Citation	Wim Aerts, Eli Biham, Dieter De Moitié, Elke De Mulder, Orr Dunkelman, Sebastiaan Indesteege, Nathan Keller, Bart Preneel, Guy A. E. Vandenbosch, Ingrid Verbauwhede, (2012), A Practical Attack on KeeLoq , Journal of Cryptology, 25(1), 136-157.
Archived version	Author manuscript: the content is identical to the content of the published paper, but without the final typesetting by the publisher
Published version	http://dx.doi.org/10.1007/978-3-540-78967-3_1
Journal homepage	http://www.springer.com/computer/theoretical+computer+science/journal/145
Author contact	ingrid.verbauwhede@esat.kuleuven.be + 32 (0)16 328625
IR	https://lirias.kuleuven.be/handle/123456789/283441

(article begins on next page)

A Practical Attack on KeeLoq^{*}

Wim Aerts^{1,2}, Eli Biham³, Dieter De Moitié¹, Elke De Mulder¹,
Orr Dunkelman⁴, Sebastiaan Indesteege¹, Nathan Keller⁴, Bart Preneel¹,
Guy A. E. Vandenbosch², and Ingrid Verbauwhede¹

¹ Department of Electrical Engineering ESAT/COSIC, Katholieke Universiteit
Leuven. Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium.
Interdisciplinary Institute for BroadBand Technology (IBBT), Belgium.

`sebastiaan.indesteege@esat.kuleuven.be`

² Department of Electrical Engineering ESAT/TELEMIC, Katholieke Universiteit
Leuven. Kasteelpark Arenberg 10, B-3001 Heverlee, Belgium.

³ Computer Science Department, Technion. Haifa 32000, Israel.

⁴ Faculty of Mathematics and Computer Science, Weizmann Institute of Science.
P.O. Box 26, Rehovot 76100, Israel.

Abstract. KeeLoq is a lightweight block cipher with a 32-bit block size and a 64-bit key. Despite its short key size, it is used in remote keyless entry systems and other wireless authentication applications. For example, there are indications that authentication protocols based on KeeLoq are used, or were used by various car manufacturers in anti-theft mechanisms. This paper presents a practical key recovery attack against KeeLoq that requires 2^{16} known plaintexts and has a time complexity of $2^{44.5}$ KeeLoq encryptions. It is based on the principle of slide attacks and a novel approach to meet-in-the-middle attacks.

We investigated the way KeeLoq is intended to be used in practice and conclude that our attack can be used to subvert the security of real systems. In some scenarios the adversary may even reveal the master secret used in an entire class of devices from attacking a single device. Our attack has been fully implemented. We have built a device that can obtain the data required for the attack in less than 100 minutes, and our software experiments show that, given the data, the key can be found in 7.8 days of calculations on 64 CPU cores.

Key words: KeeLoq, cryptanalysis, block ciphers, slide attacks, meet-in-the-middle attacks.

1 Introduction

The KeeLoq technology [21] by Microchip Technology Inc. includes the KeeLoq block cipher and several authentication protocols built on top of it. The KeeLoq block cipher allows for very low cost and power efficient hardware implementations. This property has undoubtedly contributed to the popularity of the cipher

^{*} The final publication is available at www.springerlink.com.

in various wireless authentication applications. There are indications that multiple car manufacturers use, or have used KeeLoq to protect their cars against theft [8,9,10,11,13,14,15]. According to some researchers [16,19,26], KeeLoq is also used in garage door openers. We verified these claims to the best of our ability, but no manufacturer seems eager to publicly disclose which algorithms are used. These claims seem reasonable, though, given the application areas for which KeeLoq is marketed by Microchip Inc., such as “vehicle alarm arming and disarming” and “home garage and gate door openers” [21]. The datasheet of the HCS410 KeeLoq transponder describes it as being “ideally suited to keyless entry systems, primarily for vehicles and home garage door openers” [22].

1.1 Previous Work

Despite its design in the 80’s, the first cryptanalysis of KeeLoq was only published by Bogdanov [8] in February 2007. This attack is based on the slide technique and a linear approximation of the non-linear Boolean function used in KeeLoq. The attack has a time complexity of 2^{52} KeeLoq encryptions and requires 16 GB of storage. It also requires the entire codebook, i.e., 2^{32} known plaintexts.

Courtois, Bard and Wagner [11,13] apply algebraic techniques to cryptanalyse KeeLoq. Although a direct algebraic attack fails for the full cipher, they reported various successful slide-algebraic attacks. For example, they claim that an algebraic attack can recover the key when given a slid pair in 2.9 seconds on average. As there is no way to ensure or identify the existence of a slid pair in the data sample, the attack is simply repeated 2^{32} times, once for each pair generated from 2^{16} known plaintexts. They also describe attacks requiring the entire codebook, which exploit certain assumptions with respect to fixed points of the internal state. The fastest of these requires 2^{27} KeeLoq encryptions (given the data set) and has an estimated success probability of 44% [13]. Courtois, Bard and Bogdanov show further extensions requiring the entire codebook in [14].

In [9,10], Bogdanov published an updated version of his original attack. A refined complexity analysis yields a slightly smaller time complexity, i.e., $2^{50.6}$ KeeLoq encryptions while still requiring the entire codebook. This paper also includes an improvement using the work of Courtois and Bard [11] on the cycle structure of the cipher. We note that the time complexity of the attack using the cycle structure given in [9,10] is based on an assumption from an earlier version of [11], that a random word can be read from 16 GB of memory with a latency of only one clock cycle. This is very unrealistic in a real machine, so the actual time complexity is probably much higher. In a later version of [11], this assumption on the memory latency was changed to be 16 clock cycles. In [15], Courtois exploits the self-similarity of KeeLoq to yield an optimised exhaustive key search attack. Finally, the PhD thesis of Bard [2] contains more details on several of these attacks.

Given that KeeLoq is a cipher that is widely used in practice, side-channel analysis is also a viable option for attacking chips that implement KeeLoq. In [16,26], differential power analysis (DPA) and differential electromagnetic

analysis (DEMA) were applied to several implementations of KeeLoq. It was shown that actual implementations of KeeLoq are not at all resistant to side-channel attacks, and can be broken using only 10–1000 measurement traces. Simple power analysis (SPA) was shown to be effective against an embedded software implementation of KeeLoq in [19].

1.2 Our Contribution

Our practical attack is based on slide attacks as well. However, unlike other attacks, we combine it with a novel meet-in-the-middle attack. The optimised version of the attack uses 2^{16} known plaintexts and has a time complexity of $2^{44.5}$ KeeLoq encryptions. We have implemented our attack and the total running time is roughly 500 days. As the attack is fully parallelisable, given x CPU cores, the total running time is only $500/x$ days. A variant which requires 2^{16} chosen plaintexts needs only $218/x$ days on x CPU cores. For example, for 10 000 euro, one can obtain at least 50 dual core computers, which will take about two days to find the key. Another, probably even cheaper, though illegal option would be to rent a botnet to carry out the computations.

Microchip Inc. specifies two protocols using KeeLoq: the “Code Hopping” and the “Identify Friend or Foe (IFF)” protocol [22]. In practice, the latter protocol, a simple challenge-response protocol, is the most interesting target to acquire the data that are necessary to mount the attack. Because the challenges are not authenticated in any way, an adversary can obtain as many chosen plaintext/ciphertext pairs as needed from a transponder (e.g., a car key) implementing this protocol. Depending on the transponder, it takes about 65 or 98 minutes to gather 2^{16} plaintext/ciphertext pairs. In addition to the computation phase of the attack, we have also built hardware that is capable of gathering the required data.

Finally, as was previously noted by Bogdanov [9,10], we show that one of the two suggested key derivation algorithms is blatantly flawed, as it allows an adversary to reconstruct many secret keys once a single secret key has been exposed.

Table 1 presents an overview of the known attacks on KeeLoq, including ours. In order to make comparisons possible, we have converted all time complexities to the number of KeeLoq encryptions needed for the attack.¹

The structure of this paper is as follows. In Sect. 2, we describe the KeeLoq block cipher and how it is intended to be used in practice. Our attacks are described in Sect. 3, including some related key attacks that are more of theoretical interest in Sect. 3.6. In Sect. 4 we discuss implementation aspects of the attack and experimental results and in Sect. 5 we show the relevance of our attacks in practice. Finally, in Sect. 6 we conclude.

¹ We list slightly better complexities for the attacks from [11,13] because we used a more realistic conversion factor from CPU clocks to KeeLoq rounds (i.e., 12 rather than four CPU cycles per KeeLoq round).

