Related-Key Cryptanalysis of 3-WAY, Biham-DES,CAST, DES-X, NewDES, RC2, and TEA

John Kelsey Bruce Schneier
Counterpane Systems
{kelsey,schneier}@counterpane.com

David Wagner
U.C. Berkeley
daw@cs.berkeley.edu

Abstract. We present new related-key attacks on the block ciphers 3-WAY, Biham-DES, CAST, DES-X, NewDES, RC2, and TEA. Differential related-key attacks allow both keys and plaintexts to be chosen with specific differences [KSW96]. Our attacks build on the original work, showing how to adapt the general attack to deal with the difficulties of the individual algorithms. We also give specific design principles to protect against these attacks.

1 Introduction

Related-key cryptanalysis assumes that the attacker learns the encryption of certain plaintexts not only under the original (unknown) key K, but also under some derived keys $K' = f(K)$. In a chosen-related-key attack, the attacker specifies how the key is to be changed; known-related-key attacks are those where the key difference is known, but cannot be chosen by the attacker. We emphasize that the attacker knows or chooses the relationship between keys, not the actual key values. These techniques have been developed in [Knu93b, Bih94, KSW96].

Related-key cryptanalysis is a practical attack on key-exchange protocols that do not guarantee key-integrity—an attacker may be able to flip bits in the key without knowing the key—and key-update protocols that update keys using a known function: e.g., K, $K + 1$, $K + 2$, etc. Related-key attacks were also used against rotor machines: operators sometimes set rotors incorrectly. If the operator then corrected the rotor positions and retransmitted the same plaintext, an adversary would have a single plaintext encrypted in two related keys [DH79]. Hash functions built from block ciphers can also be vulnerable to a related-key attack against the block cipher [Win84, RIPE92].

In [KSW96] we gave a summary of key-schedule attacks against block ciphers, showed practical protocols that allow related-key attacks to be mounted, and presented related-key attacks against GOST [GOST89], IDEA [LMM91] with a reduced number of rounds, SAFER K-64 [Mas94], DES with independent subkeys, G-DES [PA90a, PA90b], and three-key triple-DES. This paper continues the research undertaken in that work.

2 New Differential Related-Key Attacks

2.1 3-WAY

3-WAY is an 11-round cipher on 96-bit blocks [Dae94]. Ignoring trivialities such as the input and output transformations, the 3-WAY round function $F(x)$ has an equivalent representation as:

$$y = N(x), \qquad z = L(y), \qquad F(x) = z \oplus K \oplus C_i$$

where N is a fixed nonlinear layer built out of 32 parallel 3-bit permutation S-boxes, L is a fixed linear function, K is the 96-bit master key, and C_i is a fixed, round-dependent public constant.

3-WAY is vulnerable to a simple related-key differential attack. It is trivial to find a differential characteristic for one S-box with probability 1/4, so we can construct a characteristic $\Delta x \to \Delta y$ with probability 1/4 for the non-linear layer N by using only one active S-box. By linearity we see that $\Delta y \to \Delta z = L(\Delta y)$ with probability 1 under the linear layer L. If we pick $\Delta K = \Delta x \oplus \Delta z$, then $\Delta x \to \Delta x$ by F with probability 1/4, which is a one-round iterative differential characteristic. In this way we can derive a 9-round characteristic with probability 2^{-18} to cover rounds 1–9, and apply a 2R analysis to the last two rounds. This breaks 3-WAY with one related-key query and about 2^{22} chosen plaintexts.

2.2 DES-X

DES-X is a DES variant proposed by Rivest [Riv95] to strengthen DES against exhaustive attacks. The DES-X encryption of P with key (K_1, K_2, K_3) is simply

$$C = K_1 \oplus \mathrm{DES}_{K_2}(K_3 \oplus P)$$

where K_3 is the pre-whitening key and K_1 is the post-whitening key. DES-X has many complementation properties. Furthermore, every DES-X key (K_1, K_2, K_3) has another equivalent key $(\overline{K_1}, \overline{K_2}, \overline{K_3})$. Therefore, DES-X cannot be used in a Davies-Meyer-like hash function construction.

This complementation property leads to an attack which requires roughly $2^{56+64-n}$ trial encryptions when 2^n chosen plaintexts are available [Dae91]. Note that Kilian and Rogaway [KR96] have proven that this attack is theoretically approximately optimal when DES is viewed as a black box, so any better (non-related-key) attack would have to take advantage of the internal structure of DES. However, their proof doesn't deal with related-key attacks. We give a related-key differential attack on DES-X, using key differences modulo 2^{64} and plaintext differences modulo 2. The attack requires 64 chosen key relations to recover the key, with one plaintext encrypted under each new key.

We start with a simple intuition. Suppose we have some unknown number Z. We are allowed to add any number we like modulo 2^{64}, and then XOR it with another number of our choosing. We are told whether or not the result of our calculation is equal to Z. Thus, we choose T and U, and test whether

$$(Z + T \bmod 2^{64}) \oplus U = Z$$

It is clear that we can learn the value of Z with enough queries. This is essentially the position we are in with DES-X. We can add T to K_1, and XOR U into our plaintext block or visa versa. If the resulting ciphertext block is the same as the ciphertext that results from encrypting the unaltered plaintext block under the unaltered DES-X key, then we can restrict the list of possible values for K_1. With enough such restrictions, we recover all of K_1 except for its high-order bit. This then allows attacks against the remainder of DES-X.

