A Privacy Preserving Reputation System for Mobile Information Dissemination Networks

***Abstract*— In a mobile information dissemination network mobile users, equipped with wireless devices, exchange information in a spontaneous manner whenever they come into communication range. Users have to specify what kind of information they are looking for and what kind of information they can offer. A priori there is no relation between users, literally spoken, they don't know each other and confidence in newly collected information might be low.**

This work presents two reputation schemes, a simple and an extended version, for mobile information dissemination networks that, based on user ratings, increase a user's confidence in some information source. As reputation systems collect sensitive personal information and monitor users' behavior, privacy is an essential requirement — especially in a mobile scenario — that is neglected by many existing approaches. Using cryptographic group signatures and the concept of an observer, our extended reputation scheme guarantees high user privacy.

I. Introduction

The increasing spread of mobile wireless devices with ad hoc or short range networking capabilities, e.g. bluetooth-enabled mobile phones or handheld computers (PDAs) with integrated 802.11 WiFi technology, opens the path for new applications for mobile Ad-hoc Networks.

Motivated by the way information spreads by word-of-mouth communication between humans, we use mobile devices to form *mobile information dissemination networks*. In a mobile information dissemination network, users share information as they come into communication range to each other. Their devices match information they want to share and information they are looking for. If there is a match, information is passed without any further user interaction. This can be viewed as *information infection*.

We assume no relation between the participating users, in the worst case they are anonymous to each other. In this situation a reputation system could improve the usefulness of information dissemination networks. Consider the following example:

> Bob's favorite Italian restaurant in town is *Da Pino* and he thinks that they serve delicious pasta for reasonable prizes. Bob provides this information via his device. Now Alice is interested in all kinds of restaurant information and while passing by Bob's device, her device learns about *Da Pino*. But how could she tell that this information is trustworthy?

If Bob had acquired some good reputation from others who have already checked out Bob's restaurant recommendations, Alice could benefit from this extra information and give it a try.

We define *reputation* as the collected information about one entity's former behavior as experienced by others. *Trust* expresses an entity's willingness to proceed with an action that might be harmful based on information like the risk, benefit and reputation of involved entities. Reputation systems collect ratings, process and consolidate this information and make it available upon request. This has two effects: They provide information that allows users to predict how someone may behave in future. Simultaneously, they influence the future by giving the incentive to behave well.

At a first glance it seems unreasonable to base trust in information spread by a user on her reputation, because she may only pass on what she has learned from someone else. However, our reputation systems should motivate users to verify information before spreading it further or to mark it as uncertain.

Most existing approaches to use reputation systems for mobile ad-hoc networks concentrate on solving routing issues raised by misbehaving nodes. They also neglect the dynamics of a scenario with highly mobile nodes where the chance to meet someone again is low. In the majority of cases privacy issues are completely ignored. Especially when a human user as the owner of a device is involved, reputation is sensitive personal information. Often it contains detailed information about the entity's former transactions including the subject and a list of all peers that have been involved. Additionally, it is critical for an entity not to lose reputation by false accusation.

Therefore, it is important to design reputation systems that do not jeopardize privacy. Essential requirements are rater and ratee[1] anonymity and that an entity has control over its own reputation information.

The contribution of this work are two reputation schemes for mobile information dissemination networks. The first one should illustrate what can be achieved without relying on a centralized trusted party for reputation management. A key characteristic of the second scheme guarantees high user privacy by using a trusted local component called observer. In this work we use iClouds as a reference architecture for an information dissemination network.

This work is structured as follows: The next part of the introduction will present the notation and cryptographic primitives we use throughout the text. Section II introduces the iClouds project and outlines important key characteristics and system components in iClouds. Section III has the focus on reputation in iClouds. Then in Section IV we present the first reputation scheme and extend this in Section V to a reputation scheme with high user privacy properties. In Section VI we present related work on the topics information dissemination networks, reputation systems and recommender and collaborating filtering systems. We conclude our work in Section VII.

A. Notation and Cryptographic Primitives

Since our approach makes use of standard public key cryptography, we denote $S_{priv\_k_A}(some\_data)$ as the digital signing operation on *some_data* carried out by user A using her private key and $V_{pub\_k_A}(some\_data)$ as the corresponding verification operation.

In the same manner we define $C_{pub\_k_A}(some\_data)$ and $D_{priv\_k_A}(some\_data)$ as the encryption respectively decryption operation on *some_data* using A's public-key (respectively private-key).

To certify that a public key belongs to a registered user we rely on a simple public key infrastructure (SPKI) [1] with a Certification Authority (CA) that issues a certificate $Cert_{pub\_k_A}$ for user A with public key pub_k_A. In contrast to other certificate standards like X.509 SPKI certificates do not contain any additional data apart from the public key.

We have chosen an ElGamal signature and encryption scheme [2] as implementation of the public key operations introduced above. We do not want to give a formal description of this schemes here. They are based on the problem that computing discrete logarithms in a carefully chosen cyclic group G is infeasible. A's public key is $y = g^x$ where $g \in G$ is a generator of G and x is A's private key. This characteristic allows us to create a blinded representation of a public key as described later on.

