HELSINKI
INSTITUTE FOR
INFORMATION
TECHNOLOGY

# Is Your Neighbor a Traitor? Distributed Reputation Management In Member-Initiated Virtual Communities

Katri Sarkio, Martti Mäntylä

March 15, 2006

HIIT
Technical
Report
2006-1

# Is Your Neighbor a Traitor? Distributed Reputation Management in Member-Initiated Virtual Communities

Katri Sarkio, Martti Mäntylä

HIIT Technical Reports 2006-1

ISSN 1458-9478

Notice: The HIIT Technical Reports series is intended for rapid dissemination of articles and papers by HIIT authors. Some of them will be published also elsewhere.

# Is Your Neighbor a Traitor? Distributed Reputation Management in Member-Initiated Virtual Communities

Katri Sarkio and Martti Mäntylä
Helsinki Institute for Information Technology
P.O.Box 9800,02015 HUT, Finland
{Katri.Sarkio; Martti.Mantyla}@hiit.fi

March 15, 2006

## Abstract

Misbehaving users and fraud are problems plaguing virtual communities and limiting their growth and success. To protect the shared interests of the community members and to reveal the traitors, various kinds of *reputation systems* has been developed to record users' past behaviour and make it transparent to others. Mostly, these systems are centralised: a privileged node sees all transactions between the users, and records them for reputation computing. This requires that the party providing the privileged node has an incentive for gathering and storing the data, such as in many client-server based services where the service provider can cover the cost of reputation management by income from some other source. Unfortunately, in peer-to-peer environments this does not apply: in a community of equals, there may not exist an *a priori* peer with such an incentive. If so, centralised reputation systems are not applicable. In this paper, we argue for *distributed reputation management* that matches the scenario where all peers are (initially) treated equal. Thus, we view reputation management as a service of the P2P platform that maintains a distributed record of the transactions between the peers, and provides reputation information to them at request. We aim to show that (i) this concept matches the special characteristics and requirements of real-world scenarios of member-initiated virtual communities that would be outside the reach of centralised systems, and thus has social value; (ii) it facilitates solving inherent problems of privacy that otherwise may render centralised reputation systems unattractive; and (iii) the concept is also technically feasible on the basis of current and developing Internet technology. We conclude with the agenda for further research in the direction of distributed reputation management.

KEYWORDS: Reputation management, virtual community, peer-to-peer, P2P

# 1 Introduction

Misbehaving users and attempts at fraud threaten the growth and success of virtual communities. Ephemeral identities and spoofed transactions are an issue and fraudulent behavior can, in the most severe cases, lead to financial losses. For instance, in a community where users promote their professional skills, if a user's demonstration of skills ends up in the name of someone else the original content producer may lose a job opportunity.

To record users' past behavior and to reveal the misbehaving members, a variety of systems have been developed to produce reputation estimates. A typical reputation system currently in use has a trusted third party (TTP) controlling the reputation management. Conversely, a peer-to-peer (P2P) based community has no central authority available and the users themselves control the interactions. In the symmetric setting of equals, the members have no a priori roles and hierarchy. The members' roles emerge from the activities. A P2P community is distributed in its character and the peers contribute, for example by offering storage capacity. The users have similar kinds of interests, they get to know each other through the digital interactions and tend to form social networks. The members of these communities hide their true identity, but are aware of each other, because of the used pseudonyms, public keys or similar that are, in fact, the users' identifiers.

In this paper, we examine what are the design principles of a distributed reputation management approach if it is to support the character and requirements of virtual P2P communities. We present the principles with respect to the social value of the community, the users' privacy, the relevance of the reputation information to the requestor's interests and scalability when new users join the community. We examine the problem area from the perspective of a system designer in the context of a self-organising and socially oriented member-initiated virtual communities and thus we do not consider the organizational aspects.

The main contribution of this paper are the criteria when the distributed approach to reputation management is justified and when the centralised approach is more useful. Our conclusion is that principally, the centralised approach is well-founded in communities where the members' roles are inherently asymmetrical, such as consumers/content providers and customers/a bank. In contrast, the distributed approach is justified in a community of equals.

