Countermeasures against Government-Scale Monetary Forgery*

Alessandro Acquisti[1], Nicolas Christin[2], Bryan Parno[3] and Adrian Perrig[3,4]
[1] *Carnegie Mellon University, Heinz School of Public Policy and Management*
[2] *Carnegie Mellon University, Information Networking Institute and CyLab Japan*
[3] *Carnegie Mellon University, Department of Electrical and Computer Engineering*
[4] *Carnegie Mellon University, Department of Engineering and Public Policy*

Extended version
Carnegie Mellon University CyLab Technical Report CMU-CyLab-07-016, December 2007

Abstract

Despite cryptographic breakthroughs in the area of digital cash and the rapid advance of information technology, physical cash remains the dominant currency: it is easy to use and its exchanges are largely independent of computing devices. However, physical cash is vulnerable to rising threats - such as large-scale, government-mandated forgeries - that digital cash may protect against more effectively. We study mechanisms to combine physical cash with digital cash to remove their respective shortcomings and obtain their combined advantages. We discuss initial mechanisms, ranging from cryptographic signatures embedded in 2-D barcodes, to physical one-way functions coupled with online verification systems, and examine their cost and benefit trade-offs.

Keywords: Economics of security, Monetary forgeries, Secure payment systems

1 Introduction

Counterfeiting money is arguably as old as minting money. Fake coins have been discovered dating back to the 4th century BC [13]. Recently, possible evidence of a nation-state issuing large amounts of nearly perfect counterfeit US dollars has surfaced [19]. Even though the evidence is circumstantial (and debatable), the mere possibility of large-scale forgeries mandated by hostile governments suggests a reevaluation of the traditional threat model used to design anti-counterfeiting techniques.

Government-scale monetary forgery differs from traditional forgery (e.g., that perpetrated by organized crime) in scale, motivation, and perception. First, a counterfeiting government has access to manufacturing resources and capabilities that can be considered equivalent - in quality and production levels - to that of the national bank whose currency is being faked. Second, while these counterfeits may simply be used to increase the purchasing power of the nation-state producing the forgeries, the forged bills may also be used to finance hostile activities, such as weapons purchases, or terrorism sponsorship. As a result, targeted countries may be willing to consider relatively expensive defenses against government-mandated forgeries.

The core contribution of this paper is to introduce and outline the main technical and economic challenges that stem from the design and deployment of possible countermeasures against government-scale monetary forgery.

*Authors listed in alphabetic order. Short version to appear in the Proceedings of the Twelfth International Conference on Financial Cryptography and Data Security (FC'08). Cozumel, Mexico. January 28–31, 2008.

Despite major developments in paperless currency over the past decade, physical cash remains widely used throughout the world. An appealing aspect of physical cash is that people can trade it without the assistance of computing devices. People expect that simple visual and tactile inspection reveals fake bills. Physical cash can survive extreme situations: it can be washed in a washing machine and it can survive extreme temperatures that would render any smartcard unusable. Although not perfectly anonymous [17], physical cash, especially smaller and widely circulated bills, provides a reasonable level of privacy.

Cryptographic digital cash offers numerous benefits too, and provides two key advantages over physical money. First, an adversary cannot forge digital cash, assuming the security of the cryptographic mechanisms and the secrecy of the associated cryptographic information. Second, replication of digital cash is easy, so that one can easily safeguard against loss or theft of digital cash through digital backups.

However, we cannot simply switch to digital cash and abandon physical cash: we want to preserve the appealing aspects of physical cash, and we need to support the legacy business practices built around it. Indeed, in spite of the advance of cell phones and credit cards, we are still far from a cashless society, especially in many developing nations.

A natural approach to preventing forgery of physical cash is to combine it with digital cash, yielding *physical digital cash*. Essentially, physical digital cash consists of regular bills[1] in which the issuing government embeds an easily verifiable cryptographic value. The main goal is to devise a monetary system resilient to forgery, without requiring drastic changes to the existing monetary infrastructure.

Devising physical digital cash leads to a number of design trade-offs between the security properties achieved, the technological complexity involved, and the economic costs incurred. The core contribution of the present paper is to explore these trade-offs in search of deployable techniques against counterfeiting.

After surveying related work in Section 2, we analyze design requirements for physical digital cash in Section 3. We contrast the advantages and disadvantages of several schemes, including our own proposals, for physical digital cash in Section 4. These schemes offer various levels of protection against both basic theft and attempts at government-scale forgery. We then analyze general security threats against physical digital cash in Section 5. Finally, we conclude in Section 6.

2 Related Work

Many researchers have proposed and studied implementations of digital cash schemes; Asokan et al. [2] provide an overview article of electronic payment systems.

Based on seminal works on blind signatures [8], one line of research focuses on cryptographic digital cash systems, e.g., [6]. Similarly, several micropayment systems have also been proposed to pay for very small amounts, e.g., [9, 14, 23, 24]. Instead of looking at how one could replace cash by a novel digital cash payment system, this paper discuss the trade-offs in attempting to enhance the security of physical cash.