Table 1. An overview of attacks on KeeLoq.

Attack Type	Complexity			Reference
	Data	Time	Memory	
Time-Memory Trade-Off	2 CP	$2^{42.7}$	≈ 100 TB	[18]
Exhaustive Search	2 CP	2^{56}	16 GB	[15]
Slide/Algebraic	2^{16} KP	$2^{65.4}$	?	[11,13]
Slide/Algebraic	2^{16} KP	$2^{51.4}$	?	[11,13]
Slide/Guess-and-Determine	2^{32} KP	2^{52}	16 GB	[8]
Slide/Guess-and-Determine	2^{32} KP	$2^{50.6}$	16 GB	[9,10]
Slide/Algebraic	2^{32} KP	2^{40}	?	[14]
Slide/Cycle Structure	2^{32} KP	$2^{39.4}$	16.5 GB	[11]
Slide/Cycle/Guess-and-Det. ^a	2^{32} KP	(2^{37})	16.5 GB	[9]
Slide/Fixed Points	2^{32} KP	2^{27}	> 16 GB	[13]
Side-Channel Analysis (DPA)	10–1000 traces	-	-	[16,26]
Side-Channel Analysis (SPA)	1 measurement	-	-	[19]
Slide/Meet-in-the-Middle	2^{16} KP	$2^{45.0}$	≈ 2 MB	Sect. 3.3
Slide/Meet-in-the-Middle	2^{16} KP	$2^{44.5}$	≈ 3 MB	Sect. 3.4
Slide/Meet-in-the-Middle	2^{16} CP	$2^{44.5}$	≈ 2 MB	Sect. 3.5
Time-Memory-Data Trade-Off	68 CP, 34 RK	$2^{39.3}$	≈ 10 TB	[5]
Related Key	66 ACP, 34 RK $\gg$	negligible	negligible	Sect. 3.6
Related Key	512 CP, 2 RK $\gg$	2^{32}	negligible	Sect. 3.6
Related Key/Slide/MitM	2^{17} CP, 2 RK $\oplus$	$2^{41.9}$	≈ 16 MB	Sect. 3.6

Time complexities are expressed in full KeeLoq encryptions (528 rounds).

KP: known plaintexts; CP: chosen plaintexts; ACP: adaptive chosen plaintexts

RK $\gg$: related keys (by rotation); RK $\oplus$: related keys (flip LSB)

^a The time complexity for this attack is based on very unrealistic memory latency assumptions and hence will be much higher in practice.

2 Description and Usage of KeeLoq

In this section, we give a short description of the KeeLoq block cipher, and discuss two protocols built on it, “KeeLoq Hopping Codes” and “KeeLoq Identify Friend or Foe (IFF)”. The descriptions given here are based on the information in [23].

2.1 The KeeLoq Block Cipher

The KeeLoq block cipher has a 32-bit block size and a 64-bit key. It consists of 528 identical rounds each using one bit of the key. A round is equivalent to an iteration of a non-linear feedback shift register (NLFSR), as shown in Fig. 1.

More specifically, let $Y^{(i)} = (y_{31}^{(i)}, \dots, y_0^{(i)}) \in \{0, 1\}^{32}$ be the input to round i ($0 \leq i < 528$) and let $K = (k_{63}, \dots, k_0) \in \{0, 1\}^{64}$ be the key. The input to round 0 is the plaintext: $Y^{(0)} = P$. The ciphertext is the output after 528

Fig. 1. The i th KeeLoq encryption cycle.

rounds: $C = Y^{(528)}$. The round function can be described as follows (see Fig. 1):

$$\begin{aligned} \varphi^{(i)} &= \text{NLF} \left(y_{31}^{(i)}, y_{26}^{(i)}, y_{20}^{(i)}, y_9^{(i)}, y_1^{(i)} \right) \oplus y_{16}^{(i)} \oplus y_0^{(i)} \oplus k_{i \bmod 64} , \\ Y^{(i+1)} &= (\varphi^{(i)}, y_{31}^{(i)}, \dots, y_1^{(i)}) . \end{aligned} \quad (1)$$

The non-linear function NLF is a Boolean function of five variables with output vector **3A5C742E_x**—i.e., $\text{NLF}(i)$ is the i th bit of this hexadecimal constant, where bit 0 is the least significant bit. We can also represent the non-linear function in its algebraic normal form (ANF):

$$\begin{aligned} \text{NLF}(x_4, x_3, x_2, x_1, x_0) &= x_4 x_3 x_2 \oplus x_4 x_3 x_1 \oplus x_4 x_2 x_0 \oplus x_4 x_1 x_0 \oplus \\ &\quad x_4 x_2 \oplus x_4 x_0 \oplus x_3 x_2 \oplus x_3 x_0 \oplus x_2 x_1 \oplus x_1 x_0 \oplus \\ &\quad x_1 \oplus x_0 . \end{aligned} \quad (2)$$

Decryption uses the inverse round function, where i now ranges from 528 down to 1.

$$\begin{aligned} \theta^{(i)} &= \text{NLF} \left(y_{30}^{(i)}, y_{25}^{(i)}, y_{19}^{(i)}, y_8^{(i)}, y_0^{(i)} \right) \oplus y_{15}^{(i)} \oplus y_{31}^{(i)} \oplus k_{i-1 \bmod 64} , \\ Y^{(i-1)} &= (y_{30}^{(i)}, \dots, y_0^{(i)}, \theta^{(i)}) . \end{aligned} \quad (3)$$

There used to be some ambiguity about the correct position of the taps. Our description agrees with the “official” documentation [8,9,13,23]. Additionally, we have used test vectors generated by an actual HSC410 chip [22], manufactured by Microchip Inc., to verify that our description and implementation of KeeLoq are indeed correct. Finally, we note that our attacks are unaffected by this difference.

2.2 Protocols built on KeeLoq

A device like the HCS410 by Microchip Technology Inc. [22] supports two authentication protocols based on KeeLoq: “KeeLoq Hopping Codes” and “KeeLoq

Identify Friend or Foe (IFF)”. The former uses a 16-bit secret counter, synchronised between both parties. In order to authenticate, the encoder (e.g., a car key) increments the counter and sends the encrypted counter value to the decoder (e.g., the car), which verifies if the received ciphertext is correct. In practice, this protocol would be initiated by a button press of the car owner.

The second protocol, “KeeLoq Identify Friend or Foe (IFF)” [22], is a simple challenge-response protocol. The decoder (e.g., the car) sends a 32-bit challenge. The transponder (e.g., the car key) uses the challenge as a plaintext, encrypts it with the KeeLoq block cipher² under the shared secret key, and replies with the ciphertext. This protocol is executed without any user interaction whenever the transponder receives power and an activation signal via inductive coupling from a nearby decoder. Hence, no battery or button presses are required. It could for instance be used in vehicle immobilisers by placing the decoder near the ignition. Inserting the car key in the ignition would place the transponder within range of the decoder. The latter would then activate the transponder and execute the protocol, all completely transparent to the user. The car would then either disarm the immobiliser or activate the alarm, depending on whether the authentication was successful.

Of course both protocols can be used together in a single device, thereby saving costs. For example, the HCS410 chip [22] supports this combined mode of operation, possibly using the same secret key for both protocols, depending on the configuration options used.

3 Our Attacks on KeeLoq

This section describes our attacks on KeeLoq. We combine a slide attack with a novel meet-in-the-middle approach to recover the key from a slid pair. First we explain some preliminaries that are used in the attacks. Then, we proceed to the description of the attack scenario using known plaintexts and a generalisation thereof. Finally, we show how chosen plaintexts can be used to improve the attack. Section 3.6 explores related key attacks on KeeLoq that are more of theoretical interest.

3.1 The Slide Property

Slide attacks were introduced by Biryukov and Wagner [6] in 1999. The typical candidate for a slide attack is a block cipher consisting of a potentially very large number of iterations of an identical key dependent permutation F . In other words, the subkeys are repeated and therefore the susceptible cipher can be written as

$$C = F(\underbrace{F(\dots F(P))}_r) = F^r(P) . \quad (4)$$

² This corresponds to what is called the “HOP algorithm” in [22]. The other option, the so-called “IFF algorithm”, uses a reduced version of KeeLoq with 272 rounds instead of 528. Our attacks are also applicable to this variant, without any change.

Fig. 2. A slide attack.

The permutation $F(\cdot)$ does not necessarily have to coincide with the rounds of the cipher, i.e., F might combine several rounds of the cipher.

