

Note on the Robustness of CAESAR Candidates

Daniel Kales, Maria Eichlseder, and Florian Mendel

Graz University of Technology, Austria
daniel.kales@student.tugraz.at

Abstract. Authenticated ciphers rely on the uniqueness of the nonces to meet their security goals. In this work, we investigate the implications of reusing nonces for three third-round candidates of the ongoing CAESAR competition, namely Tiaoxin, AEGIS and MORUS. We show that an attacker that is able to force nonces to be reused can reduce the security of the ciphers with results ranging from full key-recovery to forgeries with practical complexity and a very low number of nonce-misuse queries.

Keywords: Cryptanalysis · Nonce-misuse attacks · CAESAR

1 Introduction

The CAESAR authenticated encryption competition was initiated to encourage the design and analysis of authenticated encryption ciphers. Almost all authenticated encryption schemes are nonce-based schemes [4]. A nonce is a public, unique and (usually) fixed length number, which is necessary to hide plaintext equality and introduce freshness into internal state parts. While nonces are not necessarily random or unpredictable, these are two desirable properties, as they can further reduce attack vectors. It is most important that a nonce is used only once for a specific key. If this property is violated, many new attack vectors can surface. For example, a single known plaintext-ciphertext pair with the same nonce is enough to forge valid ciphertext-tag pairs for arbitrary plaintexts when using AES-GCM [1].

In this brief note, we analyze the impact of using a repeated nonce for three of the third-round candidates of the CAESAR competition. We target the ciphers Tiaoxin [3], AEGIS [6], and MORUS [5], which have a similar structure and thus present similar angles of attack. We consider a chosen-plaintext attacker who can force a small number of encryptions under the same nonce. Note that such a nonce-misuse attacker voids the security claims of the analyzed designs. We show that all three candidates investigated in this work are vulnerable to the proposed attack scenario, and propose forgery, state-recovery, or key-recovery attacks. We tested all attacks with practical implementations [2]. Finally, we emphasize that none of our attacks threatens the security of the schemes in a normal setting as specified by the designers.

Outline. We briefly describe the ciphers in Section 2, but refer to the respective design documents for a full specification. A detailed description of the attacks can be found in Section 3, and a summary of the results in Section 4.

2.2 Description of AEGIS

AEGIS is a family of authenticated encryption algorithms by Wu and Preneel [6]. We only focus on the AEGIS-128 version of the algorithm, but the idea of the attack can be extended easily to other versions. The internal state of AEGIS-128 consists of 5 words of 16 bytes each. The state update function consumes a plaintext block of 16 bytes and produces the new internal state as in Figure 2.




Fig. 2: The StateUpdate function of AEGIS-128. (A) represents one AES round. A more detailed explanation can be found in the design document [6].

Like Tiaoxin, the internal state is initialized with the secret key and the public nonce. To complete the initialization, the state is updated with a fixed number of iterations of the StateUpdate function with constants in place of the message words. The encryption phase of AEGIS-128 can then be summarized as follows:

$$\begin{aligned} & \text{for } i = 1 \text{ to } m \\ & \quad C_i = M_i \oplus S_{i,1} \oplus S_{i,4} \oplus (S_{i,2} \wedge S_{i,3}) \\ & \quad S_{i+1} = \text{StateUpdate}(S_i, M_i) \end{aligned}$$

In contrast to the state update function of Tiaoxin-346 (Figure 1), the state update function of AEGIS-128 is not easily invertible. This means that recovering the internal state does not allow us to recover the secret key.

2.3 Description of MORUS

MORUS is a family of authenticated encryption ciphers by Wu and Huang [5]. We only focus on the MORUS-640-128 version of the cipher, but as with the two other ciphers before, the ideas of the attack are easily extended to the other versions of the cipher. The internal state of MORUS-640-128 consists of 5 words of 16 bytes each. The state update function consumes a plaintext block of 16 bytes and updates the internal state. A graphical representation of the state update function of MORUS-640-128 can be seen in Figure 3. In contrast to Tiaoxin-346 and AEGIS-128, the state update function is not based on AES, but instead is a custom construction using ANDs, rotations and XORs (ARX). It uses two kind of rotations internally: A rotation of one full message word with different constant rotation offsets (denoted by $\lll w_i$) and a second kind of rotation, where one 16 byte internal state is split into four 4-byte blocks, which are rotated independently (denoted by $\lll^* b_i$).

