SCIENCE CHINA
Information Sciences

• RESEARCH PAPER •

July 2014, Vol. 57 072110:1–072110:11
doi: 10.1007/s11432-014-5096-6

What is the effective key length for a block cipher: an attack on every practical block cipher

HUANG JiaLin & LAI XueJia*

Cryptography and Information Security Lab, Department of Computer Science, Shanghai Jiaotong University, Shanghai 200240*, China*

Received December 4, 2013; accepted March 6, 2014; published online May 5, 2014

Abstract Recently, several important block ciphers are considered to be broken by the brute-force-like cryptanalysis, with a time complexity faster than the exhaustive key search by going over the entire key space but performing less than a full encryption for each possible key. Motivated by this observation, we describe a meet-in-the-middle attack that can always be successfully mounted against any practical block ciphers with success probability one. The data complexity of this attack is the smallest according to the unicity distance. The time complexity can be written as $2^k(1-\epsilon)$, where $\epsilon > 0$ for all practical block ciphers. Previously, the security bound that is commonly accepted is the length k of the given master key. From our result we point out that actually this k-bit security is always overestimated and can never be reached because of the inevitable loss of the key bits. No amount of clever design can prevent it, but increments of the number of rounds can reduce this key loss as much as possible. We give more insight into the problem of the upper bound of effective key bits in block ciphers, and show a more accurate bound. A suggestion about the relationship between the key size and block size is given. That is, when the number of rounds is fixed, it is better to take a key size equal to the block size. Also, effective key bits of many well-known block ciphers are calculated and analyzed, which also confirms their lower security margins than thought before. The results in this article motivate us to reconsider the real complexity that a valid attack should compare to.

Keywords block cipher, effective key bits, meet-in-the-middle, brute-force attack

Citation Huang J L, Lai X J. What is the effective key length for a block cipher: an attack on every practical block cipher. Sci China Inf Sci, 2014, 57: 072110(11), doi: 10.1007/s11432-014-5096-6

1 Introduction

As one of the fundamental primitives in symmetric cryptography, block ciphers play an important role in today's secure communication. They protect data against unauthorized access and tampering in an insecure communication channel. Also, the design of many cryptographical schemes, such as secure encryption modes and authentication modes, is based on the security of block ciphers. Therefore, their security evaluation has been a hot research issue over the decades, giving rise to different analysis techniques. One line of research is the so-called provable security approach [1,2], such as indistinguishability analysis. This approach usually studies design principles or cipher structures by assuming the pseudo-randomness of some components. Another line of research focuses on the practical security, that is, if

*Corresponding author (email: lai-xj@cs.sjtu.edu.cn)

any cryptanalytic attacks can be mounted successfully on a block cipher, such as differential attacks, linear attacks, meet-in-the-middle attacks, related-key attacks, as well as other existing cryptanalysis techniques [3–5]. A block cipher is considered secure when it can resist against all known attacks. Traditionally, the strength of a cryptanalytic attack is measured by comparing it to the exhaustive search over the entire key space. Hence, the security of a block cipher highly relies on the key length.

Recently, the full version of AES has been called broken because of the biclique attack [6], which performs faster than the exhaustive search. In [7], the authors proposed a complex meet-in-the-middle attack on KASUMI using various subtle weaknesses of the cipher. In [8], the authors proposed several techniques to speed up the exhaustive key search on the Full IDEA (by combining the BD-relation), KASUMI, and GOST. All of the above attacks have the following in common: by going over the entire key space with performing less than a full encryption for each possible key, the full rounds of the ciphers are targeted with a time complexity slightly faster than the exhaustive key search (for AES-256 is $2^{254.4}$, for KASUMI is $2^{125.8}$, for IDEA is $2^{126.8}$). These results are far from being any threat to the use of ciphers in practice. However, they motivate us to consider the realistic complexity that an attack should be compared to. That is, in a real world context, what should the time complexity of a valid attack at the most be.

1.1 Related work

In [8], Biham et al. recalled two well-known techniques to marginally reduce the time complexity of the exhaustive key search for almost any block ciphers. One is the distributive technique, which extracts the key bits that are not used in the first (or last) few operations of the encryption process. Another is the early abort technique referred in [9], which is to discard a wrong key before computing the full ciphertext. Assume that a subset $K(1)$ of the key bits is not used in the first few operations, and a (possibly different) subset $K(2)$ is not used in the last few operations. Then, Biham et al. proposed a more advanced algorithm using the meet-in-the-middle technique, as follows.

For each value of the bits in $K \setminus K(1) \setminus K(2)$, perform the following:

1. For each value of the bits in $K(2) \setminus K(1)$, perform the first few operations of the encryption process for the given plaintext. Store the intermediate value and the corresponding value in $K(2) \setminus K(1)$ in a table.

