The Fall of a Tiny Star

Flavio D. Garcia[1(✉)] and Bart Jacobs[2]

[1] School of Computer Science, University of Birmingham, Birmingham, UK
f.garcia@bham.ac.uk
[2] Institute for Computing and Information Sciences, Digital Security Group, Radboud University Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands
www.cs.bham.ac.uk/∼garciaf, www.cs.ru.nl/∼bart

Abstract. This short paper gives a combined technical-historical account of the fate of the world's most-used contactless smart card, the MIFARE Classic. The account concentrates on the years 2008 and 2009 when serious security flaws in the MIFARE Classic were unveiled. The story covers, besides the relevant technicalities, the risks of proprietary security mechanisms, the rights and morals wrt. publishing security vulnerabilities, and eventually the legal confrontation in court.

1 Introduction

Contactless smart cards (often called RFID tags) are tiny electronic devices that communicate wirelessly with a reader. The functionality of these tags ranges from simply sending a serial number to doing complex (public key) cryptographic operations in a fully programmable manner. These tags are used for identification mainly as replacement for barcodes, but they are also widely used in access control and transport ticketing systems. Many countries have incorporated RFID tags in their electronic passports and identity cards [15] and many office buildings and secured facilities such as airports and military bases use them for access control.

The MIFARE Classic was introduced in the market back in 1994 by Philips Semiconductors (later NXP) and quickly became the industry standard for access control in buildings and payment in public transport ticketing systems all over the world, such as the Oyster card in London and the OV-Chipkaart in the Netherlands, among others. According to the manufacturer, two billion MIFARE cards had been sold by 2008. The OV-Chipkaart was, back in 2007, in a test phase in the city of Rotterdam and, if successful, it would be extended nationwide. The Digital Security group at Nijmegen has been investigating software and protocols for smart cards since the late 1990s. Naturally, there was an interest at the moment a chip card was about to be introduced that should end up in the pockets of almost all Dutch citizens.

This is an inside story of the demise of the MIFARE Classic. This story involves a mix of technical and historical details. The authors have been directly involved in this story, one on the technical side (FG), and one on the more organisational (management) side (BJ). The story is thus told by insiders, which

P.Y.A. Ryan et al. (Eds.): Kahn Festschrift, LNCS 9100, pp. 69–87, 2016.
DOI: 10.1007/978-3-662-49301-4_5

has both advantages and disadvantages. The added value lies in having details that are unknown to outsiders. On the downside, the authors may not always have the most detached perspective on these developments.

Throughout this article, the pronoun 'we' refers to the MIFARE team[1] within the Digital Security Group from Nijmegen, and not specifically to the authors. Whenever it is inappropriate not to mention the role of the authors individually, initials (FG and BJ) will be used.

2 MIFARE Ultralight, Cardis and Before

Back in November 2004, the development of a device, dubbed 'Ghost', was initiated within the Digital Security research group at Nijmegen. The Ghost was planned as a programmable device, capable of emulating an RFID tag. Since the group lacked the necessary background on electronics, Peter Dolron, from the faculty's Technical Center, was asked for help. Developing and debugging hardware is a very tedious and time consuming task. By the end of 2006 the project really started taking off when Roel Verdult, then a student looking for a topic for his master's thesis [23] asked FG for supervision. Verdult invested the time and patience necessary to get things running and by mid-2007 there was a working prototype. In order to have a bold and appealing goal, Verdult was challenged by his supervisor to get unauthorized access to the parking system of the university building, which uses MIFARE cards. It was slightly shocking to see that the system did not really use the security mechanisms on the card but simply its serial number. Thus, the beam of the parking lot could be raised by waiving the Ghost, programmed to replay that serial number, in front of the reader.

In May 2007, another student, Gerhard de Koning Gans started to work on his master's thesis project [7] under the supervision of Jaap-Henk Hoepman and co-supervised by FG. The initial idea was for de Koning Gans to focus on the OV-chipkaart system using the Ghost tool developed by Verdult. As the development of the Ghost was slow and tedious de Koning Gans started looking for alternatives and ordered a Proxmark III. This device, much more advanced than the Ghost, had not only tag emulating capabilities but it was also capable of doing reader emulation. The drawback of the Proxmark was that it was programmed to communicate using Manchester encoding, which is a different way of communicating bits than the one used in the MIFARE Classic, called Miller encoding. Therefore,

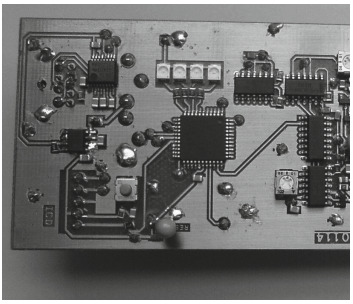

Fig. 1. The Ghost, a programmable RFID tag emulator developed at Nijmegen [22].

[1] Including: Flavio Garcia, Jaap-Henk Hoepman, Bart Jacobs, Ravindra Kali, Vinesh Kali, Gerhard de Koning Gans, Ruben Muijrers, Peter van Rossum, Wouter Teepe, Roel Verdult.

de Koning Gans had to program the Miller encoding on the Proxmark himself, which was also a tedious and time consuming task. Verdult and de Koning Gans started collaborating immediately, working as a team, using one device to debug the other.

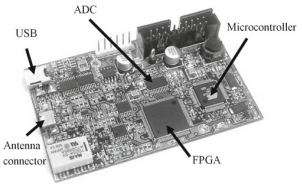

Fig. 2. The Proxmark III

The next challenge for Verdult was the payment system for public transport in NL, the OV-Chipkaart (while de Koning Gans continued working on the Proxmark). The system had basically two types of cards: disposable and multiple use; the latter can be further subdivided into personalised and anonymous ones. The disposable cards are mainly targeted on tourists and infrequent travelers while the re-usable cards are mainly used for subscriptions and frequent travelers. Verdult quickly found out that the disposable cards were MIFARE Ultralight. This kind of card does not support any cryptography and the only security mechanism it has is a write-once memory. This security mechanism is ineffective against an emulator device like the Ghost. Verdult quickly managed to mount a replay attack on the MIFARE Ultralight, in which the Ghost device acted as an un-used card, that ignored the command to change its status to 'used'. This could already grant free public transport, see the section below.

