Design Options for Privacy-Respecting Reputation Systems within Centralised Internet Communities

Sandra Steinbrecher

Technische Universität Dresden, Fakultät Informatik
D-01062 Dresden, Germany
steinbrecher@acm.org

Abstract. Reputation systems play an important role in Internet communities like eBay. They allow members of the community to estimate other members' behaviour before an interaction. Unfortunately the design of current reputation systems allows to generate user profiles including all contexts the user has been involved in. A more privacy-enhancing design of reputation systems is needed while keeping the trust provided to the members by the use of reputations. We will present design options for such a system and analyse the privacy it provides with common information-theoretic models. The analysis of our reputation system also allows to analyse similar aspects of privacy in other systems, especially privacy-enhancing identity management.

1 Introduction

With the growth of the Internet more and more people spend a lot of their spare time in so-called Internet communities instead with friends or relatives at their domicile. Most of these virtual friends they have neither met in the past nor will meet them in the future. The spectrum of Internet communities reaches from mailing lists, newsgroups and discussion forums to role-playing and electronic marketplaces. Most of these communities are implemented in a centralised way. There is a community provider which offers the technical system where the community members meet. The members typically want to specify the security requirements - especially confidentiality and integrity of information and actions - such a system should fulfill. But even more than on the systems they have the requirements on the other (usually initially unknown) members they are interacting with.

Example 1 (Self help forums). Someone seeking advice might get technical integrity of other users' answers, but if they give him false advice, technical integrity is meaningless for his problem. Someone who wants his question within the community to be kept confidential from others than legitimated readers might get this guaranteed by the system, but this is insufficient if legitimate readers distribute this information manually.

The measures to reach confidentiality and integrity within a technical system are well-known: By adding digital signatures integrity of digital information or actions can be guaranteed. This needs appropriate public-key infrastructures. But by signatures people only get evidences for others' misbehaviour. Every dispute between individuals has

Author(s) [insert Last name, First-name initial(s)], 2006, in IFIP International Federation for Information Processing, Volume 201, Security and Privacy in Dynamic Environments, eds. Fischer-Hübner, S., Rannenberg, K., Yngström, L., Lindskog, S., (Boston: Springer), pp. [insert page numbers].

to be solved outside the Internet community in a legal process based on national and international law. Confidentiality of digital information can be reached by using encryption. Confidentiality of actions needs appropriate anonymity systems e.g., anonymity on the IP level can be reached by Web mixes [1] or Tor [9], both more or less based on Chaum's Mixes [3].

But beneath these technical measures members of Internet communities need trust in each other that other members do not make technical measures (if applied or even more than not applied) obsolete by their misbehaviour.

When becoming a member of an Internet community an individual develops a new partial digital (or virtual) identity within this community. He starts with a new pseudonym and has to gain a reputation for this pseudonym within the community depending on his (mis)behaviour and its valuation by other members, especially how trustworthy he has been.

Reputation systems can collect the experiences members made with interactors in a technically efficient way. These experiences may help other members to estimate the future behaviour of interactors they have no personal experience with. But it does not prevent any member from making bad experiences with interactors because e.g., reputation usually is context-dependent and subjective. Although 'social attacks' (e.g., members may lie about others' behaviour [6] or members may suddenly change their behaviour) are possible, a usually large number of reputations and an honest majority of members will hopefully reach that dissatisfied members are the exception. For the case that two members are dissatisfied with an interaction, technical measures (like e.g. digital signatures under agreements made) could still give them the possibility to reach legal enforceability of the other's behaviour. So reputation systems do not make other technical security measures obsolete, but hopefully reduce the cases where expensive legal enforceability might become necessary. Note the social and legal aspect of Internet communities cannot be discussed in further detail here.

A very-popular example of Internet communities are marketplace communities whose security was studied in [20].

Example 2 (Marketplace communities). The members of a marketplace community are allowed to sell and buy arbitrary items within the community. One of the greatest providers is eBay (`http://www.ebay.com/`). After an item within the community has been sold the respecting seller and buyer have to exchange item purchased and money. This is usually done by bank transfer and conventional mail. Many of these exchanges are successful, but unfortunately some are not. Although the money lost might be little and fraud seems to occur only rarely (for instance an eBay representative indicates [25] that 'Fewer than 0.01 percent of all listings on eBay result in a confirmed case of fraud'), the nuisance perceived by the customer is high, and can hamper the further development of marketplace communities. Reputation systems were introduced to most providers' service to collect the experience sellers and buyers made. They are used as a cheap alternative and unfortunately not as an additional option to expensive public-key measures and infrastructures. After every exchange the respective members may give comments or/and marks to each other that are added to the members' public reputation (usually together with the annotator and the exchange considered as context information).

