

Chip and PIN is Broken

Steven J. Murdoch, Saar Drimer, Ross Anderson, Mike Bond
University of Cambridge
Computer Laboratory
Cambridge, UK
<http://www.cl.cam.ac.uk/users/{sjm217,sd410,rja14,mkb23}>

Abstract—EMV is the dominant protocol used for smart card payments worldwide, with over 730 million cards in circulation. Known to bank customers as “Chip and PIN”, it is used in Europe; it is being introduced in Canada; and there is pressure from banks to introduce it in the USA too. EMV secures credit and debit card transactions by authenticating both the card and the customer presenting it through a combination of cryptographic authentication codes, digital signatures, and the entry of a PIN. In this paper we describe and demonstrate a protocol flaw which allows criminals to use a genuine card to make a payment without knowing the card’s PIN, and to remain undetected even when the merchant has an online connection to the banking network. The fraudster performs a man-in-the-middle attack to trick the terminal into believing the PIN verified correctly, while telling the card that no PIN was entered at all. The paper considers how the flaws arose, why they remained unknown despite EMV’s wide deployment for the best part of a decade, and how they might be fixed. Because we have found and validated a practical attack against the core functionality of EMV, we conclude that the protocol is broken. This failure is significant in the field of protocol design, and also has important public policy implications, in light of growing reports of fraud on stolen EMV cards. Frequently, banks deny such fraud victims a refund, asserting that a card cannot be used without the correct PIN, and concluding that the customer must be grossly negligent or lying. Our attack can explain a number of these cases, and exposes the need for further research to bridge the gap between the theoretical and practical security of bank payment systems. It also demonstrates the need for the next version of EMV to be engineered properly.

Keywords—EMV; Chip and PIN; card fraud; bank security; protocol failure; security economics; authentication

I. INTRODUCTION

Smart cards have gradually replaced magnetic strip cards for point-of-sale and ATM transactions in many countries. The leading system, EMV [1], [2], [3], [4] (named after Europay, MasterCard, and Visa), has been deployed throughout most of Europe, and is currently being rolled out in Canada. As of early 2008, there were over 730 million EMV-compliant smart cards in circulation worldwide [5]. In EMV, customers authorize a credit or debit card transaction by inserting their card and entering a PIN into a point-of-sale terminal; the PIN is typically verified by the smart card chip, which is in turn authenticated to the terminal by a digital certificate. The transaction details are also authenticated by a cryptographic message authentication code (MAC), using




Figure 1. Fraud statistics on UK-issued cards [6]

a symmetric key shared between the payment card and the bank that issued the card to the customer (the issuer).

EMV was heavily promoted under the “Chip and PIN” brand during its national rollout in the UK. The technology was advertised as a solution to increasing card fraud: a chip to prevent card counterfeiting, and a PIN to prevent abuse of stolen cards. Since its introduction in the UK the fraud landscape has changed significantly: lost and stolen card fraud is down, and counterfeit card fraud experienced a two year lull. But no type of fraud has been eliminated, and the overall fraud levels have actually risen (see Figure 1). The likely explanation for this is that EMV has simply moved fraud, not eliminated it.

One goal of EMV was to externalise the costs of dispute from the issuing bank, in that if a disputed transaction has been authorised by a manuscript signature, it would be charged to the merchant, while if it had been authorised by a PIN then it would be charged to the customer. The net effect is that the banking industry, which was responsible for the design of the system, carries less liability for the fraud. The industry describes this as a ‘liability shift’.

Security economics teaches us that such arrangements create “moral hazard,” by insulating banks from the risk of their poor system design, so it is no surprise when such plans go awry. Several papers have documented technical attacks on EMV. However, it is now so deeply entrenched that changes can be very hard to make. Fundamental pro-

protocol changes may now require mutual agreement between banks, merchants, point-of-sale hardware manufacturers, and international card schemes (Visa, MasterCard, and American Express), all of which lobby hard to protect their interests. As with the Internet communications protocols, we are stuck with suboptimal design decisions made a decade ago. So few system changes have been made, and meanwhile the volume of customer complaints about disputed transactions continues to rise. A June 2009 survey revealed that one in five UK victims of fraud are left out of pocket [7].

In the past few years, the UK media have reported numerous cases where cardholders' complaints have been rejected by their bank and by government-approved mediators such as the Financial Ombudsman Service, using stock excuses such as 'Your card was CHIP read and a PIN was used so you must have been negligent.' Interestingly, an increasing number of complaints from believable witnesses indicate that their EMV cards were fraudulently used shortly after being stolen, despite there having been no possibility that the thief could have learned the PIN.

In this paper, we describe a potential explanation. We have demonstrated how criminals can use stolen "Chip and PIN" (EMV) smart cards without knowing the PIN. Since "verified by PIN" – the essence of the system – does not work, we declare the Chip and PIN system to be broken.

II. PROTOCOL FAILURE

EMV is both a protocol suite and a proprietary protocol framework: a general toolkit from which protocols can be built. In practice, it works as follows. A bank that issues EMV cards selects a subset of the EMV protocols, choosing for instance between digital signature methods, selecting a MAC algorithm, and deciding on hundreds of customisable options regarding authentication and risk management. Their selection must comply with card scheme rules as well as the EMV framework. Meanwhile merchants and acquiring banks (who receive payments on behalf of merchants) simply procure EMV-compliant hardware and software and connect it to the payment networks (operated by card schemes).

Since we cannot enumerate the many possible protocols, we mainly describe the protocol as it is deployed within the UK. However, it is implemented similarly in many other countries. In particular, the attack we introduce in this paper results both from a protocol failure of the EMV framework, and a failure of the proprietary MAC protocols that are used by issuing banks (and approved by the card schemes).

As Figure 2 shows in detail, the EMV protocol can be split into three phases:

Card authentication:

Assures the terminal which bank issued the card, and that the card data have not been altered

Cardholder verification:

Assures the terminal that the PIN entered by the customer matches the one for this card

Transaction authorization:

Assures the terminal that the bank which issued the card authorizes the transaction

1) *Card authentication:* EMV smart cards may contain multiple separate applications with different cryptographic keys, such as a debit or credit card for use at shops, ATM functionality, and MasterCard Chip Authentication Programme (CAP) applications for online banking. Thus when a card is inserted into a point of sale terminal, the terminal first requests a list of supported applications (by reading the file "1PAY.SYS.DDF01") and selects one of them. The actual transaction is then initiated by sending the *Get Processing Options* command to the card.

Next, the terminal reads cardholder information from the card by sending a *Read Record* command with the appropriate file identifiers. These records include card details (e.g. primary account number, start and expiry date), backwards compatibility data (e.g. a copy of the magnetic strip), and control parameters for the protocol (e.g. the cardholder verification method list, and card data object lists, which will be discussed later).

