ELSEVIER

Discrete Applied Mathematics 111 (2001) 103–116

DISCRETE
APPLIED
MATHEMATICS

On the structure of Skipjack

Lars Knudsen[a,*], David Wagner[b]

[a] *Department of Informatics, University of Bergen, N-5020 Bergen, Norway*
[b] *University of California Berkeley, Soda Hall, Berkeley, CA 94720, USA*

Received 3 September 1999; revised 10 May 2000; accepted 13 July 2000

Abstract

In this paper the structure of the NSA-designed block cipher Skipjack is examined. By cryptanalysing a large number of variants of the algorithm, we give plausible arguments for several principles behind the Skipjack design. The conclusion is that the algorithm seems to be carefully designed. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Cryptology; Cryptanalysis; Block cipher; Skipjack

1. Introduction

Skipjack is a 64-bit block cipher that is used in the Clipper Chip [7,14]. The design principles of Skipjack have not been published and the algorithm itself was only recently made public by the NSA [12,13].

Skipjack is a remarkably simple cipher, and one interesting feature is the use of two different types of rounds. These are referred to as A, and B-rounds and encryption with Skipjack consists of first applying 8 A-rounds, then 8 B-rounds, once again 8 A-rounds and finally 8 B-rounds. Earlier papers have demonstrated that the number of rounds was apparently not chosen with a large margin of security [2,3,10], but they did not focus on the high-level structure of Skipjack.

In this paper we examine the structure of Skipjack, focusing especially on understanding the rationale behind the design choices embodied in the cipher. A central motivation is the observation that Skipjack is just one representative from a very large design space of related ciphers, and in particular there are many parameters which could easily be changed to get a different construction with presumably different properties. Consequently, it is natural to wonder whether the designers of Skipjack missed any opportunities to improve the cipher by selecting one of the other alternatives in this design space.

[*] Corresponding author.
E-mail addresses: lars.knudsen@ii.uib.no (L. Knudsen), daw@cs.berkeley.edu (D. Wagner).

PII: S0166-218X(00)00347-4

In an attempt to shed light on these issues, we propose the following working hypothesis: *Skipjack appears to be essentially the only strong cipher in a very large family of weaker constructions.* We provide a large body of evidence for this conjecture by showing that many other natural choices of the design parameters lead to apparently weaker cryptosystems. [1]

There are several natural ways to vary the design of Skipjack without affecting its basic 'look and feel' too much:

1. We can vary the relationship between Skipjack's two round-types.
2. We can vary the ordering of the two types of rounds used in Skipjack.

Each of these design elements allows for many plausible choices of parameters, and each design element can be varied independently. In this way, we get a large space of Skipjack-variants. In each case, we show that changing the definition of Skipjack can lead to certificational weaknesses not present in the real Skipjack. This suggests that any simple change to Skipjack risks introducing new vulnerabilities, and in particular that Skipjack may have been carefully chosen as the optimal candidate from this family of ciphers.

Of course, there are many other possible variations of Skipjack, e.g., we can vary the way the round counter is injected into each round, and we can vary the key schedule and the S-box used. However, the attacks we present in this paper on variants of Skipjack would work more or less with unchanged complexities for any of these other modifications.

This paper is organized as follows. First, Section 2 describes Skipjack and introduces some other important background material. Section 3 examines the relationship between the A- and B-rounds. Then in Section 4, we study the round ordering of Skipjack and show that Skipjack's ordering seems to be better than many natural alternatives. Next, we consider the importance of the round counter in Section 5. Afterwards, we sum up and discuss our results in Section 6.

2. Background

Skipjack is a block cipher that supports a 64-bit block size and a 80-bit key. The block is internally divided into four 16-bit words, where each round applies a keyed non-linear permutation to one word from the block.

Skipjack uses two different types of round functions, the A-rounds and the B-rounds. Each encryption consists of a total of 32 rounds, applied in a specific order: first we apply 8 A-rounds, then 8 B-rounds, then another 8 A-rounds, and finally we finish with 8 more B-rounds. We repeat the definitions of the A-rounds and the B-rounds here for

[1] Of course, we cannot prove that Skipjack is strong, but so far it appears to resist attacks where its cousins cannot.