# 2 Problem statement

In the traditional approach, centralised reputation management is motivated because the members' roles are characteristically asymmetrical. The TTP has the incentive to manage the users' reputation information as it can cover the costs by income from other sources. Yet, the traditional setting has major shortcomings. Firstly, the TTP typically has one fixed mechanism to present the members' reputation estimates, which are based on their past behavior. The fixed approach ignores the fact that the members of a community need the reputation information in different situations. But, the TTP has little or no incentive to provide tailored reputation estimations if that requires investments for more storage and processing capacity, in addition to the investment of new pieces of software. Secondly, the reputation information consists of data which describes the user as a person, similarly to the definition, for instance, in the EU data protection directive (Article 2) [9]. Consequently, the possible misuse of the reputation information is a privacy concern, even though the information must be somewhat public to be of any use. From the TTP's point of view the publicity is important because the more transactions the members of the community conduct producing the reputation information, the bigger figures the TTP has for advertising and promoting its volume. This dilemma of privacy and publicity requires awareness of the reputation system designers. Thirdly, the TTP can be the bottleneck of the system when new members join the system and increase the required processing and storage capacity.

The traditional situation of a centralised TTP and users is opposite to the initial setting of a member-initiated virtual community. The P2P community is composed of equals and may not have naturally a priori defined roles. If the reputation system is centralised on a P2P community, one of the members is forced to take the role of the administrator and the TTP. However, the members cannot be sure whether the user acting as a TTP is trustworthy or a traitor, which raises e.g. the risks related to the users' privacy. Also, the scalability of the system is an issue, when the community does not inherently have a member who has the incentive to make the investment in the equipment required for the administrator's and TTP's roles. Accordingly, if the P2P community lacks of a suitable reputation management, the benefit of transparency in members' behavior is lost.

If a reputation management system is to support the character of a P2P community, it should introduce personalisation, effective information lookup systems and scalability. Taking the reputation requestor's current situation into account is paramount, which would be possible, if the reputation estimates were produced locally. Tailoring the estimates to the requestor's needs requires storage and processing capacity only from the requestor's own device. However, getting more personalised information is a clear benefit and further, an incentive to a user to load his device. Producing the personalised estimates locally is not an issue as it sets no special processing requirements for a user's device [36]. Then, to improve the members' privacy, the freely available massive amounts of reputation information should be made more difficult, which would be possible if the information requests were limited. Also, in P2P networks scalability would not be an issue, because if new users join the system, the processing and storage capacity equally increases. P2P networks provide a robust and resilient ground for the interactions [5], even if the missing TTP highlights the importance of the management of the members' identifiers. Overall, the distributed reputation management have to fulfill requirements, which we set out to define.

## 3 Background: Member-initiated virtual communities

Constructing distributed systems that support users' interactions requires understanding of the communities basis and character. To describe the character of virtual communities, C. E. Porter [26] reports a variety of categorisation approaches that researchers have taken. These categorisation approaches are based on business type the purpose of revenue generation, in the discipline of information systems the supporting communication technology and, in sociology, the structure and location of the interactions.

A multidisciplinary approach is necessary when examining a system that relates to many of these aspects. Therefore, we follow the multidisciplinary typology that Porter [26] proposes. In addition to the member-initiated establishment type and the social orientation, we find important the divisions of small/large and open/closed communities, which are similar to Porter's idea of pattern of interaction. These divisions have an impact on e.g. activities to join the community, the factors that affect the reputation gaining as well as the relation between the role of the community and the user's own reputation. For instance, a small number of users can initiate a closed discussion forum, which can be joined only by invitation, whereas some communities have no requirements for joining the group.

In addition to the typological differences, virtual communities also have similarities. Porter identifies five common characters of these communities, which help us to outline the constraints having impact on the P2P based reputation management. The first, the *purpose* means the content for interaction. The second is the *place*, which is the extent of technology mediation of interaction. This means that the interaction is guided by protocols and rules of actions at the same time as the users have the sense of place. This sense is further composed of the community's physical structure and socio-cultural properties. Thereupon the third, the *platform* enables the users interaction. The fourth is the *population interaction structure*, which means the pattern of interaction in small groups, larger networks and publics. Finally, the last common character is the *profit model*.

