

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

Jochen Mundinger  
EPFL-IC-LCA  
BC203, Station 14  
CH-1015 Lausanne, Switzerland  
jochen.mundinger@epfl.ch

Jean-Yves Le Boudec  
EPFL-IC-LCA  
BC205, Station 14  
CH-1015 Lausanne, Switzerland  
jean-yves.leboudec@epfl.ch

June 26, 2007

## Abstract

The application of decentralized reputation systems is a promising approach to ensure cooperation and fairness, as well as to address random failures and malicious attacks in Mobile Ad-Hoc Networks. However, they are potentially vulnerable to liars. With our work, we provide a first step to analyzing robustness of a reputation system based on a deviation test. Using a mean-field approach to our stochastic process model, we show that liars have no impact unless their number exceeds a certain threshold (phase transition). We give precise formulae for the critical values and thus provide guidelines for an optimal choice of parameters.

**Keywords:** Mobile Ad-Hoc Network, reputation system, robustness against liars, mean-field approach, performance optimization

## 1 Introduction

The 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. Forwarding their peers' packets, for example, uses up limited battery lifetime without giving any immediate reward. Users themselves do not need to be technology experts; it would suffice if a company produced devices that do not cooperate and therefore have a longer runtime.

Free-riding is a well-known problem in economics and often occurs in situations involving *public goods* [12], that is, goods that are both non-rivalrous (meaning that consumption by one does not limit consumption by others) and non-exclusive (meaning benefits cannot be excluded). In Mobile Ad-Hoc Networks, nodes free-ride in that they increase their utility by not shouldering their fair share of the costs.

On the other hand, altruism has been observed. *Altruism* is the practice of being helpful to other people with little or no interest in being rewarded for one's efforts. The concept has a long history in philosophical and ethical thought. For recent references see [41, 19, 4]. However, we cannot typically rely on altruistic nodes to ensure collaboration in Mobile Ad-Hoc Networks. Moreover, we would like to protect the network from malicious attacks and

random failures. In the absence of suitable countermeasures, performance of the network can dramatically decrease [31, 8, 1].

One solution approach addressing all of these problems is the application of reputation systems. By *reputation* we mean an estimate about a node's actual quality in terms of its behaviour in the network, which is sometimes also referred to as *trust*. Nodes keep track of their peers' behaviour and exchange this information with others in order to compute a reputation value about the other nodes. Nodes with a good reputation are then favoured. For more general articles on reputation systems see [42] and [32]. However, in the context of Mobile Ad-Hoc Networks a reputation system must necessarily be fully decentralized to fit the architecture and to preserve its advantages such as scalability.

By using second-hand information in addition to own observations, an accurate estimate of some *subject's* (the node under consideration) behaviour can be obtained faster. Moreover, a node can have a reputation value about a subject without ever having interacted with it itself. 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 as part of their CONFIDANT Protocol (Cooperation of Nodes, Fairness In Dynamic Ad-hoc NeTworks) [8, 9]. 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 [9] is more complex. It also allows for using second-hand information from trusted peers, where here trust refers to agreement in the reputation values. Therefore each node maintains both a reputation and a trust value about each of its peers.

The system appears to work well, although performance has only been evaluated through simulations of a network with a particular set of assumptions (e.g. on the routing protocol). Further simulations suggested that the deviation test on its own without the trust component performs nearly 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.

We thus consider an abstract model of a reputation system based on the deviation test. We are primarily concerned with the formation of reputation rather than with the detection and response components of a reputation system. The detection component depends on the given Mobile Ad-Hoc Network. For example, most reputation systems suggested for Mobile Ad-Hoc Networks have been assuming Dynamic Source Routing (DSR) [24] and have based the detection component on its properties. In contrast, we merely assume that bad behaviour can be told apart from good behaviour. We also assume that if reputation values can be computed accurately, then there exists a response mechanism using them to obtain the desired effects. Typically, this might mean exclusion of benefits for the misbehaving node.

In an earlier paper [39] we dealt with a simplified one-dimensional model. We now consider the original two-dimensional system. This paper is an extended version of the conference paper [38]. Whereas in the one-dimensional model there is only a notion of good and bad reputation, in the two-dimensional model there is an additional notion of the degree of certainty about the reputation. We analyzed in detail the case of two nodes, one honest and the other a liar. Moreover, we obtained evidence that the results carry through to the general case qualitatively, and to some extent even quantitatively. Our new results confirm the results of the simplified one-dimensional model.

Our model is also relevant in the context of social networks. If a person in the network is confronted with information that is not verifiable, they will probably believe it only if, to

them, it seems likely. However, they will ignore it if, to them, it seems unlikely. Work in this context can be found in [40].

The methodology we have developed for this problem is as follows. We first formulate a stochastic process that models the system based on certain assumptions. We then derive the ‘mean’ ordinary differential equation (ODE) by averaging the dynamics and passing to a *fast-time scaling* limit. We solve the ODE and study its fixed points. Our approach is thus a *mean-field* approach. Full details are given in Sections 4 and 5. Finally, we verify the analytical results by means of simulation.

We find that there is a *phase transition*. That is, there is a threshold proportion of lying below which the lying has no impact. Above it, the lying does have an impact and corrupts the reputation system. We give precise formulae and quantify the impact, thereby providing a performance evaluation of a reputation system based on the deviation test. Alternatively, this phase transition can be phrased in terms of the system parameter controlling the deviation test. We thus provide guidelines for a good choice of parameters and hence a good system design.

This paper is organized as follows. We present related work in Section 2. The precise modelling assumptions are listed in Section 3 and the stochastic process model is formulated in Section 4. We give details of our methodology and provide mean-field results in Section 5 and verify them by means of simulation in Section 6. We conclude in Section 7.