2. For each value of the bits in $K(1) \setminus K(2)$, perform the last few operations in the decryption direction for the given ciphertext. Then, guess the value of the remaining bits in $K(2)$, and complete the rest of the computation up to the intermediate value. Check the match with the values in the table.

The above algorithm (called the Biham's algorithm in this article) is enhanced further with the splice-and-cut technique by considering the common key bits that are not used in the operations between the plaintext and a pre-chosen intermediate value and in the last few operations, at the cost of increasing the data complexity (we call this the splice-and-cut version of Biham's algorithm). Based on the cipher structures and weaknesses of the key schedules, Biham et al. showed the speedup for IDEA, GOST, and KASUMI.

1.2 Our contribution

Most block ciphers in common use are designed to have security equal to their key length (an exception is Triple-DES). Given that a key consists of k bits, the exhaustive search of the key space would take 2^k encryption[1)], with success probability of one when the number of plaintext–ciphertext pairs satisfies the unicity distance.

In this article, by giving a universal attack which has a time complexity of $2^k(1-\epsilon)$ where $\epsilon > 0$, we point out that the previously thought bound of the effective key size k can never be achieved for almost all practically used block ciphers. The data complexity of this attack is the smallest according to the unicity distance, and the success probability is about one. We present a formulated description, measuring the effective key length explicitly with some general parameters, such as block size, key size, and number of

1) Note that in the average case this complexity is 2^{k-1}, but in this article we have considered the worst case.

rounds. Also, our algorithm is applied to many well-known block ciphers and their effective key bits are calculated. As predicted, the effective key bits of these ciphers are all less than the master key size k.

Compared with previous work, our analysis is basing on more general structures and weaker assumptions, which have nothing to do with the specifics of key schedules. No matter how clever and secure a practical block cipher is, our algorithm is always available to the cryptanalyst. The data complexity of our algorithm reduces greatly. Only three instances are given for the splice-and-cut version of Biham's algorithm: IDEA, KASUMI, and GOST. This indicates that weak key schedules (all these three ciphers have simply linear key schedules) and large amount of data complexity are required for the attacks. No instances are given for the basic Biham's algorithm. We do a partial match in the middle, instead of using early abort technique in the ciphertext. More details, such as the computational complexity, and not just a rough description about the algorithm are presented. By the explicit quantization of the real bound of effective key lengths, the relationship between the key size and block size and the effect of increasing the number of rounds can be considered from a new point.

This article is organized as follows. In Section 2, we introduce the basic notations and construction of block ciphers. In Section 3, a generic attack is proposed and its computational complexities are studied. The upper bound of effective key bits is also investigated in this section. In Section 4, we give several widely used block ciphers as examples to show their effective key lengths. Section 5 discusses and concludes with our results.

2 The construction, notations, and conventions

Based on Shannon's conception of confusion and diffusion, most modern block ciphers have been designed to use many iterations of substitution (nonlinear layer) and permutation (linear layer) to obtain enough security (each iteration is referred to as one round). We give the following notations first.

- P: plaintext
- C: ciphertext
- n: the block size
- K: master key
- k: the master key size
- R: the number of rounds
- S: the nonlinear layer
- L: the linear layer
- K^r: the subkey used in round r, K_i^r is the ith sub-block in K^r
- X^r: the input block to round r where $X^0 = P$, X_i^r is the ith sub-block in X^r
- Y^r: the output block of the key mixing in round r, Y_i^r is the ith sub-block in Y^r
- Z^r: the output block of the nonlinear layer in round r, Z_i^r is the ith sub-block in Z^r

For almost all block ciphers used in practice, their R-round generic structure is depicted in Figure 1.

There can be more than one nonlinear or linear transformations in each round function. Usually the key mixing layer adds the subkey to the current state block using linear operations, such as XOR and modulo addition. Note that a round function in practice cannot be designed as random permutations.

For a block cipher with key size k, the easiest and universal attack an adversary can mount is to simply try and guess each possible key. The probability of correctly guessing the key at the first attempt is 2^{-k}. Adding an additional bit to the length of the key halves the probability that the key is correctly guessed. The time required to exhaust the whole key space is proportional to the time required to perform 2^k encryption operations.