In the meantime de Koning Gans managed to get the Miller encoding working on the Proxmark device, making it possible to eavesdrop and impersonate both tag and reader messages. With this powerful tool he started to study the MIFARE Classic. After a few experiments he observed that some random numbers generated by the card repeated surprisingly often. This weakness, even without knowing the whole cipher, quickly led to an attack on the MIFARE Classic [8]. It allows an attacker to read and modify the contents of a card; see Technical 1 for more details.

Technical 1. Sketch of the first "keystream shifting" attack [8]

The MIFARE Classic documentation states that whenever a reader tries to read the secret key of a particular memory sector, the card returns a sequence of zeros. An attacker proceeds as follows:

1. Record a legitimate trace where the reader reads a sector on the card of which the key is known (for instance because it is a default key);
2. When the card repeats the nonce, replay the messages but change the sector number to the one of an unknown key. In this way the card answers with a sequence of zeros, XOR-ed with the keystream—enabling an adversary to get plain keystream;
3. Use this keystream to decrypt the recorded message. This keystream can also be used to read other sectors or modify the data on the card.

Even though this attack is serious and harmful from a security point of view, the authentication protocol of the MIFARE Classic was not broken. Thus it was

not possible, for instance, to get access to buildings using MIFARE Classic for access control, such as our own university building.

3 Dismantling MIFARE Classic

3.1 Cracks Appearing

Although the MIFARE Classic is old and widely used, the research community (both scientists and hackers) have been slow in taking it up as a target of investigation. The first independent public review of the chip, as far as we know, was announced at the yearly meeting of the German Chaos Computer Club (CCC)[2] in Berlin, late December 2007. Karsten Nohl and Henryk Plötz, at the time affiliated with Virginia University and Humboldt University Berlin, respectively, presented their analysis of the card [19]. It was hardware-based and involved peeling of, layer-by-layer, the protective shielding of the chip, until the chip layout was visible. They thus derived the schematics of the chip and were able to reconstruct part of the algorithms involved. At the CCC meeting they did not present all of their findings, for fear of legal action. Hence it is hard to assess what they precisely knew at that stage. But for sure, they were aware of the structure of the generating polynomial of the LFSR used in the cipher and the weakness in the pseudo-random generator on the tag. They also claimed to have knowledge of the filter function but for some reason they decided not to make it public.

Early January 2008 the media in NL reported on this CCC presentation and pointed to its relevance for the national OV-chipkaart project. The original plan with Verdult was to postpone publicity until after finishing the master thesis. The CCC presentation led to a reconsideration of this intention. When RLT journalist Koen de Regt contacted Nijmegen with some questions, he was informed about the local research results. The journalist immediately saw the high relevance and publicity value of the topic. He made an appointment with Verdult to meet the next week for a on-site recording in Rotterdam, the only place where the OV-chipkaart was operational, at the time.

That weekend BJ had to leave for a workshop in Spain. He discussed the media strategy with Verdult (stick to your expertise, make a clear point, and don't let journalists seduce you to make far-reaching political statements). BJ also asked Wouter Teepe, who had some previous journalistic experience, to be available as back-up and to inform the company Trans link Systems (TLS), running the OV-chipkaart project, on the day of broadcast. FG, as a non-Dutch speaker, had a minor role in these matters.

On Monday 14 Jan, 2008 RTL news opened its evening edition with a long item showing that the OV-chipkaart has been broken[3]. It involved an interview

[2] The CCC is a large, influential association of computer enthusiasts, hackers and digital rights activists in Germany.

[3] This was a premature statement, since only the throw-away version was broken at that time.

with Verdult, and a demonstration where he walks many circles, entering and exiting entrance gates of the Rotterdam metro using his ghost device to emulate a MIFARE Ultralight, see Fig. 3. The imagery is powerful. It is played in the back while the (poor) spokeswoman of TLS explains that nothing is wrong with the OV-chipkaart. A publicity disaster for TLS begins to enroll. A media wave results (handled jointly by Teepe and Verdult), setting a political reaction in motion: on Wednesday there is hearing by the relevant Parliamentary subcommittee (involving besides Teepe and Verdult also Amsterdam colleague Melanie Rieback and long time hacker Rop Gonggrijp) and on Thursday there is a meeting with the responsible junior minister Tineke Huizinga. The message is that TLS should have used open, publicly scrutinised algorithms, with an undertone of frustration about the privacy-unfriendliness of the system and a hint of triumphalism. The distinction between the MIFARE Ultralight (Verdult's target, which has no cryptographic protection and is used only for day-cards) and the MIFARE Classic (for regular, multiple use cards) is not always clearly made. Formally, the minister is powerless in this matter, because the OV-chipkaart is operated by private companies in the transport sector (united in TLS). However, questions are being asked in Parliament, which she has to answer. Although public transport has been largely privatised in NL, the fact that many people depend on it and have no real alternative explains why a substantial level of government regulation and steering is expected.

In the weeks ahead research at Nijmegen continues to further understand the cryptographic protection of the MIFARE Classic (see below). At the same time contacts are established with the relevant security people at NXP (represented by Hans de Jong). There are also contacts at management level with the Transport Ministry, which is mainly trying to understand the technicalities and the political impact. Jeroen Kok, the chairman of TLS visits Nijmegen, with a shocked, but open mind, trying to understand "what makes these guys tick".

Fig. 3. Screenshot from RTL news, 14 Jan. 2008

3.2 Before the Breakthrough

Mid February 2008 there is sudden excitement among the MIFARE researchers. Verdult has found a non-trivial bug: he managed to make a commercially available, official MIFARE reader believe that it was talking to a MIFARE Card while in fact it was just talking to the Ghost/Proxmark. At this stage more staff members get involved, notably Peter van Rossum.