## 2 Related Work

A number of reputation mechanisms have been suggested and studied. A comprehensive survey and more detailed overview of reputation systems for Mobile Ad-Hoc Networks can be found in [10]. Reputation systems are classified according to (1) representation of information and classification, (2) use of second-hand information, (3) trust and (4) redemption and secondary response.

In addition to the CONFIDANT protocol already discussed in the introduction, the Collaborative REputation mechanism (CORE) was introduced in [34]. Here, reputation also takes into account a task-specific functional reputation. OCEAN [7] and SORI [23] are also discussed in more detail in the survey. These systems have all been developed for a fairly specific set of assumptions, in particular assuming Dynamic Source Routing (DSR) [24].

Reputation systems have also been considered in other scenarios, in particular Internet-based Peer-to-Peer (P2P) systems. Aberer and Despotovic [2] suggest a mechanism for P-Grid, a P2P system, that spreads negative information only. Collaboration enforcement in P2P systems has also been considered by Moreton and Twigg [37]. Another mechanism is PeerTrust as introduced by [48]. The reader is referred to [25] for the EigenTrust algorithm, a method to compute global trust values in the presence of pre-trusted peers. Another mechanism is PeerTrust as introduced by [48]. A comprehensive survey and more detailed overview of reputation systems for Internet-based P2P systems can be found in [3]. Another review with a focus on the artificial intelligence literature is given in [43].

Our work is the first analytical approach to evaluate robustness against liars. [50] consider the problem of liars via some models of deception. Their approach is based on the weighted majority technique where the last second-hand information is tested by comparing it to the next direct interaction. The analysis is based only on simulation. [47] is also concerned with filtering out manipulated second-hand information that seems unlikely. However, they

consider quantiles of the Beta distribution rather than distance. This paper, too, is based merely on simulation. In the context of centralized reputation systems, [35] consider incentive mechanisms not in order to stimulate good behaviour in the network, but to stimulate honest reports within the reputation system. They show that honest reporting is a Nash equilibrium.

Alternatives to reputation systems have also been considered, e.g. incentive mechanisms. They make it advantageous for nodes to act in such a way that the resulting social welfare is optimal (e.g., [6]). In general, this can either be achieved by means of pricing schemes [13] or rules [5]. As for pricing schemes, they might involve payments in kind, or virtual or real payments. For example, a virtual currency called nugglets is suggested in [11] and a pricing scheme based on notional credit is used in [15], both in the context of Mobile Ad-Hoc Networks. In [22], the authors examine a micro-payment mechanism for P2P file sharing networks. A simple example of a rule is to force nodes to contribute in order to consume. Rules are used in [14] in the context of P2P Wireless LAN Consortia. Incentive mechanisms typically give rise to game-theoretic problems. Unlike reputation systems, incentive mechanisms do not address malicious nodes and random failures.

A more recent alternative to reputation systems is artificial immune systems. They are aimed primarily at misbehaviour detection and designed so that they adapt to normal behaviour, but they also recognize *new* misbehaviour patterns that were not anticipated in the system design phase [45]. Moreover, artificial immune systems use mechanisms for faster detection of *repeated* misbehaviour [28] and [44]. An important potential advantage of such systems is their inherent randomness that provides diversity at the population level. Even if some computers are vulnerable to an attack, there should be many others that are resistant to the same attack. However, implementations seem to depend very much on the particular application and at the moment there is no scenario in which artificial immune system have proven superior to other approaches.

All three approaches – incentive mechanisms, reputation systems and artificial immune systems – need to take into account both the economic side and the engineering side. For example, identity is an issue in all these systems. For more information, the reader is referred to [21, 18].

## 3 Assumptions

### 3.1 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. In particular, this allows all nodes to be subjects themselves, as is the case in a Mobile Ad-Hoc Network.

### 3.2 Reputation

Each node  $i$  maintains a *reputation value*  $R^i = R_n^i$  about the subject that reflects its belief about  $\theta$  at the time of the  $n$ th interaction. The initial value is  $R_0^i$ . Reputation values are based on two counters, as will be explained in Section 4. 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 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  $\omega$ . Finally, we account for discount factor  $u$  so that the system gradually forgets about old observations. This also allows for tracking subject behaviour that changes.

For tractability, we shall restrict attention to the case of two nodes in the modeling part of the analysis, an honest one 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 because we can focus on one out of the honest nodes by symmetry. Furthermore, it looks like ignoring the other honest ones can be accounted for by increasing the proportion of (necessarily truthful) direct interactions. Moreover, we provide simulation results suggesting that the results carry through to the general case qualitatively, and to some extent even quantitatively.

### 3.3 Interaction Model

In the general case, the nodes make observations when interacting with the subject or with a peer. Interactions depend on the topology of the network and the mobility model for the nodes. We take a high-level view and directly specify the interactions rather than topology and mobility. More specifically, we assume that interactions occur at certain points  $T_n$  in time and that a given resulting observation is direct or indirect with fixed probabilities  $p$  and  $q = 1 - p$  respectively. As a natural model, let the  $T_n$  be the points of a Poisson process and consecutive interactions be independent. Without loss of generality we take the Poisson process to have rate 1, otherwise we can scale time suitably. Whereas for a given Mobile Ad-Hoc Network the values of  $p$  and  $q$  are determined by topology and mobility, we examine the whole range of parameters  $p = 1 - q \in [0, 1]$ .  $p$  large accounts for Mobile Ad-Hoc Networks where direct observations are frequent,  $p$  small accounts for networks where there are predominantly indirect observations.

For ease of notation, we shall often refer to the number of the observation  $n$  rather than its point in time  $T_n$ .

Finally, note that in the case of two nodes with one honest node and one liar, the honest node's second-hand information is necessarily from its lying peer.

### 3.4 Adversary Model

One needs to make precise assumptions about the adversaries' abilities in order to give performance guarantees. We 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.

### 3.5 Performance

As explained in the introduction, we focus on the formation of reputation rather than the detection and response components, assuming merely that that bad behaviour can be told apart from good behaviour (detection) and that – based on accurate reputation values – there exists a response mechanism using them to obtain the desired effects (response).

