PRIVACY PRESERVING ONLINE REPUTATION SYSTEMS

Marco Voss
Darmstadt University of Technology, IT Transfer Office (ITO), Wilhelminenstr. 7, 64283 Darmstadt, Germany
voss@ito.tu-darmstadt.de

Abstract Reputation systems evolve as a mechanism to build trust in dynamic electronic societies. However, they are also a danger to privacy because they monitor a user's behavior. At the same time reputation systems offer the possibility to limit the information a user has to give away during a transaction to ensure accountablity. Privacy preserving reputation systems solve the conflict between anonymity and accountability. This paper examines privacy problems of current reputation systems and classifies them with respect to the location of stored information. Requirements for reputation systems that provide privacy protection are derived from this analysis. As result a new privacy preserving online reputation system is presented that uses locally stored coins to represent reputation information.

Keywords: Anonymity, privacy, accountability, trust, reputation

1. INTRODUCTION

Reputation systems [24] evolve as a mechanism to build trust in dynamic electronic societies. Especially peer-to-peer systems, anonymous remailer networks, online marketplaces, auction sites and web logs rely more and more on reputation to improve their performance and security or to eliminate unwanted behavior. Classical mechanisms to build trust fail in these scenarios.

Reputation systems monitor the behavior of an entity and provide this information upon request. Then this information can be used to make a decision about the trustworthiness of an entity.

By storing data about former transactions of an entity, reputation systems can represent a danger to a user's privacy. However, reputation can be also a mechanism for privacy, because the amount of information that must be disclosed during a transaction can be limited.

This paper investigates the privacy aspects of reputation systems. These systems are classified with respect to the location of stored information. We present requirements for privacy preserving reputation systems and make a proposal for such a system.

Recommender systems and collaborative filtering are related topics, but their focus is more personalization and the filtering of a list of alternatives than to deal with the behavior of entities. Some research has been undertaken in the field of privacy in recommender systems [23]. Similar problems exist here to prevent faking of ratings.

Throughout this paper the following terminology is used: The subject of reputation is an *entity*. The terms entity, peer, user or node are used synonymously. An entity assumes an *identity* for transacting with others. Consequently, an entity can have more than one identity. *Ratings* or *votes* are the committed opinion of one entity about another and are also used synonymously. *Reputation* is the result of collected ratings after consolidation.

The layout of this paper is as follows: The next chapter summarizes some facts about reputation and reputation systems with emphasis on privacy aspects. Then eBay's feedback system is examined as an example of a well known reputation system in use with a lack of privacy protection. After that a classification of reputation systems is presented. This is followed by listing the identified requirements of a privacy preserving reputation system. Finally a sketch of two proposals of such systems is drawn and the paper is concluded with a summary of our results.

2. REPUTATION

More and more interactions involving humans or companies are handled over the Internet. Its openness allows everyone to participate by opening a web site and offering services. The biggest problems appear if one leaves a closed and established group of users. Then there is insufficient information to decide about the trustworthiness of a unknown peer. There are also no regulations to guarantee proper behavior.

Trust and reputation are objects of research in many disciplines from sociology to economics to computer science. The interpretation of these terms differs from author to author and a common definition is still missing.

Mui et al. [21] give a short overview of different notions of reputation and trust from various disciplines. They also derive a computational model of trust and reputation. This discussion is not repeated here, but a more simple and intuitive definition is used:

> Reputation is the collected information about one entity's former behavior as experienced by others. Trust is a decision made on the basis of reputation.

Reputation systems have two effects: they predict the future based on past behavior and they simultaneously influence the future by giving people the incentive to behave well.

Reputation is tightly coupled with an identity. It has no meaning without this identity. From this point of view reputation is not a tradeable asset. After a restaurant is purchased, the name is rarely changed in order to profit from its former reputation. However, there is no guarantee that the new owner will provide the same quality of service.