Another line of research, e.g. [3, 20, 29], focuses on trusted hardware-based payment systems, for instance electronic wallets. In contrast with these architectures, physical digital cash does not rely on external hardware to store balances and perform payments.

In the area of physical protection against counterfeits, each currency-printing nation has developed its own secret techniques. However, a number of public features enable people to visually inspect and verify the authenticity of each bill. For example, the U.S. Bureau of Printing and Engraving publishes details about some of the features of new U.S. dollars, such as color-shifting ink, a new watermark, a metallic security thread, and the use of micro print [31]. Euro bank notes also provide numerous security features,

[1] We will only discuss bills, although these principles could be translated to coins as well. However, given the lower economic value of coins and their high cost of production, counterfeiting coins is usually not viable.

including raised print, watermarks, a security thread, see-through numbers, holograms, a glossy stripe, a color-changing number, and UV-visible features [11]. The most valuable bank note in the world, the 1000 Swiss Franc bill, includes a kinegram, an irodin number, a watermark, UV-visible features, numbers visible only under oblique incident light (the Kipp effect), and the use of copper print, micro perforation, and optically variable ink [27].

In 2001, the European Central Bank considered embedding RFID tags in each Euro note [34]. As we will discuss in Section 4, there are numerous technological and economic drawbacks to such an approach.

Closer to the physical digital cash we envision, a few proposals have attempted to couple physical security, using physical one-way functions [21], with cryptographic verification of the bill [15, 28]. However, as we discuss further in Section 4, due to the cost of the required verification equipment, forgeries may travel undetected in the monetary network for considerable amounts of time. Finally, one of the applications of quantum cryptography lies in counterfeit deterrence [33], but, here again, the verification equipment required may be quite costly.

3 Physical Digital Cash Requirements

Ideally, currency should be resilient to large-scale, high-quality forgeries almost indistinguishable from real notes. As discussed in [19], such forgeries have already been encountered "in the wild," and are extremely difficult to detect even using sophisticated machinery. However, to justify any drastic changes to the current approach (combining physical security and police intervention), techniques used to prevent counterfeiting should remain economically efficient, and maintain the usability properties of traditional cash.

3.1 Economic properties

From a macroeconomic standpoint, the impact of monetary forgeries remains small. Assuming that both the gross domestic product and the rate at which money changes hands remain constant over a given interval of time, an increase of x% in the money supply will cause an approximate increase of x% in the inflation rate [12]. As of February 2006, the U.S. cash supply totaled about $780 billion [22]. Out of these, there are an estimated total of $180 million forged U.S. banknotes in circulation worldwide, that is around 0.01% of total currency. Under these values, forged money production must increase by a factor of 200 to corrupt 1% of the monetary supply of the US and have a 1% impact on the inflation rate. While this back-of-the-envelope calculation neglects important factors such as money multiplier effects, [2] even minimal security measures can prevent forgeries from threatening the value of the money itself.

Simple upgrade. As a result, the marginal cost of physical digital cash is tightly constrained. Current estimates suggest that the US government spends approximately 5.7 cents per bill produced, though recent anti-counterfeiting measures increased the cost to almost 8 cents. The extensions that we require for physical bills should impose a negligible overhead over current bill production methods. Techniques that would raise the production cost of a bill to 20 cents, for instance, are unlikely to be adopted.

Minimal cost to the users. A number of failed currency innovations, such as past efforts to popularize dollar coins, have shown that people are generally conservative when it comes to currency, and tend to resist drastic changes when they do not perceive any added value to it. Hence, to gain wide acceptance, any

[2] For instance, a forged $100 dollar bill deposited at a banking institution and then released in circulation results in $200 of "fake" money being in the system.

physical digital cash design should not put any burden on the users, while at the same time providing tangible benefits. Namely, the bill exchange process should not impose any additional transaction cost (monetary or otherwise) to the user, and any verification cost should remain negligible compared to the actual value of a given bill.

3.2 Usability properties

Currency is an extremely universal product, in that almost every single individual uses cash. Thus, any changes to currency must maintain its very high usability properties, in particular:

Universal use. Physical digital cash should provide the same usage characteristics as current physical cash, offering extreme ruggedness and enabling exchange without any digital devices.

Reusability. A single physical digital cash bill should be reusable once it is passed from one owner to another. This is in contrast to digital cash, which is used only once, then destroyed.

3.3 Security properties

To be resistant to any type of counterfeit, physical digital cash should fulfill the following security properties:

Forgery-proof. Given an electronic verification device, it must be impossible, or at least computationally infeasible, to create a bill that differs from one issued by a legitimate entity. In other words, forgers cannot create bills with new denominations or serial numbers; instead, they are limited to high-quality duplication of existing bills.

Universal verifiability. We require that bills be verifiable using a commodity electronic verification device. That way, individuals can easily start verifying the correctness of bills. For instance, one of the approaches we consider in this paper is to employ current camera-equipped smart phones as verification devices, since these phones are quickly becoming ubiquitous.

Useless duplication. Given an online electronic verification device, it must be impossible to duplicate an existing bill and successfully cash both bills. A single physical digital cash bill has at most a single owner at any given instant in time. This property does not imply that duplicating a physical digital cash bill is impossible, but merely that the duplicated bill should be useless.