Notice that liars can easily change their own reputation values to anything they want. The question is whether they can influence the reputation values of the honest nodes.

The *faster* the nodes can obtain *accurate* estimates of the subject behaviour, the better the system will work, but there is a fundamental trade-off between *speed* and *robustness*. By using more second-hand information, an estimate of the subject’s behaviour can be obtained faster – a node can even have a reputation value about a subject without ever having interacted with it itself; however, it is then more sensitive to potentially wrong information and one needs to compromise on accuracy.

We assess robustness in detail. It will then be possible, for example, to choose parameters such that the system will be as fast as possible, subject to a given accuracy condition. In fact, we show that by using a reputation system based on the deviation test, to some extent it is possible to gain in terms of speed without compromising on accuracy.

## 4 Model


We first formulate a stochastic process that models the systems based on the assumptions of the previous section. A natural scheme, suggested by the reputation system in [9] 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 stochastic process  $(\alpha_n, \beta_n)$  for  $n \geq 0$ .

$$(\alpha_{n+1}, \beta_{n+1}) = u(\alpha_n, \beta_n) + \begin{cases} (1, 0) & \text{w.p. } p\theta \\ (0, 1) & \text{w.p. } p(1 - \theta) \\ (0, \omega)\mathbf{1}\left\{\frac{\alpha_n}{\alpha_n + \beta_n} \leq d\right\} & \text{w.p. } q \end{cases} \quad (1)$$

The three possibilities correspond to a positive direct observation, a negative direct observation and an indirect observation respectively (cf. Sections 3.1 and 3.3). Direct observations are always accepted and counted with 1. Indirect observations have to pass the deviation test with parameter  $0 < d < 1$ , modelled by the indicator function, and are weighted by  $\omega > 0$  as described in Section 3.2. Indirect observations are negative, because they are obtained from the liar who is assumed to report extremely negative behaviour (cf. Section 3.4). In any case, both components are discounted individually by  $0 < u < 1$  (cf. Section 3.2). Typically,  $u$  is close to 1.

The stochastic process (1) is a homogeneous Markov Chain with the state space being a subset of the triangular area  $\{(\alpha, \beta) : \max\{1, \omega\}\alpha + \beta \leq \max\{1, \omega\}/(1 - u)\}$  in the first quadrant. This is illustrated in Figure 1. For the parameters chosen to be rational, the process

will take rational values only and the state space is easily seen to be countable, although not easy to describe.



**Figure 1: The state space of the stochastic process (1) is a subset of the triangular area. Dashed lines show points of constant reputation, increasing from bottom to top. Dotted lines show points of constant certainty, increasing from left to right.**

The important quantity of interest is  $R_n = \alpha_n / (\alpha_n + \beta_n)$ , in some sense the proportion of positive observations. We examine how well this compares to the true subject behaviour  $\theta$ . In addition, in this two-dimensional model, there is the notion of the degree of certainty about reputation values, determined by  $\alpha_n + \beta_n$ . The larger this sum, the more certain the node is of its reputation value, the more ‘locked in’ the state. Lines of constant reputation, as well as constant certainty, are also shown in Figure 1.

The initial values are  $(\alpha_0, \beta_0)$ . Note that if  $\omega = 1$  and the deviation test is passed, then starting with  $\alpha_0 + \beta_0 = 1/(1-u)$  will leave  $\alpha_n + \beta_n$  unchanged; starting elsewhere the sum will converge to  $1/(1-u)$ . It makes sense to start with such a converged value, because we would like to allow for tracking behaviour that changes. However, there can be no a priori knowledge of a change, so we cannot simply reset the system to an arbitrary starting value. On the other hand, as the fixed points are not known in advance, we cannot generally choose initial values to correspond to them. We still take initial values satisfying  $\alpha_0 + \beta_0 = 1/(1-u)$  as this choice can be viewed as an extreme case. Moreover, we shall see in Sections 5 and 6 that, the choice of initial value does not influence the process in the long run. If we observe the process settling down at a fixed point starting from such an initial state, then it will also do so from a less ‘locked in’ state with the same reputation value. 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 that differ by a constant. This is due to the discounting. So, we assess convergence (in distribution) to some limiting distribution from which we infer  $\theta$ . We will point out another advantage of the discounting in the reputation system that is essentially due this convergence in distribution.

symbol	meaning
$\theta$	prob. of positive subject behaviour (cf. Section 3.1)
$p$	prob. of an obs. being direct (cf. Section 3.3)
$q$	prob. of an obs. being indirect (cf. Section 3.3)
$\alpha_n$	positive reputation component (cf. Section 4)
$\beta_n$	negative reputation component (cf. Section 4)
$R_0$	initial reputation value (cf. Sections 3.2, 4)
$R_n$	inferred reputation value (cf. Sections 3.2, 4)
$d$	deviation test parameter (cf. Section 3.2)
$\omega$	weighting factor for indirect obs. (cf. Section 3.2)
$u$	discount factor (cf. Section 4)

Table 1: Summary of notation

## 5 Mean-field Approach

Our methodology is as follows. From the stochastic process formulation (1), we derive the ‘mean’ ordinary differential equation (ODE) by averaging the dynamics and passing to a *fast-time scaling* limit. That is, we scale time so that events occur more frequently, that is the honest node makes direct and indirect observations at a higher rate. At the same time, the impact of each observation is reduced by the same factor. A similar approach has also been used in TCP modelling (cf. [36] and [30]). One possibility to turn this type of argument into a rigorous formal limit statement is the fluid limit in [16]. We then derive the solutions of the differential equation and study its fixed points. We confirm the analytical results by means of simulation. This is the topic of the next section.