Reputation systems as data bases for community members' experiences with each other should be protected by means of technical data protection to ensure users' right of informational self-determination [15]. Beneath the members' legitimate interest in informing themselves about future interactors, numerous data collectors will be desirous to get access to such large data bases which contain information who interacted at which time with whom in which context. Unfortunately reputation systems currently in use in electronic marketplace communities [14] allow to generate interest and behaviour profiles of pseudonyms (e.g. time and frequency of participation, valuation of and interest in specific items). If the pseudonym becomes related to a real name, as it typically does for trading partners, the profile becomes related to this real name as well.

In [6] an electronic marketplace community with an alternative unidirectional reputation system (where buyers rate sellers) is outlined: Only the provider of the marketplace is able to link identities of members to their pseudonyms in transactions and reputations received in them, but it only publishes estimated reputations of members not their true identities. But this linkability and control of reputations by the provider is not desirable in all scenarios (e.g. if the provider might be corrupted by an attacker). In this paper we try to outline more privacy-respecting but also centralised alternatives of a reputation system that allows bidirectional reputation.

But before going into details of our reputation system we need to give an overview of common models, terms, and measurement methods common in privacy-enhancing technologies in section 2.

Based on this in section 3 we present a model of a centralised Internet community with design options for a privacy-respecting reputation system.

The analysis of the anonymity/unlinkability provided by our privacy-enhancing reputation system follows in section 4. Additionally in contrast to previous approaches our system tries to keep the level of trust provided to the members by the use of reputations at the same level than in not privacy-respecting approaches.

2 Privacy-Enhancing Technologies and Their Measurement

Privacy-enhancing technologies on the IP and application layer try to minimize the data necessary for applications, especially they try to provide confidentiality of circumstances of an action. In this section we give a short overview of privacy-enhancing technologies, privacy-enhancing and -respecting application design and the measurement of anonymity and unlinkability these technologies or applications reach.

2.1 Anonymity

Anonymity of a subject means it is 'not identifiable within a set of subjects, the anonymity set.' [13]. In typical examples a subject's anonymity usually is related to an action. Then the anonymity set is formed by all actors who might have executed the action. For Internet communities the anonymity set is a subset of the community. But the size of the anonymity set is not sufficient to measure a user's anonymity, the 'anonymity is the stronger, the larger the respective anonymity set is and the more evenly distributed' the action's execution 'of the subjects within that set is.', I.e., not only the size of the respective anonymity set determines the anonymity of a certain subject but also how likely a subject of the anonymity set might have executed the action.

There exist several approaches how to describe and measure anonymity. We skip approaches based on formal languages and logics (e.g., [22, 17, 16, 19, 11] because they do not include probabilism that we assume to be necessary for an analysis of anonymity in Internet Communities and only present approaches based on information theory that allow to assign probability distributions to anonymity sets. Then the optimal situation (where every subject in the anonymity set might have executed the action with the same probability) can be compared with the situation where the subjects might be assigned different probabilities because of additional information. In the following we present the information theoretic model from [7, 18] shortly and in a slightly extended version [21] where anonymity regarding arbitrary actions is considered:

Let A be a non-empty set of actions of arbitrary size and $U = \{u_1, \dots, u_n\}$ be a set of subjects (the anonymity set regarding a specific action $a \in A$) of size n. Given an action a every subject $u_i \in U$ with $i \in \{1, \dots, n\}$ executes a with the a priori probability $\frac{1}{n}$ and with the a posteriori probability $p_i = P_a(X = u_i) > 0$ (with X random variable) for a possible attacker's view on the system. Naturally $\sum_{i=1}^{n} p_i = 1$. Then there exist two possibilities to describe the global anonymity a system provides for action a:

- Serjantov and Danezis [18] define the a posteriori entropy to be the **effective size of the anonymity probability distribution** $(p_1, \dots, p_n)$:

$$H(X) = -\sum_{i=1}^{n} p_i \log_2(p_i). \tag{1}$$

- Diaz et al. [7] use the normalised information of what the attacker has learnt (max $(H(X)) - H(X))$ with $\max(H(X)) = \log_2(n)$ and define the **global degree of anonymity** a system provides as

$$d(U) := 1 - \frac{max(H(X)) - H(X)}{max(H(X))} = \frac{H(X)}{max(H(X))}. \tag{2}$$

  The normalisation has the effect that only the probability distribution not the size of the anonymity set is measured in the degree of anonymity. Both degree and size of an anonymity set have to be given to describe the anonymity a system provides.