A slide attack aims at exploiting such a periodic structure to reduce the strength of the entire cipher to the strength of F . Thus, it is independent of the number of rounds of the cipher. To accomplish this, a so-called *slid pair* is needed. This is a pair of plaintexts that satisfies the slide property

$$P_2 = F(P_1) . \quad (5)$$

We depict such a slid pair in Fig. 2. For a slid pair, the corresponding ciphertexts also satisfy the slide property, i.e., $C_2 = F(C_1)$. By repeatedly encrypting this slid pair, we can generate as many slid pairs as needed [7,17]. As each slid pair gives us a pair of corresponding inputs and outputs of the key dependent permutation F , it can be used to mount an attack against F .

KeeLoq has 528 identical rounds, each using one bit of the 64-bit key. After 64 rounds the key is repeated. So in the case of KeeLoq, F is defined as 64 rounds. However, because the number of rounds in the cipher is not an integer multiple of 64, a straightforward slide attack [6] is not possible. A solution to this problem is to guess the 16 least significant bits of the key and use this to strip off the final 16 rounds. Then, a slide attack can be applied to the remaining 512 rounds [8,11,13].

In order to get a slid pair, 2^{16} known plaintexts are used. As the block size of KeeLoq is 32 bits, we expect that a random set of 2^{16} plaintexts contains a slid pair due to the birthday paradox.³ Determining which pair is a slid pair is done by the attack itself. Simply put, the attack is attempted with every pair. If it succeeds, the pair is very likely a slid pair, otherwise it is not.

3.2 Determining Key Bits

If two intermediate states of the KeeLoq cipher, separated by 32 rounds (or less) are known, all the key bits used in these rounds can easily be recovered. This was first described by Bogdanov [8], who refers to it as the “linear step” of his attack.

³ The probability that a set of 2^{16} random plaintexts contains at least one slid pair is $1 - (1 - 2^{-32})^{2^{16}} \approx 0.63$. Hence, the attack has a success probability of about 63%. With not much higher data complexity, higher success rates can be achieved.

Fig. 3. The notation used in the attack.

Let $Y^{(i)} = (y_{31}^{(i)}, \dots, y_0^{(i)})$ and $Y^{(i+t)} = (y_{31}^{(i+t)}, \dots, y_0^{(i+t)})$ be the two known states; $t \leq 32$. If we encrypt $Y^{(i)}$ by one round, the newly generated bit is

$$\varphi^{(i)} = \text{NLF} \left(y_{31}^{(i)}, y_{26}^{(i)}, y_{20}^{(i)}, y_9^{(i)}, y_1^{(i)} \right) \oplus y_{16}^{(i)} \oplus y_0^{(i)} \oplus k_{i \bmod 64} . \quad (6)$$

Because of the non-linear feedback shift register structure of the round function and since $t \leq 32$, the bit $\varphi^{(i)}$ is equal to $y_{32-t}^{(i+t)}$, which is one of the bits of $Y^{(i+t)}$ and thus known. Hence

$$k_{i \bmod 64} = \text{NLF} \left(y_{31}^{(i)}, y_{26}^{(i)}, y_{20}^{(i)}, y_9^{(i)}, y_1^{(i)} \right) \oplus y_{16}^{(i)} \oplus y_0^{(i)} \oplus y_{32-t}^{(i+t)} . \quad (7)$$

By repeating this t times, all t key bits can be recovered. The amount of computations that need to be carried out is equivalent to t rounds of KeeLoq. This simple step will prove to be very useful in our attack.

3.3 Basic Attack Scenario

We now describe the basic attack scenario, which uses 2^{16} known plaintexts. For clarity, the notation used is shown in Fig. 3 and a pseudocode overview is given in Fig. 4. We denote 16 rounds of KeeLoq by $g_{\hat{k}}$, where $\hat{k}$ denotes the 16 key bits used in these rounds. The 64-bit key k is split into four equal parts: $k = (\hat{k}_3, \hat{k}_2, \hat{k}_1, \hat{k}_0)$, where $\hat{k}_0$ contains the 16 least significant key bits.

As already mentioned in Sect. 3.1, the first step of the attack is to guess $\hat{k}_0$ —the 16 least significant bits of the key. This enables us to partially encrypt each of the 2^{16} plaintexts by 16 rounds (P_i to P_i^+) and partially decrypt each of the 2^{16} ciphertexts by 16 rounds (C_j to C_j^-). Encrypting P_i^+ by 16 more rounds yields P_i^{++} . Similarly, decrypting P_j by 16 rounds yields P_j^- (see Fig. 3).

Observation 1 *Let X_i and X_{i+t} be two KeeLoq states separated by t rounds, $t < 32$. Then, we have that $\text{msb}_{32-t}(X_i) = \text{lsb}_{32-t}(X_{i+t})$, i.e., the $32 - t$ most significant bits of X_i are equal to the $32 - t$ least significant bits of X_{i+t} .*

Proof. Note from the description of KeeLoq in Sect. 2 that, for a single round of KeeLoq, we see that $\text{msb}_{31}(X_i) = \text{lsb}_{31}(X_{i+1})$. Indeed, the 31 most significant bits of the state are merely shifted down over one bit position. Similarly, t rounds shift the most significant bits of the state down over t bit positions. Hence, if $t < 32$, $\text{msb}_{32-t}(X_i) = \text{lsb}_{32-t}(X_{i+t})$ holds. $\square$

```

1: for all  $\hat{k}_0 \in \{0, 1\}^{16}$  do
2:   for all plaintexts  $P_i, 0 \leq i < 2^{16}$  do
3:     Partially encrypt  $P_i$  to  $P_i^+$ .
4:     Partially decrypt  $C_i$  to  $C_i^-$ .
5:   end for
6:   for all  $\text{lsb}_{16}(P_j^-) \in \{0, 1\}^{16}$  do
7:     for all plaintexts  $P_j, 0 \leq j < 2^{16}$  do
8:       Determine the key bits  $\hat{k}_3$ .
9:       Partially decrypt  $C_j^-$  to  $C_j^{--}$ .
10:      Save the tuple  $\langle P_j^-, C_j^{--}, \hat{k}_3 \rangle$  in a table.
11:    end for
12:    for all plaintexts  $P_i, 0 \leq i < 2^{16}$  do
13:      Determine the key bits  $\hat{k}_1$ .
14:      Partially encrypt  $C_i$  to  $C_i^+$ .
15:      for all collisions  $\text{msb}_{16}(C_i^+) = \text{lsb}_{16}(C_j^{--})$  in the table do
16:        Determine the key bits  $\hat{k}_2$  from  $P_i^{++}$  and  $P_j^-$ .
17:        Determine the key bits  $\hat{k}_2'$  from  $C_i^+$  and  $C_j^{--}$ .
18:        if  $\hat{k}_2 = \hat{k}_2'$  then
19:          Encrypt two known plaintexts with the key  $k = (\hat{k}_3, \hat{k}_2, \hat{k}_1, \hat{k}_0)$ .
20:          if the correct ciphertexts are found then
21:            return success (the key is  $k$ )
22:          end if
23:        end if
24:      end for
25:    end for
26:  end for
27: end for
28: return failure (i.e., there was no slid pair)

```

Fig. 4. The attack algorithm.

If P_i and P_j form a slid pair, it follows from this observation that $\text{msb}_{16}(P_i^{++}) = \text{lsb}_{16}(P_j^-)$, since the states P_i^{++} and P_j^- are separated by 16 rounds.

The next step in the attack is to apply a meet-in-the-middle approach. We guess the 16-bit value $\text{lsb}_{16}(P_j^-)$. For each plaintext P_j we can then determine $\hat{k}_3$ using the algorithm described in Sect. 3.2. Indeed, as the other bits of P_j^- are determined by P_j , we know all of P_j^- when given the plaintext. There is always exactly one solution per plaintext. Using this part of the key, we can now partially decrypt C_j^- to C_j^{--} . This result is saved in a hash table indexed by the 16-bit value $\text{lsb}_{16}(C_j^{--})$. Each record in the hash table holds a tuple consisting of P_j^- , C_j^{--} and the 16 key bits $\hat{k}_3$.

Now we do something similar from the other side. For each plaintext we use the algorithm from Sect. 3.2 to determine $\hat{k}_1$. Again this can be done because we know all of P_i^{++} , and there is exactly one solution per plaintext. Knowing $\hat{k}_1$, we partially encrypt C_i to C_i^+ .

