

Security Aspects of the Ecash™ Payment System

Berry Schoenmakers

DigiCash

Kruislaan 419, NL-1098 VA Amsterdam, The Netherlands

berry@digicash.com

Abstract. Ecash is a payment system designed and implemented for making purchases over open networks such as the Internet. In this paper we review the main cryptographic techniques used throughout the ecash system. We will focus on security aspects as well as some performance related issues. The central notion of an electronic coin is treated in detail, and the basic protocols manipulating coins are described.

1 Introduction

Behind the scenes, banks, credit card companies, and other financial institutions have been processing transactions electronically for several decades now. Two important developments are taking place however that will open up the field of electronic payment systems to the general public. First, the prospect of electronic commerce over the ever-growing Internet is creating a large demand for electronic payment methods that can be used over such an open network. Secondly, the introduction of nation-wide electronic purse schemes is creating many more places and situations where smart cards can be used for cost-effective off-line payments.

In this paper we will describe several aspects of the ecash system, mostly security related, and discuss its place among other payment technologies. Ecash finds its roots in the work by Chaum (see, e.g., [Cha83, Cha90]), who invented the notion of electronic coins as well as the basic protocols for electronic cash. Electronic coins possess similar properties as metal coins, among which is the unique feature that, due to the use of coins, a payment transaction leaves no trace about (the identity of) the payer. Currently, ecash technology (as provided by DigiCash, see <http://www.digicash.com> for more details) is used by a number of banks around the globe. These banks issue ecash to their customers, who can then spend it at affiliated merchants on the Internet.

We will focus on the core protocols that make up the ecash system. For brevity, we will not mention all the details and alternative approaches that are taken into account in the actual system. We do, however, consider some future extensions such as the extension to off-line ecash, where the use of smart cards and the Internet are combined into a highly secure and versatile privacy-protecting payment system.

2 Characteristics of Electronic Payment Systems

Although more and more consensus is building up as to which properties are required of a payment system—and in any case what the properties are supposed to mean—we are not going to list and describe these properties in this paper. Instead, we take a bottom up approach, and describe some of the basic characteristics of payment systems. From these characteristics one can then infer the possibilities and impossibilities for the numerous combinations, and what their impact is on the performance and flexibility of the system.

Payment by Instruction vs Prepaid Electronic Cash In the so-called *payment by instruction* type of systems, a payer basically orders the bank to move some sum of money from its account directly into a payee's account. Examples in this category are credit and debit cards as well as many forms of cheques. The moment at which the money is actually moved from the payer's account into the payee's account depends on the system, but at all times banks and credit card companies will try to prevent discrepancies between accounts.

The central security aspect in these systems is to ensure that only legitimate account holders are able to issue payment instructions. Of course, digital signatures are the solution for doing this over a large, open network such as the Internet. Since digital signatures only make sense if there is an infrastructure for certifying public keys, a lot of effort is devoted to just this. See, for instance, the SET (Secure Electronic Transaction) proposal, a joint effort by MasterCard, VISA, and other influential partners, which specifies a hierarchy of certification authorities on top of the payment protocols laid out in the iKP system [BGH⁺95].

Prepaid systems are conceptually close to electronic equivalents of cash. Telephone cards, smart card-based systems, as well as ecash are examples in this category. The user's account is debited as soon as the card or device is reloaded with electronic cash. During payments the electronic cash is released again, and only then the payee's account will be credited. In the mean time the issuer keeps a float corresponding to the outstanding cash.

The central security aspect in this type of system is to ensure that cards or representations of cash cannot be forged. When forgery happens, the float will ultimately be insufficient to credit the payee's accounts for received payments. Of course, it should also be ensured that only legitimate account holders can reload cash from their accounts. However, this security aspect is now limited to the infrequent withdrawal protocol, and is no part anymore of the more frequent payment protocol.

On-line vs Off-line In the field of electronic payment systems, the notions on-line and off-line refer to a specific property of the payment protocol. Although the payment protocol is functionally a protocol between two parties (payer and payee) many payment systems require that the payee contacts a third party (e.g., the bank or the credit-card company acting as an acquirer) before accepting a

payment. If that is the case, the system is called an on-line payment system; the communication between a payee and its acquirer may be using any communication medium (not necessarily the Internet). If such a contact with a third party is not required during the payment protocol, the system is called off-line. In an off-line system payees are required to contact their acquirer on a regular basis for clearing all received payments.