Users typically are not only interested in the global anonymity a system provides in the average case but in the local anonymity a specific individual $u \in U$ might reach in the worst case. A similar anonymity measure for this case can be defined, but because we want to study the global anonymity in an Internet community we omit this approach here and refer to [23] for details.

2.2 Unlinkability

In an extension to anonymity of a person the unlinkability of his actions can be measured in a similar way. Unlinkability of two items (e.g., actions) within a system means that 'within the system (comprising these and possibly other items), from the attackers perspective, these items of interest are no more and no less related after his observation than they are related concerning his a-priori knowledge '[13].

Unlinkability of items within one set (e.g. actions that might be executed by the same user) can be measured as following [21]: Let $A = \{a_1, \dots, a_n\}$ be the set of items within a given system. For someone with full knowledge of the system some items of this set are related while others are not. The notion of 'is related' should form an equivalence relation $\sim_{r(A)}$ on the set A. By this relation A is split in l $(1 \leq l \leq n)$ equivalence classes $A_1, \dots, A_l$ with $\forall i, j \in \{1, \dots, l\}$, $i \neq j$: $A_i \cap A_j = \emptyset$ and $A_1 \cup \dots \cup A_l = A$. Then items are related to each other iff they belong to the same equivalence class.

An attacker on unlinkability of items within one set knows A but a priori should not know the structure of $\sim_{r(A)}$.

For a random variable X let $P(a_i \sim_{r(A)} a_j) := P(X = (a_i \sim_{r(A)} a_j))$ denote the attacker's a posteriori probability that given two items $a_i, a_j \in A$, X takes the value $(a_i \sim_{r(A)} a_j)$ (or a_i and a_j are related). And $P(a_i \not\sim_{r(A)} a_j)$ denotes the analog probability that a_i and a_j are not related.

The degree of (i, j)-unlinkability $d(i, j)$ describing the unlinkability of two items $a_i, a_j \in A$ a system provides is

$$\begin{aligned} d(i,j) &:= H(i,j) := H(X) \\ &= -P(a_i \sim_{r(A)} a_j) \cdot \log_2(P(a_i \sim_{r(A)} a_j)) \\ &\quad -P(a_i \not\sim_{r(A)} a_j) \cdot \log_2(P(a_i \not\sim_{r(A)} a_j)). \end{aligned}$$

Both anonymity measurement and unlinkability measurement will be needed to describe the privacy our reputation system could provide.

2.3 Privacy-enhancing identity management

Within an Internet community a user develops one or several partial digital (or virtual) identities within this community, 'each represents the person in a specific context or role. [...] A pseudonym can be used as an identifier for a partial identity.' [13].

Identity management systems try to help users to manage the various partial identities they might have for specific applications like Internet communities. User-controlled privacy-enhancing identity management [4] gives the possibility of pseudonymous interaction on the Internet that tries to satisfy the security requirements of all parties involved. Typically the user-server scenario is considered. A user can protect against unauthorised access to personal information while by the use of credentials issued by identity providers the server can be sure pseudonymous users are reliable and can be made accountable for misbehaviour. E.g., the use of an identity management system is applicable to the scenario of classical e-commerce on the Internet as sketched in [5]. But it also could be applied to more general scenarios like Internet communities with equal interactors. 'Identity management is called privacy-enhancing if it does not provide essentially more linkability between partial identities' [13].

By extending [13] we define applications (e.g., Internet communities) where identity management could be used to be designed in a privacy-enhancing identity management enabling way if neither the pattern of actions nor the attributes given to entities (i.e., humans, organizations, computers) acting within the system imply more linkability than is strictly necessary to achieve the purposes of the application. A privacy-respecting application at least tries to respect the privacy of the entities using the application by not necessarily reducing the linkability reached within the application to a minimum but to a level acceptable for all entities involved. This is what our reputation system tries to be because the trust reputation systems try to achieve within an Internet community should still be achievable. Unfortunately trust is something very difficult to measure but we are able to measure at least the anonymity provided with the measures outlined above.