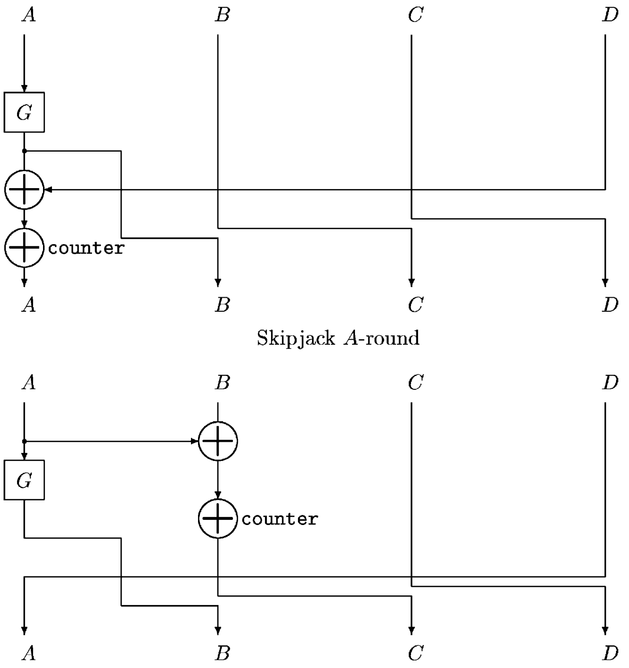

Skipjack A-round

Skipjack B-round

Fig. 1. The two rounds used in Skipjack.

convenience:

$$A(a,b,c,d) = (d + G_k(a) + \text{counter}, G_k(a), b, c),$$
$$B(a,b,c,d) = (d, G_k(a), a + b + \text{counter}, c).$$

$$A^{-1}(a,b,c,d) = (G_k^{-1}(b), c, d, a + b + \text{counter}),$$
$$B^{-1}(a,b,c,d) = (G_k^{-1}(b), c + G_k^{-1}(b) + \text{counter}, d, a).$$

In this paper, '$+$' denotes the bitwise exclusive-or operation, and 'counter' stands for a round counter that starts at 1 and counts up to 32. See Fig. 1 for an illustrated version.

The G_k box takes a 16-bit input and a 4-byte subkey k, and is itself a 4-round Feistel cipher using in each round a byte permutation F and a key-byte k_i. It follows that the inverse of G is equal to G itself using the subkeys in reverse order, except for a swap of the halves of both the input and the output. All of the remarks in this paper are independent of the specific choice of the byte permutation F and are instead directed at the high-level structure of Skipjack.

The Skipjack key schedule specifies how each 32-bit round subkey k is derived from the 80-bit key K. The 80-bit key K is split up into ten bytes, denoted K_0 through K_9. Then the key bytes are repeated in rotating order, and each round takes the next 32 bits for use as its subkey. For example, the first round uses $K_0,\ldots,K_3$, the second round uses $K_4,\ldots,K_7$, and the third round uses K_8,K_9,K_0,K_1.

We assume that the reader has some passing familiarity with the fundamentals of truncated differential cryptanalysis [9,10], but we will briefly review the basic ideas here. In a differential attack, the attacker chooses two plaintexts with a specified difference between them and attempts to predict with some probability how the difference will evolve during the encryption process. The defining feature of a *truncated* differential attack is that the cryptanalyst predicts the difference for only some portion of the block, leaving the remainder of the difference unpredicted. For this paper, the difference of two texts T, T' will refer to the exclusive-or $T + T'$ of their values.

In the analysis of Skipjack, it is natural to break the block into four 16-bit words and to predict the difference of some of those words. In this paper, our predictions will specify for each word whether that word should be zero or non-zero; note that we usually do not attempt to distinguish between the $2^{16} - 1$ non-zero values. Thus, we might use the notation $(a, b, 0, c)$ to refer to a difference that is predicted to be zero in its third word, and write $(0, 0, a, 0) \xrightarrow{1r_A} (0, 0, 0, a)$ for a 1-round truncated differential which predicts that the output difference after 1 A-round will take the form $(0, 0, 0, a)$ when the input difference has the form $(0, 0, a, 0)$.

With this notation, we easily obtain some trivial truncated differentials of probability one, such as $(a, 0, 0, 0) \xrightarrow{1r_A} (b, b, 0, 0)$ and $(0, a, 0, 0) \xrightarrow{4r_A} (b, b, 0, 0)$ and $(a, a, 0, 0) \xrightarrow{4r_B} (b, 0, 0, 0)$, as well as some non-trivial truncated differentials, e.g. $(a, 0, 0, b) \xrightarrow{1r_A} (0, b, 0, 0)$ which holds with probability $1/(2^{16} - 1)$ when a, b represent arbitrary non-zero 16-bit values. Also, the r-round differential $\delta \to \delta'$ may be concatenated with a s-round differential $\gamma \to \gamma'$ when the intermediate differences agree (i.e., $\delta' = \gamma$), and in this case we obtain a $(r + s)$-round differential $\delta \to \gamma'$ whose probability is the product of the original differentials' probabilities.