Note that if P_i and P_j are indeed a slid pair their partial encryptions and decryptions (under the correct key) must “meet in the middle”. More specifically, it must hold that $\text{msb}_{16}(C_i^+) = \text{lsb}_{16}(C_j^{--})$. So, we look for a record in the hash table for which such a collision occurs. Because the hash table is indexed by $\text{lsb}_{16}(C_j^{--})$ this can be done very efficiently. A slid pair produces a collision, provided the guesses for $\hat{k}_0$ and $\text{lsb}_{16}(P_j^-)$ are correct. Therefore, we are guaranteed that all slid pairs are found at some point. Of course, a collision does not guarantee that the pair is actually a slid pair.

Finally, we check each candidate slid pair found. We determine the remaining key bits $\hat{k}_2$ from P_i^{++} and P_j^- and similarly $\hat{k}_2'$ from C_i^+ and C_j^{--} . If $\hat{k}_2$ and $\hat{k}_2'$ are not equal, the candidate pair is not a slid pair. Note that we can determine the key bits one by one and stop as soon as there is a disagreement. This slightly reduces the complexity of the attack.

If $\hat{k}_2 = \hat{k}_2'$, we have found a pair of plaintexts and a key with the property that encrypting P_i by 64 rounds gives P_j and encrypting C_i by 64 rounds gives C_j . This is what is expected from a slid pair. It is however possible that the recovered key is not the correct key, so we can verify it by a trial encryption of one of the known plaintexts. Even if a wrong key is suggested during the attack, and discarded by the trial encryption, we are still guaranteed to find the correct key eventually, provided there is at least one slid pair among the given plaintexts.

Complexity Analysis. The complexity of the attack can be derived from the attack algorithm given in Fig. 4. Using one round of KeeLoq as a unit, the time complexity of the attack can be expressed as

$$2^{16} \left(32 \cdot 2^{16} + 2^{16} \left(32 \cdot 2^{16} + 2^{16} (32 + N_{\text{coll}} \cdot V) \right) \right), \quad (8)$$

when N_{coll} denotes the expected number of collisions for a single guess of $\hat{k}_0$, $\text{lsb}_{16}(P_j^-)$ and a given plaintext P_i , and V denotes the average cost of verifying one collision, i.e., checking if it leads to a candidate key and if this key is correct. This follows directly from the description of the attack. As the hash table has 2^{16} entries and a collision is equivalent to a 16-bit condition, $N_{\text{coll}} = 1$. In the verification step, we can determine one bit at a time and stop as soon as there is a disagreement, which happens with probability $1/2$. Only when there is no disagreement after 16 key bits, we do two full trial encryptions to check the recovered key. Of course the second trial encryption is only useful if the first one gave the expected result. Hence, due to this early abort technique, the average cost of verifying one collision is

$$V = 2 \cdot \sum_{i=0}^{15} 2^{-i} + 2^{-16} \cdot (528 + 528 \cdot 2^{-32}) \approx 4. \quad (9)$$

Thus the overall complexity of the attack is $2^{54.0}$ KeeLoq rounds, which amounts to $2^{45.0}$ full KeeLoq encryptions.

As mentioned before, the data complexity of the attack is 2^{16} known plaintexts. The storage requirements are very modest. The attack stores the plaintext/ciphertext pairs, 2^{16} values for P_i^+ and C_i^- , and a hash table with 2^{16} records of 80 bits each. This amounts to a bit over 2 MB of RAM.

3.4 A Generalisation of the Attack

The attack presented in the previous section can be generalised by varying the number of rounds to partially encrypt/decrypt in each step of the attack. We denote by t_p the number of rounds to partially encrypt from the plaintext side (left on Fig. 3) and by t_c the number of rounds to partially decrypt from the ciphertext side (right on Fig. 3). More specifically, encrypting P_i^+ by t_p rounds yields P_i^{++} , encrypting C_i by t_p rounds yields C_i^+ . On the ciphertext side, P_j^- is obtained by decrypting P_j by t_c rounds and C_j^{--} by decrypting C_j^- by t_c rounds. Also, the partial keys $\hat{k}_0$ through $\hat{k}_3$ are adapted accordingly to contain the appropriate key bits.

Let t_o denote the number of bits that, provided P_i and P_j form a slid pair, overlap between P_i^{++} and P_j^- . As P_i^{++} and P_j^- are separated by $48 - t_p - t_c$ rounds, we have that $t_o = 32 - (48 - t_p - t_c) = t_p + t_c - 16$. The t_o least significant bits of P_j^- are denoted by $\text{lsb}_{t_o}(P_j^-)$ and the t_o most significant bits of P_i^{++} are denoted by $\text{msb}_{t_o}(P_i^{++})$.

Depending on the choices for the parameters t_p and t_c , the attack scenario has to be modified slightly. If $t_c < t_o$, not all plaintexts necessarily yield a solution for a given $\text{lsb}_{t_o}(P_j^-)$ when determining $k_3 = (k_{63}, \dots, k_{64-t_c})$ because $t_o - t_c$ of the guessed bits overlap with plaintext bits. Similarly, if $t_c > t_o$, each plaintext is expected to offer multiple solutions because $t_c - t_o$ extra bits have to be guessed before all of P_j^- is known. From the other side, similar observations can be made.

In Sect. 3.3, the parameters were $t_p = t_c = 16$ which results in $t_o = 16$. It is clear that the choice of these parameters influences both the time and memory complexity of the attack.

Complexity Analysis. The generalisation leads to a slightly more complex formula for expressing the time complexity of the attack. Because of the duality between guessing extra bits and filtering due to overlapping bits, all cases can be expressed in a single formula, which is a generalisation of (8) (i.e., with $t_p = t_c = 16$, it reduces to (8)):

$$2^{16} (32 \cdot 2^{16} + 2^{t_o} (2t_c \cdot 2^{16+t_c-t_o} + 2^{16+t_p-t_o} (2t_p + N_{\text{coll}} \cdot V))) \quad (10)$$

In the generalised case, finding a collision is equivalent to finding an entry in a table of $2^{16+t_c-t_o}$ elements that satisfies a t_o bit condition, so $N_{\text{coll}} = 2^{16+t_c-t_o}/2^{t_o}$. Verifying a collision now requires an average effort of

$$V = 2 \cdot \sum_{i=0}^{47-t_p-t_c} 2^{-i} + 2^{t_p+t_c-48} \cdot (528 + 528 \cdot 2^{-32}) \quad (11)$$

KeeLoq rounds. Simplification yields that the total complexity is approximately

$$2t_c \cdot 2^{32+t_c} + 2t_p \cdot 2^{32+t_p} + 4 \cdot 2^{80-t_p-t_c} . \quad (12)$$

The optimum is found when $t_p = t_c = 15$ and thus $t_o = 14$, where the complexity reduces to $2^{53.524}$ KeeLoq rounds or $2^{44.5}$ full KeeLoq encryptions.

The memory requirements in the generalised case can also easily be evaluated. As before, 2^{16} plaintext/ciphertext pairs and 2^{16} values for P_i^+ and C_i^- are stored. The hash table now has $2^{16+t_p-t_o}$ entries of $64 + t_p$ bits each. For $t_p = t_c = 15$, the required memory is still less than 3 MB.

3.5 A Chosen Plaintext Attack

Using chosen plaintexts instead of known plaintexts, the attack can be improved. Consider the generalised attack from Sect. 3.4 in the case where $t_c < t_o$ (which is equivalent to $t_p > 16$). In this case, the $t_o - t_c$ least significant bits of the plaintext P_j are bits $(t_o, \dots, t_c + 1)$ of P_j^- . Hence, choosing the 2^{16} plaintexts in such a way that these $t_o - t_c$ least significant bits are equal to some constant, only 2^{t_c} guesses for $\text{lsb}_{t_o}(P_j^-)$ have to be made at the beginning of the meet-in-the-middle step, instead of 2^{t_o} .

Complexity Analysis. The time complexity of the chosen plaintext attack, in KeeLoq rounds, can be expressed as

$$2^{16} \left(32 \cdot 2^{16} + 2^{\min(t_c, t_o)} \left(2t_c \cdot 2^{\max(16, 16+t_c-t_o)} + 2^{16+t_p-t_o} (2t_p + N_{\text{coll}} \cdot V) \right) \right) . \quad (13)$$

The expected number of collisions is $N_{\text{coll}} = 2^{\max(16, 16+t_c-t_o)}/2^{t_o}$. The verification cost, V , is given by (11). Note that, when $t_c > t_o$, this reduces to the known plaintext case, see (10). Indeed, only when $t_c < t_o$, the chosen plaintexts can be used to accelerate the attack. After simplification and removal of negligible terms, the time complexity becomes

$$2t_c \cdot 2^{32+t_c} + 2t_p \cdot 2^{\min(48, 32+t_p)} + 4 \cdot 2^{80-t_p-t_c} . \quad (14)$$

The optimum is found when $t_p = 20$, $t_c = 13$ and thus $t_o = 17$, where the attack has a time complexity of $2^{53.5}$ KeeLoq rounds or $2^{44.5}$ full KeeLoq encryptions. It is clear that the (theoretical) advantage over the known plaintext attack from Sect. 3.4 is not significant. However, as is discussed in Sect. 4, the chosen plaintext variant can provide a significant gain in our practical implementation, because the verification cost V turns out to be higher there.

