

Improved Attacks on Full GOST

Itai Dinur¹, Orr Dunkelman^{1,2} and Adi Shamir¹

¹ Computer Science department, The Weizmann Institute, Rehovot, Israel

² Computer Science Department, University of Haifa, Israel

Abstract. GOST is a well known block cipher which was developed in the Soviet Union during the 1970's as an alternative to the US-developed DES. In spite of considerable cryptanalytic effort, until very recently there were no published single key attacks against its full 32-round version which were faster than the 2^{256} time complexity of exhaustive search. In February 2011, Isobe used the previously discovered reflection property in order to develop the first such attack, which requires 2^{32} data, 2^{64} memory and 2^{224} time. In this paper we introduce a new fixed point property and a better way to attack 8-round GOST in order to find improved attacks on full GOST: Given 2^{32} data we can reduce the memory complexity from an impractical 2^{64} to a practical 2^{36} without changing the 2^{224} time complexity, and given 2^{64} data we can simultaneously reduce the time complexity to 2^{192} and the memory complexity to 2^{36} .

Keywords: Block cipher, cryptanalysis, GOST, reflection property, fixed point property, 2D meet in the middle attack

1 Introduction

During the 1970's, the US decided to publicly develop the Data Encryption Standard (DES), which was the first standardized block cipher intended for civilian applications. At roughly the same time, the Soviet Union decided to secretly develop GOST [14], which was supposed to be used in civilian applications as well but in a more controlled way. The general design of GOST was finally published in 1994, but even today some of the crucial elements (e.g., the choice of Sboxes) do not appear in the public description, and a different choice can be made for each application.

GOST is a Feistel structure over 64-bit blocks. The round function consists of adding (modulo 2^{32}) a 32-bit round key to the right half of the block, and then applying the function f described in Figure 1. This function has an Sbox layer consisting of eight different 4×4 Sboxes, followed by a rotation of the 32-bit result by 11 bits to the left using the little-endian format (i.e. the LSB of the 32-bit word enters the rightmost entry of the first Sbox).

The full GOST has 32 rounds, and its key schedule is extremely simple: the 256-bit key is divided into eight 32-bit words (K_1, K_2, \dots, K_8). Each round of GOST uses one of these words as a round key in the following order: in the first 24 rounds, the keys are used in their cyclic order (i.e. K_1 in rounds 1,9,17, K_2 in rounds 2,10,18, and so forth). In the final 8 rounds (25–32), the round keys are used in reverse order (K_8 in round 25, K_7 in round 26, and so forth).

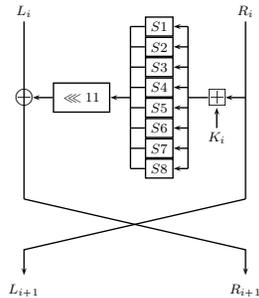


Fig. 1. One round of GOST

A major difference between the design philosophies of DES and GOST was that the publicly available DES was intentionally chosen with marginal parameters (16 rounds, 56-bit keys), whereas the secretive GOST used larger parameters (32 rounds, 256-bit keys) which seemed to offer an extra margin of security. As a result, DES was broken theoretically (by using differential and linear techniques) and practically (by using special purpose hardware) about 20 years ago, whereas in the case of GOST, all the single key attacks [1, 9, 17] published before 2011 were only applicable to reduced-round versions of the cipher.¹

The first single key attack on the full 32-round version of GOST was published by Isobe at FSE'11 [8]. It exploited a surprising reflection property which was first pointed out by Kara [9] in 2008: Whenever the left and right halves of the state after 24 rounds are equal (which happens with probability 2^{-32}), the last 16 rounds become the identity mapping, and thus the effective number of rounds is reduced from 32 to 16. Isobe developed a new key-extraction algorithm for the remaining 16 rounds of GOST which required 2^{192} time and 2^{64} memory, and used it 2^{32} times for different plaintext/ciphertext pairs in order to get the full 256-bit key using a total of 2^{32} data, 2^{64} memory, and 2^{224} time. This is much faster than exhaustive search, but neither the time complexity nor the memory complexity are even close to being practical.

Shortly afterwards, Courtois [4] published on ePrint a new attack on the full GOST. It uses a very different algebraic approach, but had an inferior complexity of 2^{64} data, 2^{64} memory, and 2^{248} time. Later, Courtois and Misztal [5] described a differential attack which again used 2^{64} data and memory, but reduced the time complexity to 2^{226} .

In this paper we improve several aspects of these previously published attacks. We describe a new *fixed point property*, and show how to use either the previous reflection property or the new fixed point property in order to reduce the general cryptanalytic problem of attacking the full 32-round GOST into an attack on 8-round GOST with two known input-output pairs. We then develop

¹ Attacks on full GOST in the stronger related-key model are known for about a decade, see [7, 10, 11, 16, 17].

a new way to extract all the 2^{128} possible values of the full 256-bit key given only two known 64-bit input-output pairs of 8-round GOST, which requires 2^{128} time and 2^{36} memory² (all the previously published attacks on 8-round GOST have an impractical memory complexity of at least 2^{64}). By combining these improved elements, we can get the best known attacks on GOST for the two previously considered data complexities of 2^{32} and 2^{64} .

Our new results on GOST (including the fixed point based attack) use well known and easy to analyze cryptanalytic techniques such as “Guess and Determine” and “meet-in-the-middle”. A month after this paper appeared on eprint [6] (and more than four months after its results were publicly disclosed in a public talk by Adi Shamir at MIT), Courtois posted to ePrint his independently discovered attacks [3], which use several different algebraic techniques. Some of his attacks are also based on the fixed point property, but all of them have higher claimed complexities: Given 2^{32} data, the best attack in [3] has a time complexity of 2^{224} and a memory complexity of 2^{128} , and given 2^{64} data, the best attack in [3] has a time complexity of 2^{216} and a negligible memory complexity. We include the results of [3] in Table 1 (which summarizes all the previously known single-key attacks on the full GOST, our new results, and Courtois’ subsequent results) for the sake of completeness.

An important observation about Isobe’s attack is that it uses in an essential way the assumption that the Sboxes are invertible. Since the GOST standard does not specify the Sboxes, and there is no need to make them invertible in a Feistel structure, Isobe’s attack might not be applicable to some valid incarnations of this standard. A similar problem occurs in most of Courtois’ attacks [3–5], as their complexities are only estimated for one particular choice of Sboxes described in [15] which is used in the Russian banking system, and it is possible that for other choices of Sboxes the complexities will be different. Our new attacks do not suffer from these limitations, since they can be applied with the same complexity to any given set of Sboxes.

