(x_n+y_n) \leq \Delta\}} & \text{w.p. } 1 - p \end{cases} \quad (1)$$

Essentially, the first component keeps track of positive observations and the second component keeps track of negative observations. Direct observations are counted with 1, indirect observations are weighted by $\omega > 0$. Moreover, we discount both components individually with a discount factor $0 < \rho < 1$, typically very close to 1. We want discounting in the model to be able to track changing behaviour. The initial conditions are $z_0 = (x_0, y_0)$.

The quantity we are interested in is $R_n = x_n/(x_n + y_n)$, the proportion of positive observations of the total number of observations that the honest peer collects during the first n events. We examine how well this compares to the true θ , that is the actual proportion of positive behaviour of the subject in question.

Rewriting the two-dimensional formulation above in terms of R_n gives an expression that depends on the unknown $x_n + y_n$. However, assuming $\omega = 1$ and replacing the neutral increment $(0, 0)$ with the also neutral $(x_n/(x_n + y_n), y_n/(x_n + y_n))$ in the case of a rejection from the deviation test, the sum $x_n + y_n$ increases deterministically

by 1 at each step and $x_n + y_n$ can be determined from the starting value (x_0, y_0) .

$$R_{n+1} = R_n + \frac{1}{\rho(x_n + y_n) + 1} \begin{cases} (1 - R_n) & \text{w.p. } p\theta \\ -R_n & \text{w.p. } p(1 - \theta) \\ -R_n 1_{\{R_n \leq \Delta\}} & \text{w.p. } 1 - p \end{cases} \quad (2)$$

Taking (x_0, y_0) with $x_0 + y_0 = 1/(1 - \rho)$ we know that $x_n + y_n = 1/(1 - \rho)$ for all n . Hence we have the following, simpler formulation for $R_n = R^1(t_n)$ where t_n is the time of the n th interaction.

$$R_{n+1} = R_n + (1 - \rho) \begin{cases} (1 - R_n) & \text{w.p. } p\theta \\ -R_n & \text{w.p. } p(1 - \theta) \\ -R_n 1_{\{R_n \leq \Delta\}} & \text{w.p. } 1 - p \end{cases} \quad (3)$$

Starting with such a value is in fact reasonable. We would like to account for the case when behaviour might change over time. But then, if such a change occurs after the system has been running for some time, we would start from a state (x, y) that nearly satisfies $x + y = 1/(1 - \rho)$. However, there would be no a priori knowledge of the change, so we could not simply reset the system to an arbitrary starting value. So, we take the state at time 0 to be a ‘fully converged’ state with $x_0 + y_0 = 1/(1 - \rho)$. We can thus describe it by r_0 , meaning $\frac{1}{1-\rho}(r_0, 1 - r_0)$. Again, see Table 1 for a summary of the notation.

The process is a homogeneous Markov Chain with the state space a subset of the interval $[0, 1]$. For rational parameters, it will take rational values only and the state space is countable, although, complicated.

Note that we have lost a degree of freedom by assuming $\omega = 1$. It is reasonable to expect, however, that the important quantity is the product $\omega\lambda$, so in this respect we do not lose generality. Moreover, one can consider a generalization of (3) which corresponds to a certain projection onto the line $\{(x, y) : x + y = 1/(1 - \rho)\}$ in (1). The advantage of formulation (3) is that it is one-dimensional, so computing the distribution of R_n is tractable (cf. Section 5).

Although we have defined the process in order to estimate θ it does not converge to a constant. Neither with probability 1, nor in L^p , nor in probability. This is

because, for all times n , there is positive probability that the next state takes either one of two values which differ by a constant. However, for convergence, we would need this difference to become arbitrarily small. This lack of convergence is due to the discounting which we require to allow for tracking of behaviour that changes over time. Another advantage will become apparent later on (cf. Section 5). So, we assess convergence (in distribution) to some limiting distribution from which we infer θ .