The memory complexity is about 2 MB as in Sect. 3.3 because the size of the hash table is the same. The data complexity remains at 2^{16} plaintext/ciphertext pairs, but note that we now require chosen plaintexts instead of known plaintexts.

Fig. 5. A related key attack using keys related by rotation.

3.6 Related Key Attacks on KeeLoq

Related key attacks [3] exploit the relations between the encryption processes under different but related keys. In this section we present two related key attacks on KeeLoq. The first attack is a very efficient attack using pairs of keys related by rotation. The second attack is an improvement of the attack presented in Sect. 3.3 using pairs of keys related by flipping the least significant bit of the key.

A Related Key Attack Using Keys Related by Rotation. The first attack exploits the extremely simple way in which the key is mixed into the state during encryption.

Denote a full encryption of a plaintext P by KeeLoq with the key K by $E_K(P)$, and encryption through a single round with the subkey bit k by $f_k(P)$. Consider a pair (K, K') of related keys, such that $K' = (K \ggg 1)$. If for a pair (P, P') of plaintexts we have $P' = f_{k_0}(P)$, where k_0 is the LSB of K , then $E_{K'}(P') = f_{k_{16}}(E_K(P))$. Indeed, in this case the encryption of P' under the key K' is equal to the encryption of P under K shifted by one round (see Fig. 5). This property, which is clearly easy to check, can be used to retrieve two bits of the secret key K .

Consider a plaintext P . We note that there are only two possible values of $f_{k_0}(P)$, i.e., $1|| (P \gg 1)$ and $0|| (P \gg 1)$. Hence, we ask for the encryption of P under the key K and for the encryption of the two plaintexts $P'_0 = 0|| (P \gg 1)$ and $P'_1 = 1|| (P \gg 1)$ under the related key K' , and check whether the ciphertexts satisfy the relation $E_{K'}(P') = f_{k_{16}}(E_K(P))$. This check is immediate, since $E_K(P)$ and $f_{k_{16}}(E_K(P))$ have 31 bits in common. Exactly one of the candidates (P'_0 or P'_1) is expected to satisfy the relation. This pair satisfies also the relation $P' = f_{k_0}(P)$. Actually, we note that it is sufficient to test only one of the two candidates (P'_0 or P'_1) as if one of them fails, then it is necessarily the other value (and the probability that both of them would seem correct by chance is negligible).

At this stage, since P' and P are known, we can infer the value of k_0 immediately from the update rule of KeeLoq, using the relation $P' = f_{k_0}(P)$. Similarly, we can retrieve the value of k_{16} from the relation $E_{K'}(P') = f_{k_{16}}(E_K(P))$. Hence, using only three chosen plaintexts encrypted under two related keys, we can retrieve two key bits with a negligible time complexity.

In order to retrieve additional key bits, we repeat the procedure described above with the pair of related keys $(K', K'' = (K' \ggg 1))$ and one of the plaintexts P'_0 or P'_1 examined in the first stage. As a result, we require the encryption of two additional chosen plaintexts (under the key K''), and get two additional key bits: k'_{016} and k'_{16} , which are equal to k_1 and k_{17} .

We can repeat this procedure 16 times to get bits $k_0, \dots, k_{31}$ of the secret key. Then, the procedure can be repeated with the 16 related keys of the form $(K \ggg 32), (K \ggg 33), \dots, (K \ggg 47)$ to retrieve the remaining 32 key bits. The attack then requires 66 plaintexts encrypted under 34 related keys (two plaintexts under each of 32 keys, and a single plaintext under the two remaining keys), and a negligible time complexity.

An option to reduce the required amount of plaintexts and related keys in exchange for a higher time complexity, is to switch to an exhaustive key search after a suitable number of key bits has been determined. For example, if 32 key bits remain to be found, a brute force search can be conducted in several hours on a PC, or even much less on FPGAs.

Another variant of the attack, requiring fewer related keys, is the following. Denote the encryption of a plaintext P through r rounds of KeeLoq with the key $k = (k_0, \dots, k_{r-1})$ by $f_k^r(P)$. Consider a pair of related keys of the form $(K, K' = K \ggg r)$. If a pair of plaintexts (P, P') satisfies $P' = f_k^r(P)$, then the corresponding ciphertexts satisfy $E_{K'}(P') = f_{k'}^r(E_K(P))$, where $k' = (k_{16}, \dots, k_{16+r-1})$. Since $E_K(P)$ and $f_{k'}^r(E_K(P))$ have $32 - r$ bits in common, this property is easy to check.

However, as r increases, it becomes more and more difficult to detect P' such that $P' = f_k^r(P)$ because there are 2^r candidates for P' . Checking all these candidate pairs increases the time complexity of the attack. The data complexity increases as well. A naive approach requires $1 + 2^r$ plaintexts: one P and all 2^r candidates for P' . But using structures, this data complexity can be improved to $2^{1+r/2}$. Consider two structures of $2^{r/2}$ plaintexts each. In the first structure, S_1 , the $32 - r$ most significant bits of each plaintext are equal to some constant C and the other bits are arbitrary. In the second structure, S_2 , the $32 - r$ least significant bits of each plaintext are equal to the same constant C and the other bits are arbitrary. By birthday paradox arguments on the 2^r possible pairs (P, P') such that $P \in S_1$ and $P' \in S_2$ we expect one pair for which $P' = f_k^r(P)$, and this pair can be used for the attack.

In the attack, we go over the 2^r possible pairs and check whether the colliding bits of the relation $E_{K'}(P') = f_{k'}^r(E_K(P))$ are satisfied. If $r \leq 16$, this check discards immediately most of the wrong pairs. After finding the right pair, $2r$ bits of the key can be found using the algorithm presented in Sect. 3.2.

By choosing different values of r , we can get several variants of the attack:

1. Using $r = 16$, we can recover 32 key bits, and then the rest of the key can be recovered using exhaustive key search. The data complexity of the attack is 512 chosen plaintexts encrypted under two related keys (256 plaintexts under each key), and the time complexity is 2^{32} KeeLoq encryptions.

2. Using $r = 8$ twice (for the pairs $(K, K \ggg 8)$, and $(K \ggg 8, K \ggg 16)$) we retrieve 32 key bits, and exhaustively search the remaining bits. The data complexity of the attack is 64 adaptive chosen plaintexts encrypted under three related keys (16 plaintexts under two keys, and 32 plaintexts under the third key), and the time complexity is 2^{32} KeeLoq encryptions.
3. Using $r = 8$ four times (for the pairs $(K, K \ggg 8)$, $(K \ggg 8, K \ggg 16)$, $(K \ggg 32, K \ggg 40)$, and $(K \ggg 40, K \ggg 48)$) we can retrieve the full key. The data complexity of the attack is 128 adaptive chosen plaintexts encrypted under six related keys (16 plaintexts under four keys, and 32 plaintexts under two keys), and the time complexity is negligible.

Other variants are also possible, and provide a trade-off between the number of chosen plaintexts and the number of related keys.

Improved Slide/Meet-in-the-Middle Attack Using Related Keys. Using a related key approach, we can improve the attack presented in Sect. 3.5. Denote the encryption of a plaintext P through 64 rounds of KeeLoq under the key K by $g_K(P)$. Denote by e_0 the least significant bit of a word. We observe that if two related keys (K, K') satisfy $K' = K \oplus e_0$, i.e., they differ in the least significant bit, and two plaintexts (P, P') satisfy $P' = P \oplus e_0$, then we have $g_K(P) = g_{K'}(P')$. Indeed, in the first round of encryption the key difference and the data difference cancel each other. As a result, after the first round the intermediate values in both encryptions are equal, and the key difference is not mixed into the data until the 65th round. Thus, the intermediate values after 64 rounds are equal in both encryptions.

