A Practical Attack on *Patched* MIFARE Classic

Yi-Hao Chiu[1], Wei-Chih Hong[2], Li-Ping Chou[3], Jintai Ding[4,5],
Bo-Yin Yang[2(✉)], and Chen-Mou Cheng[1]

[1] National Taiwan University, Taipei, Taiwan
[2] Academia Sinica, Taipei, Taiwan
by@crypto.tw
[3] Chinese Culture University, Taipei, Taiwan
[4] University of Cincinnati, Cincinnati, USA
[5] Chongqing University, Chongqing, China

Abstract. MIFARE Classic is the world's most widely deployed RFID
(radio-frequency identification) technology. It was claimed to be crypto-
graphically protected by the proprietary Crypto-1 stream cipher. However,
it proved inadequate after weaknesses in the design and implementation of
Crypto-1 and MIFARE Classic started surfacing since late 2007
[7,8,12–17]. Some operators of MIFARE Classic-based systems reacted by
upgrading to more secure alternatives such as MIFARE DESFire. How-
ever, many (especially in Asia) opted to "patch" MIFARE Classic instead.
Their risk analysis might have gone as follows: *"The most serious threat
comes from efficient card-only attacks, where the attacker only needs an
off-the-shelf reader and a PC to tamper a target tag. All efficient card-only
attacks depend on certain implementation flaws. Ergo, if we just fix these
flaws, we can stop the most serious attacks without an expensive infrastruc-
ture upgrade."* One such prominent case is "EasyCard 2.0," today accepted
in Taiwan as a means of electronic payment not only in public transporta-
tion but also in convenient stores, drug stores, eateries, cafes, supermar-
kets, book stores, movie theaters, etc. Obviously, the whole "patching"
approach is questionable because Crypto-1 is fundamentally a weak cipher.
In support of the proposition, we present a new card-only attack based on
state-of-the-art algebraic differential cryptanalytic techniques [1,2]. Still
using the same cheap reader as previous attacks, it takes 2–15 min of com-
putation on a PC to recover a secret key of EasyCard 2.0 after 10–20 h of
data collection. We hope the new attack makes our point sufficiently clear,
and we urge that all MIFARE-Classic operators with important transac-
tions such as electronic payment upgrade their systems to the more secure
alternatives soon.

Keywords: RFID security · MIFARE Classic · Card-only attack · Alge-
braic cryptanalysis

1 Introduction

MIFARE Classic, a brand owned by the NXP Semiconductors, is the most
widely used RFID technology in the world today, with billions of chips sold

worldwide. It is used in many public-transportation ticketing systems in, e.g., Beijing, Chongqing, Guangzhou, Boston, the Netherlands, London, Seoul, Taipei, etc. In recent years, it has even found its way into electronic payment systems in several Asian countries including China and Taiwan.

The proprietary Crypto-1 stream cipher is designed to provide cryptographic protection to MIFARE Classic. NXP Semiconductors has never made public the detailed algorithm of Crypto-1. Nevertheless, starting from late 2007 in a series of papers, the specifications and several weaknesses of the cipher have been found via reverse engineering and cryptanalysis [7,8,12,13,15–17]. As Courtois *et al.* concluded: "The security of this cipher is therefore close to zero" [8]. Users of MIFARE Classic around the world responded differently to this incident. Some kept silent, while others promptly announced plans of replacing MIFARE Classic—unfortunately not always with more secure technologies.

In this paper, we shall investigate in detail one such replacement being deployed in Taipei, an early adopter and aggressive user of MIFARE Classic. Branded under the name "EasyCard," more than 35 million cards have been issued in Taipei since the official release in 2002, with more than 4.6 million transactions per day in 2012. Starting from 2010, the card is also accepted as a means of electronic payment by almost all convenient store chains, as well as drug stores, eateries, cafes, supermarkets, book stores, movie theaters, etc. Similar use of MIFARE Classic is reported in several cities in China including Beijing, Chongqing, and Guangzhou.

In a nutshell, not only does Crypto-1 use way too short a key (48 bits) by today's standards, its cipher structure also allows very easy recovery of its internal state (and hence the secret key) if the attacker learns a small number of contiguous keystream bits [12]. This allows a sniffer to recover the secret key if it is placed in proximity when a pair of legitimate reader and tag are in a transaction.

In addition, there are two serious implementation flaws which also cause weaknesses: (i) *parity bits are computed over plaintext and then encrypted*; (ii) *the 32-bit tag nonces used in the authentication satisfy a degree-16 linear recurrence relation and can be controlled by appropriately timing the authentication attempts*. Furthermore, there is a convenient way for the attacker to extract information on keystream bits from (i), as a tag would respond *differently* depending on whether the parity bits are correct or not. Together, they allow extremely efficient attacks even when the attacker only has access to the tag [13].

Compared with sniffer-based attacks, these efficient card-only attacks are arguably *much more serious* because of the low entry barriers. All the attacker needs is a PC and a cheap, off-the-shelf reader, so any ordinary person can launch such an attack in private by downloading the appropriate software from the Internet.