If a reputation system is implemented by recording and evaluating the outcomes of an entity's interactions, then this means that reputation is sensitive personal information. One may not want to disclose all information about former transactions, because it may also say something about one's preferences and circumstances.

Most simple reputation systems require that an entity is identifiable at least by a pseudonym (entity B in figure 1). This may require a former registration. Then reputation information for this pseudonym is gathered. This information is processed and consolidated. Finally, it is made available upon request.

This process is represented in the figure by the three components collector, processor and emitter. Another entity, A, uses this information together with its own experience to make a trust decision. If A believes that B is trustworthy enough, it enters the interaction. After completion, A rates B according to the outcome of the interaction. This rating is fed back into the system. With further interactions the reputation data of B is steadily updated.

If reputation is not represented as a single value it may also contain detailed information about past transactions. At least for the lifetime of the pseudonym, an entity's transactions can be linked and the development of its reputation can be monitored.

It is vitally important for an entity not to lose reputation because of false accusation. Many people in the USA are concerned about identity theft. The Federal Trade Commission has already setup a web site with information about how to deal with this crime.

The opposite problem is to prevent lending of identities. If the authentication data is not tightly bound to an entity (using biometrics for instance), it is easy to hand on this data to someone else. Some systems try to prevent this by hard punishments or by including some personal information (like a credit card number) in the authentication data. An-

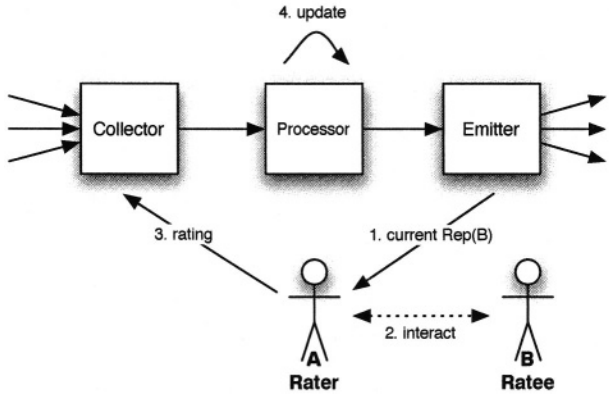

Figure 1. Basic Architecture of Reputation Systems

other issue is to prevent switching of identities without making it too hard to join a reputation system for a newcomer.

Dellarocas [6] analyzes common attacks to compromise reputation systems. He identifies four mayor problems: unfairly high ratings ("ballot stuffing"), unfairly low ratings ("bad-mouthing") and negative or positive discrimination. One of his results is that controlled anonymity can avoid unfairly low ratings and negative discrimination.

2.1 Context and Interdomain Sharing

Reputation is context and domain dependent. It is not clear how a reputation for one context can be transfered to another. A rating earned for having a good understanding of numerical mathematics says nothing about being an honest buyer. This also means that reputation cannot really be captured in a single number. It is more a vector where the elements represent a value for a specific context.

Existing reputation systems are closed and service specific. One cannot transfer a reputation gained in one system to another (e.g. from eBay to Amazon). It would be nice to be able to transfer reputation from one system or domain to another if the context is compatible. This may also involve some normalization of reputation values if the systems have different measures. Single sign-on systems like Microsoft's Passport or the Liberty Alliance project would allow to provide such an interdomain sharing. However, users have strong privacy concerns against such systems.

2.2 Privacy

In this section we will first discuss a general definition of privacy. Then relevant aspects for reputation systems will be stated.

Privacy is a widely and ambiguously used term that is often confused with secrecy, but it means more than only encrypted communication and secure data storage. Communication privacy means that a third party is not able to learn anything about the content of a message exchange between two entities or even that such a communication has happened between them.

Information privacy means that only authorized entities can access and modify information. However, it also means that the owner of personal information should have control over it and that only necessary information has to be disclosed. Goldberg [14] states:

> "Privacy refers to the ability of the individual to control the distribution of information about himself."