The simplest version of this attack uses T and U values each with the same single bit on. For each bit except the high-order bit, we try a T, U pair with the same bit on. If this results in the same ciphertext as resulted when $T = U = 0$, then we learn that that bit in K_1 was a zero. If it results in a different ciphertext, then we learn that that bit in K_1 was a one.

Some have suggested [KR96] using a DES-X variant which replaces the XOR pre- and post-whitening steps by addition modulo 2^{64}:

$$C = K_1 + \mathrm{DES}_{K_2}(K_3 + P).$$

From the discussion above, it should be clear that this would be vulnerable to a related-key attack very similar to the one that works against regular DES-X. [KR96] recommends a method of deriving DES-X keys from a single starting key, using SHA-1. This method seems to defend against related-key attacks.

2.3 CAST

CAST is a Feistel cipher whose key schedule uses nonlinear S-boxes [Ada94].[1] The key schedule for 8 round CAST with a 64 bit master key is as follows:

$$(k_1, k_2, \ldots, k_8) = \text{Master Key}$$

$$\begin{aligned}(k'_1, k'_2, k'_3, k'_4) &= (k_1, k_2, k_3, k_4) \oplus S5[k_5] \oplus S6[k_7] \\ (k'_5, k'_6, k'_7, k'_8) &= (k_5, k_6, k_7, k_8) \oplus S5[k'_2] \oplus S6[k'_4]\end{aligned}$$

$$\begin{array}{llll} K_1 = (k_1, k_2) & K_2 = (k_3, k_4) & K_3 = (k_5, k_6) & K_4 = (k_7, k_8) \\ K_5 = (k'_4, k'_3) & K_6 = (k'_2, k'_1) & K_7 = (k'_8, k'_7) & K_8 = (k'_6, k'_5) \end{array}$$

$$(K_{r,1}, K_{r,2}) = K_r \qquad r = 1, \ldots, 8$$

$$sk_r = S5[K_{r,1}] \oplus S6[K_{r,2}] \qquad r = 1, \ldots, 8.$$

where $S5$ and $S6$ are different 8-bit to 32-bit S-boxes. The r-th round subkey, sk_r, is XORed into the input of the F function as is conventional for Feistel ciphers.

[1] The variant of CAST analyzed here is an older version of CAST, not the CAST-128 that is used in Entrust products and described in Internet RFC 2144 [Ada97].

CAST is an interesting example of a cipher designed to resist Biham's rotational related-key cryptanalysis, but not differential related-key cryptanalysis. We apply a key-difference to the master key which changes only the byte k_1; this will lead to a difference only in round subkeys sk_1 and sk_6. When Δk_1 is known, there are only 256 possible differences for Δsk_1; by encrypting 2^{16} chosen plaintexts under each key, we can ensure that the first round is bypassed for some pair. Cover rounds 2–5 with the trivial differential characteristic of probability 1, and use a 2R attack. Note that sk_7 and sk_8 have only 32 bits of entropy in total, so we can try all 2^{32} possibilities for them, decrypt the last two rounds, and recognize correct guesses by 32 zero bits in the block difference. We recover the rest of the key with 2^{16} offline guesses by auxiliary techniques. In the end, we can recover the entire CAST master key with a total of about 2^{17} chosen plaintexts, one related-key query, and 2^{48} offline computations.

2.4 Biham-DES

Biham and Biryukov have suggested strengthening DES against exhaustive attacks by using extra key bits to modify the F-function slightly [BB94]. One of their modifications uses 5 key bits to select from 32 possible reorderings of the 8 DES S-boxes. We consider related keys which differ only in those 5 bits, and we apply related-key differential cryptanalysis. Specifically, suppose one key uses ordering 15642738 and another uses ordering 75642138 (both are from the 32 suggested reorderings listed in [BB94]). The only difference between the two F-functions is that S-boxes 1 and 7 have been swapped. Observe that:

$$\Pr_x \left(S1[x] \oplus S7[x \oplus 2] = 0 \right) \;=\; 14/64.$$

The input differential 2 appears only in the middle input bits of the S-box, and will not spread to neighboring S-boxes. Hence, we can construct a one-round characteristic with probability $(\frac{14}{64})^2$.

This leads to a 13-round iterative characteristic with probability $(\frac{14}{64})^{12} = 2^{-26}$. The differential techniques of Biham and Shamir [BS93] will break Biham-DES with 2^{27} chosen plaintexts when this special related-key pair is available.

If two related keys allow the above attack (i.e. differ only in the key orderings as defined above), we call them partners. There is a $\frac{1}{16}$ chance that a randomly chosen key will have a partner; if it does, this can be detected with one related-key probe. Furthermore, we can always obtain one useful pair of related-key partners from any starting key after 32 related-key queries. Therefore, when using Biham-DES with the 32 recommended DES S-box reorderings, we have a $\frac{1}{16}$ probability of success when 2^{27} chosen plaintexts and one related-key query are available; success is nearly guaranteed with 2^{31} chosen plaintexts and 32 related-key queries.

Biham and Biryukov also mention the possibility of using 2^{15} reorderings of the s^3-DES S-boxes [KPL93]. They don't present the recommended reorderings, so it is impossible to present any specific results. Still, in general, increasing the number of reorderings gives the cryptanalyst more degrees of freedom to find more efficient attacks. Therefore, using this variant is not expected to increase security against our attack.