2 Overview of Our New Attacks on the Full GOST

The 32 rounds of GOST can be described using only two closely related 8-round encryption functions. Let $G_{K_{i_1}, \dots, K_{i_j}}$ be j rounds of GOST under the subkeys K_{i_1}, \dots, K_{i_j} (where $i_1, \dots, i_j \in \{1, 2, \dots, 8\}$), and let (P_L, P_R) be a 64-bit plaintext, such its right half, P_R , enters the first round. Then $GOST_K(P_L, P_R) = G_{K_8, \dots, K_1}(G_{K_1, \dots, K_8}(G_{K_1, \dots, K_8}(G_{K_1, \dots, K_8}(P_L, P_R))))$.

Our new attacks on the full GOST exploit its high degree of self-similarity using a general framework which is shared by other attacks: the algorithm of each attack consists of an outer loop which iterates over the given 32-round

² We can reduce the memory complexity by an additional factor of 2^{17} (to 2^{19}) if we are willing to increase the time by a factor of 2^{12} (to 2^{140}). This may seem like an unattractive tradeoff since the 2^{36} memory complexity is already practical, but one can argue that 2^{19} words fit into the cache whereas 2^{36} do not, which may result in a big performance penalty.

Reference	Data (KP) ^{††}	Memory	Time	Self-Similarity Property	8-Round Attack	Sboxes
[8]	2^{32}	2^{64}	2^{224}	Reflection	-	Bijective
[4]	2^{64}	2^{64}	2^{248}	Other (unnamed)	Algebraic	Russian Banks [15]
[5]	2^{64}	2^{64}	2^{226}	None (differential attack)	-	Russian Banks [15]
[3] ^{†††}	2^{32}	2^{128}	2^{224}	Reflection	-	any
[3] ^{†††}	2^{64}	Negligible	2^{216}	fixed point	Algebraic	Russian Banks [15]
This paper	2^{64}	2^{36}	2^{192} [†]	fixed point	2DMITM	any
This paper	2^{64}	2^{19}	2^{204} [†]	fixed point	low-memory	any
This paper	2^{32}	2^{36}	2^{224} [†]	Reflection	2DMITM	any
This paper	2^{32}	2^{19}	2^{236} [†]	Reflection	low-memory	any

[†] The time complexity can be slightly reduced by exploiting GOST's complementation properties (as described in the full version of the paper [6])

^{††} Known plaintext

^{†††} Published on ePrint after the original version of this paper [6].

Table 1. Single-key Attacks on the Full GOST

plaintext-ciphertext pairs, and uses each one of them to obtain suggestions for two input-output pairs for G_{K_1, \dots, K_8} . For each suggestion of the 8-round input-output pairs, we apply an 8-round attack which gives suggestions for the 256-bit GOST key. We then verify the key suggestions by using some of the other plaintext-ciphertext pairs. The self-similarity properties of GOST ensure that the 8-round attack needs to be applied a relatively small number of times, leading to attacks which are much faster than exhaustive search.

We describe several attacks on the full GOST which belong to this common framework but differ according to the property and the type of 8-round attack we use. The two self-similarity properties are:

1. The *reflection property* which was first described in [9], where it was used to attack 30 rounds of GOST (and 2^{224} weak keys of the full GOST). This property was later exploited in [8] to attack the full GOST for all keys. We describe this property again in Section 3.1 for the sake of completeness.
2. A new *fixed point property* which is described in Section 3.2.

The two properties differ according to the amount of data required to satisfy them, and thus offer different points on a time/data tradeoff curve.

Given two 8-round input-output pairs, we describe in this paper several possible attacks of increasing sophistication:

1. A very basic meet-in-the-middle (MITM) attack [2], which is described in Section 4.1.

2. An improved MITM attack, described in Section 4.2, which uses the idea of equivalent keys (first described by Isobe in [8]).
3. A low-memory attack, described in Section 5, which requires 2^{19} memory and 2^{140} time.
4. A new *2-dimensional meet-in-the-middle* (2DMITM) attack, described in Section 6, which requires 2^{36} memory and 2^{128} time.

In order to attack the full GOST, we first select one of the two self-similarity properties to use in the outer loop of the attack according to the number of plaintext-ciphertext pairs available: in case we have 2^{64} pairs available, we select the fixed point property, and if we only have 2^{32} pairs, we select the reflection property. We then select one of last two 8-round attacks according to the amount of available memory: in case we have 2^{36} memory available, we select the 2DMITM attack, and if we only have 2^{19} memory, we select the low-memory attack. The outcome of this selection is an attack algorithm of the form:

1. For each plaintext-ciphertext pair (P, C) :
 - (a) Assuming that (P, C) satisfies the conditions of the self-similarity property, derive suggestions for two 8-round input-output pairs (I, O) and (I^*, O^*) .
 - (b) For each suggestion for (I, O) and (I^*, O^*) :
 - i. Execute the 8-round attack on (I, O) and (I^*, O^*) in order to derive suggestions for the key, and test each suggestion by performing trial encryptions on the remaining plaintext-ciphertext pairs.

The total time complexity of our attacks is calculated by multiplying the complexity of the 8-round attack by the expected number of times it needs to be applied according to the self-similarity property: An arbitrary (P, C) pair satisfies the fixed point property with probability of about 2^{-64} . Thus, it requires about 2^{64} known (P, C) pairs to succeed with high probability, and since we do not know in advance which pair satisfies the property, we need to repeat step 1 of the attack 2^{64} times. For each (P, C) pair, the fixed point property immediately suggests two 8-round input-output pairs (which are correct if the pair indeed satisfies the property). Hence, we need to perform step 1.(b) of the attack only once per (P, C) pair. In total, we need to execute the 8-round attack about 2^{64} times. On the other hand, an arbitrary (P, C) pair satisfies the reflection property with a much higher probability of about 2^{-32} . Thus, it requires about 2^{32} known (P, C) pairs, and we need to repeat the attack only 2^{32} times. However, for each (P, C) pair, the reflection property suggests a large number of 2^{64} values for (I, O) and (I^*, O^*) (out of which only one is correct if the pair indeed satisfies the property). Hence, we need to perform step 1.(b) of the attack 2^{64} times per (P, C) pair. In total, we need to execute the 8-round attack about $2^{32+64} = 2^{96}$ times.

Altogether, we obtain four new attacks on the full GOST. In three out of the four cases, we obtain better combinations of complexities than in all the previously published attacks. In the remaining case, we use the reflection property

and the low-memory 8-round attack to significantly reduce the memory requirements of Isobe’s attack [8], at the expense of a small time complexity penalty. We note that the computation required by each one of our attacks can be easily parallelized, and thus using x CPUs reduces the expected running time of the attack by a factor of x .