Secret Key vs Public Key Authentication A basic requirement of a payment protocol is that it allows a payee to receive payments from any payer. A payment can be seen as some sort of authentication of the payer towards the payee (to show that the payment is authentic). The authentication can be based on secret key cryptography or on public key cryptography. In the latter case, the payee only needs to dispose of a public key in order to verify incoming payments. Although the costs of equipping smart cards with crypto co-processors are expected to become marginal, it is important to note that the property of public verifiability can be obtained using simple smart cards only, provided one applies a method of what we call *signature transport*. In such a system, signatures are created by the issuer only, and later *endorsed* by the payer during the payment protocol, depending on a challenge from the payee. The trick is to achieve that sufficiently many payments can be made between successive reloads, which requires optimal use of the limited amount of EEPROM available on simple smart cards.

In the case authentication is based on secret key (symmetric) cryptography, however, the payer and the payee must dispose of a shared secret key in order to complete a payment. A straightforward solution is to give all users the same secret key, but this is generally considered insecure, as this would mean that breaking a single smart card (i.e., extracting its secret key) will suffice to break the complete system. The standard solution is therefore to break the symmetry between payers and payees by equipping the merchants with a highly secure tamper-proof box called a SAM that contains a masterkey. The payers' keys are derived from this master key in a process called diversification by applying a cryptographic hash (e.g., SHA-1) to the concatenation of the master key and the payer's card number. The idea is that the SAM is more difficult to break than a smart card, and also that it is possible to routinely check (as part of the maintenance) if the SAMs have not been tampered with. A compromise is the use of RSA certificates in the EMV standard. Each card carries a fixed RSA certificate to show the validity of the card number. At the start of each payment, the certificate can be verified against the public key stored in the POS terminal. The remainder of the payment protocol again relies on a secret master key stored in the SAM of the POS terminal.

Counters vs Coins The most direct way of representing electronic cash is to use a counter stored on a smart card. Clearly, this is a very efficient and flexible method, and any amount can be paid from the card as long as it does not exceed the value of the counter. A more involved way is to represent electronic cash by

a set of electronic coins. As with ordinary coins, each electronic coin has a fixed denomination. Now, any amount can be paid as long as it can be obtained as a sum of the denominations of a subset of the available coins. Hence it is possible that a certain amount cannot be paid using the available coins although the total value of the coins is larger than the required amount. In practice this problem is alleviated by choosing a suitable distribution of the coin denominations upon reloading, possibly as a function of the expected spending pattern.

The CAFE project (see, e.g., [BBC⁺94]) relies on a compromise between these two basic methods. For each payment the required amount is debited from a counter but at the same time *one* special coin is used up. The special coins have no value by itself. Upon reloading the counter is credited with the withdrawn amount and the supply of special coins is replenished.

As we will explain later on in this paper, an important property that separates coins from counters is that electronic coins are the only way to achieve a system that is secure (in the bank's interest) and at the same time protects the users' privacy in a strong sense. Another argument that indicates a separation between coin-based and counter-based systems is given by the following scenario. Consider an off-line payment system that is capable of determining "after the fact" which cards have been broken in case a fraud is detected. Suppose that for this reason an attacker decides to work with stolen cards. Now, if the system relies on a counter, it is no problem if the stolen cards are empty: the attacker only needs to manipulate the counter (e.g., set its most significant bit to 1) and may then start spending right away. For a coin system, however, it depends on how many coins the stolen cards contain: if the cards are completely empty the attack will fail completely, and in any case the attacker is limited to using the coins that were present at the time the cards were stolen. Furthermore, since payees can be instructed not to accept the same coin twice, the attacker can be forced to visit different shops in order to be successful.

3 Money Flow

We briefly describe the money flow in the ecash system. Where appropriate we will distinguish between the ecash bank (or issuer/acquirer) and the ecash mint. The mint is the component of the ecash system where coins are created and where the databases of spent coins are held. So-called ecash accounts form the interface between the bank and the mint. In practice, several ways will be provided to transfer money to and from an ecash account. For example, an ecash issuer may provide a home-banking application that allows its customers to move money between their bank accounts and their ecash accounts. Another possibility is that the bank accepts credit card payments, by means of which users can feed their ecash accounts.

We concentrate on the basic operations that manipulate ecash coins. Other important ingredients of electronic commerce protocols, such as certificates and receipts, are omitted as these parts are more or less independent of the way the core protocols are implemented.

Withdrawal By means of the withdrawal protocol, users are able to convert money from their ecash accounts into ecash coins. Access to the ecash account is only possible if the user is able to sign the withdrawal request, where the signature is checked against the public key registered with the ecash account.¹ The coins obtained in a withdrawal are stored on the user's hard disk. By default the coins are stored in a password-encrypted manner to prevent them from being stolen (copied).