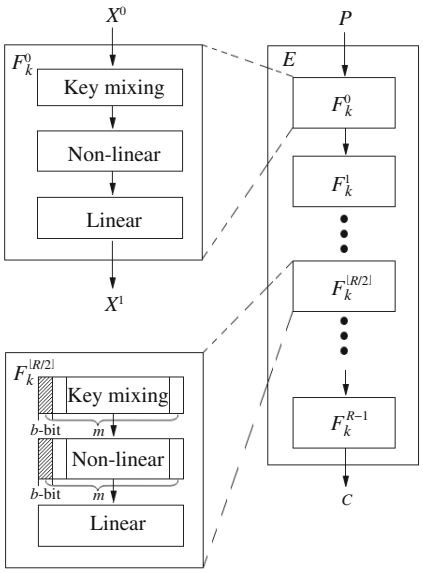

Figure 1 Structure of an R-round block cipher E.

3 A generic attack

We introduce a generic meet-in-the-middle attack that can be mounted on every practical block cipher. This attack is given in Algorithm 1.

S_1 is the internal state that can be calculated from P only with k_1 bits of subkeys, where k_1 is maximum smaller than k that can be obtained. Similarly, S_2 is the internal state that can be derived from C only with (other) k_1 bits of subkeys. For any block cipher, the states of S_1 and S_2 certainly can be found. This algorithm has two phases, the meet-in-the-middle phase that generates the candidate list containing 2^{k-M} keys where M is the size of the met intermediate value, and the check phase that examines the keys in the list.

For further discussion, we make two assumptions that are reasonable for practical block ciphers. First, nonlinear transformations are assumed to consume much more time than linear transformations. Hence, as in previous works, only nonlinear operations are counted [6,10]. Second, key schedules are assumed as negligible, since they are usually simpler than the encryption function.

Now we discuss the time complexity, data complexity, memory complexity, and success probability for Algorithm 1.

3.1 Time complexity

Based on the above assumptions, the time complexity is considered as follows. For any block ciphers, it is always smaller than 2^k.

$$\begin{aligned} T_{\text{comp}} &= 2^{k_1}\left(\frac{N_{P\to S_1}}{N_{\text{total}}}+\frac{N_{C\to S_2}}{N_{\text{total}}}+2^{k-k_1}\frac{N_{\text{total}}-N_{P\to S_1}-N_{C\to S_2}-N_{\text{disc}}}{N_{\text{total}}}\right) \\ &\quad + 2^{k-M}+2^{k-M-n}+2^{k-M-2n}+\cdots \qquad (1a) \\ &\approx 2^k\left(\frac{N_{\text{total}}-N_{P\to S_1}-N_{C\to S_2}-N_{\text{disc}}}{N_{\text{total}}}\right) \\ &= 2^k\left(1-\frac{N_{P\to S_1}+N_{C\to S_2}+N_{\text{disc}}}{N_{\text{total}}}\right), \qquad (1b) \end{aligned}$$

where N_{total} means the total nonlinear components required in a full encryption. Denote $N_{P\to S_1}$ as the required nonlinear components in the calculation from P to S_1. Denote $N_{C\to S_2}$ as the required nonlinear

Algorithm 1: the generic meet-in-the-middle attack

Data: $\lceil \frac{k}{n} \rceil + 1$ pairs of plaintext and ciphertext
Result: the output key K
for each value of the first k_1 key bits **do**
  Compute S_1 from P with these k_1 bits ;
  for each value of the remaining $k - k_1$ key bits **do**
    Compute $Z_0^{\lfloor \frac{R}{2} \rfloor}$ from S_1;
    Store $Z_0^{\lfloor \frac{R}{2} \rfloor}$ in a table corresponding to the guessed key;
  end
end
for each value of the last k_1 key bits **do**
  Compute S_2 from C with these k_1 bits ;
  for each value of the remaining $k - k_1$ key bits **do**
    Compute $Z_0^{\lfloor \frac{R}{2} \rfloor}$ from S_2;
    if $Z_0^{\lfloor \frac{R}{2} \rfloor}$ corresponding to the guessed key is in the table **then**
      add the guessed key into the candidate list;
      move onto the next guess;
    else
      move onto the next guess;
    end
  end
end
Check the keys in the candidate list with other $\lceil \frac{k}{n} \rceil$ plaintext–ciphertext;

components in the calculation from C to S_2. N_{disc} is the number of nonlinear components that do not need to be computed when partial matching techniques are used in the middle. The partial matching can filter M bits information of the key after the meet-in-the-middle phase. If $\frac{N_{P\to S_1}+N_{C\to S_2}+N_{\text{disc}}}{N_{\text{total}}}$ is written as ϵ, then (1b) is $2^k(1-\epsilon)$, where $\epsilon > 0$.

3.2 Data complexity

The required number of pairs of plaintext–ciphertext here is $U + 1$, where $U = \lceil \frac{k}{n} \rceil$ is the smallest data complexity according to the unicity distance. We use the first pair of data to filter parts of the wrong keys and generate the candidate list. Then, we require at most another U pairs of plaintext–ciphertext for finding the right key.

