Science of Computer Programming 111 (2015) 300–317

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico

Formal analysis of privacy in Direct Anonymous Attestation schemes ☆

Ben Smyth [a,*], Mark D. Ryan [b], Liqun Chen [c]

[a] *Mathematical and Algorithmic Sciences Lab, France Research Center, Huawei Technologies Co. Ltd., France*
[b] *School of Computer Science, University of Birmingham, UK*
[c] *HP Laboratories, Bristol, UK*

ARTICLE INFO

Article history:
Received 29 May 2014
Received in revised form 4 March 2015
Accepted 9 April 2015
Available online 27 April 2015

Keywords:
Accountability
Applied pi calculus
Direct Anonymous Attestation
Privacy
Trusted computing

ABSTRACT

This article introduces a definition of privacy for Direct Anonymous Attestation schemes. The definition is expressed as an equivalence property which is suited to automated reasoning using Blanchet's ProVerif. The practicality of the definition is demonstrated by analysing the RSA-based Direct Anonymous Attestation protocol by Brickell, Camenisch & Chen. The analysis discovers a vulnerability in the RSA-based scheme which can be exploited by a passive adversary and, under weaker assumptions, corrupt issuers and verifiers. A security fix is identified and the revised protocol is shown to satisfy our definition of privacy.

1. Introduction

Trusted computing allows commodity computers to provide cryptographic assurances about their behaviour. At the core of the architecture is a hardware device called the Trusted Platform Module (TPM). The TPM uses shielded memory to store cryptographic keys, and other sensitive data, which can be used to achieve security objectives, in particular, the chip can measure and report its state, and authenticate. These security objectives assume that a TPM's shielded memory protects keys and TPMs are said to be compromised if this assumption does not hold (see Tarnovsky [64] for a hardware attack that successful extracts keys from shielded memory).

Cryptographic operations, by their nature, may reveal a platform's identity. As a consequence, the TPM has been perceived as a threat to privacy by some users, for example, see Stallman [62,63] and Anderson [4,5]. In an attempt to overcome these privacy concerns, Brickell, Camenisch & Chen [14] have introduced Direct Anonymous Attestation (DAA), a historical account of DAA's development is presented by Brickell, Camenisch & Chen [15].

Direct Anonymous Attestation enables a platform to authenticate in a manner that provides privacy and accountability. The concept is based upon group signatures with stronger anonymity guarantees, in particular, the identity of a signer can never be revealed, but certain signatures can be linked (as discussed below) and signatures produced by compromised platforms can be identified. A DAA scheme considers a set of *hosts*, *issuers*, *TPMs*, and *verifiers*; the host and TPM

☆ This article is a based upon [61,60] and an earlier version appeared in [59, Chapter 4]. The ProVerif scripts which support this article are available from the following URL: http://www.bensmyth.com/publications/2012-Direct-Anonymous-Attestation-anonymity-definition/.

[*] Corresponding author at: Huawei, 20 quai du Point du Jour, 92100 Boulogne-Billancourt, France.
E-mail address: research@bensmyth.com (B. Smyth).
URL: http://www.bensmyth.com/ (B. Smyth).

http://dx.doi.org/10.1016/j.scico.2015.04.004
0167-6423/© 2015 Elsevier B.V. All rights reserved.

together form a *trusted platform* or *signer*. DAA protocols proceed as follows. A host requests membership to a *group of signers* managed by an issuer. The issuer authenticates the host as a trusted platform and grants an *attestation identity credential* (occasionally abbreviated *credential*). A verifier can authenticate trusted platforms using signatures produced from such credentials.

Brickell, Chen & Li [17,18] and Chen [29,31] characterise the following security properties[1] for Direct Anonymous Attestation schemes:

- *Anonymity.* The identity of a signer cannot be revealed from a signature.
- *Non-frameability.* An adversary cannot produce a signature associated with an honest TPM.
- *Unforgeability.* Signatures cannot be produced without a TPM.
- *User-controlled linkability.* A signer can control whether her signatures can be detected as being from the same signer.