Payment To pay a certain amount, a set of coins is selected such that the values add up to the required amount. In the on-line ecash system, this set of coins is then encrypted for the bank, using the bank's public key, to prevent that the shop or anybody else can steal (copy) the coins. The shop deposits the payment at the bank, who credits the shop's ecash account if all coins are valid and none of the coins has been spent before. Accepted coins are added to the database of spent coins so that double-spending will be detected.

Payment Deposit In the on-line ecash system this protocol is part of the payment protocol as executed by the shop. In an off-line ecash system this protocol is executed at a later moment, preferably in batch mode. An important property of the payment protocol is that the payment deposit is made specific to the payee. That is, a payment deposit for a specific payee cannot be deposited to any other account than the account of the specified payee.

Coin Redemption It is possible to return coins directly to the mint without using them in a payment. A natural restriction is that the number of coins that a user redeems is not allowed to exceed the number of coins that has been withdrawn by the user; a more refined way is to monitor this per coinage and per denomination. This protocol is used when expired coins are refreshed, or to improve the distribution of the coin denominations.

Recovery If desired, users are able to resurrect coins that have been lost, for instance, because of a hard disk crash. By means of a special recovery protocol executed between the user and the mint, all the coins that have been withdrawn by the user since the previous checkpoint can be reconstructed. However, only coins that have not been spent before will be usable for new transactions. (The same effect can be achieved when users keep backups of their coins.)

¹ A simple way to set-up ecash clients is to assume that the software and the bank's public key are presented to the user securely (e.g., on a sealed floppy disk). The user also gets an account number and a PIN code from the bank. At home the user installs the ecash client, generates its own private key/public key pair, and registers it with its ecash account by sending it to the bank together with the PIN code (everything encrypted with the bank's public key). The user's private key can be stored in a password-encrypted manner on the hard disk.

4 Preliminaries

RSA Signatures For authentication of messages we currently use the RSA public key cryptosystem. Each participant picks its own RSA key pairs at random. The public key consists of a modulus $n = pq$ of prescribed size, where p, q are two randomly generated primes of equal size, and an exponent e , which is co-prime with $\phi(n) = (p - 1)(q - 1)$. The private key d is the multiplicative inverse of e modulo $\phi(n)$, that is, the unique number d satisfying $de \equiv 1 \pmod{\phi(n)}$, which we denote by $1/e$.

To sign a message $m \in \mathbb{Z}_n^*$, where the signer has private key d and public key (e, n) , the signer computes the signature $s = m^d \pmod{n}$. To verify the signature we check that $s^e \equiv m \pmod{n}$; this identity must hold as we have the identity $m^{de} \equiv m \pmod{n}$, since $m^{\phi(n)} \equiv 1 \pmod{n}$ for all $m \in \mathbb{Z}_n^*$. Actually, it can be proven that $m^{de} \equiv m \pmod{n}$ for all $m \in \mathbb{Z}_n$.

For various reasons (e.g., because RSA signatures can be existentially forged as described above: select an arbitrary s and take $m = s^e \pmod{n}$, then s is a valid RSA signature on the “message” m), the actual message is usually first transformed into a related message by applying a one-way hash and/or redundancy-adding function f to it and then signing the result $f(m)$. In the next section we will see an example of such a function f .

Hybrid Encryption The RSA cryptosystem can be used for encryption as well. To encrypt a message m , $0 \leq m < n$, for a receiver with public key (e, n) , the sender computes the ciphertext $c = m^e \pmod{n}$ and sends c to the receiver. To decrypt the message, the receiver uses its private key d to compute $c^d \pmod{n}$, which is equal to $m^{de} \pmod{n} = m$.

In practice, the use of public key encryption is often limited to the encryption of session keys. A symmetric encryption algorithm is then used to encrypt the actual message with the session key. In the ecash system, we also use such a hybrid encryption method (or, “envelope method” as it is sometimes called), where we combine RSA with 3-DES (with 112 bit keys). For reasons of efficiency, it is often advantageous to use public exponent $e = 3$; encryption of the session key then amounts to two modular multiplications, and poses no security risk as long as some well-known attacks are taken into account.

