

Reliability of Chip & PIN evidence in banking disputes

Steven J. Murdoch

University of Cambridge Computer Laboratory

<http://www.cl.cam.ac.uk/users/sjm217/>

Smart cards are being increasingly used for payment, having been issued across most of Europe, and they are in the process of being implemented elsewhere. These systems are almost exclusively based on a global standard – EMV (named after its designers: Europay, Mastercard, Visa)¹ – and commonly known as Chip & PIN in the United Kingdom. Consequently, the reliability of the Chip & PIN system, and the evidence it generates, has been an increasingly important aspect of disputes between banks and their customers. A common simplification made by banks when deciding whether to refund a disputed transaction, is the assertion that cloned smart cards will be detected, and that the correct PIN must be entered for a transaction to succeed. The reality is more complex, so it can be difficult to distinguish the difference between customer fraud,² a third party criminal attack, and customer negligence. This article will discuss the situations which may cause disputed transactions to arise, what may be inferred from the evidence, and the effect of this on banking disputes.

The replacement of magnetic stripe cards with smart cards for credit and debit card payments has changed the nature of disputes between banks and their customers over unauthorized transactions. Previously the operation and weakness of cards was well understood, and there was ample evidence of criminal practice. Now, with the implementation of Chip & PIN, the situation has become uncertain: the system is much more complex, the level of security is less clear, and little is known about the capabilities of criminals in terms of committing fraud. This complicates the task of a bank in identifying whether a customer is entitled to be refunded.

Chip & PIN offers greater resistance to fraud when compared with the previous magnetic stripe system, and unlike earlier domestic smart card payment standards, it works across national boundaries. However, the implementation is not infallible, and its complexity increases the likelihood of flaws. In several respects there has also been a trade-off between cost and security, leading to the creation of weaknesses, some of which have been exploited by criminals, some have been demonstrated by researchers, and the remainder are currently assumed to be merely theoretical.

Customers who notify their bank of unauthorized transactions are often recompensed, but sometimes the disputed transactions are not reversed. One possible reason is that the bank believes that the customer authorized the transaction, and is attempting to defraud the bank by making a spurious complaint. Statistics on this type of fraud are not publicly reported by the banking industry, but a fraud investigator working for a major bank, speaking under the Chatham House rule,³ did perceive that levels are high. For example, a group of people have been accused of committing, with the assistance of bank insiders, fraud in the region of US\$422,000, where they opened bank accounts and then claimed their ATM cards had been lost or stolen, and that certain ATM withdrawals were not authorized by them.⁴ The bank may alternatively believe that the customer has acted negligently, in violation of the account terms and conditions, by inadequately protecting their card or PIN, or both their card and PIN. If challenged over such a decision, arguably the bank ought to be required to show that their position is defensible, and that

¹EMV Specifications for Payment Systems, available at <http://www.emvco.com/specifications.aspx>.

²The term “first-party fraud” is used within the banking industry to describe fraud by a customer.

³The Chatham House Rule reads as follows: “When a meeting, or part thereof, is held under the Chatham House Rule, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed. The Chatham House Rule may be invoked at meetings to encourage openness and the sharing of information”: <http://www.chathamhouse.org.uk/about/chathamhouserule/>.

⁴“Gang charged in \$400,000 ATM scam”, Finextra News, 31 July 2009, <http://www.finextra.com/fullstory.asp?id=20328>. See also Stephen Mason, editor, *Electronic Evidence: Disclosure, Discovery & Admissibility* (LexisNexis Butterworths, 2007), 4.04–4.15 for a discussion of cases regarding ATM fraud and banking fraud across the world, including insider fraud. The cases in this text pre-date the introduction of Chip & PIN.

the transaction was not in fact performed by a third-party criminal exploiting a security vulnerability.

The bank's decision will be based on the evidence they have regarding the disputed transaction, the value of the customer's relationship with the bank, and the perceived security of the Chip & PIN system. Much of this evidence will be in digital form, and requires processing and interpretation before it can be understood. While this was also the case with magnetic stripe payment cards, Chip & PIN increases the amount of evidence that could be made available and its level of complexity.

Almost all of this evidence will be held by the bank, as is the information necessary to interpret it. Thus during a dispute, if the bank is required or volunteers to give this evidence to the customer, there will be questions as how to verify the accuracy of the information, and what conclusions can be safely drawn from a forensic analysis. First, this article provides a simplified introduction to Chip & PIN. Then the article sets out the evidence created regarding transactions, and the interpretation of the evidence is discussed to discover whether and how card fraud has been performed.

1 Introduction

In addition to the visible security mechanisms – such as the hologram, embossing, and fluorescent ink – UK credit and debit cards incorporate a magnetic stripe. This stores the data which is visible on the face of the card (name, expiry date, card number, and such like). It also holds the CVV (Card Verification Value; not to be confused with the CVV2, which is printed on the signature strip of the card). Prior to the use of Chip & PIN, the data from the magnetic stripe would be read by the point-of-sale (PoS) terminal or automated teller machine (ATM) and sent to the bank that issued the card to their customer (the cardholder). This bank (the issuer) would be capable of verifying whether the CVV they received corresponded to the one expected for that particular card number. Thus, based only on information which is visible on the card or a receipt, a criminal should not be able to produce a cloned card which evades detection.

The data read from the magnetic stripe only offers assurance that the card is authentic. It is also necessary to confirm that the genuine card-holder has authorized the transactions. For PoS transactions, the cashier would ask the customer for their signature, which they can then compare to the one on the card. ATM transactions are authorized by PIN. Here, the customer enters their PIN at a keypad, which the ATM encrypts and sends, along with the data from the magnetic stripe, to the issuer, potentially via networks operated by parties such as Visa, Mastercard, or VocaLink. The bank can then compare the PIN entered with the one stored in their records.

Magnetic stripe cards have well-known weaknesses. Using commercially available equipment, it is easy to read details from the magnetic stripe of a card, including the CVV, and write a perfect copy of it to a blank card. Such a cloned card would work at an ATM, because only the magnetic stripe is used. Criminals have exploited this weakness in numerous ways, for example adding a “skimmer” to ATMs, which records the magnetic stripe of the card as it is inserted, and incorporates a camera to record the PIN being entered. Together, this yields enough information to make and use a clone in an ATM. To use a clone in a PoS transaction, the visible security features would also need to be copied, which takes more effort but is well within the capabilities of criminals, and has the advantage that the PIN is not required.

The explanation above has been somewhat simplified for brevity. In fact the authorization systems which verify the CVV and PIN can be quite complex, consisting of many components built and operated by different parties; there will also be significant variation between banks and even more between countries. It can be that the issuer does not authorize the transaction at all, but delegates this responsibility to a third party. Card and PIN details are also likely to pass through several different systems between the PoS terminal or ATM, and the authorization system. However, despite this complexity, the cards themselves use the same technology as video and audio tapes, which means there is good intuitive understanding of their main security vulnerability – that if someone obtains possession of the card, even briefly, they can create a perfect copy.

1.1 Chip & PIN

Chip & PIN was designed to mitigate vulnerabilities in magnetic stripe cards, albeit with increased costs, as well as requiring much infrastructure to be upgraded. The cards include a magnetic stripe and the same visible security features as before, but incorporate an additional computer chip underneath the cards' surface. A terminal can interact with the chip through electrical contacts on the face of the card. This chip

is a computer with processing power comparable to desktop computers of the 1980s, but with additional security functionality.

The chip has a program loaded into it, which is designed to follow the communication conventions (a protocol) specified by the EMV documentation, and so be able to communicate with terminals that comply with the EMV standard. This specification is complex, consisting of several thousand pages, but there also will be many thousands of additional pages which describe the design of the chip and its software. National industry bodies and industry members may also extend the specification with additional material.

The chip performs three main operations: card authentication (establishing that the card is authentic), card-holder verification (establishing that the person presenting the card is the authorized account holder), and transaction authorization (establishing that there are enough funds to complete the transaction and the card is not cancelled).

1.1.1 Card authentication

The aim of card authentication is to allow the operator of a PoS terminal (the merchant) to establish whether a card presented is legitimate, without contacting the issuer. This is important because in a small proportion of UK PoS transactions, the terminal is “offline” and does not communicate with the issuer until after the customer has left with the goods. However, since ATM transactions should always be carried out online, card authentication is not performed here. During card authentication at the PoS, the card submits a cryptographic certificate to the terminal, incorporating the card’s account number and a digital signature. The terminal can then check whether this certificate was issued by a bank recognized by a payment system (e.g. Visa or Mastercard) supported by the terminal, and validate the digital signature.