3 Modelling the Internet Community

In the following our model of a centralised Internet community and a privacy-respecting reputation systems that could be established within such a community is outlined.

3.1 Assumptions

We consider an Internet community whose users use global pseudonyms within the system. Linked to the pseudonyms there are specific global reputations. Under certain circumstances users are allowed to give ratings to other users that are added to the global reputation that is associated with the other user's pseudonym. The global reputation has to be updated in a centralized reputation data base.

The circumstances that initiate a rating can be neglected for the model. In example 2 circumstances were sales between pseudonyms.

We limit our model to global reputations and do not consider different views on reputations or inconsistent local storage. Examples of such systems can be found in P2P systems e.g., EigenTrust [12] or [8]. We assume this global and central data base model of reputation systems because it allows to consider a global attacker who has easy access to all reputation information. Every kind of distributed storage, distribution and views only needs to be protected against weaker attackers. Nevertheless our approach can be extended to analyse local storage of global reputation values that might be realised by coins as suggested in [24]. But their approach needs to guarantee that users cannot get rid of negative reputations. This was already formulated in [2] for negative credentials. A central storage guarantees that the reputations collected under one pseudonym can only be used linked to each other, positive and negative ones.

Nevertheless also central storage does not solve the problem of pseudonym changes without a transfer of (especially negative) reputations or the creation of new pseudonyms in Internet communities. Friedman and Resnick [10] propose to charge every newcomer an entrance fee or to use cryptographic mechanisms for anonymous credentials usable for specific purposes in every Internet community. A person using such a pseudonym is not able to change the pseudonym without transferring the reputation collected under this pseudonym.

Anonymity measurement within a reputation system needs the distribution of reputations on all pseudonyms. In the case of local storage these numbers can only be estimated while central storage allows an easy query from the reputation data base.

Let the set of global pseudonyms at time t be $\mathcal{P}_t = \{p_{t,1}, \dots, p_{t,m}\}$ and the set of possible reputations that might be associated with $\mathcal{P}_t$ be R with $(R, +)$ a commutative group and $+$ an operator to combine elements from R and R independent of t.

At time t_1 every pseudonym $p_{t_1,l}$ $(l \in \{1, \dots, m\})$ has a reputation $rep(t_1, p_{t_1,l}) \in R$. Let R' be a subset of R that contains the possible ratings that might be given to pseudonym $p_{t_1,i}$ by pseudonym $p_{t_1,j}$.

Then as long as the pseudonym $p_{t_1,i}$ exists globally and the rating r_{j,i,t_1} that $p_{t_1,i}$ has received at time t_1 from p_{j,t_1} was the only one since t_1 at time $t_2 \geq t_1$ the pseudonym $p_{t_1,i}$ has the new reputation $rep(t_2, p_{t_1,i}) = rep(t_1, p_{t_1,i}) + r_{j,i,t_1} \in R$.

We do not consider the concrete implementation of R, R' and $+$ here but this general model is applicable to many existing reputation systems.

Example 3 (eBay's reputation system). In eBay's reputation system an element of the set R' consists of several elements: a value from {-1,0,1}, a limited free-text field, the time the rating was given and possibly the annotator and the trade considered. The operator $+$ is the simple succession of elements from R' and accordingly it holds $R = R'^*$.

To assure the authenticity of ratings given either each rating could be signed by its issuer's pseudonym. This needs a public-key infrastructure to be established. Or the provider of the community needs to assure authentic ratings by appropriate authentication methods for the pseudonymous accounts established within the community. This needs appropriate trust in the community provider. Note both measures do not prevent an issuer from giving wrong ratings.

3.2 Usage of pseudonyms

All ratings given to the same pseudonym are linkable to each other globally because they build the reputation of one pseudonym. To increase privacy (or especially anonymity of users or unlinkability of their actions) in Internet communities the usage of pseudonyms and the possibility of changing them should follow specific rules which are explained in this section. These rules have effects on the reputation associated with these pseudonyms as also outlined in the following.

Parallel usage of pseudonyms Unlinkability between different contexts (or context types) a member of the community is involved in can be reached by using role pseudonyms regarding to the roles he has in these contexts.