5 Ecash Coins

We will now have a closer look at the internal structure of ecash coins. For each coinage (short for a “generation of coins”), the mint will randomly generate a fresh RSA modulus $N = pq$, keeping the primes p, q secret by storing them in a safe place. Preferably the mint’s private keys are only used within the boundaries of tamper-resistant devices, while backups are kept between several entities using secret-sharing techniques. In this way, it is prevented as much as possible that private keys are compromised through attacks by insiders.

v	3	5	7	11	13	17	19	23	29	31	37	41
D_v	\$0.005	\$0.01	\$0.02	\$0.04	\$0.08	\$0.16	\$0.32	\$0.64	\$1.28	\$2.56	\$5.12	\$10.24

Table 1. A binary scheme with $k = 12$ different denominations

The denominations of a coinage are encoded by using different public exponents (but the same modulus). Let k denote the number of different denominations, and let $\{v_i\}_{i=1}^k$ denote the first k odd primes. In order that each v_i is a valid RSA exponent, we have the condition that each v_i is co-prime with $\phi(N)$, that is, $\gcd(v_i, \phi(N)) = 1$ for $i = 1, \dots, k$.² We denote the denomination that is associated with public exponent v_i by D_{v_i} ; see Table 1 for an example.

To limit the storage space required for ecash coins we take full advantage of the message recovery facility of RSA signatures. This is particularly useful because ecash coins are in fact RSA signatures on small messages. Thus we get the following form, where an ecash coin C of denomination D_v consists of an RSA signature only:

$$C = f(x)^{1/v} \pmod N. \tag{1}$$

For concreteness assume that x is 160 bits long, and let \mathcal{H} denote a one-way hash function whose output length is 160 bits as well (like SHA-1). Function f is a redundancy-adding function defined by

$$f(x) = x_t \parallel \dots \parallel x_1 \parallel x_0,$$

with $x_0 = x$ and $x_{i+1} = \mathcal{H}(x_0 \parallel \dots \parallel x_i)$. Parameter t is fixed such that the total length of $f(x)$ is about equal to the size of the modulus N . (Actually, the bit string $f(x)$ is truncated such that the corresponding integer is always less than N .) Note that function f may alternatively be defined by $f(x) = y_t$, with $y_0 = x$ and $y_{i+1} = \mathcal{H}(y_i) \parallel y_i$.

In current ecash implementations RSA moduli used for ecash coins are at least 768 bits long; for this size forgery of ecash coins is considered entirely infeasible—certainly within the limited time-frame that a coinage is valid. Hence, a storage of about 100 bytes per coin on the user’s side is required. With today’s hard drives and memory chips there is absolutely no problem of storing any sensible number of coins. Over time this will only improve as the required key length for RSA is expected not to double every year or so, while the storage capacity of modern devices does.

As for the coin storage on the mint side, we have the important observation that after checking the validity of a coin signature, and checking that the coin has not been spent before, it suffices to store the coin number only. As described

² This condition may be tested efficiently by precomputing $V_k = \prod_{i=1}^k v_i$, and verifying whether $\gcd(V_k, p - 1) = 1$ and $\gcd(V_k, q - 1) = 1$. Note that this condition on the primes p, q can easily be met by choosing them from the set of safe primes ($p = 2p' + 1$, $q = 2q' + 1$ with p', q' prime), but we do not want to limit the set of candidate primes more than necessary.

above, we have fixed the size of the coin number at 20 bytes. The only condition on the size of the coin numbers is that it is large enough to prevent that the same coin number is accidentally generated for coins of the same coinage. (As a slight refinement, note that it suffices that the coin numbers are unique per denomination.) By the standard result of the birthday paradox, the probability that two coins will be equal when the coin numbers are picked uniformly at random from $\{0, 1\}^{160}$ is bounded by $e^{-B(B-1)/2^{161}}$, as long as at most B coins of the same type are generated. This shows that the probability that two coin numbers collide is truly negligible for any practical number of coins B . Since the number of coins per coinage is limited anyway (see Section 8), this analysis in fact shows that the size of the coin numbers can also be limited to 64–80 bits (8–10 bytes), say, if desired. Hence, a 1Gbyte hard disk can already store more than 100 million coin numbers.

6 Protocols

We consider the withdrawal protocol, the payment protocol (including the payment deposit), and the recovery protocol. Each protocol is described at a level of detail that permits us to explain the main security features of the ecash system.

6.1 Withdrawal

For each ecash coin to be withdrawn from a user's bank account, the user and the mint execute an instance of Chaum's blind signature protocol [Cha83]. This protocol is executed in parallel as many times as required to withdraw the desired amount. The distribution of the coin denominations is chosen in such a way that it is guaranteed that—no matter how the money is spent—at least a certain number of payments can be made. The protocol for one coin is depicted in Figure 1. Apart from the random coin number x , the user also picks a random number $r \in \mathbb{Z}_N^*$, which is used to blind the “message” $f(x)$ to be signed by the mint. Since $r \in \mathbb{Z}_N^*$ its inverse $1/r$ exists, hence the blinding factor can be removed again to obtain the coin $C = f(x)^{1/v} \bmod N$.