Waters et al. [3] describe an extension of an ElGamal encryption scheme where the owner A of a secret key can generate a number of *incomparable* public keys. Every public key can be used by a sender B to encrypt a message for A, but two of these keys are not linkable to A. This scheme allows A to receive encrypted messages anonymously using public key encryption. A simple alternative is to use everytime a new random public / privacy key pair with much higher administration effort.

Another important primitive used later on is the *group signature* scheme [4], [5]. Let $\Gamma = \{A_1, ..., A_n\}$ be a set of group members. Each group member A_i registers with a designated entity, called group manager, and receives its private key $priv\_k_{GA_i}$. The group manager publishes the group's public key pub_k_G. Every group member can sign a message m on the group's behalf using the operation denoted $S^G_{priv\_k_{GA_i}}(m)$. Everyone can verify the resulting signature s by performing $V^G_{pub\_k_G}(s)$ but no one except the group manager can find out which group member was the originator.

Further $\#n_{RP_{(t)}}$ denotes a certain amount *n* of reputation points issued at a certain time *t*.

II. The iClouds Project

In this section we present key concepts, characteristics and system components of the iClouds Project [6].

iClouds allows mobile users to form a spontaneous network by using mobile devices with short range communication capabilities (our prototype uses 802.11b WiFi (Ad-hoc-Mode) enabled PDAs). A priori, we do not assume any relation between users.

With iClouds, users are able to share information with others in a peer-to-peer manner, i.e. users publish information and subscribe to information interests. Using a kind of user profile, we call it *iWish*- and *iHave*-list, devices match information they want to share and information they are looking for. If there is a match between user's iWish-list and another user's iHave-list, information is passed without any further user interaction. By this we mimic the way information spreads by word-of-mouth between humans. For more details on the information modeling and matching in iClouds see [7], [8].

A note on privacy: In the basic setup of iClouds, the communication channel is unencrypted and shared

[1] a ratee is the target of a rating issued by a rater

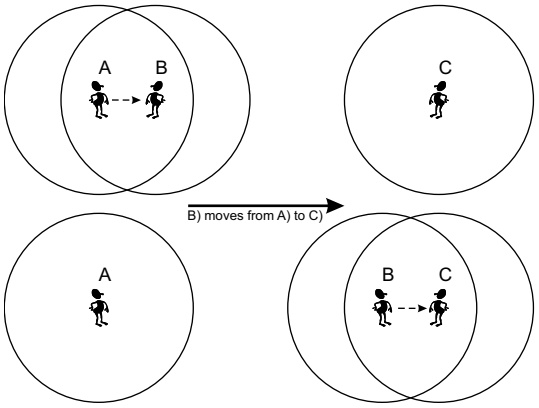

Fig. 1. Multi-hop Information Dissemination

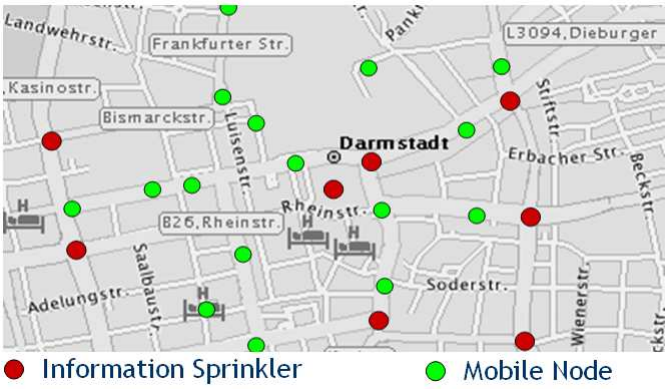

Fig. 2. iClouds System Components - City Setup

(i.e. the radio-based physical layer). Therefore it is easy to tap communication. Also devices are identifiable by their address (e.g. MAC-address). Since iClouds does not provide multi-hop communication (see below), it is feasible to choose a new random MAC-address from time to time. This would alleviate re-recognition of previously encountered devices (see [9] for details).

A. Information Dissemination in iClouds

Unlike the proposed algorithms of the MANET [10] group, iClouds deliberately does not support multi-hop communication. The reason for this derives from the fact that users are anonymous to each other. It can be broken down into the following considerations:

- In anonymous groups of users there is little to no incentive of an individual member to act as an intermediate node and provide its energy, CPU or memory resources. Especially since battery power is still a resource bottleneck in today's devices.
- Intermediate nodes are in the perfect position for all kinds of *man-in-the-middle*-attacks. Why should communicating partners trust and rely on the correct behavior of any intermediate nodes?

Therefore we believe that multi-hop is suitable for *closed*-group communication, as found in military- or rescue-scenarios, but not for anonymous groups of people.

Even though iClouds communication happens only within a user's vicinity, i.e. is *one-hop*, and is limited by the communication range of a user's device, there is a kind of *multi-hop* information dissemination, see Fig. 1.

Here user A comes into communication range to B. Suppose now that A passes some information to B. If B later encounters another user C, this information can be passed to C as well.[2]

B. iClouds System Components

There are two system components in iClouds:

1) iClouds device
   A mobile device with short range communication capabilities, e.g. 802.11b WLAN or Bluetooth, that runs the iClouds software. The current prototype implementation uses 802.11b WLAN and runs on a Compaq iPAQ.
2) Information Sprinkler (IS)
   A fixed device that broadcasts information relevant to the area where the IS is located, e.g. an IS could provide digital advertisements in a shopping mall (see [9]).
   This paper extends the IS functionality to work also as an information sink (to collect ratings). Also all sprinklers are connected by a backbone network, e.g. a metropolean area network.