4 Mean-Field Approach

4.1 Zero Drift Values

We shall first determine the values of R satisfying $R = \mathbb{E}R'$. These ‘zero drift values’ are solutions to

$$R = R + (1 - \rho) \begin{cases} p\theta - pR & \text{if } R > \Delta \\ p\theta - pR - (1 - p)R & \text{if } R \leq \Delta. \end{cases} \quad (4)$$

Thus, $R = \theta$ is a solution if and only if $\theta > \Delta$. $R = p\theta$ is a solution if and only if $p\theta \leq \Delta$.

Phrased in terms of λ we obtain the following: If $\theta > \Delta$, there is the truthful zero drift value $R = \theta$. For $\lambda < \lambda_c = (\theta - \Delta)/\Delta$ it is unique. Otherwise, there exists a second, false value $R = p\theta$. If $\theta \leq \Delta$ then the latter, false zero drift value is unique.

Furthermore, it is easy to check that the drift at other values is towards these zero drift values. The further away, the stronger the drift. In the case of two zero drift values, the change of drift occurs at $R = \Delta$. A graphical interpretation of the results is thus that θ and $p\theta$ are zero drift values only if they are above and below of Δ respectively.

Comparison with the invariant distribution of a Birth-and-Death chain (cf. [19]) with suitably chosen transition probabilities suggests that a single zero drift value maximizes the distribution, i.e. is the most likely state. Thus, what we can expect is

convergence of R_n (in distribution) to a limiting distribution which exhibits a phase transition in terms of the number of modes.

4.2 Mean Ordinary Differential Equation

Our formulation (3) can be written in stochastic approximation form. For a comprehensive reference see Kushner and Yin [20]. The basic paradigm is a stochastic difference equation where one recursively adjusts the parameter so that some goal is met asymptotically. The main concept used is to show that noise effects average out so that the actual behaviour is determined by that of a ‘mean’ ordinary differential equation (ODE). This has been applied in diverse areas, in particular in signal processing and communications.

In the stochastic approximation framework, our discounting corresponds to a constant step size parameter. This class of algorithms has been developed to allow for tracking changing parameters. The type of results is that the process spends nearly all of its time in a neighbourhood of the limit point or set. The size of the neighbourhood depends on the constant step size. Thus, results are of weak convergence type (i.e. convergence in distribution).

We will now study the deterministic mean ODE. Instead of showing that the process R_n is governed by it, we will then compute the distribution of R_n explicitly (cf. Section 5) and thus show that the process behaves as predicted from the ODE.

From (3), averaging the dynamics in the stochastic approximation form, we obtain the following ODE.

$$\dot{R}(t) = (1 - \rho)[p\theta - pR(t) - (1 - p)R(t)1_{\{R(t) \leq \Delta\}}] \quad (5)$$

Note that the ODE is discontinuous in $R(t)$. We can solve this separately for $R(t) \leq \Delta$ and $R(t) > \Delta$ to obtain the solution

$$R(t) = \begin{cases} (r_0 - p\theta)e^{-(1-\rho)t} + p\theta & \text{if } R(t) \leq \Delta \\ (r_0 - \theta)e^{-p(1-\rho)t} + \theta & \text{if } R(t) > \Delta . \end{cases} \quad (6)$$

Thus there are two possible solutions: θ and $p\theta$, the zero drift values from the previous section which have also been obtained by means of averaging. In addition, we can now assess stability of the deterministic system.

The condition for θ to be a solution is $\theta > \Delta$. The condition for $p\theta$ to be a solution is $p\theta \leq \Delta$. Both are globally asymptotically stable on their respective region. If only one exists then the trajectories in the other region also converge to it. So all trajectories converge to the solution and in this sense it is globally asymptotically stable for the whole system. Note that usually one would take global stability to imply smoothness. This is not the case here. Otherwise, both are globally asymptotically stable on their respective region only and in this sense, they are both locally stable overall.

