All You Can Eat

or

Breaking a Real-World Contactless Payment System⋆ (short paper)

Timo Kasper, Michael Silbermann, and Christof Paar

Horst Görtz Institute for IT Security, Ruhr-University Bochum, Germany
{Timo.Kasper, Michael.Silbermann, Christof.Paar}@rub.de

Abstract. We investigated a real-world contactless payment application based on MIFARE Classic cards. In order to analyze the security of the payment system, we combined previous cryptanalytical results and implemented an improved card-only attack with customized low-cost tools, that is to our knowledge the most efficient practical attack to date. We found several flaws implying severe security vulnerabilities on the system level that allow for devastating attacks including identity theft and recharging the amount of money on the cards. We practically verify and demonstrate the attacks on the commercial system.

1 Introduction

A growing number of payment systems incorporate contactless technology, as it offers additional benefits in terms of flexibility and convenience over its contact-based counterpart. With the recent trend towards issuing contactless smartcards in large companies, universities and government entities, a number of privacy- and security-related concerns have been raised.

The "ID-Card" analyzed in the following is, according to the manufacturer, used by more than a million people in Germany. The multi-purpose electronic ID and payment card is based on a dual-interface smartcard, i.e., both contact-based and contactless interfaces are provided. Besides the use for payments, e.g., for food, printing services and washing machine, the functionality of the contactless part includes access control to workplace and apartment, and automatic recording of the working hours of employees. The ID-Card can be charged at dedicated charging terminals with a maximum amount of € 150. The wireless technology implies new threats compared with contact-based systems, for instance, a card could be read out from the pocket or wallet without the owner taking note of it. Thus, we examine the security of the ID-Card as an example for a widespread payment system and intend to answer the question: How secure are today's electronic payment systems really? We thereby focus on a realization of the ID-Card e-cash system that is operational since fall 2006 at a large enterprise in Germany.

Through interacting with the card we identified it as a MIFARE Classic from NXP, for which a number of attacks have been published recently, as summarized in Sect. 2. We then used our custom-built radio-frequency identification (RFID) hardware described in Sect. 3.1 to implement the most efficient practical attack on MIFARE Classic known to date and to extract the secret keys from the card. During the subsequent analysis of the system we discovered severe flaws that enable a variety of real-world attacks, as practically evaluated in Sect. 3. The dramatic implications of the attacks are finally described in Sect. 4.

Contactless Payment Systems A typical contactless payment system consists of RFID readers at the points of sales, and contactless smartcards in the field that operate as electronic wallets and supposedly store the current balance safely. A reader generates a strong electro-magnetic (EM) field at a frequency of 13.56 MHz for supplying a card with energy for its operation and starts the communication. Contactless smartcards, as standardized in ISO 14443 [6], get activated from up to 25 cm [8], while the communication can be passively eavesdropped from a distance of several meters [11].

To allow for rapid operation, the cash registers typically work autonomously without an on-line connection to a database in the back-end, and synchronize their data, for instance, blocked cards or the balance on the account of the cash register, only infrequently, e.g., once per day. This scenario also applies to the payment system analyzed in this paper. For some systems, so-called "shadow accounts" exist for each card that contain the amount of money stored at the instant of the last synchronization and could be used for detecting fraud or functional errors in the system.

MIFARE Classic Basically, the contactless MIFARE Classic cards are memory cards, i.e., the information is stored in an internal non-volatile memory with an integrated digital control unit to handle the communication with a reader. Furthermore, the interchanged bitstreams can be encrypted. The cards generally comply to Parts 1-3 of the ISO/IEC 14443A [6], but are using a higher-level communication protocol that diverges from Part 4 of the standard. This proprietary protocol for authentication and subsequent data encryption promises to prevent replay attacks, cloning, and eavesdropping by means of the proprietary "CRYPTO1" cipher. The stream cipher is based on a 48-bit linear feedback shift register and six non-linear filter functions. In addition to the small key space, its cryptanalysis revealed several security vulnerabilities, hence CRYPTO1 is commonly regarded as cryptographically weak. A more detailed description of the cipher and its operation principle can be found in [5].

