Security and privacy issues in RFID-enabled bank notes

Mark van Cuijk

June 6, 2009

Abstract

Over the last decades, central banks around the world have invented several optical and tactile security features to prevent counterfeit bank notes to be created by criminals. Central banks in Europe and Japan have investigated the possibility of extending their arsenal by including electronic security features in the form of RFID tags. While these new security features haven't been deployed yet, research has been done on the impact of RFID tags on security and especially the privacy of bearers of bank notes. This papers reflects a privacy protection scheme designed by Juels and Pappu, discusses the attacks formulated by Avoine and adds an analysis that should be consulted for all future work in this area.

1 Introduction

There are several payment models in use around the world, that can be categorized in four groups: cash, credit card, debit card and online payment systems. Although RFID technology can be useful in a other payment models, this paper focusses on cash-based payments. In a cash-based payment model, a central finanical institute issues large amounts of payment tokens, like bank notes or coins. Each token represents a certain amount of value, called the denomination of the token. During a transaction, a payment can be settled by transferring these tokens to another party.

There is no kind of online connection between a cash token and some database or backend system, so validation of the token can only be performed by verifying the physical appearance of the token using visual and tactile senses. To prevent counterfeiting, the recipe for constructing the paper and the ink used in bank notes is kept secret and several security features are added, among others carbon fibers, difficult to recreate holographic images, semi-transparent regions, multi-layer ink and minuscule details that can only be read by using a magnifier. More about visual and tactile security features in Euro bank notes can be found in [2].

Some reported the European Central Bank had planned to incorporate RFID tags as an additional security feature in the bank notes of higher denominations by 2005 [10, 3] and Japan even planned to implant the chips on bank notes entering production [7]. Although both central banks seem to have been talking to Hitachi about integrating the Hitachi μ-chip into bank notes, the reports that these bank notes are being produced — or even designed — have been premature [8].

The most prevalent issues with RFID tags in bank notes are skimming and illicit tracking of tags, e.g. pickpockets may easily detect who is carrying bank notes of high denomination and merchants can cooperate to learn which customers they have in common. These concerns clearly violate the privacy of bearers of cash tokens. Another practical issue is the fact that bank notes are extremelly thin and must be resistent to all kind of physical situations, like folding, tearing and surviving a washing machine. External antennas to RFID tags tend to break in these situations.

1.1 Overview

After this short introduction on payment models and RFID in bank notes, section 2 describes a privacy protection scheme, designed by Juels and Pappu. This chapter is based on [6] and introduces the goals set by the designers of the scheme, describes the scheme itself and the cryptographic operations that are involved. In [1], Avoine describes several attacks that are possible in this scheme, which defeat the privacy goals set by Juels and Pappu and render the scheme useless for any practical use. I made an analysis on the attacks by Avoine to extract those vulnerabilities in the privacy protection scheme that must be addressed in any future work in this area. The analysis is shown in section 4. The paper ends with a conclusion in section 5.

2 Privacy protection scheme by Juels and Pappu

In [6], Juels and Pappu propose a scheme that aims to provide better privacy to bearers of RFID-enabled bank notes, while allowing anyone with visual contact to the bank note to query the RFID chip for verification purposes and allowing a law enforcement agency to legally track interesting bank notes. Their scheme is based on the principle of re-encryption — analogous to that of a mix network — while the entities performing the re-encryption know the plaintext, which is the serial number of a bank note. This section introduces the setting for which the privacy protection scheme was designed, the goals stated by the designers, describes the privacy protection scheme and gives a short introduction on the cryptographic operations that are involved.

Juels and Pappu distinguish four parties that are involved in handling of bank notes. A central bank — in [6] denoted by $\mathcal{B}$ — that is empowered to create and issue banknotes, a law enforcement agency $\mathcal{L}$ that is able to trace the flow of bank notes, the merchant $\mathcal{M}$ and the consumer $\mathcal{C}$. The scheme uses public key algorithms, such that $\mathcal{B}$ has a keypair for signing purposes and $\mathcal{L}$ has a keypair for encryption purposes. Although signed and encrypted information is stored on the RFID tag, no public key operations are performed inside of the tags.

2.1 The goals

Juels and Pappu stated six properties that their scheme should meet, which are repeated in this section.

Consumer privacy Only law enforcement agencies should be able to trace bank notes effectively using information transmitted by RFID tags.

Strong tracing Given interception of valid RFID information from a given bank note, law enforcement should be able to determine the associated serial number.