Regarding the security of this part of the ecash system, the important question that needs to be answered is whether it is infeasible to obtain more or other coins than prescribed by the withdrawal protocol of Figure 1. The answer is that an in-depth security analysis shows that (in a reasonable model) any way to obtain more or other coins than prescribed by the withdrawal protocol is equivalent to breaking the RSA assumption (which roughly states that it is infeasible to compute $y^{1/v}$ for randomly selected $y \in \mathbb{Z}_n^*$). For the scope of this paper, we will confine ourselves to two aspects of this analysis that are specific to the ecash system.

First, there is the fact that we use not just one, but a number of public exponents with the same RSA modulus to encode the different denominations within a coinage. This raises the question whether, for instance, coins of lower denominations cannot be combined into coins of higher denominations. More

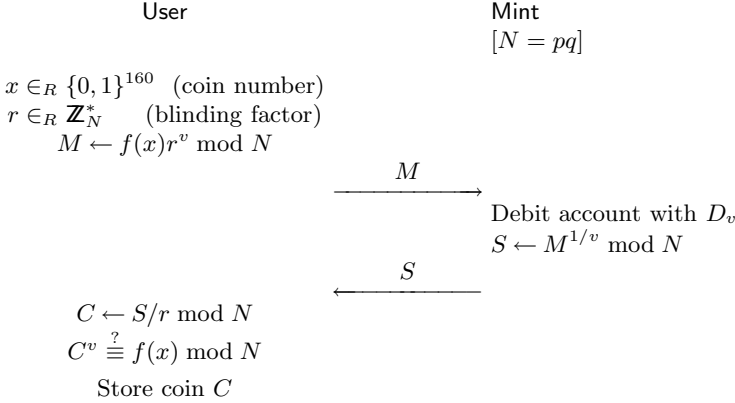


Fig. 1. Withdrawal of a coin $C = f(x)^{1/v} \bmod N$ of denomination D_v

formally, the question is whether for randomly selected $M \in \mathbb{Z}_N^*$, it is feasible to compute, say, M^{1/v_k} from the values $\{M, M^{1/v_1}, \dots, M^{1/v_{k-1}}\}$. Recall that the v_i 's are pairwise co-prime. For a very different purpose, but equally applicable to our setting, Shamir has shown that this question is answered in the negative: he showed that given the values $\{M, M^{1/v_1}, \dots, M^{1/v_{k-1}}\}$ the computation is just as difficult as if these values weren't given at all.

Secondly, there is the fact that the mint will just accept *any* message M , and return $S = M^{1/v} \bmod N$ to the user—but each time debiting the user's account with D_v , even if the user picks the message M of the wrong form. So, there is no guarantee that the message M is of the prescribed form $f(x)r^v$, where the user knows x, r , which contrasts with the usual setting for RSA signatures, where the signer will ensure that each message signed is of the required form $f(m)$, say. In fact, by means of an example we will now show that it is possible to obtain valid coins although we deviate from the prescribed protocol.

Consider the following approach. We are going to obtain two coins $C_1 = f(x_1)^{v_1}$ and $C_2 = f(x_2)^{v_2}$ of denominations v_1 and v_2 respectively. However, instead of just asking the mint to sign $f(x_1)$ and $f(x_2)$, respectively, we do it as follows.

1. Ask the mint to sign $M_1 = f(x_1)^{v_2} f(x_2)^{v_1}$ for denomination D_{v_1} , which results in $S_1 = M_1^{1/v_1}$.
2. Subsequently, ask the mint to sign $M_2 = S_1$ for denomination D_{v_2} , which results in $S_2 = M_2^{1/v_2} = M_1^{1/v_1 v_2} = f(x_1)^{1/v_1} f(x_2)^{1/v_2}$.
3. Finally, using the fact that $\gcd(v_1, v_2) = 1$, hence there exist t_1, t_2 such that $t_1 v_1 + t_2 v_2 = 1$, we extract the coins C_1 and C_2 from S_2 :

$$S_2^{t_2 v_2} f(x_1)^{t_1} f(x_2)^{-t_2} = f(x_1)^{(t_2 v_2 + t_1 v_1)/v_1} = C_1,$$

$$S_2^{t_1 v_1} f(x_1)^{-t_1} f(x_2)^{t_2} = f(x_2)^{(t_1 v_1 + t_2 v_2)/v_2} = C_2.$$

So, although we do not follow the protocol as described in Figure 1, we are able to obtain two valid coins anyway. (If desired, it is possible to blind the second request to the mint (step 2), such that the mint cannot detect the deviation.) The net result, however, is not very encouraging: we have obtained two coins for a total value of $D_{v_1} + D_{v_2}$, but also the account has been debited with this amount.