The memory of a MIFARE Classic 1K chip is partitioned into 16 sectors, each consisting of four 16-byte blocks. The read-only block 0 of the first sector contains the factory-programmed Unique Identifier (UID) of the card and the fourth block of each sector contains amongst others two 48-bit keys for the authentication to that sector.

2 Previous Attacks on MIFARE Classic

Since its invention in 1995 by Philips [10], all details of the MIFARE Classic chip had been kept secret until 2007, when the cipher was reverse-engineered [9]. Afterwards, the research on MIFARE Classic revealed numerous security-relevant vulnerabilities, providing the basis for our system break.

Keystream Recovery The first discovered weakness [9] is that a random nonce n_C generated by the card depends on the time elapsed between the power-up of the card and the issuing of the authentication command by the reader. Hence, the authors are able to reproduce the same nonce with a certain probability by controlling the timing.

A first consequence of this weakness is a keystream recovery attack [3] requiring a recorded authentication session between a genuine reader and a MIFARE Classic card. Afterwards, several queries to the card are issued by a specially prepared reader. Altogether, all 16 bytes of keystream for sector 0 and up to 12 keystream bytes for higher sectors can be recovered. The attack does not enable card cloning, as the cryptographic keys remain secret. However, it is possible to read or modify the whole sector 0 and partially the higher sectors.

Key Recovery from Genuine Authentications The possibility to recover parts of the keystream led to an improved attack [4] that allows to extract a secret key from just two eavesdropped authentications between a card and a genuine reader. No precomputation is required and, after recording the two authentication sessions, it takes only 0.1 s to recover the secret key of one sector, using ordinary computers.

The attack further exploits another weakness of MIFARE Classic cards: Instead of computing parity bits from the actual (encrypted) bits transmitted, the parity is derived from the plaintext. In addition, the same bit of the keystream that is used to encrypt the parity bit of byte N is again used to encrypt the first bit of the next byte $N + 1$ sent.

Card-Only Key Recovery In a card-only scenario, an attacker only needs to be close enough to the targeted MIFARE Classic card to activate it and communicate with it by means of a special-purpose reader, in order to recover a secret key. Garcia et al. [5] propose four different attacks to obtain the secret key of one sector. Their most time-effective approach requires a precomputed table with a size of 384 GB. On average 4096 authentication attempts are required, thereby keeping the nonce of the reader n_R constant while varying n_C, to obtain one secret key in about two minutes. In addition, if one sector key is known to the attacker, an authentication to that sector can be decrypted to obtain the corresponding n_C. Due to a weakness of the random number generator in the card, the subsequent n_C that is used for authenticating to another sector can be predicted, to recover 32 bit of the keystream generated by the secret key of the new sector. After three authentication attempts 64 bit of keystream are obtained and used in an offline computation step to find the correct key in less than 1 s on

a standard PC. The authors make use of the fact that the card answers with an encrypted `NACK = 0x5` command under certain conditions. This known plaintext allows to establish a side-channel to recover four bits of the keystream with a probability of 1/256 per authentification attempt.

The latest and most efficient card-only attack [1] requires only 300 queries to the card on average and a few seconds of off-line computation to find the sector key. It exploits a mathematical vulnerability of the filter functions used in CRYPTO1 to mount a differential attack on the cipher. The off-line computations are negligible, as on average only $2^{48-32} = 2^{16}$ key candidates need to be tested in order to recover a secret key. Note, that the authors of the above summarized papers report difficulties in fixing the card nonce n_C in practice.

3   Tampering with a Real-World System

Analyzing the security of the payment system, we observed that the ID-Card replies with `ATQA = 00 04` and `SAK = 08`, indicating that it contains MIFARE Classic chip [12]. Hence, we performed a practical key-recovery, as detailed in the following.