Minimal infrastructure Consumers should require no special equipment for the handling of bank notes. Merchants and banks should require only relatively inexpensive devices for this purpose and should not require persistant network access.

Forgery resistance A forger must at a minimum make optical contact with a bank note in order to be able to forge a copy bearing the same serial number and other data, e.g., associated digital signatures. A forger should be unable to forge new bank notes with previously unseen serial numbers and should be unable to alter the denomination associated with a given bank note.

Privilege separation So as to prevent wayward or malicious tampering with bank note information, RFID tag data should only be alterable given optical contact with bank notes.

Fraud detection If invalid law enforcement information is written to a RFID tag on a bank note, this should be widely detectable.

2.2 The scheme

The scheme expects RFID tags with two memory cells — named γ and δ — and some static data printed on the bank note to be read optically. The printed data contains a unique serial number S and a signature Σ created with a private key of $\mathcal{B}$ over the serial number S and the denomination of the bank note. This data might be printed in either machine-readable form (like a barcode), human-readable form or both. Memory cell δ contains a random value r and is a keyed read-write cell, protected by a key D. Memory cell γ is readable by anyone, keyed writeable under key D and contains the signature Σ and the serial number S, encrypted with the public key of $\mathcal{L}$ and seeded with the random value r. The access key D can be derived from the signature Σ.

When a merchant $\mathcal{M}$ receives a bank note from a consumer $\mathcal{C}$, he can verify the bank note by optically reading S and Σ and then deriving the access key D from it. Using D, he can read memory cell δ and obtain r. $\mathcal{M}$ can now use the public key of $\mathcal{L}$ to compute the value that should be in memory cell γ and compare it with the value that is actually stored. The privacy aspect of the scheme comes into play at this point; after the bank note has been verified, $\mathcal{M}$ selects a new random value r', stores it in memory cell δ and computes the new values to store in memory cell γ. The contents of the publicly readable γ cell has now changed in such a way that someone without optical access to the bank note is unable to trace is.

The exception to this claim is the law enforcement agency $\mathcal{L}$. Since the value stored in the γ cell is encrypted using the public key of $\mathcal{L}$, law enforcement officers with access to the corresponding private key can decrypt the contents of the γ cell without the need of optical access to the bank note, regardless of the random value r that is being used.

2.3 Cryptography

Section 2 introduces the γ memory cell as containing a value encrypted under the public key of the law enforcement agency $\mathcal{L}$. Juels and Pappu propose in [6] to use the ElGamal encryption scheme in an elliptic curves setting. The ElGamal scheme is setup in a group $\mathcal{G}$ with a generator P. A private key $x \in_R \mathcal{G}$ is a uniform randomly chosen group element; the public key Y is given by xP. Given a message m and a random factor r, ElGamal encryption is given by:

$$(\alpha, \beta) := \varepsilon_Y(m, r) := (m + rY, rP). \qquad (1)$$

Juels and Pappu reason that ElGamal is succeptible to an adaptive chosen-ciphertext attack and therefore propose to implement the Fujisaki-Okamoto scheme [4]. Given a message m, a random factor σ and a public key Y, the Fujisaki-Okamoto scheme combines an asymmetric encryption function $\varepsilon^a_Y(m, \sigma)$, a symmetric encryption function $\varepsilon^s_K(m)$ and two hash function h_1 and h_2 into a hybrid encryption scheme with enhanced security properties, using equation 2. In this equation, $||$ is the concatenation operator.

$$(c_1, c_2) := \varepsilon^{hy}_Y(m, \sigma) = (\varepsilon^a_Y(\sigma, h_1(\sigma||m)), \varepsilon^s_{h_2(\sigma)}(m) \qquad (2)$$

Using equation 1 for $\varepsilon^a_Y(m, \sigma)$ in equation 2 and using a one-time pad for $\varepsilon^s_K(m)$, the encryption function proposed in [6] resolves to

$$((\epsilon_1, \epsilon_2), \epsilon_3) := (r + h_1(r||m)Y, h_1(r||m), h_2(r) \oplus m), \qquad (3)$$

with Y the public key of the law enforcement agency $\mathcal{L}$, $m = \Sigma||S$ the plaintext message, r the random factor and $||$ the concatenation operator.

3 Attacks on the scheme

Avoine discusses the privacy protection scheme designed by Juels and Pappu in [1], where he describes several attacks on the scheme and actually explains how the scheme fails to establish (some of) the intented goals. The first attacks he describes is the pickpocketing attack mentioned in the introduction. The remainder of this section holds a more in-depth description of the other attacks he described.