Extending work by Shamir [Sha83] and work by Akl and Taylor [AT83], Evertse and van Heyst showed that for a very general class of deviations from the prescribed protocol, nothing is to be gained [EH92, EH93]. As shown by Chaum [Cha90], however, “deviations” as exemplified above can be used to improve the efficiency of the withdrawal protocol. For example, it is perfectly safe to collapse steps 1 and 2 in the above protocol, where the mint issues a signature w.r.t. exponent $v_1 v_2$ (charging, of course, $D_{v_1} + D_{v_2}$ for this service). This cuts the communication costs with the mint by a factor of two, and also saves the mint from performing one signature. Needless to say, this method can be extended to collapse the withdrawal of any number of coins, as long as each denomination does not occur more than once.

6.2 Payment

Before a payment actually takes place, payer and payee have to come to an agreement about what the object is that is going to be purchased and for which amount X . We assume that the result of this negotiation is recorded in the string `pay-spec`. For later reference, the payer will also assign a random transaction number ID_{trans} to the payment (the use of successive transaction numbers might enable the bank to link successive payments made by the same user).

Next, the payer will select a set of coins C_1, \dots, C_l from its coin supply, such that the total value of these coins equals the requested amount X . To pay a payee with identity ID_{shop} , the payer then assembles a payment message that consists of the concatenation of `pay-spec`, ID_{trans} and an encrypted message Y for the bank:

$$Y = E_{PK_{bank}}(ID_{trans} \parallel ID_{shop} \parallel \mathcal{H}(\text{pay-spec}) \parallel \mathcal{H}(\text{pay-code}) \parallel C_1 \parallel \dots \parallel C_l),$$

where E denotes a hybrid RSA encryption method (using 3-DES), as explained in Section 4. Upon receiving this message, the payee will then sign and forward a payment deposit message consisting of $\mathcal{H}(\text{pay-spec})$, ID_{trans} and the encrypted message Y to the bank. Finally, the bank decrypts Y and checks the values of ID_{trans} , ID_{shop} , and $\mathcal{H}(\text{pay-spec})$, and then proceeds to check the validity and freshness of the coins C_1, \dots, C_l . Only if all coins are accepted, the payment is accepted as well, and the coins are stored for future double-spending checks.

A few remarks regarding security. It is important that neither the payee nor any eavesdropper can extract the coins from the payment message. For this reason, the coins are encrypted with the bank’s public key. (Note that public key encryption is required in this case, because the bank and payer cannot use

a shared secret key, as the payer needs to remain anonymous.) Also, note that although the string `pay-spec` itself is never shown to the bank, the payee is sure that the payer used the same string, since the bank compares the hash values provided by the payer and payee, respectively. In this way the transaction details remain hidden from the bank; if required, however, either the payer or the payee can later reveal the string `pay-spec` which can then be checked against the data stored at the bank.

Finally, let us briefly describe how interrupted payments can be resolved, in case the payer and payee aren't able to recover from the interruption by normal means. Due to network problems, the payment protocol may be interrupted at several stages, but to the payer it only matters whether (i) the payment wasn't processed by the bank, or in any case it was not credited to the payee's account, or whether (ii) the payment was processed by the bank and credited to the payee's account. To find out about the status of a payment, the payer can either redeem the coins C_1, \dots, C_l on an individual basis, or find out about the complete payment by sending message Y to the bank.

In the latter case, the payer proves to be the owner of the payment by revealing the value of `pay-code`, to which the user committed by including $\mathcal{H}(\text{pay-code})$ in the message Y . If the payment turns out to be credited to the payee's account, the bank signs a statement to this effect, which the user can then show to the payee to prove that the payment arrived at the payee after all.

7 Privacy

We present the by now standard argument why the ecash system protects the privacy of its users. To fully protect the users' privacy, the system must not only satisfy the requirement of untraceability, but also the stronger requirement of unlinkability; a system in which each user gets a pseudonym is not sufficient to protect the privacy of its users.