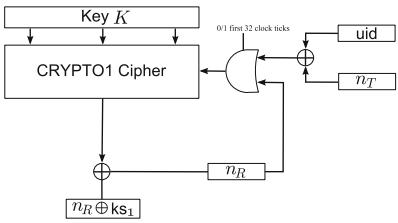

Fig. 4. Initialization diagram.

Technical 2. MIFARE Classic authentication [9]

The authentication protocol between a MIFARE Classic tag and reader is depicted in the figure below. First the tag sends its unique identifier *uid* and a nonce n_T. From this point on all communication in encrypted with the shared key k—which is either the same on many cards, or derived by the reader from the *uid* via a master key (key diversification). Next the reader answers with a nonce n_R of its own choosing and a value a_R, which is a function of n_T. To conclude, the tag answers with a_T which is a function of n_R. At this point both tag and reader believe that they have authenticated each other.

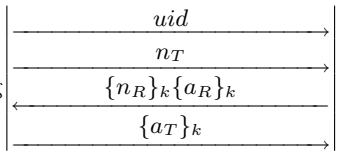

Technical 3. Description of the Random Number Generator (RNG) weakness

After power up (older) MIFARE readers will produce always the same sequence of nonces $n_R^1, n_R^2, n_R^3 \ldots$ (in successive runs of the authentication protocol, see Technical 2). This deterministic character of the RNG was not immediately recognised; at first repeated nonces made us think the RNG was weak. Such repetitions can be exploited: whenever the reader repeats its nonce n_R, an attacker playing the role of a tag is able to replay a previously recorded $\{a_T\}_k$ and, by doing so, impersonate this legitimate tag.
The RNG on the card uses an LFSR of only 16 bits, so once a list of all successive values is compiled, after observing one nonce subsequent nonces can simply be obtained by look-up. Technical 9 describes how this can be exploited.

Upon learning about the discovery the group becomes even more aware of the explosive character of the research. Actually breaking the cryptographic protection of the MIFARE Classic comes in sight. This would be a major blow for the OV-chipkaart. But more importantly, it would present a acute problem for all the organisations that use the MIFARE Classic card for controlling access to their facilities. These include military bases, banks, ministries, many companies, not only in NL but worldwide. In comparison, the OV-chipkaart is 'peanuts'.

A group meeting is planned where these sensitivities are discussed explicitly and an internal mode of operation is adopted in order to prevent accidental leakage of sensitive information (or software). It is decided that all the research takes place in one office and is done jointly by students and staff. All internal communication (and stored information) is encrypted, via PGP. Further, external contacts will be coordinated with BJ. An unintended side-effect of the concentration of efforts in a single office is a research boost. The level of excitement is high; the team smells blood.

On March 3 the Crypto 1 cipher is reconstructed, see Technical 4, and on March 7 there is a working attack, see Technical 5.

Technical 4. Reverse engineering MIFARE Classic [9]

While trying to reproduce the replay attack described in Technical 3 at the entrance of the faculty building, repeated nonces from the reader did not appear. Soon it was realised that the sequence of nonces generated by the readers repeated after power up, with each authentication attempt, but the pseudo-random generator on the reader had a full cycle. This meant that the readers had to be powered down in order to be able to carry out a replay attack. This was impractical and gave a moment of frustration within the team. There was only one option left; to fully reverse-engineer the whole cipher. It was suspected that the Crypto-1 cipher would be similar to the one in the Hitag2 tags, another RFID tag from NXP. The cipher in this tag had a software implementation and this had been reverse-engineered and released on the Internet. This cipher consists of an LFSR and a boolean filter function. FG asked Verdult to first initiate the cipher with a random state and record the first bit of the produced keystream, and then do this again with the same state but with one bit flipped. Whenever a different keystream bit appeared it could be deduced that the flipped bit is an input to the boolean function. Once the input bits to the boolean function were known, van Rossum proposed to use a similar procedure to build a boolean table in order to recover the whole filter function. On March 3, 2008 we had a software implementation of the whole cipher and authentication protocol that was fully consistent with the behavior of the MIFARE Classic. This was a moment of excitement among the team, seeing the secret that has been zealously kept for more than 15 years.

3.3 A Hectic Week, Early March 2008

When the MIFARE Classic is first cloned on Friday afternoon March 7 a pre-conceived plan is set in motion. BJ calls the chairman of the University, Roelof de Wijkerslooth, and says: "we have an emergency situation; I'm pushing a red button; please come over and have a look". Ten minutes later de Wijkerslooth arrives and sees a secured door being opened with a cloned card. He hears about (and agrees in principle to) the rest of the plan: (1) informing the national government about the card vulnerabilities, notably wrt. access control, (2) informing the card producer NXP, (3) giving a public warning to card users, and (4) publishing the results in the scientific literature (after a delay of several months). De Wijkerslooth is a former senior civil servant and decides to inform the national authorities himself, at cabinet level. The message is understood there, and a threat assessment is initiated. The task of verifying the results is given to the NLNCSA[4], a part of the national intelligence service, informally known as the government's crypto group. A manager of the NLNCSA calls BJ at home later that evening to make an appointment, possibly even the same night. A meeting is planned on Saturday afternoon at the university in Nijmegen, involving Roel Verdult, Wouter Teepe, BJ and two crypto experts from NLNCSA (Marcel and Gido). These visitors are keen to hear the results, showing not only professional interest, but also some amazement ("so it's really this bad!"). They are satisfied

[4] NLNCSA is an abbreviation of The Netherlands National Communications Security Agency, in Dutch also known as *Nationaal Bureau Verbindingsbeveiliging* (NBV); it is comparable to the British CESG, part of GCHQ.