The records also include an RSA digital signature over a subset of the records, together with a certificate chain linking the signing key to a card scheme root key known to the terminal. In one variant of EMV, known as SDA (static data authentication), the card itself is not capable of performing RSA operations, so it can only present the terminal with a static certificate. Cards employing the DDA (dynamic data authentication) variant additionally contain RSA private keys which are used to sign a nonce sent by the terminal and whose corresponding public keys are authenticated by the certificate chain.

SDA cards (which prior to 2009 all UK banks issued) are vulnerable to a trivial and well-known replay attack in which the certificate is read from a card and written to a counterfeit one (these are often called "yes cards" because they will respond "yes" to a PIN verification request, no matter what PIN is entered). The card is then used at a point-of-sale terminal which has no online connection to the banking network, and because there is no real-time interaction, the MAC produced during transaction authorization cannot be checked before the goods are handed over.

However, the vast majority of UK point-of-sale terminals maintain a permanent online connection, so yes cards could normally be detected¹. Since 2009, some UK banks have started issuing DDA cards, which resist counterfeiting even in offline transactions, by giving the cards the capability to sign a terminal-provided nonce under an asymmetric key. However the attack presented in this paper does not rely on the yes card attack; it is entirely independent of card authentication, whether by SDA or DDA.

¹There are viable criminal attack scenarios involving yes cards, and criminal business models, but these are beyond the scope of this paper.




Figure 2. A complete run of a Chip and PIN protocol.

2) *Cardholder verification*: The cardholder verification step starts with a mechanism negotiation, performed between the card and the terminal, to establish what cardholder authentication method they can (or must) use. This is driven by a data element called the cardholder verification method (CVM) list. The CVM list states the card’s policy on when to use a PIN, or a signature, or nothing at all, to authenticate the cardholder.

Protocols for negotiating an authentication mechanism are notoriously hard to get right. EMV specifies a complex negotiation algorithm by which the terminal can decide the appropriate method depending on the value of the transaction, its type (e.g. cash, purchase), and the terminal’s capabilities. The CVM list also specifies what action should be taken if cardholder verification fails, i.e., whether the next method should be tried or the transaction rejected.

In practice, however, only a small subset of these capabilities is used. UK cards we have examined specify, in descending order of preference, *PIN verification*, *signature verification*, and *no verification*. A terminal may skip an option of which it is not capable; for example, unattended terminals cannot do signature verification, and some vending machines are not equipped with PIN entry devices/keypads. There may also be scope for operator discretion. For example, the card may permit the terminal to attempt signature verification if PIN verification fails, but in practice merchants will normally reject such a transaction. In the UK there also exists a type of card known as a “Chip & Signature” card, which does not support PIN verification

at all. These cards are issued to customers who request them, normally because they are unable to remember a PIN or are visually impaired. Some customers also request such cards because they are concerned about the additional liability that PIN-based transactions would place on them.

However, the vast majority of transactions are ‘PIN verified’, which means the customer enters the PIN on a PIN entry device. The PIN is sent to the card, and the card compares it to the PIN it stores. If they match, the card returns $0x9000$, and if it fails the card returns $0x63Cx$, where x is the number of further PIN verification attempts the card will permit before locking up. Note that the card’s response is not directly authenticated.

ATM cardholder verification works differently, and uses a method known as “online PIN”, as opposed to “offline PIN” described above. Here, the PIN is encrypted by the ATM, and sent to the issuer over a payment network. The issuer then verifies the PIN centrally, and sends the result back to the ATM. The attack we present in this paper only applies to offline PIN cardholder verification.

We have observed variations between countries. While cards from Belgium and Estonia work like British cards, we have tested cards from Switzerland and Germany whose CVM lists specify either chip and signature or online PIN, at least while used abroad. The attack described here is not applicable to them. However, because UK point-of-sale terminals do not support online PIN, a stolen card of such a type could easily be used in the UK, by forging the cardholder’s signature.

3) *Transaction authorization*: In the third step, the terminal asks the card to generate a cryptographic MAC over the transaction details, to be sent to the issuing bank. The terminal calls the **Generate AC** command, to request an ARQC (authorization request cryptogram) from the card. The payload of this command is a description of the transaction, created by concatenating data elements specified by the card in the CDOL 1 (card data object list 1). Typically this includes details like the transaction amount, currency, type, a nonce generated by the terminal, and the TVR (terminal verification results), which will be discussed later.

The cryptogram sent to the bank includes a type code, a sequence counter identifying the transaction (ATC – application transaction counter), a variable length field containing data generated by the card (IAD – issuer application data), and a message authentication code (MAC), which is calculated over the rest of the message including a description of the transaction. The MAC is computed, typically using 3DES, with a symmetric key shared between the card and the issuing bank.

If the card permits the transaction, it returns an ARQC; otherwise, it returns an AAC (application authentication cryptogram) which aborts the transaction. The ARQC is then sent by the terminal to the issuing bank, via the acquirer and payment network. The issuer will then perform various cryptographic, anti-fraud and financial checks: such as whether the card has been listed as stolen, whether there are adequate funds, and whether the risk analysis algorithm considers the transaction acceptable. If the checks pass, the issuer returns a two byte ARC (authorization response code), indicating how the transaction should proceed, and the ARPC (authorization response cryptogram), which is typically a MAC over $ARQC \oplus ARC$. Both items are forwarded by the terminal to the card with the **External Authenticate** command.

The card validates the MAC contained within the ARPC, and if successful updates its internal state to note that the issuer authorized the transaction. The terminal then calls **Generate AC** again, but now using the CDOL 2, requesting that the card issues a TC (transaction certificate) cryptogram, signifying that it is authorizing the transaction to proceed. Finally, the terminal sends the TC to the issuer, and stores a copy in its own records in case there is a dispute. At this point it will typically print a receipt, which may contain the legend ‘Verified by PIN’ if the response to **Verify** indicated success. One copy of the receipt is given to the cardholder and a second copy is retained. We have also seen different receipts with ‘confirmed’ for the cardholder and ‘PIN verified’ on the merchant copy (perhaps to assure the merchant that the liability for disputes is no longer on them).

The above description assumes that the terminal chose to perform an online transaction and contacted the issuer. In the event of an offline transaction, the terminal requests that the card return TC on the first call to **Generate AC**. The




Figure 3. The man-in-the-middle suppresses the PIN Verify command to the card, and tells the terminal that the PIN has been verified correctly. A complete transaction is detailed in Appendix A.

Table I
TERMINAL VERIFICATION RESULTS (TVR) BYTE 3.

Bit	Meaning when bit is set
8	Cardholder verification was not successful
7	Unrecognized CVM
6	PIN Try Limit exceeded
5	PIN entry required and PIN pad not present or not working
4	PIN entry required, PIN pad present, but PIN was not entered
3	Online PIN entered
2	Reserved for future use
1	Reserved for future use

card may then either decide to accept the transaction offline by returning a TC, force the transaction online by returning an ARQC, or reject the transaction entirely by returning an AAC. Our attack applies just as well to the offline case.