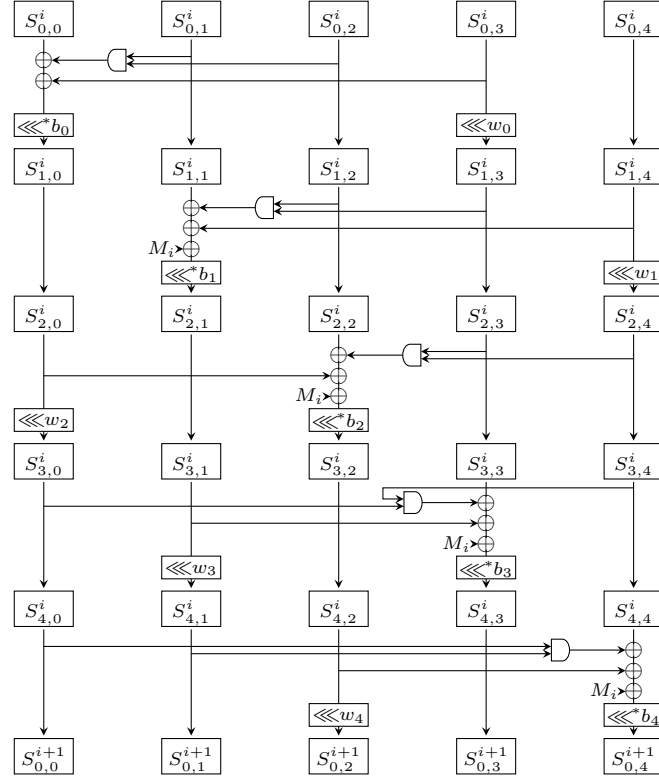


Fig. 3: The state update function of MORUS-640-128. A more detailed explanation can be found in the design document [5].

As with the two previous ciphers, the internal state is initialized with the secret key and the public nonce. The internal state is then updated with a fixed number of calls to the internal state function (with constants in place of the message words). After the final initialization round, the secret key is again XORed to one part of the internal state. This results in the initial state used for encryption. We will now shortly summarize the encryption phase of MORUS-640-128:

for $i = 1$ to m

$$C_i = M_i \oplus S_0^i \oplus (S_1^i \lll 96) \oplus (S_2^i \wedge S_3^i)$$

$$S_{i+1} = \text{StateUpdate}(S_i, M_i)$$

Like the state update function of Tiaoxin-346, the state update function of MORUS is invertible given the internal state. The initialization phase, however, is not, due to the fact that the key is XORed to a part of the state after the initialization rounds.

3 Attack on the CAESAR Candidates

3.1 Similarities between Candidates and Basic Idea

All ciphers analyzed in this note use a similar approach to encrypt: The internal state is initialized with the IV and key, and a few rounds of the state update function are applied. For encryption, parts of the internal state (a, b, c and d) are combined with a mix of linear and non-linear operations to obtain a key-stream KS that is XORed with the plaintext P to generate the ciphertext C :

$$\begin{aligned}KS &= a \oplus b \oplus (c \wedge d) \\C &= P \oplus KS.\end{aligned}\tag{1}$$

After the encryption phase, the message tag is generated by applying a few rounds of the state update function, before finally combining parts of the internal state. By keeping a, b and c constant and varying d , we can recover c with the following observation. Since XOR and AND operate on single bits, if we manage to get the values d to include both 0 and 1 for a range of trials, we can observe the following property for one bit of the xor of two trials:

$$\begin{aligned}KS_0 \oplus KS_1 &= a \oplus b \oplus (c \wedge d_0) \oplus a \oplus b \oplus (c \wedge d_1) \\&= (c \wedge d_0) \oplus (c \wedge d_1) = (c \wedge 0) \oplus (c \wedge 1) = 0 \oplus c = c.\end{aligned}$$

If the values for the bit of d are equal, the result will always be 0, since all inputs to the xor are equal. This means we need at least one pair of trials that have different values for d . In the following section, we present attacks that make use of these properties and allow recovery of the internal state.