In late 2012, the EasyCard Corporation rolled out "EasyCard 2.0," a dual-interface smart card that is compatible with existing EasyCard readers, yet with all implementation flaws *fixed*. The tag nonces seem random, both unpredictable and uncontrollable, and the tag responses are indistinguishable whether the parities sent by the reader are correct or not. This renders all existing efficient

card-only attacks [7,8,12,13,15] *ineffective*, as we have verified through experiments. This does not stop, of course, brute-force attacks, which are arguably less threatening because it takes years of computation on an ordinary PC. The attacker would need to have access to expensive supercomputers, e.g., GPU or FPGA clusters, in order to recover the keys within a reasonably short amount of time [5]. As a result, the EasyCard Corporation seems confident that Easy-Card 2.0 can be "reasonably secure," as the computational power required by brute-force attacks is way beyond the reach of an ordinary person.

In this paper, we will show that such a sense of security is *false*. Namely, we will present a new card-only attack based on state-of-the-art algebraic differential cryptanalytic techniques [1]. The attack is highly practical: it uses the same cheap reader as previous attacks [7,8,12,13,15] and takes 2–15 min on a PC to recover the secret key of EasyCard 2.0 or other similar implementations of MIFARE Classic. The extra price the attacker needs to pay for the new attack is a slightly longer time for data collection, typically 10–20 h. We note that this is not atypically long for differential attacks and still makes the new attack a serious threat because the data collection can be done by the attacker in private without needing access to a legitimate reader. Overall, this is a significant improvement over the brute-force attacks, which would take about 4 years on the same PC.

The rest of this paper is organized as follows. In Sect. 2, we will first give some background information on the cipher itself and the cryptanalytic techniques we have used to attack it. We will then present our new attack in Sect. 3 and experiment results in Sect. 4. Finally, we will discuss the implications and conclude this paper in Sect. 5.

2   Background and Related Work

2.1   Crypto-1 and the MIFARE Classic Authentication Protocol

Crypto-1 is a stream cipher used to provide cryptographic protection to MIFARE Classic tags and contactless smart cards. For more than a decade, its design was kept secret by NXP, along with the rest of MIFARE Classic. After the details of MIFARE Classic was reverse-engineered in 2007 [12,16,17], many weaknesses have been discovered, and with them many attacks. These attacks vary greatly in efficacy. The first few key-recovery attacks exploit the weaknesses of the cipher and gather the required information either by direct communication with a legitimate reader or by eavesdropping a communication session. Although some system vendors argue even today that these attacks are impractical, the cipher itself was by then considered cryptographically broken.

A few months later, better, card-only attacks were published [13]. These exploit several properties in the authentication protocol of MIFARE Classic as well as flaws in generating tag nonces.

For the sake of completeness, we include here a brief description of Crypto-1 and its use in the authentication protocol of MIFARE Classic. Crypto-1 uses a 48-bit linear feedback shift register (LFSR) with nonlinear output filter [12]. The feedback function of the LFSR is $F(s_0, s_1, \ldots, s_{47}) := s_0 \oplus s_5 \oplus s_9 \oplus s_{10} \oplus s_{12} \oplus$
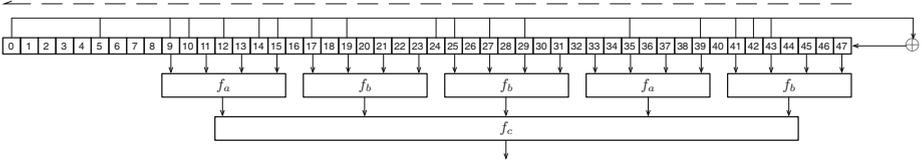
Fig. 1. The structure of the Crypto-1 stream cipher

$s_{14} \oplus s_{15} \oplus s_{17} \oplus s_{19} \oplus s_{24} \oplus s_{25} \oplus s_{27} \oplus s_{29} \oplus s_{35} \oplus s_{39} \oplus s_{41} \oplus s_{42} \oplus s_{43}$. With every tick of the clock, 20 bits from the LFSR are fed into the function f to generate one new bit of the keystream. Then the LFSR shifts one bit to the left, and the new rightmost bit is filled by the output of F—or, if the operational phase calls for inputs, F `XOR`ed with an input bit. F is *primitive*: the LFSR has a period of $2^{48} - 1$, the maximum possible.

The function f or *output filter* consists of two layers of nonlinear functions. The first layer is a mixed combination of two 4-input nonlinear functions f_a and f_b, and the second layer is a 5-input function f_c. Here, $f_a = $ `0x2c79`, $f_b = $ `0x6617`, $f_c = $ `0x7907287b` in "table form" (collating the output bits as the input goes lexicographically over its range), and f can then be expressed as

$$f(s_0, \ldots, s_{47}) := f_c(f_a(s_9, s_{11}, s_{13}, s_{15}),$$
$$f_b(s_{17}, s_{19}, s_{21}, s_{23}), f_b(s_{25}, s_{27}, s_{29}, s_{31}),$$
$$f_a(s_{33}, s_{35}, s_{37}, s_{39}), f_b(s_{41}, s_{43}, s_{45}, s_{47})). \tag{1}$$

Note that each has an equal number of 0 and 1 bits and hence outputs 0 or 1 each with probability 1/2 if input bits are independently and uniformly distributed over $\mathbb{F}_2$ [13].

