A Framework to Provide Anonymity in Reputation Systems[*][†]

Hugo Miranda Luís Rodrigues
Universidade de Lisboa

Abstract

In ubiquitous networks, the multiple devices carried by an user may unintentionally expose information about her habits or preferences. This information leakage can compromise the users' right to privacy. A common approach to increase privacy is to hide the user real identity under a pseudonym. Unfortunately, pseudonyms may interfere with the reputation systems that are often used to assert the reliability of the information provided by the participants in the network.

This paper presents a framework for combining anonymity with reputation and shows that it can be configured to provide a desired degree of balance between these two conflicting goals. The proposed solution leverages on well-known cryptographic techniques, such as public key infrastructure and blind signatures.

1 Introduction

The actions of an user in an ubiquitous network can expose different information such as newspaper of preference or current medication. A systematic analysis of this data can be used to put at risk the users' privacy. The goal of algorithms providing anonymity is to increase privacy by hiding the link between the information leaked to the network and the true identity of the user. Pseudonyms are an effective mechanism to support anonymity, but need to be changed frequently, to make harder to map the pseudonym in the real identity of the user [10]. Anonymity has been an active field of research, for example in web browsing [9], ad hoc routing [1] and e-commerce [10].

Recently, the literature has been emphasising the advantages of using reputation systems for the establishment of trust relationships in large scale networks [1, 2, 3, 6, 8]. Reputation systems provide the mechanisms to capture and spread information about which users are reliable and which users are unreliable.

Both reputation and anonymity are desirable features for large scale, ubiquitous networks. Unfortunately, reputation systems may become useless if users can frequently change their identities, invalidating any reputation information referring to a given pseudonym. On the other hand, the possibility of a user

[*] This work was partially funded by FCT project MICAS - MIddleware para sistemas adaptáveis ao contexto, POSI/EIA/60692/2004 through POSI and FEDER.

[†] An extended version of this paper appeared in Mobiquitous 2006.

to frequently change her pseudonym is a fundamental premise of anonymous networks.

This paper leverages on a number of established cryptographic techniques, like asymmetric cryptography and blind signatures, to implement a compromise solution that combines reputation and anonymity. This algorithm, named RuP (from Reputation using Pseudonyms) hides the real identity of the user, even from trusted third-parties. Additionally, our solution also prevents users from impersonating other users, a problem frequently ignored in reputation systems. Besides assuming that each reputation information is uniquely identified, our architecture makes no particular assumptions about the operation of the reputation system. Therefore, our scheme can be combined with a wide range of existing reputation systems, including the systems described in [3, 6, 8].

Blind signatures have been previously used to provide anonymity in access controlled wireless networks. In [4], they are used to mask the temporary identification numbers used by the clients of the network. Clients are authorised by an authentication server. In comparison with [4], our paper describes a method providing a number of additional features. RuP allows the server to verify additional constraints that may be imposed on the clients and provides to users a mean to prove their identity to others, without contacting the authentication server. Novel in our paper is also the possibility of securely transferring reputation information between different random pseudonyms without loosing anonymity. Furthermore, our model does not present the vulnerabilities of [4] exposed in [7].

1.1 Threat model

RuP addresses two separate categories of threats. Those to the anonymity of the user and those to the reputation system. We assume that malicious users (either internal or external to the network) try to uncover the real identity of the users by correlating known information about a participant with the actions performed on behalf of a pseudonym so that an unique mapping can be established. If the information gathered using one pseudonym is not sufficient, the malicious user attempts to track sequences of pseudonyms used by the same participant. Mechanisms that prevent devices from leaking information (such as MAC addresses) that would allow to uncover the real identity of the user are outside the scope of this paper and can be found elsewhere (for example [5]).

It is assumed that malicious users try to tamper the reputation system by forging or duplicating reputation information. Users may attempt to impersonate others either to benefit from their good reputation or to degrade it. Finally, it is also assumed that malicious users try to acquire multiple personalities to avoid punitive measures applied as a result of their bad reputation. Similarly to many reputation systems, it is assumed that users do not collude for getting additional benefits from the network.

1.2 Cryptographic building blocks

RuP relies on a number of established security algorithms. This paper assumes that the reader is familiar with the basics of digital cryptography like private/public key pairs, digital signatures and certification authorities. The interested reader may find information on the subject in surveys such as [11].