3.2 Application to Tiaoxin

The encryption of Tiaoxin follows the structure in Equation (1):

$$C_i^1 = T_6[0] \oplus T_4[2] \oplus T_3[1] \oplus (T_6[5] \wedge T_3[2]).$$

We attack the generation of C_2^1 with the following conditions: The state words $T_6[0], T_4[2]$ and $T_3[1]$ of S_3 are kept constant, while a difference is introduced in $T_3[2]$. A graphical representation of the attack can be seen in Figure 4.

However, due to the AES call in the state update function, we cannot exactly predict the final difference in $T_3[2]$ that is introduced by our chosen plaintext difference Δ . This means we need to perform a higher number of trials to achieve a high probability that at least one pair of message differences has a difference in one bit of d . To get a reliable result for the attack, we use 128 random plaintexts. Since the input to the AES block is random, and AES is not biased with respect to its output, we get a probability of $\frac{1}{2}$ for a single bit to be either 0 or 1. This means the probability that one bit of the output is the same for all 128 messages is 2^{-127} . Combining the results for each of the 128 bits of the AES output block gives us a total success probability of $(1 - 2^{-127})^{128} \approx 1$.

With smart choice of the message words, we can achieve the required properties of the attack. If we choose a difference (referred to as Δ in the following) for M_0^0 and no difference for M_0^1 , this results in the same difference Δ in M_0^2 . If we use the same values for the next message block M_1^0 , M_1^1 and M_1^2 , the differences cancel out for the most part. For the third round state S_3 , only $T_3[2]$ and $T_6[2]$ contain any differences (see Figure 4).

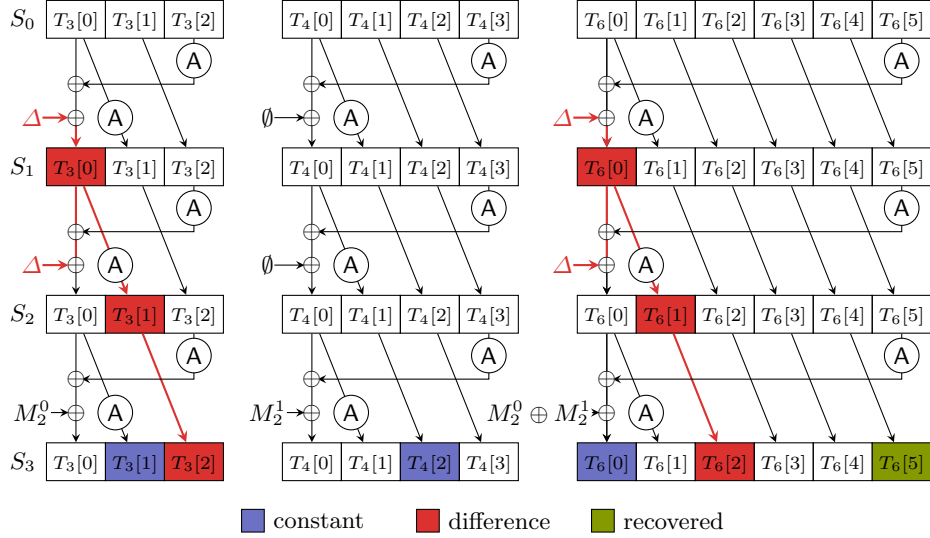


Fig. 4: Propagation of differences in the attack scenario for Tiaoxin.

Using this attack outline, we can recover one part of the internal state T_6 after three rounds of encryption. Furthermore, we can repeat our attack for later rounds: If we execute n rounds of encryption with a constant plaintext before our attack, we can recover $T_6[5]$ of state S_{3+n} . Additionally, the structure of the update function tells us that $T_6[5]$ of S_4 equals $T_6[4]$ of S_3 and a similar relation can be found for all parts of T_6 of S_3 . So we can repeat our initial attack six times with offsets $0, \dots, 5$ and recover the internal state T_6 . As explained in Section 2, this allows us to fully recover the secret key.