Anonymity. One of the most salient features of physical cash is anonymity. Even though banknotes do not ensure perfect anonymity [17], physical digital cash should provide a level of anonymity equivalent to that provided by physical cash.

In essence, the above requirements describe the properties that physical cash should ideally satisfy. However, simultaneously meeting all security, usability, and economic requirements is extremely difficult. In the reminder of this paper we contrast several approaches, based on augmenting traditional cash with cryptographic primitives, and show which designs come the closest to satisfying all of our requirements.

(a) Implementation (b) Verification

Figure 1: **Barcode Signatures.** A sample implementation of physical digital cash using 2-D barcodes to encode a signature of authentication. Anyone with an appropriate scanner or camera phone can verify that a legitimate institution issued this bill (or one identical to it).

4 Physical Digital Cash Techniques

In this section, we consider a number of techniques for designing physical digital cash, including novel proposals. We evaluate both the advantages and disadvantages of each system. While none of the techniques perfectly meet all requirements outlined in Section 3, they represent interesting and useful building blocks for future physical digital cash schemes.

4.1 Barcode Signatures

By encoding signatures in 2-D barcodes, we can 1) keep all the properties of existing physical cash, and 2) strengthen the design using cryptographic primitives to make forgery impossible. Simply stated, we propose to augment existing bills with an unforgeable cryptographic signature.

Design. Since each bill already possesses a unique serial number, N, the bill's issuing authority (e.g., federal bank) can sign the serial number and the bill's denomination, D, with its private key, R_{gov}. The associated public key, U_{gov} should be widely published. While traditional bills only contain N and D, physical digital cash bills contain $(N, D, \{N||D\}_{R_{gov}})$.

To preserve the ruggedness of physical cash, we propose to embed the digital signature on the bill using a 2-D barcode, e.g., PDF417 [16], as shown in Figure 1(a). 2-D barcodes have previously been used for cryptographic verification of metered postage [30]. They allow fast optical scans and are therefore easily verifiable.

Evaluation. Since the 2-D barcode does not require any electronic circuitry on the bill, the encoded signature will be robust under extreme physical conditions. The encoding process can also employ error-correcting codes to further enhance the robustness of the signature. Thus, barcode signatures satisfy the *universal use* property of physical digital cash.

As long as the private key R_{gov} is kept secret, and assuming a secure signature scheme, such as RSA [25] or DSA [1], the bills are *forgery-proof*.

By encoding the signature with a 2-D barcode that can be readily read by commodity camera-based smart phones (as shown in Figure 1(b)), we achieve *universal verifiability*. In general, a 2-D barcode reader is much simpler than most other verification devices, such as RFID readers. Some smart phones, especially in Japan or South Korea, are already equipped with barcode reader software. We note that users would not

need to verify *all* bills they have in their possession. However, the ability to do so, for a negligible cost, is an important asset.

The manufacturing technology for adding a barcode to a bill is trivial – current bills already contain serial numbers that are printed on each individual bill, and the same technology can be used to also print a barcode. For these reasons, the barcode satisfies the *simple upgrade* property.

Finally, a physical digital cash bill does not contain more information than a traditional bill: the signature itself can only be used to verify the authenticity of a bill. Thus, the proposed scheme satisfies our *reusability* and *anonymity* requirements.

However, used alone, signatures cannot enforce the *useless duplication* property. Indeed, a duplicated bill would have the same serial number N as the original (valid) bill, so that $(N, D, \{N, D\}_{R_{gov}})$ would remain valid. To achieve the *useless duplication* property, we must turn to additional (or alternate) techniques.

4.2 RFID-based Protection

An alternative solution, which was once considered for Euro bills [34], is to embed RFID chips in bills. Using an RFID chip offers two primary advantages over 2-D barcodes. First, an RFID chip can perform limited computations and can even interact with a reader. Second, while 2-D barcodes are read-only, some RFID chips have writable memory.

Design. If we assume the use of tamper-proof RFID chips (we discuss the strength of this assumption below), then we can design a simple protocol, similar to SiB [18], to authenticate physical digital cash. For a bill with serial number N, the issuing authority generates a public-private key pair (K_N, K_N^{-1}), stores $(K_N, K_N^{-1}, \{K_N^{-1}\}_{R_{gov}})$ on the embedded RFID chip, and prints a barcode encoding of $H(\{K_N^{-1}\}_{R_{gov}})$ on the face of the bill, where H is assumed to be a cryptographically secure hash function.

To authenticate a bill, any user with an appropriate reader can transmit a randomly chosen nonce, κ, to the RFID chip. The chip responds with a signature $\{\kappa\}_{K_N}$ on the nonce, its public key, K_N^{-1}, and the certificate, $\{K_N^{-1}\}_{R_{gov}}$, for its public key. The reader checks the signature using the public key provided and checks that the hash of the certificate matches the commitment printed on the face of the bill.