3.1   Hardware and Software Set-Up

For the security analyses we use a self-built, freely programmable device termed "RFID Tool" [7]. In contrast to commercially available products, our RFID reader allows to fully control the communication and the RF field with a high timing accuracy. This includes arbitrary chosen challenges for the authentication protocol, intentionally wrong calculated checksums and modified parity bits.

The microcontroller of the RFID Tool was programmed to support the full MIFARE Classic authentication protocol described in Sect. 1, and to allow comfortable reading, writing, and cloning of MIFARE Classic cards. The software for the key-recovery is subdivided into two parts: During an on-line step, the embedded part, running on the RFID Tool, emulates a MIFARE Classic reader which collects the data needed for the attack. The data acquisition can be performed in the field by powering the RFID Tool from a battery and storing the acquired information internally. The data is afterwards sent via USB to a desktop PC, where the second part of the software computes the sector keys off-line.

3.2   Recovering the Secret Keys

We implemented our attack based on a combination of the existing attacks [1, 5], as described in Sect. 2, and an open-source implementation of the CRYPTO1 cipher [2]. During a normal protocol run, the response times of the MIFARE card vary. Hence, using commercial readers, the same card nonce n_C can only be reproduced with a relatively low probability during an authentication. In order to precisely fix the timing, we use the capability of our reader to wait a fixed multiple of 75 ns between the power-up of the card and the authentication

command. As a result, we can force the card to generate exactly the same n_C in *every* attempt, which highly improves the efficiency of our key-recovery attack. We found that the EM field has to be turned off for approx. 70 ms to ensure a complete reset of the card – smaller time windows did not allow to fix n_C to one value. This defines a new practical lower bound for the time required for an attack and questions some theoretical estimations in the literature.

With our current hardware setup it takes less than 30 s to perform the required authentication runs for revealing one sector key of the card. Once the data is collected, it takes less than 3 s to recover the key on a standard PC. To extract all 16 sector keys from the card we need less than $16 \cdot (30 + 3) = 528 \, \text{s} \approx 9 \, \text{min}$, which is by a factor eight faster than the 80 min for 16 sectors proposed in [1]. No precomputation or other data storage is needed for our attack.

We tested our implementation on an ID-Card and successfully revealed all its secret keys, thereby noticing that the revealed keys differ from the factory-programmed default keys of MIFARE Classic. When analyzing a second card, we found that identical secret keys are used for all sectors. We conjectured that all ID-Cards in the system might have the same keys. By analyzing a dozen more cards we verified the hypothesis: *All* ID-Cards in the payment system have *identical* sector keys.

From comparing the content of several ID-Cards, before and after carrying out charges and payments, we learned which data the smartcard contains apart from the read-only UID and where it is stored in the memory. We revealed that only some bytes of the first three sectors (0,1, and 2) are affected by monetary transactions, while most parts of the memory remain unused. In more detail, sector 0 contains a card number, which is for some ID-Cards also printed on the card, and a checksum that is calculated as an XOR over all bytes of blocks 1 and 2. The credit value is stored twice in the value blocks 4 and 5 of sector 1 on the card and is not secured by any checksum or other means. Sector 2 contains some information about the last charging and payment process, e.g., a time stamp, a transaction number, and the ID of the last terminal used.

3.3 Practical System-Level Tests

To determine which threats emerge from the vulnerabilities of the system, we carried out a number of tests as detailed in the following. First, we charged an empty ID-Card with, e.g., € 5. We copied it to a blank card – which naturally has a different UID – and then pay with this (almost) cloned card: The cloned card behaves exactly as a genuine one, and the money is withdrawn (e.g., € 5 → € 2.70), implying that no checks of the UID are performed.

Three hours later, we paid again with the cloned card to see whether it is now detected or blacklisted. Still, we were able to carry out our payment (e.g., € 2.70 → € 2). Afterwards, we inserted the original ID-Card still charged with, e.g., € 5 and the cloned one (e.g., with € 2 balance) into the charging machine: The existence of two cards with identical card numbers but different credit amounts is *not* detected, one card shows € 5, the other one € 2.

