Blinded Digital Signatures

We need one more general concept before we dive into digital cash. A blinded digital signature is something like signing a blank check, or like signing the envelope containing the check. It doesnt sound like a good idea, but as we will see it is very useful. Blinded signatures are one of several techniques used to preserve privacy with digital cash. They make it harder to trace the path of some cash.

Let M be a message which is a bank note or check. A bank could sign this in the RSA signature scheme by first computing a secure hash $H(M)$ of the message, and then computing $H(M)^d (\mathrm{mod}\ \ n)$ as the signature, where d is the banks private RSA key. Anyone could verify given M and the signature, that the bank had meant to sign this note. The problem with this is that it gives banks great power to trace your spending. Since they issue the note to you in the first place, but will receive the note for collection from a merchant, they could match the two. That is, they could save M in a database indexed with your name and then match the M that the merchant brings in against it. This happens with credit card transactions, and people often prefer cash because it is not traceable this way.

The anonymity is accomplished with a blinding factor k, which is known to you only. You ask the bank to sign:

$$r = H(M)k^e (\mathrm{mod}\ \ n)$$

where e is the bank's *public* RSA key. What you get back is

$$r^d = H(M)^d k (\mathrm{mod}\ \ n)$$

and you can multiply this by k^{-1} to get back $H(M)^d (\mathrm{mod}\ \ n)$ which is the signed note. The bank only saw $H(M)k^e$ which is difficult to distinguish from random data.

But how do you convince a bank to sign a blank piece of paper? They might be commiting themselves to pay a huge amount of money. The answer is very similar to a zero-knowledge proof (c.f. the ZKP of graph coloring). Instead of one note, you present the bank with many notes, all with the same value (say $10). That is, you give the bank:

$$r_i = H(M)k_i^e (\mathrm{mod}\ \ n)$$

for $i = 1, \ldots, k$, and the same message M. The blinding factors k_i are different in each case. Then the bank picks one of the notes at random and sets it aside and asks you to unblind the other $k-1$. You send the bank the k_i factors for notes it has asked for. The bank unblinds those notes by multiplying by $k_i^{-e} (\mathrm{mod}\ \ n)$, and checking for a match with $H(M)$. The bank discovers through this process that you were honest in filling out all $k-1$ of these notes. So it agrees to sign the other note (that it set aside) without seeing it. The bank doesnt know what this note actually contains. But there can be at most one bad note, and the probability that the bank picks it is

$\frac{1}{k}$

This is not much protection, but we will see some other schemes for protecting the bank later.

Digital Checks

If we made a digital version of a classical personal check, it would look something like this:

$$\text{Signed}_{Owner}(BankName, Owner'sName, Amount, RecipientName)$$

We'll assume that the signature includes the original document, i.e. that $\text{Signed}_{Owner}(X)$ means $(X, R(X, d))$ where $R(X, d) = X^d (\text{mod } n)$ is the RSA en(de)cryption function. The Owner's signature proves to the bank that they want to withdraw this money.

In order to cash the check, someone would present it to a bank. In general, the bank where the check is collected will be different from the bank that actually provides the funds, which is normally your bank. The chain of signatures that gets applied to the check looks like this:

$$\text{Signed}_{Owners\ bank}(\text{Signed}_{Collection\ bank}(\text{Signed}_{Recipient}(\text{Signed}_{Owner}(Data..))))$$

The nesting of the signatures proves that they were applied in the correct order. This is both stronger and weaker than an ordinary check. Stronger because it is hard to forge digital signatures, but weaker because it is easy to make multiple copies of the check at any stage.

Adding Privacy: First version of digital cash

There are several reasons for having a cash-like form of digital money. The first and most important is privacy. We dont want a lot of agencies to be able to monitor how the user is spending their money. So our goal as digital cash designers is to hide the identities of the parties to a digital exchange. A first attempt might be to use a check made out to "cash".

$$\text{Signed}_{Owner}(BankName, Owner'sName, Amount, \text{“}Cash'')$$

This provides some protection for the recipient, but it doesnt prove to a merchant that that person has the money in their bank account. So we add a signature from the owner's bank that verifies that the bank is willing to pay that amount.