Evaluation. RFID chips will be less tolerant of daily wear and tear and extreme environmental conditions than the original bill. As such, an RFID-based approach may not fully satisfy the *universal use* requirement. Further, at present, an RFID approach does not satisfy the *universal verifiability* requirement, as RFID readers have not yet penetrated the consumer market. Likewise, embedding a computational device in each bill would significantly raise the cost per bill (up to $1, according to [34], that is, a 20-fold increase) and alter production methods. While improvements in RFID technology may remedy this drawback, this technique currently does not provide a *simple upgrade*.

Since the data stored on the RFID chip does not include any information about the owner of a bill, this technique achieves both *reusability* and *anonymity*. A perfectly secure RFID chip may make forgery and duplication impossible, thereby directly enforcing the desired *forgery-proof* and *useless duplication* properties. Unfortunately, trusting the security of an RFID chip is an extremely strong assumption, as has been evidenced by existing attacks [4]. It remains an open question whether similar techniques can be developed using insecure RFID chips.

Finally, another disadvantage of RFID chips is that they can be remotely read, potentially enabling a thief to determine the amount of money a potential victim is carrying. Similar to the vulnerabilities of the new RFID-based US passport [26], adding RFID tags to bills would raise numerous new vulnerabilities.

4.3 Physical One-Way Functions

A different way to ensure the useless duplication property is to embed a physical one-way function in each bill.

Design. Physical one-way functions can be implemented, for instance, by randomly sprinkling bits of optical fiber in the fabric of each banknote [28], or by using magnetic polymers [15]. Each bill has unique characteristics due to the length and orientation of the fiber strands or polymer present in its fabric, and it is extremely hard to produce a copy of the bill with an identical physical configuration.

Exposing the bill to a light (or magnetic) source under different conditions (e.g., different angles) yields a unique characterization of the structure of the bill, which can be numerically encoded and printed on the bill. Verification is a matter of exposing the bill to the same conditions and matching the information printed on the bill. Combining this scheme with a signature scheme, e.g., by signing the value characterizing the physical structure of the bill can further ensure the forgery-proof property.

Evaluation. This approach has the merit of providing enhanced security without changing the way people would use bills. Three important open problems remain, however, regardless of the physical one-way function used. First, the manufacturing cost of such bills is hard to assess, but is certainly much higher than the current production cost. Second, fibers, or polymers may break or get dirtied easily, resulting in genuine bills failing the verification process. Third, the equipment needed to verify such enhanced bills is likely to be too high an investment for most merchants, let alone individual users.

As such, physical one-way functions do not easily satisfy *universal verifiability*, *simple upgrade*, or *universal use*. However, as we discuss later, we believe physical one-way functions may be very useful when deployed in conjunction with other techniques.

4.4 Centralized Verification

Both centralized verification and online verification (discussed in Section 4.5) attempt to achieve the *useless duplication* property. While neither provides a completely satisfactory solution, both represent interesting points in the design space.

Design. One simple way of making duplication more costly for counterfeiters is to keep a database of issued serial numbers at the issuing central bank and require that all banks verify whether a given serial number has already been deposited or not. We can thus ensure that two bills with the same serial number cannot be deposited at the same time. Adding a cryptographic signature on the bill would both prevent the introduction of illegitimate serial numbers and detect the duplication legitimate serial numbers. Without the cryptographic signature, this technique directly applies to unmodified physical cash, but it offers weaker properties, since it can only detect the introduction of illegitimate serial numbers when the bills are deposited at a bank.

Evaluation. Given that centralized verification utilizes unmodified physical cash, it clearly meets our *universal use* and *reusability* goals. It imposes no additional production costs, making it a *simple upgrade* to the printing process, though it does impose costs on the central bank, which must maintain the serial number database, as well as on the member banks that must constantly monitor and report on the serial numbers entering and leaving their control. Centralized verification minimally impacts the traditional *anonymity* of

physical cash, since the bills remained unchanged, and serial number data is already available at the member banks.

Without barcode signatures, centralized verification of serial numbers is only partially *forgery-proof* and provides only limited verifiability, since only banks can perform the verification procedure. Further, duplicate bills can remain in circulation undetected for extended periods of time. In fact, until one of the bills is deposited, not even the central bank knows that duplication has occurred.

4.5 Online Verification

Ideally, we could achieve instant detection of duplicates, such that no one would accept a duplicate bill. Online verification attempts to achieve this property by enabling individuals and merchants to perform real-time validation of bills they receive. The system offers stronger properties, but it also imposes larger costs and may introduce new vulnerabilities. While it does not offer a perfect solution, it does suggest an intriguing direction for further research.

Design. At a high level, a decentralized database (perhaps hosted by various member banks or other governmental agencies) associates each bill's serial number with a cryptographic "lock bit". Once a bill is locked, only the current "owner" of the bill can unlock it. To transfer ownership of a locked bill, the current owner cryptographically unlocks it and allows the new owner to lock it. Participants can check the current state of a particular bill's lock bit and refuse to accept a locked bill.