1.1.2 Card-holder verification

Once the merchant is satisfied that the card is authentic, both the card and merchant must be assured that the person presenting the card is the legitimate account holder. This is the role of card-holder verification, which is normally achieved by using a PIN. The customer first enters their PIN on a PIN entry device attached to the PoS terminal or ATM. For PoS transactions, the PIN is sent to the card and the card compares the PIN against the one it stores, and returns the result of the comparison to the terminal. If the PIN entered is incorrect, the card will allow the PIN entry to be re-attempted, but only up to a maximum number of tries – normally three. For ATM transactions, the PIN is not sent to the card, but encrypted and sent back to the issuer, as with magnetic stripe transactions.

1.1.3 Transaction authorization

The final step is transaction authorization, where the issuer, card, and merchant are assured that the card is authentic, card-holder verification succeeded, the card has not been cancelled, and there are adequate funds in the customer’s account. Here, the terminal or ATM sends the card a summary of the transaction (amount, date, and such like). The card appends its own data, such as the result of card-holder verification, and also its application transaction counter (ATC), which is a value maintained by the card, counting how many transactions have been initiated. The card then responds with a cryptographic authentication code. For offline transactions, the authentication code (the transaction certificate – TC) is stored by the terminal for later transmission to the issuer, and the transaction is complete. However, for online transactions, the card sends a different type of authentication code, an authorization request cryptogram – ARQC. The ARQC is sent to the issuer, and it responds with a message stating whether the ARQC is valid, incorporating an authorization response cryptogram – ARPC. Finally, the ARPC is sent to the card for verification, and it responds with a TC indicating that the transaction has succeeded. Alternatively the card can at any time send an application authentication cryptogram (AAC) which means the transaction has been declined.

The issuer and card share cryptographic keys which allow them to generate and verify the cryptographic authentication codes (ARQC, ARPC, TC, and AAC). These keys are loaded during its “personalization” process. However, the merchant does not have these keys, so must rely on the issuer or card to perform the verification. This is because the digital signature keys used in transaction authorization are symmetric, meaning that the same key is used for both generation and verification, and so merchants

could not be trusted with the keys. In contrast, the cryptographic keys used for card authentication are asymmetric, meaning that one key is used for signature generation and another key for signature verification, and it is infeasible to convert the latter key into the former. Thus the merchants are all given a verification key (the public half), but the generation key (the private half) is kept by the bank.⁵

2 Security failures in Chip & PIN

As noted above, the process of a Chip & PIN transaction is much more complex than magnetic stripe transactions. In fact the description above is a simplified version, which shows how transactions should normally happen in the UK; for a variety of reasons the process may diverge from the steps above, and other countries may have different procedures. This complexity is largely for good reasons: the additional verification catches more types of fraud, and so allows transactions to proceed in situations where magnetic stripe cards could not be safely used. However, the complexity also increases the number of ways in which security failures could occur, and makes it more difficult to establish what has happened when they do. This section will summarize some of the potential security vulnerabilities in Chip & PIN, how they may come about, and what their effect might be.

2.1 Card vulnerabilities

While magnetic stripe cards merely act as storage, the security of Chip & PIN depends on the cards implementing a set of security constraints, such as not releasing cryptographic keys or the PIN, and performing card-holder verification correctly. If, due to a bug in the software running on the chip, it is possible to violate these security constraints, criminals could exploit the weakness to commit fraud. Even if the software is correct, before a card can be used, it must be configured during the personalization process. If there is a mistake or oversight in this process, the card may be left in an unlocked state in which some security constraints are not enforced.

Criminals must discover the vulnerabilities in order to exploit them. This may be achieved with the help of an insider, who learns about the vulnerability after the cards with the security vulnerability are already issued. The insider may even create the security vulnerability themselves, by interfering with the software or configuration process, or by disclosing the cryptographic keys needed to unlock a Chip & PIN card. Alternatively criminals could discover vulnerabilities on their own. One technique for doing so is “fuzzing”, where an automated process is used to discover security vulnerabilities. This does not need any knowledge of the software being tested. Fuzzing has been widely used in other contexts by both security researchers and criminals, and is a very effective technique.

A further approach to compromising card security is to attack the chip itself, rather than the software. One set of techniques are known as invasive and semi-invasive attacks, where the chip is removed from the card and manipulated using laboratory equipment.⁶ These techniques can discover confidential information or create carefully chosen failures in the enforcement of security constraints. Non-invasive attacks are also possible, which do not require the chip to be removed from the card. For example, by measuring minute variations in the power consumption of smart cards, it is possible to extract cryptographic keys.⁷ While smart cards do commonly incorporate defences against attacks, they are not always effective, and criminals regularly use these techniques to clone the smart cards used for subscription television.⁸

Regardless of how the criminal has discovered the security vulnerability, if they can extract the card's cryptographic keys used for transaction authorization, they can create a clone of the card which will be undetectable to the bank systems. A criminal does not need to know the correct PIN to use the cloned card for PoS transactions, because the PIN is verified by the card, and it can be programmed to accept any

⁵The details of the cryptography used are not important for the purposes of this article, but for further information on the design and use of digital signatures and authentication codes, refer to Ross J Anderson, *Security Engineering*, (2nd edition, Wiley, 2008).

⁶Sergei P. Skorobogatov, *Semi-invasive attacks – A new approach to hardware security analysis*, University of Cambridge Technical Report UCAM-CL-TR-630, April 2005: <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-630.html>.

⁷Stefan Mangard, Elisabeth Oswald and Thomas Popp, *Power Analysis Attacks: Revealing the Secrets of Smart Cards*, (2007, Springer-Verlag New York, Inc.).

⁸Kevin Poulsen, “DirecTV hacker sentenced to seven years”, *SecurityFocus*, 10 December 2004, <http://www.securityfocus.com/news/10103>; Kim Zetter, “From the Eye of a Legal Storm, Murdoch’s Satellite-TV Hacker Tells All”, *Wired News*, 30 May 2008 (Condé Nast Digital), <http://www.wired.com/politics/security/news/2008/05/tarnovsky>.

PIN. Another way a criminal could use a card would be if the correct PIN could be extracted from a card, or if the PIN stored on the card could be changed without the authorization of the issuer.

Other attacks against Chip & PIN do not require the exploitation of security vulnerabilities at all, but rely on inherent limitations of the cards. One such approach is the “relay attack”, which makes use of the fact that smart cards do not have a display to inform the card-holder which transaction they are authorizing.⁹ The attack works as follows: the card-holder inserts their authentic card into a compromised Chip & PIN terminal, and at approximately the same time, the criminal inserts a special relay card into a real Chip & PIN terminal or ATM. As the relay card is interrogated, it passes on messages to and from the authentic card via the compromised terminal. Thus the real terminal or ATM will believe the relay card is authentic. The customer will think they are authorizing one transaction, but actually the criminal is carrying out a far larger one, potentially on the other side of the world.

2.1.1 Personalization failures

It may not be necessary to compromise the card in order to clone a card, because all the information needed is available at the personalization bureau (where blank cards have keys and customer data loaded), and at the authorization centre (where transaction authorization messages are sent). Personalization and authorization are both performed on behalf of the issuer, but they are commonly sub-contracted (in whole or in part) to specialist service providers. If a criminal is able to interfere with or extract information from either of these processes, they could create a cloned card without having seen the real one.

When considering disputed transactions, banks commonly make the assumption that exactly one copy of each card has been produced. Therefore, if bank records show that the transaction authorization succeeded, then they infer that the particular card issued to the customer was used. The bank may then consider the customer negligent for allowing their card to be used without authorization, and therefore liable for the transaction. However, the assumption that cloned cards cannot exist is not valid, even if it is assumed that the security vulnerabilities above, which allow card cloning, cannot or have not been exploited.

This is because the personalization bureau must have the ability to produce cloned cards, because the process of personalization occasionally fails due to mechanical problems. For instance, the personalization of the chip may have failed, or the printing on the card may be imperfect. An operator should notice the failure, and if they do, they will request that a second card with the same data be produced. Procedural controls should ensure that the damaged card is destroyed, and all cards are accounted for. If these procedures are followed correctly, each customer should receive exactly one card, which complies with quality assurance standards.

However, these procedures occasionally fail. For example, two bank customers have contacted the author to inform him that they each received two identical cards in the post. This is, presumably, due to a technical or procedural failure at the personalization bureau. The author has read the data from the chip on these cards. In one case, both chips appear to contain identical information, including cryptographic keys, and therefore are perfect clones of each other. The customer had used one of the cards successfully, but had not used the second one. In the other case, one chip was functional, but the other was not active. This customer reported that he used both cards for successful Chip & PIN transactions, so it could be that the bank eventually noticed the cloned card and remotely de-activated it. In both cases the cards were visibly identical, had the same information recorded on the magnetic stripe, and had the same details printed on the card, including the CVV2.