Untraceability More precisely, this property is called payer-untraceability (or, payer-anonymity). This means that *individual* payment transactions are not traceable to the user who acted as the payer in such a transaction. The information that the payee and its acquirer obtain from the transaction details does not contain any information as to which user took part as the payer.

Unlinkability This property says that it should even be impossible to link any two payment transactions originating from the same user.

To appreciate the difference between untraceability and unlinkability consider the following scenario. When you buy a prepaid telephone card you can do this completely anonymously at a newsstand (paying the card with cash). Later when you use the card in a public phone the telephone company will have no clue that it is you making the phone call because you bought it anonymously. That is,




Fig. 2. Life-cycle of a coinage

the individual telephone calls are untraceable, as they cannot be connected to your identity. Suppose however that the telephone company gives every card a unique number, which is quite realistic as this is a basic mechanism to detect fraud (i.e., to find cards on which the total spent is larger than the card's value). Then it is easy to keep a file per card of all phone numbers called from that card (and possibly the time and date of the calls as well). Since a similar file is kept per home-phone as well, a simple pattern matching procedure will in many cases reveal the identity of a card's owner. Thus, although the card is obtained anonymously, the identity of the card's owner can be revealed anyway because all calls from the same card are linkable.

We now argue why the ecash withdrawal protocol protects the users' privacy in an information-theoretic sense. As untraceability is weaker than unlinkability we only have to show that the ecash coins are unlinkable. Consider a fixed denomination D_v . Let S_A denote a signature of denomination D_v that has been issued to some user A , and let C_B denote a coin of denomination D_v that has been spent by some user B (A and B are not necessarily different). For any such pair (S_A, C_B) there exists a unique blinding factor $r_{AB} \in \mathbf{Z}_N^*$ that satisfies $C_B = S_A/r_{AB} \bmod N$. In other words, each signature issued by the mint matches equally likely with any of the coins spent. So, in an information-theoretic sense, there is no way the bank will be able to link coins that belong to the same user.

8 Coinages

Key Schedule In practice, an ecash mint will work with several *coinages* at the same time. For each coinage the mint generates a fresh RSA modulus, as described in Section 5. Hence, there is a one-to-one correspondence between the RSA moduli and the coinages. Apart from the RSA modulus, a coinage has at least the following attributes: the identity of the mint, the sequence of denominations, the currency, and a key schedule (i.e., expiration dates) as in Figure 2. These attributes are distributed with each coinage. Note that the private key can in fact be removed as soon as withdrawals are deactivated, while the public key must be available until the coinage becomes invalid.

There are two main reasons why coinages are refreshed on a regular basis, say every three months. First, there is the standard reason for refreshing keys, namely to limit the risk that the secret key is compromised. There are two direct

ways this may happen. In the first method, the attacker will try to find the secret key from the public key only, possibly using the mint as a signing oracle. In the second method, the attacker simply tries to get hold of the secret key by breaking into the bank or into its computer network, possibly assisted by insiders.

The other important reason is to limit the size of the “spent coin” databases. Using a scheme like that of Figure 2, spent coins are first stored on hard disk to enable fast checking for duplicates. After some time the coins are moved to tape, and during that period it will still be possible to check for duplicates, but using a slower procedure (which is only available so that users can return coins when they haven’t used their ecash for a longer period of time). Eventually, a coinage will become invalid and the tapes can be removed.

Clustering Apart from a division in time, where coinages are refreshed every three months say, we can also use a division in space. That is, instead of viewing all users as one big set of users, it makes sense to divide the users in clusters. Each cluster is chosen sufficiently large such that the behaviour of the individual users is not visible. The advantage is that by limiting the size of a cluster that the corresponding parts of the “spent coin” database are independent of each other. This enhances the scalability of the system considerably. The way the bank divides its users into clusters should be publicly verifiable, and not at the discretion of the bank (to prevent the bank from introducing a few small clusters). There are many ways to accomplish such a fair division into clusters. A simple idea is to first take a cryptographic hash of the user’s identities, and then define 2^t clusters, $t \geq 0$, by looking at the first t bits of the hash value. Assuming that the bank cannot influence (the representation of) the user’s identities, the users are thus evenly spread over the clusters.

If so desired, other types of clusters like age-groups can also be encoded as different coinages. Certain shops could then be made to accept ecash from certain age-groups only. This is an example of the token functionality that can also be achieved with the same core protocols of ecash. The token functionality is in turn subsumed in the general notion of electronic credentials, which encompasses things as diverse as theatre tickets, driver’s licenses, passports, diplomas, etcetera (e.g., see [Cha92]).