Another aspect of privacy is unlinkability. As it is sometimes not possible to avoid disclosure of personal information during transactions with others, the disclosed information should be protected against profiling. Alamäki et al. [1] give a linking-based definition of privacy:

> "A system is privacy-enabled if no party is able to or has the right to link data about a user to its true identity or its other data, unless the user has explicitly allowed it."

Control and unlinkability are two important aspects of privacy that have an impact on reputation systems. The former includes both control of access to reputation information and when it is updated by new votes. Anonymity is a measure for privacy of identity and unlinkability [22]. However, full anonymity reduces accountability if no damage to one's reputation has to be feared.

To whom belongs one's reputation, the operator of the reputation system or the person concerned? Fulda [12] sees reputation as property:

> "Actually, reputation is based on our abilities, capacities, and even physiognomy as modified over the years by every action we take, every behavior we display. Thus, like personal property, reputation is formed by taking natural resources and mixing our labor with it."

In this sense defamation is damaging personal property. Only the interpretation of this record, the opinion other people have about us, belongs to them. Although this position may be extreme and conflicts with freedom of speech, we believe that reputation is sensitive personal information that must be protected and controlled by the user concerned.

This is even more true for online reputation because the amount of information that can be collected and processed is bigger in an online world. If there is only a small possibility to make money out of this, it

is guaranteed that it will be abused by marketing (taking spam mails as an example).

Centralized reputation systems like eBay's feedback system [16] collect and store a lot of data about recent transactions. In the case of eBay this data is public and allows the creation of a detailed personal profile about a pseudonym. It is not really difficult to get the email address of an eBay pseudonym to link this information to some identity.

By virtue of containing so much information, reputation systems are a danger to a user's privacy. At worst the collected information can be misused to build a profile of the user's tendencies. Cheating and discrimination can also be a problem if the user lacks control of her reputation.

On the other hand, reputation systems can be a way to provide privacy protection [8]. If one can prove to a partner in a transaction that one has a good reputation in the relevant context there is no need to ask for additional data except the data necessary for fulfilling the transaction.

Chaum [4] already had the vision that anonymous credential systems can promote the protection of privacy. Instead of authentication based access control, one proves authorization by showing possession of an anonymous credential issued by an organization. Brands [3] advances this idea by introducing techniques for a privacy preserving PKI. Any boolean formula can be proved about attributes in a certificate.

What is missing so far is a mechanism to compensate for the decreasing accountablity in systems with strong anonymity. Following the ideas above, we see a privacy protecting reputation system as an extension that allows secure updates to anonymous credentials. This can provide an balance between anonymity and accountability. A digital credential or certificate is something that changes rarely and must be reissued when changed. In contrast, reputation rises and falls according to the owners behavior. If proving and updating a reputation can be done anonymously, it is a perfect mechanism to guarantee a user's privacy. Only necessary information must be disclosed and successive transactions of a single user cannot be linked to each other.

3. CLASSIFICATION OF REPUTATION SYSTEMS

Mui et al. [20] have introduced a typology for reputation systems that concentrates on how the reputation is collected. They subclassify individual reputation into direct and indirect reputation. Direct reputation is derived from direct experiences. Indirect reputation is derived from second-hand information.