Now, recall that in Sect. 3.1, the pair (P_i, P_j) is called a slid pair if it satisfies $P_j = g_K(P_i)$. The attack searches among 2^{32} candidates for a slid pair, and then the key can be easily retrieved. Note that by the observation above, if (P_i, P_j) is a slid pair with respect to K , then the pair $(P_i \oplus e_0, P_j)$ is a slid pair with respect to $K' = K \oplus e_0$, and thus $E_{K'}(P_j) = g_{(K' \ggg 16)}(E_{K'}(P_i \oplus e_0))$. This additional slid pair can be used to improve the check of candidate slid pairs, and thus to reduce the time complexity of the attack.

Indeed, during the verification phase, each additional slid pair gives us an additional way to determine each key bit from $\hat{k}_2$. If P_i and P_j are a slid pair under the original key K , then $P_i \oplus e_0$ and P_j must also be a slid pair under the related key K' . Thus, if at any point the different ways to determine bits from $\hat{k}_2$ disagree, we can conclude that P_i and P_j are not a slid pair. In this way, the related keys allow us to detect sooner that a plaintext pair is not a slid pair, saving time.

In this case, (13) can be rewritten as

$$2^{16} \left(48 \cdot 2^{16} + 2^{\min(t_c, t_o)} \left(3t_c \cdot 2^{\max(16, 16+t_c-t_o)} + 2^{16+t_p-t_o} (3t_p + N_{\text{coll}} \cdot V) \right) \right) . \quad (15)$$

The expected number of collisions becomes $N_{\text{coll}} = 2^{\max(16, 16+t_c-t_o)} / 2^{2t_o}$. Verifying a collision now costs on average V KeeLoq rounds, where

$$V = \sum_{i=0}^{47-t_p-t_c} (2 \cdot 2^{-2i} + 2^{-2i-1}) + 2^{2t_p+2t_c-96} \cdot (528 + 528 \cdot 2^{-32}) \approx 3.33 \quad (16)$$

Simplification yields the following

$$3t_c \cdot 2^{32+t_c} + 3t_p \cdot 2^{48} + 3.33 \cdot 2^{96-2t_p-2t_c} \quad (17)$$

The optimum is situated at $t_p = t_c = 12$ where the time complexity of the attack is $2^{50.9}$ KeeLoq rounds, or $2^{41.9}$ full KeeLoq encryptions.

The plaintexts are chosen such that for each plaintext P_i the set of 2^{16} chosen plaintexts also contains $P_i \oplus e_0$. This can be accomplished easily: pick 2^{15} 31-bit values X_i at random, and let $P_{2i} = X_i || 0$ and $P_{2i+1} = X_i || 1$. Then, the total data complexity is 2^{17} chosen plaintexts encrypted under two related keys (2^{16} plaintexts under each key). The time complexity is $2^{41.9}$ KeeLoq encryptions and the memory complexity is about 16 MB.

One can extend the attack by using more related key relations, for example by considering the key $K'' = K \oplus e_1$ as well. This can even further reduce the time complexity, in exchange for increased data complexity. One can see that, in general, by using k keys, the time complexity becomes

$$2^{16} \left(16(k+1) \cdot 2^{16} + 2^{\min(t_c, t_o)} ((k+1)t_c \cdot 2^{\max(16, 16+t_c-t_o)} + 2^{16+t_p-t_o} ((k+1)t_p + N_{\text{coll}} \cdot V)) \right) \quad (18)$$

The expected number of collisions is $N_{\text{coll}} = 2^{\min(16, 16+t_c-t_o)} / 2^{k \cdot t_o}$. Due to the additional possibilities for filtering, the verification cost, V , is now given by

$$V = \sum_{i=0}^{47-t_p-t_c} 2^{-k \cdot i} \left(1 + \sum_{j=0}^{k-1} 2^{-j} \right) + 2^{k(t_p+t_c-48)} \cdot (528 + 528 \cdot 2^{-32}) \quad (19)$$

$$\approx \left(\frac{1}{1-2^{-k}} \right) \left(1 + \frac{2^{-k}-1}{2^{-1}-1} \right) = 2 + \frac{2^k}{2^k-1} \quad (20)$$

Hence the time complexity in KeeLoq rounds is given by the following simplified expression:

$$t_c(k+1) \cdot 2^{32+t_c} + t_p(k+1) \cdot 2^{\min(48, 32+t_p)} + \left(2 + \frac{2^k}{2^k-1} \right) \cdot 2^{64-k(t_p+t_c-16)} \quad (21)$$

As long as $k \leq 17$, it is still possible to choose the plaintexts in a clever way, as before, such that for each P_i , the set of 2^{16} plaintexts also contains $P_i \oplus e_0, \dots, P_i \oplus e_{k-2}$. Hence, the data complexity is $k \cdot 2^{16}$ chosen plaintexts encrypted under k related keys (2^{16} plaintexts under each key). We list in Table 2 the obtained time complexities (and optimal values) for various number of keys.

Table 2. Complexity of Related Key Attacks.

Number of keys	Complexity		Optimal t_p, t_c
	Data	Time	
1	$2^{16.0}$ CP	$2^{44.5}$	$t_p = 20, t_c = 13$
2	$2^{17.0}$ CP	$2^{41.9}$	$t_p = t_c = 12$
3	$2^{17.6}$ CP	$2^{40.8}$	$t_p = t_c = 11$
4	$2^{18.0}$ CP	$2^{40.4}$	$t_p = 10, t_c = 11$ or vice versa.
5	$2^{18.3}$ CP	$2^{40.0}$	$t_p = t_c = 10$
6	$2^{18.6}$ CP	$2^{40.0}$	$t_p = 10, t_c = 10$
7	$2^{18.8}$ CP	$2^{39.9}$	$t_p = 9, t_c = 10$ or vice versa.
8	$2^{19.0}$ CP	$2^{40.0}$	$t_p = 9, t_c = 10$ or vice versa.

Time complexities are expressed in full KeeLoq encryptions (528 rounds).
CP: chosen plaintexts

4 Implementation Aspects

We have fully implemented and tested our attacks on KeeLoq that were described in Sect. 3.3, Sect. 3.4 and Sect. 3.5. In this section, we discuss certain practical aspects of our attack implementation and give our experimental results.

4.1 An Implementation of KeeLoq

To verify the accuracy and correctness of the description of KeeLoq found in [23], we built a software implementation of the KeeLoq block cipher and the various key derivation methods, based on [23]. Then, we used an actual HCS410 chip [22] to generate test vectors. Finally, we checked these test vectors against our software implementation. The fact that all test vectors matched indicates that the information gathered from [23] indeed gives a correct description of the workings of the KeeLoq system, and that our software implementation is correct.

4.2 Implementing the Attack

The attacks were implemented in the C language, using assembly for critical sections of the code. The overall structure of the implementation is identical to the pseudocode in Fig. 4. Our implementation is parametrised as described in Sect. 3.4.

Data structure. In line 10 of Fig. 4, tuples of the form $\langle P_j^-, C_j^{--}, \hat{k}_3 \rangle$ are saved in a table. This table is read again in line 15, where all tuples having a particular value for $\text{lsb}_{t_o}(C_j^{--})$ are searched. This usage pattern motivated our decision to use a simple hash table as the data structure for the table. The tuples are indexed by the value of $\text{lsb}_{t_o}(C_j^{--})$, which allows one to perform the

1: $x_1 \leftarrow x_1 \oplus x_3$	5: $x_3 \leftarrow x_3 \oplus x_2$	9: $x_3 \leftarrow x_3 \oplus t$
2: $x_0 \leftarrow x_0 \oplus x_2$	6: $x_0 \leftarrow x_0 \vee x_1$	10: $x_3 \leftarrow x_3 \wedge x_4$
3: $x_1 \leftarrow x_1 \oplus x_0$	7: $x_0 \leftarrow x_0 \oplus x_3$	11: $x_4 \leftarrow \neg x_4 \wedge x_0$
4: $t \leftarrow x_0$	8: $x_3 \leftarrow \neg x_3 \wedge x_1$	12: $x_3 \leftarrow x_3 \vee x_4$

Fig. 6. Bitsliced implementation of the KeeLoq non-linear function $\text{NLF}(x_4, x_3, x_2, x_1, x_0)$. The output is placed in x_3 , and t is a temporary variable. Each line corresponds to a single instruction in the SSE instruction set.

lookups in line 15 using a single memory access per lookup. Singly linked lists are used to accommodate for the possibility that there are multiple tuples in the table having the same value for $\text{lsb}_{t_o}(C_j^{--})$. This does not increase the required number of memory accesses per table lookup.