Fig. 2 shows a hypothetical system setup. Information Sprinklers are positioned at cross-roads and public places (red dots). Users (green dots) move around the city.

III. Reputation in iClouds

Since iClouds works for an a priori *unrelated* group of mobile users, for example, consider users passing by each other in a pedestrian zone, the *source* of information is unknown. Therefore to evaluate some information I, a user can only rely on this own judgments.

Here a reputation scheme could improve a user's confidence in I. Also it would alleviate the spreading of *bad* or *wrong* information.

[2]We assume that A has some information to provide, both B and C are interested in.

While not being the focus of this paper, the basic idea is that the information can be filtered by taking into account a user's reputation. See section IV-B for a simple equation to calculate a user's reputation. Also we believe that the prospect to receive accurate information will motivate other users to take part in the rating scheme.

There is no reason to restrict the reputation system only to information dissemination. It could also be applied for other transactions a user performs using her mobile device. For instance, a user's reputation could also be involved when buying something or using some other services. This leads to the problem that reputation is context dependent. A number of positive ratings earned for being an expert in computer science say nothing about the quality of restaurant recommendations or being an honest customer. To simplify matters we do not deal with this problem in this paper and assume that the used contexts are compatible.

A. Reputation and Privacy

A user's privacy is the more important the more information technology invades our daily life. The iClouds system is such a technology that can have many benefits by providing information about the current environment. On the other side it represents a real danger, because it allows others to monitor our habits, preferences and movements.

We do not want to give a full definition of privacy here but only stress the two most important attributes: In the first place privacy means control over sensitive personal information. The owner of this information should be able to decide who can read, modify and distribute this information. The user's right to be unobserved is equally important.

This is especially true for a reputation system whose major objective is to collect information about the former behavior of users. As argued in [11] the location where reputation information is stored is important for the privacy a system can provide. A centralized system takes away the control from the user and does not meet the iClouds scenario which assumes only a limited fixed infrastructure.

Therefore, our approach relies on a decentralized system with local storage of reputation information. This means that every user stores her own reputation and not ratings about others. A user's reputation consists of a list of discrete reputation points. Each positive rating leads to a new reputation point. To provide as much privacy as possible is an overall design goal.

B. Requirements of a Privacy Preserving Reputation System for iClouds

After having discussed the relevance of privacy, we now summarize the requirements of a reputation system that provides privacy protection. These requiremets have already been motivated in [11].

General requirements:

- A reputation system must provide information that allows users to distinguish between trustworthy and untrustworthy peers (information).
- It should encourage entities to behave trustworthily (motivation).
- An entity must not be able to fake a reputation value (unforgeability).
- The reputation information must be bound to an entity. Lending or trading reputation must not be possible (accountability).
- Negative ratings (not only positive ones) should be supported.
- An entity must not be able to get rid of a negative reputation (integrity).
- An entity should not have 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 (ballot-stuffing attack).
- It should not be possible to defame someone without proof (defamation).

Privacy related 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 (control).
- 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 little information as possible about the transaction between rater and ratee.

IV. Simple Reputation Scheme

We now describe a first (simple) reputation system for iClouds that deals only with positive ratings. Since we do not introduce a trusted third party it is not possible to guarantee the integrity of the list of received ratings. This would be a necessary premise to include negative

ratings. Negative ratings are relevant if it is important to detect fraud immediately. However, this is not the case in our scenario. We also believe that negative ratings might discourage users to participate in spreading some information which cannot be judged objectively (i.e. taste).

In this scheme users have to register with the CA to get a digital ID. We make use of SPKI certificates [1] here that only contain the user's public key and the CA's signature. This CA is part of the iClouds infrastructure. The registration process has to make sure that it is not possible to create more than one identity per real user.

A. Passing Information and Rating

Consider user A comes into communication range to user B and A has some information to offer for user B, e.g. there is a match between an *iWish*-entry of user B and an *iHave*-entry of user A.

User A constructs the following offer O for user B regarding the information I[3]:

$$O := \boxed{\begin{array}{l} I, \\ S_{priv\_k_{Rater_i}}(\#1_{RP_{(t_i)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_i}}\ , \\ S_{priv\_k_{Rater_j}}(\#1_{RP_{(t_j)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_j}}\ , \\ S_{priv\_k_{Rater_k}}(\#1_{RP_{(t_k)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_k}}\ , \\ \ldots\ , \\ Cert_{pub\_k_A}\ , \\ \textit{Nonce} \end{array}}$$

The $Nonce$ is a random number that functions as a marker and allows A to efficiently recognize messages encrypted for her.

Let S be the set of all signed reputation points available to user A. We want to shield A against the risk that B learns about all previous satisfied[4] communication partners of A. Therefore we leave it up to A to include any subset $S^O \subseteq S$ into a certain offer O.

Step 1: A sends $S_{priv\_k_A}(O)$ to B.