Thus, our approach can be called a *mean-field* approach. This term has different interpretations in different research communities. In [49], it refers to making an independence assumption for one side of the Chapman-Kolmogorov equations. The mean-field idea was first introduced in physics and has been used to describe systems like plasma and dense gases where interaction between particles is somewhat weak, meaning that the strength of the interaction is inversely proportional to the size of the system. A given particle is seen as under a collective force generated by the other particles. In recent years, the mean-field approach has also found an appeal in the area of communication networks. To name but a few of the many publications, the reader is referred to [27], [46], [26] and [17]. The stochastic approximation framework by [29] is also a mean-field approach. 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 process is determined by a mean ODE. This has been applied in diverse areas, in particular in signal processing and communications.

The mean-field approach is very powerful in that it enables us to study a complex system analytically without losing the important features. As we shall see, the phase transition, as well as the fixed points and critical parameter values, is predicted accurately. What it cannot do, however, is to compare two simultaneous fixed points in terms of their attraction. It does not keep all the information contained in the stochastic process. For this, we would need to investigate a diffusion approximation approach as in [20]. We will gain some insight into this question via the simulations.



Specifically, we consider a family of processes indexed by  $N$ . We will now make the notation explicit again to stress that the subscripts  $n$  are the points  $T_n$  of a Poisson process at rate 1.

$$\begin{aligned}
(\alpha_{T_{n+1}/N}^N, \beta_{T_{n+1}/N}^N) &= \left(1 - \frac{1-u}{N}\right)(\alpha_{T_n/N}^N, \beta_{T_n/N}^N) \\
&+ \frac{1}{N} \begin{cases} (1, 0) & \text{w.p. } p\theta \\ (0, 1) & \text{w.p. } p(1-\theta) \\ (0, \omega)\mathbf{1} \left\{ \alpha_{T_n/N}^N / (\alpha_{T_n/N}^N + \beta_{T_n/N}^N) \leq d \right\} & \text{w.p. } q \end{cases} \quad (2)
\end{aligned}$$

Next, we consider a continuous-time rescaled version. The number of jumps of the process indexed by  $N$  in the interval  $[t, t + \epsilon)$  is Poisson( $N\epsilon$ ), the average jump is of size

$$\begin{aligned}
&-\frac{1-u}{N} \alpha_t^N + \frac{1}{N} [p\theta] \\
&-\frac{1-u}{N} \beta_t^N + \frac{1}{N} \left[ p(1-\theta) + q\omega\mathbf{1} \left\{ \frac{\alpha_t^N}{\alpha_t^N + \beta_t^N} \leq d \right\} \right] \quad (3)
\end{aligned}$$

We obtain

$$\begin{aligned}
\alpha^N(t + \epsilon) - \alpha^N(t) &= \frac{N\epsilon}{N} [(u-1)\alpha(t) + p\theta] \\
\beta^N(t + \epsilon) - \beta^N(t) &= \frac{N\epsilon}{N} \left[ (u-1)\beta(t) + p(1-\theta) + \omega q\mathbf{1} \left\{ \frac{\alpha}{\alpha+\beta} \leq d \right\} \right] \quad (4)
\end{aligned}$$

Dividing by  $\epsilon$  and taking the limit, we are thus led to consider the following deterministic ODE.

$$\begin{aligned}
\dot{\alpha}(t) &= (u-1)\alpha(t) + p\theta \\
\dot{\beta}(t) &= (u-1)\beta(t) + p(1-\theta) + \omega q\mathbf{1} \left\{ \frac{\alpha}{\alpha+\beta} \leq d \right\} \quad (5)
\end{aligned}$$

This system is discontinuous, but linear above and below the line of discontinuity  $\alpha/(\alpha+\beta) = d$ . Above it, for  $\alpha/(\alpha+\beta) > d$ , we obtain

$$\begin{aligned}
\alpha(t) &= c_1 e^{-(1-u)t} + \frac{1}{1-u} p\theta \\
\beta(t) &= c_2 e^{-(1-u)t} + \frac{1}{1-u} p(1-\theta) \quad (6)
\end{aligned}$$

for constants  $c_1, c_2$  whereas below it, for  $\alpha/(\alpha+\beta) \leq d$ , we obtain

$$\begin{aligned}
\alpha(t) &= c_3 e^{-(1-u)t} + \frac{1}{1-u} p\theta \\
\beta(t) &= c_4 e^{-(1-u)t} + \frac{1}{1-u} [p(1-\theta) + \omega q] \quad (7)
\end{aligned}$$

for constants  $c_3, c_4$ . Fixing the initial reputation counter at  $(\alpha_0, \beta_0)$  we obtain the constants and hence the solutions. We find that the system has either one or two fixed points, depending on the parameters of the model.

$$(\alpha, \beta) = \frac{p}{1-u} (\theta, 1-\theta) \quad (8)$$

is a fixed point is a fixed point if  $\theta > d$ . If it exists, it is asymptotically stable and trajectories from  $\alpha/(\alpha+\beta) > d$  are attracted to it. The corresponding reputation value is  $\theta$ .

$$(\alpha, \beta) = \left( \frac{p\theta}{1-u}, \frac{p(1-\theta)}{1-u} + \frac{\omega q}{1-u} \right) \quad (9)$$

is a fixed point if  $\theta \leq d$  or  $q \geq (\theta - d)/(\theta - d + \omega d)$ . If it exists, it is asymptotically stable and trajectories from  $\alpha/(\alpha + \beta) \leq d$  are attracted to it. The corresponding reputation value is  $\pi = \theta p/(p + \omega q)$ . If only one of the two fixed points exists then the trajectories from the other region lead into its region and thus are also attracted to it. That is all trajectories are attracted to it. Otherwise, both are asymptotically stable on their respective region. Hence, we have the following result.