Bitslicing. Note that large parts of the attack algorithm in Fig. 4 can be performed in parallel. In particular, lines 3–4, 8–9 and 13–14 can be performed in parallel for each of the plaintexts. Because these parts of the attack algorithm can be parallelised efficiently, the collision verification phase (lines 15–24) becomes more expensive in comparison. Hence, the optimal parameters for our implementation differ slightly from the theoretical ones given earlier. For the known plaintext attack from Sect. 3.4, the optimal parameters for our implementation were found to be $t_p = t_c = 16$. This means that, at least in our implementation, the best attack is the basic attack from Sect. 3.3. For the chosen plaintext attack of Sect. 3.5, the optimal parameters are $t_p = 22$ and $t_c = 13$.

In practice, we perform 128 partial KeeLoq encryptions or decryptions in parallel using a technique called bitslicing. This technique was first introduced by Biham [4] as a method for speeding up software implementations of DES. The basic idea is to treat a machine word as a vector of bits, i.e., a case of “Single Instruction Multiple Data” (SIMD). The operations we want to perform are then implemented using elementary Boolean operations. We used the SSE instruction set to operate on words of 128 bits. Each bit of these vectors corresponds to a different encryption, hence we process 128 encryptions (or decryptions) in parallel.

Most of the KeeLoq cipher can be transformed into a bitsliced implementation easily. We focus on the non-linear function NLF, see Sect. 2. A search for an efficient bitsliced implementation of the NLF function was performed, using an algorithm similar to the one proposed by Osvik [25]. Figure 6 shows the resulting implementation of the NLF function. Each line corresponds to a single instruction in the SSE instruction set. Note that the PANDN instruction allows to perform the operation $x \leftarrow \neg x \wedge y$ in a single instruction.

Prefetching. When implemented in a naive way, attacks based on the meet-in-the-middle technique suffer from the high latency of memory lookups when

the tables grow beyond the size of the available cache memory. In the attack algorithm as shown in Fig. 4, the table lookups are performed in line 15. Note that, to a large extent, these lookups can be predicted. As explained above, we perform a block of partial encryptions and decryptions simultaneously. Hence, we immediately know many—128 in our case—table lookups that will be made next by the attack algorithm.

Prefetching allows us to instruct the processor to move specific data from the main memory to the cache memory, which can be accessed much faster. This technique does not eliminate the latency of a memory lookup, but it allows the processor to perform useful computations, while data that will be used in the near future are being fetched from memory. This drastically increases the throughput when compared to a straightforward approach, where most time would be spent waiting for data to arrive.

4.3 Experimental Results

If we give the correct values for the 16 least significant key bits, the known plaintext attack completes in 10.97 minutes on average.⁴ The chosen plaintext attack needs just 4.79 minutes to complete the same task.⁵ This corresponds to 499.22 CPU days resp. 217.82 CPU days ($9.49 \cdot 10^{16}$ CPU cycles resp. $4.14 \cdot 10^{16}$ CPU cycles) for the entire attack. The large difference between both attacks can be explained by considering the impact of V , the cost of the verification step, on the time complexity of the attack. If V increases, and t_p and t_c are adapted as needed because their optimal values may change, the time complexity of the known plaintext attack increases much faster than the time complexity of the chosen plaintext attack does. Hence, even though their theoretical time complexities are the same, the chosen plaintext attack performs much better in our practical implementation because V is higher than the theoretical value.

We did not stop either of the attacks once a slid pair and the correct key were found, so we essentially tested the worst-case behaviour of the attack. This also explains the very small standard deviations of the measured running times. The machine used is an AMD Athlon 64 X2 4200+ at 2.2 GHz with 1 GB of RAM (only one of the two CPU cores was used) running Linux 2.6.17. The C code was compiled with gcc version 4.1.2 (using the `-O3` optimiser flag).

The known plaintext attack performs over 288 times faster than the fastest attack with the same data complexity from [11,13], although the actual increase in speed is probably slightly smaller due to the difference in the machines used. Courtois, Bard and Wagner used (a single core of) a 1.66 GHz Intel Centrino Duo microprocessor [12]. The chosen plaintext attack performs more than 661 times faster, but this comparison is not very fair because chosen plaintexts are used. We note that the practicality of our results should also be compared with

⁴ We performed 500 experiments. The average running time was 658.15 s and the standard deviation was 1.69 s.

⁵ We performed 500 experiments. The average running time was 287.17 s and the standard deviation was 0.55 s.

exhaustive key search due to the small key size. For the price of about 10 000 euro, one can obtain a COPACOBANA machine [20] with 120 FPGAs which is estimated to take about 1000 days to find a single 64-bit KeeLoq key.⁶ Using our attack and 50 dual core computers (which can be obtained for roughly the same price), a KeeLoq key can be found in only two days.

5 Practical Applicability of the Attacks

This section investigates to which extent our attacks on KeeLoq, that were introduced in Sect. 3, can be applied to real systems using KeeLoq. In particular, we focus on how to gather the data required for the attack, and we discuss the impact of the key derivation schemes used in KeeLoq.

5.1 Gathering Data

One might wonder if it is possible to gather 2^{16} known, or even chosen plaintexts from a practical KeeLoq authentication system. As mentioned in Sect. 2.2, a device like the HCS410 by Microchip Technology Inc. [22] supports two authentication protocols based on KeeLoq: “KeeLoq Hopping Codes” and “KeeLoq Identify Friend or Foe (IFF)”. As the initial value of the counter used in “KeeLoq Hopping Codes” is not known, it is not easy to acquire known plaintexts from this protocol apart from trying all possible initial counter values. Also, since only 2^{16} plaintexts are ever used, knowing this sequence of 2^{16} ciphertexts suffices to break the system, i.e., to be able to always send the correct hopping code, as this sequence is simply repeated.

The second protocol, “KeeLoq Identify Friend or Foe (IFF)” [22], is more appropriate for our attack. It is executed without any user interaction as soon as the transponder comes within the range of a decoder and is sent an activation signal. The challenges sent by the decoder are not authenticated in any way. Because of this, an adversary can build a rogue decoder which can be used to gather as many plaintext/ciphertext pairs as needed. The plaintexts can be fully chosen by the adversary, so acquiring chosen plaintexts is no more difficult than just known plaintexts. The only requirement is that the rogue decoder can be placed within the range of the victim’s transponder for a certain amount of time. From the timings given in [22], we can conclude that one authentication completes within 60 ms or 90 ms on average, depending on the baud rate used. This translates into a required time of 65 or 98 minutes to gather the 2^{16} plaintext/ciphertext pairs. As these numbers are based on the maximum delay allowed by the specification [22], a real chip may respond faster, as our experiments confirm. No data are given with respect to the operational range in [22], because this depends on the circuit built around the HCS410 chip. However, one can expect the range to be short, likely not more than about ten centimetres.

⁶ The estimate was done by adapting the 17 days (worst case) required for finding a 56-bit DES key, taking into consideration the longer key size, the fact that more KeeLoq implementations fit on each FPGA, but in exchange take more clocks to test a key.

5.2 Key Derivation

The impact of the attack becomes even larger when considering the method used to establish the secret keys, as was previously noted by Bogdanov [9]. To simplify key management, the shared secret keys are derived from a 64-bit master secret (the manufacturer’s code), a serial number and optionally a seed value [9,23,24]. The manufacturer’s code is supposed to be constant for a large number of products (e.g., an entire series from a certain manufacturer) and the serial number of a transponder chip is public, i.e., it can easily be read out from the chip. The seed value is only used in the case of so-called “Secure Learning”, and can also be obtained from a chip with relative ease [9,23,24]. The other option, “Normal Learning”, does not use a seed value.

In both types of key derivation mechanisms, a 64-bit identifier is constructed, which contains the serial number, the (optional) seed and some fixed padding. Then, the secret key is derived from this identifier and the master secret using one of two possible methods. The first method simply uses XOR to combine the identifier and the master key. The consequence of this is that once a single key is known, together with the corresponding serial number and the (optional) seed value, the master secret can be found very easily.

The second method is based on decryption with the KeeLoq block cipher. The identifier is split into two 32-bit halves which are decrypted using the KeeLoq block cipher, and concatenated again to form the 64-bit secret key. The master secret is used as the decryption key. Although much stronger than the first method, the master secret can still be found using a brute force search. Evidently, once the master secret is known, so are all the keys that were derived from it, and the security of the entire system is compromised. Thus, it is a much more interesting target than a single secret key. This may convince an adversary to legitimately obtain a car key, for the sole purpose of recovering the master key from its secret key.