3.3 Application to AEGIS

The attack against AEGIS works according to the same general principle as the attack against Tiaoxin. Again, the structure of the encryption of AEGIS follows the structure in Equation (1):

$$C_i = M_i \oplus S_{i,1} \oplus S_{i,4} \oplus (S_{i,2} \wedge S_{i,3}).$$

Figure 5 depicts the structure of the attack for AEGIS. We introduce a difference Δ in the message word P_1 and use the same difference Δ in P_2 to cancel most of the effect. The only remaining difference propagates to the state word $S_{2,1}$ and can be recovered from the encryption equation of C_2 , since the state words $S_{2,2}$, $S_{2,3}$ and $S_{2,4}$ are constant. The difference propagates to $S_{3,2}$ and since the difference in $S_{3,1}$ is equal to the previously recovered difference $S_{2,1}$ we can use our approach to recover the state $S_{3,3}$.

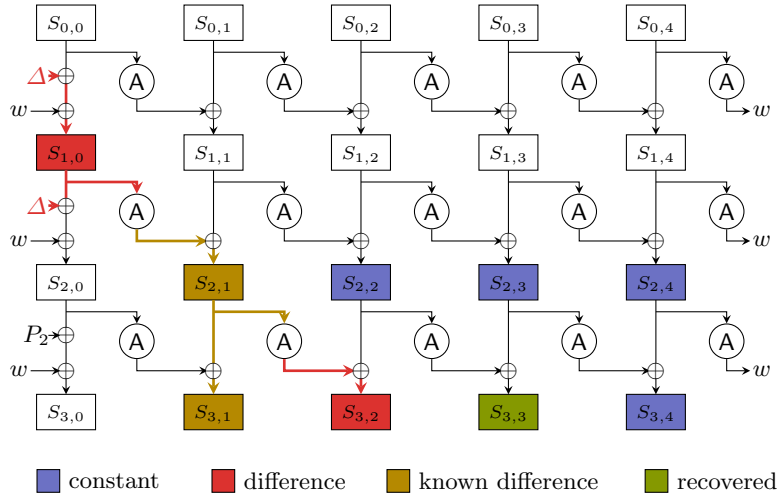


Fig. 5: Propagation of differences in the attack scenario for AEGIS.

To recover more parts of the state, we have to repeat our attack one round later. Recovering $S_{4,3}$ in addition to $S_{3,3}$ enables us to calculate $S_{3,2}$ as can be seen in Figure 6. However, the message word P_1 has an influence on $S_{3,2}$ and therefore needs to be constant. This means we need to choose a fixed P_1 for our attack and can only recover a state for this choice of P_1 . Repeating this attack for $S_{5,3}$ and $S_{6,3}$ in order to recover $S_{3,1}$ and $S_{3,0}$ in turn means we also have to fix P_2 and P_3 for our attack, meaning we can recover the state for a 3-block chosen-plaintext prefix. This enables us to create forged ciphertexts for messages starting with the same 3-block chosen-plaintext prefix under the same nonce.




Fig. 6: Combining two recovered states to recover more parts of the internal state.

3.4 Application to MORUS

The attack on MORUS is based on the same principle as the previous attacks. Again, the structure of the encryption of MORUS follows the structure in (1):

$$C_i = M_i \oplus S_0^i \oplus (S_1^i \lll 96) \oplus (S_2^i \wedge S_3^i).$$

Since the state update function does not use a round of AES, but XOR, AND and rotation operations instead, we can propagate desirable differences more easily. Instead of using about 2^7 encryption oracle calls for one part of the state, we can reduce the overall encryption oracle calls to below 32.

To recover the first part of the internal state, we choose a difference of $\Delta = 1^{128}$ for the first plaintext block. This difference spreads to $S_{0,1}^1$ and $S_{0,2}^1$ while $S_{0,3}^1$ stays constant, because the difference introduced by Δ is canceled between $S_{3,3}^0$ and $S_{4,3}^0$ (see Figure 7). This allows us to recover $S_{0,3}^1$ with only 2 encryptions.