If we store the internal states before and after the meet-in-the-middle state, we can use the first data pair for filtering another $n - M$ bits. Now the first data also can provide all its n-bit information for checking, the same as other pairs of plaintext–ciphertext. Thus, the data complexity can be reduced to U. Since the data complexity now is $U + 1$, which is small enough, this tradeoff is unnecessary.

3.3 Memory complexity

Algorithm 1 has a memory complexity of $2^k \cdot M$ bits. If more memory can be sacrificed, the data complexity can be lowered as mentioned above. The time–memory tradeoff is not our concern here.

3.4 Success probability

In the meet-in-the-middle phase, a wrong key is eliminated with a probability of $1 - 2^{-M}$. On examining with the second data pair in the candidate list, a wrong key is discarded with a probability of $1-2^{-(M+n)}$. And on examining with the third data pair (if needed), a wrong key is eliminated with a probability of $1 - 2^{-(M+2n)}$, and so on. The success probability of Algorithm 1 is the product of these probabilities for all $2^k - 1$ wrong keys, which is approximately one.

Algorithm 1 is similar with Biham's algorithm, but has several differences. First, Biham's algorithm does not mention where the intermediate value is to meet. We explicitly claim that the meet position does not influence the complexities of Algorithm 1. Without loss of generality, we fix this value as some

sub-block in the middle round. Second, instead of aborting the evaluation after computing parts of the ciphertext, we partially match in the middle before computing the full intermediate state.

3.5 More information for specific structures of block ciphers

Eq. (1a) can be made more concrete for specific cipher structures. Usually there are two major structures for block ciphers, SPN, and Feistel structure, as well as their generalized variants and combinations. The SPN structure constitutes of a layer of keyed confusion (nonlinear operations such as S-boxes) and a layer of diffusion (linear transformation), such as Serpent, AES, and ARIA that is widely used now. This structure is a direct implementation of Shannon's confusion and diffusion concepts. By iterating the round function repeatedly, the dependency between inputs and outputs of the ciphers becomes complicated. Consider n-bit internal state $W = (W_0, W_1, \ldots, W_{m-1})$ as a concatenation of m b-bit words W_i, where b is the size of a nonlinear sub-block. For most SPN ciphers, every nonlinear sub-block is keyed, and we match a b-bit word in the middle. Hence, the time complexity is written as:

$$\begin{aligned} T_{\text{comp}} &= 2^{k_1}\left(2\frac{k/b-1}{\text{Rm}} + 2^{k-k_1}\frac{\text{Rm}-(m-1+2(k/b-1))}{\text{Rm}}\right) + 2^{k-b} + 2^{k-b-n} + 2^{k-b-2n} + \cdots \\ &\approx 2^k\left(1 - \frac{m-1+2(k/b-1)}{\text{Rm}}\right) \\ &= 2^k\left(1 - \frac{1-3b/n+2k/n}{R}\right), \end{aligned} \tag{2}$$

where $m > 1$ in practical block ciphers because of the limit of the size of one nonlinear operation, as well as $k > b$. For an entire encryption, there are Rm nonlinear operations. For the first $(k/b-1)$ operations, we always do not need to guess all k bits of the key. We can compute S_1 by only searching the first $(k-b)$ bits, without guessing the remaining key bits. The time complexity of a factor of $(k/b-1)$ is saved here. The multiplication of 2 means that the computation from both the plaintext and the ciphertext should be considered. In the middle round, using the partial matching technique, we can only compute one nonlinear operation to get a b-bit filter and save another $(m-1)$ operation.

Another primary structure is Feistel, with the input to each round divided into two halves. One half is transformed by some nonlinear round function and then XORed to the other half. Then these two halves are swapped except for the last round. For the Feistel structure, time complexity can be derived in the same way and the resulting formula is very similar. Because of the half diffusion property, at least one round of computation can be saved when matching in the middle. For other more detailed structures, such as MISTY (note that it has different sizes for nonlinear components, 7 to 7 bits and 9 to 9 bits S-boxes) and Lai-Massey structures, we give examples directly in Section 4.

For the block ciphers, the meet-in-the-middle attack proposed in this article requires that the subkeys affect the round transformation with a separable pattern between different sub-blocks. That is, parts of subkeys directly act on parts of the internal state, e.g., K_i^r is mixed with X_i^r. If the round function is designed as random permutations, where the subkey can be regarded to act as a whole, then our attack will fail. Hence, our concern is of all the block ciphers existing in practice, which always satisfy this condition.