This paper also makes use of a procedure known as “blind signing”, that allows some principal to sign a document that she was not able to read because it was encrypted. However, the signature holds even when the encryption is removed. Different applications have been described for blind signatures. In the one of interest for RuP, principals use probabilities to assert that the document they sign satisfies some constraints, although they never become fully aware of its content. The signing principal defines the minimum number n of documents to be presented by the requesting principal. The requesting principal creates n such documents, which should have a different content, although all of them must respect the constraints agreed between the principals. The documents are encrypted for blind signing by the requesting principal, using a different key for each. The signing principal randomly selects $n-1$ of the documents and asks for the keys necessary to decrypt them. These $n-1$ documents are used to probabilistically confirm that the constraints are also satisfied in the remaining one. The signing principal blindly signs the document that was not decrypted.

2 RuP

In this paper we assume that reputation information refers to users and not to individual devices, i.e., users are the solely relevant source of (good or bad) reputation information. Users are identified by their unforgeable real identity and by pseudonyms, used to preserve anonymity. To combine anonymity with reputation, RuP never requires the disclosure of the real identity of the owner of a given pseudonym.

Reputation information is stored in a repository which is easily accessible for storage and query by all users. The reputation repository may be implemented in a centralised or decentralised manner; this option is not relevant for our model. Reputation information refers to pseudonyms.

Central to our architecture is the concept of *Certified Pseudonym*. At any given point in time, each user is uniquely identified in the network by a certified pseudonym valid only for a predefined time interval. A certified pseudonym has the following properties: *i*) only one user can claim its ownership, *ii*) at each moment, each user owns at most one valid certified pseudonym.

Certified pseudonyms are issued by a *pseudonym certification authority* (PCA), trusted by all participants. The PCA is implemented by a Certification Authority (CA). The format of certified pseudonyms can follow closely the format of standard public key infrastructure certificates like X.509. In this paper, we only describe the fields of the certificate that are relevant for the exposition. The PCA does not need to be permanently reachable from all nodes. Message exchanges between users and the PCA are occasional and can be scheduled in advance. The anonymity of the actions of the users is preserved even from the PCA, who although simultaneously accessing both the user' pseudonym and her real identity, is not able to map one in the other.

In RuP, time is partitioned in time slots. Certified pseudonyms are associated to exactly one time slot and must be issued before the slot begins. This paper uses T^i to denote time slot i and p^i_A to represent the pseudonym of user A in time slot i.

2.1 Certified pseudonyms and PCA

A certified pseudonym issued by a PCA in the network on behalf of principal A is defined by $\mathbb{C}^i_{p_A} = \{p_A, K_{u_{p_A}}, i\}$. The certified pseudonym creates an association between a pseudonym p_A and a public key $K_{u_{p_A}}$, both defined by the user. The certified pseudonym is valid in the time slot T^i. Note that the real identity of principal A is not present in the certified pseudonym.

Users are required to present their unforgeable real identity when requesting certified pseudonyms from the PCA. It is a responsibility of the PCA to ensure the real identity of the user and that each user owns at most one certificate for each time slot. To reduce the number of contacts to the PCA, the user can request multiple certified pseudonyms at once, one for each forthcoming time slot.

The mapping between the real identity and the pseudonym is hidden from the PCA by applying the probabilistic algorithm for blind signatures described in Section 1.2. Here, the user is the requesting principal while the PCA blindly signs the pseudonym. Each encrypted certificate tentatively presented by the user for signing should have a different public/private key pair and pseudonym. The PCA blindly signs one of the certificates, creating a certified pseudonym after verifying that the following constraints are met: *i*) the user can prove the ownership of the real identity that is presented; *ii*) the time slot is equal on all $n - 1$ disclosed certificates; *iii*) the requested time slot has not started and *iv*) a certificate valid for the same time slot has not been previously issued for the same user.

The pseudonym in the blindly signed certified pseudonym is used as the identity of the principal for time slot T^i. The certificate exhibits the following properties: *i*) the PCA is not aware of the pseudonym of the user; *ii*) the user is unable to obtain another pseudonym for the same time slot. Therefore, anonymity is preserved, although users are unable to impersonate others or assume multiple identities.