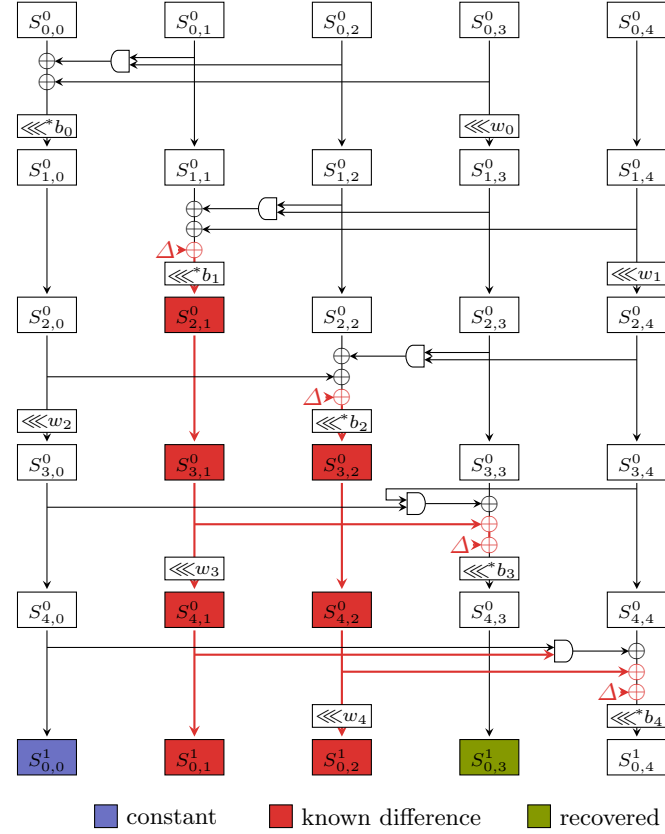


Fig. 7: Propagation of differences to recover the first part of the MORUS state.

We can now recover an additional part of the state in a similar fashion. Taking advantage of the rotations in the cipher, we are able to generate a plaintext block that keeps half of $S_{0,2}^1$ constant, while introducing a difference in the corresponding half of $S_{0,3}^1$ (see Figure 8). Repeating this process for the other half allows us to fully recover $S_{0,2}^1$ with only 2 additional encryptions.

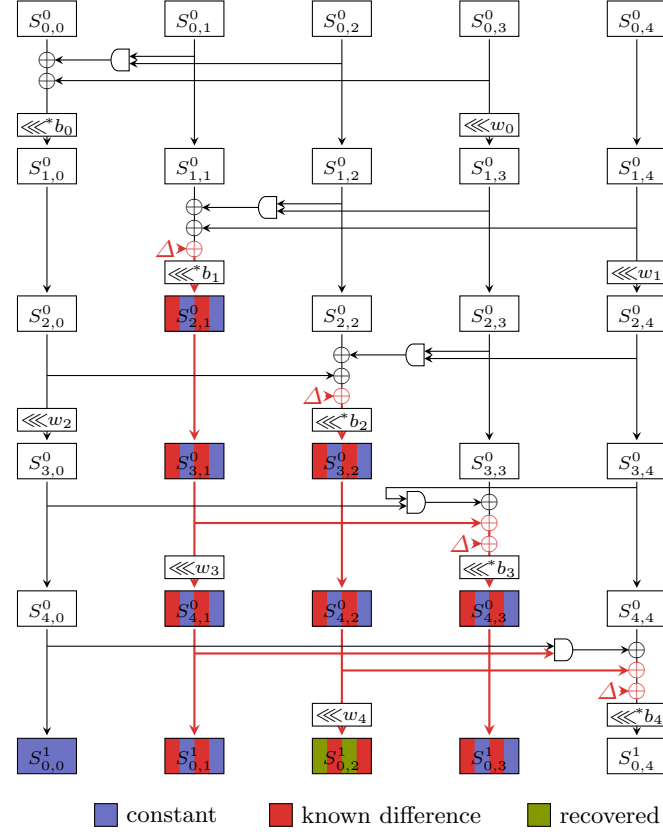


Fig. 8: Propagation of differences to recover the second part of the MORUS state.

