An Attack Against Fixed Value Discrete Logarithm Representations

Gergely Alpár[1,2*], Jaap-Henk Hoepman[1,2], and Wouter Lueks[1,2**]

[1] Institute for Computing and Information Sciences,
Radboud University Nijmegen, The Netherlands.
{gergely, jhh, lueks}@cs.ru.nl
[2] TNO Information and Communication Technology, The Netherlands.

Abstract Attribute-based credentials (ABCs) are an important building block of privacy-enhancing identity management. Since non-identifying attributes can easily be abused as the anonymity they provide hides the perpetrator, cryptographic mechanisms need to be introduced to make them revocable. However, most of these techniques are not efficient enough in practice.
ABCs with practical revocation have recently been proposed by Hajny and Malina [5]. Their ABCs make use of different discrete logarithm representations of a fixed value. Although this technique is attractive as the verification of a particular issuer's credentials is easy, it has an intrinsic weakness. Colluding users can efficiently forge new credentials that are indistinguishable from legally issued ones.

Keywords: attribute-based credentials, revocation, cryptanalysis, discrete logarithm representation

1 Introduction

Much research has focused on attribute-based credentials (ABCs) and possibilities for their revocation in particular [2,3,6]. The main problem boils down to the fact that when verifying non-identifying attributes, credentials should be completely unlinkable, whereas revocation requires by definition some sort of (escrowed) linkability of credentials.

Hajny and Malina [5] propose a new construction for revocable ABCs. The basic idea is that credentials are distinct discrete logarithm representations (DL-REPs) of the same fixed value; one of the components in such a DL-REP is used for the revocation. Credentials can only be created by the Issuer, or by the revocation authority (RA) in the proposal, who knows the corresponding secret key. Since the public key is the fixed value, anyone can verify that a DL-REP is indeed a valid credential, while new DL-REPs can only be created by the RA. In this paper we do not discuss how the actual issuing and revocation protocols work as they are not relevant to the weakness described here.

In the following section we describe an attack that shows that it is possible to create new credentials without knowing the RA's private key. In fact, any

* Supported by the research program Sentinels as project 'Mobile IDM' (10522). Sentinels is being financed by Technology Foundation STW, the Netherlands Organisation for Scientific Research (NWO), and the Dutch Ministry of Economic Affairs.

** Supported by the research program Sentinels as project 'Revocable Privacy' (10532).

group of at least two users can create a virtually infinite number of new valid credentials. We note that the authors mention that the security of their scheme relies on the tamper resistance of the card. While this is often an acceptable assumption, it is desirable to offer as much cryptographic protection as possible. Moreover, the damage should be proportional to the committed act; that is, compromising two cards should not undermine the security of the whole system.

2 DL-REPs, and how to create them

The concept of a DL-REP was introduced by Brands [1]. Let G be a finite group of prime order q. If $A = \prod_{i=1}^{k} g_k^{x_k}$, then $(x_1, \ldots, x_k) \in \mathbb{Z}_q^k$ is the DL-REP of A with respect to generators $g_1, \ldots, g_k \in G$.

Let $g_i = g^{y_i}$ for all $i \in \{1, \ldots, k\}$ where $G = \langle g \rangle$. Then

$$A = \prod_{i=1}^{k} g_i^{x_i} = g^{\sum_{i=1}^{k} x_i y_i} \in G.$$

The ability to create a DL-REP of a value with respect to $(g_i)_{i=1,\ldots,k}$ without knowing exponents $(y_i)_{i\in\{1,\ldots,k\}}$ is equivalent to the ability of breaking the discrete logarithm assumption.

However, if one knows two DL-REPs of the same value, one can compute many new DL-REPs. Indeed, let $A = \prod_{i=1}^{k} g_i^{x_i} = \prod_{i=1}^{k} g_i^{\hat{x}_i}$, where $(x_1, \ldots, x_k) \neq (\hat{x}_1, \ldots, \hat{x}_k)$. (That is, at least one entry is different, which also happens to mean that at least two entries are different.) Then we get that $(x_1 - \hat{x}_1, \ldots, \hat{x}_k - \hat{x}_k) \neq (0, \ldots, 0)$. Applying this inequality to the given DL-REPs of A, we get a non-trivial DL-REP (*i.e.,* not all zeros) of 1 in the group G: $1 = \prod_{i=1}^{k} g_i^{x_i - \hat{x}_i}$ and for any $\alpha \in \mathbb{Z}_q$, $1 = 1^\alpha = \prod_{i=1}^{k} g_i^{\alpha(x_i - \hat{x}_i)}$. So, one can generate q different DL-REPs of 1: $(\alpha(x_1 - \hat{x}_1), \ldots, \alpha(x_k - \hat{x}_k))$. Consequently, A has q different DL-REPs (where $\alpha = 0$ and $\alpha = q - 1$ correspond to the initial DL-REPs):

$$A = 1 \cdot A = 1^\alpha \cdot A = \prod_{i=1}^{k} g_i^{\alpha(x_i - \hat{x}_i)} \cdot \prod_{i=1}^{k} g_i^{x_i} = \prod_{i=1}^{k} g_i^{x_i + \alpha(x_i - \hat{x}_i)} \in G.$$

3 The collusion attack

Hajny and Malina [5] propose a DL-REP-based credential revocation technique. The underlying mathematical structure is not a prime group as in [5] but a composite group G in which the discrete logarithm problem is trapdoor one-way (by Okamoto and Uchiyama [7]). This enables the RA to take discrete logarithms which plays an important role during issuance and revocation. Since the trapdoor functionality does not affect the attack, we do not consider it further.

Each attribute has its own key pair in the system. The public key is a value $A_{seed} \in G$, the private key, only known to the RA, consists of the prime factors of the composite modulus $n = p^2 q$. We consider here only one attribute, but the description and the attack can easily be generalised to multiple attributes.