Theorem 1 *If $\theta > \Delta$, then $R = \theta$ is a solution of the mean ODE (5). For $\lambda < \lambda_c = (\theta - \Delta)/\Delta$ it is globally asymptotically stable (in the sense explained above). Otherwise, there exists a second, false solution $R = p\theta$ and both are locally stable. If $\theta \leq \Delta$ then the latter, false one is globally asymptotically stable.*

Assuming $\Delta < \theta$, the reputation system exhibits a phase transition behaviour. In the subcritical regime, that is, for lying rates below the non-zero critical value λ_c , the true reputation value θ is the unique solution. In the supercritical regime where the lying rate is above the critical rate there is a second, false value.

In practical terms, this suggests that the reputation system works and that the liar cannot achieve anything if $\Delta < \theta$ and $\lambda < \lambda_c$. 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 former condition, it is clear that if the true θ is too close to the extreme 0 behaviour, the deviation test will not filter out the lies and the liar will have an impact. In conclusion, the deviation test cannot protect a ‘very bad’ subject behaviour to be pushed by the liar to an ‘extremely bad’ perception by the honest user. However, there is a range of parameters for which the deviation test does protect the reputation system.

As mentioned in Section 3.1, the analysis can be repeated symmetrically to show that the reputation system protects against extremely positive reports rather than extremely negative ones. Combining the two, we obtain the following conditions for the true solution to be unique if both, positive and negative lying is permitted: $\min\{\theta, 1 - \theta\} > \Delta$ and $\lambda < (\min\{\theta, 1 - \theta\} - \Delta) / \Delta$.

5 Distribution of R_n

Given R_0 we can compute the distribution of R_n analytically. We will use this now to confirm the mean-field predictions from the previous section.

The state space of our process is complicated and large, making the theoretical study hard and direct computation infeasible. So, instead we compute the exact distribution of R_n for a finite state space. For a real-world reputation system we are interested in considering only finitely many states anyway, as it is not possible to store arbitrary precision values on a finite device. Alternatively, the process can be simulated, but the explicit computation of the distribution is more powerful.

The results below are obtained when the unit interval is split into $G = 1000$ boxes of size $g = 0.001$ each. We first compute the transition matrix for the 1000 state chain. Then, with starting value r_0 , we repeatedly compute the distribution at the next step until this remains unchanged (in double precision, i.e. 64 bits, i.e. to sixteen significant digits). This suggests that we have converged.

In fact, taking the midpoint of intervals as the corresponding state gives rise to very non-smooth distributions due to approximation errors which do not seem to go away with increasing grid size G . So, instead, we choose the corresponding state uniformly from the interval. This amounts to the following: Let x denote the left endpoint of the interval, then the right one is $x + g$. The possible new states from x are $\rho x + (1 - \rho)$, ρx and x . The ones from $x + g$ are $\rho x + (1 - \rho) + \rho g$, $\rho x + \rho g$ and $x + g$ which differ from the previous one by at most g . All other states that can be reached from within the interval lie in between due to monotonicity. Thus the new states lie in at most two neighbouring intervals. We split the probability flow into these two intervals according to the proportions corresponding to a uniform

distribution over the interval rather than a point mass in the middle. The effect of this is demonstrated in Figure 1, where we plot a typical distribution first with a point mass and then with a uniform distribution.

The graphs in Figure 2 show the distributions obtained for $\theta = 0.8$, $\Delta = 0.4$, $\rho = 0.99$, $r_0 = 0.4$ ($= \Delta$), a typical set of parameters, and various values of p . Thus, from the previous section, the predicted critical value is $\lambda_c = (\theta - \Delta)/\Delta = 1$, that is $p_c = \Delta/\theta = 0.5$. Since some of the values are smaller by several orders of magnitude than others, the features are obscured. So we also plot them in log-scale in Figure 3.