2.2 Transference of reputation information

Reputation information describing the level of trust of user A on user B follows the format of the reputation system in use, hereafter abstractly represented as $\tau_{p^i_A \to p^i_B}$. It is assumed that reputation information includes a field providing for a unique identification, for example, a randomly generated serial number. RuP further requires that user A signs the reputation information[1] with the private key of its current certified pseudonym and that she provides to user B a copy of her certified pseudonym. Therefore, we define a reputation declaration as $\mathbb{T}^i_{p^i_A \to p^i_B} = \{\mathbb{C}^i_{p^i_A}, \mathbb{S}_{K_{r_{p^i_A}}}(\tau_{p^i_A \to p^i_B})\}$. The combination of $\mathbb{T}^i_{p^i_A \to p^i_B}$ and $\mathbb{C}^i_{p^i_B}$, allows any participant to: *i*) assert that the user presenting the reputation item used p^i_B at time slot T^i, by challenging $\mathbb{C}^i_{p^i_B}$; *ii*) assert that the reputation item was not self-generated by user B by confirming that $\mathbb{C}^i_{p^i_B}$ and $\mathbb{C}^i_{p^i_A}$ are different; and *iii*) confirming that the reputation information was not forged by verifying the signature of $\tau_{p^i_A \to p^i_B}$. We emphasise that $\mathbb{T}^i_{p^i_A \to p^i_B}$ does not expire with

[1] The signature of some text x with the private key of some user A using pseudonym p is denoted by $\mathbb{S}_{K_{r_{p_A}}}(x)$.

time slot T^i. The same verifications can be performed later, provided that user B stored the private key associated to its past pseudonym. Furthermore, the verification does not require the participation of user A.

Presenting reputation information obtained using an old pseudonym, implicitly reveals two pseudonyms of the same user. We now describe a scheme that allows a user to transfer reputation information from one pseudonym to another, without disclosing this link or her real identity. This is a two step procedure (anonymity and endorsement). In the anonymity step, the user and the PCA cooperate to strip the reputation information of references to the old pseudonyms. In the endorsement step, the user and the PCA bind the reputation information to the new pseudonym.

In the anonymity step, the PCA is required to blindly sign a *transferable reputation voucher* (TRV). Depending on the reputation system and user preferences, the TRV may either represent one reputation item or subsume multiple reputation items in possession of the user. The TRV does not include references to pseudonyms. To be untraceable, the granularity of the reputation value should be augmented (for example by rounding numeric values). In the transference of reputation information, the blind signature is used to hide the serial number of the TRV from the PCA.

For each reputation declaration to be subsumed in the TRV, the PCA verifies that *i*) the user can claim its ownership and *ii*) it has not been presented before for revalidation (using the unique serial number of the reputation declaration). The PCA blindly signs one TRV after probabilistically asserting that the reputation value in the TRV is consistent with the subsume operation defined by the reputation system and with the mechanisms in use to prevent the TRV from being traceable.

The endorsement step prevents the user from creating multiple copies of the TRV. At some random later time, the user presents to the PCA her new certified pseudonym $\mathbb{C}^j_{p_B}$ and the TRV. The PCA confirms that the serial number included in the TRV has never been endorsed in the past and signs a new reputation item, $\mathbb{T}^j_{\_\rightarrow p^j_B} = \{\mathbb{C}_{CA}, \mathbb{S}_{K_{r_{CA}}}(\tau_{\_\rightarrow p^j_B})\}$ with the reputation value in the TRV. Because it was signed by a trusted third party (the PCA), all participants can safely accept $\mathbb{T}^j_{\_\rightarrow p^j_B}$ as valid if user B proves the ownership of pseudonym p^j_B. Subsequent transfers of the reputation information to another pseudonym follow a similar procedure, with the user presenting $\mathbb{T}^j_{\_\rightarrow p^j_B}$ at the anonymity step.

The transference of reputation information hides the link between pseudonyms as long as the PCA is unable to establish an association between the anonymity and endorsement steps. This can only be achieved if the same value of the reputation information is frequently presented to the PCA by multiple users. We anticipate that this would be the case in large-scale ubiquitous systems.