In these cases no harm was done, because the legitimate card-holder was sent both clones of the card. However, these instances raise the possibility that a malicious insider could trigger the issue of a cloned card, and retain the cloned card in order to commit fraud. Procedural controls are supposed to stop such activity, but clearly they are not infallible, otherwise cloned cards would not be seen. It is unclear what caused the cloned cards to be sent to these customers, because both had no visible problems. It was confirmed that both chips in one pair worked correctly; for the other pair, the chips presumably worked correctly, otherwise they should not be able to complete transactions. It could be that a software bug or

⁹Saar Drimer and Steven J. Murdoch, Keep Your Enemies Close: Distance Bounding Against Smartcard Relay Attacks, Proceedings of 16th USENIX Security Symposium on USENIX Security Symposium (2007, USENIX Association Berkeley, CA, USA), <http://www.usenix.org/events/sec07/tech/drimer.html>. A demonstration of the attack on a Chip & PIN terminal was also filmed by BBC Watchdog, and aired on 6 February 2007, 19:00, BBC One.

human error triggered the creation of clones, but another possibility was that a malicious insider caused it, but failed to intercept the clone before it was dispatched.

2.1.2 Attacks on hardware security modules

In both the personalization and authorization centres, cryptographic keys and PINs are processed within hardware security modules (HSM). These are computers running specialized software, which will disable themselves should unauthorized interference be detected by their enclosure. Their storage capabilities are limited, and it would be infeasible for them to store separate cryptographic keys for every card that was issued. Therefore a single master key is stored, and a cryptographic procedure called “key derivation” is used to generate a different key for each card. The key derivation procedure takes as input the master key and identifying information of the card (account number, sequence number), and produces a unique derived key (UDK). The procedure is designed so that if a person did find out the UDK for one or more cards, such knowledge would not provide any help in discovering the UDK for any other cards.

The security of HSMs is therefore of critical importance to the integrity of Chip & PIN. In addition to their tamper resistance, the software running in HSMs must enforce security constraints, such as permitting cryptograms to be verified, but not allowing cryptographic keys to be extracted. However, the complexity of the software has led to the discovery of numerous security vulnerabilities. Initially these were found only by academic researchers,¹⁰ but more recently criminals have been exploiting security vulnerabilities in order to commit fraud. In one case reported by Verizon, criminals had extracted cryptographic keys from an HSM and used these to decrypt customer PINs as they were being processed, presumably at an authorization centre.¹¹ If it was possible for a criminal to extract the cryptographic keys used during the personalization or authorization process of Chip & PIN cards, they could create undetectable cloned cards or discover the correct PIN for a card, or both. For example, in February 2008, Citibank reported to the FBI that one of their ATM authorization systems was compromised by criminals, account details collected, and US\$750,000 of fraudulent ATM transactions carried out.¹²

There have been persistent rumours of a system sometimes termed “Bergamot” which is claimed to allow criminals to obtain the PIN for a stolen card. Neither the operation nor the existence of such a device has been verified, although reports of its use exist.¹³ A journalist for ARD Germany also investigated 18 cases of unauthorized ATM withdrawals committed in La Palma in January 2005. In all cases the cards were stolen, but the customers claimed that their PIN was not written down. Nevertheless, the bank records show the correct PIN was used, and the customers were considered liable. One of the criminals responsible (they were convicted in January 2009) was interviewed by a journalist, but refused to say how they used the card without a PIN. Files from the Guardia Civil assumed that the criminals used a “dispositivo” (device) to obtain the PIN, but did not give further details.¹⁴

2.2 Design constraints of EMV

The discussion above has assumed that the authorization system will always detect cloned cards and an incorrect PIN. However, this is not always the case; this section will describe some scenarios in which the authorization system will fail to detect even imperfect cloned cards. Some of the vulnerabilities discussed below may be a consequence of errors made during design and implementation, which remained undiscovered until the system was in use. However, others may have been identified earlier, but permitted remain because the risk of the vulnerability was perceived to be smaller than the cost to resolve it, after other checks and balances were put in place. In designing a system, the bank will try to find an appropriate compromise by applying only those security measures that will at least reduce fraud by their

¹⁰Mike Bond and Ross Anderson, API-Level Attacks on Embedded Systems, Computer, Volume 34, Issue 10, (October 2001), 67–75, available at <http://www.cl.cam.ac.uk/~rja14/Papers/API-Attacks.pdf>.

¹¹Kim Zetter, “PIN Crackers Nab Holy Grail of Bank Card Security”, Wired News, 14 April 2009 (Condé Nast Digital), <http://www.wired.com/threatlevel/2009/04/pins/>.

¹²Kevin Poulsen, “Citibank Hack Blamed for Alleged ATM Crime Spree”, Wired News, 18 June 2008, (Condé Nast Digital), <http://www.wired.com/threatlevel/2008/06/citibank-atm-se/>.

¹³Steve Gold, “A PIN to go with that stolen card sir”, IT Pro Portal, 16 August 2006, <http://www.itproportal.com/security/news/article/2006/8/16/a-pin-to-go-with-that-stolen-card-sir/>.

¹⁴Sabina Wolf, “Sicherheitsrisiko EC-Karten: Wie Banken mit geschädigten Kunden umgehen”, Report MÜNCHEN, 15 June 2009, <http://www.br-online.de/das-erste/report-muenchen/report-sicherheit-eckarten-ID1244812929699.xml>.

cost. For example, cards issued in the UK are vulnerable to attack in offline transactions, but the UK banks agreed that the cost of the more secure cards would be higher than the fraud they would resolve. This is because it is cheaper to put high-value transactions online, and put fraud detection algorithms in place. The problem from the perspective of the customer is that these trade-offs may not be known by the fraud investigation team. For instance, while the weakness of UK cards in offline transactions is well known, other vulnerabilities might exist as a consequence of one department making a cost saving, without informing other departments.

Cost-benefit trade-offs may have been considered during the design of EMV if the vulnerability was identified then, but if the vulnerability was discovered during the implementation phase, a decision would have been made not to fix it, because it was not cost effective for the banks. Other vulnerabilities are not inherent to all systems that implement EMV, but are a property of a particular implementation; again, these may be because of an oversight or due to a deliberate design decision.

2.2.1 Static PIN

One inherent design decision in EMV is to use a single PIN for the card, which must be entered in its entirety. This means that if someone can see the PIN being entered by the customer, and subsequently steals the card, the criminal can easily commit fraud. Other countries, for example Brazil, have adopted a different approach. In addition to the PIN, an ATM prompts the customer for a number of letters from a password. This makes it unlikely that if a person is able to look over a customer's shoulder, that they will obtain enough information to commit fraud.

It may also be possible for a criminal to guess the right PIN for a card. From a single guess, if it is assumed that every combination of PIN is equally likely, a thief who steals a card has a 1 in 10,000 chance of guessing the PIN (and perhaps even more because some PINs are much more popular and the issuer may not permit certain easy-to-guess combinations of numbers). But the thief actually has six guesses because the card permits three tries before the card will lock, and an ATM will permit a further three, making the chance of success 1 in 1,666. If the customer has multiple cards with the same PIN (a practice recommended by banks to prevent customers having to write down their PIN), the odds for criminals can be even better. With four cards in a stolen wallet, the thief could have five attempts on each card without locking them, and thus have a 1 in 500 chance of finding the PIN, and if successful, they will then be able to use all the cards.

These estimates are assuming that each card only has one PIN which will be accepted. This is certainly the case (assuming the card functions correctly) for PoS transactions where the card verifies the PIN. However, this is probably not the case for ATM transactions if the PVV (PIN verification value) technique of verifying PINs is used. This approach is used to reduce the risk that a compromise of the authorization system will lead to PINs being discovered. Rather than storing the PIN, the customer's PIN is encrypted and then truncated to 4 digits of ciphertext – the PVV. The PINs entered at the keypad are also encrypted, truncated, and compared to the PVV. If the correct PIN is entered, the two will match, but if an incorrect PIN is entered, there is still a chance of a match. Most cards will have two or more PINs which will trigger a PVV match, and some will have as many as ten.¹⁵

2.2.2 Yes cards

During the process of card authentication, the card presents a cryptographic certificate to prove that it is a legitimate card. This certificate can be verified with information which is available publicly, and therefore can be carried out even by PoS terminals which are offline. However, in order to allow anyone to verify the certificate, the card must also permit anyone to read its certificate and all the other information it presents. Therefore anyone who can read the certificate can, for most UK cards, produce a cloned smart card that will present identical information, and so pass card authentication. Criminals could produce such a clone by reading data from a Chip & PIN card and writing it to a generic smart card. Equipment and software to achieve this, along with programmable smart cards, are commercially available, and cloned smart cards created in this way have already been found in Europe.¹⁶

¹⁵Mike Bond and Jolyon Clulow, Encrypted? Randomised? Compromised? (When Cryptographically Secured Data is Not Secure), Workshop on Cryptographic Algorithms and their Uses, July 2004, (Queensland University of Technology), <http://www.cl.cam.ac.uk/~mkb23/research/Enc-Rand-Comp.pdf>.