Example 4 (Different contexts in Internet Communities). By the parallel usage of unlinkable pseudonyms the contexts 'offering goods within the community' or 'giving advice regarding a specific topic' or 'chatting about a hobby' could be separated.

This has the positive side effect that reputations for different roles are collected separately. This should even increase the trust in the reputation system because members might be different trustworthy depending on the context. The definition of a context and the distinction between contexts has to be made in the reputation system to make the reputations collected under a pseudonym sensible. All members with access to the reputation system have the opportunity to link all context information regarding the respective pseudonym.

Consecutive usage of pseudonyms Beneath using role pseudonyms for different contexts users should change the pseudonyms they use within these contexts from time to time. This gives members the possibility to determine the linkability of their actions within the community. They might even use their reputation with different sequenced pseudonyms. In a global system like in our model there are several attacker models imaginable under which unlinkability can be guaranteed that are similar to the attacker models of anonymous communication:

The users might trust the provider of the reputation data base that he changes his pseudonym on demand, then for outsiders the pseudonyms are unlinkable but for the provider they are still linkable.

Example 5 (eBays reputation system). eBays reputation system already allows users to change pseudonyms but the history of all previous pseudonyms is available globally as well. The minimum necessary to reach unlinkability against outsiders would mean that eBay keeps this history secret.

A stronger attacker model (that tries to limit the trust in the provider) would be to include several third parties in the change of a pseudonyms that make several consecutive changes. But this still needs trust that the third parties do not collaborate resp. none observes the communication between them. This is similar to the use of anonymous proxies for anonymous communication.

The use of convertible credentials [2] issued by identity providers enables users to transform statements made about one of his pseudonyms to statements about another one of his pseudonyms while the pseudonyms are still unlinkable to each other for everyone except himself. But in the case of a misuse the identity providers can reveal a pseudonym's user. This would mean that in the case of a pseudonym change the user asks the identity provider to issue a credential on his reputation value that he can convert to another credential himself and send it to the provider of the global reputation data base.

3.3 Frequency of pseudonym changes

Pseudonym changes naturally only make sense if the reputation related to the pseudonym is the same many other members have as well, the pseudonym's possible anonymity set for the pseudonym change. Usually the possible anonymity sets in Internet communities are quite large. If the number of possible reputations is limited, e.g. by a numerical sum of ratings many members will have the same reputation and thus the anonymity set of one single member could contain all members with the same reputation after a pseudonym change of all members with the same reputation. If the reputation system allows the members to give additional comments regarding their rating, the possibility for the formulation of comments has to be limited as well to guarantee appropriate anonymity sets. Especially digital signatures of issuers have to be omitted but be substituted by signatures of the community provider or identity providers.

Because the change of a pseudonym and the corresponding reputation usually is costly and needs many members to participate, there has to be made a trade-off between the costs of a pseudonym change and the linkability of information regarding a pseudonym. In the following measurement of anonymity for pseudonym changes we concentrate on the optimal points for pseudonym changes regarding the anonymity sets considered. Note that we could also measure unlinkability of pseudonyms instead of anonymity regarding a pseudonym change. Due to the lack of space we omit this measurement and also neglect that and how the number and kind of ratings collected under one pseudonym will influence the times of pseudonym changes as well.

4 Measuring Anonymity within an Internet Community

Using the terms common in identity management reputation is an attribute that in contrast to many other attributes varies frequently over the time. Regarding this attribute we analyse the user's unlinkability regarding parallel usage of pseudonyms and his anonymity regarding pseudonym changes, the measures suggested above to reach a privacy-respecting design of reputation systems. For many applications this examination should be extended to more attributes a user might have.

Unlinkability of parallel used pseudonyms If users do not reveal additional information the parallel usage of pseudonyms is not linkable, this means for all linkability relations $\sim_{r(\mathcal{P}_{t_1})}$, pseudonyms $p_{t_1,i}, p_{t_1,j} \in \mathcal{P}_{t_1}$ and $i \neq j$ the (i,j)-unlinkability $d(i,j)$ the reputation system provides is

$$\begin{aligned} d(i,j) = &-P(p_{t_1,i} \sim_{r(\mathcal{P}_{t_1})} p_{t_1,j}) \cdot \log_2(P(p_{t_1,i} \sim_{r(\mathcal{P}_{t_1})} p_{t_1,j}) \\ &-P(p_{t_1,i} \not\sim_{r(\mathcal{P}_{t_1})} p_{t_1,j}) \cdot \log_2(P(p_{t_1,i} \not\sim_{r(\mathcal{P}_{t_1})} p_{t_1,j}) = \frac{1}{2}. \end{aligned}$$

As far as no additional information becomes known (e.g. how many users have a specific number of pseudonyms) the unlinkability has the maximum value above.