2.5 RC2

RC2 is a block cipher designed by Ron Rivest [Riv97]. The RC2 key schedule takes an arbitrary length master key and expands it to 128 bytes with the help of a public non-linear 8-bit permutation ρ; the result is converted to 64 16-bit round subkeys.[2] We have analyzed RC2, and found single-bit differential characteristics which pass through most rounds with probability $\frac{1}{2}$.

Consider a 64-byte master key $K = (x_0, x_1 ..., x_{63})$; its related-key partner will be $K^* = (x_0^*, x_1, ..., x_{63}^*)$. In other words, K and K^* differ only in their first and last bytes. We choose x_0, x_{63}, x_0^*, and x_{63}^* so that $\rho[x_0 + x_{63}] = \rho[x_0^* + x_{63}^*]$. This is easy—we just subtract t from x_0 and add it to x_{63} to obtain K^*, where t is a byte quantity to be carefully chosen below. The RC2 key schedule expands K to the 128-byte expanded key $xk_{0..127}$ as follows:

$$xk_{0..63} = x_{0..63} \qquad\qquad xk_i = \rho[xk_{i-1} + xk_{i-64}] \quad \forall\, i \geq 64.$$

We observe that $xk_{0..127}$ and $xk_{0..127}^*$ differ only in positions 0, 63, and 127.

Next, note that we know the difference t between xk_0 and xk_0*. This makes it very easy to bypass the subkey difference entering round 0 in a chosen plaintext attack by using a suitable plaintext pair P, P^*. P^* is just P with t added to its high byte. Let P_i be P after i rounds, where each round is $\frac{1}{4}$ of a cycle. We have

$$\begin{aligned} P_0^* &= P_0 + 2^{56}t \\ P_i^* &= P_i \qquad i = 1, \ldots, 31 \\ P_{32}^* &= P_{32} + t \end{aligned}$$

If we choose a difference t with only one bit set, then we've just dropped a one-bit difference into the middle of the cipher. Note that there is an iterative four-round (one-cycle) differential characteristic with this one-bit difference as input and probability 2^{-4}. This leads to a 28-round characteristic with probability 2^{-28}, which can be used in a 4R attack.

The probability of the characteristic is slightly decreased by two different cycles in the middle of encryption processing. There are eight such rounds; each has a 2^{-5} chance of hitting one of the two changed key words and destroying the propagation of a right pair. The chance of successfully missing all of these of $(1-2^{-5})^8 \approx 0.775$. Furthermore, one of those variant rounds adds a quantity with difference 0 to a quantity with a one-bit difference, which halves the probability of our characteristic. Finally, a subsequent variant takes the low 6 bits of a quantity

[2] There is also an optional key-weakening stage, intended for export control use. For our purposes, we will assume it is not used.

with a one-bit difference as input; a careful choice of t can ensure that the one-bit difference falls in the high 2 bits, so that the characteristic is not disrupted. We have to multiply the earlier estimate by $0.775 \cdot 0.5$, obtaining a total probability of $2^{-29.4}$ for our characteristic. With this technique, RC2 can be broken with one related-key query and about 2^{34} chosen plaintexts.

2.6 NewDES

NewDES [Sco85] is a 17-round 64-bit block cipher with a 120-bit key. The key schedule is simple: each cycle (which consists of 2 rounds) uses 56 bits from the key and then shifts the key by 56 bits. NewDES succumbs to standard rotation related-key techniques: it can be broken with 2^{32} known plaintexts, one related key, and about 2^{56} offline trial encryptions.

When informed of this attack, Scott modified the NewDES key schedule to resist rotational related-key cryptanalysis [Sco96]. NewDES-1996 in turn falls to differential related-key cryptanalysis.

The NewDES-1996 key schedule expands 15 bytes $K0 \ldots K14$ of the master key K into 60 round subkey bytes $SK0 \ldots SK59$ according to the following pattern:

$$\begin{array}{lllll} K0 & K1 & K2 & \ldots & K14 \\ K0 \oplus K7 & K1 \oplus K7 & K2 \oplus K7 & \ldots & K14 \oplus K7 \\ K0 \oplus K8 & K1 \oplus K8 & K2 \oplus K8 & \ldots & K14 \oplus K8 \\ K0 \oplus K9 & K1 \oplus K9 & K2 \oplus K9 & \ldots & K14 \oplus K9 \end{array}$$

When $K7, K8, K9$ are all non-zero, this updated key schedule defeats rotational related-key cryptanalysis, as the sequence of round subkeys no longer repeats.[3]

Note that the NewDES-1996 key schedule is completely linear and exhibits poor avalanche. In fact, it falls to a differential related-key attack we call the *double-swiping* attack.

The double-swiping attack is somewhat involved, with technical and notational distractions, so we first describe the basic flow of the attack. We derive three related keys K', K^*, and $K^{*\prime}$ from the original key K according to a differential quartet structure. We take an arbitrary ciphertext P and apply a plaintext difference to it to obtain P^*; for a right pair P, P^* the attack will succeed, and a right pair occurs with very high probability. "Swipe" P back and forth through the NewDES-1996 cipher: encrypt P under K to obtain C, and decrypt $C' = C$

[3] There are weak keys—namely those where $K7 = K8 = K9 = 0$—that succumb easily to rotational related-key cryptanalysis given 2^{32} known plaintexts, one related key, and about 2^{56} offline trial encryptions.