Here, we consider the 'purpose' as the members' shared interest in a particular topic area such as users that share self-made movies. The users essentially are aware of each other in the 'place', because of their digital identifiers - even though, this does not exclude the fact that the users may know each other in true life. Nevertheless, they interact with each other in different roles, which can be for example a moviemaker and an evaluator. The interactions can be movie downloads and evaluations, which we term transactions similar to S. Marti and H. Garcia-Molina [23]. Then, the 'platform' and the underlying technology solutions provide the necessary services for the interaction. These services are for example content and reputation information search and retrieval. In the other way around, the computers and other personal devices enable and support the social ties among the members in the 'population interaction structure'. Moreover, we consider the community as a network that has a finite number but variable types of memberships with more and less active members. The reputation information is of value as not all users know each other. But, the size of the community must also be small enough so that the members have the possibility to meet again in the future [4]. Lastly, the 'profit model' of a community can be revenue oriented in the sense that, e.g., a user's reputation has an impact on the on-line auction prices [14]. Also, the reputation can help a user to market his professional skills and further earning income.

To sum up, we term these five constraints later all together as a *context* meaning an interaction environment of a member-initiated and socially oriented community. The context has a set of users that are the members of the community, they share a similar area of interest and they communicate and conduct transactions in different roles by digital means. They use different applications and devices, whereas the underlying platform and network provide services that are required for the communication.

To better understand gaining reputation, the emergence of roles and generated social value in distributed reputation management, we shall next look at three examples of P2P communities.

## 3.1 Example: Teachers' community

In our first example, teachers from a variety of schools form a P2P community, where they produce and share educational material. The teachers are equal members of the community, but the educational material of high-quality stands out from the rest. Here, other members' positive feedback on a teacher's contribution of good quality and her increasing reputation on the community incentives her to contribute more.

In this kind of a community, the distributed reputation management is necessary, because then none of the members is forced to take the central role of an administrator and a TTP. Instead, the members' different roles are transparent and emerge from each teachers' activities; who is more active in producing new material and who is more active in commenting others' work. As all teachers need to produce the educational material in their work, starting the work from a semi-finished version that a peer provides is of value. Additionally, the available material is of different value to teachers in different subjects, which supports personalisation. Moreover, producing the educational material together, the same topic can be taught in different subjects, for instance geography in a lesson of a foreign language. This again increases to the social value of the teachers' community.

## 3.2 Example: Micro-movies community

Our second example is semi-professionals that produce micro-movies as demonstrations of skills. No centralised party organises the activities and therefore, the producers share the movies in a P2P network. Actually, the P2P community is composed of the moviemakers and a smaller group of audience specialised in the area. The primary goal of the moviemakers is to catch the audience's interest and get feedback. In addition to the P2P sharing, the moviemakers present their work in a variety of film festivals. But, these festival organisers have no incentives to take the role of a TTP in sharing the movies. Even if they would, they would most probably take a benefit of it and make profit, which does not increase the social benefit of the community.

Compared to the previous example, this community is more competitive in character as the moviemakers compete for public recognition and awards. In this case, the movies are the means to gain reputation and it is important to maintain the information who has produced the movie, who has been the writer, the cutter and so forth. In fact, the community has inherent roles as well as feedback of different importance, i.e. evaluations from the competing moviemakers and the target audience. Taking the nuances of roles and importance of feedback into account in producing reputation estimates requires an advanced reputation management mechanism, which could be implemented locally. In this kind of a community, a distributed reputation management approach also is motivated as no party has a reason to take the role of a TTP, but still the reputation is of importance to the community.

### 3.3 Example: Ubimedia community

The last example is Ubimedia, which is a future scenario and means an electronic extension of an urban area. In this case, cafés, shops, barbers, etc. and the consuming audience form the P2P community. The goal of Ubimedia is to give life to the urban area to support media intensive mobility of consumers with interaction forums that are bound to location and context. The consumers join the community via ad hoc networks and can, e.g., recommend a café to others, which increases the café's reputation.

In this case, the P2P community is scattered and the tradespeople compete for the customers who come and go. Therefore, any of the party cannot reasonably be chosen to be a trustworthy TTP. Also, the businesses change in sale rooms, which makes the community members' roles dynamic and not possible to determine a priori. Essentially, the social benefit of having distributed reputation management is that it enriches the understanding of the urban area and the consumers' lives.