**Theorem 1** *If  $\theta > d$ , (8) is a fixed point of the mean ODE (5). For  $q < q_c = (\theta - d)/(\theta - d + \omega d)$  it is asymptotically stable and all trajectories are attracted to it. The corresponding reputation value is the true  $\theta$ . Otherwise, there is a second, false fixed point (9) and both are asymptotically stable, attracting trajectories from  $\alpha/(\alpha + \beta) > d$  and  $\alpha/(\alpha + \beta) \leq d$  respectively. If  $\theta \leq d$  then only the latter, false one is asymptotically stable and all trajectories are attracted to it. The corresponding reputation value is  $\pi = (p\theta)/(p + \omega q)$ .*

Thus, the reputation system exhibits a phase transition. Assuming  $\theta > d$ , we find a bifurcation in terms of the parameter  $q$ . In the *subcritical regime*, that is, for  $q < q_c$ , the fixed point corresponding to the true  $\theta$  is unique. In the *supercritical regime* where  $q \geq q_c$  there is a second, false fixed point.


In practical terms, this suggests that the reputation system works and that the liar cannot achieve anything if  $\theta > d$  and  $q < q_c$ . 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.

If there are predominantly indirect observations in the network, that is for small values of  $p$  (or large values of  $q$ ), the condition for the second, false fix point to exist as well as the true one is more likely satisfied and the reputation value corresponding to the false fixed point (9) is close to 0.

As mentioned in Section 3, 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 fixed point to be unique when both, positive and negative lying are permitted:  $\min\{\theta, 1 - \theta\} > d$  and  $q < (\min\{\theta, 1 - \theta\} - d)/(\min\{\theta, 1 - \theta\} - d + \omega d)$ .

Alternatively, the phase transition can be phrased in terms of the system parameter  $d$ . For small  $d$  there is only one fixed point, the true one (8). For intermediate  $d$ , there are both the true and the false fixed points and for large  $d$  there is only the false fixed point (9). The exact conditions are given in Corollary 1 and an illustration is provided in Figure 2.

**Corollary 1** *If  $d < d_{c1} = \pi = (p\theta)/(p + \omega q)$ , (8) is the unique fixed point of the mean ODE (5). It is asymptotically stable and all trajectories are attracted to it. The corresponding reputation value is the true theta. If  $d_{c1} \leq d < d_{c2} = \theta$  there is a second, false fixed point (9). Both are asymptotically stable, attracting trajectories from  $\alpha/(\alpha + \beta) > d$  and  $\alpha/(\alpha + \beta) \leq d$  respectively. If  $d_{c2} \leq d$ , then only the latter, false one is asymptotically stable and all trajectories are attracted to it. The corresponding reputation value is  $\pi = (p\theta)/(p + \omega q)$ .*



**Figure 2:** Bifurcation plot showing the existence of fixed points for each  $d$ : As  $d$  increases from 0 to 1 (from left to right along the horizontal axis), the number of fixed points changes from 1 (the true fixed point at reputation value  $\theta$ ) to 2 (the true fixed point at  $\theta$  and the false fixed point at  $\pi$ ) and back to 1 (the false fixed point at  $\pi$ ).


In practical terms, this suggests that the reputation system works as long as  $d$  is sufficiently small. It might not work otherwise, however, and if  $d$  is sufficiently large it will not work at all. Although this qualitative inverse dependence between  $d$  and robustness against liars is entirely as one would expect, the exact dependence might seem surprising. It can be exploited for system optimization as follows. If we are interested in a system that is as fast a possible subject to being accurate, we would choose  $d$  just below the first critical value  $\pi$ . Then this is the largest value of  $d$  for which the true  $\theta$  is still the unique fixed point reputation value. As such it allows for maximal gain in terms of speed using second-hand information while still being accurate. This will protect the reputation systems from liars for a reasonable range of parameters. Given a more general cost function with arbitrary weights on accuracy and speed, we could again compute the optimal choice of the system parameter  $d$ . Thus, we provide guidelines for a good choice of parameters and hence system design.

## 6 Simulations

In this section we report our simulation results. They confirm the analytical results from the previous section. Moreover, we report our simulation results for more than two users. They provide evidence that the analytical results for the case of two users carry through this more general case qualitatively, and to some extent even quantitatively.

We used formulation (1) to compute  $10^5$  steps and then plot  $R_n = \alpha_n / (\alpha_n + \beta_n)$  against  $n$ . The program code was written in Java (Version 1.4.1). We used both the standard Java random number generator and the Mersenne Twister (MT19937) [33] to generate pseudo-random numbers. The lower and upper boundaries in the graphs correspond to reputation values 0 and 1 respectively. The lower and upper intermediate lines correspond to the possible fixed point reputation values  $\pi = (p\theta)/(p+\omega q)$  and  $\theta$  respectively. 100 independent runs were carried out in each case and a typical sample path is shown. We consider the following set of parameters and variations thereof.

**Parameter set 1**  $\theta = 0.8$ ,  $d = 0.4$ ,  $u = 0.99$  and  $\omega = 1$  for various values of  $p$  and  $q = 1 - p$ . We used both the extreme initial values  $R_0 = 0$  and  $R_0 = 1$ . Thus, from the



**Figure 3:**  $10^5$  steps of a typical sample path for parameter set 1 with  $q = 0.2$  and  $R_0 = 0$ . We plot  $R_n$  against  $n$ , the upper line corresponding to  $\theta$  and the lower line to  $\pi = (p\theta)/(p + \omega q)$  (cf. Figure 2).  $R_n$  increases from 0 past  $\pi$  to  $\theta$  and then remains close to  $\theta$ . Thus,  $\theta$  is confirmed as the unique fixed point reputation value as expected from the analytical results.

previous section, the predicted possible fixed point reputation values are 0.8 and  $0.8p$ . The critical lying probability is  $q_c = p_c = 0.5$ .

In addition, we carried out simulations for the case of more than two users. Ten independent runs were carried out for ten users with corresponding formulation, graphs and parameter set. In the graphs we now show the sample paths of two different honest users, one in black, the other in grey, computing  $10^5$  steps per honest user, that is  $10^6$  steps in total.