Nothing we've done so far stops a merchant or the owner from double-spending. To do that, we add a serial number chosen by the owner at random from a very large set (say 128 bits) so that the probability of collision with other notes is negligible. The cash note now looks like this:

$$\text{Signed}_{Bank}(\text{Signed}_{Owner}(BankName, SerialNumber, Amount))$$

When the bank receives notes for collection, it checks the serial numbers against the ones that were presented before. The bank's signature verifies that the person has the money in their account. The bank may choose to debit the account when the note is created, or simply keep track of how

many notes have been debited.

We have removed the owner's name in an attempt to protect their privacy. The bank has to check the owner's digital signature to make sure they really want to create this money. And the serial number allows them to keep track of whether the note was spent multiple times. So this note is reasonably safe. The drawback is that the serial numbers give the bank an easy way to keep track of what each user is paying for. When an owner gives their bank a note to sign, the bank can record the serial number. Then when someone deposits that note later, the bank can look up the serial number and match that person with the owner.

Anonymous Digital Cash

To have real cash, we need better protection of the owner's identity so that their spending cannot be tracked. One solution is to use the blinded signature technique described earlier. Through that scheme, the bank signs a note for which it has no record of the Owner's name (you would probably have to give the bank real cash in order for them to do this). In more detail, you create k notes of the form

$$K_i^d \text{Signed}_{Owner}(BankName, Serial_i, Amount) \text{mod } n$$

where K_i is a random blinding factor which is different for each note, d is the bank's RSA decryption key, and the serial number is different for each note.

The bank sets one of these notes aside, say the lth note, and asks you to unblind the other $k-1$. You give the bank the blinding factors K_i for those notes only, and the bank checks that the unblinded notes are all of the form

$$\text{Signed}_{Owner}(BankName, Serial_i, Amount)$$

The bank takes this as strong evidence that the lth note is also of this form, so it signs that note and returns it to you. You remove the blinding factor, and you now have a note of the form

$$\text{Signed}_{Bank}(\text{Signed}_{Owner}(BankName, SerialNumber, Amount))$$

for which the bank does not know the serial number. This note still has your signature on it, which is a weakness for privacy. Q: How could a bank that cashes this note figure out who you are from the signature? How could you remove the signature during the creation of the note?

This method still gives quite good privacy protection for the owner and recipient. But because of that, it is more tempting for one of those parties to cheat, especially by double spending. The serial numbers provide some book-keeping to show that someone has cheated, but its impossible to tell who duplicated the note.

Traceable Anonymous Cash

Traceable anonymous cash sounds like an oxymoron. But in the digital world, remarkable things are possible. Specifically, we can hide someone's identity in a note in such a way that if they spend the note once, they will remain anonymous. But if they try to spend multiple times, their identity will almost certainly be revealed.

Not surprisingly, we use secret sharing to do this. Recall that it is possible to take an n-bit message M and make two n-bit messages each of which contain no information about M but which together completely define M. Here is the procedure:

1. The customer makes j copies of his/her identity and splits each one in two halves. That is, the copies are $\{Id_1, \ldots, Id_j\}$ and $\{Id'_1, \ldots, Id'_j\}$. The owner's identity will be computable from any pair (Id_i, Id'_i) but from no other combination.

2. These Identities are encrypted each with a different key and become part of the users "Identity info". Each result would look like $f(Id_iG, K_i)$ or $f(Id_iG, K'_i)$, where f is an encryption function, and G is a "recognizable" string like the name of the bank. It is needed because Id_i is itself just a random string, and decrypting it with a false key would be hard to detect. The keys K_i and K'_i are bit-committed (e.g. by hashing) but kept secret.

3. A customer requesting a unit of cash creates k samples containing:
   $(Bank'sName, Amount, SerialNumber, Identityinfo)$
   where the serial number is unique to each bill, and the identity info is computed separately for each note using steps 1 and 2 above. The k units of cash are blinded using blinding factors $b_1, \ldots, b_k$ and presented to the bank.

4. The bank selects one unit at random and sets it aside. Then it asks the customer to unblind the other $k - 1$. The customer provides the blinding factors and all the keys K_i and K'_i for those notes, and the bank looks inside them. For each note, the bank checks that the blinding factor works, that all the keys K_i match their bit committed values (and that they correctly decrypt Id_iG), that all the identity halves Id_i and Id'_i gives the owner's identity, that all the bills have distinct serial numbers.

5. The bank signs the unit it had set aside, and gives it to the customer.

6. The customer unblinds this unit and is ready to spend.