User B uses the information provided by A according to her needs. If B finds I useful and therefore B wants to provide A with a positive rating, B constructs the following rating R at time t_0

$$R := \boxed{\begin{array}{l} S_{priv\_k_B}(\#1_{RP_{(t_0)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_B} \end{array}}$$

Step 2: B sends ($C_{pub\_k_A}(R)$, $Nonce$) to A directly or to the next IS if A is out of range. This data is distributed to all Information Sprinklers in the network.

Step 3: If A comes again into communication range to an IS he asks for new ratings using the $Nonce$ as a reference. An IS can look this up and delivers $C_{pub\_k_A}(R)$ to A. A is able to decrypt and validate this packet and next time another user asks for the information I, A can offer I by constructing O^* that includes the new rating of user B, see below:

$$O^* := \boxed{\begin{array}{l} I, \\ S_{priv\_k_{Rater_i}}(\#1_{RP_{(t_i)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_i}}\ , \\ S_{priv\_k_{Rater_j}}(\#1_{RP_{(t_j)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_j}}\ , \\ S_{priv\_k_{Rater_k}}(\#1_{RP_{(t_k)}} \quad \text{for} \quad pub\_k_A)\ , \\ Cert_{pub\_k_{Rater_k}}\ , \\ \ldots\ , \\ \boldsymbol{S_{priv\_k_B}(\#1_{RP_{(t_0)}} \quad \text{for} \quad pub\_k_A)}\ , \\ \boldsymbol{Cert_{pub\_k_B}}\ , \\ Cert_{pub\_k_A}\ , \\ \textit{Nonce} \end{array}}$$

B. Reputation Usage

Now consider user A meets another party, user C. The informal protocol that includes the usage of the previously gained reputation looks like the following:

Step 1: A meets C and C shows interest in I.

Step 2: A sends $S_{priv\_k_A}(O^*)$ to C. C is able to evaluate the following:

- Since this data is signed by A, user C is able to validate (by using the certificate $Cert_{pub\_k_A}$, that all reputation points where issued for A.
- C is able to verify that all raters are regular members of the system.
- Using the time-stamps C knows the *age* of the information and can judge on this.
- Since the reputation points are signed, C is able to distinguish between ratings of the same and a different user. A simple ballot-stuffing attack with multiple ratings from the same user can be easily discovered.

Our scheme does not prevent a group of friends to rate each other. This would require knowledge about the

[3]Note that I is not bound to the ratings, i.e. we do not support *per information rating*.

[4]Obviously, B will not be able to learn about the unsatisfied communication partners. Or communication partners with a selfish behaviour that do not rate at all.

social network of a rated user. However, we assume that the number of interactions with unknown people is much bigger [5] than the ordinary number of friends.

A naïve evaluation algorithm would simply calculate

$$X = \sum_{i=1}^{n} \frac{(\#1_{RP_{(t_i)}} \text{ from } Rater_j) * f(t_i)}{m(j)}$$

where n = number of ratings for user A, $m(j)$ is the number of ratings issued by $Rater_j$ for user A and $f(t_i)$ is a function that weights the reputation point by its timestamp, i.e. older reputation points get a lower factor. User C (in the example above) then could specify a lower bound for X in order to accept information I.

C. Privacy Properties

The presented scheme does only provide very limited privacy for the participants. The registered public key functions as a user's pseudonym. The use of SPKI certificates guarantees that no real username is involved. All messages between A and B contain this pseudonym. Consequently, users cannot issue an information offer or a rating anonymously.

Moreover the receiver of an offer from A learns the pseudonyms of all former transaction partners contained in the rating list.

The information sprinkler does not learn the contents of a submitted rating, because it is encrypted with the ratee's public key. Nevertheless, the IS (and everyone else) can still listen to all messages in communication range and can notice all offers sent between users.

An advantage of the local storage of reputation is that A has full control over the list of reputation points. She can decide to whom she shows her reputation and how the list of ratings is composed. Since there are no negative ratings, integrity for the list of ratings is not necessary.

V. Extended Reputation System with Observer

In this section an extended version of the protocol is presented. Its aim is to provide more privacy for both rater and ratee. All explicit identification information is removed from the rating message and the list of reputation points. However, a totally anonymous system is open for the ballot-stuffing attack where a group of friendly users rate each other over and over again.

[5]Imagine how many people you meet at a train station on your way to work.

To prevent this a trusted component called *observer* is introduced on the client device. This term was originally used by Chaum and Pedersen [12]. An observer is a tamper-proof module issued by an organization. Together with the client device the observer forms an electronic wallet that stores a database of user credentials. The observer guarantees the correctness and integrity of this data. This means that the user is not able to change the data illegally. At the same time the observer maintains the user's privacy by revealing only necessary information.

This seems to be he a quite heavy approach with restricted feasibility. Management and distribution of observer components requires a costly organizational infrastructure. However, we envision the mobile phone as a possible target device to run the iClouds software. This already provides a perfect environment for the observer: the smart card based Subscriber Identity Module (SIM card).