This leads to a more general rotational-based attack on NewDES-1996. For any key K, after 2^{24} related-key probes one can find a weak key K' of known relation to K, recover K' by the above attack on NewDES-1996 weak keys, and thus find K. However, this attack requires about 2^{25} related-key queries, 2^{56} known plaintexts, and 2^{80} offline trial encryptions in general; therefore, we have disregarded this attack on NewDES-1996 as impractical.

under K' to obtain P'. Next swipe P^* back and forth: encrypt P^* under K^* to obtain C^*, and decrypt $C^{*\prime} = C^*$ under $K^{*\prime}$ to obtain $P^{*\prime}$. For a right pair, it turns out that the quartet key structure ensures that P' and $P^{*\prime}$ will be nearly the same, differing only in the action of $SK0'$ and $SK0^{*\prime}$; a final analysis stage reveals $SK0$ from P' and $P^{*\prime}$. Now we peel off the effect of $K0$ and iterate to find the rest of the key bytes.

The double-swiping attack is an optimization of a more conventional (*single-swiping*) related-key differential attack. The more conventional attack proceeds by decrypting $C = C'$ under both K and K' to obtain P, P'; the problem is that (with NewDES-1996) the single-swiping attack requires a 4R analysis stage on P, P', which appears rather tricky to perform as it must take into account the effect of 15 round subkey bytes $SK0 \ldots SK14$. The intuition is that the double-swiping attack allows us to insert a difference much closer to the end of the cipher, so the analysis stage depends only on $SK0$ and thus becomes much easier. The single-swiping related-key attack is already a big improvement over non-related-key attacks, but we can do even better by double-swiping.

We now present the technical details of the double-swiping attack. Fix any two byte values x, y, and take three related keys $K', K^*, K^{*\prime}$ according to the quartet structure

$$\begin{aligned} K' &= K \oplus (x, x, x, \ldots, x) \\ K^* &= K \oplus (y, 0, 0, \ldots, 0) \\ K^{*\prime} &= K \oplus (x \oplus y, x, x, \ldots, x). \end{aligned}$$

The related keys can be obtained under the differential related-key assumption. Note that, with these definitions, we have

$$\begin{aligned} SK'i &= SKi \oplus \begin{cases} x \text{ if } i = 0, \ldots, 14 \\ 0 \text{ if } i = 15, \ldots, 59 \end{cases} \\ SK^*i &= SKi \oplus \begin{cases} y \text{ if } i = 0, 15, 30, 45 \\ 0 \text{ otherwise} \end{cases} \\ SK^{*\prime}i &= SKi \oplus SK'i \oplus SK^*i. \end{aligned}$$

For some plaintext $P = P0$, we will use the notation Pi to indicate the intermediate value of the block after encryption with the first i subkey bytes; for instance, $P15$ is the output after the first two rounds, and $P60 = C$ is the final ciphertext block. When we "swipe" the first time to obtain $C = P60 = P60' = C'$ and $P' = P0'$, in general we have $P0' \neq P0$. However, since SKi and $SK'i$ differ only for $i < 15$, note that $P15' = P15$. We define $P^* = P0^* = P0 \oplus \Delta = P \oplus \Delta$, where Δ is carefully chosen to bypass [BS93] the key difference $SK0 \oplus SK0^* = y$ entering in the first step of the first round. Define a right pair as a pair P, P^* where $P1^* = P1$; examination of the NewDES F function reveals that the carefully-chosen values $x \oplus y = 224$ and $\Delta = 18$ cause right pairs to occur with probability $\frac{12}{256} \approx 1/21.3$. After the second swipe, we have $P15^* = P15^{*\prime}$, since SKi^* and $SKi^{*\prime}$ differ only for $i < 15$. Furthermore, the quartet structure of the related keys ensures that $P15 = P15' = P15^* = P15^{*\prime}$ for a right pair. In particular, we have $P1' = P1^{*\prime}$ for a right pair. Note that $P0', P0^{*\prime}$ are known, and they differ from $P1', P1^{*\prime}$ only in the application of a 8-bit to 8-bit F function keyed by $SK0', SK0^{*\prime}$. Therefore, we can apply a standard differential 1R analysis stage [BS93] to P' and $P^{*\prime}$; one can filter out wrong pairs very effectively, so recovering $SK0$ should be possible with just one right pair.

This double-swiping differential attack finds one subkey byte $SK0$ with a quartet of differentially related keys and about 88 chosen-plaintext/ciphertext queries. Now we can peel off the effect of the first subkey byte $SK0$ and iterate the attack to recover $SK1$, etc. Thus we can recover all 15 key bytes $(K0, \ldots, K14) = (SK0, \ldots, SK14)$ and completely break NewDES-1996 with total complexity of about 24 related-key probes and 530 chosen plaintext/ciphertext queries.

2.7 TEA

TEA [WN95] is a Feistel block cipher with a 128-bit master key, $K[0..3]$, and a simple key schedule: odd rounds use $K[0,1]$ as the round subkey, and even rounds use $K[2,3]$. Two rounds of TEA applied to the block Y_i, Z_i consists of:

$$c = c + \delta \qquad Y_{i+1} = Y_i + F(Z_i, K[0,1], c) \qquad Z_{i+1} = Z_i + F(Y_{i+1}, K[2,3], c)$$

where the round function F is defined by

$$F(z, K[i,j], c) = (SL_4(z) + K[i]) \oplus (z + c) \oplus (SR_5(z) + K[j]).$$

Here $SL_4(z)$ denotes the result of shifting (not rotating) z to the left by 4 bits, and $SR_\cdot(\cdot)$ denotes a shift to the right. In this description, c is a value which perturbs the F function so that it is different in each round.[4] Before each cycle, c is incremented by a fixed constant $\delta = \lfloor(\sqrt{5} - 1)2^{31}\rfloor$; c is initially 0. The designers of TEA mention that 32 Feistel rounds (i.e. 16 cycles) may be enough, though they recommend using 64 rounds (32 cycles) [WN95].