As described in the full version of this paper [6], the time complexity of all these attacks can be slightly reduced by exploiting GOST’s complementation properties. However, in some of these improved attacks we have to use chosen rather than known plaintexts, which reduces their attractiveness.

3 Obtaining Two 8-Round Input-Output Pairs for GOST

In this section, we describe the two self-similarity properties of GOST which we exploit in order to obtain two 8-round input-output pairs.

3.1 The Reflection Property [8, 9]

Assume that the encryption of a plaintext P after 24 rounds of GOST results in a 64-bit value Y , such that the 32-bit right and left halves of Y are equal (i.e. $Y_R = Y_L$). Thus, exchanging the two halves of Y at the end of round 24 does not change the intermediate encryption value. In rounds 25–32, the round keys K_1 – K_8 are applied in the reverse order, and Y undergoes the same operations as in rounds 17–24, but in the reverse order. As a result, the encryption of P after 32 rounds, which is the ciphertext C , is equal to its encryption after 16 rounds (see Figure 2). By guessing the state of the encryption of P after 8 rounds, denoted by the 64-bit value X , we obtain two 8-round input-output pairs (P, X) and (X, C) . For an arbitrary key, the probability that a random plaintext gives such a symmetric value Y after 24 rounds is 2^{-32} , implying that we have to try about 2^{32} known plaintexts (in addition to guessing X) in order to obtain the two pairs. Note that the reflection property actually gives us another “half pair” (\hat{C}, Y) , where the 64-bit word \hat{C} is obtained from C by exchanging the right and left 32-bit halves of C , and the 32-bit right and left halves of Y are equal.³ However, it is not clear how to exploit this additional knowledge in order to significantly improve the running time of our attacks on the full GOST which are based on the reflection property.

³ In our attacks, we use 8-round input-output pairs whose encryption starts with K_1 and thus need to apply the Feistel structure in the reverse order (starting from round 32) for input-output pairs obtained for rounds 25–32. Since in Feistel structures the right and left halves of the block are exchanged at the end (rather than at the beginning) of the round function, we exchange the right and left sides of the input and the output of the input-output pairs obtained for rounds 25–32. We call (\hat{C}, Y) a “half pair” since we have to guess only 32 additional bits in order to find it, once (P, C) is known.

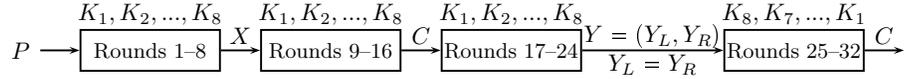


Fig. 2. The Reflection Property of GOST

3.2 The Fixed Point Property

Assume that for a plaintext P , $G_{K_8, \dots, K_1}(P) = P$. Since rounds 9–16 and 17–24 are identical to rounds 1–8, we obtain P after 16 and 24 rounds as well. In rounds 25–32, the round keys K_1 – K_8 are applied in the reverse order, and we obtain some arbitrary ciphertext C (see Figure 3). The knowledge of P and C immediately gives us the 8-round input-output pairs (P, P) and (\hat{C}, \hat{P}) (in which the right and left 32-bit halves of P and C are exchanged).

For an arbitrary key, the probability that a random plaintext is a fixed point is about 2^{-64} , implying that we need about 2^{64} known plaintexts to have a single fixed point, from which we obtain the two input-output pairs needed in our attack. If we have only $c \cdot 2^{64}$ known plaintexts for some fraction c , we expect this fixed point to occur among the given plaintexts with probability c , and thus the time complexity, the data complexity, and the success probability are all reduced by the same linear factor c . Consequently, it makes sense to try the fixed point based attack even when we are given only a small fraction of the entire code book of GOST. Such a graceful degradation when we are given fewer plaintexts (which also occurs for the reflection property) should be contrasted with other attacks such as slide attacks, in which we have to wait for some random birthday phenomenon to occur among the given data points. Since the existence of birthdays has a much sharper threshold, the probability of finding an appropriate pair of points goes down quadratically rather than linearly in c , and thus they are much more likely to fail in such situations.

We note that our fixed point property is closely related to a previously published property which (in addition to the assumption the P is an 8-round fixed point) also assumes that the right and left halves of P are equal. Such a plaintext exists for an arbitrary key with probability 2^{-32} and thus was used in [9] to attack 2^{224} weak keys of the full GOST. The same property was also used later in [13] in cryptanalysis of the GOST hash function.

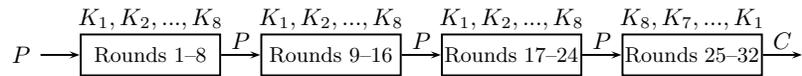


Fig. 3. The fixed point property of GOST

4 Simple Meet-in-the-middle Attacks on 8 Rounds of GOST

Meet-in-the-middle (MITM) attacks can be efficiently applied to block ciphers in which some intermediate encryption variables (bits, or combinations of bits) depend only on a subset of key bits from the encryption side and on another subset of key bits from the decryption side: the attacker guesses the relevant key bits from the encryption and the decryption sides independently, and tries only keys in which the values suggested by the computed intermediate variables match. While the full 32-round GOST resists such attacks, 8-round GOST uses completely independent round keys. Thus, the full 64-bit value after 4 encryption rounds depends only on round keys K_1 – K_4 from the encryption side and on round keys K_5 – K_8 from the decryption side.

4.1 The Basic Meet-in-the-middle Attack

We describe how to mount a simple meet-in-the-middle attack on 8 rounds of GOST given two 8-round input-output pairs and several additional 32-round plaintext-ciphertext pairs:

1. For each of the 2^{128} possible values of K_1 – K_4 , encrypt both inputs and obtain two 64-bit intermediate encryption values after 4 rounds of GOST (i.e., 2^{128} intermediate values of 128 bits each). Store the intermediate values in a list, sorted according to these 128 bits, along with the corresponding value of K_1 – K_4 .
2. For each of the 2^{128} possible values of K_5 – K_8 , decrypt both outputs, obtain two 64-bit intermediate values and search the sorted list for these two values.
3. For each match, obtain the corresponding value of K_1 – K_4 from the sorted list and derive a full 256-bit key by concatenating the value of value of K_1 – K_4 with the value of K_5 – K_8 of the previous step. Using the full key, perform a trial encryption of several plaintexts and return the full key, i.e., the one that remains after successfully testing the given 32-round pairs.

