

**AEGIS:
A Fast Authenticated Encryption
Algorithm (v1.1)**

Designers and Submitters: Hongjun Wu¹, Bart Preneel²

¹Division of Mathematical Sciences
Nanyang Technological University
wuhongjun@gmail.com

²Dept. Elektrotechniek-ESAT/COSIC
KU Leuven and iMinds, Ghent
bart.preneel@esat.kuleuven.be

2016.09.15

Contents

1	Introduction	3
2	Specification	5
2.1	Recommended parameter sets	5
2.2	Operations, Variables and Functions	5
2.2.1	Operations	5
2.2.2	Variables and constants	6
2.2.3	Functions	6
2.3	AEGIS-128	7
2.3.1	The state update function of AEGIS-128	7
2.3.2	The initialization of AEGIS-128	7
2.3.3	Processing the authenticated data	8
2.3.4	The encryption of AEGIS-128	8
2.3.5	The finalization of AEGIS-128	8
2.3.6	The decryption and verification of AEGIS-128	9
2.4	AEGIS-256	9
2.4.1	The state update function of AEGIS-256	9
2.4.2	The initialization of AEGIS-256	10
2.4.3	Processing the authenticated data	10
2.4.4	The encryption of AEGIS-256	10
2.4.5	The finalization of AEGIS-256	11
2.5	AEGIS-128L	11
2.5.1	The state update function of AEGIS-128L	11
2.5.2	The initialization of AEGIS-128L	12
2.5.3	Processing the authenticated data	12
2.5.4	The encryption of AEGIS-128L	12
2.5.5	The finalization of AEGIS-128L	13
3	Security Goals	14
4	Security Analysis	15
4.1	The security of the initialization	15
4.2	The security of the encryption process	16
4.3	The security of message authentication	17

4.3.1	Recovering key or state.	17
4.3.2	Internal collisions.	17
4.4	Other attacks	18
5	Features	20
6	The Performance of AEGIS	21
6.1	Software Performance	21
6.2	Hardware Performance	22
7	Design Rationale	23
7.1	Reply to the Comments of CAESAR Committee Members	24
8	Changes in the Third Round Submission	25
9	No Hidden Weakness	26
10	Intellectual property	27
11	Consent	28
12	Acknowledgements	29

Chapter 1

Introduction

The protection of a message typically requires the protection of both confidentiality and authenticity. There are two main approaches to authenticate and encrypt a message. One approach is to treat the encryption and authentication separately. The plaintext is encrypted with a block cipher or stream cipher, and a MAC algorithm is used to authenticate the ciphertext. For example, we may apply AES [20] in CBC mode [21] to the plaintext, then apply AES-CMAC [25] (or Pelican MAC [7] or HMAC [22]) to the ciphertext to generate an authentication tag. This approach is relatively easy to analyze since the security of authentication and encryption can be analyzed almost separately. Bellare and Namprempre have performed a detailed analysis of this type of authenticated encryption for randomized encryption [2]. Another approach is to apply an integrated authenticated encryption algorithm to the message; one can expect that this is more efficient since authentication and encryption can share part of the computation.

There are three approaches to design an integrated authenticated encryption algorithm. The first approach is to use a block cipher in a special mode (the block cipher is treated as a black box). The research on this approach started about ten years ago [11, 14, 16]. There are now two NIST recommended modes of operation for authenticated encryption, namely, CCM [23] and GCM [24]. OCB [27, 28, 17] is a widely known authenticated encryption mode, and OCB2 is an ISO standard. The second approach is to use a stream cipher (the stream cipher is treated as a black box). The keystream is divided into two parts: one part for encryption and another part for authentication. A typical example of this approach is Grain-128a [1]. The third approach is to design dedicated authenticated encryption algorithms. In this approach, a message is used to update the state of the cipher, and message authentication can be achieved almost for free. Two examples of this approach are Helix [10] and Phelix [29]. The attack against Phelix [30] shows that it is unlikely that this type of authenticated encryption algorithm can withstand nonce-reuse attacks if it requires much less computation than a block cipher.