TEA admits several related-key attacks which arise from the severe simplicity of its key schedule.

Attack One For a differential related-key attack, consider the effect of simultaneously flipping bit 30 (the next most significant bit) of $K[2]$ and $K[3]$. With probability nearly $\frac{1}{2}$, the output of the F function in the even rounds will remain the same. This immediately yields a 2-round iterative differential characteristic with probability $\frac{1}{2}$, and thus a 60-round characteristic with probability 2^{-30}. Our analysis indicates that a 4R differential related-key attack can break 64-round (32-cycle) TEA with one related-key query and about 2^{34} chosen plaintexts. This is only one of several of this type of characteristic.

[4] This perturbation is crucial to avoid degenerate attacks. Indeed, R. Fleming found a known-plaintext attack on a TEA variant weakened to use a constant c [Fle96]. (His variant also differs from TEA in that the the precedence of addition and XOR are reversed [Ber97], but a modification of his attack will work without this reversal.)

Attack Two The second differential related-key attack is very similar in spirit to the first. We request the encryption of (Y, Z) under key $K[0..3]$ and the encryption of $(Y, Z \oplus 2^{31})$ under key $K^*[0..3] = K[0..3] \oplus (0, 2^{31} \oplus 2^{26}, 0, 0)$. Examining the three terms of $F(Z, K[0,1], c)$ when bit 31 of Z is flipped along with bits 26 and 31 of $K[1]$, we see

$SL_4(Z) + K[0]$	Neither change has any effect.
$Z + c$	The high bit is always changed.
$SR_5(Z) + K[1]$	Half the time, only the high bit is changed.

This gives us a one-cycle (2-round) iterative differential characteristic with probability $\frac{1}{2}$, when we can choose one key difference. We can pass 30 rounds with probability 2^{-30}.

Attack Three The third attack is complicated. Therefore, we briefly point out the approach and intuition behind the attack, leaving the technical details of the full attack to be described in Appendix A. We write P_j to represent the value of the block after j rounds of encryption, and write K_j to represent the round subkey value used to compute P_{j+1} from P_j; the block is enciphered with a round function F as $P_{j+1} = F(K_j, P_j)$, where (P_0, P_{64}) represents a plaintext/ciphertext pair for 64-round TEA.

In Biham's standard key rotation attack [Bih94], we succeed when

$$K'_j = K_{j+1} \qquad P'_j = P_{j+1} \qquad\qquad j = 0, \ldots, 63.$$

This condition is achieved by choosing suitable related keys K, K' and searching over P_0, P'_0 to find a pair with $P'_0 = P_1$; the birthday paradox ensures that a match will occur with a reasonable number of known texts. Note that

$$P'_{j+1} = F(K'_j, P'_j) = F(K_{j+1}, P_{j+1}) = P_{j+2} \tag{1}$$

for all j, so by induction we see that a match $P'_0 = P_1$ will propagate down to the ciphertexts, where we can recognize it.

Our extended attack combines the ideas of both rotational and differential related-key attacks. We require that

$$K'_j = K_{j+1} + \Delta K_{j+1} \qquad P'_j = P_{j+1} + \Delta P_{j+1} \qquad\qquad j = 1, \ldots, 63.$$

In the extended attack, we need a generalization of (1) to hold

$$\begin{aligned} P'_{j+1} &= F(K'_j, P'_j) = F(K_{j+1} + \Delta K_{j+1}, P_{j+1} + \Delta P_{j+1}) \\ &= F(K_{j+1}, P_{j+1}) + \Delta P_{j+2} = P_{j+2} + \Delta P_{j+2} \end{aligned}$$

with significant probability p_{j+2}; this generalization has a strong differential feel to it. Suppose the 63-round differential related-key characteristic that is patched into the rotational attack has probability $p = \prod_j p_j$. In the extended attack, we search for about $\frac{1}{p}$ matches $P'_0 = P_1 + \Delta P_1$ with the birthday paradox. Each such match has a probability p of leading to a right pair that is recognizable from the known ciphertexts, so we expect to see one right pair.

Specifically, in our third attack on TEA, we take ΔP_{j+1} to be a fixed constant $(\delta, 0)$ independent of j, set $\Delta K_j = \Delta K_{j \bmod 2}$, and choose $\Delta K_{0,1}$ to maximize p. We can thus obtain a full 63-round characteristic of probability $p = (\frac{25}{32})^{31} \approx 2^{-11}$ by repeating a 2-round iterative characteristic many times.

This improved attack combines ideas from both Biham's key-rotation attack [Bih94] and differential related-key cryptanalysis [KSW96] to break TEA with just 2^{23} chosen plaintexts and one related-key query.