Technical 5. First key recovery attack [9]

After having reverse engineered the cipher, the first key recovery attack against a MIFARE reader followed almost immediately. Verdult had the idea of splitting the 48 bit search space in an online and offline search. Pretending to be a tag, an attacker sends several authentication attempts to the target reader. On each attempt the attacker selects a special nonce. The idea is that one of these nonces will produce a specific pattern in the internal state of the cipher (e.g., a sequence of 12 zeros which would take 2^{12} authentication attempts). Then, offline, the attacker builds a table of all possible internal states with this pattern (e.g., of length $2^{48-12} = 2^{36}$) together with the keystream they produce. When you get a match on the keystream, you can simply lookup in the table the internal state of the cipher. Since the secret key is the initial state of the cipher, all we have to do then is to run the cipher backwards to recover the key.

Technical 6. Second key recovery attack [9]

Soon after the first key recovery attack, Ronny Wichers Schreur noticed that the filter function only uses the odd numbered bits of the LFSR as input to the filter function. This is a serious design flaw. It means that the 48 bit internal state of the cipher can be seen as two small ciphers of 24 bits each. One of these small ciphers producing the even numbered bits of keystream (called even cipher) and the other one the odd numbered bits (called odd cipher). These two (small) ciphers can be run independently. Since there are only 2^{24} possible small cipher states, it is feasible (and very fast) to try them all and discard those that do not match the corresponding (even/odd) keystream bits. This drastically reduces the amount of candidate states for both small ciphers. Next, one can combine these small cipher states (one even with one odd) in order to reconstruct the original 48 bit internal state of the cipher. In fact, given 64 bits of keystream that an attacker can obtain from a single authentication attempt, there will be only one candidate state for the even cipher that can be combined with another candidate state for the odd cipher to form a valid 48 bit internal state.

to learn that the technical details will not be published immediately, but only after some delay. On the way back they inform their superiors, who report to the interior minister and the prime minister. That weekend the country goes to a higher level of alert, and the access procedures for sensitive facilities are strengthened immediately. Also, friendly agencies are notified internationally.

On Sunday NXP is informed, via Hans de Jong, who is invited to Nijmegen to see the results for himself. On Monday morning he listens to what the NLNCSA people heard two days earlier. de Jong is understandably more defensive, immediately trying to delineate NXP's responsibility and accountability; he urges to keep things secret as long as possible. Clearly, he is not amused, also not because he is the second in line to be informed. Later that day Hans de Jong has a meeting at TLS, where he reports on the recent developments. TLS wishes to assess the impact for their systems, which happens in the course of the week.

Since a parliamentary debate about the OV-Chipkaart was already planned later in the week, the government, being in the know, could not hide what had happened. The interior minister Guusje ter Horst decides to inform Parliament

via a letter on Wednesday, March 12. The content is coordinated with Nijmegen, via the NLNCSA. It is decided that Nijmegen will go public after release of this official letter, with its own press conference, press statement [27], and YouTube video. In advance NLNCSA and NXP get to see a draft version of Nijmegen's press statement. NLNCSA is comfortable with the text, but NXP complains (without effect) that it gives away too many technical details and helps malicious hackers.

The letter to Parliament and the press conference at Nijmegen (and subsequent demonstration) on Wednesday lead to broad media coverage. The press statement, also available in English, helps journalists to get the story right.

On Friday March 14 a high level meeting takes place between NXP and the university, involving among others Fred Rausch (director NXP NL), Hans de Jong, de Wijkerslooth and BJ. Rausch brings a large bottle of wine and congratulates the researchers with their results. He tells that NXP wants to cooperate closely with the research team in order to improve its products and its advise to customers. He insists that such cooperation with universities is normally done under NDA (non-disclosure agreement). BJ refuses to sign any NDA, because he does not wish to restrict his academic freedom, and also because he senses political motives: such NDA could be used to prevent him from talking to the media (or to others, such as members of Parliament). Additionally, the university does not simply wish to give away its carefully built knowledge position for free. The matter is not resolved at this meeting.

3.4 Implications for the OV-Chipkaart

After the CCC presentation of December 2007 on the MIFARE Classic (see Subsect. 3.1) the company TLS that operates the OV-Chipkaart asked the research institute TNO to assess the situation. End of February 2008 TNO delivers its report, of which only the conclusions are published [2]. TNO writes that card manipulation requires advanced equipment. It sees no criminal business case in public transport ticketing fraud, and advises to replace the cards within the next two years. TNO turned out to be right on these last two points, but not on the advanced equipment. The report is criticised, also by Nijmegen, but with hindsight the criticism is too harsh. The transport ministry asks the Smart Card Center of Royal Holloway University London (RHUL) to investigate the matter. RHUL reports [1] mid April, after the breaking of the MIFARE Classic. It is more critical: with a nationwide system fraud is more likely than with a regional system (like in London); card replacement should be started immediately, using open designs and independent reviews.

The parliamentary committee for transport is closely following the matter and organising several hearings. The junior minister for transport, Tineke Huizinga, is often critised in Parliament over her way of handling the issue. This even leads to a no-confidence motion; it is rejected, but it does damage her political position and reputation. In the end she forces TLS to develop a migration plan (towards a successor card) that needs to be approved by RHUL. The ministry also pushes the use of open cryptographic designs and communication

standards. It eventually leads to the foundation http://openticketing.eu and to a closer collaboration with academia.

In the political debate on the introduction of the OV-Chipkaart containing a broken chip an often-used argument is: London's Oyster card works well with the same chip, so why would it not work in NL? This kind of reasoning motivated in particular the students involved in the MIFARE team to show that also the Oyster card could be manipulated. After extensive deliberations, it was decided that it was worth taking the risk, and so mid-April Roel Verdult, Gerhard de Koning Gans and Ruben Muijrers departed to London. The first day of their visit was spent traveling around London looking for a quiet station, in order to try their attacks. In the end they found a small Docklands Light Railway (DLR) with card readers not covered by security cameras. They used the Proxmark device to obtain traces of the communication between card and reader [11,24], from which the keys of (more than 20) memory sectors could be obtained. With these keys the contents of the sectors could be changed at will. After checking in with a card, a decrease of balance could be seen. They restored the balance and used this manipulated card for another trip without any problems. They thus made their point. They video-taped their actions, but the clip has never been released publicly.