If we are able to find a truncated differential of sufficiently large probability that covers all 32 rounds of Skipjack, we may be able to distinguish Skipjack from an ideal cipher. Moreover, a good truncated differential for all but the last round often allows us to recover key material by guessing the last-round subkey and checking for *right pairs* that follow the differential.

The technique of structures is useful in many truncated differential attacks since it allows us to generate large pools of candidate pairs using a relatively small number of chosen texts. For instance, we can obtain 2^{2n-1} pairs of texts with input difference of the form $(0, a, b, c)$ by requesting the encryption of 2^n chosen plaintexts whose first word is fixed (for $n \leqslant 48$). These 2^n texts are collectively called a *structure*, and if we want more than 2^{95} pairs we may obtain many structures by using different values of the first word for each structure.

Similarly, many techniques are known for efficiently identifying right pairs and for recovering key material from them. See the literature [9,10] for more details.

3. The relationship between A- and B-rounds

The Skipjack A- and B-rounds are closely related in their internal structure: the structure of a B-round is almost the structure of the inverse of an A-round, differing only by a word-wise swap before and after. However, it may not be immediately obvious why this particular relationship was chosen in Skipjack or that this choice has important implications for the security of Skipjack. In this section, we explore this subject more closely, develop several apparent design principles, and show why all the other choices lead to serious weaknesses in the cipher.

3.1. The equivalence of encryption and decryption

The encryption and decryption functions of Skipjack may be written as

$$E_K(x) = B^8 \circ A^8 \circ B^8 \circ A^8(x), \tag{1}$$

$$D_K(y) = (A^{-1})^8 \circ (B^{-1})^8 \circ (A^{-1})^8 \circ (B^{-1})^8(y). \tag{2}$$

Beware: the symbol A should be thought of not as a specific transformation — because of the round constants and the key schedule, each keyed instance of the cipher will apply different transformations in different rounds — but rather as a shorthand for the high-level structure of the cipher. In other words, (1) should be interpreted as saying that the structure of Skipjack encryption uses 8 A-rounds followed by 8 B-rounds, and so on.

Let A^{-1} denote decryption in an A-round. Define $\mu(a,b,c,d) = (b,a,d,c)$, or equivalently, in cycle notation, $\mu = (1\,2)(3\,4)$. Note that $\mu = (1\,2)(3\,4)$ is an involution, so that $\mu^{-1} = \mu$. Then a B-round can be expressed in terms of μ and the inverse of an A-round,

$$B(a,b,c,d) = (\mu \circ A^{-1} \circ \mu)(a,b,c,d).$$

(Beware: the G_k box is not an involution, so the symbol A^{-1} in the formula should be thought of as a shorthand reference to the structure of the inverse A-round, but using G_k rather than G_k^{-1}.) In other words, encryption through a B-round may be implemented as decryption through an A-round preceded and followed by a μ-swap (except for inverting G_k). The encryption function of Skipjack may then be written as

$$\begin{aligned} E_K &= (\mu \circ A^{-1} \circ \mu)^8 \circ A^8 \circ (\mu \circ A^{-1} \circ \mu)^8 \circ A^8 \\ &= \mu \circ (A^{-1})^8 \circ \mu \circ A^8 \circ \mu \circ (A^{-1})^8 \circ \mu \circ A^8. \end{aligned}$$

In this notation the decryption function of Skipjack can be written as

$$D_K = (A^{-1})^8 \circ \mu \circ A^8 \circ \mu \circ (A^{-1})^8 \circ \mu \circ A^8 \circ \mu.$$

Note the strong similarity: the structure of encryption and decryption differ only by a μ-swap before and after the cipher.

Thus, the design of the A- and B-rounds seems to be an implementation feature. (See also Section 4 for another example of how this implementation consideration bears on the round ordering.) As a result, we get the following observation.

Fact 1. *The decryption function in Skipjack has exactly the same structure as the encryption function except for the application of a byte-wise swap in both the plaintext and the ciphertext.*

This property also has the following obvious implication.

Fact 2. *With uniformly chosen keys, Skipjack has equal security against chosen plaintext attacks and chosen ciphertext attacks.*

These two properties seem to be an important design principle for Skipjack-like ciphers, and in particular this illustrates an important facet of the relationship between the A- and B-rounds in Skipjack.

3.2. The choice of μ

The above-described similarity between encryption and decryption in Skipjack would also be valid for any other involution μ, including the identity function. In the following we give a possible explanation why $\mu = (1\,2)(3\,4)$ was chosen in the official Skipjack cipher.

To understand the implications of other choices for μ, we must look at the interaction between A- and B-rounds. Recall that $B = \mu \circ A^{-1} \circ \mu$. Clearly, an A-round followed by an inverse A-round is unfortunate, since three of four words will be left unencrypted (through G) in two rounds of encryption. This suggests that taking $\mu = \mathrm{id}$ could significantly weaken the cipher and thus would be a poor choice.