For the other state blocks we have to attack the second round of encryption. With the knowledge of the recovered state $S_{0,2}^1$ we can now calculate a plaintext block M_1 , so that the resulting state of $S_{0,2}^1$ is equal to 0^{128} . This disables the AND gate in the first sub-step of the round update, meaning $S_{0,1}^1$ has no influence on $S_{0,0}^2$. If we now calculate a second plaintext M_1 , so that $S_{0,2}^1$ is equal to $1^{32}||0^{96}$, we only allow one quarter of $S_{0,1}^1$ to be propagated into $S_{0,0}^2$. Due to the nature of the state rotation and its different rotation values for each state part, the newly introduced quarter block of ones only influences a specific quarter of $S_{0,0}^2$, but does not influence the same quarter in the other state parts $S_{0,1-3}^2$. This

means by comparing the output of the second-round encryptions of these two plaintexts, we can recover one quarter of $S_{0,1}^1$, and repeating this process 3 more times while setting different quarters of $S_{0,2}^1$ to 1^{32} allows us to fully recover $S_{0,1}^1$. The state $S_{0,0}^1$ is then easily recovered with one known plaintext-ciphertext pair, since it is the only unknown part of the encryption equation of MORUS.

The final state part $S_{0,4}^1$ is then recovered by looking at the first and second round of encryption. We can recover the state $S_{0,1}^2$ by performing our attacks with an offset of one constant plaintext block. Now we can obtain $S_{0,4}^1$ by considering the equation given by the second column of the state update function, in which we already know the value of every variable except $S_{0,4}^1$:

$$S_{0,1}^2 = S_{0,1}^1 \oplus (S_{0,2}^1 \wedge (S_{0,3}^1 \lll w_0)) \oplus S_{0,4}^1.$$

We can then apply the inverse round update function to get the state S^0 , but we cannot recover the secret key, since the initialization phase cannot be reversed without previous knowledge of the secret key. Nevertheless, knowledge of the internal state allows forgery of ciphertext-tag pairs for arbitrary plaintexts.

3.5 Practical Implementation

We implemented all three attacks and verified the results using the designers’ reference implementations of the respective cipher. The run-time of all of the three attack implementations is negligibly short (< 1 second each) and the success probability is essentially 100%. The source code is available on GitHub [2].

4 Summary and Conclusion

Our results are summarized in Table 1, and underline the importance of correct nonce usage. The designers of all three algorithms stress that their security claims are void in case of nonce misuse and this note confirms this assessment. Clearly, misconfiguration or wrong usage of a cipher can weaken the security immensely, and even a very small number of nonce misuses can suffice for key recovery or forgery. Only a few CAESAR candidates make any security claims in a nonce misuse scenario. However, among the nonce-based schemes, there seems to be significant variance in the practical impact of “small-scale” nonce misuse, ranging from drastic global impact (e.g., key recovery and universal forgery as in the presented attacks) to much more local impact (e.g., confidentiality loss for some blocks of the misuse plaintext in various duplex-based schemes).

Table 1: Summary of attack complexity and results.

Cipher	Nonce-Misuse Queries	Attack
Tiaoxin-346	2^{10}	Key recovery
AEGIS-128	2^9	Almost universal forgery
MORUS-640-128	2^5	Universal forgery

References

1. Böck, H., Zauner, A., Devlin, S., Somorovsky, J., Jovanovic, P.: Nonce-disrespecting adversaries: Practical forgery attacks on GCM in TLS. In: USENIX Workshop on Offensive Technologies – WOOT 16. USENIX Association (2016)
2. Kales, D.: Implementation of the nonce-misuse attacks for the CAESAR competition candidates AEGIS, Tiaoxin and MORUS. GitHub repository (January 2017), <https://github.com/dkales/caesar-nonce-misuse>
3. Nikolić, I.: Tiaoxin-346. Submission to the CAESAR Competition (2016), <http://competitions.cr.yp.to/round3/tiaoxinv21.pdf>
4. Rogaway, P.: Nonce-based symmetric encryption. In: Roy, B.K., Meier, W. (eds.) Fast Software Encryption – FSE 2004. LNCS, vol. 3017, pp. 348–359. Springer (2004)
5. Wu, H., Huang, T.: The authenticated cipher MORUS (v2). Submission to the CAESAR Competition (2016), <http://competitions.cr.yp.to/round3/morusv2.pdf>
6. Wu, H., Preneel, B.: AEGIS: A fast authenticated encryption algorithm (v1.1). Submission to the CAESAR Competition (2016), <http://competitions.cr.yp.to/round3/aegisv11.pdf>