**Parameter set 2**  $\theta = 0.8$ ,  $d = 0.4$ ,  $u = 0.99$  and  $\omega = 1$  for various values of  $p$  and  $q$  that now no longer add up to 1; the remaining probability corresponds to an observation being indirect from an honest peer. We consider the case of ten honest users and various numbers of liars (corresponding to the probabilities). Again, we used both the extreme initial values  $R_0 = 0$  and  $R_0 = 1$ . If the theoretical results in the previous section for the two user case carry through to this more general case, we expect the fixed point reputation values to be 0.8 and  $0.8p$ .

In Figure 3 we show a typical sample path for parameter set 1 with  $q = 0.2$ , i.e.  $p = 0.8$ , and  $R_0 = 0$ . We note that  $R_n$  increases from 0 past  $\pi$  to  $\theta$  and then remains within its neighbourhood until the end of the simulation. All 100 independent runs showed the same qualitative sample paths. As  $q = 0.2 < q_c$ , this is a subcritical scenario. As expected,  $\theta$  is confirmed as the unique fixed point reputation value. Indeed, we obtained the same results when  $q = 0.2$  is replaced by  $q = 0.4$  and  $q = 0.45 < q_c$ . 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 results.

Qualitatively the same results are obtained in the ten user case, as is shown in Figure 4. Here  $p = 0.6$  and  $q = 0.2$ , corresponding to the case of nine liars and the ten honest users.

In Figure 5 we illustrate the effect of the discount parameter  $u$ . A typical sample path is shown for parameter set 1 except  $u = 0.999$ . The variability around  $\theta$  is smaller and it takes longer for the process to approach  $\theta$ . Again, qualitatively the same results are obtained in the case of ten honest users when  $u$  is changed to 0.999. This is shown in Figure 6.



Figure 4:  $10 \times 10^5$  steps of two honest users' typical sample paths for parameter set 2 with  $q = 0.2$ ,  $p = 0.6$  and  $R_0 = 0$ . We plot reputation values  $R_n$  of two users (one in black, one in grey).  $R_n$  increases from 0 past  $\pi$  to  $\theta$  and then remains close to  $\theta$ . Again,  $\theta$  is confirmed as the unique fixed point reputation value.




Figure 5:  $10^5$  steps of a typical sample path for the same parameters as in Figure 3 except  $u = 0.999$ . The variability is smaller and it takes longer for the process to approach  $\theta$ .




Figure 6:  $10 \times 10^5$  steps of two honest users' typical sample paths for the same parameters as in Figure 4 except  $u = 0.999$ . As in case of two users in Figure 5 the variability is smaller and it takes longer for the process to approach  $\theta$ .

In Figure 7 we consider the supercritical case with now  $q = 0.8$ . Also, unlike in parameter set 1, we take  $u = 0.95$  for clearer illustration. As a side effect, variability is increased. Increasing from 0,  $R_n$  settles down for some time in the neighbourhoods of  $\pi = (p\theta)/(p+\omega q) = 0.16$  and  $\theta = 0.8$  in an alternating fashion. This is in agreement with the mean-field prediction that both  $\pi$  and  $\theta$  are fixed point reputation values. It is due to the discounting that we do not have convergence to a constant, but there is always a positive probability of moving from one fixed point to the other. This can be viewed as another advantage of discounting, because the process cannot get stuck at the false fixed point forever. The same qualitative results are observed in all 100 independent runs and also for  $q = 0.6$  and  $q = 0.55 > q_c$ , only the proportion of time spent near  $\theta$  is higher. Note also that the false fixed point reputation value depends on  $q$  and  $\omega$ :  $\pi = 0.32$  and  $\pi = 0.36$  respectively.

The case of ten honest users is shown in Figure 8. Here,  $p = 0.2$  and  $q = 0.5\bar{3}$ , corresponding to the case of 18 liars and the ten honest users. Again, the results are qualitatively the same and show that both  $\pi$  and  $\theta$  are fixed point reputation values.

In Figure 9 we demonstrate that the starting value  $R_0$  does not influence the process in the long run. Whereas for the solution of the ODE each fixed point has a basin of attraction (Theorem 1), there is always positive probability for the stochastic process (1) to move from one to the other, so the independence of the starting value is as expected. The same is demonstrated in Figure 10 for the case of ten honest users, where, again, now  $R_0 = 1$ .

With a different choice of parameters, the prediction of only one fixed point at  $\pi = (p\theta)/(p + \omega q)$  for the case  $\theta \leq d$  can also be verified. This is shown in Figures 11 and 12 for the case of two users and ten honest users respectively.

Thus, we have verified the results of the mean-field approach in Section 5. In addition, Figures 7 and 9 give us an idea of the proportion of time near the false  $\pi = (p\theta)/(p + \omega q)$  when there are two fixed points. This time increases with  $q$ .

Moreover, simulations have demonstrated that qualitative results as well as the reputation values themselves are the same in the more general case. On the other hand, the condition  $q < q_c$  does not seem to carry through to this more general case; rather, it appears that the corresponding condition for the general case depends on  $p$  as well. This, however, will have to be addressed in detail by extending the analytical approach to the general case.

## 7 Conclusions and Further Work

The use of decentralized reputation systems is a promising approach to addressing free-riding, as well as random failures and malicious attacks, in Mobile Ad-Hoc Networks. In this paper, we analyze a reputation system based on the deviation test independently of specific implementation assumptions. We consider the case of two nodes, one honest and the other a liar. Moreover, we provide evidence that the results carry through to the general case qualitatively, and to some extent even quantitatively. Our results confirm the results obtained for the simplified one-dimensional model addressed in [39].