We expect to try about $2^{128+128-128} = 2^{128}$ full keys in step 3 of the attack, out of which only the correct key is expected to pass the exhaustive search of step 3. Including the 2^{128} 8-round encryptions which are performed in each of the first two steps of the attack, the total time complexity of the attack is slightly more than 2^{128} GOST encryptions. The memory complexity of the attack is about 2^{128} words of 256 bits.⁴

⁴ Note that it is possible obtain a time-memory tradeoff: we partition the 2^{128} possible values of K_1 – K_4 into 2^x sets of size 2^{128-x} (for $0 \leq x \leq 128$), and run the second and third steps of the attack independently for each set. Thus, the memory complexity decreases by a factor 2^x to 2^{128-x} , and the time complexity increases by a factor of 2^x to 2^{128+x} .

4.2 An Improved Meet-in-the-middle Attack Using Equivalent Keys

In this section, we use a more general variant of Isobe’s equivalent keys idea [8] to significantly improve the memory complexity of the attack. Both our and Isobe’s MITM attacks are based on a 4-round attack that uses one 4-round input-output pair to find all the 2^{64} possible values of subkeys K_1 – K_4 that yield this pair. However, our MITM attack is more general since we can attack all possible incarnations of the GOST standard, whereas Isobe’s attack works only on those which use bijective Sboxes.⁵ An additional advantage of our MITM attack over Isobe’s one, is that our attack can use any two input-output pairs for 8-round GOST, regardless of how they are obtained. We can thus use the same algorithm to exploit both the reflection and the fixed point properties. On the other hand, Isobe’s attack is restricted to the case of a single input-output pair obtained for the first 16 rounds of GOST (by guessing the intermediate values obtained after 4 and 12 rounds) and thus can be combined with the reflection property, but cannot be directly applied to the two input-output pairs produced by the fixed point property.

We now describe Isobe’s 4-round attack procedure: Denote the 4-round input (divided into two 32-bit words) by (X_L, X_R) and the output by (Y_L, Y_R) . Denote the middle values (after the second round) by (Z_L, Z_R) (see Figure 4). Then:

$$Z_L = X_L \oplus f(X_R \boxplus K_1)$$

$$Z_R = Y_R \oplus f(Y_L \boxplus K_4)$$

$$Y_L \oplus Z_L = f(Z_R \boxplus K_3)$$

$$X_R \oplus Z_R = f(Z_L \boxplus K_2)$$

Isobe’s attack assumes bijective Sboxes (making f invertible), and finds the equivalent keys as follows:⁶ for each value of K_1, K_2 , compute Z_L from the first equation and Z_R from the fourth equation. From the second equation we have: $K_4 = f^{-1}(Z_R \oplus Y_R) \boxplus Y_L$ and from the third equation: $K_3 = f^{-1}(Z_L \oplus Y_L) \boxplus Y_R$.

Our 8-round attack is a variant of Isobe’s MITM attack, given two 8-round input-output pairs (I, O) and (I^*, O^*) :

1. For each possible value of the 64-bit word $Y = (Y_L, Y_R)$ obtained after 4 encryption rounds of I :
 - (a) Apply the 4-round attack on (I, Y) to obtain 2^{64} candidates for K_1 – K_4 .
 - (b) Partially encrypt I^* using the 2^{64} candidates and store $Y^* = (Y_L^*, Y_R^*)$ in a list with K_1 – K_4 .

⁵ The Feistel structure of GOST does not require bijective Sboxes and the published standard does not discuss this issue, but all the known choices of Sboxes happen to be bijective (perhaps due to the weakness of non-bijective Sboxes against differential cryptanalysis).

⁶ In case f is not bijective, then for a random (X_L, X_R) and (Y_L, Y_R) there exist an average of 2^{64} equivalent keys which can be found using a simple preprocessing MITM algorithm that requires about 2^{64} time and memory.

- (c) Apply the 4-round attack on (Y, O) to obtain 2^{64} candidates for K_5-K_8 .
- (d) Partially decrypt O^* using each one of the 2^{64} candidates and obtain $Y^* = (Y_L^*, Y_R^*)$.
- (e) Search the list obtained in step (b) for Y^* , and test the full 256-bit keys for which there is a match.

The expected time complexity of steps (a–d) is about 2^{64} (regardless of the algorithm that is used to find the equivalent keys). The time complexity of step (e) is also about 2^{64} since we expect to try about $2^{64+64-64} = 2^{64}$ full keys. Steps (a–e) are performed 2^{64} times, hence the total time complexity of the attack is about 2^{128} GOST encryptions, which is similar to the first attack. However, the memory complexity is significantly reduced from 2^{128} to slightly more than 2^{64} words of 64 bits.

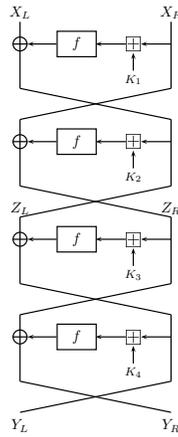


Fig. 4. Four Rounds of GOST

5 A New Attack on 8 Rounds of GOST with Lower Memory Complexity

Simple meet-in-the-middle attacks, such as the ones described in Sections 4.1 and 4.2 are much faster than exhaustive search for the entire 256-bit key. However, they do not fully exploit the slow diffusion of the key bits in 4 rounds of GOST. As a result, these MITM attacks use a large amount of memory to store the many intermediate encryption values obtained for all the possible values of large sets of key bits. In this section, we describe an improved 8-round attack which exploits the slow diffusion properties of 4 rounds of GOST in order to reduce the memory complexity from the impractical value of 2^{64} to the very practical value of 2^{19}

words of memory, with a very small time complexity penalty. The main idea of this attack is to guess the 4 round keys K_5 – K_8 and apply an optimized “Guess and Determine” attack on the remaining 4 rounds using two input-output pairs. In the 4-round attacks we have 128-bits of unknown key and 128 bits of input-output pairs. Thus, we expect that only one value for K_1 – K_4 exists (although there are likely to be input-output pairs for which the encryptions of the inputs does not match the outputs for any of the keys, and input-output pairs for which the encryptions of the inputs matches the outputs for several values of K_1 – K_4).

In the rest of this section we describe the algorithm for deriving the 32 bits of K_1 and the 32 bits of K_4 . Afterwards, deriving the values of K_2 and K_3 is immediate using the third and fourth equations of Section 4.2 (Z_L and Z_R are known from the first and second equations).

5.1 Overview of the “Guess and Determine” Attack on 4-Round GOST

Now that we deal with 4-round GOST, we apply a typical “Guess and Determine” attack which traverses a tree of partial guesses for the round keys K_1 and K_4 and intermediate encryption values. The tree is composed of layers of nodes ℓ_i for integral $0 \leq i \leq k$, where each layer contains nodes that specify the potential values (i.e. guesses) for a certain subset of key and intermediate encryption values. In each layer we expand each node by guessing the values of a small number of additional key bits and state bits that are needed to calculate some intermediate encryption bits, both from the encryption and the decryption sides. We then calculate the bits by evaluating the Feistel structure from both sides on a small number of bits, compare the values obtained, and discard guesses in which the values do not match (i.e., we discard child nodes that do not satisfy a predicate which checks the consistency of intermediate encryption values).