It is important that there is no direct communication between the observer and the outside. All messages are mediated through the client device. The protocols have to be designed in a way that there is no uncontrolled information *inflow* or *outflow* between the observer and a third party. For instance the observer is not allowed to choose random values alone but only with the help of the client device. Otherwise the observer can encode secret information into these random numbers. An ideal observer only stores a part of the database that is not readable without the user's help. In this case capturing an observer and even breaking its tamper-resistance does not leak any usable information. However, in the presented scheme a simple observer is used that stores fully readable data entries.

A. Setup

As before, every user has to register with a central authority and receives her certified public key and the observer that has to be plugged into her client device. The observer OS is personalized with the user's public key. It has installed a private key $priv\_k_{GOS}$ to issue group signatures. All observers share the same group public key pub_k_G. Additionally, the user and the observer use the technique of [3] to generate incomparable public keys if required. To prevent linkability on communication level the users choose a new random hardware address for their device every time this is appropriate.

B. Role of the Observer

The task of the observer is to prevent ballot-stuffing, to assure the rater's privacy and to attest the authenticity of reputation points. A rating message still contains the identity of the rater, but visible only for the ratee's observer. An additional inner envelope is added to a rating that has to be removed by the observer before it can be used by the ratee. This means that a rating has to be activated by the observer. The envelope contains the same information about the identity of the rater as in the former scheme (signature and certificate of B). The observer stores this identity to prevent multiple ratings from the same user during the lifetime of a rating. More than one rating from the same rater is not activated by the observer. The observer removes all identifying information from the rating before it returns it as a reputation point to the client. The authenticity is attested by a group signature of the observer.

As already mentioned the observer maintains a table of rater identities with active ratings. The observer may not have a build-in clock, for instance when it is implemented as a smart card. In this case it is not possible to give reputation points a time based lifetime. Therefore, we use a solution that counts the received ratings. The observer adds a counter to every entry in the table and increases it with every received rating. When the counter reaches a predefined limit the entry is removed from the table. Consequently, a new rating from the same rater will be accepted by the observer but also the former reputation point will not be included anymore when the user asks for an updated list of active ratings.

To guarantee anonymity of the client it is necessary to make it not identifiable by its observer. To achieve this, the group signature scheme is used. The signatures of all observers can be verified with the same public key. However, a signature is not linkable to a concrete instance of an observer. Two signatures of the same observer cannot be linked as well.

C. Phases of the Protocol

As described in the simple scheme, it is assumed that A has some information to offer to B. Since A cannot use its certified public key when she wants to stay anonymous she uses a new incomparable public key $rpub\_k_A$.

For the same reason she asks the observer also to generate a new incomparable public key $rpub\_k_{OS}$. To ensure integrity of the protocol B has to know that this key really comes from an observer. Therefore, the observer signs $rpub\_k_{OS}$ with its group signature key.

User A constructs the following offer O for user B regarding the information I:

$$O := \boxed{\begin{array}{l} I, \\ \text{[attested list of reputation points]}, \\ S^G_{priv\_k_{GOS}}(rpub\_k_{OS}) \ , \\ rpub\_k_A, \\ \textit{Nonce} \end{array}}$$

The contents of the list of reputation points will be described later. Again the $Nonce$ functions as a marker that allows A to recognize encrypted rating messages intended for A.

Step 1: A sends O to B.

User B uses the information provided by A according to her needs. If B finds I useful and therefore B wants to provide A with a positive rating, B constructs the following rating R at time t_0

$$R := \boxed{\begin{array}{l} S_{priv\_k_B}(\#1_{RP_{(t_0)}}) \ , \\ Cert_{pub\_k_B} \end{array}}$$

In contrast to the simple scheme R does not contain the public key of A any more. The binding between R and A is achieved by encrypting it for A's observer. B verifies that $S^G_{priv\_k_{GOS}}(rpub\_k_{OS})$ is a valid signature.

Step 2: B constructs the envelope for the observer $E := C_{rpub\_k_{OS}}(R)$ and sends $(C_{rpub\_k_A}(E), Nonce)$.

Step 3: A receives the rating directly or via IS as described in the simple scheme. A is able to decrypt the packet but cannot use the rating because of the inner envelope. A passes E to its observer.

Step 4: The observer decrypts the received data and verifies B's signature of the rating. If B is not in the table of active raters the observer adds a new entry for B and constructs a new reputation point for A. To bind this reputation point to A, the observer includes the blinded public key of A in this structure. In an ElGamal signature setting this can be achieved as follows: x is secret key of A and $y = g^x$ is public key. The observer chooses a random r and includes $z = h^r y$ into the result and gives r to A. h is another generator of the underlying group.

The observer returns a new reputation point

$$P := \boxed{S^G_{priv\_k_{GOS}}(z, \#1_{RP})}$$

Step 5: A adds this data to her list of reputation points.

D. Showing Reputation

To use the reputation point, A has to prove that it knows the corresponding private key to the public key encoded in z. According to [5] such a proof of

knowledge protocol is constructed as follows: A has to give a proof of knowledge of numbers α and β such that $z = h^{\alpha}g^{\beta}$ holds (short $PK\{(\alpha,\beta) : z = h^{\alpha}g^{\beta}\}$). From this proof of knowledge protocol a signature scheme on a message m can be derived, denoted by $SPK\{(\alpha,\beta) : z = h^{\alpha}g^{\beta}\}(m)$. This scheme does not leak any information about A's public key.