On being powered up by the reader's electromagnetic field, the tag sends its unique identifier uid to the reader to start the anti-collision phase. The reader may then request to authenticate a specific block. On receiving the request, the tag loads the secret key for the block as the initial state of the cipher and sends a randomly chosen challenge nonce n_T to the reader. Meanwhile, $n_T \oplus$ uid is shifted into the LFSR. All subsequent communication is encrypted with the keystream, and we will use the notation $\{X\}$ to represent the ciphertext of X, i.e., $X \oplus$ keystream. Next, the reader picks its challenge n_R, which will also be shifted into the LFSR, and sends $\{n_R\}$ followed by the answer $\{a_R\}$ to the tag's challenge. Finally, the tag replies with its answer $\{a_T\}$ to conclude the authentication procedure (see Fig. 2). If the tag and the reader used the same secret key for the initial state of their ciphers, this authentication procedure should bring the ciphers on either side to the same internal state, and the two keystreams generated by both ends will be henceforth in synchronization.

2.2   Existing Card-Only Attacks Against MIFARE Classic

The best known attacks have been summarized by Garcia *et al.* [13], which we will recapitulate here for the sake of completeness. The card-only attacks mainly exploited the following weaknesses.
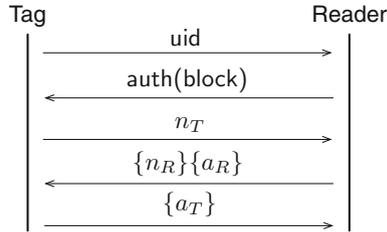
Fig. 2. The authentication protocol in MIFARE Classic

1. The communication of MIFARE Classic follows the ISO 14443-A standard, which requires that a parity bit be sent after every 8 bits of transmission. However, in MIFARE Classic, these parity bits are computed over the plaintext, and the keystream bit used to encrypt the parity bits is reused to encrypt the next bit of plaintext. Furthermore, during authentication, the tag would not reply anything if the received messages have incorrect parity, i.e., the tag checks the authenticity of the reader's answer only if the parity bits are correct.
2. If all parity bits are correct but the encrypted answer $\{a_R\}$ to the tag's nonce cannot be correctly verified, the tag responds with an encrypted 4-bit NACK code. Since the NACK code is fixed, this leaks 4 keystream bits.
3. The 32-bit tag nonce is actually generated by a 16-bit LFSR that runs in a deterministic cycle after it powers up, i.e., timing is used as the source of randomness to the internal random number generator (RNG). Therefore, controlling or measuring when the reader sends every authentication request basically gives us control or a very good guess to the next tag nonce.
4. When a reader is already communicating with a tag (i.e., having authenticated to certain sector), the protocol of a subsequent authentication for a new sector differs slightly from the initial one in that the tag nonce will be encrypted by the new sector key before transmitted to the reader. Since the first tag nonce was sent in plaintext, and the timing between two authentication attempts is known, the attacker can guess the second tag nonce and recover 32 bits of keystream with high accuracy.

Taking advantage of the weakness in the parity bits, the attacker can ask to authenticate for a sector of the tag at hand and answer the tag's challenge with random $\{n_R\}$ and $\{a_R\}$ (totally 8 bytes) accompanied with 8 random parity bits. On average, one out of 256 trials will the attacker receive the encrypted NACK code from the tag. Each such trace reveals 12 bits of information (8 from parity bits and 4 from NACK code) on the secret key. In practice, six traces are enough for the offline brute-force check of the secret key. It takes $6 \cdot 256 = 1536$ trials on average to gather these traces and can be accomplished within a minute. The offline part of this attack is to check which key out of the 2^{48} possible keys generates all "correct" parity and NACK code bits in these traces, and the computing

time depends on the implementation realized by the attacker. Pessimistically, the run-time of checking on a powerful FPGA cluster like COPACOBANA is around half an hour.

Two other attacks try to trade online communication for the offline computing time using the weakness that tag nonces n_T can be controlled precisely by timing the authentication requests. The attacker may substantially reduce offline search space by fixing either tag or reader nonce while varying the other, and look for specific properties.

In the second card-only attack [13], the tag nonce n_T is fixed. The attacker searches for a reader nonce n_R such that flipping the last bit in each byte of n_R also flips the following encrypted parity bit, which averages around 28500 authentication attempts or 15 min. Such an n_R let us cut approximately 15 bits from the offline search space, enabling a standard desktop to finish the computation in around 1 min.

For the third attack [13], the attacker fixes response of the reader to $\{n_R\} = 0 = \{a_R\} = 0$, and searches for an n_T such that the tag responds with the desired encrypted NACK code. For example, it might be desired that the keystream bits are all zero, which means that the ciphertext would be identical to the plaintext. Such search takes 4096 attempts on average since we need 12 bits (8 parity plus 4 keystream bits) to be exactly zero. The direct offline search in a huge precomputed table (with around 2^{36} entries) of the cipher states that could lead to such pattern may take about one day. However, with some further attempts to find the parity bits that correspond to the same n_T but different n_R and a_R (e.g., $\{n_R\} = \{a_R\} = \texttt{0xffffffff}$), one can split the table into 4096 parts. This not only makes it easier to store and read the table but also speeds up the offline search significantly.

A fourth attack [13] tries to derive from a known sector key 32 keystream bits generated by another unknown sector key. Because Crypto-1 is structured such that the internal state can be separated into odd- and even-numbered bits, this allows us to further reduce the search space in exploiting the parity-bit weakness. As a result, the attacker can determine the second sector key in less than a second of computation time after about three nested authentication attempts.