III. THE ATTACK

The central flaw in the protocol is that the PIN verification step is never explicitly authenticated. Whilst the authenticated data sent to the bank contains two fields which incorporate information about the result of the cardholder verification – the Terminal Verification Results (TVR) and the Issuer Application Data (IAD), they do not together provide an unambiguous encoding of the events which took place during the protocol run. The TVR mainly enumerates various possible failure conditions for the authentication, and in the event of success does not indicate *which particular method was used* (see Table I).

Therefore a man-in-the-middle device, which can intercept and modify the communications between card and terminal, can trick the terminal into believing that PIN verification succeeded by responding with 0×9000 to **Verify**, without actually sending the PIN to the card. A dummy PIN must be entered, but the attack allows any PIN to be accepted. The card will then believe that the terminal did not support PIN verification, and has either skipped cardholder

Table II
IAD FORMAT, BYTE 5 (BITS 4–1) FROM A VISA VERSION 10
CRYPTOGRAM [8, APPENDIX A-13, P222].

Bit	Meaning when bit is set
4	Issuer Authentication performed and failed
3	Offline PIN performed
2	Offline PIN verification failed
1	Unable to go online

verification or used a signature instead. Because the dummy PIN never gets to the card, the PIN retry counter is not altered. The modified protocol flow is shown in Figure 3.

Neither the card nor terminal will spot this subterfuge because the cardholder verification byte of the TVR is only set if PIN verification has been attempted *and* failed. The terminal believes that PIN verification succeeded (and so generates a zero byte), and the card believes it was not attempted (so will accept the zero byte).

The IAD (Table II) does often indicate whether PIN verification was attempted. However, it is in an issuer-specific proprietary format, and not specified in EMV. Therefore the terminal, which knows the cardholder verification method chosen, cannot decode it. The issuer, which can decode the IAD, does not know which cardholder verification method was used, and so cannot use it to prevent the attack. Because of the ambiguity in the TVR encoding, neither party can identify the inconsistency between the cardholder verification methods they each believe were used. The issuer will thus believe that the terminal was incapable of soliciting a PIN – an entirely plausible yet inaccurate conclusion.

For offline transactions, the issuer will not be contacted until after the transaction has been completed, so has even less ability to detect the attack. Some cards may refuse to authorize an offline transaction without having successfully verified the PIN. This however is no obstacle to the attack, because the man-in-the-middle can simply change the cryptogram type field in the response to the **Generate AC** call, turning an ARQC or AAC into a TC. This modification will possibly cause the cryptogram verification to fail, but this would only be detected after the cardholder has left with the goods.

In the UK, PIN-based cardholder verification is mandatory and all cards support offline PIN verification. Although the CVM list permits merchants to fall back to signature, they rarely offer this (they become liable for fraud if they do). Therefore, unless a thief can somehow discover the PIN, using a stolen card is difficult. Here, our attack could be used by criminals to carry out a point-of-sale transaction.

In fact, the authors are regularly contacted by bank customers who have had fraudulent transactions carried out shortly after their card has been stolen, and who state that they did not write down their PIN, but found that their bank accused them of negligence and refused to refund the losses.

The attack we describe in this paper may explain some of these cases.

IV. ATTACK DEMONSTRATION

We successfully executed the attack using several different Chip and PIN cards at a live terminal. The schematic and a photograph of the equipment used is shown in Figure 4. Stills from a video of us carrying this attack out are in Figure 5; a film by BBC Newsnight of us carrying out the attack is also available [9]. The hardware for the attack was made of cheap off-the-shelf components and required only elementary programming and engineering skills.

The man-in-the-middle circuit connects to the terminal through a fake card. This card has thin wires embedded in the plastic substrate, which connect the card’s contact pads to an interface chip (\$2 Maxim 1740 [10]) for voltage level-shifting. This is connected to a general-purpose FPGA board (\$189 Spartan-3E Starter Kit [11]) that drives the card and converts between the card and PC interfaces. Through a serial link, the FPGA is connected to a laptop, which is in turn connected to a standard smart card reader from Alcor Micro (\$8) into which the genuine card is inserted. A Python script running on the laptop relays the transaction while waiting for the **Verify** command being sent by the terminal; it then suppresses it to the card, and responds with `0x9000`:

```
if VERIFY_PRE and command[0:4] == "0020":
    debug("Spoofing VERIFY response")
    return binascii.a2b_hex("9000")
```

The rest of the communication is unaltered.

Where the merchant colludes with the attacker for a cut of the profit, the hardware bulk is not a factor. When the merchant is unwitting, the security measures introduced to protect the customer from a corrupt merchant skimming the magnetic strip work in the attackers’ favour. Cardholders are instructed not to hand their card to the merchant, and the merchant is under social pressure to look away during a transaction while the cardholder enters their PIN. The attack could easily be miniaturized: it can be ported to smaller hardware devices, and would not require a PC at all if the FPGA or microcontroller is programmed to parse the transaction and interface with the card. Miniaturized hardware could be entirely hidden in a coat sleeve and used immediately after the card is stolen.

Finally, we can envision a carrier card that hosts a cutout of the original card, which interfaces with a microcontroller that communicates with the terminal. This way, the attack is entirely encapsulated in a card form factor and can be moderately industrialized. Miniaturized “shims” with an embedded microcontroller have already been created for SIM cards for unlocking phones from a particular network [12]; the simple code required for our attack can be ported to





Figure 4. Components of the attack.

run on a similar device. Miniaturization is mostly a mechanical challenge, and well within the expertise of criminal gangs: such expertise has already been demonstrated in the miniaturised transaction interceptors that have been used to sabotage point of sale terminals and skim magnetic strip data. Miniaturization is not critical, though, as criminals can target businesses where a card can be used with wires running up the cashout operative’s sleeve, while a laptop and FPGA board can be hidden easily in his backpack. There are firms such as supermarkets and money changers whose terminals are located on the other side of a barrier from the checkout staff, who therefore do not scrutinise the cards their customers use.

V. CAUSES

The failure we identify here might be patched in various ways which we will discuss later. But at heart there is a protocol design error in EMV: it compartmentalises the issuer-specific MAC protocol too distinctly from the negotiation of the cardholder verification method. Both of the parties who rely on transaction authentication – the merchant and the issuing bank – need to have a full and trustworthy view of the method used to verify the cardholder; and because the relevant data cannot be collected neatly by either party, the framework itself is flawed.

A key misconception of the designers was to think of the TVR and card verification results primarily as separate lists

of possible failures represented by a bit mask, rather than as a report of the authentication protocol run.