3 Balancing reputation and anonymity

The duration of each time slot, which dictates the validity of a pseudonym, trades off the efficiency of the reputation system with the efficiency of the measures to preserve anonymity. Recall from the threat model presented in

Section 1.1 that the true identity of an user will be as much at risk as the number of actions performed using one, or a traceable sequence of pseudonyms, increases. Therefore, the smaller the validity of a pseudonym, the harder will be for a malicious user to infer the real identity of a participant. On the other hand, many reputation systems implement a transitive function that allows participants to infer their trust on unknown nodes from a trust value provided by a third party. A fundamental requisite for the application of this function is that nodes develop relationships with each other. In RuP, the frequent change of pseudonym limits the applicability of this function by constraining the lifetime of the relations to the duration of the pseudonyms. Furthermore, the transference of reputation information between pseudonyms is a voluntary process driven by the user. It is possible for a misbehaving user to become detached from bad reputation associated to one of her pseudonyms by not converting it for a new pseudonym.

We assume that users will not revalidate information describing bad reputation. Therefore, malicious users are more likely to present a lower amount of revalidated reputation information. However, since RuP prevents the tracking of sequences of pseudonyms, a malicious user will be indistinguishable from a user that recently joined the network. In some models, for example, for routing in ad hoc networks [2], it is possible to evaluate the reputation of an user by the lack of past references if the time at which the user joined the network is known. In such scenarios, we propose to add a field to the certified pseudonyms indicating the time slot at which the user joined the network. An undesirable side effect of this field is the disclosure of additional information that could lead to reveal the identity of the user.

RuP requires the consumption of additional resources at the devices. Both bandwidth (due to the transference of certificates and challenges) and computational power are expected to increase. However, it should be noted that the most resource demanding operations, like acquiring a new pseudonym and transferring reputation information can be performed off-line (with respect to the ubiquitous network) by personal workstations. Asymmetric cryptography operations to be performed by ubiquitous devices can follow common practice and be applied over small digests of the messages.

4 Conclusions

This paper has described RuP, an algorithm that allows to put together two desirable properties for ubiquitous networks: reputation and anonymity. RuP relies on a certification authority. Its role is to authenticate the participants, assuring that they can not impersonate others or acquire multiple identities. The framework uses pseudonyms to preserve the anonymity of the users. Users are not required to reveal their pseudonyms at any point, even for the certification authority.

When integrated with a reputation system, RuP provides the following properties: *i*) it prevents users from detaching from bad reputation while a pseudonym is valid, for example, by changing their identity or by impersonation, *ii*) it provides anonymity of the users, making it more suitable and appealing in some application domains and *iii*) allows users to continue to benefit from the positive reputation acquired with different pseudonyms.

References

[1] A. Boukerche, K. El-Khatib, Li Xu, and L. Korba. Anonymity enabling scheme for wireless ad hoc networks. In *Works. of the Global Telecommunications Conference, GlobeCom 2004*, pages 136–140, 2004.

[2] S. Buchegger and J.-Y. Le Boudec. Nodes bearing grudges: towards routing security, fairness, and robustness in mobile ad hoc networks. In *Proc. of the 10th Euromicro Work. on Parallel, Distributed and Network-based Processing*, pages 403–410, 2002.

[3] V. Cahill et al. Using trust for secure collaboration in uncertain environments. *IEEE Pervasive Computing*, 2(3):52–61, 2003.

[4] Q. He, D. Wu, and P. Khosla. The quest for personal control over mobile location privacy. *IEEE Communications*, 42(5):130–136, 2004.

[5] M. Jacobsson and I. Niemegeers. Privacy and anonymity in personal networks. In *Proc. of the 3rd IEEE Int'l Conf. on Pervasive Computing and Communications Work. (PerCom 2005)*, pages 130–135, 2005.

[6] J. Keane. Trust based dynamic source routing in mobile ad hoc networks. Master's thesis, University of Dublin, 2002.

[7] R.C.-W. Phan. Security limitations of an authorized anonymous id-based scheme for mobile communication. *IEEE Communications*, 43(5):149–153, 2005.

[8] A. A. Pirzada, A. Datta, and C. McDonald. Propagating trust in ad-hoc networks for reliable routing. In *Proc. of the 2004 Int'l Work. on Wireless Ad-hoc Networks (IWWAN'04)*, 2004.

[9] M. K. Reiter and A. D. Rubin. Crowds: anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1):66–92, 1998.

[10] M. Schmidt. Subscriptionless mobile networking: anonymity and privacy aspects within personal area networks. In *Proc. of the Wireless Comm. and Networking Conference (WCNC2002)*, volume 2, pages 869–875, 2002.

[11] P. Veríssimo and L. Rodrigues. *Distributed Systems for System Architects*, volume 1 of *Advances in Distributed Computing and Middleware.* Kluwer Academic Publishers, 2001.