## 4 Requirements analysis for distributed reputation management

The research field of reputation management, especially in the P2P area is somewhat scattered. To assist reputation system design in the P2P environments, S. Marti and H. Garcia-Molina [23] propose a taxonomy of a three phase division. They break the systems down to *information gathering*, *scoring and ranking* and *response* components. They also identify the necessary properties of these functionalities. However, in reputation management, trust and reputation are often mixed and the computational models do not take the sociological foundation of these concepts into account [24]. In fact, our approach mirrors Marti's and Garcia-Molina's approach as we bridge the technical functionalities they present and the requirements derived initially from the member-initiated virtual communities needs. We build this bridge to motivate distributed reputation management from the following three essential aspects:

- Interacting parties: a virtual community and individual users;
- Identity management; and
- Underlying architecture.

We emphasise the identity management due to its fundamental role in reputation management, i.e. reputation is bound to users' identifiers and above all, it is personal data. Next, we discuss requirements that these three aspects deriver to reputation management before looking at the related concerns.

### 4.1 Common requirements for reputation management

A vital requirement for the existence of the virtual communities is the members' cooperation. P. Kollock [16] summarises the fundamental work that R. Axelrod

[4], E. Ostrom [25] and M. Godwin [11] have conducted on examining the users' cooperation and a variety of design principles for making a virtual community successful. Moreover, e.g., P. Resnick et al. [29] present desired properties more specific to reputation systems. In the following, we summarise the most important cooperation promoting requirements:

- A community requires reoccurring interaction and continuity promotion;
- Information about the members' past behavior must be available and aggregated;
- The members of the community must be identifiable; and
- Long-term identities are required to motivate the existence of future interaction, i.e. the belief that the member's might meet again [4].

A reputation management mechanism provides transparency in the members' behavior and further, supports the community's goal and vital requirement of having cooperative members. Essentially, the mechanism provides an incentive structure for the members of a virtual community to behave in an expected manner. The incentives mean that a user can gain reputation with a reasonable amount of successful transactions and the unsuccessful transactions have a decreasing effect. In the first place, the members of a community can gain reputation when they conduct transactions with each other. Each of these transactions generate a piece of reputation information. The information can be based only on the metadata of the transaction. It also can contain the interacting users' subjective evaluations on each other and the success of the conducted transaction.

In fact, reputation information is important in trustworthiness estimations [29, 17] and the similarity of interests is the most valuable pieces of it [15]. Also, in reality, people tend to trust people they know more than strangers. Accordingly, much research effort has been put on collaborative or social filtering, where a mechanism matches a user's interests to others' recommendations and opinions. A recommendation system can assist users in finding the content that matches their tastes and interests. Then, when a users needs to evaluate whether to conduct a transaction with his peer or not, the reputation management mechanism generates an overall figure that describes the peer's reputation. For example, C. Dellarocas [7] and G. Zacharia et al. [37] propose to use the collaborative filtering approach to assist the users in evaluating each others' reputation. The evaluation is done by comparing and evaluating the recommendations that others have given in the context of electronic trading. Additionally, the social network that the user belongs to can be used as a basis for evaluating the reputation such as, e.g. L. Mui [24] proposes. In addition to the similarity of others' opinions and the social network, the available reputation information is related to different kinds of transactions and has different importance depending on the information requestor's current interests. Actually, an advanced reputation system has a reputation management mechanism that examines the reputation information that it processes in more detail, see e.g. [6, 21, 34, 35].

A reputation system must make it more difficult for the users to get easily rid of bad reputation by switching their identifiers [37]. For instance, if the reputation of a fraudulent member decreases under the newcomers' then switching the identifier after the frauds is beneficial. Actually, a properly designed mechanism separates the members with same reputation value - whether they are newcomers or they have a long history with successes and frauds. Moreover, a well designed reputation mechanisms finds out the endeavours of gaining reputation e.g. by selling matches and then cheating in a big deal. We identify a following set of requirements to a reputation management mechanism.