This is not to say that issuing banks cannot in future implement secure proprietary schemes within the EMV framework: because the internal protocols are proprietary anything is possible, and some potential options will be discussed in Section VI. But such schemes must make ever more complex and intricate analysis of the transaction data returned, driving up the complexity and fragility of the existing EMV card authorization systems. Essentially, they will have to ignore the framework, and without a change in the framework itself, the authorization calculations will remain so complex and dependent on external factors that further mistakes are very likely. Also, as the protocol becomes more customized by the issuer, the introduction of new system-wide features sought for other purposes will become progressively more difficult and expensive.

The failure of EMV has many other aspects which will be familiar to security engineers. There was a closed design process, with no open external review of the architecture and its supporting protocols. The protocol documentation appeared eventually in the public domain – nothing implemented by 20,000 banks could have been kept secret – but too late for the research community to give useful feedback before a lot of money was spent on implementation.

The economics of security work out not just in the interaction between banks, customers and merchants – with

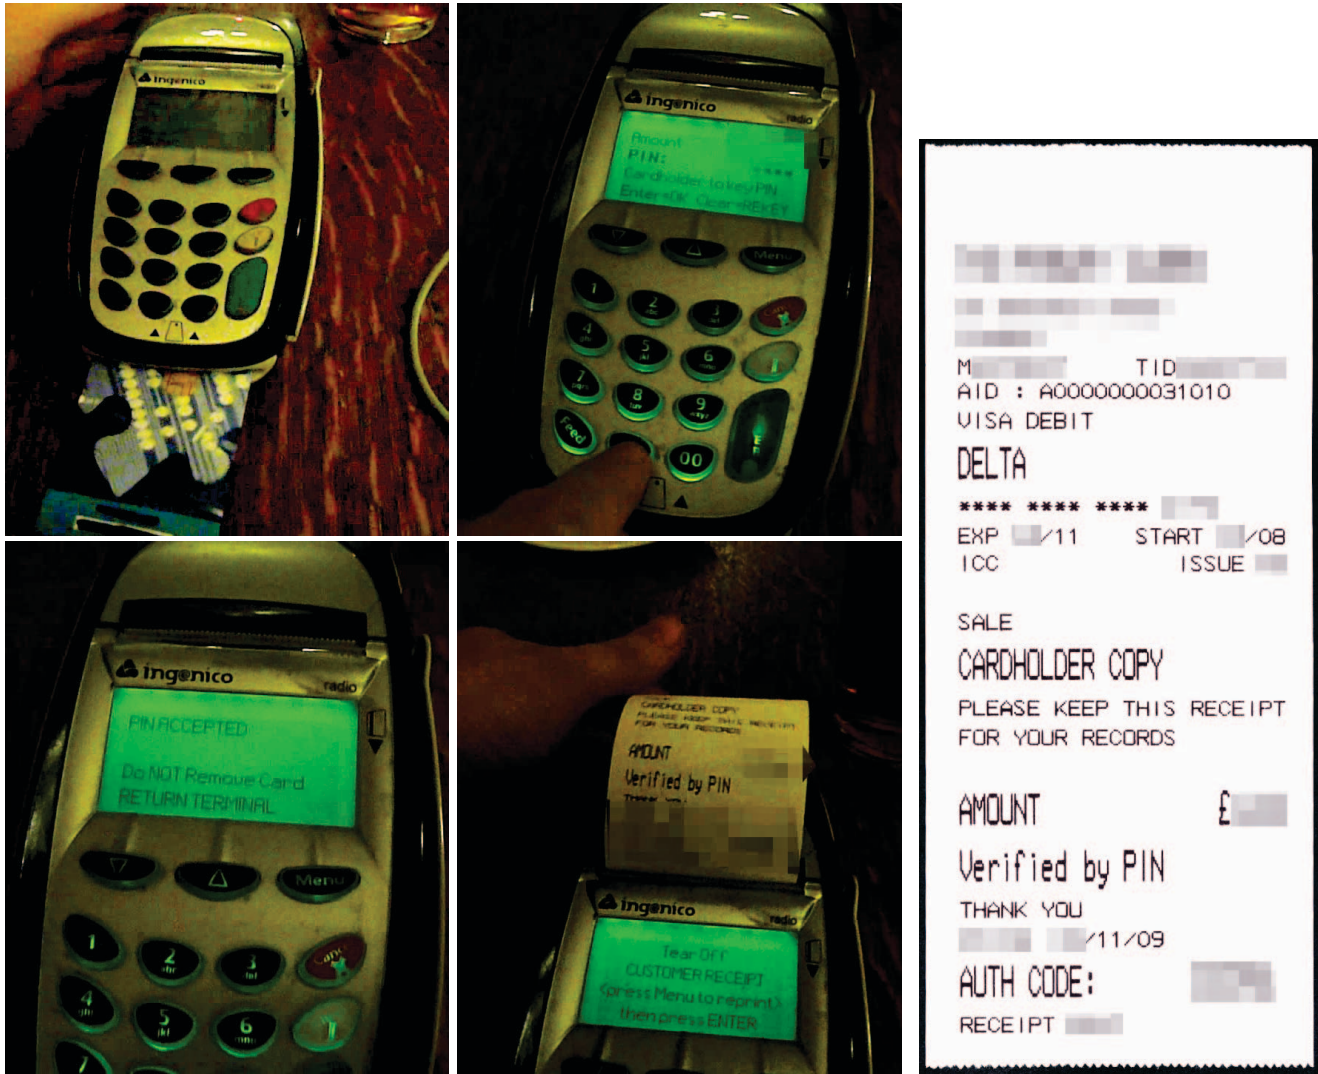


Figure 5. Carrying out the attack. Although we entered the wrong PIN, the receipt indicates that the transaction was “Verified by PIN”.

the banks using their control of the system to dump liability, and thus undermining their own incentive to maintain it. There are also mismatches between acquirer and issuer banks, with only the latter feeling any real incentive to remediate security failures; between banks and suppliers, with the latter being squeezed on costs to the point that they have little incentive to innovate; and between banks and the facilities management firms to whom much of the business of card personalisation, network operation, and so on gets outsourced. The industry as a whole suffers from a significant collective action problem. It will be interesting to see which of the dozens of national bank regulators, or which of the three card schemes, will initiate action to deal with those aspects of the problems described here that cannot be tackled by issuer banks acting alone. It may be worth bearing in mind that the smart card industry spent

some twenty years pitching its products to the banks before it managed to overcome the collective action problem and get the industry to move. In the absence of a catastrophe, changes that require everyone to act together are going to be slow at best.

A major contributing factor to the fact that these protocol flaws remained undiscovered is the size and complexity of the specification, and its poor structure. The core EMV protocols are now 707 pages long, there are a further 2126 pages of testing documentation, and card schemes also specify extensions (Visa publishes 810 pages of public documentation, and there is more which is secret). Many options are given, and a typical implementation mixes some of the functionality from the published manuals with some issuer-specific enhancements. Security critical details are scattered throughout, and there is no one section which is

sufficient to understand the protocol, the threat model, or the security policy. In fact, much detail is not specified at all, being left to implementation decisions by individual issuers.