3.1 Data recovery attack

Avoine discusses that, given the encryption function from equation 3, knowledge of the random factor r is sufficient to obtain the message m and therefore the serial number S. He describes a two-step data recovery attack: first obtain r and than compute m.

Step 1 As part of the verification procedure, the reader has optical contact with the bank note and is able to compute the access key D. It transmits this value to the bank note to read the δ cell and obtain r. Since the forward channel uses much more power, it is relatively easy for an adversary to eavesdrop on this channel to obtain D. Using D, the attacker can contact the bank note to read the δ cell and thus obtain r.

Step 2 The attacker reads the γ cell to obtain $((\epsilon_1, \epsilon_2), \epsilon_3)$ and extracts the message $m = \epsilon_3 \oplus h_2(r)$. Since the message $m = \Sigma||S$ contains the serial number, he is able to retrieve it.

3.2 Ciphertext tracking attack

The third attack that Avoine describes incorporates a cooperation among merchants. Without optical access, one can query the γ cell of bank notes and track the value stored in there, for as long it isn't changed. In particular, Avoine describes a scenario where two merchants store all values they read. By comparing the databases, bank notes that have been in the vicinity of both readers probably mean that one consumer has been at both places. If one of these bank notes was actually used during a transaction with one of the merchants, linking the bank note to a particular person, the other merchant can be very confident that that person has been in his store.

A more aggressive variant of this attack is when a merchant that returns change to a customer can re-encrypt the bank notes with r set to a fixed number r_0. Now, the other merchant — that has knowledge of the value r_0 — can read the publicly readable γ cell and compute Σ_0, using r_0 and the method described in section 3.1. In a genuine situation, a merchant should read the Σ value using optical contact and use it to compute access key D, but in this situation the merchant can use Σ_0 to compute a value D_0 and try to read memory cell δ. If this read succeeds, then he knows that the bank note was returned as change by the first merchant. In fact, there is no backend communication necessary for this attack to work.

Both variants of the attack won't work after the bank note has been re-encrypted. The important aspect in these attacks is that re-encryption only takes place by merchants and not by consumers. The consequence is that re-encryption doesn't take

place very often, such that bank notes keep the same (encrypted) value in the publicly readable γ memory cell.

3.3 Access-key tracking attack

Avoine also describes a side-channel attack that allows an adversary to illicitly track bank notes, even after re-encryption took place. The attack relies on the fact that the access key D to read the δ memory cell is a fixed value. After optically reading the signature Σ once, an adversary can compute and store the access key D. From this moment on, he is able to track the bank note by attempting a keyed read using the stored access key. When the read succeeds — regardless of the returned value r — he can link the bank note to the first encounter. Instead of requiring optical contact once, the attack is also possible by eavesdropping the forward channel, as is done in the data recovery attack described in section 3.1.

3.4 Other attacks

Two more attacks are described in [1]. One is a denial of service attack, where an adversary write two unrelated values to the γ and δ cells. While this doesn't compromise verification correctness or bearer privacy, the attack is mounted very easily and might flood the law enforcement agency with failed verification reports. The other attack is a cookie attack, that allows an adversary to store data in the δ cell. In a genuine situation, this data is completely random, and it is therefore not noticed by legal readers. The information that is stored will be cleared in the next re-encryption round, but may contain some information that can compromise bearer privacy.

4 Attack analysis

This section shows the result of the work I did, based on the privacy protection scheme designed by Juels and Pappu in [6] and the attacks described by Avoine in [1]. I describe how the goals set by Juels and Pappu are defeated by these attacks and I have analysed the attacks to extract the vulnerabilities that made the attacks possible. All future work on this topic should incorporate this list to prevent these kind of attacks to happen.

Refering to the goals in section 2.1, I can split the goals in two groups. The first group contains two goals that are not defeated by the attack: *minimal infrastructure* and *fraud detection*. The four other goals are defeated by the attacks discussed in section 3:

Consumer privacy Although consumer privacy is the reason why Juels and Pappu designed their scheme, it is actually compromised by almost all attacks described by Avoine. The data recovery attack allows an adversary to recover the serial number of a bank note without optical contact, the ciphertext tracking attack allows merchant to cooperate to profile consumer behaviour and the access-key tracking side-channel attack actually defeats the purpose of the re-encryption.