Good choices for μ. To examine the good choices of involutions μ, consider four consecutive rounds. If all rounds are either A-rounds or inverse A-rounds, all four words input to the four rounds will have been encrypted (through G). Therefore, consider the following choices of four rounds with a transition from A-rounds to inverse A-rounds:

1. $A^{-1} \circ A^{-1} \circ \mu \circ A \circ A$,
2. $A^{-1} \circ \mu \circ A \circ A \circ A$,
3. $A^{-1} \circ A^{-1} \circ A^{-1} \circ \mu \circ A$.

Since 2 and 3 are the inverses of each other it suffices to consider cases 1 and 2. Let the four input words be (a, b, c, d) and find the possible involutions μ, such that after four rounds of encryption, in each of the cases 1. and 2., none of a, b, c, d appears unencrypted (through G). It is easily checked that after two A-rounds the input words b and c have not yet been input to the G-function. With e.g., $\mu = \mathrm{id}$ it is also easily checked that these two words are unencrypted after an additional two inverse A-rounds. It turns out that for 1, the only possible involutions with the desired property are $(1\,2)(3\,4)$ and $(1\,3)(2\,4)$. However, for the latter, case 2 leaves the input word b, unencrypted. And indeed, the only μ which has the desired property is $(1\,2)(3\,4)$, as chosen in Skipjack. Any other choice of μ would weaken the cipher by introducing a 4-round subsequence that leaves one input word unencrypted under G.

We conclude that the involution μ was presumably introduced to strengthen Skipjack and to minimize the negative interaction of A and A^{-1}.

Still, we are yet to answer the question of why use 8 A-rounds followed by 8 B-rounds and not, e.g., 4 A-rounds followed by 4 B-rounds. This question is answered in the following section.

Table 1
Attacks on some Skipjack variants with a different round ordering[a]

Round ordering	Rounds	Attack complexity Texts	Time
$(A,B)^*$	∞	2 CP or $2^{16.5}$ KP	—
$(2A,2B)^*$	∞	2 CP or $2^{16.5}$ KP	—
$(3A,3B)^*$	32	2^{33} CP	2^{33}
$(3A,4B)^*$	∞	2 CP or 2^8 KP	—
$(4A,3B)^*$	∞	2 CPor 2^8 KP	—
$(16A,16B)^*$	28	2^{56} KP	2^{56}

[a] Notation: KP = known plaintexts, CP = chosen plaintexts.

4. The round ordering: why $(8A, 8B)$?

In this section the interaction between A-rounds and B-rounds is closely examined. The structure of Skipjack variants using n A-rounds followed by m B-rounds will be referred to as an $(nA, mB)^*$ structure.

We begin by introducing a simple design goal:

Design principle 1. *Encryption and decryption should be symmetrical.*

This goal is clearly desirable, because it simplifies implementation efforts and ensures equivalent security against chosen plaintext and chosen ciphertext attacks; see Section 3.1. The above design principle immediately has the following implication.

Corollary 1. *One should consider only the symmetrical structures* $(nA, mB)^*$, *i.e., those where* $n = m$.

This design principle allows us to narrow down the search space by eliminating all structures with $n \neq m$. Moreover, it explains why the official Skipjack cipher uses a $(nA, mB)^*$ round structure where $n = m$: any other choice would forfeit the advantages of a symmetrical cipher. Still, it does not explain why $n = 8$ is chosen in the official Skipjack standard, or help us understand whether other round orderings might not improve on the standard construction.

In this section, we examine several natural alternative round orderings for Skipjack and show that many of them allow for better attacks than those available in the real Skipjack construction. For the remaining constructions, truncated differentials have been identified, which have much higher probabilities than for real Skipjack. It is left as an open question to apply these differentials in cryptanalytic attacks. See also Table 1 for a summary of our attacks on the variants with a different round ordering.

We start by examining why it is better to start with A-rounds and end with B-rounds; this reduces the search space by a factor of two. Then, we eliminate most of the remaining round orderings by exhibiting truncated differential attacks.

In the following we write $G_k(\cdot)$ for every application of G, where the subscript k is fixed to indicate that the function is keyed, and we omit the round counter.

4.1. Why A before B?

One natural question when looking at Skipjack is why use the A-rounds before the B-rounds. Does the ordering make a difference to the security of Skipjack? We argue that the answer is yes, that Skipjack's ordering is preferable to the alternative, and that this is a crucial design principle with important ramifications on security.

Diffusion considerations hint at why this might be so. A-rounds exhibit better diffusion (in the encryption direction) than B-rounds, so it is natural to suspect that it might be better to start with the best diffusion available. In the following, we show how truncated differential cryptanalysis may be used to develop more persuasive evidence for this design principle.