In this report, we propose a dedicated authenticated encryption algorithm

AEGIS following the third approach above. AEGIS is constructed from the AES encryption round function (not the last round). AEGIS-128L uses eight AES round functions to process a 32-byte message block (one step). AEGIS-128 processes a 16-byte message block with 5 AES round functions, and AEGIS-256 uses 6 AES round functions. The computational cost of AEGIS is about half that of AES. AEGIS is very fast. On the Intel Sandy Bridge processor Core-i5, the encryption speeds of AEGIS-128L, AEGIS-128 and AEGIS-256 are about 0.48 cpb, 0.66 cpb and 0.70 cpb, respectively. On the Intel Haswell processor Core-i7, the encryption speeds of AEGIS-128L, AEGIS-128 and AEGIS-256 are about 0.37 cpb, 0.60 cpb and 0.62 cpb, respectively. The speed of AEGIS-128L is much faster than that of AES in counter (CTR) mode, and are about 8 times that of AES encryption in CBC mode. AEGIS offers a very high security. As long as the nonce is not reused, it is impossible to recover the AEGIS state and key faster than exhaustive key search (under the assumption that a 128-bit authentication tag is used, and the forgery attack is not successful by repeating the attack). AEGIS is suitable for network communication since AEGIS can protect a packet while leaving the packet header (associated data) unencrypted.

The specifications of AEGIS-128 and AEGIS-256 were published at SAC 2013 [31]. AEGIS-128L is introduced into this submission.

Chapter 2

Specification

The specification of AEGIS-128, AEGIS-256 and AEGIS-128L are given in this chapter.

2.1 Recommended parameter sets

- Primary Recommendation: AEGIS-128L
128-bit key, 128-bit nonce, 1024-bit state, 128-bit tag
Reason: AEGIS-128L is the fastest AEGIS algorithm.
Use Case 2: High-performance applications
- Secondary Recommendation: AEGIS-128
128-bit key, 128-bit nonce, 640-bit state, 128-bit tag
Reason: The state of AEGIS-128 is smaller than that of AEGIS-128L.
Use Case 2: High-performance applications
- Tertiary Recommendation: AEGIS-256
256-bit key, 256-bit nonce, 768-bit state, 128-bit tag
Reason: AEGIS-256 uses 256-bit secret key.
Use Case 2: High-performance applications

2.2 Operations, Variables and Functions

The operations, variables and functions used in AEGIS are defined below.

2.2.1 Operations

The following operations are used in AEGIS:

\oplus	:	bit-wise exclusive OR
$\&$:	bit-wise AND
\parallel	:	concatenation
$\lceil x \rceil$:	ceiling operation, $\lceil x \rceil$ is the smallest integer not less than x

2.2.2 Variables and constants

The following variables and constants are used in AEGIS:

AD	:	associated data (this data will not be encrypted or decrypted).
AD_i	:	a 16-byte associated data block (the last block may be a partial block).
$adlen$:	bit length of the associated data with $0 \leq adlen < 2^{64}$.
C	:	ciphertext.
C_i	:	a 16-byte ciphertext block (the last block may be a partial block).
$const$:	a 32-byte constant in the hexadecimal format; $const = 00 \parallel 01 \parallel 01 \parallel 02 \parallel 03 \parallel 05 \parallel 08 \parallel 0d \parallel 15 \parallel 22 \parallel 37 \parallel 59 \parallel 90 \parallel e9 \parallel 79 \parallel 62 \parallel db \parallel 3d \parallel 18 \parallel 55 \parallel 6d \parallel c2 \parallel 2f \parallel f1 \parallel 20 \parallel 11 \parallel 31 \parallel 42 \parallel 73 \parallel b5 \parallel 28 \parallel dd$. This is the Fibonacci sequence modulo 256.
$const_0$:	first 16 bytes of $const$.
$const_1$:	last 16 bytes of $const$.
IV_{128}	:	128-bit initialization vector of AEGIS-128.
IV_{256}	:	256-bit initialization vector of AEGIS-256.
$IV_{256,0}$:	first half of IV_{256} .
$IV_{256,1}$:	second half of IV_{256} .
K_{128}	:	128-bit key of AEGIS-128.
K_{256}	:	256-bit key of AEGIS-256.
$K_{256,0}$:	first half of K_{256} .
$K_{256,1}$:	second half of K_{256} .
$msglen$:	bit length of the plaintext/ciphertext with $0 \leq msglen < 2^{64}$.
m_i	:	a 16-byte data block.
P	:	plaintext.
P_i	:	a 16-byte plaintext block (the last block may be a partial block).
S_i	:	state at the beginning of the i th step.
$S_{i,j}$:	j -th 16-byte element of the state S_i . For AEGIS-128, $0 \leq j \leq 4$; for AEGIS-256, $0 \leq j \leq 5$; for AEGIS-128L, $0 \leq j \leq 7$.
T	:	authentication tag.
t	:	bit length of the authentication tag with $64 \leq t \leq 128$.
u	:	$u = \lceil \frac{adlen}{128} \rceil$.
v	:	$v = \lceil \frac{msglen}{128} \rceil$.
u_L	:	$u_L = \lceil \frac{adlen}{256} \rceil$.
v_L	:	$v_L = \lceil \frac{msglen}{256} \rceil$.

2.2.3 Functions

The AES encryption round function (not the last round) is used in AEGIS:

$AESRound(A, B)$: A is the 16-byte state, B is the 16-byte round key. This function mapping 2 16-byte inputs to a 16-byte output can be implemented efficiently on recent x86 processors using the AES instruction `_m128_aesenc_si128(A, B)`, where A and B are two 128-bit integers `_m128i`.

2.3 AEGIS-128

We first describe AEGIS-128 since its structure is the simplest among those three algorithms. With a 128-bit key and a 128-bit initialization vector, AEGIS-128 encrypts and authenticates a message. The associated data length and the plaintext length are less than 2^{64} bits. The authentication tag length is less than or equal to 128 bits. We strongly recommend the use of a 128-bit tag.

2.3.1 The state update function of AEGIS-128

The state update function updates the 80-byte state S_i with a 16-byte message block m_i . $S_{i+1} = \text{StateUpdate128}(S_i, m_i)$ is given as follows:

$$\begin{aligned} S_{i+1,0} &= AESRound(S_{i,4}, S_{i,0} \oplus m_i); \\ S_{i+1,1} &= AESRound(S_{i,0}, S_{i,1}); \\ S_{i+1,2} &= AESRound(S_{i,1}, S_{i,2}); \\ S_{i+1,3} &= AESRound(S_{i,2}, S_{i,3}); \\ S_{i+1,4} &= AESRound(S_{i,3}, S_{i,4}); \end{aligned}$$

The state update function is shown in Fig. 2.1 :




Figure 2.1: The state update function of AEGIS-128. R indicates the AES encryption round function without XORing the round key and w is a temporary 16-byte word.

2.3.2 The initialization of AEGIS-128

The initialization of AEGIS-128 consists of loading the key and IV into the state, and running the cipher for 10 steps with the key and IV being used as message.

1. Load the key and IV into the state as follows:

$$\begin{aligned}
S_{-10,0} &= K_{128} \oplus IV_{128}; \\
S_{-10,1} &= const_1; \\
S_{-10,2} &= const_0; \\
S_{-10,3} &= K_{128} \oplus const_0; \\
S_{-10,4} &= K_{128} \oplus const_1;
\end{aligned}$$

2. For $i = -5$ to -1 , $m_{2i} = K_{128}$; $m_{2i+1} = K_{128} \oplus IV_{128}$;

3. For $i = -10$ to -1 , $S_{i+1} = \text{StateUpdate128}(S_i, m_i)$;

2.3.3 Processing the authenticated data

After the initialization, the associated data AD is used to update the state.