For example, to confirm the existence of the security vulnerability discussed in this paper, we needed to establish:

- Lack of authentication in transport layer (EMV Book 1 [1])
- Encoding of Verify (EMV Book 3 [3, p71])
- Encoding of the TVR (EMV Book 3, Annex C [3, p171])
- Recommended generation algorithm for the ARPC (EMV Book 2 [2, p89])
- Recommended transaction data items to be included in the ARQC and TC (EMV Book 2 [2, p88])
- Absence of cardholder verification result in ARQC and TC requests (EMV Book 2 [2, p73], EMV Book 3 [3, p58])
- Encoding of the CVM list (EMV Book 3, Annex C [3, p168])
- Algorithm for selecting cardholder verification method (EMV Book 3 [3, p103])
- Transaction flow (EMV Book 3 [3, p83])
- Values of the TVR for signature transaction (EMV Book 4 [4, p49])
- Whether the actual cardholder verification method used is included in the CDOL (unspecified, found by experiment)
- Whether the issuer checks value of IAD in online transactions (unspecified, found by experiment)
- Whether the terminal attempts to decode the IAD (unspecified, found by experiment)
- Encoding of the IAD (proprietary, specified in Visa Integrated Circuit Card Specification, Appendix A [8, p222])

Ultimately EMV is a compatibility system and protocol toolkit. It allows interoperable protocols to be built, but following the specification – even including the optional recommendations – does not ensure a secure protocol. This may explain why there has been little analysis of EMV. The specification does not contain enough detail to support any claims about the security of implementations, as they depend on proprietary, and often unpublished, details. It is necessary to do experiments, as we did. But researchers, and merchants who assist them, may be afraid of retribution from the banking industry, which makes experimentation difficult.

VI. SOLUTIONS AND NON-SOLUTIONS

Core protocol failures are difficult to fix. None of the security improvements already planned by banks will help: moving from SDA to DDA will not have any effect, as these are both methods for card authentication, which occurs before the cardholder verification stage. Neither will a further proposed enhancement – CDA (combined data authentication) – in which the transaction authorization stage

additionally has a digital signature under a private key held by the card. This is because the attack we present does not interfere with either the input or output of transaction authentication, so replacing a transaction MAC with a digital signature will not help.

One possible work-around is for the terminal to parse the IAD, which does include the result of PIN verification (Table II). This will only be effective for online transactions, and offline transactions where CDA is used, otherwise the man-in-the-middle device could tamper with the IAD as it is returned by the card. It would also be difficult to implement because the IAD was intended only for the issuer, and there are several different formats, without any reliable method to establish which one is used by a particular card. However a solution along these lines would require the acquiring banks and the terminal vendors to act together, which for the incentive reasons discussed above would be both slow and difficult.

The realities of security economics mean that we have to look for a fix which requires changes only to customer cards or to the issuer's back-end systems. Such a repair may in fact be possible: the card can change its CDOL to request that the CVMR (cardholder verification method results) be included in the payload to the Generate AC command. This specifies which cardholder verification method the terminal believes was used, and so should allow the card and issuer to identify the inconsistency. Out of many, we have only seen one EMV card which requests this field, and it is not clear that the issuer actually validates the CVMR against the IAD. Whether this fix works for a given bank will depend on its systems; we have not been able to test it, and given that it involves reissuing the card base it would take years to roll out.

In addition to the global EMV specifications, and ones from card-scheme operators such as Visa and MasterCard, there are also country-specific standards. In the UK, the standard for communications between merchant terminal and acquirer is APACS 70, Book 2 [13], which specifies that both the IAD and CVMR must be sent. This is sufficient information for the issuer to detect the attack, but our results clearly show that they are not currently doing so. One possible reason is that the data items are dropped or corrupted between the acquirer and issuer (industry experts disagree over whether this is the case). Another possibility is that some terminals do not set the CVMR correctly, resulting in too many false positives if it were compared against the IAD. In any case, unless the CVMR is included in the CDOL it may not be integrity-protected, so a second man-in-the-middle between terminal and acquirer (perhaps installed with co-operation of a corrupt merchant staff member) could tamper with it too.

These workarounds should resolve the particular flaw discussed in this paper, but there are likely to be more. A more prudent approach would be to follow established

design principles for robust security protocols. For example, adopting the “Fail-stop” [14] principles would prevent this attack; so would the explicitness principle, of ensuring you authenticate all data that might be relied on. Either approach would be likely to prevent other attacks too, and it would also make the protocols easier to analyze. Alternatively an industry standard transport-layer confidentiality and authenticity standard, such as TLS [15], could be wrapped around the existing command set. However, it’s important that a fix should not just be an ad-hoc hack. The next version of EMV needs a proper security engineering exercise; regulators should insist that a threat model, security policy and protocol specification are published for open review.

VII. EVIDENCE IN CHIP AND PIN DISPUTES

Even if it turns out to be too expensive in the short term to prevent the attack we present in this paper, it is important to detect whether it occurred when resolving cases where a customer disputes a transaction. While assisting fraud victims who have been refused a refund by their bank, we have requested the IAD so as to discover whether the card believes PIN verification succeeded, but have almost always been refused. This paper illustrates that while the IAD can be considered trustworthy (after its MAC has been verified), the TVR and merchant receipt must not.

In fact, dispute resolution processes we have seen in the UK are seriously flawed, even excluding the protocol failure described here. In one disputed transaction case we assisted in, the customer had his card stolen while on holiday, and then used in an EMV transaction. The issuer refused to refund this customer on the basis that their records showed the PIN was used. Luckily, the customer managed to obtain the merchant receipts, and these contained the TVR. This indicated that the PIN was not used, and the merchant opted to fall back to signature. We decoded the TVR and informed the customer, who was then able to get a refund.

Other customers are less fortunate: it is unusual for the TVR to be included on the receipt, and often the merchant receipt has been destroyed by the time the dispute is being considered. In these cases we have not been able to obtain the TVR, IAD, or even a statement by the bank as to how they established that the cardholder was verified through the correct PIN being entered.

Our demonstration therefore exposes a deeper flaw in EMV and the associated systems: they fail to produce adequate evidence for dispute resolution or litigation. Procedures are also a mess. For example, once a transaction is disputed a typical bank either destroys the card or asks the customer to do so, preventing information from being extracted which might show whether the card was actually used. Transaction logs are commonly only kept for 120 days, and by the time the dispute is being heard the bank may have destroyed most of the records. (This was the case in the well-known Job v. Halifax trial: even though the Halifax

had been notified that the transaction was being disputed, the logs were then destroyed in defiance of Visa guidelines [16].)