We also note that an earlier work [10] already shows some indications of this design principle. This paper analyzed two 16-round Skipjack variants, a $(8A, 8B)$ variant and a $(8B, 8A)$ variant, finding attacks on the former with 2^{17} chosen plaintexts and 2^{34} work, while the latter can be broken with just 2–3 chosen texts and 2^{46}–2^{29} work. The dramatic difference in security between the two variants suggests that the order of the rounds may be very important.

A deeper inspection of [10] reveals strong evidence that reversing the ordering of the A- and B-rounds can be harmful in two important ways:

Starting with B-rounds makes it easier to find long truncated differentials The authors of [10] found a 12-round truncated differential of probability 1 for the variant starting with B-rounds, whereas the first 12 rounds of the differential used to analyze the variant starting with A-rounds has probability 2^{-32}. In general, B-rounds exhibit poor diffusion in the forward direction, and thus it is easier to pass through B-rounds with a truncated differential.

Ending with A-rounds makes it easier to use these truncated differentials in a key-recovery attack. A second factor that made the attack on $(8B, 8A)$ so efficient is the relative ease of peeling off the final rounds of a cipher which ends with A-rounds. That attack may be viewed as a 7-R attack which needs to guess a surprisingly small amount of key material to be able to look deep inside the cipher. (Compare to the $(8A, 8B)$ cipher, where only a 1-R analysis was possible because of the quantity of key material that must be guessed to peel off more than one round [10].) In general, this property is to be expected: A-rounds exhibit poor diffusion in the reverse direction, so peeling off a few final A-rounds is not difficult.

All in all, the effect is that shorter truncated differentials may be used when analyzing a cipher that ends with A-rounds; this, of course, usually increases the efficiency of the attack.

For Skipjack-like constructions, we conclude that starting with A-rounds and ending with B-rounds is likely to be stronger than the alternative.

Additional evidence for this proposition may be found below, where we show that the $(8B,8A)^2$ Skipjack-variant is more vulnerable to truncated differential cryptanalysis than the standard $(8A,8B)^2$ version.

4.2. Truncated differentials

Structure $(A,B)^*$**.** Consider an A-round followed by a B-round. The encryption function over these two rounds is (omitting the counter)

$$f(a,b,c,d) = (c, G_k(G_k(a)+d), d, b).$$

As seen, the two outer output words depend only on the two middle input words, and the two middle output words depend only on the two outer input words. This means that the variant $(A,B)^*$ is very weak, since there are truncated differentials of probability 1 for any number of rounds. The following 4-round truncated differential holds with probability one and can be iterated to any number of rounds:

$$(a,0,0,b) \overset{1r_A 1r_B}{\rightarrow} (0,c,d,0) \overset{1r_A 1r_B}{\rightarrow} (e,0,0,f),$$

where $(a,b) \neq (0,0), (c,d) \neq (0,0)$, and $(e,f) \neq (0,0)$. In other words, the $(A,B)^*$ cipher exhibits poor diffusion: changing any of the middle bits of the plaintext does not affect the outer bits of the ciphertext.

As a result, we can distinguish the $(A,B)^*$ cipher from a random permutation with two chosen texts or with $2^{16.5}$ known texts. Moreover, this result holds for any number of rounds.

Structure $(2A,2B)^*$**.** Consider 2 A-rounds followed by 2 B-rounds. The encryption function over these four rounds is (omitting the counter)

$$f(a,b,c,d) = (G_k(a), G_k(b), h(a,b,c,d), c),$$

where h is some function dependent on all four input words. It is seen, that the first two words of the output do not depend on the third and fourth words of the input. Thus in this variant, there exist 4-round iterative truncated differentials with probability one:

$$(0,0,a,b) \overset{2r_A 2r_B}{\rightarrow} (0,0,c,d),$$

where $(a,b) \neq (0,0)$ and $(c,d) \neq (0,0)$.

One consequence is that the $(2A,2B)^*$ cipher — with any number of rounds — can be broken with two chosen texts or with $2^{16.5}$ known texts.

As in the case of the $(A,B)^*$ cipher, this weakness may also be viewed as a simple diffusion failure. However, we will see shortly that truncated differentials provide a more sensitive measure of diffusion than is available with conventional techniques, so we will find it useful to apply the machinery of truncated differential cryptanalysis throughout.