3.6 An upper bound of the effective key length

For any practical block cipher, Algorithm 1 is always available to the cryptanalyst. This indicates that a more accurate effective key length can be considered by taking the logarithm of the time complexity for this universal algorithm. For convenience, we can focus on (2) here, and the results for other structures are similar.

The effective key size is $k + \log(1 - \frac{1-3b/n+2k/n}{R})$, which is always smaller than k for any block ciphers. Usually the size of nonlinear sub-block b is much smaller than n and k, and only takes a few fixed values. For conventional block ciphers, the routine size of S-boxe is 4, 8, or 16 bits (for MISTY and KASUMI this is 7 or 9 bits). And for lightweight block ciphers, the routine size is 3 or 4 bits. $1 - 3b/n + 2k/n$

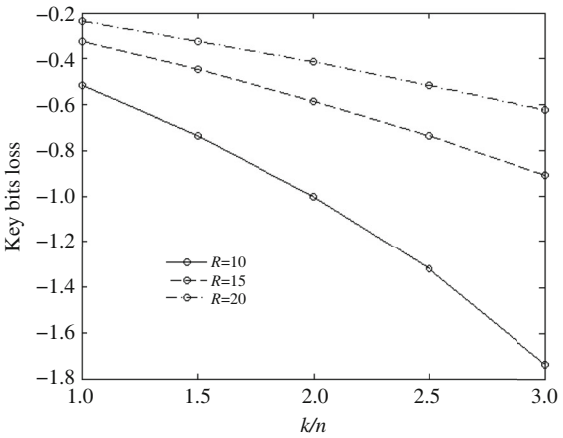

Figure 2 The relationship of key bits loss with key size, block size, and number of rounds.

is always larger than zero, and so $\log(1-\frac{1-3b/n+2k/n}{R})$ is smaller than zero. The previously accepted security bound of k key bits actually cannot be achieved.

Denote $\log(1-\frac{1-3b/n+2k/n}{R})$ as the loss of effective key bits. We ignore the factor of $3b/n$ and draw the function of $\log(1-\frac{1+2k/n}{R})$ with fixed different R (see Figure 2). When R is the same and larger the k/n, the more the key bits loss. When k/n is the same and the more rounds a cipher iterates, the lesser the key bits loss. Thus, we can avoid the loss of key bits as much as possible by increasing the number of rounds or shortening the ratio of k and n. This indicates what a relation of key size and block size should be in a secure design. Although it is not very precise, it still can be a rough guidance for most block ciphers. Also, from the formula we can conclude that when the number of rounds is sufficiently large, the key bits loss can approximately be zero.

4 Effective key lengths for block ciphers

In this section, many practical block ciphers are analyzed for their actual effective key lengths. For a clear exhibition of (1a), we consider conventional block ciphers and lightweight block ciphers, respectively[2)].

4.1 Conventional block ciphers

We take AES, which is currently the most widely used block cipher, as an example. It is selected as the new standard for replacing DES by NIST in 2000. As we assumed before, we count the computational complexity for S-boxes. This is also taken in [6]. For AES-128, the key size k and block size n are both 128 bits, the size of non-liner sub-blocks (S-box) b is 8, and the number of rounds R is 10. There are 16 sub-blocks in the state, that is $m=16$. Refer to [11] for more details about AES. Note that the whitening key K^0 should also be considered here. The detailed application of Algorithm 1 is as follows. Z_0^5 is the intermediate value for meeting. Choose $(Z_0^1, Z_1^1, \ldots, Z_{14}^1)$ as S_1. Compute S_1 from P by guessing $(K_0^0, K_1^0, \ldots, K_{14}^0)$, 120 bits totally. Then, for each guess of K_{15}^0, the last 8 bits of the master key K complete the encryption operations from S_1 to Z_0^5. This requires a calculation of Z_{15}^1, Z^2, Z^3, Z^4, Z_0^5, 50 S-boxes totally. Store Z_0^5 in a hash table corresponding to the guessed key. Choose $(X_1^{10}, X_2^{10}, \ldots, X_{15}^{10})$ as S_2. Compute S_2 from C by guessing $(K_1^{10}, K_2^{10}, \ldots, K_{15}^{10})$, 120 bits again. Then, for each guess of K_0^{10}, the last 8 bits of a mapping of the master key K complete the decryption operations from S_2 to Z_0^5. This needs a computation of X_0^{10}, X^9, X^8, X^7, X^6, 65 S-boxes totally. There are $10\times 16=160$ S-boxes for the full AES-128. Thus, for each guess of 128-bit of K only 115 S-boxes need to be computed. This means 115/160 of one full 10-round encryption for each guess. The time complexity of Algorithm 1 here is about $2^{128}\times 115/160=2^{127.5}$ (this value also can be directly derived from (2)). Thus, the effective key

2) Without specification, the notation in this section is as mentioned in Sections 2 and 3.