3 Prudent Rules of Thumb for Key-Schedule Design

There is much overlap between the requirements for strong key schedules and cryptographic hash functions. Firstly, key schedules should be hard to invert—given some of the round keys, it should be difficult to recover any new information about other bits of the key—and hash functions are supposed to be one-way. Secondly, to avoid equivalent keys, key schedules should possess some form of collision-freedom; collision-freedom is a standard hash function property as well. Finally, it should not be possible to produce controlled changes in the round keys. The key schedules of Blowfish [Sch94] and SEAL [RC94] were designed according to this principle.

One should typically avoid generating round subkeys as a (fixed, public) linear transformation of the seed. While some cryptosystems have successfully incorporated linear key schedules (e.g. DES), designing this type of key schedule appears to be a subtle and difficult task. Many ciphers' linear key schedules have been shown to be quite weak: we have cryptanalyzed TEA, 3-WAY, and GOST [KSW96], and others have cryptanalyzed LOKI [Knu93a], LOKI91 [Knu93b], Lucifer [BB93], and SAFER [Knu95].

To protect against the known related-key attacks, we propose several attack-oriented design goals. To avoid the "subkey rotation" attacks [Bih94], round subkeys should be generated differently, so that each key bit affects nearly every round, but not always in the same way. Key schedules should be specifically designed to resist differential related-key attacks. And, when related-key queries are cheap, the master key should be long enough to avoid generic black box attacks, as the key length is effectively halved under these attacks [WH87, KSW96].

Avoid dead spots; ensure that every key bit is about equally powerful in terms of its effect on the round keys. Beware of equivalent representations, for they can expose new avenues of attack to an adversary. Our analysis of 3-WAY bears witness to this recommendation.

Avoid independent round subkeys. It has commonly been assumed that a cipher's key length (and strength) can be increased by allowing round keys to be specified independently, but we have shown that this dramatically lowers the cipher's resistance to related-key attacks [KSW96]. In general, when independent round subkeys are in use, the strength of a cipher against related-key attacks will be approximately proportional to the strength of one round standing on its own. Additionally, avoid multiple encryption with independent keys; a construction like [DK96] is much more secure.

And finally, protocol designers should be aware of related-key attacks. Key-exchange protocols should exchange a short master key rather than exchanging expanded keys. Design tamper-resistant devices so that it is not possible to change the subkeys without such changes being detected.

References

[Ada94] C. Adams, "Simple and Effective Key Scheduling for Symmetric Ciphers," *Workshop on Selected Areas in Cryptography: SAC '94*, 1994, pp 129-133.

[Ada97] C. Adams, "Constructing Symmetric Ciphers Using the CAST Design Procedure," *Designs, Codes and Cryptography*, v 12, n 3, 1997, to appear.

[BB93] I. Ben-Aroya and E. Biham, "Differential Cryptanalysis of Lucifer," *Advances in Cryptology—CRYPTO '93*, Springer-Verlag, 1994, pp. 187–199.

[Ber97] D. Bernstein, personal communication, 1997.

[Bih94] E. Biham, "New Types of Cryptanalytic Attacks Using Related Keys," *Advances in Cryptology—EUROCRYPT '93*, Springer-Verlag, 1994, pp. 398–409.

[BB94] E. Biham and A. Biryukov, "How to Strengthen DES Using Existing Hardware," *Advances in Cryptology—ASIACRYPT '94*, Springer-Verlag, pp. 398–412.

[BS93] E. Biham and A. Shamir, "Differential Cryptanalysis of the Full 16-round DES," *Advances in Cryptology—CRYPTO '92*, Springer-Verlag 1993, pp. 487–496.

[Dae91] J. Daemen, "Limitations of the Even-Mansour Construction," *Advances in Cryptology—ASIACRYPT '91*, Springer-Verlag, 1992, pp. 495–498.

[Dae94] J. Daemen, "A New Approach to Block Cipher Design," *Fast Software Encryption, Cambridge Security Workshop Proceedings*, Springer-Verlag, 1994, pp. 18–32.

[DK96] I.B. Damgard and L.R. Knudsen, "Multiple Encryption with Minimum Key," *Cryptography: Policy and Algorithms*, Springer-Verlag, 1996, pp. 156–164.

[DH79] W. Diffie and M.E. Hellman. "Privacy and Authentication: An Introduction to Cryptography". *Proceedings of the IEEE*, vol 67 no 3, March 1979.

[Fle96] R. Fleming, "An attack on a weakened version of TEA," post to the `sci.crypt` newsgroup, October 1996.

[GOST89] GOST, Gosudarstvennyi Standard 28147-89, "Cryptographic Protection for Data Processing Systems," Government Committee of the USSR for Standards, 1989.

[KSW96] J. Kelsey, B. Schneier, and D. Wagner, "Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES," Advances in Cryptology—CRYPTO '96, Springer-Verlag, 1996, pp. 237–251.

[KPL93] K. Kim, S. Park, and S. Lee, "Reconstruction of s^2DES S-Boxes and their Immunity to Differential Cryptanalysis," *Proceedings of the 1993*

Japan-Korea Workshop on Information Security and Cryptography, Seoul, Korea, 24-26 October 1993, pp. 282-291.

[Knu93a] L.R. Knudsen, "Cryptanalysis of LOKI," *Advances in Cryptology—ASIACRYPT '91*, Springer-Verlag, 1993, pp. 22–35.

[Knu93b] L.R. Knudsen, "Cryptanalysis of LOKI91," *Advances in Cryptology—AUSCRYPT '92*, Springer-Verlag, 1993, pp. 196–208.