We traverse the partial guess tree starting from the root using DFS (which requires only a small amount of memory). In our attack, the nodes of the last layer of the tree ℓ_k contain guesses for the full key, which can be verified using trial encryptions.

The total number of operations performed during the traversal is proportional to the total number of nodes in the tree. However, the operations performed when expanding a single node work only on a few bits (rather than on full words). At the same time, when expanding a full path of nodes in the tree from the root to the last layer, we work on the full-size Feistel structure to obtain a guess for the full key. Hence, we estimate the time complexity of expanding a full path by a single Feistel structure evaluation on a full 64-bit input. Using this estimation, we can upper bound the time complexity of the tree traversal (in terms of Feistel structure evaluations) as the width of the tree, or the size of the layer which contains the highest number of nodes. Note that when counting the number of nodes in a layer for the time complexity analysis, we must also include nodes that were expanded and discarded since they do not satisfy the predicate of the previous layer.

5.2 Notations

Assume that we have two input-output pairs for 4 encryption rounds of GOST under the subkeys K_1, K_2, K_3, K_4 . Similarly to Section 4.2, denote the input, output and middle values (after using K_2) for the first pair by $(X_L, X_R), (Y_L, Y_R)$ and (Z_L, Z_R) , respectively. For the second pair, denote these values by $(X_L^*, X_R^*), (Y_L^*, Y_R^*)$ and (Z_L^*, Z_R^*) respectively.

Since our attack analyzes 4-bit words (which are outputs of single Sboxes), we introduce additional notations: Define the functions f^0, f^1, \dots, f^7 where each f^i takes a 4-bit word as an input, and outputs a 4-bit word by applying Sbox i to the input. Denote by W^i the i 'th bit of the 32-bit word W , and by $W^{i,j}$ the $(j - i + 1)$ -bit word composed of consecutive bits of W starting from bit i and ending at bit j . We treat W as a cyclic word, and thus $W^{24,3}$ contains 12 bits which are bits 24 to 31 and 0 to 3 of W .

5.3 An Attack on 4 Rounds of Simplified GOST

We start by describing an attack on 4 rounds of a simplified variant of GOST (which we call S-GOST), in which the round-key addition is replaced by XOR, and the 11-bit rotation is replaced by 12-bit rotation. The simplified variant is easier to analyze since it provides much slower diffusion of the key bits compared to full GOST: unlike addition, the XOR operation does not produce carries, and since 12 is a multiple of 4, rotating by 12 bits implies that the output of any Sbox effects the input of only a single Sbox in the next round.

We now describe the basic procedure performed by a node in layer 0 of our guess tree for S-GOST. The procedure requires the value of $K_1^{0,3}$ (whose value we guess before executing the procedure), and expands nodes in the next layer, which suggest a value for the additional 4 bits of $K_4^{20,23}$. The steps of this procedure can be easily verified using a variant of Figure 4 where the addition is replaced by XOR.

1. Given $K_1^{0,3}$ and $X_R^{0,3}$, compute $Z_L^{12,15} \equiv f^0(X_R^{0,3} \oplus K_1^{0,3})$ for both pairs (i.e., given $K_1^{0,3}$ and $X_R^{0,3}$, compute $Z_L^{12,15} \equiv f^0(X_R^{0,3} \oplus K_1^{0,3})$).
2. Obtain $f^0(Z_R^{0,3} \oplus K_3^{0,3}) \equiv Z_L^{12,15} \oplus Y_L^{12,15}$ for both pairs. Then, invert⁷ f^0 to obtain $Z_R^{0,3} \oplus K_3^{0,3}$ and $Z_R^{*0,3} \oplus K_3^{0,3}$.
3. XOR the two expressions calculated in step 2, to eliminate $K_3^{0,3}$, and obtain the value of $Z_R^{0,3} \oplus Z_R^{*0,3}$.
4. XOR the 4-bit difference obtained in step 3 to the difference $Y_R^{0,3} \oplus Y_R^{*0,3}$ and obtain the value of $T = Z_R^{0,3} \oplus Y_R^{0,3} \oplus Z_R^{*0,3} \oplus Y_R^{*0,3} \equiv (f(Y_L \oplus K_4) \oplus f(Y_L^* \oplus K_4))^{0,3}$ (from the encryption side).
5. For each of the 2^4 possible values of $K_4^{20,23}$:
 - (a) Allocate a node in the next layer.

⁷ We expect one solution on average. However, in case the inversion has more than one solution, we need to try each one. In case the inversion has no solution, we can discard the node.

- (b) Evaluate the expression $f^5(Y_L^{20,23} \oplus K_4^{20,23}) \oplus f^5(Y_L^{*20,23} \oplus K_4^{20,23})$ from the decryption side by plugging the current value of $K_4^{20,23}$ into the expression. Discard nodes which do not agree with the value T .

Note that given $K_1^{0,3}$, we expect the procedure above to process a single child in the next layer: in step 5 we have a 4-bit condition on 4 bits of the key $K_4^{20,23}$, and thus we expect one node to satisfy the predicate. Moreover, step 5 can be optimized by using a small amount of precomputation and memory in order to calculate in advance the solutions to the 4-bit condition (as described in the full version of this paper [6]).

We now generalize the procedure above in order to derive more key bits in a similar way:

- Since encryption and decryption are completely symmetric (except the order of the subkeys), steps 1–5 can also be performed from the decryption side: in steps 1–5 we use the value of $K_1^{0,3}$ in order to obtain the value of $K_4^{20,23}$, and thus we define the symmetric steps 6–10 which use the value of $K_4^{20,23}$ in order to obtain the value of $K_1^{20+20,23+20}$, i.e. $K_1^{8,11}$.
- Given any integer $0 \leq i \leq 7$, we can rotate the indices of all the 32-bit words in steps 1–10 by $4i$ bits. Namely, given i , we define analogues steps 1–10 which use the value of $K_1^{4i,4i+3}$ to obtain the value of $K_4^{4i+20,4i+23}$ and $K_1^{4i+8,4i+11}$.

In order to derive the full 32-bit values of K_1 and K_4 , we define a tree which contains 9 layers $\ell_0, \ell_1, \dots, \ell_8$ (and an additional root node). The nodes of each layer are expanded using the generalized procedure which uses 4 bits of K_1 in order to derive 4 additional bits of K_1 and 4 additional bits of K_4 . Since the 10 steps of the procedure for expanding the nodes of layers 0–7 are basically the same, we call this procedure an *iteration*, and index it according to the value of i (which determines the 4-bit chunks that we work on).