Anonymity of a pseudonym change Usually the set of pseudonyms at time t_1 is split into several subsets regarding the relation 'has the same reputation'. Now we consider one of these subsets, the set $\mathcal{P}_{t_1,r}$ with the pseudonyms having the reputation $r \in R$:

$$\mathcal{P}_{t_1,r} = \{p_{t_1,i} \quad | \quad rep(t_1, p_{t_1,i}) = r.\}$$

If a global pseudonym change within this group is announced a subset of this set will change its pseudonyms with the function $c : \mathcal{P}_{t_1,r} \rightarrow \mathcal{P}_{t_2,r}$ with $\mathcal{P}_{t_2,r}$ the resulting pseudonym set after the pseudonym change. None of the members has an advantage regarding the anonymity of his own profile in not participating in the pseudonym change but concurrently he decreases the other pseudonyms' anonymity. This might be a motivation for an attacker.

Every pseudonym in $\mathcal{P}_{t_2,r}$ is 'related' to exactly one pseudonym in $\mathcal{P}_{t_1,r}$ because no new pseudonym could be added to $\mathcal{P}_{t_2,r}$ without having reached reputation r at time t_1 and thus belonging to $\mathcal{P}_{t_1,r}$. If there exist users who have not participated in the pseudonym change phase it holds $\mathcal{P}_{t_2,r} \cap \mathcal{P}_{t_1,r} \neq \emptyset$.

According to section 2 the effective size of the anonymity set probability distribution regarding the anonymity $p_{t_2,j}$ has for the pseudonym change c can be calculated (for X the random variable that $p_{t_2,j}$ resulted from $p_{t_1,i}$ through c):

$$H(X) = - \sum_{p_{t_1,i} \in \mathcal{P}_{t_1,r}} P(c(p_{t_1,i}) = p_{t_2,j}) \log_2(P(c(p_{t_1,i}) = p_{t_2,j}))$$

The continual growth of reputations within the system produces different sizes of anonymity sets.

If the pseudonym change is initiated for all $r \in R$ at the same time, the same degree of anonymity for all pseudonyms is reachable at time t_1 when it holds $|\mathcal{P}_{t_1,r_1}| = |\mathcal{P}_{t_1,r_2}| \; \forall r_1, r_2 \in R$ and all pseudonyms participate in the pseudonym change.

But the goal of every user is to maximise his own degree of anonymity of a pseudonym change. He is interested in maximising his own anonymity set.

The other possibility for the provider is to fix minimal sizes of the anonymity sets as security parameter that should be reached for a pseudonym change and initiate pseudonym changes for each anonymity group separately. But this may lead to more pseudonym changes for single pseudonyms who change the anonymity set meanwhile while other pseudonyms often might 'miss' the pseudonym changes by changing the anonymity set.

5 Summary and Future Work

We have presented privacy-respecting design options of centralised reputation systems for Internet communities that keep the level of trust provided to the members by the use of reputations. Common information-theoretic models were used to evaluate the unlinkability/anonymity of user reputation profiles the proposed privacy measures provide. The proposed measurements allow to study similar aspects of privacy in other systems, especially single changing user attributes in privacy-enhancing identity management systems.

In the near future we will study the linkability regarding the information collected under one pseudonym in interplay with the provider-initiated points of pseudonym changes that are determined by the anonymity sets regarding reputation. Especially we will make simulations on the privacy-enhancing reputations systems to allow the user scalability of reputation and privacy and especially give him feedback and suggestions on both factors. Further in future research the model has to be extended for privacy-enhancing identity management to allow the measurement of combinations from more than one user attribute than just reputation.

References

1. Oliver Berthold, Hannes Federrath, and Stefan Köpsell. Web mixes: A system for anonymous and unobservable internet access. Designing Privacy Enhancing Technologies. Proc. Work-

shop on Design Issues in Anonymity and Unobservability, LNCS 2009, Springer-Verlag, Heidelberg 2001, pp. 115–129.