Table 1 Time complexity of Algorithm 1 for conventional block ciphers, which also indicates effective key lengths

Block cipher	n	k	R	Time complexity of Algorithm 1	Previously best time complexity on full rounds
AES-128	128	128	10	$2^{127.2}$	$2^{126.1}$ [6]
AES-192	128	192	12	$2^{191.1}$	$2^{189.7}$ [6]
AES-256	128	256	14	$2^{255.1}$	$2^{254.4}$ [6]
SHACAL2[1)]	256	512	64	2^{511}	NO
MISTY1 [12]	64	128	8	$2^{127.6}$	NO
ARIA-128 [13]	128	128	12	$2^{127.4}$	NO
ARIA-128	128	192	14	$2^{191.4}$	NO
ARIA-128	128	256	16	$2^{255.3}$	NO
IDEA [14]	64	128	8.5	$2^{127.4}$	$2^{126.1}$ [10]
KASUMI	64	128	8	$2^{127.4}$	$2^{125.8}$ [7]

1) Handschuh H, Naccache D. SHACAL: a family of block ciphers. Submission to the NESSIE project, 2002.

length can be regarded as 127.5 bits. We consider a little more of the structure of AES, that is, its branch number in the diffusion layer. Only four bytes knowledge of Z^4 is needed for computing Z_0^5, and four bytes knowledge of X^6 is needed. This can save additional 24 S-boxes, and the time complexity of Algorithm 1 is reduced to $2^{127.2}$. Similarly, the time complexity of Algorithm 1 for AES-256 is $2^{255.1}$. Compared with our upper bound of key bits, the best attack result so far on AES-256 with a time complexity of $2^{254.4}$ has much less gain than expected, since the effective key bits of AES-256 is actually only 255.1 bits.

We compute effective key lengths for other well-known block ciphers listed in Table 1.

We briefly explain KASUMI [15]. Assume that the most time consuming sub-functions are three FI in each round for KASUMI. Only seven 16-bit words of the key require to be guessed before going to the third FI of round 1. Also, there is no need to guess all 128 bits of the key when the three FI operations are completed in round 8. Besides, the Feistel structure saves one more round in the middle, such that there are 16 FI calculated for each guessed key. The time complexity is given as $2^{128} \times \frac{16}{24} = 2^{127.4}$.

The above ciphers are all recommended as standards or used by the industry for secure communications. According to our analysis, their security margin needs to be reconsidered. For example, if an attack on SHACAL2 has a time complexity larger than 2^{511}, then this attack should be regarded as invalid. The best attack on full IDEA that was thought to have optimized 1.9 bits now should be regarded as only 1.3 bits optimization.

4.2 Lightweight block ciphers

Secure communication on extremely constrained devices has been developing, such as RFID tags and sensor nodes. The constraints are mainly driven by cost and result in highly limited computing power, chip area, and power supply, which mean that we must leave behind much of our conventional block ciphers. Thus, the development of lightweight block ciphers is progressing greatly, resulting in more and more aggressive designs that often show two features. First, innovative techniques are used to improve existing ciphers. Second, the security margins that the block ciphers are traditionally equipped with are reduced as much as possible to optimize the cipher performance. Because of these differences in conventional block ciphers, we discuss the application of Algorithm 1 on lightweight block ciphers separately.

Take GOST as an example. GOST is known as the former Soviet encryption standard GOST 28147-89 which was standardized as the Russian encryption standard in 1989. It is well-suited for compact hardware implementations because of the simple structure, and the most compact implementation requires only 651 GE [16]. Therefore, GOST is considered as ultra lightweight. GOST has a 32-round Feistel structure with a 64-bit block size n and 256-bit key size k. The F-function consists of eight S-boxes. Refer to [17] for more details. The application of Algorithm 1 is as follows. Because of the Feistel structure, we can

Table 2 Time complexity of Algorithm 1 for lightweight block ciphers, which also indicates effective key lengths

