# TATA: Towards Anonymous Trusted Authentication

Daniele Quercia, Stephen Hailes, and Licia Capra

Department of Computer Science, University College London, London, WC1E 6BT, UK
{D.Quercia, S.Hailes, L.Capra}@cs.ucl.ac.uk

**Abstract.** Mobile devices may share resources even in the presence of untrustworthy devices. To do so, each device may use a computational model that on input of reputation information produces trust assessments. Based on such assessments, the device then decides with whom to share: it will likely end up sharing only with the most trustworthy devices, thus isolating the untrustworthy ones. All of this is, however, theoretical in the absence of a general and distributed authentication mechanism. Currently, distributed trust frameworks do not offer an authentication mechanism that supports user privacy, whilst being resistant to "Sybil attacks". To fill the gap, we first analyze the general attack space that relates to anonymous authentication as it applies to distributed trust models. We then put forward a scheme that is based on blinded threshold signature: collections of devices certify pseudonyms without seeing them and without relying on a central authority. We finally discuss how the scheme tackles the authentication attacks.

## 1 Introduction

To produce reliable assessments, distributed trust frameworks must be able *uniquely* to authenticate their users. To see why, consider the following example. Samantha's and Cathy's devices exchange recommendations about shops in their local area. After the exchange, as they know (have authenticated) each other, Samantha's device values Cathy's recommendations based on Cathy's reputation as recommender (i.e., whether her past recommendations have been useful), and vice versa. If it was able to easily generate a new pseudonym, Cathy's device could produce fake recommendations without being traceable. In general, to trace past misbehavior, users should not be able easily to change their pseudonyms - ideally, each user should have one and only one pseudonym.

On the other hand, to protect their privacy, users should *anonymously* authenticate each other, i.e., authenticate without revealing real identities. For example, Samantha may wish to buy kinky boots. She thus uses her mobile device to collect the most useful recommendations from the most trustworthy sources. The recommendation sharing service requires devices to use trust models that, in turn, require users to authenticate. Thus, Samantha's device has to authenticate in order to ask for recommendations; as the subject (kinky boots) is sensitive, the device authenticates itself without revealing Samantha's identity (anonymously).

Existing research in distributed reputation-based trust models does not offer any general solution for *unique* and *anonymous* authentication without relying on a central authority. Some distributed trust models [1] allow the use of anonymous pseudonyms

K. Stølen et al. (Eds.): iTrust 2006, LNCS 3986, pp. 313–323, 2006.

that, however, suffer from "Sybil attacks" [7]. Others tackle such attacks, but mostly with either centralized solutions [15] or approaches that only apply to limited scenarios [11] [12] [13] [17].

Our contribution lies in: firstly, systematically analyzing the general attack space that relates to anonymous authentication as it applies to distributed trust models; secondly, proposing a scheme that is decentralized, yet general enough to be applied to most of the existing trust models. More specifically, the scheme meets appropriate security requirements and supports desirable features. Security requirements include: (i) *anonymity* to prevent privacy breaches; (ii) *non-repudiation* to prevent false accusation; (iii) *unique identification* to avoid attacks caused by disposable pseudonyms; (iv) *pseudonym revocation* to cope with stolen pseudonyms. Desirable features include: (i) *general applicability*, in that our scheme is general-purpose so that any reputation-based system benefits from it; (ii) *off-line authentication* between two users without relying on anyone else; (iii) *distributed pseudonym issuing*, in that valid pseudonyms are issued without relying on a central authority.

The remainder of the paper is structured as follows. Section 2 discusses related work. Section 3 introduces a scenario that we will use to exemplify our model. Section 4 describes the attacks that relate to anonymous authentication. Starting from both those attacks and the general problem space, section 5 draws security requirements and desirable features for a protection scheme. Section 6 details our proposition and section 7 critically analyzes how it meets the security requirements and supports the desirable features. Section 8 concludes.

## 2 Related Work

Over the course of nearly five years, cooperation and authentication have begun to diverge: authentication has relied on central authorities, while cooperation has migrated to decentralized solutions. Only recently, authentication for cooperative mechanisms started to be decentralized.

Disposable pseudonyms facilitate anonymity, yet hinder cooperation in the absence of a central authority. To see why, consider a collection of actors cooperating. If each actor authenticates himself with an anonymous pseudonym, then he does not have to disclose his real identity and, thus, he can remain anonymous. However, an actor may profit from ease of creating pseudonyms. For example, an actor may authenticate himself with a pseudonym, misbehave, create a new pseudonym, authenticate himself with the new pseudonym (pretending to be new actor), and misbehave again. As a result, the actor misbehaves without being traceable. Resnick and Friedman [15] formally laid down such a problem, presenting a game theoretical model for analyzing the social cost of allowing actors to freely change identities. They concluded that, if actors generate pseudonyms by themselves, all unknown actors should be regarded as malicious. To avoid mistreating all unknown actors, they proposed the use of free but unreplaceable (once in a lifetime) pseudonyms, which a *central* authority certifies through blind signature. A couple of years later, Doucer put similar ideas to test in P2P networks. He discussed the attacks resulting from P2P users who could use multiple identities and named them "Sybil attacks" [7]. He concluded with a critical take on decentralized au-