- The user's own experiences are important;
- Management of the user's social networks, i.e. the reputation information received from a friend or a fried-of-a-fried (FOAF) is more important than information received from strangers;
- Processing of received recommendations and own experiences;
- A user must gain reputation with a reasonable amount of successful transactions;
- Unsuccessful transactions must have a decreasing effect to the reputation;
- The mechanism must separate newcomers and members with long history even if they have the same reputation value;
- A newcomer must be assigned a lowest possible reputation value [10];
- A fraud must not decrease the reputation below the newcomers' value; and
- The mechanism must separate the endeavours of gaining reputation on false basis.

## 4.2 Requirements specific to the distributed approach

In this study, we consider reputation management as a technical solution, which demands "little or nothing in the way of change in human values [...] of morality [12]". Essentially, the reputation information is data that is available on the P2P network and is based on the users' past transactions. With this in mind, we raise here two main requirements that a user derives to reputation management mechanism in the distributed approach, even if the research area of human-computer interaction is wide.

- The reputation management mechanism must adapt to the user's current situation; and
- Users must be able to make their decisions separate from others' decisions [4].

For a P2P system to be accepted in a community, J. Pouwelse et al. [28] list four important properties from the file sharing systems' point of view. The first is the *high availability* of the requested information and the second is that the user gets *no fake files.* The third is the *ability to deal with flashcrowds*, i.e. situations where the a number of users are suddenly interested in one particular downloadable file. Then, the last property is a relatively *high download speed.* In addition to the file downloads' point of view, these properties play a key role in P2P systems also from the reputation management aspect. This is because in the distributed reputation approach, the users similarly share the reputation data.

In a P2P context, the peers typically join and leave the network dynamically, which is often termed as churn. Therefore, the design of the underlying architecture must take care of the trustworthiness in reputation information storage, share, search and processing; thus many proposals are available on distributed reputation systems, e.g. [1, 31, 34], and P2P based recommendation systems, e.g. [27]. Essentially, the architecture must support the processes related to gaining reputation and handling with the recommendation requests and responses.

According to the findings of K. Lai et al. [20], a system that relies only on the user's own information on others' past actions fail as the population size increases. But, with supporting infrastructure the information can be shared and it scales better to larger communities. This underlines the fact that in distributed systems well designed information storage and retrieval solutions are essential. We list the number of features as follows, which the architecture has to take care of to be able to support successful distributed reputation management.

- A mechanism to bootstrap the system in the initial phase;
- Routines for the new members to join the community and establish an identifier;
- Identity management;
- Efficient information search and retrieval; and
- The integrity of reputation information.

One fundamental question relates to the ownership of the reputation information; who is in charge of the information? In distributed reputation management, the reputation information gained in one context actually is the property of that community. This is because none of the member's of the community can be identified as the only responsible data controller [22, 36]. Here, we enlarge the Marti's and Garcia-Molina's [23] three division of the reputation system functionalities from the information ownership aspect. In a distributed context, when a user needs to evaluate his peer's reputation he requests recommendations from the other members of the community - not from the particular peer whose reputation is under evaluation. Essentially, the peer whose reputation information is requested cannot decide if the requestor from the same community receives the information.

The reputation information requests and responses relates closely to identifying the members' of the community: who requests and receives the reputation information and to whom the information relates? We summarise the list that S. Holtmanns et al. [13] present from a P2P system's point of view and extend it from a user's point of view. Accordingly, we list the requirements specific to a user's identifier.

- A new identifier must be possible to create;
- The system must have low costs because of the user creates a new identifier;

- Creating a new identifier must generate reasonable costs to a user;
- The identifier must be hard to forge; and
- The identifiers must be spoof resistant, i.e. other users must be able to verify that the particular user is behind the communication [23].

To put it briefly, the character of member-initiated virtual communities sets special requirements for reputation management that relate to: the members' cooperation, tailoring the reputation estimates for the information requestor's current interests, protection of the user's privacy and robust but still adaptable underlying architecture.

# 5 Challenges

Meeting the cooperation promoting requirements of the community is not straightforward and the amount of concerns increases along with the population interaction structure and size. Benefits of having the reputation management mechanism, must be in line with the community's goals. Like Porter [26] reports, small groups tend to have strong social ties and limited amount of members that are highly interactive. Consequently, this kind of communities have only little worries from the reputation management point of view. But, when the size of the group increases, the social ties weaken as not all of the members know each other and the concerns of others willingness to cooperate increases.