1. If the last associated data block is not a full block, use 0 bits to pad it to 128 bits, and the padded full block is used to update the state. Note that if $adlen = 0$, the state will not be updated.
2. For $i = 0$ to $\lceil \frac{adlen}{128} \rceil - 1$, we update the state:

$$S_{i+1} = \text{StateUpdate128}(S_i, AD_i);$$

2.3.4 The encryption of AEGIS-128

After processing the associated data, at each step of the encryption, a 16-byte plaintext block P_i is used to update the state, and P_i is encrypted to C_i .

1. If the last plaintext block is not a full block, use 0 bits to pad it to 128 bits, and the padded full block is used to update the state. Note that for the last block, only the original plaintext bits will be encrypted. Note that if $msglen = 0$, the state will not be updated, and there is no encryption.
2. Let $u = \lceil \frac{adlen}{128} \rceil$ and $v = \lceil \frac{msglen}{128} \rceil$. For $i = 0$ to $v - 1$, we perform encryption and update the state:

$$\begin{aligned}
C_i &= P_i \oplus S_{u+i,1} \oplus S_{u+i,4} \oplus (S_{u+i,2} \& S_{u+i,3}); \\
S_{u+i+1} &= \text{StateUpdate128}(S_{u+i}, P_i);
\end{aligned}$$

2.3.5 The finalization of AEGIS-128

After encrypting all the plaintext blocks, we generate the authentication tag using seven more steps. The length of the associated data and the length of the message are used to update the state.

1. Let $tmp = S_{u+v,3} \oplus (adlen \parallel msglen)$, where $adlen$ and $msglen$ are represented as 64-bit integers.

- For $i = u + v$ to $u + v + 6$, we update the state:

$$S_{i+1} = \text{StateUpdate128}(S_i, tmp) ;$$

- We generate the authentication tag from the state S_{u+v+7} as follows:

$$T' = \bigoplus_{i=0}^4 S_{u+v+7,i} ;$$

The authentication tag T consists of the first t bits of T' .

2.3.6 The decryption and verification of AEGIS-128

The exact values of key size, IV size, and tag size should be known to the decryption and verification processes. The decryption starts with the initialization and the processing of authenticated data. Then the ciphertext is decrypted as follows:

- If the last ciphertext block is not a full block, decrypt only the partial ciphertext block. The partial plaintext block is padded with 0 bits, and the padded full plaintext block is used to update the state.
- For $i = 0$ to $v - 1$, we perform decryption and update the state.

$$\begin{aligned} P_i &= C_i \oplus S_{u+i,1} \oplus S_{u+i,4} \oplus (S_{u+i,2} \& S_{u+i,3}) ; \\ S_{u+i+1} &= \text{StateUpdate128}(S_{u+i}, P_i) ; \end{aligned}$$

The finalization in the decryption process is the same as that in the encryption process. We emphasize that if the verification fails, the ciphertext and the newly generated authentication tag should not be given as output; otherwise, the state of AEGIS-128 is vulnerable to known-plaintext or chosen-ciphertext attacks (using a fixed IV). This requirement also applies to AEGIS-256.

2.4 AEGIS-256

In this section, we describe AEGIS-256. With a 256-bit key and a 256-bit initialization vector, AEGIS-256 encrypts and authenticates a message. The associated data length and the plaintext length are less than 2^{64} bits. The authentication tag length is less than or equal to 128 bits. We strongly recommend the use of a 128-bit tag.

2.4.1 The state update function of AEGIS-256

The state update function updates the 96-byte state S_i with a 16-byte message block m_i . $S_{i+1} = \text{StateUpdate256}(S_i, m_i)$ is illustrated as follows:

$$\begin{aligned}
S_{i+1,0} &= \text{AESRound}(S_{i,5}, S_{i,0} \oplus m_i); \\
S_{i+1,1} &= \text{AESRound}(S_{i,0}, S_{i,1}); \\
S_{i+1,2} &= \text{AESRound}(S_{i,1}, S_{i,2}); \\
S_{i+1,3} &= \text{AESRound}(S_{i,2}, S_{i,3}); \\
S_{i+1,4} &= \text{AESRound}(S_{i,3}, S_{i,4}); \\
S_{i+1,5} &= \text{AESRound}(S_{i,4}, S_{i,5});
\end{aligned}$$

2.4.2 The initialization of AEGIS-256

The initialization of AEGIS-256 consists of loading the key and IV into the state, and running the cipher for 16 steps with the key and IV being used as message.