A signer defines whether her signatures are *linkable* (that is, can be detected as being from the same signer) or *unlinkable* (that is, cannot be detected as being from the same signer) at the time of construction.

Our security properties aim to balance the privacy (anonymity and unlinkability properties) demands of users with the accountability (linkability, non-frameability and unforgeability properties) needs of issuers and verifiers. The distinction between privacy and accountability properties is reflected in our trust model: anonymity and unlinkability assume that two signers are honest, whereas, linkability, non-frameability and unforgeability assume that an issuer is honest. (The issuer must be honest for linkability, since a dishonest issuer can provide an adversary with a new credential for every signature, thereby ensuring that two signatures are never linked.) In addition, DAA schemes must be *correct*: valid signatures can be verified and, where applicable, linked.

Brickell, Camenisch & Chen [14] propose the first concrete instance of a Direct Anonymous Attestation scheme. Their scheme is based upon RSA and support for this scheme is mandated by the TPM specification version 1.2 [65], which has been defined as an ISO/IEC international standard [46]. Moreover, TPM version 1.2 is estimated to have been embedded in over 500 million computers [66] (however, the Trusted Computer Group acknowledges that the opt-in policy – whereby, users must choose to enable the TPM – has hindered development [67], moreover, Martin claims that only 5% of these TPMs have been turned on [53, §6] and we suspect significantly fewer are in active use). Furthermore, the RSA-based DAA scheme has also been included in the ISO/IEC anonymous digital signature standard [47]. A brief review of other DAA schemes appears in Appendix A.

1.1. Contribution

We formalise Direct Anonymous Attestation protocols in the applied pi calculus and present a definition of privacy as an equivalence property which is suited to automated reasoning using ProVerif (Section 4). Informally, the security definition asserts that an adversary cannot distinguish between signatures produced by two distinct signers, even when the adversary controls the issuer and has observed signatures produced by each signer. The application of the definition is demonstrated by analysing privacy in the RSA-based DAA protocol (Section 5). The analysis discovers a vulnerability in the protocol which allows privacy to be violated by a passive adversary and, under weaker assumptions, corrupt issuers and verifiers. A fix is identified, and the revised RSA-based DAA protocol is shown to be secure in the symbolic model. We examine the balance between privacy and accountability offered by DAA and propose extensions to DAA (Section 6): we propose a stronger notion of privacy which is intuitively satisfied by the fixed RSA-based scheme, address an issue which can prevent linkability, and provide some practical guidelines for basenames to help resolve a flaw in unlinkability.

1.2. Related work

In the computational model, Brickell, Camenisch & Chen [14] introduce simulation-based models of security and Brickell, Chen & Li [17,18] propose a game-based security definition; the relationship between simulation-based models and game-based definition is unknown [34, pp. 158]. Bernhard et al. [10] argue that the simulation-based definitions and the game-based definition are insufficient for accountability due to informal handling of identities and propose an alternative game-based security definition, moreover, Bernhard et al. show that the simulation-based model by Chen, Morrissey & Smart [36] is unsatisfiable (for all protocols there trivially exists a distinction between the ideal- and real-world). We consider a symbolic definition for privacy, based upon the game-based definition by Brickell, Chen & Li (we stress that the criticisms from Bernhard et al. relate to the accountability game and not the privacy game, hence, their concerns are not relevant to our work). Backes, Maffei & Unruh [8] formalise an earlier notion of privacy (informally described in [14]) for the RSA-based DAA protocol. This formalisation is tightly coupled with their model of the RSA-based protocol and it is unclear whether other DAA schemes can be analysed or, indeed, how to analyse alternative models of the RSA-based protocol. In addition, their formalisation pre-dates the privacy definitions by Brickell, Chen & Li and considers a conceptually weaker

[1] The necessity for non-frameability was highlighted by Backes, Maffei & Unruh [8] and formalised by Chen [29,31], the remaining properties were formalised by Brickell, Chen & Li [17,18].