(Going back, on their way out of the city they saw a special box in which tourists could deposit their used Oyster cards, thus donating the left-over value on the card to charity. The three students were tempted to top-up an Oyster card to £100.000 and drop it in the box. However, they had to catch a flight and had too little time for such a "charity prank".)

3.5 Litigation and Publication

In the course of March 2008 the research team prepares a scientific publication, called Dismantling MIFARE Classic, on the MIFARE algorithms and their vulnerabilities. Early April the paper is submitted to the *European Symposium on Research in Computer Security* (ESORICS'08), a respectable security conference series, to be held in October 2008 in Malaga, Spain. The chairs of the program committee, Sushil Jajodia and Javier Lopez, were informed about the sensitivity of the submission and asked to make sure nothing would leak out during the refereeing process.

As an aside, there were some sensitive authorship issues. The first, submitted version of the dismantling-mifare paper had six authors, namely: Garcia, de Koning Gans, Muijrers, van Rossum, Verdult, and Wichers Schreur. These are the people that did the actual scientific work of analysing the MIFARE protocol and encryption. Teepe and Jacobs were not listed as authors, because their contribution was non-scientific, involving external (media) contacts, negotiations within the university, hearings etc. After the paper got accepted and the relations with NXP deteriorated (see below), the chairman of the university insisted that BJ, as research leader of the group, be added as author; in the published version [9] he appears last in the list of authors. In a follow-up paper [13] he is not an author. Jaap-Henk Hoepman occurs as author of the very first paper [8].

He is affiliated to both Nijmegen University and the research institute TNO, putting him right in the middle of controversies. Because of the delicacy of the matter, he was excluded from MIFARE work on both sides, at Nijmegen and at TNO. Sadly for him, this meant that his early work on MIFARE could not be continued.

In conversation and in writing NXP expresses its strong objection against the intended publication after half a year. NXP argues for publication in 2010, after a delay of about two years. NXP makes clear that it will hold the university and its researchers responsible for any damages resulting from publication. In the course of March 2008 the university assembles a legal team, consisting of the rector Bas Kortman (a legal scholar himself), the university's own internal lawyer (Berthe Maat) and its external lawyer Dirkzwager, represented by Jaap Kronenberg and Mark Jansen. BJ has regular meetings with this team, plus de Wijkerslooth to discuss the case. Academic freedom was at stake, but possibly also the very existence of the university, once substantial claims were made. It was non-trivial for the lawyers to grasp the technical issues in sufficient detail and to appreciate the computer security tradition of publishing vulnerabilities as a contribution to security itself.

Mid-June the notification of acceptance of the ESORICS submission is received. The team is of course very happy with the scientific recognition (the referee reports are all very positive), but soon realises that the university leadership could still try and stop the publication. This possibility gives rise to strong emotions because it is felt as unjustified obstruction of highly relevant research. Some members of the team express (internally) that they will leave the university in Nijmegen in case publication is forbidden. In a meeting with the university's rector and chairman it is decided that a copy of the paper will be sent to NXP and to NLNCSA. Also, the "point of no return" is clearly communicated, namely the date when the final version of the paper has to be sent to ESORICS, for inclusion in the (printed) conference proceedings. The date was July 7 at first, but later postponed to July 14 (by ESORICS), and then again to July 18 (to await the outcome of the court case, see below). This transparency gives both NXP and the national authorities time to assess the publication and its possible impact, and the opportunity to react in time. In the weekend of 21–22 June BJ travels to Japan, for a three-week research visit in Kyoto that had been planned quite some time earlier.

On June 25 the NXP director Fred Rausch sends a letter to BJ personally in response to the article that is due to be published. The tone is formal and threatening. He writes (in English) that publication violates NXP's intellectual property rights. Rausch further writes:

> "Publishing the secret information (or substantial parts thereof) will most likely cause substantial damage also to NXP, for which damage NXP will hold all those responsible for the publication liable. Also, the publication is deemed to be irresponsible, as it will jeopardize the security of assets protected with the systems incorporating the MIFARE IC. Furthermore, this might induce others to commit criminal acts (to which

> the party publishing the material could be aiding and abetting). Needless to say that – in addition – third parties using systems incorporating the MIFARE Classic IC will have their own claims under tort vis-à-vis those responsible for the publication (also for the damages that they would suffer). NXP therefore kindly requests, and in as far as necessary hereby demands, that you withdraw the publication from the conference and that you do not publish it in any other way or distribute it to others."

A copy of the letter is sent to the Wijkerslooth and to the ESORICS program chairs. A reply is expected before June 30. De Wijkerslooth summons BJ to return from Japan, since the communication lines are too long.

In the meantime it becomes clear via informal channels that the national authorities (read: NLNCSA) do not object publication by October. NXP complains to the interior ministry about the intention to publish—and about spending their tax money on destroying their own products. The minister, under whose responsibility NLNCSA operates, is not impressed. The education minister, Ronald Plasterk, is a former scientist himself and defends "his scientists" and their academic freedom. This is, understandably, important for the university's leadership. Several legal scholars are consulted, notably about the risks of claims, both in NL and abroad. Then, the rector and the chairman decide to refuse to give in to NXP's demands to withdraw the article. They do offer NXP mediation as an instrument to resolve the dispute. NXP turns it down and decides to start legal action in order to get a publication ban (via an injunction). NXP not only takes the university to court, but also BJ personally: a clear case of legal intimidation.

A court meeting (called *Kort Geding* in Dutch) takes place on July 10, 2008, at Arnhem, presided by Judge Boonekamp. At NXP's request, the meeting takes place behind closed doors. On the university's side Kortman, de Wijkerslooth and BJ are present, represented by Dirkzwager, and on NXP's side Fred Rausch, represented by De Brauw, Blackstone en Westbroek. NXP pleads that publication violates its intellectual property rights and is irresponsible because of the resulting risks and damage. The university refers to article 10 of the European Convention on Human Rights (ECHR) on freedom of expression, and argues that banning publication is not socially beneficiary since it would protect companies selling faulty products, and since it leaves people with a false, unjustified sense of security. Part of the discussion focuses on whether the mathematically phrased article is an actual guide for (malicious) hackers.