5.4 Extending the Attack to 4 Rounds of the Real GOST

In order to extend the iteration procedure from S-GOST to full GOST, we need to make several adjustments. The most significant adjustments are given below:

- Since the round keys are added (and not XORed) to the state, we have to guess the carry bits into the LSBs of several addition operations of 4-bit words. For example, in the expression $f^5(Y_L^{20,23} \boxplus K_4^{20,23}) \oplus f^5(Y_L^{*20,23} \boxplus K_4^{20,23})$ evaluated in step 5, we have to guess two carry bits (one for $Y_L^{20,23}$ and one for $Y_L^{*20,23}$).
- GOST uses 11-bit rotation (instead of 12-bit rotation), and thus the 4-bit chunks that we work on in each iteration are not aligned. Consequently, we have to guess additional state bits in order to compare the evaluation of the 4-bit predicates from both sides. For example, since $20 + 11 = 31$, in step 5 of the iteration we actually calculate $(f(Y_L \oplus K_4) \oplus f(Y_L^* \oplus K_4))^{31,2}$ from the decryption side. Thus, we additionally guess bit 31 of this expression from the encryption side.

These adjustments create strong dependencies between iterations with consecutive indexes (i.e., i and $i + 1$), namely:

- The carry bits required by iteration $i + 1$ are known after iteration i . For example, iteration 1 requires the carry into bit 24 of the addition operation $Y_L \boxplus K_4$ (in order to calculate $f^6(Y_L^{24,27} \boxplus K_4^{24,27}) \oplus f^6(Y_L^{*24,27} \boxplus K_4^{24,27})$ in step 5). This bit can be calculated after step 5 of iteration 0, where the 4-bit value of $Y_L^{20,23} \boxplus K_4^{20,23}$ is calculated in order to evaluate the predicate.
- The state bits required by iteration $i + 1$ are known after iteration i . For example, iteration 1 requires calculation of bit 3 of the expression $f(Y_L \boxplus K_4) \oplus f(Y_L^* \boxplus K_4)$ from the encryption side. However, this bit is already guessed in step 4 of iteration 0.

This suggests that we perform the iterations in their natural order, namely assign layer ℓ_i iteration i for $0 \leq i \leq 7$. As a result, we need to guess carry and state bits only in the first iteration. Afterwards, the required carry and state bits for each iteration can be calculated by the knowledge from the previous one. On the other hand, we pay a (relatively small) penalty on key bit guesses since key bits required by iteration $i + 2$ are derived in iteration i (and not in iteration $i + 1$). Since iteration i requires key bits $K_1^{4i, 4i+3}$, we need to guess 4 key bits in both iterations 0 and 1 ($K_1^{0,3}$ and $K_1^{4,7}$). For iterations $i \geq 2$, the required key bits are already derived in previous iterations (as shown in Table 2).

We note that since there is no carry into the LSBs of addition operations, starting the process with iteration 0 has the advantage that we do not need to guess the carries for all the addition operations (e.g., we do not need to guess the carry into the addition $f^0(X_R^{0,3} \boxplus K_1^{0,3})$ in step 1).

The full details and analysis of the ‘‘Guess and Determine’’ attack are given in the full version of this paper [6], most of which is not required in order to understand the rest of this paper. It shows that the expected number of nodes in the widest layer of the partial guess tree is 2^{14} , and it is obtained at iterations 1 to 5 (this was also verified using simulations performed on a PC). Basically, the number 2^{14} is obtained due to the 8 key-bit guesses ($K_1^{0,3}$ and $K_1^{4,7}$) and 6 additional carry and state bit guesses in iteration 0. This gives an expected time complexity of about 2^{14} 4-round Feistel structure evaluations for two input-output pairs, which is equivalent to about 2^{12} full GOST evaluations. Since we apply this 4-round attack 2^{128} times, the time complexity of the 8-round attack is about $2^{128+12} = 2^{140}$ GOST evaluations. In terms of memory, the attack has a completely practical complexity of 2^{25} bits, which is equivalent to 2^{19} 64-bit words.

6 A New 2-Dimensional Meet-in-the-middle Attack on 8 Rounds of GOST

In this section, we present a new attack on 8 rounds of GOST given two input-output pairs, which combines the ideas of the ‘‘Guess and Determine’’ attack

Iteration	0	1	2	3	4	5	6	7
K_1 bits derived	0-3	4-7	<u>8-11</u>	<u>12-15</u>	<u>16-19</u>	<u>20-23</u>	<u>24-27</u>	<u>28-31</u>
K_4 bits derived	8-11	12-15	16-19	20-23	24-27	28-31	<u>0-3</u>	<u>4-7</u>
	20-23	24-27	28-31	0-3	4-7	8-11	12-15	16-19

The key bits which are already known from previous iterations are underlined.

Table 2. The key bits derived in each iteration

(which progresses horizontally across the state) and the MITM attack (which progresses vertically across the rounds). Unlike the attack of the previous section, we do not guess the last 4 round keys in advance. Instead, we divide the 8-round Feistel structure horizontally by splitting it into a *top part*, which uses round keys K_1 – K_4 , and a *bottom part*, which uses round keys K_5 – K_8 .

Our main observation is that due to the slow diffusion of the data bits into the state, we can run a substantial part of the “Guess and Determine” attack of Section 5 with very partial knowledge of Y and Y^* (obtained after 4 rounds of encryption). This allows us to split the “Guess and Determine” attack into two partial 4-round attacks which we run a relatively small number of times (once for each value of the bits of Y and Y^* that it requires). Our full 4-round attacks on the top and bottom parts combine the suggestions of the partial attacks in order to suggest values for the 4-round keys. Finally, we use an 8-round attack which joins the suggestions of the two partial attacks in order to obtain suggestions for the full 256-bit key.

Schematically, we split the top and bottom parts of the block cipher vertically into two (potentially overlapping) cells, such that on each cell we execute an independent partial attack to obtain suggestions for some subset of key bits. We then join all the suggestions to obtain suggestions for the full key using three MITM attacks. This can be visualized using a 2×2 matrix (as shown in Figure 5), where the horizontal line separates the four initial and final rounds of the 8-round block cipher, and the dashed vertical line separates the left and right cells in each one of the top and bottom parts.