Still it is a problem that showing the list of reputation points more than once to the same peer (or to a group of peers that exchange the collected data) makes it linkable. In fact, it is a kind of implicit pseudonym of A. There are two solutions to solve this. The first makes use of the fact that the offer O has to include an authenticated public key of the observer. This means that the observer has to perform a signing operation anyway. Consequently, we adapt the structure of O as follows:

$$O' := \boxed{\begin{array}{l} I, \\ S^{G}_{priv\_k_{GOS}}(z, \#1, \#1, ..., rpub\_k_{OS}) \ , \\ rpub\_k_{A}, \\ Nonce \end{array}}$$

Step 1 is changed to: A sends O', $SPK\{(\alpha,\beta) : z = h^{\alpha}g^{\beta}\}(O')$ to B.

This means that the observer builds a new reputation list whenever A wants to send an offer. Instead of using a list of single reputation points it is much easier then to ask the observer to attest a statement like 'A's reputation is bigger than n'. The drawback of this solution is that the observer is much more involved in constructing an offer.

The second solution is to combine the group signature scheme of the observer with an anonymous credential system. Showing such a credential is unlinkable, because either everytime a new representation of the credential is calculated or zero knowledge proofs of knowledge are used. Unfortunately, the authors do not know if there are such systems that allow this combination. A very primitive solution is to use the same RSA keypair for every observer (as replacement for a real group signature) and a RSA based anonymous credential system like [13].

We will investigate this problem further in our future work.

E. Privacy Properties

In comparison with the simple scheme presented in section IV the extended version provides nearly perfect privacy for the participants. The prize for this is a higher cryptographic effort and the need for a local observer. External third parties learn no information about the identity of the owner of a reputation and cannot link the actions by the presented reputation points. The ratee learns nothing about the rater because she sees only an encrypted rating. Unless the observer denies to activate a rating the ratee knows that it comes from someone it has already interacted with before.

A major drawback is that the observer participates in every transaction and learns the identities of all peers involved. Therefore it is really critical to guarantee that no unwanted information can leak from the observer. When the observer is lost, there is a danger that someone is able to tamper it and retrieve all stored information. Our future work will concentrate on solving this disadvantage by storing only a necessary part of this data inside the observer.

VI. Related Work

Information sharing and dissemination in mobile networks is the subject of the 7DS architecture [14]. Mobile nodes share storage and connectivity to the Internet by forming an Ad-hoc Network. A caching approach is used to make data from websites available within nodes that are in communication range to each other, even if the access to an Internet gateway is lost, e.g. due to mobility reasons.

The Proem platform [15] supports user collaboration in wireless networks in a more general form and by information sharing only. Proem is a quite big architecture and not suitable for PDA-like devices.

Other examples that wireless communication, collaboration and information sharing among mobile users make sense, are the *Shark* system [16] to share knowledge, *Usenet-on-the-fly* [17] for leisure information dissemination and a taxi sharing scenario described in [18].

Reputation systems [19] evolve as a mechanism to build trust in dynamic electronic societies and are already used in many applications. Trust and reputation have gained a lot of attention of researchers from many disciplines. Mui et al. [20] give a short overview of different notions of these terms within different disciplines from sociology to economics to computer science. The most popular reputation system is for sure the feedback mechanism of *eBay* [21]. The empirical study of [22] shows the effect of a high reputation value on the achieved auction prize.

Existing approaches to implement reputation systems for mobile ad-hoc networks like CONFIDANT [23] and CORE [24] concentrate on routing issues raised by misbehaving nodes.

Most reputations systems do not deal with the user's need of privacy. A detailed discussion about reputation

systems and privacy can be found in [11]. The author argues that only a distributed reputation system with local storage of reputation gives the owner control over her rating information. Fahrenholtz and Lamersdorf [25] propose such a system with local storage that relies on a central portal to guarantee integrity of aggregated ratings by maintaining a counter for each user. Since the ratings contain the raters' identities and signatures, this scheme does not provide much privacy.

Pavlov et al. [26] focus on rater anonymity in additive reputation systems and use secure multiparty protocols to collect ratings privately. Ismail et al. [27] present a centralized reputation system that guarantees rater and ratee anonymity by distributing the functionality of monitoring a transaction and collecting the ratings to two different entities. In [28] the same authors extend the distributed scheme of [25] to achieve rater anonymity. It has some similarities with our extended approach but relies on a dedicated external trusted third party. Kinateder and Pearson [29] propose a distributed recommender system[6] on top of a trusted computing platform. The Trusted Computing Group (TCG) [30] is a consortium to develop and promote standards for trusted computing building blocks. In contrast to the observer approach presented, later on the TCG tries to establish a minimal trusted hardware platform. On top of this platform trusted software agents should be executed side by side with untrusted software.

Recommender or collaborative filtering systems [31] are strongly related to reputation systems especially in an information sharing and dissemination scenario. However, the research in this area focuses on finding algorithms that allow to predict additional topics or products a new user might like based on a database about user preferences and recommendations. Privacy in recommender systems is an accepted issue [32]. Solutions are data mining techniques that handle only aggregated data and not individual records [33].