From the log-scale plots in Figure 3 we note that the distribution is unimodal for $p > p_c = 0.5$ with a mode at $\theta = 0.8$. It is bimodal for $p < 0.5$ with a second mode at a lower value $p\theta$, i.e. at 0.16 and 0.32 respectively. This is all as predicted from the previous section. In fact, consulting the output for $p = 0.45$ and 0.55, we find that the predicted critical value of $p_c = 0.5$ is confirmed further. Moreover, with a different choice of parameters the prediction of only one mode at $p\theta$ for the case $\theta \leq \Delta$ can also be confirmed.

In addition, the graphs in Figure 2 give us good idea of the total mass near the false $p\theta$ compared to the total mass near the true θ . The latter decreases in the lying rate. Still, even for $p = 0.2$, the process is more likely to be right than wrong asymptotically. Only for very small p it becomes significant.

The lightly coloured distributions in Figure 2 are obtained for a different choice of $\rho = 0.9$. The discount factor controls the variability around the solutions. The further it is from 1, the less the probability mass is concentrated near these values. In the supercritical case, if there was no discounting, we would converge to one value or the other with certain probabilities and then remain there forever. However, with the discounting, this does not happen. There is always a small but positive probability of moving from one solution to the other. The proportion of the time the process spends around the true solution corresponds to the probability of converging to the true solution in the case of no discounting.

So far, with the parameter set above, we have only considered the case $r_0 = \Delta$.

However, the corresponding distributions obtained for the two extreme cases $r_0 = 0$ and $r_0 = 1$ are essentially the same as the ones for $r_0 = \Delta = 0.4$. This is shown in Figure 4. For $p = 0.4, 0.6$ and 0.8 they agree at least to within 10^{-14} for each state. For $p = 0.2$ they differ, however, distributions had not converged fully here before the computations were stopped after 5×10^6 iterations. We expect them to agree when computations are allowed to run until completion.

This suggests that the process is independent of its initial state, which is as expected for a unique attractor. Moreover, if there are two attractors, there is positive probability of moving from one to the other for all times and we start in a fully converged state. Thus, in this case, too, the initial state should not matter.

6 Conclusions and Further Work

We have seen that the reputation system exhibits a phase transition. Via a mean-field approach, we have derived a closed form expression for the critical lying rate λ_c . We have verified the mean-field results by direct computation.

In the subcritical regime, when the lying rate is sufficiently small, the liar has no impact on the honest user. In the supercritical regime, the liar does have an impact. Thus we can give precise conditions under which the deviation test makes the reputation system robust. We can further predict the false reputation value and with what probability this will be obtained rather than the true value in the supercritical regime.

The reputation system will be most robust against lying if Δ is chosen very small. We have quantified the effect on the robustness due to a change in Δ . This is important for the fundamental trade-off, because smaller Δ means less use of second hand information. We have been interested in a system that is as fast a possible subject to being accurate. 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 a cost function with arbitrary weights on accuracy and speed, we could com-

pute the optimal choice of the system parameter Δ . One might also want to think about individually controlled Δ^i , $i = 1, 2, \dots, N$, based on the users' current information.

We have illustrated the effect of the discount parameter. The closer it is to 1, the more accurate the process can estimate the parameter. However, it takes longer to track changing behaviour.

The scenario of two peers that we have considered thus far can also be viewed as an extreme case. Even if all other users are malicious so that all second hand information is manipulated, the reputation systems protects against the lying if the aggregate lying rate is below a threshold. In a real-world scenario one would typically be able to assume that at least some if not most users are honest. To examine this in more detail, the next step is to consider the case of three peers: one honest user making direct observations at rate $\mu = 1$, indirect ones originating from the liar at rate $\lambda > 0$ and indirect ones originating from the honest peer at rate $\nu > 0$.

We have assumed independent subject behaviour. It might be interesting to consider the case when direct observations are correlated.

The next extension is then to consider strategic lying. This deals with the case when adversaries do not simply lie maximally, but attempt something more subtle. It would also be interesting to consider random noise instead of fake reports. This would model random failures in components or transmission.