After the MITM attacks on the top and bottom parts of the Feistel structure, we obtain 2^{128} suggestions for K_1 – K_4 and 2^{128} suggestions for K_5 – K_8 , each accompanied by corresponding 128-bit values of Y and Y^* . Note that so far we did not filter out any possible keys, and thus the final MITM attack, which compares the 128-bit values of Y and Y^* to obtain about 2^{128} suggestions for the full key, is essentially the basic MITM attack of Section 4.1, which would normally require 2^{128} memory.

To reduce the memory consumption, we guess many of the 128 bits of Y and Y^* in advance (in the outer loop of the 8-round attack). For each possible value of those bits, we execute the 2DMITM (2-dimensional MITM) attack described above, but obtain fewer suggestions for the key which we have to store. This increases the number of times that we execute the partial 4-round attacks and could potentially increase the overall time complexity of the full 8-round attack.

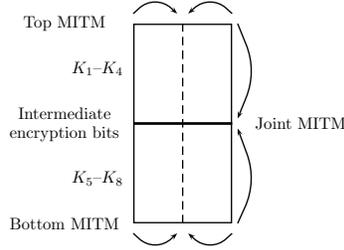


Fig. 5. The general framework of the 2-dimensional meet-in-the-middle attack

However, this is not the case, as the partial 4-round attacks are relatively efficient (the time complexity of each one is at most 2^{18}) and is executed only 2^{82} times. Thus, the partial 4-round attacks are not the bottleneck of the time complexity of the attack.⁸

6.1 Details of the 8-Round Attack

Formally, we define the following sets which contain bits of Y and Y^* :

- S_1 is the set of bits that we guess in the outer loop of the 8-round attack.
- S_2 is chosen such that $S_1 \cap S_2 = \emptyset$, and $S_1 \cup S_2$ is the minimal set that contains all the bits of Y and Y^* which are required by the partial 4-round attack on the left cell of the top part.
- S_3 is chosen such that $S_1 \cap S_3 = \emptyset$, and $S_1 \cup S_3$ is the minimal set of bits which are required by the partial 4-round attack on the right cell of the top part.

For the bottom MITM attack, we define S_4 and S_5 in a similar way to S_2 and S_3 , respectively. Note that since the 4-round attacks on both the top and bottom parts require all the 128 intermediate bits, $S_2 \cup S_3 = S_4 \cup S_5$.

The details of the 4-round attacks are given in the next section. We now refer to them as black boxes, and give the algorithm of the full 8-round attack:

1. For each value of the bits of the set S_1 :
 - (a) Perform the 4-round attack on the top part of the Feistel structure, and obtain a list with values of K_1-K_4 , sorted according to the value of the bits of $S_2 \cup S_3$.

⁸ Note again that we expect about 2^{128} keys to fulfill the filtering conditions of the two input-output pairs. Thus, the time required for the attack to list all of them cannot be reduced below 2^{128} (without exploiting additional filtering conditions).

- (b) Perform the 4-round attack on the bottom part of the Feistel structure. For each value of $S_4 \cup S_5 = S_2 \cup S_3$ (given along with the value of K_5 – K_8), search the list obtained in the previous step of matches. For each match test the full key K_1 – K_8 .

6.2 Details of the 4-Round Attacks

We concentrate first on the top part of the 8-round Feistel structure: each one of the two partial 4-round attacks on the top part sequentially executes a subset of the iterations defined in Section 5, and is called an iteration *batch*. The first (left) iteration batch executes iterations 0–3, and the second (right) executes iterations 4–7.

After performing iteration batches 0–3 and 4–7 independently, we get suggestions for the values of some key bits, along with some carry and state bits. We then discard inconsistent suggestions by comparing the values of the common bits that are derived by batches. We partition these bits into three groups (which are fully specified in the full version of this paper [6]):

- G_1 contains 16 key bits which are derived by both of the left and right batches.
- G_2 contains 6 carry and state input bits that we guess in iteration 0. These bits are also contained in the set of output bits of iteration 7 (of the right batch), and can thus be used to discard inconsistent suggestions made by the two batches.
- G_3 contains 10 carry and state input bits that we guess in iteration 4. This bits are also contained in the set of iteration output bits of iteration 3 (of the left batch), and can thus be used to discard inconsistent suggestions made by the two batches.

Assume that the values of all the bits of S_1 are known. We now give the algorithm of the MITM attack performed on the top part of the 8-round Feistel structure:

1. For each value of the bits of S_2 , perform the left batch. Save all the nodes of the final layer in a list. These nodes contain the values 40 bits of K_1 and K_4 (including the values of the bits of G_1), and also the values of the bits of G_3 . In addition to the information obtained by each node, save the value of the initial guess of the bits of G_2 , and the value of the bits of S_2 per node. Sort the list according to the values of G_1, G_2 and G_3 .
2. For each value of the bits of S_3 , perform the right batch. For each node in the final layer obtain the value of the bits of G_1, G_2 and G_3 and search the list obtained in the first step for their value. For each match, save the value of the full K_1 – K_4 in a sorted list according to the value of the bits of $S_2 \cup S_3$.

The iteration batches of the MITM attack on the bottom part of the Feistel structure are performed from the decryption side and are completely analogous

to the iteration batches on the top part (i.e. in iteration 0, we start by guessing $K_8^{0,3}$, and derive $K_5^{20,23}$ and $K_8^{8,11}$). We also define analogous sets to G_1, G_2 and G_3 for the bottom part.

The specific choices of S_1 – S_5 are given in the full version of this paper [6]. This choice of sets satisfies $|S_1| = 92$ and $|S_2| = |S_3| = |S_4| = |S_5| = 18$.

We now analyze the complexity of the MITM attack on the top part of the Feistel structure: as specified in the full version of this paper [6], when starting the iteration batch from iteration 0, the expected maximal size of the tree is 2^{14} . It is obtained after iteration 1, and is maintained until the end of iteration 5 (even though we do not perform 5 consecutive iterations in this attack). The time complexity of the first step of the attack is thus about $2^{|S_2|+14} = 2^{14+18} = 2^{32}$, and this is also the size of the sorted list at the end of the first step. The maximal size of the tree of the iteration batch 4–7 is $2^{14+4} = 2^{18}$ (as described in the full version of this paper [6], we have to guess 4 more carry bits compared to iterations 0–3). Thus, the time complexity of expanding the tree in the second step is $2^{|S_3|+18} = 2^{36}$. The time and memory complexities of the remainder of step 2 (in which we match the batches) are $2^{|S_2|+|S_3|+14+18-(|G_1|+|G_2|+|G_3|)} = 2^{|S_2|+|S_3|+14+18-(16+6+10)} = 2^{|S_2|+|S_3|} = 2^{36}$. Note that it is not surprising that the time and memory complexities of the matching part of the attack reduce to $2^{|S_2|+|S_3|}$, since given the full 128-bit intermediate value, we expect that only one key survives the filtering conditions. Altogether, the memory complexity of the top MITM attack is about 2^{36} 64-bit words. The time complexity is dominated by step 2 and is equivalent to about 2^{36} 4-round Feistel structure evaluations, which is equivalent to about 2^{33} evaluations of the full GOST cryptosystem. For the bottom MITM attack, we obtain the same time and memory complexities, since the sizes of S_4 and S_5 are equal to the sizes of S_2 and S_3 , and the sets corresponding to G_1, G_2 and G_3 are completely symmetrical.