Structure $(3A,3B)^*$**.** There are 21-round truncated differentials of probability one and there are impossible differentials for at least 30 rounds. The following differential has probability zero:

$$(0,a,0,0)\overset{(3r_A3r_B)^5}{\rightarrow}(0,0,0,b).$$

This differential was found by concatenating two 15-round differentials, each of probability one. A pair of plaintexts with difference $(0,a,0,0)$ leads to a difference in the ciphertexts after 15 rounds of $(c,d,e,0)$. Similarly, a pair of ciphertexts of difference $(0,0,0,b)$ decrypts back in 15 rounds to ciphertexts of difference $(f,g,0,h)$, where $h\neq 0$. Thus, there is a miss-in-the-middle [3] and a differential of probability 0 over 30 rounds. This situation is much better for the attacker than in the case of the original Skipjack, where only 24 rounds can be covered with such differentials [3]. In addition, the 30-round differential may be concatenated with a truncated differential of probability one over 2 A-rounds: $(0,0,0,b)\overset{2r_A}{\rightarrow}(c,c,0,0)$, where $c\neq 0$. In total, this yields a 32-round differential of probability zero, and thus, 32 rounds of $(3A,3B)^*$ can be distinguished from random with about 2^{33} chosen plaintexts. Using structures of 2^{16} plaintexts different in the second words only, one can form about 2^{31} pairs with the difference $(0,a,0,0)$. With 2^{17} such structures one obtains totally 2^{48} pairs of plaintexts. Using this Skipjack variant, the resulting pairs of ciphertexts will never have the difference $(c,c,0,0)$, whereas for a randomly chosen permutation one expects one such pair.

Structure $(3A,4B)^*$**.** Although we have argued already that one should use only structures $(nA,mB)^*$ with $n=m$, we include this variant because it is extremely weak. There is a 7-round iterative, truncated differential of probability one:

$$(0,a,b,c)\overset{3r_A4r_B}{\rightarrow}(0,d,e,f).$$

As a result, any number of rounds of the $(3A,4B)^*$ construction can be distinguished from a randomly chosen permutation with two chosen texts or with $2^{8.5}$ known texts. Because of the symmetry of A- and B-rounds (see Section 3) a similar result holds for $(4B,3A)^*$.

Structure $(4A,4B)^*$**.** In this variant there exist 31-round truncated differentials with probability 2^{-48}. For example,

$$(0,a,b,c)\overset{4r_A4r_B}{\rightarrow}(d,0,e,f)\overset{4r_A4r_B}{\rightarrow}(0,g,h,i)\overset{4r_A4r_B}{\rightarrow}(j,0,k,m)\overset{4r_A3r_B}{\rightarrow}(n,0,p,q).$$

The first and third 8 rounds have probabilities one. The second 8 rounds have a probability of 2^{-32} and the last 7 rounds a probability of 2^{-16}. There are several truncated differentials of this form, so the probability is higher than stated.

The probabilities of the found differentials are much higher than for similar differentials found for (unmodified) Skipjack. This provides evidence that Skipjack's $(8A,8B)^*$ is preferable to the $(4A,4B)^*$ alternative.

Structure $(8B, 8A)^2$**.** There is a 28-round truncated differential with probability 2^{-32}.

$$(0, a, 0, 0) \overset{8r_B}{\to} (0, b, c, a) \overset{8r_A}{\to} (d, d, e, f) \overset{8r_B}{\to} (0, e, f, g) \overset{4r_A}{\to} (h, h, i, j).$$

The first 8 rounds and the last 4 rounds have probabilities 1. The second and third 8 rounds have each a probability of 2^{-16}.

We have already seen theoretical reasons to be very skeptical of ciphers that start with B-rounds or end with A-rounds. The above differential has a probability higher than what have been identified for real Skipjack. This provides additional evidence that starting with B-rounds or ending with A-rounds may reduce security.

Structure $(16A, 16B)^1$**.** This is the only remaining natural alternative to $(8A, 8B)^2$. Truncated differential cryptanalysis does not seem as powerful against the $(16A, 16B)^1$ structure as for the other alternatives, and we do not know of any compelling attacks against the full 32-round $(16A, 16B)$ cipher.

However, the middle 28 rounds of $(16A, 16B)$ may be distinguished from a random cipher using *known* plaintexts and truncated differentials. By easy calculations one finds the following truncated differential of probability one:

$$(a, a, 0, 0) \overset{14r_B}{\to} (b, c, d, e), \quad c \neq 0.$$

Because of the symmetry between A- and B-rounds there is a similar differential for decryption through 14 A-rounds. With 2^{56} known texts, we get 2^{111} pairs, of which 2^{63} will have the difference $(a, a, 0, 0)$ in the middle. We filter, looking for pairs whose difference has second ciphertext word and first plaintext word both non-zero. For a random cipher, we expect $n = (1 - 2^{-16})^2 \times 2^{111}$ such pairs, with standard deviation about 2^{48}. For the Skipjack-variant, we expect $(1 - 2^{-16})^2 \times (2^{111} - 2^{63}) + 2^{63} \approx n + 2^{48}$ such pairs, which is about one standard deviation above the mean for a random cipher.