Acknowledgment The authors would like to thank Sonja Buchegger for valuable discussions and Richard Weber for comments on an earlier draft.

References

- [1] E. Adar and B. Huberman. Free riding on gnutella. *First Monday*, 5(10), 2000.

- [2] S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *Proceedings of Multimedia Computing and Networking (MMCN'02)*, 2002.
- [3] J. Chu, K. Labonte, and B. N. Levine. Availability and locality measurements of peer-to-peer file systems. In *ITCom: Scalability and Traffic Control in IP Networks, Proceedings of SPIE*, volume 4868, 2002.
- [4] C. Courcoubetis and R. R. Weber. *Pricing Communication Networks : Economics, Technology and Modelling*. Wiley Europe, 2003.
- [5] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. *Lecture Notes in Computer Science*, 2232, 2001.
- [6] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring. Modelling incentives for collaboration in mobile ad hoc networks. In *Proceedings of the workshop WiOpt'03*, 2003.
- [7] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, 46(1), 2004.
- [8] S. Sarafijanovic and J. Y. Le Boudec. An artificial immune system for misbehavior detection in mobile ad-hoc networks with virtual thymus, clustering, danger signal and memory detectors. In *Proceedings of ICARIS-2004, 3rd International Conference on Artificial Immune Systems*, 2004.
- [9] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. *Advances in Applied Microeconomics: The Economics of the Internet and E-Commerce*, 2002.
- [10] 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*, 2004.
- [11] P. Michiardi and R. Molva. Core: A collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In *Sixth IFIP Communication and Multimedia Security Conference 2002*, 2002.

- [12] 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)*. IEEE, 2002.
- [13] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM)*, 2001.
- [14] T. Moreton and A. Twigg. Enforcing collaboration in peer-to-peer routing services. In *Proceedings of the First International Conference on Trust Management*, 2003.
- [15] M. Carbone, M. Nielsen, and V. Sassone. A formal model for trust in dynamic networks. In *IEEE International Conference on Software Engineering and Formal Methods (SEFM'03)*, 2003.
- [16] R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, 2003.
- [17] A. Fernandes, E. Kotsovinos, S. Östring, and B. Dragovic. Pinocchio: Incentives for honest participation in distributed trust management. In *Proceedings of iTrust 2004*, 2004.
- [18] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the Twelfth International World Wide Web Conference 2003*, 2003.
- [19] J. R. Norris. *Markov Chains*. Cambridge University Press, 1997.
- [20] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, second edition, 2003.




Figure 1: Effect of smoothing for $\theta = 0.8$, $\Delta = 0.4$, $\rho = 0.99$, $r_0 = 0.4$. The distribution in the top graph has been obtained with a point mass, the bottom one with a uniform distribution.