Four days later we tried to pay with the original ID-Card charged with € 5. The money was withdrawn (e.g., € 5 → € 2.20), indicating that shadow accounts are not used or at least no action is taken based on inconsistencies.

Finally, we changed the card number, created some fake credit value (e.g., € 2.30) and wrote the content to a blank card, thereby taking the check byte at the end of block 2 into account. Even this newly produced card was recognized as genuine and accepted for payments.

3.4 Summary of the Vulnerabilities

We conclude that the security of the analyzed ID-Card payment system is extremely low based on today's cryptanalytical knowledge, since it relies solely on weak MIFARE Classic cards with identical keys. Once the secret keys are extracted from one card, all cards of the system can be read and written to wirelessly from up to 25 cm [8] in less than a second. According to our practical results, the one-time key-extraction is easily performed in minutes. All data on the cards is stored in plaintext and almost no integrity checks are performed. Neither the UID of cloned cards is verified, nor are cards with the same card number but a different balance detected. Seemingly, no additional checks are performed, neither in the front-end, nor in the back-end of the system, allowing for various attacks as illustrated in the following.

4 Resulting Attacks and their Implications

The implications of the security flaws described in Sect. 3 are dramatic. Once the secret keys are recovered from one card, a variety of devastating attacks can be mounted for any card in the system. The adversary does not need to be an employee to gain access to a card, as an anonymous ID-Card can be obtained at any cashpoint for a deposit of € 10. Blank MIFARE Classic cards can be purchased for less than € 0.50 on the Internet. All our attacks, including the one-time key extraction, require no special knowledge, but can be performed by any attacker who possesses our public-domain reader and the corresponding software. In court, the actions would be difficult to prove, because no physical traces are left when wirelessly modifying cards.

Impersonation An adversary needs about 40 ms to covertly extract the card number, and the credit value from the ID-Card of a victim from a distance. Then she has two options: For *digital pickpocketing*, she will lower the money on the card of the victim (requiring another 40 ms) and produce a new card with the same card number and the "stolen" amount of money on it. The adversary can now pay with an ID-Card that is known to the system, using the money of the victim. Hence, the fraud will not be detected in the back-end. This option would harm only the victim, while no damage is caused for the issuing institution.

As a second option, the attacker *impersonates* the victim by generating a new card with the extracted card number and the extracted credit value, to obtain a

duplicate. In addition, she can now increase the amount of money on her card, which will (according to our tests in Sect. 3) still not be noticed by the current realization of the payment system. This time, the card of the victim remains unchanged, and all losses are on the side of the issuing institution.

Trafficking ID-Cards An adversary can produce counterfeit ID-Cards with a new, random card number and a fake credit value of, e.g., € 100, and sell them to others. It is also conceivable that a criminal offers a service to charge the cards of other users. Both option imply high losses for the issuing institution.

Denial of Service For a Denial of Service (DoS) attack, we assume an attacker who wirelessly resets credit balances of ID-Cards in the field to a zero value, while leaving the rest of the data unchanged. Performing this attack would harm and confuse all affected legitimate card owners, and result in considerable costs for the customer service and loss of credibility for the contactless payment system.

Distributed All-You-Can-Eat In contrast to a DoS attack, an attacker can wirelessly set credit balances of ID-Cards in the field to a high value, e.g., € 100. For both DoS and this rather anarchistic "all-you-can-eat" attack, modifying the two value blocks storing the credit value takes again just about 40 ms. Hence, the attacks can be carried out easily in practice, e.g., by setting up a disguised reader device near a waiting line of the point of sales that modifies the balance of anyone who gets close enough to it accordingly. It is unlikely that a legitimate user of the payment system complains about having too much money on his card, hence the fraud might be detected very late and the financial losses for the institution would be dramatic. In the long term, the attack would render the payment system inoperative.