These general issues are discussed by Murdoch [17], but the vulnerability described in this paper poses a problem for such banks. If they have indeed destroyed all record of the IAD, they will be unable to show that disputed transactions actually used the correct PIN. So our findings might help banks understand that it is in their interest to retain evidence rather than destroy it.

Another evidential issue is that even if the issuer were able to establish whether the attack we present here had occurred, this may not help customers because the typical receipt still states that the PIN was verified. Although this may be false, many people evaluating evidence (adjudicators, judges, and jury members) will not know this. In one particular case, from 2009, the issuing bank, and government-approved adjudicator, explicitly relied upon the “Verified by PIN” indicator on the merchant receipt, in concluding that the transaction was PIN-verified and therefore the customer was liable. For this reason we propose that terminals no longer print “Verified by PIN” unless the protocol actually supports this assertion.

VIII. RELATED WORK

EMV has been available for 14 years and is now widely deployed despite little published research on its security. In 1999, Herreweghen and Wille [18] evaluated the suitability of EMV for Internet payments and identified the problem of not being able to determine if the `Verify` command was ever executed because it is not authenticated. In their proposed Internet-based payment scheme, they suggested that the ARQC should only be generated if the `Verify` command has been successful. But their paper did not consider that the result of PIN verification is included in the IAD, nor that the `Verify` message could be tampered with by a man-in-the-middle in a point-of-sale transaction.

More recently, interest in EMV has increased since it was widely deployed in 2005, but perhaps due to the specification’s complexity and incompleteness, the closed user community, and difficulties in carrying out experiments, researchers have not done much work on it. Anderson *et al.* [19] described how bank customers might have difficulty in obtaining refunds once transactions were authorized by PIN. That paper also outlined some potential attacks against Chip and PIN, such as cloning SDA cards for use in offline transactions, and the likelihood that criminals would migrate towards cross-border fraud if and when legacy magnetic strip transactions were disabled at domestic ATMs. It also briefly considered the attack line described in this paper, but did not follow through at the time with detailed analysis or performing experiments.

Another potential EMV weakness outlined in [19] was the relay attack, which was refined and demonstrated by Drimer and Murdoch [20]. Here, the criminal sets up a tampered

Chip and PIN terminal, which the victim uses to make a small transaction (e.g. buying a meal at a restaurant). Rather than placing the transaction, the terminal relays the session to a fake card which is being used for a far larger transaction elsewhere (e.g. buying diamonds at a jewellery shop). The authors also described a defence against this attack, in which the terminal and card engage in a cryptographic exchange which not only establishes authenticity but also a maximum distance bound, either eliminating or greatly limiting the applicability of the attack.

Another attack is to tamper a terminal to merely record card details, and then use them for a fraudulent transaction later. Drimer *et al.* [21] demonstrated that current Chip and PIN payment terminals have inadequate tamper resistance, and a tapping device can be surreptitiously added to record the customer's PIN and enough details to allow a cloned magnetic strip card to be created. Criminals are now known to have carried out variants of this attack, so banks are now taking action: the chip no longer has a copy of the magnetic strip (one data field is replaced), and magnetic strip fallback transactions are gradually being phased out.

The work presented in this paper is a significant advance in our understanding of attacks against EMV because it is applicable to online transactions (unlike cloned SDA “yes cards”); it does not require criminals to synchronise their fraudulent purchase with that of an unwitting customer (as the relay attack does); and it does not depend on magnetic strip fallback (unlike the payment terminal tampering attacks). As a consequence, it may be one of the most realistic and attractive attacks for criminals, if and when magnetic strip transactions are no longer permitted. It could even be used at the moment, by criminals who wish to make purchases in countries which now mandate EMV transactions at point of sale. It may explain a number of the transaction dispute cases reported to us.

If this attack becomes more widely used, its net effect will be that criminals can use stolen cards in shops without the cardholder being negligent – exactly as was the case with magnetic strip cards before the introduction of EMV. However, so long as the public is not aware of this, the banks will be able to get away with blaming cardholders for fraud. We have therefore decided on a policy of responsible disclosure, of publishing this paper some time after informing bank regulators in the UK, Europe and North America of the vulnerability.

At present, we understand that there is a lot of pressure on the US Federal Reserve from the banks it regulates to countenance a move from magnetic strip cards to EMV. This paper shows that such a move may be premature. It's not reasonable for the smart card industry to foist a broken framework on the US banking industry and then leave it to individual issuer banks to come up with patches. The EMV consortium should first publish its plans for fixing the framework, presumably with the next version (v 5) of

the EMV specification. The Fed should then satisfy itself of three things.

First, will the fix work technically? For this, only open peer review will do. Second, will the high level of consumer protection so far enjoyed by US cardholders be preserved? Third, will the introduction of the remediated system introduce any systemic risks, because of moral hazard effects? For these last two questions to be answered in the affirmative, we believe that there must be no associated ‘liability shift’ as there has been in Europe and Canada.

IX. RESPONSE

The response to our paper has been largely positive, with most knowledgeable respondents agreeing that the attack works. However there was substantial discord regarding our conclusion that “Chip and PIN is broken”, which can mainly be explained by differences in the way that respondents define and measure success. In this section we summarise and comment only on the discordant responses: the positive responses speak for themselves.

Respondents who measure success differently have argued that Chip and PIN is *de facto* successful because its deployment has reduced lost and stolen card fraud; others argued that it is successful because the chip itself still has not been fully cloned by criminals.

We measure the success of Chip and PIN by its two core goals: first, to prevent counterfeit card fraud using the chip, and second to prevent lost and stolen card fraud using the PIN. Because stolen cards can be used without knowing the PIN, by our definition, Chip and PIN is broken. We do not believe that the system is broken beyond repair, but neither is it the case that a simple fix will suffice, due to the unmanageable complexity of EMV. This has been demonstrated by the spirited disagreement among experts discussing the attack on our blog [22] and proposing different favoured solutions, and by the continued absence of a fix at the time of writing, almost three months since the industry was notified.

Some of our respondents argued that Chip and PIN was a success on economic grounds, claiming that it saved more money from fraud than it cost to deploy. However they did not present figures to back up this claim. And counterfactual history is hard: how would one show that in the absence of EMV, fraud would have increased even more than it in fact has? Other respondents agreed that Chip and PIN simply pushed fraud to other areas such as card-not-present fraud, undermining the argument of economic success.

Some respondents argued that our attack would be difficult to deploy, for instance because of the bulk of the equipment and because of the narrow window of opportunity between theft of a card and its cancellation once the cardholder reports it stolen. Some even insisted on characterising it as theoretical, despite the fact it was deployed against live terminals at real merchants at three different sites. Whilst our demonstration equipment was indeed bulky, miniaturization

is straightforward and well within the capabilities of criminals who already miniaturize hi-tech point-of-sale skimmers and ATM skimmers. Skimmers perform far more complex actions than blocking a single command from a protocol run. Those who argued that the window of opportunity for abuse is small fail to recognise that the very reason the PIN is used is to prevent abuse of lost/stolen cards, so clearly the threat must have been substantial enough to justify investment in PIN technology in the first place. A larger window for abuse can also be achieved by postal interception of replacement cards, by stealing the victim's mobile phone at the same time as a card, or by pickpocketing rather than mugging: there really is no shortage of opportunity to abuse stolen cards.