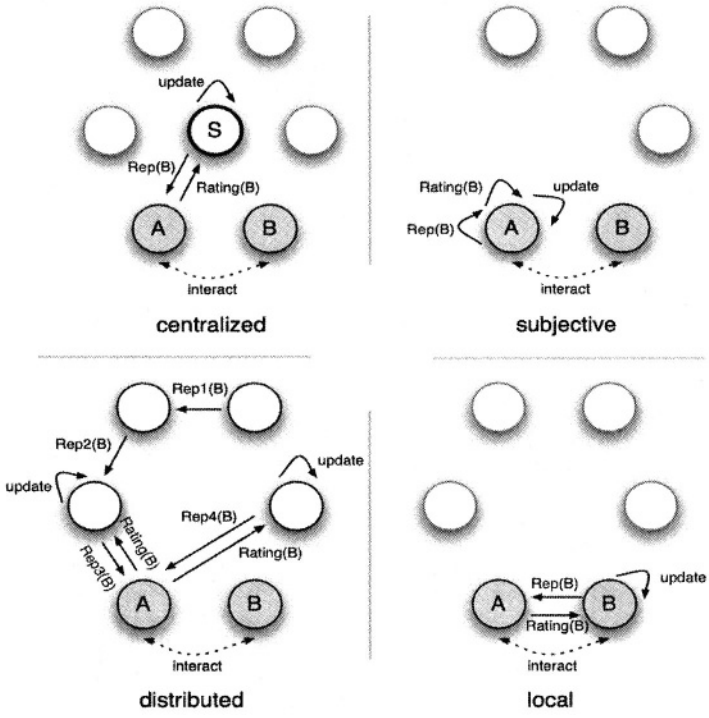

Figure 2. Categorization of Reputation Systems

Their work allows to classify a lot of reputation systems, but it says nothing about the privacy properties. We propose a new categorization according to the place where reputation information is stored. This makes sense if one looks at the privacy aspects resulting from who has control over this information. We found four different classes of storage:

- centralized,
- subjective with no propagation,
- distributed with propagation and
- local storage.

In the following we will describe each class in more detail and state an example system for illustration. Figure 2 gives an overview of this categories.

3.1 Centralized

The centralized approach is very common for reputation systems used for Internet services. Examples are slashdot, web logs, web forums, and auction sites (eBay, Amazon). In the simple case one single central server is responsible for storing and managing all users' reputation data. These systems are easy to implement and follow the need of service providers. They allow users to remain pseudonymous. Anonymity is possible between two users but not between a user and the system.

In scenarios where no central servers exists, like peer-to-peer systems, the role of a central reputation server can be distributed over a certain number of nodes. We still call this a centralized reputation system, because from a privacy perspective it makes no real difference. The propagation of reputation information is restricted to these nodes for synchronization. In [9] each peer node has a number of k designated supervisors that store an updated reputation value after each transaction. One of these act as a main supervising node whereas the others function as a backup. The authors describe how this can be mapped to an overlay peer-to-peer network with a ring structure like Chord [27]: the supervisor nodes can be the immediate successors in the ring topology. [28] describes the same idea but their KARMA system is more like a credit system.

The major drawback is that the central authority has absolute control over the collected data. It can profile the collected data and disclose to a third party. It can swindle or discriminate against users. Even if it is fair, it represents a single point of failure and attack. The latter can be also true in a distributed case. In [9] it is enough to take over the main supervising node to get control over the user's reputation data.

3.2 Subjective

In a subjective model every node stores its own view of the world. For every peer it has interacted with, it creates a record. This record is updated after every transaction according to its result. Nodes do not propagate their opinions about other nodes in order to keep communication low. This also means that every node has a restricted view. Other nodes have to be identifiable over a longer time period in order to track their behavior.

Subjective reputation systems without propagation are good for nodes that interact with a stable number of other known nodes. The number of newly appearing nodes or nodes that change their identification must not be too high. No trusted authority is needed.

GNUnet [13] is a peer-to-peer framework for anonymous distributed file-sharing that uses a subjective reputation system. There are only two operations in GNUnet's protocol which are relevant: queries and replies. A query consumes network bandwidth and is seen as a loss in reputation, whereas a reply is a way to earn reputation. It calls this an economic model, because nodes can attach a priority to a query which is used as the amount of reputation value which is earned or lost. Each node only evaluates the behavior of other nodes it communicates with.