Impact and current countermeasures. The last attack is particularly critical as it takes very little time to recover additional keys once a first key is known, making it feasible to "pickpocket" a card wirelessly if a deployed system leaves unused sectors with default keys or does not diversify keys. In response to the attacks outlined above, several countermeasures have been implemented in newer versions of MIFARE Classic cards, such as EasyCard 2.0, that are still compatible with legacy systems.

First and most importantly, the generator of the tag nonce is replaced by a better RNG such that we can no longer control or predict n_T. From our experiments, it seems a true 32-bit RNG instead of having a period of 2^{16}. This improvement breaks almost all efficient card-only attacks depicted above except the brute-force attack, as all the techniques to reduce search space make use of the flaw in tag nonce generation. Furthermore, these new cards now always

reply with encrypted NACK code if the authentication fails, whether the parity bits are correct or not. This closes the last loophole of the brute-force attack described above, as it is no longer possible to gather the required information to attack the parity-bit weakness.

2.3   Algebraic Differential Cryptanalysis

Algebraic cryptanalysis brings the concept of applying algebraic techniques to attack various cryptographic primitives, e.g., block ciphers, stream ciphers, and public key systems. It is usually done in two major steps. First, a set of multivariate polynomial equations over a finite field is constructed to describe the cryptographic scheme. This system of equations is formulated in a way that its solutions correspond to certain secret information of the cryptographic scheme. The second step is then to solve the system using techniques such as SAT solvers or Gröbner-basis algorithms. As a result, the efficiency of this category of attacks is strongly related to the quality of the constructed equations as well as the performance of the system-solving technique in use.

The idea of algebraic cryptanalysis is not new. Back in 1949, Shannon already noted the relationship between breaking a good cipher and solving a complex system of equations [18]. Shannon was probably thinking about how to build a good cipher, but this concept gives us a hint of checking possible weaknesses of cryptosystems using algebraic techniques. However, it was not until the huge progress in the efficiency of system solving, especially the solving of multivariate polynomial systems, that people started to consider system-solving as legitimate attacks. The invention of F_4 [10], XL [6], F_5 [11], and their variants greatly boosted the speed of solving multivariate polynomial systems. Also, the substantial advances in the performance of SAT solvers [9,19] provides us with an alternative, namely to transform problems into boolean formulas and search for solutions.

Differential cryptanalysis exploits information leaked by special pairs of input and output differences, called *differentials*, in a block cipher to distinguish its output from random or to recover (some of) its key bits [3]. Such an attack is statistical in nature and usually requires a large number of plaintext-ciphertext pairs, especially in the context of known-plaintext or ciphertext-only attacks, for which the attacker cannot freely choose the plaintexts. Even before its publication, differential cryptanalysis has played a very important role in cipher design. It is so successful that today's standard procedures for designing a new cipher include checking differential immunity.

In the recent seminal work [1], Albrecht and Cid tried to incorporate the information obtained from differential characteristics into algebraic attacks. They proposed three methods, labeled simply as Attack A, B, and C, of obtaining and using such information. Even though Attack B was not very effective against the PRESENT cipher [20], it did inspire our attack, which we will describe in more detail in Sect. 3.

3   Construction of Attacks

In this section, we illustrate the proposed attack against the patched MIFARE Classic cards based on modern cryptanalysis techniques. The critical step of a successful algebraic attack is to construct as many informative multivariate equations as possible. Three types of algebraic equations are collected in our attack, namely NACK equations, differential equations, and filter equations.

3.1   NACK Equations

As mentioned in Sect. 2.2, the patched MIFARE Classic cards, e.g., EasyCard 2.0, blocks all existing efficient card-only attacks by incorporating better RNG and replying every authentication error with encrypted NACK while maintaining compatibility with legacy readers. Since the plaintext of the 4-bit NACK code is fixed (`0x0` for EasyCard 2.0), replying encrypted NACK codes leaks four keystream bits per authentication failure. The data collected in each failed authentication attempt is called a *trace* and can be used to construct a set of four algebraic equations as the following.

Let $\mathbf{x} = (x_0, \ldots, x_{47})$ denote the initial state of the LFSR, i.e., the secret key. The new state of the LFSR after an input of n-bit sequence $\mathbf{i}$ can be written in a form like:

$$\mathbf{A_i}(\mathbf{x}) = \mathbf{L}^n \mathbf{x} + \mathbf{v_i}, \tag{2}$$

where $\mathbf{L}$ is a linear transformation that depends only on the LFSR's feedback function F, and $\mathbf{v_i}$ is a 48-bit vector that depends on the input $\mathbf{i}$ (and, of course, the LFSR's feedback function F). Here the important thing to note is that $\mathbf{v_i}$ *does not* depend on the secret key and hence can be computed based on the information available in a trace. Then the keystream bit generated by the nonlinear filter right after the input $\mathbf{i}$ can be obtained by

$$a_{\mathbf{i}} = f(\mathbf{A_i}(\mathbf{x})). \tag{3}$$

Both uid and n_T are transmitted in plaintext. It is then easy to express the LFSR state after the input of $\text{uid} \oplus n_T$ in terms of the unknowns $x_0, \ldots, x_{47}$ using Eq. (2). Although only the encrypted reader nonce is available (in fact, it is generated by the attacker in card-only attacks), it is still possible to *decrypt* $\{n_R\}$ using the keystream bits obtained by Eq. (3) and derive subsequent LFSR states and keystream bits in the form of polynomials of $x_0, \ldots, x_{47}$. By equating the 4 keystream bits to their corresponding polynomials, we get 4 *NACK equations* per trace of failed authentication session. It is then possible, at least in theory, to collect sufficiently many equations (≥ 12) and solve the resulted system using Gröbner-basis or SAT solvers.