Users are assumed to have smart cards to perform trusted computations for the users. During the issuing phase, two random values x_1, x_2 (from the set of possible exponents) are chosen by the card and x_3 is computed by the RA such that the master key (x_1, x_2, x_3) is a DL-REP of the value A_{seed} with respect to three public generators (g_1, g_2, g_3). The values x_1, x_2 are private to the user, they are not disclosed at issuance or verification. The RA can compute x_3 from $A'_{seed} = g_1^{x_1} g_2^{x_2} \bmod n$ (without learning x_1, x_2) using its secret key (the prime factors of n) as it can compute the discrete logarithm x_3 of A_{seed}/A'_{seed} in G with respect to g_3. Since x_3 is known to the RA this value can be used to revoke the credential.

The authors claim that the system is secure: "Users are stuck with their keys and they are unable to compute other valid keys". However, they do not consider the problem of colluding users in their model. Applying the technique from the previous section, any two colluding users can make as many master keys as they want. Assuming that two users have (different) secret keys (x_1, x_2, x_3) and $(\hat{x}_1, \hat{x}_2, \hat{x}_3)$, for any α (from the possible set of exponents):

$$A_{seed} = g_1^{x_1+\alpha(x_1-\hat{x}_1)} g_2^{x_2+\alpha(x_2-\hat{x}_2)} g_3^{x_3+\alpha(x_3-\hat{x}_3)} \bmod n.$$

Therefore, they can create new DL-REPs $(x_1 + \alpha(x_1 - \hat{x}_1), x_2 + \alpha(x_2 - \hat{x}_2), x_3 + \alpha(x_3 - \hat{x}_3))$ of A_{seed} with respect to (g_1, g_2, g_3), that is, they can forge new valid credentials. The resulting exponent of g_3 is completely random, it cannot be traced back to either of the credentials and can easily be chosen in such a way that it is not yet on any revocation list. Note that the attack itself is independent of the trapdoor functionality, it works in any group where the discrete logarithm problem is hard for the users.

4 Discussion

We have described a collusion attack against the Hajny–Malina scheme [5] that does not even require any exponentiations from users after they learnt the master keys from their smart cards. They can forge a virtually infinite number of untraceable credentials. It is interesting to consider why the security analysis overlooked this security problem.

The authors worked in a special group [7] in which the discrete logarithm (DL) problem is a trapdoor one-way function. Knowing a trapdoor enables a party to compute DL with respect to any base. Since hardness of the DL-REP problem relies on the hardness of the DL problem, the DL-REP problem is also just conditionally hard in this group. In their security proof, the authors state that "it is hard to compute integers a, b such that $1 \equiv g^a c^b \mod n$" where $a, b \neq 0$. There are two problems worth mentioning. First, this weakness has nothing to do with the trapdoor functionality of the group which enables the computation of the discrete logarithm. Second, the theorem does not consider the case when not only g, c, n are known but also at least one pair $(\hat{a}, \hat{b})$ for which $1 \equiv g^{\hat{a}} c^{\hat{b}} \mod n$. In practice, this additional information is often provided; *e.g.*, one master key gets revealed. Assuming that $(\hat{a}, \hat{b})$ is a solution for the equation above, any scalar multiple of it is also a solution. Furthermore, two such linearly independent pairs allows us to compute all possible solutions (a, b).

A possible solution for the problem described here is to use only a restricted set of all DL-REPs of a fixed value as credentials. Such DL-REPs would include some additional non-linear relation among the exponents. Even though users can then compute new DL-REPs, they cannot produce new valid credentials. This notion is closely related to authentication; the RA does not only give a simple DL-REP, but one that authentically belongs to it.

5 Conclusion

A security proof that reduces a problem to another one has to carefully model reality with practical attacks in mind. Even if a proof is correct, a simple—and often practical—modification of the model (input values, in particular) can destroy the security of the scheme. We showed that an additional input vector makes the scheme in [5] insecure and credentials easily forgeable. The assumption that the secret key never leaves the smart card is too strong considering that the corruption of only two cards destroys the security of the whole system.

References

1. Stefan A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy.* MIT Press, Cambridge, MA, USA, 2000.
2. Jan Camenisch and Anna Lysyanskaya. An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In *Advances in Cryptology – EUROCRYPT 2001*, pages 93–118. Springer-Verlag, May 2001.
3. Jan Camenisch and Anna Lysyanskaya. Dynamic accumulators and application to efficient revocation of anonymous credentials. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *LNCS*, pages 101–120. Springer-Verlag, August 2002.
4. F. Garcia, G. de Koning Gans, R. Muijrers, P. Van Rossum, R. Verdult, R. Schreur, and B. Jacobs. Dismantling mifare classic. *Computer Security-ESORICS 2008*, pages 97–114, 2008.
5. Jan Hajny and Lukas Malina. Unlinkable attribute-based credentials with practical revocation on smart-cards. In Stefan Mangard, editor, *Smart Card Research and Advanced Application Conference (CARDIS 2012)*, LNCS. Springer, 2012.
6. Jorn Lapon, Markulf Kohlweiss, Bart De Decker, and Vincent Naessens. Analysis of Revocation Strategies for Anonymous Idemix Credentials. In Bart De Decker, Jorn Lapon, Vincent Naessens, and Andreas Uhl, editors, *Communications and Multimedia Security*, volume 7025 of *Lecture Notes in Computer Science*, pages 3–17. Springer Berlin / Heidelberg, 2011.
7. T. Okamoto and S. Uchiyama. A new public-key cryptosystem as secure as factoring. In *Eurocrypt 1998*, LNCS. Springer Berlin / Heidelberg, 1998.