3.3 Distributed

A distributed reputation system extends the former by propagation of reputation information. It allows a more global view and there is information accessible even for nodes that are not members of the direct neighborhood of the request. Nodes combine their own experiences with the observations of other nodes. Other nodes have to be identifiable over a longer time period in order to track their behavior. More elaborate systems use a weighting of votes according to the voter's reputation.

P2PRep [5] extends the Gnutella protocol [19] by a distributed reputation system. Before a user decides from where to download a resource, it can enquire about the reputation of the offers by polling its peers. This is done by broadcasting a message into the p2p network asking for opinions about a specified peer.

EigenTrust [10] is the most extreme form of a distributed reputation system. It integrates the idea of transitivity. They show that global reputation values correspond then to the left principal eigenvector of a matrix of normalized local reputation values.

3.4 Local Storage

In a reputation system with local storage, every entity saves its own reputation data. It can prove to others that it has received this reputation value. After a transaction the value can be updated by another entity in a secure manner. That means that the receiver of a rating has to accept it even if it is negative. She must not be able to fake votes or to use outdated reputation data and drop unpleasant/negative votes. Additionally, it must be difficult for a group of entities to mutually improve their reputation values.

In [18],[11]and [15] schemes that follow this category are presented.

Agent. Gupta et al. introduce a reputation computation agent (RCA) in a Gnutella-like P2P network [15]. This RCA acts as a trusted third party for issuing an updated reputation of a peer. The RCA is not involved in peer-to-peer interaction. Each peer stores its current repu-

tation signed by the RCA. Requesters create signed receipts for peers serving content that can be exchanged for reputation credit with the RCA. To prevent misuse a RCA keeps a list of all processed receipts. This means that a RCA learns about all the recent transactions of a peer.

Portal. In [11] each peer stores context dependent reputation lists. Each entry in one of these lists consists of a signed questionnaire, the IDs of the involved peers and a signed portal ticket. During negotiation between two peers, a requester checks the received reputation list by verifying the signatures and identities. The completeness of the list is confirmed by asking the portal for a counter of processed transactions. This also guarantees that a peer cannot drop negative votes.

After interaction peers must rate each other. They receive an empty questionnaire from the portal to be filled out. After that they exchange the signed questionnaires and challenges contained in the portal tickets. These challenges are send back to the portal. If the portal has received both values it is assured that both questionnaires have been transmitted. The portal updates the transaction counters.

The portal is involved during the negotiation and rating phase. It only stores counters for processed transactions, but learns also the context and the IDs of involved peers. However, it does not know anything about the ratings.

RCert. In [18] the main concept is a certificate called RCert that contains a signed list of former ratings. A RCert consists of a header and units. The header contains information about the owner. A RCert unit contains a timestamp, rating, id of rater and signature. Updates are appended to the end of the RCert and the whole new content signed by the rater. The integrity (but not completeness) of a RCert can be checked by verifying the signature of the last rater. In the simple version the dropping of an entry can only be discovered by analyzing the transaction frequency.

An extended protocol RCertPX includes the previous raters to guarantee the completeness of the presented reputation certificate. Therefore the current rater contacts the previous rater and asks for a last timestamp. When commiting a new rating the rater must contact the previous rater and request that the last timestamp be revoked.

Except for the PKI used for authentication, no central instance is needed. The whole transaction history is contained in the reputation certificate. This also means that every transaction partner learns all about the former transactions.

Currency. Another approach to implementation is to treat reputation as a digital currency. This can guarantee more privacy if the coins

are made in a way such that one sees only their value. Dingledine et al. have [7–8] already stated some open questions related to this scenario. Especially, where this currency comes from and who controls its worth are unsolved topics. Currency is a zero-sum game, but reputation is not. Currency is transferable, whereas reputation is bound to an identity.

3.5 Conclusion

None of the described reputation systems has privacy as its main focus. Especially, centralized and distributed systems hold a lot of sensitive personal data accessible for every participant. Only systems with local storage can provide the owner full control over reputation information. But also the three presented reputation systems have some severe privacy shortnesses.