In practice, however, the main difficulty of the algebraic attack described above lies in the last step, namely, solving the resulted polynomial system. In fact, the degree of such systems saturates due to the nonlinearity introduced by the recurrent decryption of $\{n_R\}$. In order to speed up the solving procedure, we need to extract more information from the traces using algebraic differential cryptanalytic techniques.

3.2   Differential Equations

From Eq. (2), we can see that the difference of two LFSR states that descend from a common initial state would be $\mathbf{A_i}(\mathbf{x}) \oplus \mathbf{A_j}(\mathbf{x}) = \mathbf{v_i} \oplus \mathbf{v_j}$ if $\mathbf{i}$ and $\mathbf{j}$ are two input bit streams of the same length. It means that we can *know* the LFSR state difference after two different tag nonces even though we *cannot control* them. This is easy to circumvent, however, as one can keep authenticating with a card at hand, hoping that the desired differences will eventually show up. For example, we are interested in those pairs that have only one bit difference in the LFSR state after an input of $\mathsf{uid} \oplus n_T$, especially when the different bit lies at the leftmost possible position. As will be clear later, the reason why we are interested in such pairs is because such pairs of difference are easy to "cancel," which gives more information in recovery of the secret key. More specifically, let $\mathbf{y}_{n_T}$, $\mathbf{y}_{n'_T}$ denote the LFSR states after inputting two different tag nonces n_T and n'_T, then our targets would be the pairs such that $\Delta\mathbf{y} = \mathbf{y}_{n_T} \oplus \mathbf{y}_{n'_T} = \mathsf{0x000080000000}$. Since n_T has only 32 bits, one bit difference at position 16 (cf. Fig. 1) is the furthest we can get. Thanks to the birthday paradox, it does not take too long to gather sufficiently many such pairs.

   Once such a pair is observed, we then try to "cancel" the state difference by properly manipulating the reader's nonce in the second trace. This could be done by carefully selecting and guessing $\{n'_R\}$ in the second trace according to $\{n_R\}$ in the first trace because the reader's nonces are transmitted in ciphertext and we, as the attacker, do not have the secret key to produce the correct keystream bits. More specifically, our goal in this stage is to keep pushing bits with zero difference into the LFSR. Since there is only one bit of difference at position 16 of the LFSR state at the beginning of this stage, we only need to keep our eyes on it. When it is shifted to a position that is not part of input to the nonlinear filter function f, the output keystream bits of these two traces should be the same. In this case, we can obtain the exact difference in the corresponding bits of $\{n_R\}$ by simply inspecting the feedback function of LFSR and then cancel it by altering the input accordingly. However, for positions 15, 13, 11, and 9 (cf. Fig. 1), we need to guess the output keystream bits of the nonlinear function f. If all four guesses are correct, we will arrive at a target pair of traces with identical LFSR states.

   Figure 3 demonstrates how the first three bits of $\{n'_R\}$ are decided or guessed by showing the differential view of the LFSR states of target pairs. Let $\mathbf{z}_k$, $\mathbf{z}'_k$ be the LFSR states of the pair after shifting in k bits (including n_T). For any target pair of traces, we have

$$\Delta\mathbf{z}_{32} = \mathbf{z}_{32} \oplus \mathbf{z}'_{32} = \mathbf{v}_{n_T} \oplus \mathbf{v}_{n'_T} = \mathsf{0x000080000000}. \tag{4}$$

In this state, both the outputs of the feedback function and the filter function are identical for these two traces, so the first bit of $\{n'_R\}$ should be the same as $\{n_R\}$ (cf. the zero values in Fig. 3a). When the bit difference is pushed to position 15, we can only choose the second bit of $\{n'_R\}$ by guessing since the difference of $f(\mathbf{z}_{33})$ and $f(\mathbf{z}'_{33})$ cannot be obtained purely from this differential
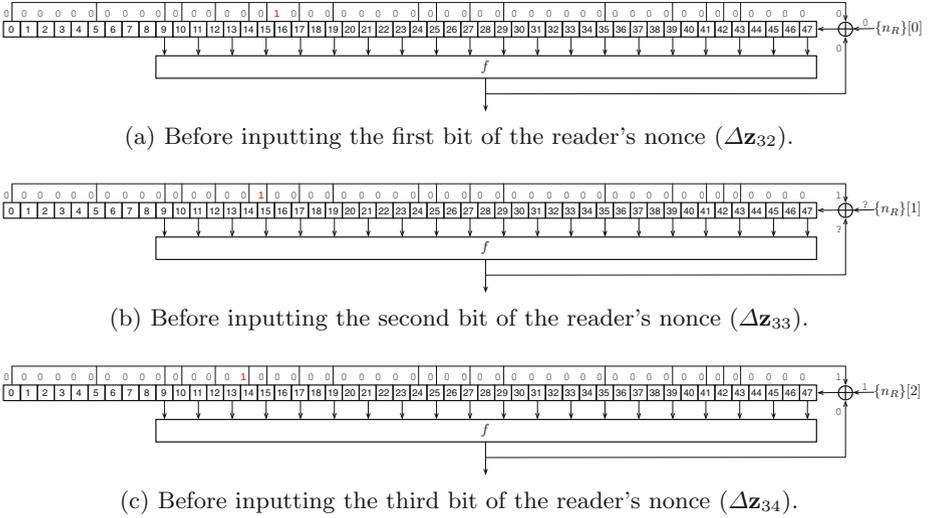
(a) Before inputting the first bit of the reader's nonce ($\Delta \mathbf{z}_{32}$).