[Knu94] L.R. Knudsen, "Block Ciphers—Analysis, Design, Applications," Ph.D. dissertation, Aarhus University, Nov 1994.

[Knu95] L.R. Knudsen, "A Key-schedule Weakness in SAFER K-64," *Advances in Cryptology—CRYPTO '95*, Springer-Verlag, 1995, pp. 274–286.

[KR96] J. Kilian and P. Rogaway, "How to protect DES against exhaustive key search," *Advances in Cryptology—CRYPTO '96*, Springer-Verlag, 1996, pp. 252–267.

[LMM91] X. Lai, J. Massey, and S. Murphy, "Markov Ciphers and Differential Cryptanalysis," *Advances in Cryptology—CRYPTO '91*, Springer-Verlag, 1991, pp. 17–38.

[Mas94] J.L. Massey, "SAFER K-64: A Byte-Oriented Block-Ciphering Algorithm", *Fast Software Encryption, Cambridge Security Workshop Proceedings*, Springer-Verlag, 1994, pp. 1–17.

[PA90a] A. Pfitzmann and R. Abmann, "Efficient Software Implementations of (Generalized) DES," *Proc. SECURICOM '90*, Paris, 1990, pp. 139–158.

[PA90b] A. Pfitzmann and R. Abmann, "More Efficient Software Implementations of (Generalized) DES," Technical Report PfAb90, Interner Bericht 18/90, Fakultat fur Informatik, Universitat Karlsruhe, 1990. http://www.informatik.uni-hildesheim.de/~sirene/lit/abstr90.html#PfAss_90

[RIPE92] Research and Development in Advanced Communication Technologies in Europe, *RIPE Integrity Primitives: Final Report of RACE Integrity Primitives Evaluation (R1040)*, RACE, Jun 1992.

[Riv95] R. Rivest, personal communication.

[Riv97] R. Rivest, "A Description of the RC2(r) Encryption Algorithm," Internet-Draft, work in progress, June 1997, ftp://ds.internic.net/internet-drafts/draft-rivest-rc2desc-00.txt

[RC94] P. Rogaway and D. Coppersmith, "A Software-Optimized Encryption Algorithm," *Fast Software Encryption, Cambridge Security Workshop Proceedings*, Springer-Verlag, 1994, pp. 56–63.

[Sch94] B. Schneier, "Description of a New Variable-Length Key, 64-Bit Block Cipher (Blowfish)," *Fast Software Encryption, Cambridge Security Workshop Proceedings*, Springer-Verlag, 1994, pp. 191–204.

[Sco85] R. Scott, "Wide Open Encryption Design Offers Flexible Implementations," *Cryptologia*, v. 9, n. 1, Jan 1985, pp. 75–90.

[Sco96] R. Scott, "Revision of NewDES," personal communication, also posted to the `sci.crypt` newsgroup on the Internet, May 1996.

[WN95] D. Wheeler and R. Needham, "TEA, a Tiny Encryption Algorithm," *Fast Software Encryption, Second International Workshop Proceedings*, Springer-Verlag, 1995, pp. 97–110.

[Win84] R. Winternitz, "Producing One-Way Hash Functions from DES," *Advances in Cryptology: Proceedings of Crypto 83*, Plenum Press, 1984, pp. 203–207.

[WH87] R. Winternitz and M. Hellman, "Chosen-key Attacks on a Block Cipher," *Cryptologia*, v. 11, n. 1, Jan 1987, pp. 16–20.

A Improved Attack on TEA

This attack combines ideas from Biham's key-rotation attack and differential cryptanalysis. It requires only 2^{23} chosen plaintexts and one related-key query. See Section 2.7 for a gentler introduction to the ideas behind the attack.

If $K[0\ldots4]$ is one TEA key value, its related key partner is defined to be $K'[0\ldots4]$ according to the following relations:

$$K'[0]=K[2] \quad K'[1]=K[3] \quad K'[2]=K[0]-SL_4(\delta) \quad K'[3]=K[1]-SR_5(\delta)-1.$$

(Refer to Section 2.7 for a definition of $SL(), SR()$, and other notation.) Fix a particular plaintext y, z which is encrypted via $K[\,]$; its related plaintext partner (which is encrypted with $K'[\,]$) will be offset from y, z by 1/2 cycle, as in rotational related-key cryptanalysis. Typically, in related-key cryptanalysis, we search for a partnered plaintext pair by the birthday paradox, and the right choice leads to a recognizable match in the corresponding ciphertexts with probability 1. In this generalization, we will consider the case where right choices of plaintext pairs leads to recognizable matches in the ciphertext with some non-trivial probability, via a differential characteristic.

The following table shows the encryption of y, z under key $K[\,]$ as well as the encryption of its offset plaintext partner $y', z' = z+\delta, y$ under key $K'[\,]$. The left half of the table depicts the left and right halves of the block when encrypting y, z; the right half of the table depicts the encryption of y', z'. (We consider the swap of the block halves to be included in each round.) We have placed y_{j+1}, z_j (respectively y_{j+1}, z_{j+1}) on the same line as y'_j, z'_j (resp. y'_{j+1}, z'_j) to suggest that the two propagate similarly. As described in Section 2.7, $F(z, K[i,j], c)$ denotes the value the round F function with input z, key values $K[i], K[j]$ with the round-dependent perturbation variable equal to c; c is incremented by δ before each cycle to make the F function different for each round.