As the community is fundamentally composed of individuals, the users' behavior and expectations have a great impact on the group itself. Marti and Gracia-Molina [23] list a number of possible fraudulent behavior, where the main division is between *selfish and malicious peers.* The selfish peers wish to minimise their contribution and maximise their benefits, thus leading to a problems termed free riding [2] and the tragedy of commons [12]. Actually, P. Kollock and M. Smith [18] state that the root problem of cooperation is a *social dilemma*, which means that in many situations a user's rational behavior is not consistent with the group's best outcome. Axelrod [4] has examined this problem from the users' decisions point of view. He presents that in the P2P type of environment the users cannot be forced to behave in a certain manner and the user cannot be sure whether the others decide to cooperate or defect next. He also points out that the users choices are not necessarily rational or conscious decisions. This particularly crystallyses one of the key problems in reputation management: modelling human notion of trust in a computation format is complicated. Even if the division of the user's own expectations, recommendations received from friends, a FOAF and strangers is similar to true life, implementing it into a distributed environment is not straightforward.

The other type of misbehavior, the malicious peers, wish to cause harm to other users or the whole system alone or as joint forces. Alone they can, e.g., share corrupted files and together further attack the system, e.g. by collusion, which means misleading increments in reputation via cooperation [23]. Moreover, malicious peers can disrupt a reputation system use via, e.g. denial of service (DoS) attacks [23].

A user's reputation is linked to the identifier that he gets when joining the community. These identifiers vary from different non-unique identifiers, such as nicknames, to globally unique cryptographic keys. The procedure for getting the identifier varies from easily created e-mail accounts to, e.g., trusted computing platforms and the authentication keys in Subscriber Identity Module (SIM) on a user's mobile phone. However, producing the globally unique identifiers is difficult, even if some proposals are available, such as the Third Generation Partnership Project's (3GPP) [32] Generic Authentication Architecture (GAA) in generating certificates to mobile users [33].

In the reputation queries, a concern is that a malicious peer can monitor a user's reputation queries, profile his activities and thus violate his privacy. In the case of one globally unique identifier, the user can always be traced and profiled, when someone knows this identifier. Actually, to be more precise, this is not true within closed forums. In these communities, before receiving the reputation information linked to the user's unique identifier the requestor has to be accepted to the community or get the information from a user who is already a member of the closed community. But, in the case of non-unique identifiers, the user is in control over revealing his different identifiers. Basically, the user can have different identifiers in different communities, but also within one community. Nevertheless, like our earlier privacy analysis [36] presents, preventing the tracing of a single user is more difficult than the tracing of users on a large scale.

The underlying architecture has a number of concerns related to privacy, availability and reliability of the reputation information. We list the main concerns as follows:

- The users are not in control of their personal information;
- The others' willingness to cooperate is an uncertainty regardless of their past behavior [4];
- Collusion (misleading increments in reputation via cooperation);
- Churn (users join and leave the network dynamically);
- Free riding;
- Tragedy of the commons;
- Availability of the requested information; and
- The available information has different relevance related to the requestor's current situation and concerns.

A fundamental procedure that the virtual communities must deal with is the newcomers. The problem the newcomers procedure has, relates closely to the members who intentionally defect and then join over again. One concern related to the reputation management is users who have multiple identifiers and who generate self-recommendations by these identifiers recommending each other, also known as the sybil attack [8]. To deal with this issue, e.g., J.-M. Seigneur et al. [30] propose that pieces of the reputation information are own observations and the recommendations also are counted. Moreover, unforgeable cryptographic identifiers offer a good protection against malicious endeavours such as whitewashing [23], which means that some misbehaving peers switch their identifiers constantly and rejoin the community with an innocent reputation. To conclude, we list the main concerns related to newcomers and identity management as follows:

- Newcomers have no history information available;
- Intentional reputation purge by rejoining as a newcomer;
- Users can have discrete identifiers and related reputations in a variety of communities; and
- Obtaining evaluations from the users is challenging [29].
- Whitewashing and ephemeral identifiers;
- Identity thefts; and
- Sybil attack and spoofed transactions.