¹⁶Dave Birch, "I didn't want to write about fraud yet again, but...", Digital Money Forum, 15 October 2008, http://digitaldebateblogs.typepad.com/digital_money/2008/10/i-didnt-want-to.html.

For PoS transactions, verification of the card-holder is performed by the card. The terminal sends the PIN entered to the card, and the card responds whether it is correct. Therefore a criminal does not need to know the correct PIN when using a cloned card, because clones can be made which simply respond that any PIN is correct – known as “Yes-Cards”. Clones such as this would not contain the correct keys for generating the ARQC or TC, and so could be detected by the issuer. However, the TC is not sent to the bank until long after the transaction for offline transactions, so by that stage the thief will have left with the goods, although the fraud can be detected afterwards. For online transactions, the incorrect ARQC should be detected and the transaction declined.

Copying the certificate to circumvent card authentication, as described above, is possible in cards which support static data authentication (SDA). As of 2009, most UK cards are of this type, but some banks are distributing out a more secure alternative – dynamic data authentication (DDA). This provides some resistance against card cloning, but was not issued in the UK, in part due to concerns about the increased costs of the cards and longer transaction times. DDA works by adding an additional step to card authentication, where the card proves that it is the legitimate owner of the certificate it presents. This feature requires giving the card an asymmetric key (both the public and private half), and the ability to produce its own digital signatures, which is more expensive, because asymmetric cryptography is much more complex than the symmetric cryptography used in transaction authorization.

2.2.3 The wedge attack

However, DDA does not prevent yes-cards completely, because card authentication can occur before card-holder verification, and so may not include the result of the PIN verification. A simple yes-card cannot be used, because it would fail card authentication, but an alternative technique might still be effective against offline transactions. Here, a stolen card is plugged into a device (a “wedge”) that can modify the data as it flows between the terminal and the card. The terminal is permitted to communicate directly with the legitimate card during card authentication, which will therefore be successful. But during card-holder verification, the wedge suppresses the messages as they are sent to the card and, regardless of the PIN entered by the thief, the wedge tells the terminal that the PIN was correct. The wedge can either pass through the TC from the real card, or create a fake one of its own. In this way, a criminal who has stolen a Chip & PIN card (SDA or DDA) can use it in offline transactions without knowing the correct PIN.¹⁷

The wedge attack also works against online transactions, due to an oversight in the design of the transaction authorization stage. In the EMV specification, the ARQC and TC message includes the result of card-holder verification. However, the result only indicates whether the verification was attempted but failed; it does not distinguish between whether the verification succeeded or whether it was not attempted. Therefore a wedge could suppress card-holder verification, and then relay the ARQC and TC between the legitimate card and terminal. The issuer would receive these cryptograms, and since they were from the legitimate card, the authorization would succeed and the bank would accept the transaction.

This flaw was eventually identified, and banks produced a proprietary extension to EMV which included an additional result in the ARQC and TC, stating whether PIN verification was attempted; a similar extension was later included in the revised EMV specification, but has yet to be widely implemented. However, these only allow the issuer to establish that PIN verification was not attempted; in the wedge attack, the merchant’s PoS terminal will still believe PIN verification succeeded, even though the wrong PIN was entered. A further extension – combined DDA and application cryptogram generation (CDA) – can prevent the wedge attack even in offline transactions by combining card authentication and transaction authorization, but this further extension has yet to be adopted, at least in the UK.

2.2.4 Stand-in authorization

As noted in the discussion above, transaction authorization is of critical importance: it is the only way to reliably detect cloned cards and whether the correct PIN has been entered. Most transactions in the UK (estimates of 80–90 per cent have been given) are processed online. Despite this, the issuer may not process the authorization message, because of the possibility of “stand-in authorization”. Here, if the issuer cannot be contacted in sufficient time, an intermediate party such as the payment system or an outsourced processing centre may authorize the transaction on behalf of the issuer. The party that

¹⁷Chris Mitchell, “Payment and e-commerce applications (Part B2)”, Lecture notes for IY5601, 2005 (Royal Holloway, University of London), http://www.isg.rhul.ac.uk/cjm/IY5601/IY5601_B_06020_83-156.pdf.

provides the authorization is sometimes contractually obliged to accept liability for the transaction if it is fraudulent. However, if there is an equipment failure, it still may be more cost-effective to authorize the transaction and accept the risk without performing all the checks. Issuers may not be aware of their own policy (or that of any outsourced provider) on how authorizations are handled when equipment fails, or the times at which such failures may have occurred. They may even fail to disclose this information to customers who are disputing a transaction.

Where the transaction value is low, and the costs of communications are high, it may be cost-effective to not attempt to contact the issuer at all. This is especially likely to happen in international transactions, but the prevalence is decreasing because of improved reliability and the lower cost of data communications. Each type of intermediate party is able to check different aspects of the transaction. For instance, some have the keys to verify the ARQC and TC, some can verify the PIN (for ATM transactions), some can check if the card is reported stolen, and some may not be able to check any of these. Issuers will have different policies on which types of intermediate party is able to perform stand-in authorization. In addition to establishing liability, contracts will also impose service level agreements that will set out the speed by which an authorization message must be processed, and in such a manner control the circumstances in which stand-in processing is appropriate.

2.2.5 Fallback

Another set of vulnerabilities exist because the magnetic stripe system is still operational, even with Chip & PIN cards. UK cards continue to have magnetic stripes, to enable them to work in terminals and ATMs without chip readers (e.g. outside the UK), or when the chip or chip reader has failed. UK PoS terminals also have magnetic stripe readers for use with foreign cards or as a backup when the chip cannot be read. This means a criminal who cannot clone a chip can simply copy the magnetic stripe from a smart card, and produce a magnetic stripe clone. From the perspective of an ATM or PoS terminal, this clone will appear to be a legitimate card, but the chip on the card might be damaged, or the chip reader in the terminal might have failed. Since chips regularly break and chip readers frequently get dirty and fail, this is not very suspicious, and the transaction may be permitted to proceed regardless. This is known as a “fallback transaction”.

The criminal does not have to read the magnetic stripe to clone the card, because the chip contains a copy of the data on the magnetic stripe. This data is also commonly sent to the issuer during a transaction. Therefore a criminal who can read the chip or intercept the communication between a terminal and the issuer, can also copy the magnetic stripe. A criminal who can intercept the communication between the PoS terminal and chip can copy the same data, and also can obtain the PIN entered by the customer, as it is sent to the chip during card-holder verification. PoS terminals have tamper resistance measures to prevent this, but due to design errors, it is quite simple to circumvent the protection in place and connect a “tap” built with off-the-shelf electronic components.¹⁸ This device reads all the information necessary to produce a cloned magnetic stripe card and use it in an ATM. Chip & PIN terminals have even been discovered with taps having been added during or soon after manufacture.¹⁹ For these reasons, more recent cards do not store the full CVV on the chip, instead replacing it with an alternative termed the “iCVV”.

This general approach has been widely exploited for committing both fraudulent ATM and PoS transactions. For instance, Maxwell Parsons was convicted in November 2006 of having collected card details by connecting a MP3 player to the back of ATMs. With this information he was able to produce cloned cards, and use them to perform unauthorized transactions.²⁰ In October 2008, Anup Patel was convicted of committing fraud to the value of £2 million. This was achieved by defeating the physical protection put in place to protect Chip & PIN terminals, and to record both PINs and card details. This attack was

¹⁸Saar Drimer, Steven J. Murdoch, and Ross Anderson, Thinking Inside the Box: System-Level Failures of Tamper Proofing, Proceedings of the 2008 IEEE Symposium on Security and Privacy, (2008, IEEE Computer Society Washington, DC, USA), 281–295, available at <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-711.pdf>. A demonstration of this attack on a Chip & PIN terminal was filmed by BBC Newsnight, and aired on 26 February 2008, 22:30, BBC Two.