Encrypt($K[\,], y_0z_0$)		Encrypt($K'[\,], y'_0z'_0$)	
$y_0=y$	$z_0=z$		
z_0	$y_1=y_0+F(z_0,K[0,1],\delta)$	$y'_0=z_0+\delta$	$z'_0=y_1$
y_1	$z_1=z_0+F(y_1,K[2,3],\delta)$	z'_0	$y'_1=y'_0+F(z'_0,K'[0,1],\delta)$
z_1	$y_2=y_1+F(z_1,K[0,1],2\delta)$	y'_1	$z'_1=z'_0+F(y'_1,K'[2,3],\delta)$
y_2	$z_2=z_1+F(y_2,K[2,3],2\delta)$	z'_1	$y'_2=y'_1+F(z'_1,K'[0,1],2\delta)$
…		…	
y_{32}	$z_{32}=z_{31}+F(y_{32},K[2,3],32\delta)$	z'_{31}	$y'_{32}=y'_{31}+F(z'_{31},K'[0,1],32\delta)$
		y'_{32}	$z'_{32}=z'_{31}+F(y'_{32},K'[2,3],32\delta)$

We define a right pair for the differential characteristic to be a pair (y_0, z_0), (y'_0, z'_0) satisfying

$$y'_j = z_j+\delta \qquad z'_j = y_{j+1} \qquad\qquad j=0,\ldots,31.$$

Since $K[2,3]=K'[0,1]$, we see from the table that we will never deviate from the right-pair condition in an odd round if it holds at the start of the odd round.

Therefore we have a right pair just if the condition holds for all even rounds; the table shows that the required condition is

$$F(z_j, K[0,1], (j+1)\delta) = F(y'_j, K'[2,3], j\delta) \qquad j = 0, \ldots, 31. \qquad (2)$$

Expanding the right-hand-side and then simplifying, we obtain

$$\begin{aligned}&(SL_4(z_j+\delta) + K[0] - SL_4(\delta)) \oplus (z_j + \delta + j\delta)\\ &\quad \oplus (SR_5(z_j+\delta) + K[1] - SR_5(\delta) - 1)\\ &= (SL_4(z_j) + K[0]) \oplus (z_j + (j+1)\delta) \oplus (SR_5(z_j) + K[1] + \Omega_j - 1)\end{aligned}$$

where $\Omega_j = SR_5(z_j+\delta) - SR_5(z_j) - SR_5(\delta)$, i.e. Ω_j is the carry bit from the addition of the low 5 bits of z_j and δ. Comparing to the right-hand-side of (2), we see that condition (2) is equivalent to the requirement that $\Omega_j = 1$ for $j = 0, \ldots, 31$. A quick check of the low 5 bits of δ shows that $\Omega_j = 1$ with probability $\frac{25}{32}$ when z_j is random.

In other words, the differential characteric carries through one cycle with probability $\frac{25}{32}$, and through 31 cycles with probability $\frac{25}{32}^{31} = .00047 = 2^{-11}$. Now we use the differential characteristic in the rotational related-key attack; we find it increases the number of plaintexts required by a factor of $2^{-11/2}$ over the number that would be required for a standard probability 1 attack.

Here is the attack in more detail. First fix a value for z_0. Now generate $2^{16+11/2} = 2^{21.5}$ values of $y_0^{(m)}$, for $m = 1 \ldots 2^{21.5}$, and encrypt the resulting value $y_0^{(m)}, z_0$ under $K[\,]$ to obtain the ciphertext $y_{32}^{(m)}, z_{32}^{(m)}$. Next set $y'_0 = z_0 + \delta$, and generate $2^{21.5}$ values of ${z'_0}^{(n)}$. For each ${z'_0}^{(n)}$, with $n = 1 \ldots 2^{21.5}$, encrypt $y'_0, {z'_0}^{(n)}$ under $K'[\,]$ to obtain the ciphertext ${y'_{32}}^{(n)}, {z'_{32}}^{(n)}$. Look for matches of the form $z_{32}^{(m)} = {y'_{32}}^{(n)}$. We expect to see one right match formed from a right pair of the differential characteristic combined with a right partnership ${z'_0}^{(n)} = y_1^{(m)}$ for the rotational attack; there will also be approximately $2^{21.5 \cdot 2}/2^{32} = 2^{11}$ matches formed by chance. Each right match allows you to recover roughly 64 key bits: it suggests about 2^{32} possible values for $K[0,1]$ and about 2^{32} possible values for $K[2,3]$.

One could repeat the attack a few more times and use a counting technique to recover the full key values with a bit more work. In more detail, each match suggests a value for $F(z_0, K[0,1], 0)$; we can now construct y_0, z'_0 pairs which are guaranteed to form a right partnership for the rotational attack, when used with the same z_0 value as before. For each guess at $F(z_0, K[0,1], 0)$, we can perform 2^{11} chosen plaintext queries; then we can recognize the true value of $F(z_0, K[0,1], 0)$ because it will cause another right pair and matching ciphertext pair. Thereafter, we can perform 2^{20} chosen plaintext queries and obtain 2^9 right pairs for the differential characteristic. This will be more than enough to recover the true value of $K[2,3]$ and find 2^{32} possible values for $K[0,1]$, so a simple search will suffice to recover the entire key.

In total, this attack needs 2^{23} chosen plaintexts, one related-key query, and roughly 2^{32} offline computations to recover the entire TEA key.

This article was processed using the LaTeX macro package with LLNCS style