Dealing with legacy users (i.e., those that cannot check a bill's lock status) requires additional measures. In general, before transferring a locked bill to a legacy user, the current owner must unlock it so that the legacy user can make use of it. For example, by default, all bills dispensed by an ATM to a legacy user would be unlocked (or locked with a null value) by the issuing bank. Participating users would then take ownership of the bills by immediately locking them.

On a related note, since a legacy user cannot check the status of a bill's lock bit, a participating user might accidentally or maliciously provide them with a locked bill. A similar problem arises if an participating user loses the cryptographic material necessary to unlock their own bills. To address this problem, the online verification service must be backed by the central bank. We assume that the central bank can distinguish a duplicate from a real bill through some, possibly costly, verification process. For instance, physical one-way functions described above could assist in the verification process on the bank side. Indeed, used as a back-up verification system, physical one way functions do not need to have the same level of robustness as when used as the primary mechanism to prevent duplication.

With this online verification system in place, a user could deposit a locked bill at a bank in a procedure similar to that used for checks today. The bank would send the locked bill back to the treasury to verify its authenticity. If the bill is authentic then the bank will credit the value of the bill to the user's account, regardless of its lock status.

Implementation. The "bank" (e.g., the central bank or the treasury), denoted B, maintains a distributed database that contains an entry for each bill in circulation. Each entry is of the form (N, λ), where N represents the bill's serial number and λ indicates the lock status of that bill. If $\lambda = \emptyset$, the bill is unlocked, whereas any non-zero value indicates that it is locked. To facilitate the automation of the steps described below, each bill's serial number should be encoded in a machine-readable form such as a 2-D barcode.

To lock an unlocked bill with serial number N, a principal (e.g., an individual or merchant) A picks a random value μ_A and computes $\lambda_A = H(\mu_A)$, where H is a one-way hash function assumed to be secure, i.e.,

at least weak-collision resistant. Using the bank's public key,[3] A securely transmits (N, λ_A) to the bank. The bank will update the database appropriately. We summarize these steps below:

1. $A \rightarrow B$: $\{N, \emptyset, \lambda_A\}_{U_{gov}}$
2. B: Retrieve (N, λ), check $\lambda = \emptyset$, store (N, λ_A)

To transfer the bill to another principal, C, A will unlock the bill and simultaneously lock it under C's lock value. To simplify the presentation, assume A and C have established a secret key K_{AC}, and let $\{M\}_{K_{AC}}$ denote the authenticated encryption of a message M. When the transaction is about to take place, C picks a secret random value μ_C, and computes its hash $\lambda_C = H(\mu_C)$. The following bill transfer protocol takes place:

1. $C \rightarrow A$: $\{\lambda_C\}_{K_{AC}}$
2. $A \rightarrow B$: $\{N, \mu_A, \lambda_C\}_{U_{gov}}$
3. B: Retrieve (N, λ_A), check $\lambda_A = H(\mu_A)$, store (N, λ_C)
4. $B \rightarrow A$: $\{N, \lambda_C\}_{R_{gov}}$
5. $A \rightarrow C$: $\{N, \lambda_C\}_{R_{gov}}$

That is, C gives A the lock value λ_C, which A forwards to the bank along with her unlocking value μ_A. The bank replaces λ_A with λ_C, effectively updating the "owner" of the bill, before communicating the change back to A. Finally, A relays this information to C, proving that the lock value has been updated, and physically transmits the bill to C.

The key feature of this scheme is that, if the values μ_A and μ_C are truly chosen at random, bills can be locked to a given individual without making this individual traceable. Basically, (μ_A, λ_A) and (μ_C, λ_C) are used as one-time public-private key pairs.

The above exchange protocol assumes that both A and C are able to participate in an online exchange. If C, for example, is unable to participate in an online exchange, because it does not have a bill scanner or does not wish to use it, then A simply unlocks the bill and leaves it in the unlocked state. This can be accomplished with a protocol similar to the locking protocol, namely:

1. $A \rightarrow B$: $\{N, \mu_A, \emptyset\}_{U_{gov}}$
2. B: Retrieve (N, λ), check $\lambda = H(\mu_A)$, store $(N, \emptyset)$.

Evaluation. Given that the only modification of the actual physical currency is the encoding of each bill's serial number in a machine-readable form, online verification achieves the same strong *universal use* property as the barcode signatures, and as far as the production process is concern, only requires a *simple upgrade*. While transfers between participants become more complicated than with standard physical cash, physical digital cash with online verification can still be used by and exchanged with legacy users that do not have the appropriate electronic devices. This also implies that this technique satisfies the *reusability* requirement.

Both the locking procedure described above (and any checks on the lock status) will fail if the serial number provided does not exist, so anyone with a scanner can determine the authenticity of a particular bill, making the currency *forgery-proof*. Since anyone with an online connection can query the lock status of a particular bill, this technique also provides *universal verifiability*. Current smart phones have access to a high-speed Internet network enabling them to establish a secure communication channel with the bank. Short-range wireless communication capabilities can be secured using known techniques [18], and used to transfer bills between participants.