2. David Chaum. Showing credentials without identification - signatures transferred between unconditionally unlinkable pseudonyms. Advances in Cryptology - EUROCRYPT 85, LNCS 219, Springer-Verlag Berlin 1986, pp. 241–244.
3. David Chaum. Untraceable electronic mail, return addresses and digital pseudonyms. Communications of the ACM, 24(2), 1981, pp. 84–88.
4. Sebastian Clauß, Andreas Pfitzmann, Marit Hansen, and Els Van Herreweghen. Privacy-enhancing identity management. The IPTS Report 67 (September 2002), pp. 8-16.
5. Sebastian Clauß and Marit Köhntopp. Identity management and its support of multilateral security. Computer Networks 37 (2001), Special Issue on Electronic Business Systems; Elsevier, North-Holland 2001; 205-219.
6. Chrysanthos Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. ACM Conference on Electronic Commerce, 2000, 150-157.
7. Claudia Díaz, Stefaan Seys, Joris Claessens, and Bart Preneel. Towards measuring anonymity. In Roger Dingledine and Paul Syverson, editors, *Proceedings of Privacy Enhancing Technologies Workshop (PET 2002)*. Springer-Verlag, LNCS 2482, April 2002.
8. Roger Dingledine, Nick Mathewson, and Paul Syverson. Reputation in P2P Anonymity Systems. In *Proceedings of Workshop on Economics of Peer-to-Peer Systems*, June 2003.
9. Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.
10. Eric Friedman and Paul Resnick. The social cost of cheap pseudonyms. Journal of Economics and Management Strategy, Aug. 1999.
11. Dominic Hughes and Vitaly Shmatikov. Information hiding, anonymity and privacy: A modular approach. *Journal of Computer Security*, 12(1):3–36, 2004.
12. Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molinal. The eigentrust algorithm for reputation management in p2p networks. Proceedings of the Twelfth International World Wide Web Conference, May, 2003.
13. Marit Köhntopp and Andreas Pfitzmann. Anonymity, unobservability, and pseudonymity - a proposal for terminology. Draft v0.26, December 2005, `http://dud.inf.tu-dresden.de/Literatur_V1.shtml`.
14. Peter Kollock. The production of trust in online markets. Advances in Group Processes (Vol. 16), Greenwich, CT: JAI Press., 1999.
15. Tobias Mahler and Thomas Olsen. Reputation systems and data protection law. eChallenges e-2004 Conference, Vienna, October 2004.
16. M. K. Reiter and A. D. Rubin. Crowds: Anonymity for web transactions. ACM Transactions on Information and System Security 1(1), November 1998, pp. 66–92.
17. Steve Schneider and Abraham Sidiropoulos. CSP and anonymity. ESORICS 1996, LNCS 1146, Springer-Verlag Berlin 1996, pp. 198–218.
18. Andrei Serjantov and George Danezis. Towards an information-theoretic metric for anonymity. Privacy Enhancing Technologies 2002, LNCS 2482, Springer-Verlag Berlin.
19. Vitaly Shmatikov. Probabilistic analysis of anonymity. Proc. 15th IEEE Computer Security Foundations Workshop (CSFW) 2002, pp 119–128.
20. Sandra Steinbrecher. Balancing privacy and trust in electronic marketplaces. DEXA Conference on Trust and Privacy in Digital Business 2004, LNCS 3184, Springer Verlag Berlin, pp. 70-79.
21. Sandra Steinbrecher and Stefan Köpsell. Modelling unlinkability. In Roger Dingledine, editor, *Proceedings of Privacy Enhancing Technologies workshop (PET 2003)*. Springer-Verlag, LNCS 2760, March 2003.

22. Paul F. Syverson and Stuart G. Stubblebine. Group principals and the formalization of anonymity. FM'99 – Formal Methods, Vol. I, LNCS 1708, Springer-Verlag Berlin 1999, pp. 814–833.
23. Gergely Tóth, Zoltán Hornák, and Ferenc Vajda. Measuring anonymity revisited. In Sanna Liimatainen and Teemupekka Virtanen, editors, *Proceedings of the Ninth Nordic Workshop on Secure IT Systems*, pages 85–90, Espoo, Finland, November 2004.
24. Marco Voss. Privacy preserving online reputation systems. In *International Information Security Workshops*, pages 245–260. Kluwer, 2004.
25. Graeme Wearden. Judge raps ebay over fraud. December 7, 2004, available from `http://news.com.com/2102-1038_3-5481601.html`.