4. REQUIREMENTS OF A PRIVACY PROTECTING REPUTATION SYSTEM

After having discussed important properties of reputation and the relevance of privacy, we now summarize the requirements of a reputation system that provides privacy protection.

General requirements:

- It must provide information that allows users to distinguish between trustworthy and untrustworthy peers.
- It should encourage entities to behave trustworthily.
- An entity must not be able to fake a reputation value.
- Negative ratings (not only positive ones) should be supported.
- An entity must not be able to get rid of a negative reputation.
- An entity should have not an interest to switch its identity to cover misbehavior. Switching should not give any advantage.
- A group of colluding entities should not be able to give each other a high reputation value.
- It should not be possible to defame someone without proof.

Privacy releated requirements:

- The amount of additional data contained in the reputation information should be as limited as possible.

- An entity should have control over its reputation information. This includes access control but also control about when this information is updated.
- The identity of a rater should be protected. If possible a rater should be anonymous.
- Also the identity of a ratee should be kept secret.
- Other parties should learn as less information as possible about the transaction between rater and ratee.

5. A COIN-BASED PRIVACY PROTECTING REPUTATION SYSTEM

We are currently working on two approaches to privacy protecting reputation systems. Both favor the local storage of reputation information. The first approach uses local storage with coin-like reputation and a trusted third party. The trusted third party is used to guarantee a correct update of the reputation information even when a rating has been negative. Because this is still ongoing work, we will only give a sketch of the main ideas. The second one relies on local storage with trusted hardware (smart card or TPM based) and will not be presented here.

5.1 Overview

Reputation is represented as coins. The more coins an entity has the more peers have been satisfied by its performance during former transactions. The coins are issued by a trusted third party for positive ratings and are personalized for the receiver. Consequently, coins do not represent a currency: they cannot be traded or cashed.

At the beginning of a transaction an entity gives coins to its peer as collateral. These coins cannot be used by the peer, but only invalidated if handed over to the trusted third party. By this negative ratings are accomplished. If the rater is satisfied by the outcome of the interaction she can give a positive rating. Thereupon the ratee receives a fresh coin from the trusted third party. To prevent ballot stuffing an entity can give another entity a positive rating only once. This is also guaranteed by the trusted third party.

Different contexts can be mapped to different kinds of coins. For instance a seller can have two kinds of coins: one for the quality of goods she sells and another for reliability and timely delivery. Or a person can have received different coins for being a credit-worthy buyer

and another kind for being a Linux expert. To simplify the description we deal only with one kind of coins in the following. The protocol can easily be extended to the case with different kinds of coins. We have implemented a first prototype of this scheme using the project JXTA [26] peer-to-peer infrastructure.

5.2 Online Protocol

This paragraph describes the steps of the protocol. Involved are three parties: the rater (A), the ratee (B) and a trusted third party (T). It is an online protocol because interaction with T is required during the transaction.

1 As a first and optional step B proves that its reputation is bigger than a value x. This is done by proving possession of a sufficient number of coins.

2 Before A agrees to start the interaction it requests a number c of coins as collateral.

3 If B agrees it gives the requested number of coins to A. These coins are not tradeable or cashable. The only possible action for A is to invalidate them.

4 A (together with T) has to verify that the presented coins are still valid and not already used in another transaction.

5 The following interaction between A and B is out of scope of the reputation system. For instance B may provide some service to A.

6 According to the outcome of this interaction A decides about a voting for B. This voting should represent A's satisfaction, but A cannot be forced to give a fair rating. Possible values are elements of {+1, 0, -1, ..., -c}.

7 A communicates this rating to T. If the rating is negative A asks T to invalidate r coins.

8 In case of a positive rating T checks whether A has given B a positive rating before. If not T creates a new coin and sends it to B.