The verdict comes on July 18. The Judge turns down NXP's request for a publication ban. He states that the university acted with due care, and that damage, if any, is not the result of publication, but of apparent deficiencies in the cards. NXP decides not to appeal. The same day, the paper is sent off, to be printed in the ESORICS proceedings (due to appear publicly in Oct. 2008).

In the evening of the day of the verdict Hans de Jong from NXP calls BJ privately to congratulate him with the outcome. He says that NXP is of course unhappy, but he expresses his hope to be able to cooperate on a technical level. This is indeed what happens. For instance, later that year Verdult finds another

MIFARE issue that could cause problems in NXP's successor card MIFARE Plus (when used in backward compatibility mode) and warns NXP in time take measures. With hindsight it is our own interpretation that NXP went to court mainly in order to strengthen its own position in case its customers would start suing NXP. NXP can now say: "Hey, we did everything we could to try and stop these guys". Still, it is unprecedented in Dutch legal and academic tradition that a company takes a university to court over an unwelcome publication.

Also looking back, it seemed easier to convince the judge than to convince the university board. However the rector and chairman had quite different responsibilities, covering the entire academic community at Nijmegen. They were genuinely concerned that substantial damage claims (hundreds of millions) could lead to closure of the university itself. In the end, after careful deliberations, they took the courageous decision to support their scientists and to stand up to defend academic freedom. It helped enormously that the rector, Kortman, is a practising legal scholar himself who is used to deal with legal arguments and pressure.

4 Card-Only Attack

Immediately after the ESORICS publication in Oct. 2008, people (sometimes from obscure origins) started asking if we were able to read the contents of their MIFARE Classic cards, without having access to a legitimate reader. The answer was no. For carrying out the attacks described in the ESORICS paper communication with a legitimate reader had to be intercepted. Several system integrators used this fact to argue that the reported attacks (see Technicals 5 and 6) were not practical because they require first communication with a reader to get the secret keys and then communication with a card, in order to be able to read its contents. Even though this argument had little grounds—from a cryptographic perspective the MIFARE Classic was completely broken—it was decided to work on another attack that could be performed having access to just a card. This was challenging, since the reader authenticates first to the card, before the card sends any ciphertext. For this, the team used a combination of four weakness discovered during the reverse engineering process. For more details see Technical 7. These weaknesses allow an attacker to recover a secret key from the card by just communicating with it for less than half a minute. For more details see Technical 8.

Even though waiting for half a minute in order to retrieve a secret key is acceptable, a MIFARE Classic 1K has 64 secret sector keys, which makes it impractical for an attacker to wirelessly pickpocket a card without being noticed. In order to speed up the process it is possible to use another two weaknesses in MIFARE Classic (see Technical 9). Once an attacker has recovered one secret key, either by using the previously described card-only attack or because the tag has a default key in some sector, she can perform a very fast re-authentication attack (see Technical 10). This attack recovers, within seconds, all remaining keys from the card.

Technical 7. Description of the weaknesses used in the card-only attacks [13]

weakness 1 While communicating with a reader, a MIFARE Classic card sends one parity bit after each byte of data in order to detect communication errors. These parity bits however, are computed over the plaintext instead of over the ciphertext.

weakness 2 Additionally, the bit of keystream used to encrypt the parity bit is reused to encrypt the next bit of plaintext, see figure below. This is a serious weakness that leaks one bit of information per byte of data sent over the air.

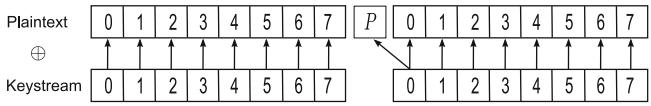

weakness 3 When the card receives a message, during the authentication protocol, it first checks whether the parity bits are correct or not before answering to the reader. If the parity bits are incorrect, the tag does not respond at all. When the parity bits are correct though, it answers either 'authentication failure' or it proceeds with the authentication protocol if the reader has authenticated successfully.

weakness 4 The error code for 'authentication failure' is sent encrypted by the card, even though in this case it cannot be assumed that the reader is able to decrypt. This leaks 4 extra bits of keystream.

5 Did the World Collapse?

It is rather uncomfortable that this embarrassingly badly designed MIFARE Classic could become the world's most-used contactless smart card. What went wrong? We don't pretend to have a definitive answer, but we do point to a number of factors (see also [17]).

1. Lack of evaluation. The MIFARE Classic has never gone through an evaluation procedure like Common Criteria. This was not normally done for smart cards in the early nineties, like it is today.
2. Use of proprietary technology. Since the design of the MIFARE Classic has been kept secret, independent expert review never happened. Nowadays cryptographic primitives like AES are established via open competition.

   The MIFARE Classic chip was designed in the early 1990s, when computing resources on a microchip were still scarce. It has been argued that the designers were aware of the limitations and thought at the time that "security by obscurity" would give them an additional layer of protection. One can also argue that this obscurity layer was quite counter-productive because it covered up mistakes and delayed a realistic view on the existing protection level.
3. Lack of re-evaluation of existing products. The MIFARE Classic was a commercially successful product, first for Philips and then for NXP. There was no incentive for the producer to look critically at what was being sold.