(b) Before inputting the second bit of the reader's nonce ($\Delta \mathbf{z}_{33}$).

(c) Before inputting the third bit of the reader's nonce ($\Delta \mathbf{z}_{34}$).

Fig. 3. Differential view of Crypto-1's LFSR states

view (cf. Fig. 3b). For the third bit, the difference of $\{n_R\}$ and $\{n_R'\}$ should be 1 because of the feedback function (cf. Fig. 3c).

Following similar procedure, we push the bit difference out of the LFSR and hope to cancel the differences in the input bits. Such cancellation of input difference could be examined by checking whether the tag responds with an identical encrypted NACK code. In other words, the four keystream bits obtained after each authentication failure are used as an oracle for confirming whether our guesses in $\{n_R'\}$ successfully produce the desired differential or not. There are, of course, false positives from this oracle due to collision in practice, and we leave the discussion about this issue to Sect. 3.4.

If the guessed bits successfully cancel the differences, the following 4 *differential equations*, corresponding to the four guessed bits, should hold.

$$f(\mathbf{z}_k') \oplus f(\mathbf{z}_k) = f(\mathbf{z}_k \oplus \mathbf{e}_{48-k}) \oplus f(\mathbf{z}_k) = \frac{\partial f}{\partial z_{48-k}}(\mathbf{z}_k) = \delta_k, \qquad (5)$$

where $\mathbf{e}_k$ is the 48-bit vector with 1 in the k-th position and 0 elsewhere, and δ_k is the guessed difference, for $k = 33, 35, 37, 39$.

3.3 Filter Equations

In addition to Eq. (5), by taking a closer look at state $\mathbf{z}_{33}$, we devise the following formula as the *filter equation* to further reduce the search space of our attack.

$$\left(\frac{\partial f}{\partial z_{15}}(\mathbf{z}_{33}) \oplus \delta_{33} \right) \left(\frac{\partial^2 f}{\partial z_{15} \partial z_{47}}(\mathbf{z}_{33}) \oplus 1 \right) = 0. \qquad (6)$$

Any state assignment that does not satisfy Eq. (6) would result in $\frac{\partial f}{\partial z_{15}}(\mathbf{z}_{33}) \neq \delta_{33}$ and $\frac{\partial^2 f}{\partial z_{15} \partial z_{47}}(\mathbf{z}_{33}) = 0$ at the same time. This means, no matter what the value of the newly input bit (z_{47}) is, the output of f would not be equal to our guessed value, which contradicts with the fact that we have already reached the same LFSR state in both traces at the end of the authentication session. As a result, we can include Eq. (6) from each successful pair of traces to the final system of equations, and each such equation is expected to work as a filter that eliminates 1/4 of the solution space.

Empirically, there is a high degree of dependency among the equations acquired from different traces, but it does not take too long to collect sufficiently many pairs such that only a few candidate solutions can pass all filters. Based on our experience, these filters help tremendously in solving the nonlinear system.

3.4   Dealing with False Positives

Up until this point, we have assumed that our oracle can 100 % accurate in telling whether two internal states are the same or not. This does not hold in practice: as we can only observe four keystream bits, it is possible for two traces to have the same keystream bits yet different internal states. In our experiments, around 26 % of the cases where the four keystream bits agree are actually false positives. As a result, not all the collected differential relations, i.e., Eqs. (5) and (6), can be incorporated in the final system to solve. In this section, we will describe how we deal with this problem by more aggressive filtering.

We note that only 18 bits ($z_9, z_{11}, z_{13}, z_{17}, z_{19}, \ldots, z_{45}$) might have an effect on the evaluation of Eq. (6). Random assignments to these 18 bits should be in the solution space of the filter equation with probability $q = 3/4$, given that the filter function f is unbiased. Additionally, the correct assignment should be a solution to those filter equations collected from the true positive results. If we collect sufficiently many pairs and rank all 2^{18} possible assignments by their number of correct evaluations to the collected filter equations, the correct assignment should be very close to the top of the list with high probability. Therefore, the list serves as a good guide for guessing the 18 bits in the resultant system of equations. We can substitute the 18 variables with the bit assignments according to the list and try solving the system using SAT solvers. Note that we should eliminate the equations derived from the traces where Eq. (6) evaluates to `false` while trying each 18-bit assignment. According to our empirical results, it takes around 2 to 15 min for CryptoMiniSat to solve the system if the 18 bits are assigned with correct values.