The stored information for each bill consists of the double (N, λ). With about 20 billion bills currently in circulation [32], and the conservative assumption that each double (N, λ) requires 64 bytes, the total size

[3] As before, the bank's public key is U_{gov} and its private key is R_{gov}. These keys need not be identical to the keys used to authenticate bills through the 2-D barcode. The bank's signature on message M is given by $\{M\}_{R_{gov}}$, and public-key encryption of M is denoted by $\{M\}_{U_{gov}}$.

of the database is about 1 TB, a small number compared to other existing highly-available databases like web indexes [5].

Online verification provides a reasonable level of protection against duplication by using a distributed network of verifiers to enforce the principal of *useless duplication*. A participant in the system that receives an unlocked duplicate will immediately lock it, preventing any of the copies from being locked (and hence accepted) by other participants. Transferring a duplicate to another participant has a similar effect. If a forgery does occur, it drives all bills back to the bank, since merchants will not accept duplicates of a bill once the first bill has been locked. This allows easier monitoring and can yield clues for enforcement.

The locking mechanism does potentially introduce new vulnerabilities. Assume that the adversary can create duplicates of existing bills at will. For a nation-scale adversary, this can be done relatively easily, for instance, by asking a large number of people to take pictures of valid bills, or to have a few spies take pictures of a large number of bills stored in banks. Now, consider one legitimate note L with serial number N, and its copy F, which has the same serial number N. L is unlocked as soon as it is passed from a merchant, bank, or individual with the proper equipment to a "legacy principal" which does not have any means to lock bills. The attacker can figure out if L is unlocked by repeatedly trying to lock the note using a null value as the current locking value. As soon as the note L is detected to be unlocked, the attacker issues F. If F is locked before L, L becomes impossible to spend *even though it is a valid bill*. The only way for the unfortunate owner of L to get his money is to confirm with the treasury that L is, in fact, a valid bill, relying on physical features of the bill, e.g., a physical one-way function.

The central bank may then decide to recall the serial number N, but this gives the attacker a way of destroying money, which can lead to sabotage operations. For instance, the attacker may start issuing many copies of bills to disrupt the monetary system by having a large number of users requesting that the treasury check their bills, and having, as a final result, vast amounts of serial numbers destroyed. While the attacker does not gain any money from such a destructive scheme, this type of attack may exert significant pressure on the monetary system targeted.

While potentially serious, this vulnerabilities already exist with physical cash. The presence of an on-line verification system does improve the situation, by making it easier and faster to detect criminal activity. Although the issue of locking a bill held by a legacy principal seems cumbersome at first glance, since the principal will need to deposit the bill at a bank for verification, this action is always due to criminal activity. This should be fairly infrequent, and actually does provide an incentive for people to adopt verification devices.

Finally, one of the most attractive features of physical cash lies in its anonymity. As shown above, we can implement the exchange protocol using only transient random numbers that cannot be matched to any real-world identity. As such, the transfer protocol does not in itself appear to pose any privacy threat.

A thornier issue is that of accesses to the online database. In the exchange protocol we propose, the bank B knows when user A wants to spend the bill N, since A contacts B directly. By extension, as long as the bills are passed between principals that use bill scanners and locking primitives, B has a way of reconstructing the whole transaction chain. Because the communications between A and B never involve the names of the principals (no message include the names A or C), the problem can be solved by using anonymous communication primitives (e.g., [7, 10]) that make it impossible for the bank to identify A. This system could achieve reasonable levels of *anonymity*, possibly at the expense of added latency.

Online verification, thanks to the (un)locking primitives, can also help combat theft. A wallet full of locked bills is useless to a thief. Ownership has not been relinquished, and the money cannot be deposited or exchanged with any participant in the system. Also, the owner of the locked bills retains the serial numbers and unlocking codes for the stolen bills, and can provide this information to the authorities: The thief cannot deposit the money at a bank by claiming to have lost the unlocking codes. These benefits may

encourage adoption, since only participants in the system will have this protection.

5 Security Analysis

The various techniques outlined above for implementing physical digital cash raise a number of questions regarding possible vulnerabilities of physical digital cash.

5.1 Compromised private keys

If the private key R_{gov} used for signing the bills is compromised, physical digital cash is not forgery-proof anymore, and the security level degrades to that of physical cash. Unfortunately, the public may rely on the cryptography as hard evidence that a bill is legitimate, rather than also checking other security signs, such as physical watermarks.

While the issuing government should immediately replace the key pair R_{gov}, U_{gov}, recalling all bills signed with the compromised key may prove problematic. Massive recalls have been shown possible in practice, e.g., by the recent shift from all national European currencies to the Euro, but large-scale recalls are costly, and takes several years to be effective. A possible way to mitigate the risk of a key compromise is to use keys applying to a unique denomination, e.g., $20 bills, produced at a given facility, and with a limited lifetime. Limiting the number of bills involved would facilitate a relatively rapid recall in case of a key compromise.

5.2 Fake signatures

Another class of attack consists of attacks on the signature itself. We are not concerned by cryptographic attacks here, but by physical attacks on the signature information. For instance, fake bills may be produced with missing or incorrect digital signatures. A missing signature is very easy to notice, but while an incorrect signature can be easily detected using a bill scanner, it is not easy to detect in the off-line realm: there is no obvious visual distinction between a good and a bad signature.