Therefore, with 2^{56} known texts and very little computation we may distinguish the middle 28 rounds of $(16A, 16B)$ from a randomly chosen permutation. This does not prove that the $(16A, 16B)^1$ structure is weaker than $(8A, 8B)^2$, but it may indicate a perhaps-undesirable property of the $(16A, 16B)$ structure.

5. The round counter and the key schedule

The key-schedule of Skipjack is remarkably simple when compared to the key-schedules of other modern block ciphers, e.g., the AES candidates. The five candidates of the final round of the AES process have complex key-schedules and the round keys are often encrypted. In Skipjack the round key bytes are computed by a simple rotation of the input key bytes. The above Facts 1 and 2 look like design principles of Skipjack and could also explain the role of the counter. Without the counter there will be keys for which encryption and decryption will be very similar (with the exception of an application of μ to both plaintext and ciphertext), and pairs of keys for which encryption with one key is very similar to decryption with another key [2]. Such keys are present in the DES [11] and regarded as weak keys. With the counter, these problems seem to disappear.

6. Discussion

These observations suggest several general principles for the design of Skipjack-like ciphers:

Seek symmetry. A cipher is *symmetrical* if encryption and decryption have the same structure. Symmetry is a beneficial implementation property which simplifies the task of implementation and which may reduce resource requirements, since one can support both encryption and decryption at once with just one cryptographic engine. Also, a symmetrical cipher has identical resistance to chosen-plaintext and chosen-ciphertext attacks, which is a desirable security property. This is not a new observation — see, e.g., [1,6,15] — but it does not seem to be as widely known as it perhaps could be [8,16].

Avoid too much symmetry. Of course, symmetry sometimes allows for clever cryptanalytic attacks. In Skipjack, symmetry is broken with round counters, and this helps avoid, e.g., complementation properties and slide attacks. While this design principle may seem to contradict the previous one, the conflict is not so bad as it may first appear. The Skipjack round counters appear to prevent most attacks that attempt to exploit their symmetry, while retaining the equivalence between chosen-plaintext and chosen-ciphertext security against a very large class of attacks (namely, the class of all attacks that ignore the round counters — and note that this class includes differential and linear cryptanalysis as well as their variants).

Minimize bad interactions between the round-types. We have seen considerable evidence that the A- and B-rounds interact poorly where they are applied in consecutive rounds. In general, transitions between round-types appear to reduce security. Moreover, if μ is not chosen carefully, this bad effect can be much worse than necessary.

We expect that these principles will also be applicable to other ciphers with similar structure, including FROG [8] and CAST-256 [1].

We have also answered a number of detailed questions about the design and structure of Skipjack. We summarize some of those contributions here:

Why two types of rounds? The reason Skipjack does not use just A-rounds or B-rounds is that both round types exhibit considerable asymmetry: for example, A-rounds have excellent diffusion in the forward (encryption) direction, but very poor diffusion in the reverse (decryption) direction; B-rounds behave dually.

Why are B-rounds a near-inverse of A-rounds? This arrangement allows us to build a symmetrical cipher out of asymmetrical round functions: if $B{=}A^{-1}$, then the $(nA,nB)^*$ structure is self-inverse. More generally, if $B{=}\mu^{-1}\circ A^{-1}\circ\mu$ where $\mu{=}\mu^{-1}$, the resulting construction is also self-inverse (up a bit-permutation before and after the cipher). This also explains why μ is an involution in Skipjack.

Why use $\mu=(1\,2)(3\,4)$? Skipjack uses a word permutation which may be expressed in cycle notation as $\mu=(1\,2)(3\,4)$. It turns out that this is the unique permutation which maximizes security and minimizes the bad interaction at the boundary between A- and B-rounds. See Section 3.

Why apply A*-rounds before* B*-rounds*? This ordering makes it much harder to peek deep inside the cipher. With B-rounds before A-rounds, one may mount a 7-R attack and look as far as 7 rounds, into the cipher, without too much effort; but applying A-rounds first makes it very expensive to look more than one round deep (i.e., 2-R attacks are too costly). Also, the $(8B, 8A)^*$ structure is weaker against truncated differential cryptanalysis than the real $(8A, 8B)^*$ cipher. See Section 4.

Why use $(8A, 8B)^*$? All of the alternatives we have examined appear to weaken the security of Skipjack. In particular, the (nA, mB) constructions with $n, m < 8$ are especially vulnerable to truncated differential cryptanalysis; and the $(nA, mB)^*$ constructions with $n, m > 8$ allow the cryptanalyst to find better truncated differentials for each half of the cipher. Also note that symmetry considerations suggest we should use $n = m$.