The next question is how many pairs are sufficient to put the correct assignment at the top of the list with high probability. Assume that in total N such differential pairs are collected, among which $\tilde{N}$ are true positives. The number of filter equations that the correct assignment would evaluate to true, denoted by N_1, should have the following probability mass function.

$$Pr[N_1 = n] = \binom{N - \tilde{N}}{n - \tilde{N}} q^{n - \tilde{N}} (1 - q)^{N - n}, n = \tilde{N}, \ldots, N. \qquad (7)$$

Furthermore, if we denote the rank of the correct assignment in the list by ρ, then we have

$$Pr[\rho = k|N_1 = n] = \binom{M-1}{k-1}a(n)^{M-k}[1 - a(n)]^{k-1}, \qquad (8)$$

where $M = 2^{18}$ and $a(n) = \sum_{i=0}^{n-1}\binom{N}{i}q^i(1-q)^{N-i}$ is the probability that an incorrect assignment evaluates less than n filter equations to true. Using Eqs. (7) and (8), it is straightforward to compute the probability function of the rank of the correct assignment by

$$Pr[\rho = k] = \sum_{n=\tilde{N}}^{N} Pr[\rho = k|N_1 = n]Pr[N_1 = n]. \qquad (9)$$

We compute the percentiles of the rank of the correct bit assignments for various numbers of filter equations and summarize the most useful results in Table 1. This gives us an estimate of how many pairs would be sufficient to substantially reduce the expected number of trials we have to perform before finally solving the system. For example, given 150 filter equations collected, we are able to solve the system in less than 7 trials with a probability of 99 %. This is a very good result because in most cases, we just need to repeat the computation a few times before we can recover the key.

3.5   The Complete Attack

We summarize our attack procedure as follows.

1. Initiate (failing) authentication sessions with the target tag and record in a database each n_T received, $\mathbf{v}_{n_T}$, and four keystream bits $\mathbf{s}$ used to encrypt the returned NACK code.
2. For each $(n_T, \mathbf{v}_{n_T}, \mathbf{s})$ received, check whether $\mathbf{v}_{n_T}\oplus\texttt{0x000080000000}$, matches any $\mathbf{v}_{n_T'}$ already recorded. If so go to Step 3, having found a pair of n_T's that produce the state difference $\Delta\mathbf{y} = \texttt{0x000080000000}$. Otherwise, repeat Step 1.
3. Guess four δ_k's and manipulate $\{n_R\}$ accordingly. Check whether we see the same four keystream bits. If so, record the four differential relations (cf. Eq. (5)) thus found.

Table 1. The percentiles of ρ ($\tilde{N}/N = 74\%$)

Number of filter equations (N)	Percentile	Rank(ρ)
90	75	11
110	90	8
130	95	4
150	99	6

4. Repeat Steps 1–3 until we have collected enough differential relations (about 600–1000, or 150 to 250 successful attempts), then we use the method from Sect. 3.4 to remove the false positives.
5. Feed the differential equations, along with (i) some NACK equations, to a Gröbner-basis or SAT solver, and (ii) the 18-bit solution to the filter equations (cf. Sect. 3.4) as hint bits, to the solver and solve for the key. Empirically, we need about 1 NACK equation for every 4–5 differential equations.

4   Empirical Results and Discussion

We have tried several different solvers including the built-in Gröbner-basis solver in Maple, as well as PolyBoRi [4]. Empirically, CryptoMiniSat outperforms the other solvers by a large margin. Hence we only report the timings obtained using CryptoMiniSat for the rest of the paper. The results also show that the hint bits are extremely helpful to CryptoMiniSat, usually resulting in a tremendous speed-up.

We also note that the differential relations, as a system of equations, tend to be highly redundant and have multiple solutions. It is to avoid ending up with such a wrong solution, that in step 5 we must add a few equations on keystream bits in order to obtain a unique solution with high probability.

A submarine patch. We had bought a fair number of EasyCards on the streets of Taipei between 2009 and 2012 trying to track there were different editions of EasyCards. Surprisingly, we discovered that EasyCard 2.0 was actually not the first "patch" attempted by the EasyCard Corporation. There is actually another different kind of EasyCard, that we shall refer to as EasyCard 1.5, which has been surreptitiously in circulation since late 2010 or early 2011.

Although to all outward appearances EasyCard 1.5 is identical to EasyCard 1.0, it has a better RNG which makes n_T neither predictable nor controllable based on timing. This already defeats some (but not all) existing card-only attacks, even though EasyCard 1.5 performs otherwise identically to the original. For example, since the parities attack relies on the capability of controlling n_T, such an improved RNG already makes the attack time much longer if still possible at all. We represent this fact using a question mark in Table 3, in which we summarize the time required to carry out various attacks. It is perhaps surprising that the EasyCard Corporation managed to resist the temptation of announcing a security upgrade and kept this modification under wraps for so long. The differences among the three types of EasyCards are summarized in Table 2.

From Table 3, it is clear that our attack is the most practical one among the effective attacks against EasyCard 2.0 in the sense that our attack can be carried out by an ordinary person in private with an off-the-shelf reader and a PC.

In Table 3, the GPU result is taken from Chih *et al.* [5], while all other experiments are all carried out on a PC with 2.3 GHz AMD CPU. The data collection, on the other hand, is performed on a laptop PC with 2.0 GHz Intel

Table 2. Types of EasyCards attacked in our experiments

Card type	Parities checked	n_T generation
EasyCard 1.0	Yes	Predictable
EasyCard 1.5	Yes	Somewhat random
EasyCard 2.0	No (always `0x0`)	Random

Table 3. Timing comparison of all known attacks

Attack type	Online time	Compute time	1.0	1.5	2.0
Sniffing attack	2 s	< 2 s	√	√	√
GPU brute-force [5]	5 s	14 h	√	√	√
CPU brute-force	5 s	4 years	√	√	√
Parities attack	> 3 min	< 30 s	√	?	
Nested authentications	15–75 s	25–125 s	√	√	
Our attack	10–20 h	2–15 min			√

CPU. For CPU brute-force attack, we obviously have not run it to completion but extrapolate based on the timing result of a partial run instead. We use open-source software whenever possible, but we have also implemented and optimized some of the attacks.