VII. Conclusion

We have presented two reputation systems for mobile information dissemination networks. The first and simple approach provides no anonymity for rater and ratee. The emphasis of our extended reputation system is to provide as much privacy for the participants as possible. The current discussion about the usage of RFID tags in consumer market [34] shows that privacy in information systems achieves more and more public attention.

Only a distributed reputation system with local storage can give the user control over her reputation information. Therefore, we have introduced an observer as local trusted entity that assures correctness of anonymous ratings. No central authority is involved in the processing of ratings.

This approach fulfills most of the requirements that we have made up for a privacy preserving reputation system. The only thing that we have skipped intentionally are secure negative ratings. This would lead to a much more complicated transaction between rater and ratee and does not fit our scenario.

However, our future work will concentrate on this topic and on how to limit the information the observer learns during its operations. Investigating alternative solutions to the linkability problem of showing reputation lists is also an open issue. We have already started to implement simulations of the proposed schemes to evaluate their efficiency.

[6]Although the title of their work claims that they deal with a reputation system, reputation is only used to weight recommendations in their approach.

Acknowledgment

This work was sponsored in part by the Deutsche Forschungsgemeinschaft (DFG) as part of the PhD program "Enabling Technologies for Electronic Commerce".

References

[1] C. Ellison, "SPKI/SDSI Certificate Documentation - Website." [Online]. Available: http://world.std.com/~cme/html/spki.html

[2] T. ElGamal, "A Public Key Cryptosystem and a Signature Scheme based on Discrete Logarithms," in *Advances in Cryptology - CRYPTO'84, Proceedings*, ser. Lecture Notes in Computer Science, G. R. Blakley and D. Chaum, Eds., vol. 196. Springer, 1985, pp. 10–18.

[3] B. R. Waters, E. W. Felten, and A. Sahai, "Receiver Anonymity via Incomparable Public Keys," in *Proceedings of the 10th ACM conference on Computer and communication security*. ACM Press, 2003, pp. 112–121.

[4] D. Chaum and E. van Heyst, "Group Signatures," in *EUROCRYPT'91, Proceedings*, ser. Lecture Notes in Computer Science, vol. 547. Springer, 1991, pp. 257–265.

[5] J. Camenisch, *Group Signature Schemes and Payment Systems Based on the Discrete Logarithm Problem*. Hartung-Gorre Verlag, 1998.

[6] A. Heinemann, "The iClouds Project - Website." [Online]. Available: http://iClouds.tk.informatik.tu-darmstadt.de

[7] A. Heinemann, J. Kangasharju, F. Lyardet, and M. Mühlhäuser, "iClouds – Peer-to-Peer Information Sharing in Mobile Environments," in *Euro-Par 2003. Parallel Processing, 9th International Euro-Par Conference*, ser. Lecture Notes in Computer Science, H. Kosch, L. Böszörményi, and H. Hellwagner, Eds., vol. 2790. Klagenfurt, Austria: Springer, 2003, pp. 1038–1045.

[8] A. Heinemann, J. Kangasharju, F. Lyardet, and M. Mühlhäuser, "Ad Hoc Collaboration and Information Services Using Information Clouds," in *Proceedings of the 3rd Workshop on Applications and Services in Wireless Networks, (ASWN 2003)*, T. Braun, N. Golmie, and J. Schiller, Eds. Bern, Switzerland: Institute of Computer Science and Applied Mathematics, University of Bern, 2003, pp. 233–242.

[9] T. Straub and A. Heinemann, "An Anonymous Bonus Point System For Mobile Commerce Based On Word-Of-Mouth Recommendation," in *Applied Computing 2004. Proceedings of the 2004 ACM Symposium on Applied Computing*, L. M. Liebrock, Ed. New York, NY, USA: ACM Press, 2004, pp. 766–773.

[10] "IETF Working Group: Mobile Ad-hoc Networks - Website." [Online]. Available: http://www.ietf.org/html.charters/manet-charter.html

[11] M. Voss, "Privacy Preserving Reputation Systems," in *Proceedings of 19th IFIP International Information Security Conference (SEC2004), Toulouse, France*. Kluver Academic Publishers, 2004, accepted for publication.

[12] D. Chaum and T. P. Pedersen, "Wallet Databases with Observers," in *Advances in Cryptology - CRYPTO '92*, ser. Lecture Notes in Computer Science, E. F. Brickell, Ed., vol. 740. Berlin, Heidelberg: Springer, 1993, pp. 89–105.

[13] P. Persiano and I. Visconti, "An Anonymous Credential System and a Privacy-Aware PKI," in *Information Security and Privacy, 8th Australasian Conference, ACISP 2003, Wollongong, Australia, July 9-11, 2003, Proceedings*, ser. Lecture Notes in Computer Science, R. Safavi-Naini and J. Seberry, Eds., vol. 2727. Springer, 2003.

[14] M. Papadopouli and H. Schulzrinne, "Seven Degrees of Separation in Mobile Ad Hoc Networks," in *Proceedings of the IEEE Conference on Global Communications (GLOBECOM)*. San Francisco, USA: IEEE Computer Society, 2000, pp. 1707–1711.