5.3 Building Blocks

Registration Authority To participate in the reputation system an entity has to register with a central authority. By this registration the

entity receives a pseudonymous identity and a secret to prove possession of this identity. This identity can be used for authentication, but it is only required by the system to personalize coins and to prevent ballot stuffing (see below).

Coin System We require a coin system that allows to personalize coins on generation. It must also provide anonymity and unlinkability. This means that coins cannot be transfered to another identity without sharing the whole secrets of this identity, but possession of a coin can be proved without authentication. T must be able to invalidate coins.

Anonymous Communication There has already been done a lot of research in the field of anonymous communication [2]. Proposed solutions include mix networks and anonymizing proxies. Although not implemented in our scheme we believe that this solutions can be easily integrated as they use common internet protocols as interface.

Anonymous Recognizing To prevent ballot stuffing T must be able to recognize when A and B interact once again. To preserve privacy T should not be able to identify a single peer but only the tuple of A and B. Currently, we have implemented this by using another trusted party V which only function is to help T recognize a tuple A,B. A and B submit a proof of identity encrypted with the public key of V. V checks whether it has already been presented A,B and gives T the corresponding answer. This scheme provides poor privacy if V is corrupt.

Therefore, we are working on a solution that uses commitments and proofs of knowledge. A and B compute together some value and present this with a proof that this value is dependent on A's and B's certified pseudonyms.

5.4 Privacy Aspects

We will now take a closer look at the privacy aspects of our proposed system. Coins limit the amount of information disclosed: only the outcome of former transactions is revealed. Not the partners, not the subject, not the time (if coins don't have a time stamp). Showing possession of coins can be implemented anonymously and unlinkable. An entity can show the same coin several times to the same partner without giving the possibility to link these actions.

The rater only learns that B possesses a number of coins. The trusted third party only learns that someone (B) has given someone else (A) a number of coins as collateral and how many coins should be returned to B, but it does not learn anything about the identities of A and B. This

means that the gained information is only useful to generate a profile about the overall behavior in the reputation system.

6. CONCLUSION AND FUTURE WORK

We have presented a new classification of reputation systems which differentiates by where the reputation data is stored. We have also discussed privacy aspects of reputation systems and the identified classes. This was followed by a summary of important requirements for privacy protecting reputation systems. Especially centralized and distributed reputation systems have privacy problems, because sensitive data about an entity's transactions is stored with inadequate protection in these systems. Local storage of reputation provides a solution to this problem and gives the control back to the user. We have sketched the approach we are currently working on that uses locally stored coin-like reputation. Our future work will focus on finalizing the implementation of this system. A simulation based comparison with existing systems will evaluate its efficiency.

References

[1] Alamäki, T., Björksten, M., Dornbach, P., Gripenberg, C., Gyorbíró, N., Márton, G., Németh, Z., Skyttä, T., Tarkiainen, M.: *Privacy Enhancing Service Architectures,* In Proceedings of Privacy Enhancing Technologies 2002, pp. 99-109

[2] Berthold, O., Federrath, H. and Köpsell, S., *Web MIXes: A System for Anonymous and Unobservable Internet Access,* Lecture Notes in Computer Science 2009, pp. 115, 2001

[3] Brands, S.: *Rethinking Public Key Infrastructures and Digital Certificates,* MIT Press, 2000

[4] Chaum D.: *Security Without Identification: Transaction Systems to Make Big Brother Obsolete,* Communications of the ACM, vol. 28 no. 10, October 1985 pp. 1030-1044

[5] Cornelli, F., Damiani, E., Vimercati, S. D. C. D. , Paraboschi, S. and Samarati, S.: *Choosing Reputable Servents in a P2P Network,* In Proceedings of the 11th World Wide Web Conference, Hawaii, USA, May 2002

[6] Dellarocas, C.: *Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior,* ACM conference on Electronic Commerce, 2000

[7] Dingledine, R., Mathewson, N., Syverson, P.: *Reputation in P2P Anonymity Systems,* Workshop on economics of p2p systems, June 2003

[8] Dingledine, R, Mathewson, N., Syverson, P.: *Reputation in Privacy Enhancing Technologies.,* Computers, Freedom, and Privacy, Apr 2002.