Block cipher	n	k	R	Time complexity of Algorithm 1	Previously best time complexity on full rounds
GOST	64	256	32	$2^{254.8}$	2^{224} [18]
PRESENT-80 [19]	64	80	31	$2^{79.7}$	NO
PRESENT-128	64	128	31	$2^{127.6}$	NO
KATAN [20]	32/48/64	80	254	$2^{79.4}$	NO
KTANTAN [20]	32/48/64	80	254	$2^{79.4}$	$2^{75.2}$ [21]
HIGHT [22]	64	128	32	$2^{127.1}$	NO
XTEA [23]	64	128	64	$2^{127.7}$	NO
Piccolo-80 [24]	64	80	25	$2^{79.7}$	NO
Piccolo-128	64	128	31	$2^{127.6}$	NO

check if R^{15} equals to L^{16} (R^i and L^i are the right and left part of the input to round i). Compute P to S_1 by guessing the seven 32-bit subkeys in the first seven rounds, and the least significant 28 bits of the subkey in round 8, 252 bits totally. Then, for each guess of the most significant 4 bits of the subkey in round 8, complete the encryption from S_1 to R^{15}. This requires a calculation of 6 rounds and the last S-box in round 8, 49 S-boxes totally. Store the first 4 bits of R^{15} in a hash table corresponding to the guessed key. Similarly, compute C to S_2 by guessing the seven 32-bit subkeys in the last seven rounds, as well as the least significant 28 bits of the subkey in round 25, 252 bits totally. Then, for each guess of the most significant 4 bits of the subkey in round 25, complete the decryption operations from S_2 to L^{16}. This needs a computation from round 24 to round 16, and the last S-box in round 25, 73 S-boxes totally. Thus, for each guess of the 256-bit master key, 122 S-boxes require to be computed, which is 122/256 of a full 32 rounds encryption (there are $8 \times 32 = 256$ S-boxes for the full GOST). The time complexity of Algorithm 1 is about $2^{256} \times 122/256 = 2^{254.9}$, and so the effective key length is 254.9 bits. Also, we can only match part of R^{15} with part of L^{16}, e.g., their least significant 4 bits. To compute these 4 bits of R^{15}, only two S-boxes require to be calculated in round 14. Similarly, only two S-boxes are needed in round 16 for the matched 4 bits of L^{16}. Twelve S-boxes are saved now, so the time complexity is slightly reduced to $2^{254.8}$. Note that previous attacks on full GOST make use of its self-similarity property and relatively simple key schedule. We only consider the basic structure, which means that even the key schedule is much more complicated, Algorithm 1 still cannot be avoided.

Other results of lightweight block ciphers are summarized in Table 2. Some lightweight block ciphers have no nonlinear components, e.g., XTEA. In this situation, different linear operations in the round function are considered to cost the same time, or we can simply take the round function as a unit when computing the time complexity.

5 Discussion and conclusion

Recently, there are significant improvements on meet-in-the-middle attacks, as well as other brute-force-like cryptanalysis. This makes us consider a universal attack on all block ciphers, except the traditional exhaustive search method. In a practical cryptographic primitive, there are always some independent sub-modules. Computing these sub-modules with an independent pattern, instead of a combinational pattern, will save the overall time complexity. We describe a generic meet-in-the-middle attack that can always be mounted against any practical block ciphers. No amount of clever design can prevent it, no matter how many rounds, or how complicated structure and key schedule the cipher has. Note that having many rounds is still an important and expedient way to protect against it, since a larger number of rounds brings a higher complexity for Algorithm 1. We indicate a more accurate upper bound of effective key lengths for practical block ciphers, and claim that no ciphers can reach their expected security margin, the given length of their master keys. Previously, exhaustive key search is generally considered as the

benchmark with which other attacks are measured. A theoretical break (or academic break) against a block cipher is an attack with time complexity less than that of exhaustive key search, i.e., 2^k. Our analysis shows that tiny sacrifice of key bits is inevitable. Thus, if an attack has the computational complexity larger than Algorithm 1 (even still faster than exhaustive search), it cannot be regarded as a valid attack. Algorithm 1 is also used in many well-known block ciphers and their effective key lengths are calculated. As predicted, the effective key bits of these ciphers are all less than the master key size k. However, our attack will not create a real threat to the existing block ciphers, because of its limit caused by having to perform at least one operation for each possible key.

Another interesting discussion is about the relationship between the block size with the master key size. Shannon's work on information theory shows that to achieve the perfect secrecy, it is necessary for the key size to be at least as large as the block size. That is, $k \geqslant n$. According to our analysis in Section 3, when the number of rounds is fixed and the larger the k/n is, the more loss of effective key bits there is. Hence, $k = n$ is the best solution in the block cipher design in this context.

In the exhaustive key search, having to go through the entire key space before finding the correct key would be very unlucky, while being correct on the first guess would be very lucky. Thus, the expected time to recover a k-bit key is 2^{k-1} encryptions. Note that most of the effective key lengths we calculate for existing block ciphers are larger than this average case, although some are still smaller. Given that the time complexity of Algorithm 1 in this article is for the worst case, considering the average case and then comparing the result with 2^{k-1} can be undertaken for future work..