Other respondents argued that the problem was not significant because systems could be patched to prevent it. Commenters proposed various cross-checking measures that might be performed by the issuing banks: checking the correspondence between CVMR and IAD (we also proposed this ourselves), or checking terminal capabilities and various acquirer fields such as POS data entry mode (defined in standard ISO 8583) against the IAD. But definite suggested fixes are generally remarkable by their absence. Indeed, some respondents claimed that card schemes were aware of this attack as far back as 2002; so if any straightforward cross-checks could fix the problem, surely they would have been implemented either now, or within the three months during which our paper circulated privately in the industry. Others argued that banks might simply move to online PIN at point-of-sale – in essence to abandon Chip and PIN in favour of an older approach – or move to CDA, where proprietary card checks such as the “terminal erroneously considers PIN OK” flag might help detect the subterfuge. Unfortunately none of these patches are easy. They require either a card re-issue, or re-engineering of the POS acquirer networks in those countries not set up to support online PIN. Both would be expensive. In particular, CDA has not been widely adopted because it is very sensitive to cryptographic errors: because more data are authenticated, it is more likely that a bug or incompatibility will cause an authentication failure. Many countries simply do not have the quality of engineering in their payment networks to be able to use CDA – a symptom of the excessive complexity of EMV.

A third class of respondents admitted the attack worked, but argued that because the IAD would be the most trusted source, and since this would not record PIN use, customers would never be liable for the losses. Unfortunately, in nearly all the disputes where we have assisted, banks have been extremely reluctant to provide any cryptographic evidence at all. Instead they have relied upon summary records of the transaction (not on any raw transaction data), or even on the printed receipts from the merchant, which we have proven to be untrustworthy.

Finally, some respondents agreed there was a problem but felt we had misattributed the blame. They argued that

it was not EMV that was at fault, or the card schemes' specifications, but the issuing banks. When contacted for comment, these same issuing banks referred us back to central bodies such as card schemes or trade associations. No-one wants to take responsibility. It is true that EMV is a protocol framework, and that its scope does not extend to issuer checks. We would argue that for any protocol specification to be valid, it must necessarily include statements of checks that must be performed by each party on the protocol messages. In the absence of named specification authors who accept responsibility, we feel it is fair to attribute responsibility to the “Chip and PIN” system, which is after all a marketing term that covers a whole specification stack.

X. CONCLUSION

We have shown how the PIN verification feature of the EMV protocol is flawed. A lack of authentication on the PIN verification response, coupled with an ambiguity in the encoding of the result of cardholder verification as included in the TVR, allows an attacker with a man-in-the-middle to use a card without the correct PIN. This attack can be used to make fraudulent purchases on a stolen card. We have shown that the live banking network is vulnerable by placing a transaction using the wrong PIN, with every major UK bank and foreign banks too. The records indeed falsely show that the PIN was verified, and the money was actually withdrawn from an account.

Attacks such as this could help explain the many cases in which a card has supposedly been used with the PIN, despite the customer being adamant that they have not divulged it. So far, banks have refused to refund such victims, because they assert that a card cannot be used without the correct PIN. This paper shows that their claim is false.

We have discussed how this protocol flaw has remained undetected; not only are the public specifications complex, but they also fail to specify security-critical details. Finally, we have discussed ways in which this vulnerability may be fixed by issuer banks, while maintaining backwards compatibility with existing systems. However, it is clear that the EMV framework is seriously flawed. Rather than leaving its member banks to patch each successive vulnerability, the EMV consortium must start planning a redesign and an orderly migration to the next version. In the meantime, the EMV protocol should be considered broken. We recommend that the Federal Reserve should resit pressure from banks to allow its deployment in the USA until it is fixed.

ACKNOWLEDGMENTS

We thank the anonymous reviewers, Colin Whittaker, and the contributors to the Light Blue Touchpaper blog for their comments. We also thank the merchants and cardholders who allowed us to carry out experiments, and Markus Kuhn for photography assistance. Steven Murdoch is funded by the Tor Project and employed part-time by Cronto Ltd.

REFERENCES

- [1] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 1: Application Independent ICC to Terminal Interface Requirements*, Version 4.2 ed., EMVCo, LLC, June 2008.
- [2] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 2: Security and Key Management*, Version 4.2 ed., EMVCo, LLC, June 2008.
- [3] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 3: Application Specification*, Version 4.2 ed., EMVCo, LLC, June 2008.
- [4] *EMV – Integrated Circuit Card Specifications for Payment Systems, Book 4: Cardholder, Attendant, and Acquirer Interface Requirements*, Version 4.2 ed., EMVCo, LLC, June 2008.
- [5] EMVCo, “About EMV,” November 2009. [Online]. Available: http://www.emvco.com/about_emv.aspx
- [6] APACS, “2008 fraud figures announced by APACS,” March 2009. [Online]. Available: http://www.ukpayments.org.uk/media_centre/press_releases/-/page/685/
- [7] Which?, “Fraud victims struggle to get money back,” June 2009. [Online]. Available: <http://www.which.co.uk/news/2009/06/fraud-victims-struggle-to-get-money-back-179150.jsp>
- [8] *Visa Integrated Circuit Card – Card Specification*, Version 1.4.0 ed., Visa International, October 2001.
- [9] S. J. Murdoch, S. Drimer, R. Anderson, and M. Bond, “EMV PIN verification “wedge” vulnerability,” February 2010. [Online]. Available: <http://www.cl.cam.ac.uk/research/security/banking/nopin/>
- [10] Maxim Integrated Products, Inc., *MAX1740, MAX1741 SIM/smart-card level translators in μ MAX*, January 2001, <http://datasheets.maxim-ic.com/en/ds/MAX1740-MAX1741.pdf>. [Online]. Available: <http://datasheets.maxim-ic.com/en/ds/MAX1740-MAX1741.pdf>
- [11] Xilinx Inc., “Spartan-3E starter kit,” November 2009. [Online]. Available: <http://www.xilinx.com/products/devkits/HW-SPAR3E-SK-US-G.htm>
- [12] “BLADOX Turbo SIM,” March 2010. [Online]. Available: <http://www.bladox.com/>
- [13] *Standard 70, Book 2 – Card Acceptor to Acquirer Interface Standards: Messages, Data Elements and Code Values for Real-time Systems*, The UK Cards Association, October 2009.
- [14] L. Gong and P. Syverson, “Fail-stop protocols: An approach to designing secure protocols,” in *International Working Conference on Dependable Computing for Critical Applications*, September 1995, pp. 44–55.
- [15] T. Dierks and C. Allen, “The TLS protocol,” IETF, RFC 2246, January 1999.
- [16] “Job v Halifax PLC, case number 7BQ00307, commentary by Alistair Kelman,” in *Digital Evidence and Electronic Signature Law Review*, vol. 6. Pario Communications, November 2009, ISBN 0-9543245-9-5; see also <http://www.alikelman.com/jobhbos.pdf>.
- [17] S. J. Murdoch, “Reliability of chip & PIN evidence in banking disputes,” in *Digital Evidence and Electronic Signature Law Review*, vol. 6. Pario Communications, November 2009, pp. 98–115, ISBN 0-9543245-9-5.
- [18] E. V. Herreweghen and U. Wille, “Risks and potentials of using EMV for Internet payments,” *USENIX Workshop on Smartcard Technology*, May 1999.
- [19] R. Anderson, M. Bond, and S. J. Murdoch, “Chip and spin,” March 2005. [Online]. Available: <http://www.chipandspin.co.uk/spin.pdf>
- [20] S. Drimer and S. J. Murdoch, “Keep your enemies close: Distance bounding against smartcard relay attacks,” in *USENIX Security Symposium*, August 2007. [Online]. Available: <http://www.usenix.org/events/sec07/tech/drimer/drimer.pdf>
- [21] S. Drimer, S. J. Murdoch, and R. Anderson, “Thinking inside the box: system-level failures of tamper proofing,” in *IEEE Symposium on Security and Privacy (Oakland)*, May 2008, pp. 281–295. [Online]. Available: http://www.cl.cam.ac.uk/~sd410/papers/ped_attacks.pdf
- [22] R. Anderson, “Chip and PIN is broken,” Light Blue Touchpaper, February 2010. [Online]. Available: <http://www.lightbluetouchpaper.org/2010/02/11/chip-and-pin-is-broken/>