9 Off-line Ecash

It is interesting to consider the development of privacy-protecting off-line cash protocols. Since the introduction of Chaum’s double-spending paradigm for such protocols and the first solutions for the problem ten years ago [CFN90], the protocols have developed into an efficient and relatively simple system for off-line cash. Chaum’s double-spending paradigm says that the privacy of a user is completely protected as long as the user spends each coin not more than once. However, if a user is able to manipulate its payment device such that some coins are used more than once, the protocols are such that the identity of the double

spender can be computed. In other words, the protocols satisfy the property “once concealed, twice revealed” (courtesy Franklin and Yung [FY93]).

[FY93], [Fer94] eliminates withdrawal protocol, [Bra94a, Bra94b] introduces, and [Sch95]

Recently, several extensions to the basic protocols have been proposed, such as an efficient method of making the privacy revokable by a trusted third party.

References

- [AT83] S. Akl and P. Taylor. Cryptographic solution to a problem of access control in a hierarchy. *ACM Transactions on Computer Systems*, 1:239–248, 1983.
- [BBC⁺94] J.-P. Boly, A. Bosselaers, R. Cramer, R. Michelsen, S. Mjølsnes, F. Muller, T. Pedersen, B. Pfitzmann, P. de Rooij, B. Schoenmakers, M. Schunter, L. Vallée, and M. Waidner. The ESPRIT Project CAFE – High Security Digital Payment Systems. In *Computer Security – ESORICS 94*, volume 875 of *Lecture Notes in Computer Science*, pages 217–230, Berlin, 1994. Springer-Verlag.
- [BGH⁺95] M. Bellare, J. Garay, R. Hauser, A. Herzberg, H. Krawczyk, M. Steiner, G. Tsudik, and M. Waidner. iKP — a family of secure electronic payment protocols. In *First USENIX Workshop on Electronic Commerce*, 1995.
- [Bra94a] S. Brands. Untraceable off-line cash in wallet with observers. In *Advances in Cryptology—CRYPTO ’93*, volume 773 of *Lecture Notes in Computer Science*, pages 302–318, Berlin, 1994. Springer-Verlag.
- [Bra94b] S. Brands. Off-line cash transfer by smart cards. In V. Cordonnier and J.-J. Quisquater, editors, *Proceedings First Smart Card Research and Advanced Application Conference*, pages 101–117, 1994. Also as report CS-R9455, Centrum voor Wiskunde en Informatica.
- [CFN90] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *Advances in Cryptology—CRYPTO ’88*, volume 403 of *Lecture Notes in Computer Science*, pages 319–327, Berlin, 1990. Springer-Verlag.
- [Cha83] D. Chaum. Blind signatures for untraceable payments. In D. Chaum, R.L. Rivest, and A.T. Sherman, editors, *Advances in Cryptology—CRYPTO ’82*, pages 199–203, New York, 1983. Plenum Press.
- [Cha90] D. Chaum. Online cash checks. In *Advances in Cryptology—EUROCRYPT ’89*, volume 434 of *Lecture Notes in Computer Science*, pages 288–293, Berlin, 1990. Springer-Verlag.
- [Cha92] D. Chaum. Achieving electronic privacy. *Scientific American*, pages 96–101, August 1992.
- [EH92] J.-H. Evertse and E. van Heyst. Which new RSA-signatures can be computed from certain given RSA-signatures? *Journal of Cryptology*, 5(1):41–52, 1992.
- [EH93] J.-H. Evertse and E. van Heyst. Which new RSA-signatures can be computed from RSA-signatures, obtained in a specific interactive protocol? In *Advances in Cryptology—EUROCRYPT ’92*, volume 658 of *Lecture Notes in Computer Science*, pages ??–??, Berlin, 1993. Springer-Verlag.
- [Fer94] N. Ferguson. Single term off-line coins. In *Advances in Cryptology—EUROCRYPT ’93*, volume 765 of *Lecture Notes in Computer Science*, pages 318–328, Berlin, 1994. Springer-Verlag.

- [FY93] M. Franklin and M. Yung. Secure and efficient off-line digital money. In *Automata, Languages and Programming, ICALP 93*, volume 700 of *Lecture Notes in Computer Science*, pages 265–276, Berlin, 1993. Springer-Verlag.
- [Sch95] B. Schoenmakers. An efficient electronic payment system withstanding parallel attacks. Report CS-R9522, Centrum voor Wiskunde en Informatica, March 1995.
- [Sha83] A. Shamir. On the generation of cryptographically strong pseudorandom sequences. *ACM Transactions on Computer Systems*, 1:38–44, 1983.