# 6 Discussion

We have argued that distributed reputation management assists in increasing the social value of a P2P community. Reputation management must be distributed in communities, where the members' roles are symmetrical and no party has an a priori role of a TTP. The situation of a missing TTP changes the traditional TTP/users trust relations. But, the virtual community actually has an inherent trust hierarchy as the members' roles emerge from the activities and cooperation. In P2P communities, the user's friends and some peers who have a good reputation can become more trustworthy and the user considers their opinions and evaluations more important than others'. This moves their roles closer to the characteristics of the TTPs.

In addition to the *inherently emerging roles* the following arguments support distributing the reputation management. Firstly, the distributed reputation management *does not require prespecified infrastructure.* A user can download a peer software to his device and join the community. However, the procedure for joining the community is a policy issue, if the community can be joint only via an invitation and the membership must be applied for.

The second argument in favor of distribution is the possibility to produce *personalised trustworthiness estimates with low costs.* In fact, Ostrom's [25] proposition that in robust communities the rules that describe the use of common resources match well into local conditions, supports our proposition of processing the reputation information locally.

Thirdly, the distributed approach *enhances the users' privacy.* A centralised database with large amounts of personal information is always a tempting target for misuse. A malicious peer can follow the user because the reputation information is linked to the user's identifier, which is linked to the user's terminal. But, distributing the information makes attacks against the users' personal data more difficult as the information is scattered in the network and thus improves privacy. Moreover, cryptography offers robust solutions to identity management and privacy protection in a distributed system. More discussion on privacy in distributed interaction environments can be found, e.g. in [19, 36] and on preventing the tracking of a user, e.g. in [3].

Finally, interesting topics for future research, but which are out of the scope of this paper, are the transfer of reputation between contexts and the interference between different kinds of reputation. When two different contexts are in question, especially the issue of the reputation information ownership, the rights to request and receive it, is more complicated than inside one context.

## Acknowledgements

The authors wish to thank Ronja Addams-Moring for feedback, Yki Kortesniemi for a discussion on the focus of the paper and Juhana Mattila for LaTeX support. In particular, the authors are grateful to Eero Sarkio, Timo Ylitalo and Jukka Ylitalo for recovering the working files of this paper when the hard disk crashed.

## References

[1] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *Proc. of the 10th International Conference on Information and Knowledge Management (ACM CIKM)*, pages 310–317. ACM Press, 2001.

[2] E. Adar and B. A. Huberman. Free riding on Gnutella. *First Moday*, 5(10), October 2000.

[3] J. Arkko, P. Nikander, and M. Näslund. Enhancing privacy with shared pseudo random sequences. In *Proc. of the Security Protocols Workshop (to appear)*, April 2005.

[4] R. Axelrod. *The evolution of cooperation.* New York: Basic Books, 1984.

[5] H. Balakrishnan, M. F. Kaashoek, D. Karger, R. Morris, and I. Stoica. Looking up data in P2P systems. *Communications of the ACM*, 46(2), February 2003.

[6] R. Conte and M. Paolucci. Social factors of unfair ratings in reputation reporting systems. In *Proc. of the IEEE/WIC International Conference on Web Intelligence (WI'03)*, pages 316–322, October 2003.

[7] C. Dellarocas. Immunizing online reputation systems against unfair ratings and discriminatory behaviour. In *Proc. of the 2nd ACM Conference on Electronic Commerce (EC)*, pages 150–157, October 2000.

[8] J. R. Douceur. The sybil attack. In *Proc. of the International Workshop on peer-to-peer systems*, 2002.

[9] EU data protection directive 95/46/EC of the European Parliament and of the Council, 24 October 1995.

[10] E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(1):173–199, 2001.

[11] M. Godwin. Nine principles for making virtual communities work. *Wired 2.06*, pages 72–73, June 1994.

[12] G. Hardin. The tragedy of the commons. *Science. New Series*, 162(3859):1243–1248, 13 December 1968.

[13] S. Holtmanns, S. Lakshmeshwar, S. Moloney, and K. Sarkio. Securing reputation systems using mobile trust. Submitted for publication, 2005.

[14] D. E. Houser and J. Wooders. Reputation in internet auctions: Theory and evidence from eBay. University of Arizona, February 2000.

[15] C. Jensen, J. Davis, and S. Farnham. Finding others online: Reputation systems for social online spaces. *CHI Letters*, 54(1), April 2002.