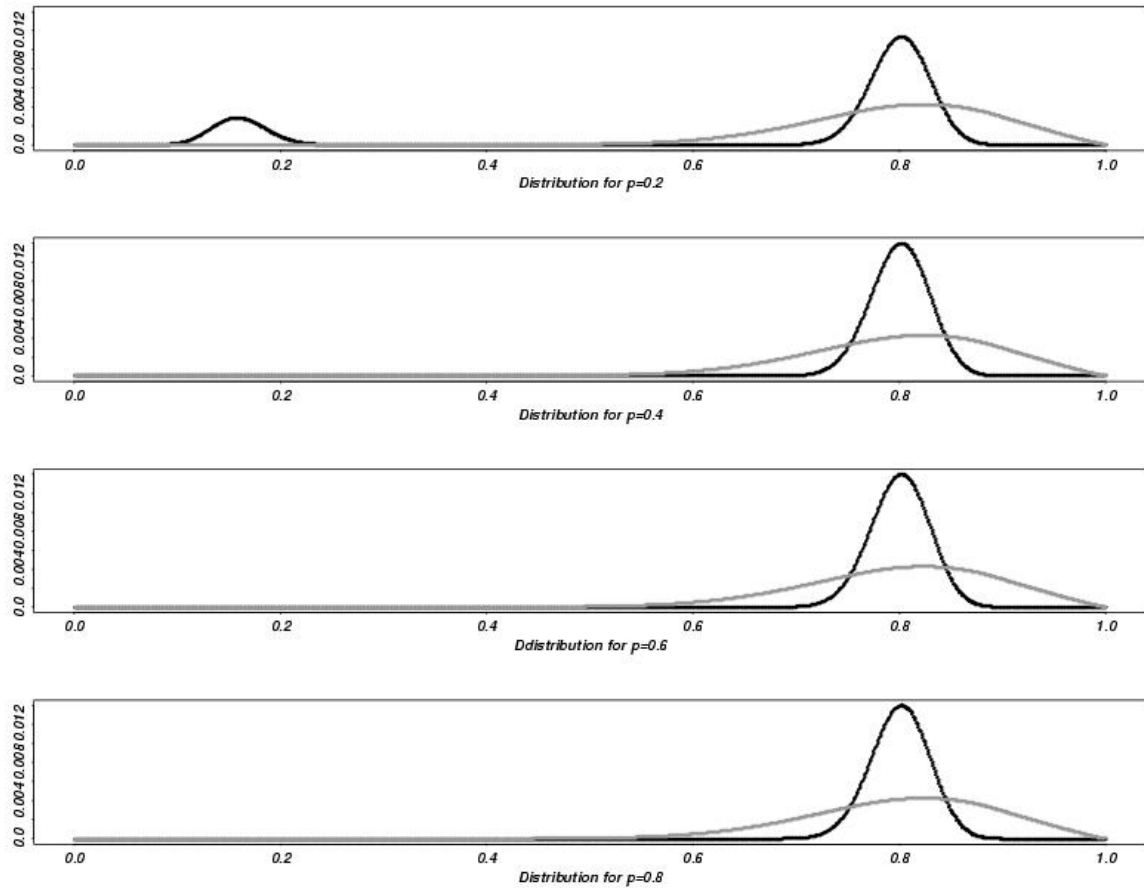


Figure 2: Distribution of R_n for $p = 0.2, 0.4, 0.6$ and 0.8 respectively. The lightly coloured graphs are obtained when $\rho = 0.99$ is replaced by $\rho = 0.9$.




Figure 3: Log-scale plots of the distributions shown in Figure 2. There are two modes for $p = 0.2$ and 0.4 and there is only one for $p = 0.6$ and 0.8 .




Figure 4: Log-scale plots with $r_0 = 0.4$ as in Figure 3 (medium) together with the corresponding distribution for $r_0 = 1$ (dark) and $r_0 = 0$ (light). They coincide except in the first plot for $p = 0.2$ where they are still close.

θ	probability of positive subject behaviour	Section 3.1
μ	direct observation rate of honest user	Section 3.1
λ	indirect observation rate of honest user, i.e. lying rate	Section 3.1
p	probability of an observation being direct	Section 3.1
R_n	reputation value of honest user after n^{th} interaction	Section 3.2
r_0	initial reputation value	Section 3.2
Δ	deviation test parameter	Section 3.1
ω	weighting factor attached to indirect observations	Section 3.2
ρ	discount factor	Section 3.2

Table 1: Notation

Jochen Mundinger graduated from the University of Cambridge in 2001 and received his M.Phil. in Statistical Science in 2002. He is currently pursuing his Ph.D. at the Statistical Laboratory. His interests are in the modeling, performance analysis and optimization of communication systems.

Jean-Yves Le Boudec graduated from Ecole Normale Supérieure de Saint-Cloud, Paris, received the agrégation in mathematics in 1980 and the doctorate in 1984 from the University of Rennes, France. In 1987 he joined Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory where he was manager of the Customer Premises Network group. In 1994 he became professor at EPFL, where he is now full professor and head of the Institute for Communication Systems. He is co-author of the book "Network Calculus". His interests are in the architecture and performance of communication systems.