¹⁹Henry Samuel, “Chip and pin scam “has netted millions from British shoppers””, The Daily Telegraph, 10 October 2008, <http://www.telegraph.co.uk/news/newstopics/politics/lawandorder/3173346/Chip-and-pin-scam-has-netted-millions-from-British-shoppers.html>.

²⁰“Cash machine bug scam expert jailed”, Manchester Evening News, 15 November 2006, http://www.manchestereveningnews.co.uk/news/s/228/228286_cash_machine_bug_scam_expert_jailed.html; Stephen Mason, editor, Electronic Evidence: Disclosure, Discovery & Admissibility, 4.10.

so successful that it enabled cloned magnetic stripe cards to be produced.²¹

2.3 Back-end failures

The description of the vulnerabilities in the section above assumed that the back-end systems controlled by the card issuer, together with other processing infrastructure, operate correctly. However, if the bank systems are not perfectly designed and correctly operated, these assumptions will not be true. It is likely that there may be weaknesses in the card processing infrastructure because of the complexity of the system, and because it is continually being upgraded to accommodate new equipment and additional operational requirements. A recent illustration of a failure was demonstrated where a couple in Essex, UK, discovered that they could withdraw cash from a particular ATM without the transaction being recorded against their account. Even though these transactions should have been declined because the couple's account was overdrawn, a failure at some point in the processing allowed them to be accepted. This failure eventually became public when the couple were convicted, having withdrawn over £61,000 in this way.²²

For online transactions, if the ARQC or TC is wrong, the issuer should decline a transaction, and for offline transactions an incorrect TC should be detected when the terminal goes online at a later time, the fraud discovered, and the customer refunded. In this way, customers should not lose money from the use of yes-card attacks, although for offline fraud, the merchant will probably have to pay once the fraudulent transaction is reversed. But if the issuer fails to detect an ARQC or TC which was generated with the wrong key, or contains the wrong information, the use of a yes-card or wedge attacks will not be detected for online or offline transactions. This failure could occur simply because of a programming error, but it could also be an intentional decision; for example if the HSMs which validate the ARQC and TC are overloaded, the issuer may decide to accept transactions without checking. Also, in some circumstances, the ARQC or TC will be corrupted before being sent to the issuer; in these situations the issuer may decide to accept the transaction rather than risk insulting the customer by declining it, and accept the risk of fraud. Such corruption can occur due to random errors, or because of a format translation error such as the one believed to be the reason that Visa debited customers' accounts by US\$23 quadrillion.²³

Because almost all UK cards, ATMs and PoS terminals have been upgraded to support Chip & PIN, it is common for the issuer to automatically decline fallback transactions on cards which have a chip when used in the UK. However, a failure in processing may also allow a fallback transaction to succeed when it should be declined. This could be because the issuer is informed that a fallback transaction has occurred, but a software bug causes it to be accepted. Alternatively, a bug in the ATM or PoS terminal may cause it to identify a fallback transaction as a chip transaction, and the authorization system does not decline the transaction on the basis of it having a missing or incorrect ARQC or TC. Alternatively, a criminal could modify the "service code" on the magnetic stripe to indicate the card does not have a chip, and hence a fallback transaction should be permitted. The issuer should detect the tampered service code, but in 2005 someone working for the London Programme made a magnetic stripe clone or a smart card, altered the service code, and successfully used it for an ATM cash withdrawal.²⁴

2.3.1 Logging failures

Most banking systems produce extensive log files that record their actions, which makes it easier to identify malfunctions and understand the cause. When transactions are disputed by the customer, these logs may be examined. However, the systems which produce the log files are complex, and their output often requires further processing before it can be easily understood. It is therefore common to have reporting systems, which take the raw log input (potentially from multiple sources), interpret them, and produce a new file which is intended to be easier to understand. A failure in either logging or reporting systems could cause the result of a transaction to be misinterpreted; for example the operator may believe it to be a Chip & PIN transaction when in fact it was fallback. If a malicious person has gained access to logging

²¹Steve Bird, "“Catch me if you can,” said student behind biggest chip and PIN fraud", The Times, 19 October 2008, <http://www.timesonline.co.uk/tol/news/uk/crime/article5034185.ece>.

²²Stephen Bates, "Couple who took £61,000 from faulty ATM sentenced", The Guardian, 21 April 2009, <http://www.guardian.co.uk/uk/2009/apr/21/cash-machine-theft-essex>.

²³Dan Goodin, "Reg readers crack case of the \$23 quadrillion overcharge", The Register, 16 July 2009, http://www.theregister.co.uk/2009/07/16/visa_programming_error_cracked/.

²⁴Chip and PIN security flaw uncovered, London Programme, ITV1 London, 15 March 2005, 19:30–20:00.

or reporting systems, they could also tamper with the result in order to cover their tracks, because these systems are commonly less well protected than authorization systems.

Even after a reporting system has processed a log file, it can still be difficult to interpret the output. Output is generally presented using terse codes, and their meaning must be found within documentation. Sometimes this documentation is not available, or it may be out-of-date following changes to the system concerned. Therefore, the operator may interpret them by comparing the log output with similar output observed in the past, and then they may draw conclusions. For example, in the case *Job v Halifax plc*,²⁵ the witness for the bank examined the format of the log entry of the disputed transaction, and pointed out that it was similar to other legitimate Chip & PIN transactions, and different from other legitimate fallback transactions. From this, the witness inferred the disputed transaction must have been a legitimate Chip & PIN transaction. The bank did not refer to any documentation on the meaning of the data, or discuss what the log entry would look like should one or more security checks have failed.

2.3.2 PIN verification

As discussed with respect to the yes-card attack, for PoS transactions, the card is responsible for verifying the PIN, and if a cloned card is used, the criminal need not know the correct PIN. In contrast for ATM transactions, the PIN is sent back to the issuer or a stand-in processor for verification. Even so, the criminal would not need to know the correct PIN if a stand-in processor that cannot verify the PIN authorizes the transaction. If there is a malfunction which allows transactions to be authorized if the PIN verification is not attempted or fails, a card could be used without the correct PIN. An insider may also try to trigger such failures, for example by gaining access to the authorization system.

2.4 Terminal failures

The discussion above has been about failures of the processing and authorization systems. These are very important, because the correct functioning of these systems is of critical importance to the integrity of the Chip & PIN system. PoS terminals and ATMs are relied upon to a lesser extent because they are under the control of potentially untrustworthy merchants, and their correct functioning cannot be guaranteed. It is for this reason they are tamper resistant, to prevent malicious people from extracting confidential information (although these measures can easily be overcome as noted above, and criminals have been caught doing so). Nevertheless it is still possible to commit fraud because the terminal fails to operate properly.

During transaction authorization, the PoS terminal or ATM generates an unpredictable number. The number is sent to the card, and incorporated into the cryptographic process which generates the TC and ARQC. If this number is predictable, a criminal could clone a card by asking the legitimate card for a number of TC and ARQC cryptograms, then writing these values to a generic smart card. This clone could then be used for a transaction, and provided the thief guessed a correct value for the unpredictable number, the clone can produce a TC and ARQC which will pass the check by the issuer. In this way, the criminal can put through online Chip & PIN transactions at both PoS and ATMs, given only temporary access to the legitimate card. The criminal may, however, need to know the correct PIN if cryptograms are required that indicate that card-holder verification succeeded.

While there are well established techniques for securely generating unpredictable numbers, it is notoriously difficult to verify whether a generator is working correctly, so failures do regularly occur. For example, one version of Linux had a feature which was supposed to generate unpredictable random numbers, but was in fact relatively easy to predict. This flaw was introduced in 2006, but remained undetected until 2008.²⁶ Linux is used in both ATMs and PoS terminals, but it is not clear whether such devices ran an affected version. A criminal can also tamper with an ATM or PoS terminal to reduce the unpredictability of the random number generator. Research has shown that it might not even be necessary to open the device to do so; manipulating the power supply or transmitting a radio signal may be sufficient.²⁷

²⁵*Job v Halifax plc*, Nottingham County Court (case number 7BQ00307), 30 April 2009, the judgment is published in *Digital Evidence and Electronic Signature Law Review*, 6 (2009).

²⁶Robert Jaques, *Debian flaw exposes communications breakdown*, V3, 28 May 2008, (Incisive Media Ltd), <http://www.v3.co.uk/vnunet/news/2217710/linux-security-flaw-should-wake>.