Strong tracing Interception of valid RFID information enables law enforcement agencies to determine the associated serial number. However, the denial of service attack described in section 3.4 demonstrate that it is easy to invalidate the stored information and therefore criminals that wish to remove the strong tracing property of the RFID chips can easily do so.

Forgery resistance The data recovery attack allows a forger to retrieve all information necessary to forge a copy with a given serial number, without requiring optical contact. However, the scheme does succeed in disallowing a forger to forge new bank notes with previously unseen serial numbers.

Privilege separation Changing the contents of the memory cells on the bank note requires an access key. However, an adversary can eavesdrop the relatively high power forward communication channel to obtain this access key.

4.1 Vulnerabilities

When determining the vulnerabilities in the scheme that allowed these attacks to be possible, I formed the following list:

Cryptographic weakness The Fujisaki-Okamoto scheme has been applied to ElGamal encryption in an insecure way, enabling the data recovery attack. This in turn enables the aggressive variant of the ciphertext tracking attack.

Static response between re-encryptions In particular the ciphertext tracking attack has been enabled by the static response from the tag when reading the γ memory cell.

Static access key The static access key that is used to read the δ meory cell and to write to both memory cells allows it to be stored in a database for later use or allows an eavesdropper to query the tag afterwards.

Lack of on-tag input validation Because the tag doesn't validate any information that is written to either memory cells, it is open to the denial of service attack.

Keeping this list in mind, any future attempts to design a privacy protection scheme that meets the goals set by Juels and Pappu should take into account that proper cryptographic algorithms are selected. A more important observation is that a privacy protection scheme can not be achieved using tags that don't properly authenticate a legitimate reader, resulting in the requirement that tags must include some cryptographic processing functionality. For example, instead of a static access key derived from only optically readable data, the key derivation could take a nonce from the tag into account.

Also, the data that is returned from the tag should ideally be different on every read. This again probably requires cryptographic operations to be performed by the tag. Several other re-encryption based schemes have been proposed, e.g. [9] and [5], but these acknowledge the fact that RFID tags are not strong enough to perform the cryptographic functions themselves. In no way can they fix the ciphertext tracking attack.

5 Conclusion

RFID is a technology that has the potential to be adopted by central banks as a security feature for bank notes to battle counterfeit productions. However, privacy issues have prevented a widespread adoption of the technology. In an attempt to solve these issues, Juels and Pappu designed a privacy protection scheme for bank notes, involving re-encryption to change the values stored in the bank note and prevent illicit tracking, while allowing law enforcement agencies to trace bank notes without having optical contact. Avoine has described several attacks that will work on bank notes implementing this scheme. As a result the privacy protection scheme can be considered useless for all practical purposes. An analysis of the attacks reveal the vulnerabilities in the scheme that made these attacks possible, which gives a list that future privacy protection scheme designers should take into consideration.

References

[1] Gildas Avoine. Privacy issues in rfid banknote protection schemes. In *Smart Card Research and Advanced Applications CARDIS*, pages 33–48. IFIP, Kluwer Academic Publishers, 2004.

[2] European Central Bank. Security features. `http://www.ecb.int/euro/banknotes/security/html/index.en.html`.

[3] Fleur de Coin. Rfid banknotes. `http://www.fleur-de-coin.com/eurocoins/rfid.asp`.

[4] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes, 1999.

[5] Philippe Golle, , Markus Jakobsson, Ari Juels, and Paul Syverson. Universal re-encryption for mixnets. In *In Proceedings of the 2004 RSA Conference, Cryptographer's track*, pages 163–178. Springer-Verlag, 2002.

[6] Ari Juels, Ravikanth Pappu, and Thingmagic Llc. Squealing euros: Privacy protection in rfid-enabled banknotes. In *Financial Cryptography 03*, pages 103–121. Springer-Verlag, 2002.

[7] John Leyden. Japan yens for rfid chips. `http://www.theregister.co.uk/2003/07/30/japan_yens_for_rfid_chips/`, July 2003.

[8] Mark Roberti. The money trail. `http://www.rfidjournal.com/article/articleview/523/1/2/`, August 2003.

[9] Junichiro Saito, Jae cheol Ryou, and Kouichi Sakurai. Enhancing privacy of universal re-encryption scheme for rfid tags. In *Embedded and Ubiquitous Computing – EUC 2004, LNCS 3207*, pages 879–890. Springer-Verlag, 2004.

[10] Junko Yoshida. Euro bank notes to embed rfid chips by 2005. `http://www.eetimes.com/story/OEG20011219S0016`, December 2001.