APPENDIX

A. Transaction Log of MITM Attack

The following log was collected during one of our man-in-the-middle experiments, where we used one of our own cards to purchase goods in a online Chip and PIN transaction, while using the incorrect PIN. Data items which could be used to identify the merchant who assisted us with the experiments has been redacted (xx), and unnecessary detail has been removed for brevity (...). Principals are Terminal (T), Card (C), and man-in-the-middle (M).

T → C	00 a4 04 00 0e 31 50 41 59 2e 53 59 53 2e 44 44 - 46 30 31	Select file "1PAY.SYS.DDF01"
C → T	6f 1a 84 0e 31 50 41 59 2e 53 59 53 2e 44 44 46 - 30 31 a5 08 88 01 02 5f 2d 02 65 6e 90 00	Opened "1PAY.SYS.DDF01" (language EN)
T → C	00 b2 01 14 00	Read Record
C → T	70 40 61 1e 4f 07 a0 00 00 00 29 10 10 50 10 4c - 49 4e 4b 20 20 20 20 20 20 20 20 20 20 20 87 - 01 01 61 1e 4f 07 a0 00 00 00 03 10 10 50 10 56 - 49 53 41 20 44 45 42 49 54 20 20 20 20 20 20 87 - 01 02 90 00	Available applications: "LINK" and "VISA DEBIT"
T → C	00 a4 04 00 07 a0 00 00 00 03 10 10	Select file "VISA DEBIT"
C → T	6f 25 84 07 a0 00 00 00 03 10 10 a5 1a 50 10 56 - 49 53 41 20 44 45 42 49 54 20 20 20 20 20 20 87 - 01 02 5f 2d 02 65 6e 90 00	Opened "VISA DEBIT" (language EN)
T → C	80 a8 00 00 02 83 00	Get Processing Options
C → T	80 0a 5c 00 08 01 01 00 10 01 04 01 90 00	Transaction started, 5 records available
T → C	00 b2 01 0c 00	Read Record
C → T	70 3e 57...5f 20...9f 1f...90 00	Record (Track 2 Equivalent Data, Cardholder Name, Track 1 Discretionary Data)
T → C	00 b2 01 14 00	Read Record
C → T	70 49 5f 25...5f 24...9f 07...5a...5f - 34...9f 0d...9f 0e...9f 0f...8e 10 00 00 - 00 00 00 00 00 00 41 03 1e 03 02 03 1f 03 90 00	Signed record (Application Effective Date, Application Expiration Date, Application Usage Control, Application Primary Account Number, Application Primary Account Number Sequence Number, Issuer Action Code – Default, Issuer Action Code – Denial, Issuer Action Code – Online, Cardholder Verification Method List)
T → C	00 b2 02 14 00	Read Record
C → T	70 81 93 93...90 00	Record (Signed Static Application Data)
T → C	00 b2 03 14 00	Read Record
C → T	70 81 c0 8f...9f 32...92...90 00	Record (Certification Authority Public Key Index, Issuer Public Key Certificate, Issuer Public Key Exponent, Issuer Public Key Remainder)
T → C	00 b2 04 14 00	Read Record
C → T	70 48 8c 15 9f 02 06 9f 03 06 9f 1a 02 95 05 5f - 2a 02 9a 03 9c 01 9f 37 04 8d 17 8a 02 9f 02 06 - 9f 03 06 9f 1a 02 95 05 5f 2a 02 9a 03 9c 01 9f - 37 04 9f 08...5f 30...5f 28...9f 42...9f - 44...90 00	Record (Card Risk Management Data Object List 1 (CDOL1), Card Risk Management Data Object List 2 (CDOL2), Application Version Number, Service Code, Issuer Country Code, Application Currency Code, Application Currency Exponent)

T → C 80 ca 9f 17 00
C → T 9f 17 01 03 90 00

T → M **00 20 00 80 08 24 00 00 ff ff ff ff ff**
M → T **90 00**

T → C 80 ae 80 00 1d xx xx xx xx xx xx 00 00 00 00 00 -
00 08 26 00 80 00 80 00 08 26 xx 11 09 00 xx xx -
xx xx

C → T 80 12 80 xx xx xx xx xx xx xx xx xx xx 06 01 0a -
03 a0 00 10 90 00

T → C 00 82 00 00 0a xx xx xx xx xx xx xx xx 30 30

C → T 90 00

T → C 80 ae 40 00 1f 30 30 xx xx xx xx xx xx 00 00 00 -
00 00 00 08 26 00 80 00 80 00 08 26 xx 11 09 00 -
xx xx xx xx

C → T 80 12 40 xx xx xx xx xx xx xx xx xx xx 06 01 0a -
03 60 00 10 90 00

Get Data (PIN try counter)
Remaining PIN tries = 3

Verify PIN "0000"
PIN correct

Generate AC (ARQC)

ARQC

External Authenticate
External authenticate successful

Generate AC (TC)

TC