²⁷A. Theodore Marketos and Simon W. Moore, *The Frequency Injection Attack on Ring-Oscillator-Based True Random Number Generators*, Workshop on Cryptographic Hardware and Embedded Systems, LNCS 5747, September 2009 (Springer).

2.5 Whistleblowing and insiders

Examples of security failures in banking systems are hard to find for a variety of reasons, such as the restrictions imposed by non-disclosure agreements; the banks are reluctant to admit vulnerabilities and the complexity of systems, thus making it challenging to discover vulnerabilities in the first place. The examples discussed above became public either because a customer noticed, or because of legal proceedings. In other fields, most notably safety critical systems such as aerospace and medical equipment, there are legal requirements on companies to report failures, and to investigate serious failures. The banks are not subject to such requirements; only whistleblowers and researchers acting outside the banking system can notify the public of problems.

However, there are substantial obstacles to this. For example, a French engineer, Serge Humpich, discovered a way to make forged banking smart cards, and reported this to the banks involved. Having demonstrated his technique worked by purchasing ten Paris Metro tickets at the request of the banks, he was arrested and convicted for counterfeiting.²⁸ Also, a journalist in the UK whom the author assisted with reporting on vulnerabilities in Chip & PIN was threatened by his own bank that they might cancel his mortgage (though the bank in question eventually withdrew the threat).²⁹ Cases such as this create a chilling effect, preventing people from testing the existence of vulnerabilities or reporting those they become aware of.

Banking insiders who are aware of security weaknesses may decide to exploit the vulnerability for fraudulent purposes. They may discover the problem directly, or be notified of it in a security testing report. Discovering a vulnerability does not, however, require insider information; while one of the criminals in the Essex case referred to above worked for a bank, it is believed they discovered the vulnerability by accident. The ATM in question was old, which offers a possible explanation as to how the fraud happened. In order to integrate legacy infrastructure with new systems, an established technique developed by Brodie and Stonebreaker is to build a “gateway”,³⁰ which translates between the old and new conventions, but potentially loses information in the process. This component is then modified to fix bugs until it passes the necessary tests.

Integrating legacy systems is notoriously difficult, and tests cannot be guaranteed to find every problem. Although it can only be hypothesized that this class of flaw allowed the fraud, another example of where integrating a legacy component causes failure is with the Ariane 5 satellite launch vehicle. Here, despite extremely rigorous testing, an integration problem between software originally designed for Ariane 4 and the Ariane 5 navigation system remained undiscovered, until it caused a failure soon after launch, destroying the rocket.³¹ The writers of the software made assumptions that were true at the time the component was designed, but were invalidated when the surrounding system was upgraded, leading to its catastrophic failure.

3 Evidence in Chip & PIN cases

When a transaction is disputed, there may be disagreement between the customer and the bank as to who should be liable for it. A common example is where the bank believes that a customer’s real card and PIN have been used, and hence argues that either the customer performed the transaction (and is attempting to defraud the bank), or has been negligent in protecting their card or PIN or both card and PIN. The customer may believe that they did not carry out the transaction, and they were not negligent with their card or PIN, and argue that the transaction is due to an error having been made by the bank, or a security vulnerability having been exploited by criminals. Evidence may be requested to corroborate each party’s position, but drawing conclusions from it must be performed with care. The evidence itself could be insufficient, and mistakes in interpretation might be made. Also, bank employees might try to cover up embarrassing security failures, and criminals may attempt to tamper with evidence.

²⁸Cedric Ingrand, “French credit card hacker convicted”, The Register, 26 February 2000. http://www.theregister.co.uk/2000/02/26/french_credit_card_hacker_convicted/.

²⁹Personal communication.

³⁰Michael Stonebraker and Michael L. Brodie, *Migrating Legacy Systems: Gateways, Interfaces & the Incremental Approach*, (Morgan Kaufmann Publishers, 1995).

³¹Professor Jacques-Louis Lions and others, Ariane 501 Inquiry Board report, ESA, 19 July 1996: <http://esamultimedia.esa.int/docs/esa-x-1819eng.pdf>.

Evidence can be collected from all the systems that have been discussed in this article, commonly in the form of log files. These include the manufacture and personalization facilities, where the chips are produced, placed on cards and loaded with data. The PoS terminal or ATM will also contain a log, generally on paper and informally called the “till-roll”, which records transactions and other important events. The card itself also contains useful information, such as the ATC, but depending on the bank, there may also be summary information available about transactions. The most important logs are kept at the authorization centre, recording the type of transaction and result of authorization. However, because there is potential for errors in all of these items due to mistakes or tampering, it is prudent to collect as much evidence as possible in order to show that records are consistent.

3.1 Audit and compliance

Another way to establish the reliability of evidence is to examine whether the system producing the logs was operating correctly. This is commonly achieved through an audit, where experts (possibly internal to the bank or external) will examine documentation or the system itself (or both the documentation and the system), and check it against requirements. The result of the audit is a report, which can be very informative as to the dependability of the system. If any potential problems are identified, it will highlight these, and where an external audit is undertaken (by a payment system for instance), the auditor may require the bank to undertake changes. However, if a vulnerability discovered by a payment system is considered to only affect the bank itself and not other members of the scheme, it may be possible for the bank not to deal with the problem and accept the risk.

Not only is the report itself important, but also the changes which were carried out in response to the report. For each change, there should be information on how and when the modification was applied, and how it was established that the modification properly fixed the issue that was identified. The methodology for performing the audit is also significant, because approaches vary in how effective they are at identifying problems. One very effective technique is penetration testing, where a skilled team are given access to the system and given the task of finding security vulnerabilities in any way they see fit. At another extreme, some audits only examine documentation and not the system itself, and so will miss implementation errors. However, regardless of the methodology, audits do miss critical vulnerabilities; for example the OpenSSL cryptographic library was subjected to an extensive audit under the FIPS 140-2 scheme, and passed, even though a serious security vulnerability existed in the random number generator.³²

Logs of changes to bank systems are important, even if modifications were not the result of an audit report. This is especially true if the modification was known to have an effect on security, but even apparently innocuous changes can cause security vulnerabilities which are hard to identify. Obtaining documentary evidence is important because there may be a dispute, even within a bank, as to when a particular change is made. For example, in the case of *Job v Halifax plc*, a number of dates were given as to when magnetic stripe transactions were disabled. The witness for Halifax, Ian Brown, stated in paragraph 5.4 of his statement dated 6 February 2008, that “If the transaction is presented in “fallback” mode, then the transaction will be declined.” The inference of this comment was that Halifax had disabled fallback before February 2006 (the date of the disputed transactions). The expert witness for Halifax plc, David Baker (Head of the APACS Cards Technical Unit), stated in paragraph 5.1 of his second expert’s report dated 14 February 2008 that “To our knowledge all UK issuers will routinely decline transactions flagged as magnetic stripe read and have been doing so since 2005.” When Mr Brown gave evidence, his barrister questioned him about the statement he made in paragraph 5.1, and he amended his evidence to the effect that the comments he made were correct in September 2006. The author has confirmed that some banks permitted fallback as late as May 2007.³³

One explanation for this discrepancy is that there can be a delay between when a change is mandated, and when it is actually applied; this has also been seen with iCVV, which APACS announced to be fully in place by 2008:³⁴

“All UK issued cards issued after 1 January 2008 include an updated iCVV (Integrated Circuit Card Verification Value) which means that if one of these cards were compromised in the method described, the

³²OpenSSL FIPS Object Module Vulnerabilities, 29 November 2007, http://www.openssl.org/news/secadv_20071129.txt.

³³Demonstration on ITV Manhunt, aired on ITV1, 29 May 2007.

³⁴Statement from APACS, in response to BBC Newsnight, 26 February 2008, <http://news.bbc.co.uk/1/hi/programmes/newsnight/7265888.stm>.

data would be useless to the fraudster (i.e. a fake magnetic stripe card created via a compromise of this type would not work in a cash machine, even overseas in a non-chip and PIN country).”

The author has confirmed that several banks, including Halifax Bank of Scotland (with a card issued in March 2008) and Barclays (with a card issued in February 2008), have not implemented iCVV, and they were still issuing cards that do not comply with the iCCV standard past this date.

Before a card transaction can take place, the chip must be manufactured and the software loaded. Logs from this process will be useful to establish which version of the chip hardware and EMV software was used. Internal audit reports may indicate if any versions were known to be vulnerable to attack. Even if these audit reports do not exist, it would be informative to examine the “change log” documentation accompanying the chip software, as this should indicate the differences between versions. If this documentation indicates that a new version of software was released in order to fix a bug in a prior version, the possibility that such a bug would allow a fraudulent transaction should be investigated.