Why use a round counter? Eliminating the round counter from Skipjack introduces complementation properties that can be used to speed up exhaustive keysearch [2,3]. Also, constructions without round counters appear to be more susceptible to slide attacks [4].

Why use a 80-bit key? With a longer key, differential-style attacks would have a lower complexity than exhaustive keysearch [3], so it is natural to choose 80 bits as an upper bound on the effective keylength of the algorithm. At the same time, keys much shorter than 80 bits are too weak to resist exhaustive keysearch for long [5].

Our ‘explanations’ of why Skipjack looks the way it does, must (of course) be categorized as unproven conjectures. Nonetheless, we feel that the body of the evidence supports our view.

References

[1] C. Adams, The CAST-256 Encryption Algorithm, Submitted as an AES Candidate, Available at `http://csrc.nist.gov/encryption/aes`.

[2] E. Biham, A. Biryukov, O. Dunkelman, E. Richardson, A. Shamir, Initial Observations on Skipjack: Cryptanalysis of Skipjack-3XOR, in: S. Tavares, H. Meijer (Eds.), Selected Areas in Cryptography, 5th Annual International Workshop, Springer, Berlin, 1998, pp. 362–375. Also available at `http://www.cs.technion.ac.il/~biham/Reports/SkipJack/`.

[3] E. Biham, A. Biryukov, A. Shamir, Cryptanalysis of Skipjack reduced to 31 rounds using impossible differentials, in: J. Stern (Ed.), Advances in Cryptology — Eurocrypt ’99, Lecture Notes in Computer Science, Vol. 1592, Springer, Berlin, 1999, pp. 12–23. Also available at `http://www.cs.technion.ac.il/~biham/Reports/SkipJack/`.

[4] A. Biryukov, D. Wagner, Slide attacks, in: L.R. Knudsen (Ed.), Fast Software Encryption, Sixth International Workshop, Rome, Italy, March 1999, Lecture Notes in Computer Science, Vol. 1636, Springer, Berlin, 1999, pp. 245–259.

[5] M. Blaze, W. Diffie, R. Rivest, B. Schneier, T. Shimomura, E. Thompson, M. Wiener, Minimal Key Lengths for Symmetric Ciphers to Provide Adequate Commercial Security, January, 1996.

[6] C. Burwick, D. Coppersmith, E. D’Avignon, R. Gennaro, S. Halevi, C. Jutla, S. Matyas Jr., L. O’Connor, M. Peyravian, D. Safford, N. Zunic. MARS — A Candidate Cipher for AES, Submitted as an AES Candidate. Available at `http://csrc.nist.gov/encryption/aes`.

[7] W. Diffie, S. Landau, Privacy on the Line, MIT Press, 1998.

[8] D. Georgoudis, D. Leroux, B.S. Chaves, The “FROG” Encryption Algorithm, submitted as an AES Candidate, Available at http://csrc.nist.gov/encryption/aes.

[9] L.R. Knudsen, Truncated and higher order differentials, in: B. Preneel (Ed.), Fast Software Encryption — Second International Workshop, Leuven, Belgium, Lecture Notes in Computer Science, Vol. 1008, Springer, Berlin, 1995, pp. 196–211.

[10] L.R. Knudsen, M.J.B. Robshaw, D. Wagner, Truncated differentials and Skipjack, in: M. Wiener (Ed.), Advances in Cryptology: CRYPTO’99, Lecture Notes in Computer Science, Vol. 1666, Springer, Berlin, 1999, pp. 165–180.

[11] National Bureau of Standards, Data encryption standard. Federal Information Processing Standard (FIPS), Publication 46, National Bureau of Standards, U.S. Department of Commerce, Washington DC, January 1977.

[12] National Security Agency, NSA Releases Fortezza Algorithms. Press Release, 24 June 1998, Available at http://csrc.ncsl.nist.gov/encryption/nsa-press.pdf.

[13] National Security Agency, Skipjack and KEA algorithm specifications. May 1998. Available at http://csrc.ncsl.nist.gov/encryption/skipjack-1.pdf.

[14] B. Schneier, D. Banisar, The Electronic Privacy Papers, Wiley, New York, 1997.

[15] D. Wagner, N. Ferguson, B. Schneier, Cryptanalysis of FROG, Second AES Candidate Workshop, 1999.

[16] G. Yuval, Reinventing the Travois: Encryption/MAC in 30 ROM Bytes, in: E. Biham (Ed.), Fast Software Encryption, 4th International Workshop Proceedings, Lecture Notes in Computer Science, Vol. 1267, Springer, Berlin, 1997, pp. 205–209.