In regards to the use of second-hand information, there is a fundamental trade-off between speed and robustness. By using second-hand information, an estimate of the subject's behaviour can be obtained faster – a node can even have a reputation value about a subject without ever having interacted with it itself; however, it is then subject to potentially wrong information and one needs to compromise on accuracy. The reputation system based on the




Figure 7:  $10^5$  steps of a typical sample path for parameter set 1 except  $u = 0.95$  with  $q = 0.8$  and  $R_0 = 0$ .  $R_n$  increases from 0 to  $\pi = (p\theta)/(p + \omega q)$  and then settles down for some time in a neighbourhood of  $\pi$  and  $\theta$  in an alternating fashion.




Figure 8:  $10 \times 10^5$  steps of two honest users' typical sample paths for parameter set 2 except  $u = 0.95$  with  $p = 0.2$ ,  $q = 0.5\bar{3}$  and  $R_0 = 0$ .  $R_n$  increases from 0 to  $\pi$  and then settles down in neighbourhoods of  $\pi$  and  $\theta$  in an alternating fashion.




Figure 9:  $10^5$  steps of a typical sample path for the same parameters as in Figure 7 except  $R_0 = 1$ .  $R_n$  decreases from 1 to  $\theta$  and then settles down for some time in a neighbourhood of  $\pi$  and  $\theta$  in an alternating fashion as before (cf. Figure 7). Thus, the starting value  $R_0$  does not influence the process in the long run.






Figure 10:  $10 \times 10^5$  steps of two honest users' typical sample paths for the same parameters as in Figure 8 except  $R_0 = 1$ . Again, the starting value  $R_0$  does not influence the process in the long run.



**Figure 11:**  $10^5$  steps of a typical sample path for the same parameters as in Figure 5 except now  $\theta = 0.4$ ,  $d = 0.8$ . Moreover,  $R_0 = 1$ .  $R_n$  decreases from 1 past  $\theta$  to  $\pi$  and then remains close to  $\pi$ . Thus,  $\pi$  is confirmed as the unique fixed point reputation value as expected from the analytical results.



**Figure 12:**  $10 \times 10^5$  steps of two honest users' typical sample paths for the same parameters as in Figure 6 except now  $\theta = 0.4$ ,  $d = 0.8$ . Moreover,  $R_0 = 1$ . As in the case of two users,  $\pi$  is confirmed as the unique fixed point reputation value.

deviation test resolves this issue to a certain degree. There is a phase transition in that 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. We give precise formulae and quantify the impact, thereby providing a performance evaluation of the reputation system.

Alternatively, this phase transition can be phrased in terms of the system parameter  $d$  controlling the deviation test. We thus provide guidelines for a good choice of parameters and hence system design. For example, for a system that is as fast as possible subject to being accurate, we would choose  $d$  just below the first critical value  $d_{c1} = \pi$ . The deviation test then protects the reputation systems from liars for a reasonable range of parameters.

We have illustrated another fundamental trade-off in terms 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 subject behaviour. It is also due to the discount parameter that we do not have convergence to a constant, but there is always a positive probability of moving from one fixed point to another. This can be viewed as another advantage, because the process cannot get stuck at the false fixed point forever.



We have not yet looked in detail at the rate of convergence. Independent runs of the same simulation showed about the same rate of convergence most of the time. We could compare this to the exponential decay terms in the mean ODE and thereby check whether the analytical results are also correct in this respect.

The mean-field approach turned out to be powerful, as well as tractable. We obtained the fixed points as well as the critical points. It cannot, however, compare two simultaneous fixed points in terms of their attraction. For this, we would need to investigate a diffusion approximation approach. Still, even here we gained insight by looking at the time spent near each of the fixed points in the simulation.

Moreover, simulations for ten honest users suggest that our analytical results carry through to the general case qualitatively, and to some extent even quantitatively; namely, the reputation values themselves are the same, however, the condition  $q < q_c$  appears to depend on  $p$  as well. It is current work to address this in detail by extending the analytical approach to the general case.

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

Another extension is then to consider strategic lying, that is adversaries attempting something more subtle than simply lie maximally. The latter assumption makes it easier to filter out lies with a simple mechanism such as the deviation test, because they are more easily discovered. 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. In this case, as a counter measure, one might also want to think about individually controlled  $d^i$ ,  $i = 1, 2, \dots, N$ , based on the nodes' current information. It would also be interesting to consider random noise instead of fake reports.

Finally, it would be of interest to consider a different interaction model, in particular one with a non-homogeneous population of nodes. There might even be some nodes that never interact with the subject directly.

## References

- [1] I. Aad, J.-P. Hubaux, and E. W. Knightly. Denial of service resilience in ad hoc networks. In *Proceedings of ACM Mobicom*, 2004.
- [2] 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.
- [3] K. Aberer, Z. Despotovic, W. Galuba, and W. Kellerer. *Computational Intelligence, Theory and Application*, chapter The complex facets of reputation and trust. Springer, to appear 2006.
- [4] J. Andreoni and J. H. Miller. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70:737–753, 2002.
- [5] P. Antoniadis, C. Courcoubetis, E. Efstathiou, G. Polyzos, and B. Strulo. Peer-to-peer wireless LAN consortia: Economic modeling and architecture. In *Proceedings of the Third IEEE International Conference on Peer-to-Peer Computing (P2P 2003)*, 2003.
- [6] G. Bamberg and K. Spremann. *Agency Theory, Information, and Incentives*. Springer, 1989.
- [7] S. Bansal and M. Baker. Observation-based cooperation enforcement in ad hoc networks. Technical report, Arxiv preprint cs.NI/0307012, 2003.
- [8] 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.