5   Concluding Remarks

In this paper, we have demonstrated a highly practical attack against the Easy-Card 2.0, which is marketed as having patched the vulnerabilities of previous implementations of MIFARE Classic. By applying algebraic differential crypt-analysis techniques, our card-only attack can recover the secret key of EasyCard 2.0 within one day. This includes the time for online data collection and offline computation, both of which can be carried by a working platform that costs no more than a few hundreds of US dollars and is affordable even to the least wealthy attacker. This again shows the weakness of the Crypto-1 cipher, and highlights the unfortunate the fact that "security" protocols based on unsound ciphers, such as MIFARE Classic, is not suitable for important transactions such as electronic payment.

References

1. Albrecht, M., Cid, C.: Algebraic techniques in differential cryptanalysis. In: Dunkelman, O. (ed.) FSE 2009. LNCS, vol. 5665, pp. 193–208. Springer, Heidelberg (2009)

2. Albrecht, M., Cid, C., Dullien, T., Faugère, J.-C., Perret, L.: Algebraic precomputations in differential and integral cryptanalysis. In: Lai, X., Yung, M., Lin, D. (eds.) Inscrypt 2010. LNCS, vol. 6584, pp. 387–403. Springer, Heidelberg (2011)

3. Biham, E., Shamir, A.: Differential cryptanalysis of DES-like cryptosystems. In: Menezes, A., Vanstone, S.A. (eds.) CRYPTO 1990. LNCS, vol. 537, pp. 2–21. Springer, Heidelberg (1991)

4. Brickenstein, M., Dreyer, A.: PolyBoRi: a framework for Gröbner-basis computations with boolean polynomials. J. Symb. Comput. **44**(9), 1326–1345 (2009)

5. Chih, M.-Y., Shih, J.-R., Yang, B.-Y., Ding, J., Cheng, C.-M.: MIFARE Classic: Practical attacks and defenses. In: 19th Cryptology and Information Security Conference (CISC 2010), Hsinchu, Taiwan May 2010

6. Courtois, N., Klimov, A., Patarin, J., Shamir, A.: Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 392–407. Springer, Heidelberg (2000)

7. Courtois, N.T.: The dark side of security by obscurity and cloning MiFare Classic rail and building passes anywhere, anytime. Cryptology ePrint Archive, Report 2009/137 (2009). http://eprint.iacr.org/

8. Courtois, N.T., Nohl, K., O'Neil, S.: Algebraic attacks on the Crypto-1 stream cipher in MiFare Classic and Oyster cards. Cryptology ePrint Archive, Report 2008/166 (2008). http://eprint.iacr.org/

9. Davis, M., Logemann, G., Loveland, D.: A machine program for theorem-proving. Commun. ACM **5**(7), 394–397 (1962)

10. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases (F_4). J. Pure Appl. Algebra **139**(1), 61–88 (1999)

11. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F_5). In: Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation, ISSAC '02, pp. 75–83. ACM, New York (2002)

12. Garcia, F.D., de Koning Gans, G., Muijrers, R., van Rossum, P., Verdult, R., Schreur, R.W., Jacobs, B.: Dismantling MIFARE Classic. In: Jajodia, S., Lopez, J. (eds.) ESORICS 2008. LNCS, vol. 5283, pp. 97–114. Springer, Heidelberg (2008)

13. Garcia, F.D., van Rossum, P., Verdult, R., Schreur, R.W.: Wirelessly pickpocketing a Mifare Classic card. In Proceedings of the 2009 30th IEEE Symposium on Security and Privacy, S&P '09, pp. 3–15. IEEE Computer Society, Washington, DC (2009)

14. Golić, J.D.: Cryptanalytic attacks on MIFARE Classic protocol. In: Dawson, E. (ed.) CT-RSA 2013. LNCS, vol. 7779, pp. 239–258. Springer, Heidelberg (2013)

15. de Koning Gans, G., Hoepman, J.-H., Garcia, F.D.: A practical attack on the MIFARE Classic. In: Grimaud, G., Standaert, F.-X. (eds.) CARDIS 2008. LNCS, vol. 5189, pp. 267–282. Springer, Heidelberg (2008)

16. Nohl, K., Evans, D., Starbug, S., Plötz, H.: Reverse-engineering a cryptographic RFID tag. In: Proceedings of the 17th Conference on Security Symposium, Security '08, pp. 185–193. USENIX Association, Berkeley (2008)

17. Nohl, K., Plötz, H.: Mifare: Little security, despite obscurity. In: 24th Chaos Communication Congress, 24C3, Berlin, Germany, Dec 2007

18. Shannon, C.E.: Communication theory of secrecy systems. Bell Syst. Tech. J. **28**(4), 656–715 (1949)

19. Soos, M.: CryptoMiniSat. https://github.com/msoos/cryptominisat

20. Wang, M., Sun, Y., Mouha, N., Preneel, B.: Algebraic techniques in differential cryptanalysis revisited. In: Parampalli, U., Hawkes, P. (eds.) ACISP 2011. LNCS, vol. 6812, pp. 120–141. Springer, Heidelberg (2011)