Worse, the visible presence of a digital signature (e.g., the presence of a 2-D barcode) may convince users that the bill is good, even in the absence of verification. From a psychological standpoint, a bill may look more trustworthy just because of the apparent presence of a digital signature, even though other physical indicators, e.g., the quality of the paper, or the presence of a watermark, may be questionable.

5.3 Rogue financial institutions

Serious problems may arise when a rogue financial institution (e.g., bank, foreign currency exchange shop) participates in exchanges. One whole class of attacks can be characterized as "money laundering," that is, in the context of counterfeit money, exchanging fake bills for good bills. The simplest instance of such an attack is that performed by a dishonest merchant who tries to pass on bad bills to customers. This type of attack is not new, and in fact already affects the existing physical cash network. Countermeasures are simple: in the physical cash network, individuals are supposed to check the physical properties of a given bill. In the physical digital cash network, individuals can use readers (e.g., applications on their smart phones for barcode signatures, miniaturized RFID readers, etc., depending on the technique employed) to thwart this problem.

A more elaborate version of money laundering involves an attacker colluding with a rogue bank, which cashes counterfeited bills produced by the attacker without checking them. Then, the counterfeited bills are

sent to the currency exchange office of the bank, where they are exchanged for good foreign currency bills from unsuspecting tourists. As long as bills are not verified and no one attempts to lock them, they may travel in the network. Monitoring banks is a plausible countermeasure against such an attack. Compared to the large number of bill users, there are relatively few banks in the world, so a centralized authority (e.g., a treasury department) could monitor them effectively. Recent events [19] indicate that such monitoring already exists in practice.

Another variant on the money laundering scheme is that used by a rogue foreign exchange shop that does not just accept, but also gives out popular foreign currency (e.g., U.S. dollars) in a different country (e.g., Japan). These shops are much less regulated and less controllable than banks. However, for a popular currency, we expect the flow of money to mostly be from the tourists to the foreign exchange shops (e.g., backpackers exchanging US dollars for local currency), so that the impact of this attack should be limited. Further, in all money laundering attacks, counterfeit bills are detected as soon as the bill is deposited at a legitimate institution, or passed to an individual equipped with a bill scanner.

5.4 Localized injection

Massive, localized, injection of forged banknotes may cause serious economic problems if the forgeries cannot be immediately detected. Consider a scenario where an attacker flies a small plane over Manhattan, and drops millions in fake currency over the streets. If the forgeries look real enough, people may be tempted to try to spend this money falling from the sky. Due to the density of population and shops in the area, the impact on the local economy may be significant, which, given the importance of the New York market itself, may have a ripple effect on the national economy.

The only way to counter such an attack is to make the fake bills impossible to spend; that is, to ensure that bills can be immediately verified, and that useless duplication can be readily enforced. Conversely, any method requiring expensive verification devices will have the adverse effect of letting the fake money travel in the network for a longer time period, and possibly to be spent multiple times. Among the techniques we discussed in this paper, inexpensive online verification coupled with a 2-D barcode signature seems more robust against this type of attack than alternative proposals.

6 Conclusion

With the objective to significantly strengthen current bills against government-scale monetary forgery, we establish a set of requirements that are needed for a viable solution. We then look at possible ways to implement these requirements, by augmenting bills with cryptographic material directly embedded in the bill. We consider optically verifiable cryptographic signatures expressed as 2-D barcodes, RFID chips, physical one-way functions, centralized verification and distributed online verification.

None of the techniques we investigate or propose, when used in isolation, satisfies all the properties we would like to enforce. However, a combination of these techniques, for instance, coupling our online verification protocol with optical signatures, and with physical one-way functions serving as back-up, come very close to implementing all the requirements we set out to achieve.

To avoid deployment issues, online verification is designed to accommodate legacy users who do not wish to participate in the online verification scheme. More importantly, deployment needs not be universal. By driving forgeries back to the banks quickly, the proposed system should work very effectively as a deterrent against counterfeiting, even in the absence of wide deployment. Likewise, it is also possible that implementing only a subset of the techniques discussed in the paper may be enough to discourage most

fraud. A design solely based on 2-D barcodes will limit forgeries to duplication of existing bills, and even such duplication would be readily detected.

In that respect, a deeper consideration of the economics at stake in the production and deployment process of counterfeit-resistant bills warrants further research. We believe that our initial approaches will encourage additional efforts in this important area.

References

[1] A proposed federal information processing standard for digital signature standard. Technical Report 910907-1207, RIN 0693-AA86, Nat. Inst. Sci. Tech., August 1991.

[2] N. Asokan, P. Janson, M. Steiner, and M. Waidner. State of the art in electronic payment systems. In *Advances in Computers*, volume 43, pages 425–449. Academic Press, March 2000.

[3] Jean-Paul Boly, Antoon Bosselaers, Ronald Cramer, Rolf Michelsen, Stig Mjølsnes, Frank Muller, Torben Pedersen, Birgit Pfitzmann, Peter de Rooij, Berry Schoenmakers, Matthias Schunter, Luc Vallée, and Michael Waidner. The ESPRIT project CAFE - high security digital payment systems. In *Proc. ESORICS'94*.