Acknowledgements

This work was supported by National Natural Science Foundation of China (Grant Nos. 61073149, 61272440), Research Fund for the Doctoral Program of Higher Education of China (Grant No. 20090073110027), State Key Laboratory of ASIC & System (Grant No. 11KF002), and Key Lab of Information Network Security, Ministry of Public Security (Grant No. C11603).

References

1 Luby M, Rackoff C. How to construct pseudo-random permutations from pseudo-random functions. In: Proceedings of Advances in Cryptology. Berlin/Heidelberg: Springer, 1986. 447–447
2 Even S, Mansour Y. A construction of a cipher from a single pseudorandom permutation. In: Proceedings of Advances in Cryptology. Berlin/Heidelberg: Springer, 1993. 210–224
3 Zhang B, Jin C H. Practical security against linear cryptanalysis for SMS4-like ciphers with SP round function. Sci China Inf Sci, 2012, 55: 2161–2170
4 Lv J Q. Differential attack on five rounds of the SC2000 block cipher. J Comput Sci Technol, 2011, 26: 722–731
5 Su B Z, Wu W L, Zhang W T. Security of the SMS4 block cipher against differential cryptanalysis. J Comput Sci Technol, 2011, 26: 130–138
6 Bogdanov A, Khovratovich D, Rechberger C. Biclique cryptanalysis of the full AES. In: Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security, Berlin/Heidelberg: Springer-Verlag, 2011. 344–371
7 Jia K, Yu H, Wang X. A meet-in-the-middle attack on the full KASUMI. Cryptology ePrint Archive, Report 2011/466, 2011
8 Biham E, Dunkelman O, Keller N, et al. New data-efficient attacks on reduced-round IDEA. Cryptology ePrint Archive, Report 2011/417, 2011
9 Lu J, Wei Y, Kim J, et al. Cryptanalysis of reduced versions of the Camellia block cipher. IET Inf Secur, 2012, 6: 228–238
10 Khovratovich D, Leurent G, Rechberger C. Narrow-Bicliques: cryptanalysis of full IDEA. Lect Note Comput Sci, 2012, 7237: 392–410
11 Daemen J, Rijmen V. AES proposal: Rijndael. In: Proceedings of the 1st Advanced Encryption Standard (AES) Conference, Ventura, 1998
12 Matsui M. New block encryption algorithm MISTY. Lect Note Comput Sci, 1997, 1267: 54–68
13 Kwon D, Kim J, Park S, et al. New block cipher: ARIA. Lect Note Comput Sci, 2004, 2971: 432–445
14 Lai X J, Massey J L, Murphy S. Markov ciphers and differential cryptanalysis. Lect Note Comput Sci, 1991, 547: 17–38

15 3rd Generation Partnership Project. Technical Specification Group Services and System Aspects, 3G Security, Specification of the 3GPP Confidentiality and Integrity Algorithms: KASUMI Specification. V3.1.1. 2001
16 Poschmann A, Ling S, Wang H. 256 bit standardized crypto for 650 GE: GOST revisited. In: Proceedings of Proceedings of the 12th International Conference on Cryptographic Hardware and Embedded Systems. Berlin/Heidelberg: Springer-Verlag, 2010. 219–233
17 National Soviet Bureau of Standards. Information Processing System—Cryptographic Protection—Cryptographic Algorithm GOST 28147-89. 1989
18 Dinur I, Dunkelman O, Shamir A. Improved attacks on full GOST. In: Proceedings of Fast Software Encryption. Berlin/Heidelberg: Springer, 2012. 9–28
19 Bogdanov A, Knudsen L R, Leander G, et al. PRESENT: an ultra-lightweight block cipher. Lect Note Comput Sci, 2007, 4727: 450–466
20 Cannière C D, Dunkelman O, Knezevic M. KATAN and KTANTAN—a family of small and efficient hardware-oriented block ciphers. Lect Note Comput Sci, 2009, 5747: 272–288
21 Bogdanov A, Rechberger C. A 3-subset meet-in-the-middle attack: cryptanalysis of the lightweight block cipher KTANTAN. Lect Note Comput Sci, 2010, 6544: 229–240
22 Hong D, Sung J, Hong S, et al. HIGHT: a new block cipher suitable for low-resource device. Lect Note Comput Sci, 2006, 4249: 46–59
23 Needham R M, Wheeler D J. TEA Extensions. Technical Report, Cambridge University, Cambridge, 1997
24 Shibutani K, Isobe T, Hiwatari H, et al. Piccolo: an ultra-lightweight block cipher. Lect Note Comput Sci, 2011, 6917: 342–357