[9] Dutta, D., Goel, A., Govindan R., Zhang H.: *The Design of A Distributed Rating Scheme for Peer-to-peer Systems,* Workshop on Economics of Peer-to-Peer Systems, 2003.

[10] Kamvar, S., Schlosser, M., Garcia-Molina, H.: *The EigenTrust Algorithm for Reputation Management in P2P Networks,* 2003

[11] Fahrenhols, D., Lamersdorf, W. *Transactional Security for a Distributed Reputation Management Sytem,* EC-Web 2002, LNCS 2455, pp. 214-223, 2002

[12] Fulda, J.: *Reputation as Property, and its relation to privacy,* Computers and Society, pp. 27-28, March 2001

[13] Grothoff, C.: *An Excess-Based Economic Model for Resource Allocation in Peer-to-Peer Networks,* WI 32003

[14] Goldberg, I.: *A Pseudonymous Communications Infrastructure for the Internet,* 2000

[15] Gupta, M., Judge, P., Ammar, M.: *A Reputation System for Peer-to-Peer Networks,* NOSSDAV'03, June 2003

[16] Houser, D., Wooders, J.: *Reputation in Auctions: Theory, and Evidence from eBay,* 2000

[17] Kinateder, M., Pearson, S.: *A Privacy-Enhanced Peer-to-Peer Reputation System* EC-Web 2003, LNCS 2738, pp. 206-216, 2003

[18] Liau, C. Y., Zhou, X. Bressan S., Tan, K.: *Efficient Distributed Reputation Scheme for Peer-to-Peer Systems,* HSI 2003, LNCS 2713, pp. 54-63, 2003

[19] NN.: *The Gnutella Protocol Specification v0.4*

[20] Mui, L., Halberstadt, A., Mohtashemi, M.: *Notions of Reputation in Multi-Agents Systems: A Review,* In Proc. of Int'l Conf. on Autonomous Agents and Multi-Agents Systems (AAMAS-02), pp. 280-287

[21] Mui, L., Halberstadt, A., Mohtashemi, M.: *A Computational Model of Trust and Reputation,* 35th Hawaii International Conference on System Science (HICSS), 2002

[22] Pfitzmann, A., Köhntopp, M.: *Anonymity, Unobservability, and Pseudonymity - A Proposal for Terminology,* In Designing Privacy Enhancing Technologies – International Workshop on Design Issues in Anonymity and Unobservability 2000, LNCS 2009, pages 1-9. Springer-Verlag, 2001

[23] Ramakrishnan, N., Keller, B., Mirza, B., Grama, A., and Karypis, G.: *Privacy risks in recommender systems,* IEEE Internet Computing, pages 54 62, November 2001

[24] Resnick, P., Zeckhauser, R.: *Reputation Systems,* Communications of the ACM 43, pp. 45-48, 2000

[25] Resnick, P., Zeckhauser, R.: *Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System,* In The Economics of the Internet and E-Commerce, volume 11 of Advances in Applied Microeconomics. Elsevier Science, 2002.

[26] Sun Microsystems, Inc.: *Project JXTA: An Open, Innovative Collaboration,* 2001

[27] Stoica, I., Morris, R., Liben-Nowell, D., Karger, D., Kaashoek, M., Dabek, F., Balakrishnan, H.: *Chord: A Scalable Peer-to-peer Lookup Protocol for Internet Applications,* To Appear in IEEE/ACM Transactions on Networking

[28] Vishnumurthy, V., Chandrakumar, S., Sirer E. *KARMA : A Secure Economic framework for Peer-to-Peer Resource Sharing,* Workshop on Economics of Peer-to-Peer Systems, 2003