Technical 8. Card-only attack [13]

The team has proposed a number of card-only attacks. For the simplest of them, an attacker proceeds as follows. First the card starts communication and sends its challenge nonce n_T as indicated in Technical 2. Then, pretending to be a reader, the attacker sends a constant bitstring (e.g., all zeros) to the card as answer to the challenge of the tag. In most cases the tag will not answer at all, since the parity bits will not be correct. On average one out of $124 = 2^{8-1}$ attempts will have correct parity bits and then the card will send an encrypted 'authentication failure' message. The attacker keeps on doing this until the encrypted error message is also equal to a constant (e.g., all zeros). Before starting the attack, the attacker has pre-computed a table with all cipher states that have this property, i.e., $\{n_R\}_k = \{a_R\}_k = 0$ then the encrypted parity bits and the four bit encrypted error message are also zeros. This table contains approximately $2^{48}/2^{12} = 2^{36}$ elements. When the attacker receives the desired answer from the tag, she knows that the internal state of the cipher after sending n_T is one of the states in the pre-computed table. Then she can test these states with another authentication trace and in this way recover the secret key.

Technical 9. Weaknesses used in the re-authentication attack [13]

weakness 1 Once the reader has successfully authenticated for one sector and then it request to authenticate for another sector, the tag nonce is sent encrypted with the key corresponding to the new sector. This deviates from the authentication protocol described in Technical 2.

weakness 2 The pseudo-random number generator in the tag iterates over time and it has a cycle of size $2^{16} = 65536$. This means that, by precise timing, it is possible to predict what the next tag nonce will be.

(We even believe that in March 2008, when the security flaws became known, there was hardly anyone left within NXP who knew the MIFARE internals; the company's cryptographers had to go back to their libraries to find the old manuals.)

4. Vulnerabilities are valuable, as long as they are secret. The security weaknesses in the MIFARE Classic first became publicly known in 2008 via academic work. We are the first to publish them, but we are not so sure we are also the first who became aware of these vulnerabilities: intelligence organisations, illegal hardware cloners, or even large criminal organisations may have been well aware of the weaknesses, of course without revealing them, but possibly using them for their own benefit. In a similar manner so-called zero day exploits are valuable today; apparently the Stuxnet worm contained four of them in Windows.

Back in 2008 there seemed to be agreement that replacing the MIFARE Classic would make the world more secure. But there was much less agreement whether publication of the workings of the chip would also make the world more secure. NXP argued that it would not.

Once we were aware of the vulnerabilities we followed an approach that is often called "responsible disclosure": Notify the public about the vulnerabilities,

Technical 10. Re-authentication attack [13]

Assume that an attacker knows a secret key k_a of a MIFARE Classic tag, then she proceeds as follows. First, she repeatedly authenticates using k_a and measures the time between two consecutive authentications. Then she sends an authentication request for a target key k_t to the tag. The tag answers with a nonce n_T encrypted. Since the attacker is able to see the previous nonce, by taking into account the time between two consecutive authentications she can guess what the new encrypted nonce is. Then she can use this nonce to retrieve 32 bits of keystream that she can use to perform the attack described in Technical 6.

give the producer access to the details, and publish the details after a delay. We chose a delay of six months. For software vulnerabilities a much shorter delay is common, because the patch cycle for software is much shorter (*e.g.* one month for Microsoft products). It is impossible to replace all MIFARE Classic cards within six months. But six months is enough to do a security review, and introduce additional security measures, if needed.

Currently, at the time of writing (early 2014), more than six years have past since the emergence of security vulnerabilities in MIFARE Classic chip cards. Most of the public attention has focused on the use of these cards in e-ticketing. Migration plans have been developed in public transport (*e.g.* in NL and the UK), new cards often based on AES encryption have been widely adopted which in our opinion is a step in the good direction. Although, most systems are still phasing out existing MIFARE Classic cards and therefore still vulnerable. Manipulated MIFARE Classic cards are detected and blocked (roughly a few dozen per day), but fraud levels are much lower than with the old, paper-based system without entry/exit gates. Unfortunately, the Dutch public transport system has opted to migrate to a new MIFARE Classic chip (instead of a AES capable chip) which has a better pseudo-random generator (but uses the same weak cipher). The new chip prevents the nested-authentication attack described in Technical 10, but it remains vulnerable to all the other (slower) attacks.

In the access control sector the necessary migrations to successor cards are cumbersome, but possibly a bit less so than in e-ticketing. Card migrations have happened, but mostly for the more sensitive facilities. Sometimes, before this migration, additional entry checks have been implemented (like at ministries). It seems fair to say that despite all these security vulnerabilities the world has not collapsed.

It also seems fair to say that this MIFARE fiasco ranks among the bigger security failures (together with, for instance, DVD protection, and GSM encryption). Companies and governments have become more acutely aware of the importance of getting the details right in computer security, and of not just relying on someone else saying: “trust us, it’s OK”. They have also become more aware of the role played by independent investigators, doing their own reviews. Hopefully, it is also realized that trying to suppress such reviews via legal means is not an easy route (see also [5]). **Do we still want to keep this statement? Depending on the timing,**

should the VW case be resolved, we could add something about Megamos here

6 Related Work

At the end of July 2008, Nohl, Evans, Starbug and Plötz published their results on how they reverse engineered the MIFARE Classic at USENIX Security [18]. They describe how they sliced the MIFARE Classic chip and recognized some crypto related functions. They also mention that it is possible to recover a secret key from a tag by building a large rainbow table. In their paper, the filter function is kept secret.

When the ESORICS paper got published in October 2008, the full details of the CRYPTO1 cipher became public. This gave rise to some more research in this area (apart from our own card-only attacks [13]). Courtois [6] exploited linear relations in the cipher to improve the attack described in Technical 8 in such a way that pre-computation is no longer necessary. Kaspe et al. [16] broke a popular payment system in Germany that uses MIFARE Classic. Tan, in his master thesis [20], surveyed and reproduced the attacks on the MIFARE Classic chip from the literature. He also brought these attacks to practice, taking as case studies the Imperial College's access control system and London's Oyster card. Van Deursen, Mauw and Radomirovic developed a methodology for the analysis and reverse engineering of sequences of card memory dumps [21]. They have applied this methodology to reverse engineer the e-go transport ticketing system of Luxembourg which also uses MIFARE Classic.