The personalization process is another area in which logs would be valuable, for example whether a cloned card was produced because the first one might apparently fail quality assurance. Procedures should also be examined, to ensure that cryptographic keys are being safely handled. Also, audit reports and change logs for the software which configures cards should be examined, to ensure it properly locks the card to protect confidential data.

Logs of the transaction and authorization process will be available from a number of places, and the information they contain will be somewhat different. The ATM or PoS terminal will be able to indicate whether the transaction that is in dispute actually took place, and how it was authorized. However, interpretation of these can be difficult. For example, in the Essex fraud case, the bank that operated the ATM originally believed that members of staff were stealing the money. This suspicion must have arisen because logs are maintained of how much money was loaded into the ATM and how much was withdrawn, and the totals were inconsistent. In fact, the CCTV surveillance put in place to catch the thief, actually showed the couple who were later convicted.

During the authorization process, messages are sent via a number of intermediate parties: the acquiring bank who is contacted by the merchant; the issuer (or stand-in processor) which generates the authorization; and the payment system which allows acquiring banks and issuers to communicate and transfer funds. All of these parties probably keep logs, especially the payment system. This is because they are responsible for providing the communication infrastructure, and they also offer dispute resolution services between their members. In these cases, where there is an inconsistency in the records of two members, the payment system can examine their logs to establish the facts. The logs of payment systems are particularly valuable in this case, because they are from a neutral party. Similarly, in the case of customer disputes, collecting logs from a third party increases the chance of detecting insider attacks.

Transaction authentication is the most important step of the EMV transaction, and for this stage the logs are kept with the issuer. These should include a description of the transaction, the result of authorization, and the cryptograms involved (ARQC, TC, and ARPC). The description should include all the usual information, such as the amount and date, but also will have the result of card-holder verification, and importantly whether PIN verification was attempted and whether it succeeded. Two versions of this are given: one by the terminal in its description of the transaction, and one by the card as part of the ARQC and TC, in its issuer application data section (IAD). These should be compared to establish consistency. Because of the cryptographic processing, these logs can be subjected to enhanced verification, but they do not replace the other logs discussed above, because they do not contain all the necessary information and still can be tampered with.

3.2 Principles for design of secure systems

In the field of computer security, the Trusted Computing Base (TCB) is the part of a system which must be relied upon in order for the overall system to function securely. A widely accepted principle of security engineering is to minimize the size of the TCB, in order to improve the robustness of the system.³⁵ Following this principle means that during the design and implementation of a system, the available testing resources can be focused more intensely on the TCB, increasing the chances of identifying bugs. Additionally, this principle aids forensic analysis, because if a component can be shown to be outside the TCB, there is no

³⁵Butler Lampson, Martin Abadi, Michael Burrows and Edward Wobber, “Authentication in Distributed Systems: Theory and Practice”, ACM Transactions on Computer Systems, Volume 12, Issue 1 (February 1994) (ACM Press), <http://research.microsoft.com/en-us/um/people/blampson/45-authenticationtheoryandpractice/acrobat.pdf>.

need to waste effort in establishing whether it is functioning correctly. In EMV, there is no clearly defined TCB, but analyzing the system from this perspective is a helpful way of deciding what system components should be examined and what they can be relied upon for.

In disputed transaction cases, the issuer will typically have almost all the evidence that is presented to the court or adjudicator. Sometimes audit reports are made public, as occurs for the banking smart cards issued in Germany and evaluated under the Common Criteria scheme.³⁶ However, in banking, it is more common to keep audit reports and system documentation secret than in other areas of security engineering.

In discussing what evidence should be presented, there may be a question as to whether a bank, by giving an opposing expert witness access to an item of information, would harm the security of the system. An accepted best practice in security engineering is that the security of a robust system should not depend on secrecy of its design. This is because it is difficult to keep design documents secret, and if the detailed functionality of a system cannot be described, questions may be raised as to whether it is in fact secure. It is for this reason that it is common to openly publish details of security systems, often including the source code from which they are built. Even if source code is not published (e.g. Microsoft Windows), lists of known security flaws are publically available.

This practice is historically known as Kerckhoffs' principle,³⁷ where it was applied to military communication systems. With respect to banking systems, the same principle is described by APACS (the UK banking industry representative body), in their PIN Administration Policy:³⁸

“The PIN Administration process must not only be secure, but also be demonstrably secure. If PIN Security is publicly challenged, either in the media or in a court of law, it must be possible to respond to such a challenge and for the response to be supported with evidence. Furthermore, the use of that evidence in the public domain must not in itself compromise security.”

3.3 Verifying authorization logs

If there is a disputed Chip & PIN transaction, it can safely be assumed that the bank authorization system shows that the correct card and PIN were used (otherwise the customer would be immediately refunded). In which case, the next step in examining the evidence would be to establish whether these logs can be relied upon in concluding that the customer's card and PIN were used. Some generic approaches have been described above, such as corroborating different items of evidence and examining documentation relating to the systems relied upon. However, there is one particularly useful set of techniques which can be applied to authorization system logs, because the EMV cryptograms act as a audit log, allowing their authenticity to be established without having to rely on the authorization system.

3.3.1 Validating the ATC

The simplest item to validate is the ATC, which is sent along with each cryptogram. It will therefore be stored in the authorization system, and may also be recorded by the payment system and at the PoS terminal or ATM. The ATC is a number stored by the card, and incremented by one each time a transaction is initiated. Therefore, logs of transactions should show the ATC increasing by one for each transaction, in chronological order of the transaction time. The ATC may pass over values if a transaction is initiated but aborted before the cryptogram is sent to the issuer, but it should never decrease. Large jumps should be viewed with suspicion.

It is important to examine the ATC sequence for both disputed and non-disputed transactions, because if a clone is being used, and a criminal is not very careful, there will be inconsistencies in the pattern. For example, suppose the criminal creates a cloned card and uses it for a transaction. If the ATC produced for this fraudulent transaction is lower or equal to that of the last legitimate one, logs of ATC values would show up a discrepancy. Even if the criminal is able to guess the correct value of the ATC to use, the logs will still show a discrepancy when the customer next uses the legitimate card, unless it happened to leave out a value due to an aborted transaction.

³⁶Certification Report for ZKA SECCOS Sig v1.5.2, BSI-DSZ-CC-0341-2006, 13 June 2006 (BSI), <http://www.commoncriteriaportal.org/files/epfiles/0341a.pdf>.

³⁷Auguste Kerckhoffs, “La cryptographie militaire”, Journal des sciences militaires, 9 January 1883, <http://www.petitcolas.net/fabien/kerckhoffs/>.

³⁸APACS PIN Administration Policy, January 2004 (APACS), http://www.apacs.org.uk/resources_publications/documents/PIN_Administration_Policy.pdf.

While the authorization system should detect grossly irregular sequences of ATC values, when investigating disputed transactions, it is advisable to perform more rigorous examination of the information than the authorization system would normally perform. This is because criminals will generally attempt to circumvent the fraud detection measures, but no more (so as not to waste effort). If the process of analyzing logs for disputed transactions is merely to repeat the same checks which it would have had to pass in order for the transaction to succeed, no new type of fraud would ever be detected. Authorization systems might also not enforce tight constraints; for example the author has tested cards which have worked despite large gaps in the ATC sequences.

However, the criminal can still circumvent the process if he has access to the legitimate card. First, the criminal uses the cloned card a few times while the customer is not using the legitimate one. Then the thief obtains the legitimate card, increments the ATC the same number of times that he used the cloned one, adds a few more additional ones, and then returns it to the customer. Finally, the criminal can use the cloned card more times, provided that its ATC remains less than the one he set on the legitimate card. In this way, the thief can interleave two groups of fraudulent transactions, without causing disruption to the pattern of ATC values. With more regular access to the legitimate card, the criminal could effect further fraudulent transactions.

EMV cards can, optionally, contain a record of the ATC value when the card last successfully completed online transaction authentication. This value can be used to help detect whether a criminal has incremented the ATC as described above; in such a case, there would be a significant gap between ATC and the last online ATC. Many UK cards have this feature enabled, and it has proved a useful forensic tool. Another optional feature which would be especially useful for investigating disputed transactions is the transaction log. Here, the card maintains a record of recent transactions, and will return the list when requested. Unfortunately, the author is not aware of any UK bank which has adopted this feature.