[15] G. Kortuem, J. Schneider, D. Preuitt, T. G. C. Thompson, S. Fickas, and Z. Segall, "When Peer-to-Peer Comes Face-to-Face: Collaborative Peer-to-Peer Computing in Mobile Ad-hoc Networks," in *Proceedings of the First International Conference on Peer-to-Peer Computing (P2P'01)*. Linköping, Sweden: IEEE Computer Society, 2001, pp. 75–94.

[16] T. Schwotzer and K. Geihs, "Shark – a System for Management, Synchronization and Exchange of Knowledge in Mobile User Groups," *Journal of Universal Computer Science*, vol. 8, no. 6, pp. 644–651, 2002.

[17] C. Becker, M. Bauer, and J. Hähner, "Usenet-on-the-fly - Supporting Locality of Information in Spontaneous Networking Entironments," *Liscano, Ramiro (ed.); Kortuem, Gerd (ed.): Workshop on Ad Hoc Communications and Collaboration in Ubiquitous Computing Envirnments*, November 16-20 2002.

[18] C. Seitz and M. Berger, "Towards an Approach for Mobile Profile Based Distributed Clustering," in *Proceedings of the 9th International Euro-Par Conference, (Euro-Par 2003)*, ser. Lecture Notes in Computer Science, H. Kosch, L. Böszörményi, and H. Hellwagner, Eds., vol. 2790. Klagenfurt, Austria: Springer, 2003, pp. 1109–1117.

[19] R. P. and Z. R., "Reputation Systems," *Communications of the ACM*, vol. 43, pp. 45–48, 2000.

[20] L. Mui, M. Mohtashemi, and A. Halberstadt, "Notions of Reputation in Multi-Agents Systems: a Review," in *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*. ACM Press, 2002, pp. 280–287.

[21] "ebay - Website." [Online]. Available: http://www.ebay.com

[22] R. P. and Z. R., "Trust Among Strangers in Internet Transactions: Empirical Analysis of ebay's Reputation System. The Economics of the Internet and E-Commerce," *Advances in Applied Microeconomics*, vol. 11, 2002.

[23] S. Buchegger and J.-Y. L. Boudec, "Performance analysis of the CONFIDANT protocol: Cooperation of nodes — fairness in dynamic ad-hoc networks," in *Proceedings of IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*. Lausanne, CH: IEEE, June 2002. [Online]. Available: citeseer.ist.psu.edu/buchegger02performance.html

[24] P. Michiardi and R. Molva, "Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks," in *Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security*. Kluwer, B.V., 2002, pp. 107–121.

[25] W. L. D. Fahrenholtz, "Transactional Security for a Distributed Reputation Management System," in *Proceedings of the 3rd International Conference on Electronic Commerce and Web Technologies*, ser. Lecture Notes in Computer Science, G. Q. K. Bauknecht, A. Min Tjoa, Ed., vol. 2455. Springer, 2002, pp. 214–223.

[26] E. Pavlov, J. S. Rosenschein, and Z. Topol, "Supporting Privacy in Decentralized Additive Reputation Systems," in *Proceedings of the Second International Conference on Trust Management, Oxford, United Kingdom, March 2004*, ser. Lecture Notes in Computer Science, vol. 2995. Springer, 2004, pp. 108–119.

[27] R. Ismail, C. Boyd, A. Jøsang, and S. Russel, "Strong Privacy in Reputation Systems," in *Proceedings of the 4th International Workshop on Information Security Applications, WISA 2003*, ser. Lecture Notes in Computer Science, vol. 2908. Springer, 2004.

[28] R. Ismail, C. Boyd, A. Jsang, and S. Russel, "Private Reputation Schemes for P2P Systems," in *Proceedings of the Second International Workshop on Security in Information Systems (WOSIS-2004)*. Auerbach Publications, 2004.

[29] M. Kinateder and S. Pearson, "A Privacy-Enhanced Peer-to-Peer Reputation System," in *Proceedings of the 4th International Conference on Electronic Commerce and Web Technologies: EC-Web 2003, Prague, Czech Republic, September 2003*, ser. Lecture Notes in Computer Science, K. Bauknecht, A. M. Tjoa, and G. Quirchmayr, Eds., vol. 2738. Springer, 2003.

[30] "Trusted Computing Group - Website." [Online]. Available: http://www.trustedcomputing.org

[31] J. S. Breese, D. Heckerman, and C. Kadie, "Empirical Analysis of Predictive Algorithms for Collaborative Filtering," in *Proceedings of the Fourteenth Annual Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, 1998, pp. 43–52. [Online]. Available: citeseer.ist.psu.edu/breese98empirical.html

[32] N. Ramakrishnan, B. J. Keller, B. J. Mirza, A. Y. Grama, and G. Karypis, "Privacy Risks in Recommender Systems," *IEEE Internet Computing*, vol. 5, no. 6, pp. 54–62, 2001.

[33] R. Agrawal and R. Srikant, "Privacy-preserving Data Mining," in *Proc. of the ACM SIGMOD Conference on Management of Data*. ACM Press, May 2000, pp. 439–450.

[34] "RFID Privacy Workshop @ MIT - Website." [Online]. Available: http://www.rfidprivacy.org