Emulation of Arbitrary Cards Employing an electronic device that can emulate a MIFARE card, e.g., the "fake tag" developed in [7], an adversary can emulate any ID-Card including its UID, e.g., to behave like an exact clone of a card. Thus, the fake tag can be used (hidden in the wallet) to pretend the presence of a new card with a random UID, a random card number and a random credit value for every payment, thus making detection of fraud and blacklisting impossible. For criminals, selling this device could be very profitable, since it allows for unlimited payments.

5   Conclusion

We summarized the existing attacks on the MIFARE Classic and implemented the most efficient practical card-only attack to date using our custom-built equipment at a cost of below € 40. Testing it at hand of the widespread ID-Card payment system we show that recovering the relevant secret keys takes less than 2 min. Once the keys are compromized, the security of the whole system collapses instantaneously, as it turns out that no additional cryptographic mechanisms or other checks are implemented.

Our subsequent analysis of the ID-Card payment application reveals obvious vulnerabilities that pose a great threat to its overall security. An adversary can, in 40 ms and imperceptibly for the victim, read out a card or write to it, increase or decrease its credit balance, impersonate the victim or simply clone his card. Furthermore, a criminal can sell counterfeit cards or electronic devices that emulate a new random card, and hence permit an unlimited amount of payments. Prosecution of the adversary or the victims is not promising. Since the attacks leave no physical traces, it is difficult (or in some cases impossible) to prove that a crime has been committed.

The security flaws do not rely solely on weaknesses of the contactless smartcards used, but are mainly caused by the realization on the system level. The demonstrated attacks, that can be performed by non-specialists, would become a lot harder or even infeasible in practice, if the system integrator would address the problems detailed in this paper. Using basic cryptographic knowledge, countermeasures could be implemented to obtain a higher security level that would probably be sufficient in the context of micropayments, even on the basis of the weak MIFARE Classic cards.

References

1. N. Courtois. The Dark Side of Security by Obscurity - and Cloning MiFare Classic Rail and Building Passes, Anywhere, Anytime. In *SECRYPT*, pages 331–338. INSTICC, 2009.
2. Crapto1. Open Implementation of CRYPTO1. `http://code.google.com/p/crapto1/`, 2008.
3. G. de Koning Gans, J.-H. Hoepman, and F. D. Garcia. A Practical Attack on the MIFARE Classic. In *CARDIS*, volume 5189 of *LNCS*. Springer, 2008.
4. F. D. Garcia, G. de Koning Gans, R. Muijrers, P. van Rossum, R. Verdult, R. W. Schreur, and B. Jacobs. Dismantling MIFARE Classic. In *ESORICS*, volume 5283 of *LNCS*. Springer, 2008.
5. F. D. Garcia, P. van Rossum, R. Verdult, and R. W. Schreur. Wirelessly Pickpocketing a Mifare Classic Card. In *IEEE Symposium on Security and Privacy*, pages 3–15. IEEE, 2009.
6. ISO/IEC 14443-A. Identification Cards - Contactless Integrated Circuit(s) Cards - Proximity Cards - Part 1-4. www.iso.ch, 2001.
7. T. Kasper, D. Carluccio, and C. Paar. An Embedded System for Practical Security Analysis of Contactless Smartcards. In *WISTP*, volume 4462 of *LNCS*, pages 150–160. Springer, 2007.
8. I. Kirschenbaum and A. Wool. How to Build a Low-Cost, Extended-Range RFID Skimmer. In *USENIX Security Symposium*, 2006.
9. K. Nohl, D. Evans, Starbug, and H. Plötz. Reverse-Engineering a Cryptographic RFID Tag. In *USENIX Security Symposium*, pages 185–194, 2008.
10. NXP. Mifare Classic. `http://www.nxp.com`, 1995.
11. NXP. AN200701: ISO/IEC 14443 Eavesdropping and Activation Distance. Technical report, 2007.
12. NXP. Mifare Classic 1K MF1 IC S50 Functional Specification. `www.nxp.com`, 2008.