Even without the optional additions, the ATC is a useful tool in validating transaction logs, and the interleaving of disputed transactions with non-disputed ones, with a consistent ATC pattern, was used by the First Trust Bank as evidence against their customer in a disputed ATM withdrawal case.³⁹ While the ATC logs are held by the bank (and potentially other parties), the customer can partially validate this information himself, because ATC values are sometimes printed on receipts. Additionally, if the customer has retained the card which his bank states was used for the disputed transactions, he or someone acting for him can read the current ATC value using specially designed software. The author has attempted to do this in three cases so far, but in two the issuer instructed that the card be destroyed (in one case by the customer, and in the other by the bank which had retained the card in an ATM), and in the third, the bank sent a message to the card instructing it to permanently disable itself before the author could obtain access to the card.

3.3.2 Validating the cryptogram

A further item that can be validated is the cryptogram (ARQC or TC or both). First, having a cryptogram contributes towards evidence that it was a Chip & PIN transaction, not fallback. The transaction data which accompanies the cryptogram includes the type of transaction, date, value, etc. as seen by the card, which should be compared against the version that was sent to the issuer. Most important is the IAD, which is generated by the card and incorporates details on whether the PIN was entered correctly, and if the card has detected any unusual activity. This is the only way to verify whether card-holder verification succeeded; because of the wedge attack, the PoS terminal may have been misled. The detailed meaning of the IAD is specific to the issuer, but it generally follows one of the standards produced by Visa, Mastercard, or the EMV consortium.

However, the records of both the IAD and ATC could be manipulated by the authorization, reporting, or logging systems and networks, so they cannot be trusted unless the reliability of these can be assured. But following from the principle of minimizing the trusted computing base, it is possible to eliminate consideration of these systems by validating the authentication code using independently implemented software, based on the public standards for cryptogram generation (such as one written by the author).⁴⁰

³⁹The author assisted in this case by attempting to read the ATC from the card; however the bank had electronically disabled the card before returning it to the customer. The customer did not take the case further than the bank's internal dispute resolution process.

⁴⁰The author wrote this software in order to be able to verify any cryptograms that might have been produced in *Job v Halifax plc*. It is not, as yet, publicly available.

Checking a cryptogram requires the UDK of the card, which needs to be kept confidential while the card is active, but after the card is cancelled it can be safely disclosed. This is because knowing the UDK of one card is of no assistance in discovering the UDK of another. This key could, for example, be obtained by requesting the HSM, which generates keys for personalizing newly issued cards, to generate a key for just one card. The key can also be validated by checking an ARQC generated by the card (if the customer still holds it), or receipts which show the ARQC or TC.

4 Nature of disputes

From the above description, it is clear that the complexity of EMV substantially changes the nature of disputes between customers and banks over unauthorized transactions. While the addition of cryptography offers greater resistance to fraud, this also makes it more likely that customers will be denied refunds by their bank. Not many of these cases make it to court in the UK, because the sums the customers claim for are typically a few hundred to a few thousand pounds, and the claimant risks an order to pay costs that can be significant, should they lose. For example, in *Job v Halifax plc*, the disputed transaction was £2,100, but the bank proved their case to the satisfaction of the judge, and Mr Job was ordered to pay £15,000 in costs. Prosecutions also occur, such as that of Jane Badger, who disputed a transaction and was subsequently charged with making a false statement. She was acquitted, but at the time of writing, the bank (Egg) continues to refuse to refund the disputed transaction.

Despite only a few cases making it to court, the consumer rights organization Which? reports that 20 per cent of customers are not refunded after claiming to be the victim of fraud.⁴¹ There are difficulties with the way customers can seek a resolution in respect of disputed transactions. Initially, they have to defer to the bank's internal dispute resolution process, and then consider adjudication by the Financial Services Ombudsman. However, the customer is in a fairly weak position, because neither the bank nor the Financial Services Ombudsman produces the evidence, unless the customer makes a request under the provisions of the Data Protection Act 1998. The Banking Code is also not very helpful. It states that banks are liable for fraudulent transactions, but this only applies if the bank believes the customer has been either negligent nor is acting fraudulently. A common position taken by banks over disputed transactions is that if a transaction is Chip & PIN, and it does not match the standard patterns of known frauds, then the customer is considered liable. However, the criteria banks use for identifying patterns are not subject to public scrutiny, and may vary between banks and individual fraud investigators.

Another frequent problem during disputes is that evidence is destroyed by the time the case is adjudicated or when legal proceedings are initiated. As mentioned above, in the cases where the author has attempted to read the ATC from cards, the bank has requested that this evidence be destroyed. This appears to be standard procedure, but seems to be unwise now that cards can contain useful forensic evidence. Similarly, in the case of *Job v Halifax plc*, the transaction logs which included the ARQC were destroyed by Halifax after 180 days, even though the transactions were in dispute. Since there was only one log of the transaction presented as evidence, any inconsistency which might have existed would not have been detected. While the judgment in *Job v Halifax plc* went in the bank's favour, the judge cautioned that in future cases, the fact that a bank destroys evidence may be considered differently by another judge in different circumstances.

Obtaining evidence held by third parties can also be problematic, such as CCTV footage. A common scenario is that upon reporting a disputed transaction to their bank, a customer is immediately refunded. The customer is then satisfied with the outcome, and does not take the case further. Simultaneously, an internal investigation is initiated by the bank, which could take many weeks. If this investigation decides against the customer, the refund will be reversed. At this point, the customer will be motivated to obtain CCTV evidence and logs from third parties, but by this stage they may have been deleted. Even if they still exist, the CCTV owner may only respond to an application by the police, and since April 2007 the police will only investigate if requested to do so by the bank. From the bank's perspective, a case in which the customer has had their refund denied is resolved, so they are unlikely to take any further action.

These problems have led to many customers contacting the press, and stories on Chip & PIN attract high levels of interest from their audiences. Investigative journalists have worked with researchers in order to discover and demonstrate security vulnerabilities. In some cases they have also contacted bank

⁴¹ "Fraud victims struggle to get money back: One in five financial fraud victims not reimbursed", Which?, 25 June 2009, <http://www.which.co.uk/news/2009/06/fraud-victims-struggle-to-get-money-back-179150.jsp>.

insiders and reformed criminals to ask for assistance. In this respect, the press performs a valuable role by protecting sources from potential recrimination. The media can also be helpful in obtaining refunds for disputed transactions. For example, Barclays refunded Suzanne Lewis £1,400 following the intervention of BBC Watchdog in February 2007.⁴²

The Financial Ombudsman has been criticized for accepting assurances from the banks that Chip & PIN cards cannot be cloned.⁴³ The banks' opinion is based on their experience that criminals have not been caught either using cloned Chip & PIN cards, or exploiting failures in authorization systems. Care must be taken to ensure that such arguments are not circular: if the definition of a cloned card is one which will evade detection by the bank's anti-fraud measures, then of course they will not have been caught by banks. Similarly, in this article, a number of examples of failures in bank computer systems and procedures have been given, that have become public only because either the customer reported the problem, or there was an associated criminal prosecution. Even though these cases are not complete explanations for how cloned smart cards could be produced, it might be that others exist which have not become public, and which could be exploited by criminals.

5 Conclusion

A theme throughout this article has been that Chip & PIN greatly increases the complexity of banking systems. This helps deter criminals, but also greatly increases the amount of preparation work necessary when disputed transactions involving Chip & PIN are the subject of litigation. The fact that logs, CCTV footage, and other useful information may be destroyed suggests that requests to preserve evidence should be sent and pursued quickly, even if the disputed transaction is initially reversed. For this reason, and so that opportunities to challenge the evidence are not missed, it is also prudent for customers disputing transactions to obtain legal representation early on in their case.

On the technical side of disputes, the complexity of Chip & PIN offers both advantages and difficulties. The extra evidence available can, potentially, help support a particular interpretation, but the technical nature of the evidence is such that it needs greater precision and effort to interpret and analyze. However, for the evidence to be subject to analysis and interpretation, it must be disclosed. In addition, it is also necessary to adduce sufficient information to establish its reliability, and what conclusions may safely be drawn from it. This presents challenges both to litigators and expert witnesses. It is anticipated that this article provides assistance to both these audiences, should they be involved in such a case.

Dr Steven J. Murdoch is a researcher at the University of Cambridge Computer Laboratory. His areas of interest include cryptography, privacy, and banking security. His publications on these topics are available on his website. Steven has also acted as an expert witness in civil and criminal cases involving Chip & PIN.

Steven.Murdoch@cl.cam.ac.uk

<http://www.cl.cam.ac.uk/users/sjm217/>

⁴²BBC Watchdog, 6 February 2007, 19:00, BBC One.

⁴³Submission to the Hunt Review of the Financial Ombudsman Service, Foundation for Information Policy Research, 16 January 2008, <http://www.fipr.org/080116huntreview.pdf>.