6.3 The Complexity of the 8-Round Attack on GOST

We can now analyze the complexity of the attack described in Section 6.1: The time complexities of each of the MITM attacks on the bottom and top parts in steps (a) and (b) are equivalent to about 2^{36} 4-round Feistel structure evaluations, as calculated above. The number of expected matches for which we run the full cipher in step (b) is $2^{36+36-36} = 2^{36}$. Hence, the time complexity of these steps is equivalent to a bit more than 2^{36} full GOST evaluations. Since $|S_1| = 92$, the total time complexity of the attack is equivalent to about $2^{92+36} = 2^{128}$ GOST evaluations. The total memory complexity of the attack is about 2^{36} 64-bit words, and is dominated by the sorted list calculated in step (a).

7 Conclusions and Open Problem

In this paper we introduced several new techniques such as the fixed point property and two dimensional meet in the middle attacks, and used them to greatly

improve the best known attacks on the full 32-round GOST. In particular, we reduced the memory complexity of the attacks from an impractical 2^{64} to a practical 2^{36} (and to an even more practical 2^{19} complexity, which can fit into the cache of modern microprocessors, with a small penalty in the running time). The lowest time complexity of our attacks is 2^{192} , which is 2^{32} times better than previously published attacks but still very far from being practical. Consequently, we are concerned about the demonstrated weaknesses in the design of GOST (especially in its simplistic key schedule), but do not advocate that its current users should stop using it right away.

The main open problems left in this paper are whether it is possible to find faster attacks, and how to better exploit other amounts of available data (in addition to the 2^{32} and 2^{64} complexities considered in this paper, which are the natural thresholds for our techniques).

Acknowledgements: The authors thank Nathan Keller, Pierre-Alain Fouque and Charles Bouillaguet for useful discussions on this work, and the anonymous referees for their helpful comments on this paper which greatly improved the presentation of our results.

References

1. Eli Biham, Orr Dunkelman, and Nathan Keller. Improved Slide Attacks. In Alex Biryukov, editor, *FSE*, volume 4593 of *Lecture Notes in Computer Science*, pages 153–166. Springer, 2007.
2. David Chaum and Jan-Hendrik Evertse. Cryptanalysis of DES with a Reduced Number Of Rounds: Sequences of Linear Factors in Block Ciphers. In *Advances in Cryptology, CRYPTO 85*, pages 192–211. Springer-Verlag, 1986.
3. Nicolas T. Courtois. Algebraic Complexity Reduction and Cryptanalysis of GOST. Cryptology ePrint Archive, Report 2011/626, 2011. <http://eprint.iacr.org/>.
4. Nicolas T. Courtois. Security Evaluation of GOST 28147-89 in View of International Standardisation. Cryptology ePrint Archive, Report 2011/211, 2011. <http://eprint.iacr.org/>.
5. Nicolas T. Courtois and Michał Misztal. Differential Cryptanalysis of GOST. Cryptology ePrint Archive, Report 2011/312, 2011. <http://eprint.iacr.org/>.
6. Itai Dinur, Orr Dunkelman, and Adi Shamir. Improved Attacks on Full GOST. Cryptology ePrint Archive, Report 2011/558, 2011. <http://eprint.iacr.org/>.
7. Ewan Fleischmann, Michael Gorski, Jan-Hendrik Huehne, and Stefan Lucks. Key Recovery Attack on full GOST Block Cipher with Negligible Time and Memory. Presented at Western European Workshop on Research in Cryptology (WEWoRC), 2009.
8. Takanori Isobe. A Single-Key Attack on the Full GOST Block Cipher. In Antoine Joux, editor, *FSE*, volume 6733 of *Lecture Notes in Computer Science*, pages 290–305. Springer, 2011.
9. Orhun Kara. Reflection Cryptanalysis of Some Ciphers. In Dipanwita Roy Chowdhury, Vincent Rijmen, and Abhijit Das, editors, *INDOCRYPT*, volume 5365 of *Lecture Notes in Computer Science*, pages 294–307. Springer, 2008.
10. John Kelsey, Bruce Schneier, and David Wagner. Key-Schedule Cryptanalysis of IDEA, G-DES, GOST, SAFER, and Triple-DES. In Neal Kobitz, editor,

- CRYPTO*, volume 1109 of *Lecture Notes in Computer Science*, pages 237–251. Springer, 1996.
11. Youngdai Ko, Seokhie Hong, Wonil Lee, Sangjin Lee, and Ju-Sung Kang. Related Key Differential Attacks on 27 Rounds of XTEA and Full-Round GOST. In Bimal K. Roy and Willi Meier, editors, *FSE*, volume 3017 of *Lecture Notes in Computer Science*, pages 299–316. Springer, 2004.
 12. Florian Mendel, Norbert Pramstaller, and Christian Rechberger. A (Second) Preimage Attack on the GOST Hash Function. In Kaisa Nyberg, editor, *FSE*, volume 5086 of *Lecture Notes in Computer Science*, pages 224–234. Springer, 2008.
 13. Florian Mendel, Norbert Pramstaller, Christian Rechberger, Marcin Kontak, and Janusz Szmidi. Cryptanalysis of the GOST Hash Function. In David Wagner, editor, *CRYPTO*, volume 5157 of *Lecture Notes in Computer Science*, pages 162–178. Springer, 2008.
 14. National Bureau of Standards. Federal Information Processing Standard-Cryptographic Protection - Cryptographic Algorithm. GOST 28147-89, 1989.
 15. OpenSSL. A Reference Implementation of GOST. <http://www.openssl.org/source/>.
 16. Vladimir Rudskoy. On Zero Practical Significance of Key Recovery Attack on Full GOST Block Cipher with Zero Time and Memory. Cryptology ePrint Archive, Report 2010/111, 2010. <http://eprint.iacr.org/>.
 17. Haruki Seki and Toshinobu Kaneko. Differential Cryptanalysis of Reduced Rounds of GOST. In Douglas R. Stinson and Stafford E. Tavares, editors, *Selected Areas in Cryptography*, volume 2012 of *Lecture Notes in Computer Science*, pages 315–323. Springer, 2000.