5.3 An Actual Transceiver

We have constructed a transceiver for KeeLoq for the purposes of validating the last aspect of the attack—the actual communications with the transponder. The transceiver is based on a Xilinx Virtex II FPGA, and we verified the ability to gather 2^{16} pairs of a chosen query and the corresponding reply. For details on the transceiver, we refer to Appendix E of [1].

We used queries of 32-bit long, starting from the all zero sequence, 00000000_x , and we increased the challenge by one until $0000FFFF_x$ was probed. We used the slower encoding method, which results in a maximal transmission time of 38.4 ms per challenge (for the all 1’s challenge, $FFFFFFF_x$).

Our transceiver also took care of decoding the signals sent back from the transponder. Again, using the slower encoding method, the maximal time for transmitting a reply is 19.2 ms, resulting in a total running time of the query process (along with the required pauses and opcode transmission) of 94.8 ms. Hence, in the slower method of communications, collecting 2^{16} pairs can be finished in less than 91 min.

6 Conclusion

In this paper we have presented a slide and meet-in-the middle attack on the KeeLoq block cipher which requires 2^{16} known plaintexts and has a time complexity of $2^{44.5}$ KeeLoq encryptions, and a variant using 2^{16} chosen plaintexts with the same theoretical time complexity.

We have fully implemented and tested both attacks. When given 16 key bits, the known plaintext attack completes successfully in 10.97 minutes. Due to implementation details, the chosen plaintext attack requires only 4.79 minutes when given 16 key bits. To the best of our knowledge, this is the fastest known attack on the KeeLoq block cipher.

Finally, we have shown that our attack can be used to attack real systems using KeeLoq due to the way it is intended to be used in practice. Moreover, one of the two suggested ways to derive individual KeeLoq keys from a master secret is extremely weak, with potentially serious consequences for the overall security of systems built using the KeeLoq algorithm.

Acknowledgments.

Orr Dunkelman was partially supported by the France Telecom chaire. Sebastian Indesteege ('Aspirant F.W.O. Vlaanderen') was supported by the Fund for Scientific Research Flanders (F.W.O. Vlaanderen). Nathan Keller was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities.

This work was supported in part by the Israel MOD Research and Technology Unit, by the Concerted Research Action (GOA) Ambiorics 2005/11 of the Flemish Government, by the IAP Programme P6/26 BCRYPT of the Belgian State (Belgian Science Policy), and by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II.

References

1. Wim Aerts, *Application Specificities of Array Antennas: Satellite Communication and Electromagnetic Side Channel Analysis*, PhD thesis, K.U.Leuven, 2009.
2. Gregory Bard, *Algorithms for Solving Linear and Polynomial Systems of Equations over Finite Fields with Applications to Cryptanalysis*, PhD thesis, Department of Mathematics, University of Maryland, College Park, 2007.
3. Eli Biham, *New Types of Cryptanalytic Attacks Using Related Keys*, Journal of Cryptology, Vol. 7, No. 4, pp. 229–246, Springer-Verlag, 1994.
4. Eli Biham, *A Fast New DES Implementation in Software*, Proceedings of Fast Software Encryption '97, Lecture Notes in Computer Science 1267, pp. 260–272, Springer-Verlag, 1997.
5. Alex Biryukov, Sourav Mukhopadhyay and Palash Sarkar, *Improved Time-Memory Tradeoffs with Multiple Data*, Proceedings of Selected Areas in Cryptography 2005, Lecture Notes in Computer Science 3897, pp. 245–260, Springer-Verlag, 2006.

6. Alex Biryukov and David Wagner, *Slide Attacks*, Proceedings of Fast Software Encryption '99, Lecture Notes in Computer Science 1636, pp. 245–259, Springer-Verlag, 1999.
7. Alex Biryukov and David Wagner, *Advanced Slide Attacks*, Advances in Cryptology, Proceedings of EUROCRYPT 2000, Lecture Notes in Computer Science 1807, pp. 586–606, Springer-Verlag, 2000.
8. Andrey Bogdanov, *Cryptanalysis of the KeeLoq block cipher*, Cryptology ePrint Archive, Report 2007/055, 16 February 2007, <http://eprint.iacr.org/2007/055/>.
9. Andrey Bogdanov, *Attacks on the KeeLoq Block Cipher and Authentication Systems*, 3rd Conference on RFID Security 2007 (RFIDSec 2007), available online at <http://rfidsec07.etsit.uma.es/slides/papers/paper-22.pdf>.
10. Andrey Bogdanov, *Linear Slide Attacks on the KeeLoq Block Cipher*, Proceedings of Inscrypt 2007, Lecture Notes in Computer Science 4990, pp. 66–80, Springer-Verlag, 2008.
11. Nicolas T. Courtois and Gregory V. Bard, *Algebraic and Slide Attacks on KeeLoq*, Cryptology ePrint Archive, Report 2007/062, 8 May 2007, <http://eprint.iacr.org/2007/062/>.
12. Nicolas T. Courtois, personal communication, 31 May 2007.
13. Nicolas T. Courtois, Gregory V. Bard and David Wagner, *Algebraic and Slide Attacks on KeeLoq*, Proceedings of Fast Software Encryption 2008, Lecture Notes in Computer Science 5086, pp. 97–115, Springer-Verlag, 2008.
14. Nicolas T. Courtois, Gregory V. Bard, and Andrey Bogdanov, *Periodic Ciphers with Small Blocks and Cryptanalysis of KeeLoq*, Tatra Mountains Mathematical Publications, vol. 41, pp. 167–188, 2008.
15. Nicolas T. Courtois, *Self-similarity Attacks on Block Ciphers and Application to KeeLoq*, International Workshop on Coding and Cryptography (WCC 2009), 2009.
16. Thomas Eisenbarth, Timo Kasper, Amir Moradi, Christof Paar, Mahmoud Salmasizadeh and Mohammad T. Manzuri Shalmani, *On the Power of Power Analysis in the Real World: A Complete Break of the KeeLoq Code Hopping Scheme*, Advances in Cryptology, Proceedings of CRYPTO 2008, Lecture Notes in Computer Science 5157, pp. 203–220, Springer-Verlag, 2008.
17. Soichi Furuya, *Slide Attacks with a Known-Plaintext Cryptanalysis*, Proceedings of Information and Communication Security 2001, Lecture Notes in Computer Science 2288, pp. 214–225, Springer-Verlag, 2002.
18. Martin E. Hellman, *A Cryptanalytic Time-Memory Trade-Off*, IEEE Transactions on Information Theory, vol. 26, pp. 401–406, 1980.
19. Markus Kasper, Timo Kasper, Amir Moradi, Christof Paar, *Breaking KeeLoq in a Flash: On Extracting Keys at Lightning Speed*, Proceedings of AFRICACRYPT 2009, Lecture Notes in Computer Science 5580, pp. 403–420, Springer-Verlag, 2009.
20. Sandeep Kumar, Christof Paar, Jan Pelzl, Gerd Pfeiffer, Manfred Schimmler, *Breaking Ciphers with COPACOBANA—A Cost-Optimized Parallel Code Breaker*, Proceedings of Cryptographic Hardware and Embedded Systems 2006, Lecture Notes in Computer Science 4249, pp. 101–118, Springer, 2006.
21. Microchip Technology Inc. *KeeLoq[®] Authentication Products*, <http://www.microchip.com/keeloq/>
22. Microchip Technology Inc., *HCS410 KeeLoq[®] Code Hopping Encoder and Transponder Data Sheet*, <http://ww1.microchip.com/downloads/en/DeviceDoc/40158e.pdf>
23. Microchip Technology Inc., *AN642: Code Hopping Decoder using a PIC16C56*, <http://www.keeloq.boom.ru/decryption.pdf>

24. Microchip Technology Inc., *TB001: Secure Learning RKE Systems using KeeLoq Encoders*, <http://ww1.microchip.com/downloads/en/AppNotes/91000a.pdf>
25. Dag Arne Osvik, *Speeding up Serpent*, The Third Advanced Encryption Standard Candidate Conference, pp. 317–329, 2000.
26. Christof Paar, Thomas Eisenbarth, Markus Kasper, Timo Kasper, Amir Moradi, *KeeLoq and Side-Channel Analysis-Evolution of an Attack*, Fault Diagnosis and Tolerance in Cryptography (FDTC 2009), pp. 65–69, 2009.

Dienst Onderzoekscoördinatie

Huis Bethlehem
Schapenstraat 34
B-3000 Leuven
België

Tel.: +32 16 32 40 65

Fax: +32 16 32 41 98

onderzoek@kuleuven.be

www.kuleuven.be