- [9] S. Buchegger and J.-Y. Le Boudec. The effect of rumor spreading in reputation systems for mobile ad-hoc networks. In *Proceedings of WiOpt 2003: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2003.
- [10] S. Buchegger and J.-Y. Le Boudec. Self-policing mobile ad hoc networks by reputation systems. *IEEE Communications Magazine*, 43(7):101–107, July 2005.
- [11] L. Buttyán and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications*, 8(5):579–592, 2003.
- [12] R. Cornes and T. Sandler. *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press, 2nd edition, 1996.
- [13] C. Courcoubetis and R. R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley Europe, 2003.
- [14] C. Courcoubetis and R. R. Weber. Asymptotics for provisioning problems of peering wireless LANs with a large number of participants. In *Proceedings of WiOpt 2004: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, 2004.
- [15] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring. Modelling incentives for collaboration in mobile ad hoc networks. *Performance Evaluation*, 57:427–439, 2003.
- [16] J. G. Dai. Stability of fluid and stochastic processing networks. Author’s class notes, 1999.
- [17] D. A. Dawson, J. Tang, and Y. Q. Zhao. Balancing queues by mean field interaction. *To appear in Queueing Systems*, 2005.
- [18] J. R. Douceur. *Peer-to-Peer Systems: First International Workshop, IPTPS 2002*, volume 2429 / 2002 of *Lecture Notes in Computer Science*, chapter The Sybil Attack, pages 251–260. Springer Berlin / Heidelberg, 2002.
- [19] U. Ebert and O. von dem Hagen. Altruism, redistribution and social insurance. *Review of Economic Design*, 5:365–385, 2000.
- [20] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence*. Wiley, 1986.
- [21] E. J. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.
- [22] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. *Lecture Notes in Computer Science*, 2232:75–87, 2001.
- [23] Q. He, D. Wu, and P. Khosla. SORI: a secure and objective reputation-based incentive scheme for ad hoc networks. In *IEEE Wireless Communications and Networking Conference (WCNC 2004)*, March 2004.
- [24] Y. Hu, D. Johnson, and D. Maltz. The dynamic source routing protocol for mobile ad hoc networks (DSR). <http://www.ietf.org/internet-drafts/draft-ietf-manet-dsr-10.txt>, July 2004.
- [25] 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.
- [26] W. Kang, F. P. Kelly, N. H. Lee, and R. J. Williams. Fluid and Brownian approximations for an Internet congestion control model. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, 2004.
- [27] F. I. Karpelevich, E. A. Pechersky, and Y. M. Suhov. Dobrushin’s approach to queueing network theory. *Journ. Appl. Math. Stoch. Anal.*, 9(4):373–398, 1996.
- [28] J. Kim and P. J. Bentley. Towards an artificial immune system for network intrusion detection: An investigation of dynamic clonal selection. In *Congress on Evolutionary Computation (CEC-2002)*, pages 1015–1020, 2002.
- [29] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer Verlag, 2nd edition, 2003.
- [30] M. A. Marsan, M. Garetto, P. Giaccone, E. Leonardi, E. Schiattarella, and A. Tarello. Using Partial Differential Equations to Model TCP Mice and Elephants in Large IP Networks. *To appear in IEEE/ACM Transactions on Networking*, 2005.
- [31] S. Marti, T. J. Giuli, K. Lai, and M. Baker. Mitigating routing misbehavior in mobile ad hoc networks. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM)*, Boston, 2000.

- [32] H. Masum and Y.-C. Zhang. Manifesto for the reputation society. *First Monday*, 9(7), 2004.
- [33] M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modelling and Computer Simulation*, 8(1):3–30, 1998.
- [34] P. Michiardi and R. Molva. Core: A Collaborative REputation mechanism to enforce node cooperation in Mobile Ad Hoc Networks. In *Proceedings of the Sixth IFIP Communication and Multimedia Security Conference 2002*, 2002.
- [35] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359–1373, 2005.
- [36] V. Misra, W. Gong, and D. Towsley. A Fluid-based Analysis of a Network of AQM Routers Supporting TCP Flows with an Application to RED. In *Proceedings of ACM SIGCOMM 2000*, 2000.
- [37] T. Moreton and A. Twigg. Enforcing collaboration in peer-to-peer routing services. In *Proceedings of the First International Conference on Trust Management*, 2003.
- [38] J. Munding and J.-Y. Le Boudec. Analysis of a reputation system for mobile ad-hoc networks with liars. In *Proceedings of WiOpt 2005: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, pages 41–46, 2005.
- [39] J. Munding and J.-Y. Le Boudec. Analysis of a robust reputation system for self-organised networks. *European Transactions on Telecommunications, Special Issue on Self-Organisation in Mobile Networking*, 16(5):375–384, October 2005.
- [40] J. Munding and J.-Y. Le Boudec. The impact of liars on reputation in social networks. In *Proceedings of Social Network Analysis: Advances and Empirical Applications Forum*, Oxford, UK, July 2005.
- [41] D. C. North. *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, 1990.
- [42] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems: Facilitating trust in internet interactions. *Communications of the ACM*, 43(12):45–48, 2000.
- [43] J. Sabater and C. Sierra. Review on computation trust and reputation models. *Artificial Intelligence Review*, 24:33–60, 2005.
- [44] 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. *International Journal of Unconventional Computing*, 1:221–254, 2005.
- [45] A. Somayaji and S. Forrest. Automated response using system-call delays. In *Proceedings of the 9th USENIX Security Symposium*. The USENIX Association, 2000.
- [46] Y. M. Suhov and N. D. Vvedenskaya. Dobrushin’s mean-field approximation for a queue with dynamic routing. *Markov Processes and Related Fields*, 3(4):493–526, 1997.
- [47] A. Whitby, A. Jösang, and J. Indulska. Filtering out unfair ratings in bayesian reputation systems. In *Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS’04)*, 2004.
- [48] L. Xiong and L. Liu. PeerTrust: Supporting Reputation-Based Trust in Peer-to-Peer Communities. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 16(7):843–857, July 2004.
- [49] B. Ycart. *Modèles et algorithmes markoviens*, volume 39. Springer Verlag, 2002.
- [50] B. Yu and M. Singh. Detecting deception in reputation management. In *Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS’03)*, 2003.