After the MIFARE hype, some members of the team started wondering whether other proprietary ciphers, developed by different manufacturers would also have so many weaknesses in their designs. This question led to an investigation into the security of the Atmel product family SecureMemory, CryptoMemory and CryptoRF. It resulted in a research paper [14] exposing serious vulnerabilities in these products as well. This story repeated in [3,4,10,12,25,26] reinforcing that proprietary cryptography and protocols often results in insecure constructions, and that 'security by obscurity' does not provide an extra layer of security but rather covers negligent designs.

References

1. Undisclosed authors: Counter expertise review of the TNO security analysis of the Dutch OV-Chipkaart. Technical report, Royal Holloway, University of London (2008). http://tinyurl.com/5wnqvrk
2. Undisclosed authors: Security analysis of the Dutch OV-Chipkaart. Technical report 34643, TNO (2008). http://www.translink.nl/media/bijlagen/nieuws/TNO_ICT_-_Security_Analysis_OV-Chipkaart_-_public_report.pdf
3. Balasch, J., Gierlichs, B., Verdult, R., Batina, L., Verbauwhede, I.: Power analysis of atmel cryptomemory – recovering keys from secure EEPROMs. In: Dunkelman, O. (ed.) CT-RSA 2012. LNCS, vol. 7178, pp. 19–34. Springer, Heidelberg (2012)

4. Blom, A., de Koning Gans, G., Poll, E., de Ruiter, J., Verdult, R.: Designed to fail: a USB-connected reader for online banking. In: Jøsang, A., Carlsson, B. (eds.) NordSec 2012. LNCS, vol. 7617, pp. 1–16. Springer, Heidelberg (2012)
5. Cho, A.: University hackers test the right to expose security concerns. Science **332**, 1322–1323 (2008)
6. Courtois, N.: The dark side of security by obscurity - and cloning Mifare Classic rail and building passes, anywhere, anytime. In: Fernández-Medina, E., Malek, M., Hernando, J. (eds.) SECRYPT, pp. 331–338. INSTICC Press (2009)
7. de Koning Gans, G.: Analysis of the MIFARE Classic used in the OV-Chipkaart project. Master's thesis, Radboud University Nijmegen (2008)
8. 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)
9. 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)
10. Garcia, F.D., de Koning Gans, G., Verdult, R.: Exposing iClass key diversification. In: 5th USENIX Workshop on Offensive Technologies (WOOT), pp. 128–136. USENIX Association, Berkeley (2011)
11. Garcia, F.D., de Koning Gans, G., Roel, V.: Tutorial: Proxmark, the swiss army knife for RFID security research. Technical report, Radboud University Nijmegen (2012)
12. Garcia, F.D., de Koning Gans, G., Verdult, R., Meriac, M.: Dismantling iClass and iClass Elite. In: Foresti, S., Yung, M., Martinelli, F. (eds.) ESORICS 2012. LNCS, vol. 7459, pp. 697–715. Springer, Heidelberg (2012)
13. Garcia, F.D., van Rossum, P., Verdult, ,R., Schreur, R.W.: Wirelessly pickpocketing a Mifare Classic card. In: IEEE Symposium on Security and Privacy (S&P), pp. 3–15. IEEE (2009)
14. Garcia, F.D., van Rossum, P., Verdult, R., Schreur, R.W.: Dismantling SecureMemory, CryptoMemory and CryptoRF. In: 17th ACM Conference on Computer and Communications Security (CCS), pp. 250–259. ACM (2010)
15. Hoepman, J.-H., Hubbers, E., Jacobs, B., Oostdijk, M., Schreur, R.W.: Crossing borders: security and privacy issues of the european e-passport. In: Yoshiura, H., Sakurai, K., Rannenberg, K., Murayama, Y., Kawamura, S. (eds.) IWSEC 2006. LNCS, vol. 4266, pp. 152–167. Springer, Heidelberg (2006)
16. Kasper, T., Silbermann, M., Paar, C.: All you can eat or breaking a real-world contactless payment system. In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, pp. 343–350. Springer, Heidelberg (2010)
17. Mayes, K.E., Cid, C.: The Mifare Classic story. Inf. Secur. Tech. Rep. **15**(1), 8–12 (2010)
18. Nohl, K., Evans, D., Starbug, S., Plötz, H.: Reverse-engineering a cryptographic RFID tag. In: USENIX Security 2008, pp. 185–193 (2008)
19. Nohl, K., Plötz, H.: Mifare, little security despite obscurity. Presentation at Chaos Computer Congress (2007)
20. Tan, W.H.: Practical attacks on the Mifare Classic. Master's thesis, Imperial College London (2009)
21. van Deursen, T., Mauw, S., Radomirović, S.: mCarve: Carving attributed dump sets. In: Proceedings of 20th USENIX Security Symposium, pp. 107–121. USENIX Association, August 2011
22. Verdult, R.: Proof of concept, cloning the OV-chip card. Technical report, Radboud University Nijmegen (2008)

23. Verdult, R.: Security analysis of RFID tags. Master's thesis, Radboud University Nijmegen (2008)
24. Verdult, R., de Koning Gans, G., Garcia, F.D.: A toolbox for RFID protocol analysis. In: 4th International EURASIP Workshop on RFID Technology (EURASIP RFID). IEEE Computer Society (2012)
25. Verdult, R., Garcia, F.D., Balasch, J.: Gone in 360 seconds: Hijacking with Hitag2. In: 21st USENIX Security Symposium (USENIX Security 2012). USENIX Association (2012)
26. Verdult, R., Kooman, F.: Practical attacks on NFC enabled cell phones. In: 3rd International Workshop on Near Field Communication (NFC), pp. 77–82. IEEE (2011)
27. Schreur, R.W., van Rossum, P., Garcia, F.D., Teepe, W., Hoepman, J.-H., Jacobs, B., de Koning Gans, G., Verdult, R., Muijrers, R., Kali, R., Kali, V.: Security flaw in MIFARE Classic. Press release, Digital Security group, Radboud University Nijmegen, The Netherlands, March 2008