[4] S. Bono, M. Green, A. Stubblefield, A. Juels, A. Rubin, and M. Szydlo. Security analysis of a cryptographically-enabled RFID device. In *Proc. USENIX Security*, pages 1–16, Baltimore, MD, August 2005.

[5] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Michael Burrows, Tushar Chandra, Andrew Fikes, and Robert Gruber. Bigtable: A distributed storage system for structured data. In *Proc. ACM/USENIX OSDI'06*, pages 205–218, 2006.

[6] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In *Proc. CRYPTO'88*, pages 319–327, 1988.

[7] David Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Comm. ACM*, 4(2), February 1981.

[8] David Chaum. Blind signatures for untraceable payments. In *Proc. CRYPTO'82*, pages 199–203, 1982.

[9] Benjamin Cox, J. D. Tygar, and Marvin Sirbu. NetBill security and transaction protocol. In *Proc. 1st USENIX Workshop on E-Commerce*, New York, NY, 1995.

[10] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, August 2004.

[11] European Central Bank. Euro banknotes – security features. `http://www.ecb.int/bc/banknotes/security/html/index.en.html`.

[12] Irving Fisher. *The purchasing power of money: Its determination and relation to credit, interest and crises*. New York: Macmillan, 1911.

[13] G. Giovanelli, S. Natali, B. Bozzini, D. Manno, G. Micocci, A. Serra, G. Sarcinelli, A. Siciliano, and R. Vitale. A puzzling mule coin from the parabita hoard: a material characterisation. In *Proc. Cavallino Archaeometry Workshop*, Lecce, Italy, May 2006.

[14] Steve Glassman, Mark Manasse, Martín Abadi, Paul Gauthier, and Patrick Sobalvarro. The MilliCent protocol for inexpensive electronic commerce. In *Proc. WWW'95*, Boston, MA, December 1995.

[15] H. Hoshino, I. Takeuchi, M. Yoda, M. Komiya, and T. Sugahara. Object to be checked for authenticity and a method for manufacturing the same, February 1997. US Patent nr. 5,601,931.

[16] S. Itkin and J. Martell. A PDF417 primer: a guide to understanding second generation bar codes and portable data files. Technical Report Monograph 8, Symbol Tech., April 1992.

[17] D. Kügler. On the anonymity of banknotes. In *Proceedings of the 4th International Workshop on Privacy Enhancing Technologies (PET'04)*, pages 108–120, Toronto, Canada, May 2004.

[18] Jonathan M. McCune, Adrian Perrig, and Michael K. Reiter. Seeing-is-believing: Using camera phones for human-verifiable authentication. In *Proc. IEEE Security and Privacy*, May 2005.

[19] S. Mihm. No Ordinary Counterfeit. *New York Times Magazine*, page 36, July 2006. Issue of July 23, 2006.

[20] Mondex. `http://www.mondex.com/mondex/home.htm`.

[21] R. Pappu, B. Recht, J. Taylor, and N. Gershenfeld. Physical one-way functions. *Science*, 297:2026–2030, 2002.

[22] Raphael F. Perl and Dick K. Nanto. North Korean counterfeiting of U.S. currency. Congressional Research Service Report for Congress, March 22, 2006, 2006.

[23] Ronald L. Rivest. Electronic lottery tickets as micropayments. In *Proc. Financial Crypto.'97*, pages 307–314, Anguilla, BWI, February 1997.

[24] Ronald L. Rivest and Adi Shamir. PayWord and MicroMint: Two simple micropayment schemes. In *Proc. Int'l Workshop on Security Protocols*, pages 69 – 88, Cambridge, UK, April 1997.

[25] Ronald L. Rivest, Adi Shamir, and Leonard M. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.

[26] Bruce Schneier. Renew your passport now! `http://www.schneier.com/crypto-gram-0610.html`, October 2006.

[27] Schweizerische Nationalbank, Banque Nationale Suisse. Die aktuelle banknotenserie. `http://www.snb.ch/d/banknoten/aktuelle_serie/aktuelle_serie.html`.

[28] G. J. Simmons. Identification of data, devices, documents and individuals. In *Proc. 25th Ann. Intern. Carnahan Conference on Security Technology*, pages 197–218, Taipei, Taiwan, ROC, October 1991. IEEE.

[29] Sony Corporation. Overview of FeliCa. `http://www.sony.net/Products/felica/abt/dvs.html`.

[30] J. D. Tygar, B. Yee, and N. Heintze. Cryptographic postage indicia. In *Proc. ASIAN'96*, pages 378–391, Singapore, December 1996.

[31] U.S. Bureau of Printing and Engraving. The new currency – about the new notes. `http://www.moneyfactory.gov/newmoney/main.cfm/currency/aboutNotes`.

[32] U.S. Dept. of Treasury. Treasury bulletin, June 2007. Available from `http://www.fms.treas.gov/bulletin/`.

[33] S. Wiesner. Conjugate coding. *SIGACT News*, 15(1):78–88, 1983.

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