[16] P. Kollock. Design principles for online communities. In *Harward Conference on the Internet and Society.* (Also published in PC Update 15(5):58-60, June 1998), 1996.

[17] P. Kollock. The production of trust in online markets. *Advances in Group Processes*, 16, 1999.

[18] P. Kollock and M. Smith. Managing the virtual commons: Cooperation and conflict in computer communities. In Susan Herring, editor, *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*, pages 109–128. Amsterdam: John Benjamins, 1996.

[19] L. Korba. Privacy in distributed electonic commerce. In *Proc. of the 35th Annual Hawaii International Conference on System Scineces (HICSS-35)*, January 2002.

[20] K. Lai, M. Friedman, I. Stoica, and J. Chuang. Incentives for cooperation in peer-to-peer networks. In *Workshop on Economics of Peer-to-Peer Systems*, 2003.

[21] J. Liu and V. Issarny. Enhanced reputation mechanism for mobile ad hoc networks. In *Proc. of the Second International Conference on Trust Management (iTrust*, pages 48–62, 2004.

[22] T. Mahler and T. Olsen. Reputation systems and data protection law. In P. Cunningham and M. Cunningham (eds.), eAdoption and The Knowledge Economy: Issues, Applications, Case studies, part 1, ISBN 1 58603 470 7, IOS Press, Amsterdam 2004, pp. 180-187. Originally in Proc. of eChallenge, October 2004.

[23] S. Marti and H. Garcia-Molina. Taxonomy of trust: Categorising P2P reputation systems. In *COMNET Special Issue on Trust and Reputation in Peer-to-Peer Systems, Preprint submitted to Elsevier Science*, 16 March 2005.

[24] L. Mui, M. Mohtashemi, and A. Habelstadt. A computational model of trust and reputation. In *Proc. of the 35th Annual Hawaii International Conference on System Scineces (HICSS-35)*, January 2002.

[25] E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action.* New York: Cambridge University Press, 1990.

[26] C. E. Porter. A typology of virtual communities: a multi-disciplinary foundation for future research. *Journal of Computer-Mediated Communication (JCMC), Article 3*, 10(1), November 2004.

[27] J. Pouwelse, M. van Slobbe, J. Wang, M. J. T. Reinders, and H. Sips. P2P-based PVR recommendation using friends, taste buddies and super peers. In *Beyond Personalization, Workshop on the Next Stage of Recommender Systems Research*, 9 January 2005.

[28] J. A. Pouwelse, P. Garbacki, D. H. J. Epema, and H. J. Sips. The bittorrent P2P file-sharing system: measurements and analysis. In *Proc. of the 4th International Workshop on Peer-to-Peer Systems (IPTPS'05)*, February 2005.

[29] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, December 2000.

[30] J.-M. Seigneur, A. Gray, and C. D. Jensen. Trust transfer: Encouraging self-recommendations without sybil attack. In *Proc. of the third International Conference on Trust Management (iTrust)*. Springer-Verlag, LNCS 3477, May 2005.

[31] A. A. Selcuk, E. Uzun, and M. R. Pariente. A reputation-based trust management system for P2P networks. In *Proc. of the fourth IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid04)*, Chicago, Illinois, USA, April 2004.

[32] The Third Generation Partnership Project (3GPP). At http://www.3gpp.org/ (Visited 16 November 2005).

[33] The Third Generation Partnership Project (3GPP), Generic Authentication Architecture (GAA); support for subscriber certificates; ts 33.221, v6.2.0, November 2004.

[34] L. Xiong and L. Liu. A reputation based trust model for peer-to-peer eCommerce communities. In *the IEEE International Conference on E-Commerce (CEC)*, June 2003.

[35] K. Ylitalo and S. Holtmanns. Tailored trustworthiness estimations in peer-to-peer networks. In *Proc. of the IEEE/Create-NET SECURECOMM workshops*, September 2005.

[36] K. Ylitalo and Y. Kortesniemi. Privacy in decentralised reputation management. In *Proc. of the IEEE/Create-NET SECURECOMM workshops*, September 2005.

[37] G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *Proc. of the 32nd Hawaii International Conference on Systems Sciences (HICSS-32)*, Cambridge, UK, January 1999.