1. Load the key and IV into the state as follows:

$$\begin{aligned}
S_{-16,0} &= K_{256,0} \oplus IV_{256,0}; \\
S_{-16,1} &= K_{256,1} \oplus IV_{256,1}; \\
S_{-16,2} &= \text{const}_1; \\
S_{-16,3} &= \text{const}_0; \\
S_{-16,4} &= K_{256,0} \oplus \text{const}_0; \\
S_{-16,5} &= K_{256,1} \oplus \text{const}_1;
\end{aligned}$$

2. For $i = -4$ to -1 ,

$$\begin{aligned}
m_{4i} &= K_{256,0}; \\
m_{4i+1} &= K_{256,1}; \\
m_{4i+2} &= K_{256,0} \oplus IV_{256,0}; \\
m_{4i+3} &= K_{256,1} \oplus IV_{256,1}.
\end{aligned}$$

3. For $i = -16$ to -1 , $S_{i+1} = \text{StateUpdate256}(S_i, m_i)$;

2.4.3 Processing the authenticated data

After the initialization, the associated data AD is used to update the state.

1. If the last associated data block is not a full block, use 0 bits to pad it to 128 bits, and the padded full block is used to update the state. Note that if $adlen = 0$, the state will not be updated.
2. For $i = 0$ to $\lceil \frac{adlen}{128} \rceil - 1$, we update the state.

$$S_{i+1} = \text{StateUpdate256}(S_i, AD_i);$$

2.4.4 The encryption of AEGIS-256

After processing the associated data, at each step of the encryption, a 16-byte plaintext block P_i is used to update the state, and P_i is encrypted to C_i .

1. If the last plaintext block is not a full block, use 0 bits to pad it to 128 bits, and the padded full block is used to update the state. Note that for the last block, only the original plaintext bits will be encrypted. Note that if $msglen = 0$, the state will not be updated, and there is no encryption.
2. Let $u = \lceil \frac{adlen}{128} \rceil$ and $v = \lceil \frac{msglen}{128} \rceil$. For $i = 0$ to $v - 1$, we perform encryption and update the state:

$$\begin{aligned} C_i &= P_i \oplus S_{u+i,1} \oplus S_{u+i,4} \oplus S_{u+i,5} \oplus (S_{u+i,2} \& S_{u+i,3}) ; \\ S_{u+i+1} &= \text{StateUpdate256}(S_{u+i}, P_i) ; \end{aligned}$$

2.4.5 The finalization of AEGIS-256

After encrypting all the plaintext blocks, we generate the authentication tag using seven more steps. The length of the associated data and the length of the message are used to update the state.

1. Let $tmp = S_{u+v,3} \oplus (adlen \parallel msglen)$, where $adlen$ and $msglen$ are represented as 64-bit integers.
2. For $i = u + v$ to $u + v + 6$, we update the state:

$$S_{i+1} = \text{StateUpdate256}(S_i, tmp) ;$$

3. We generate the authentication tag from the state S_{u+v+7} as follows:

$$T' = \bigoplus_{i=0}^5 S_{u+v+7,i} ;$$

The authentication tag T consists of the first t bits of T' .

2.5 AEGIS-128L

In this section, we describe AEGIS-128L with 128-byte state. Its key, IV and tag sizes are the same as that of AEGIS-128. AEGIS-128L encrypts and authenticates a message with length less than 2^{64} bits.

2.5.1 The state update function of AEGIS-128L

The state update function AEGIS-128L updates the 128-byte state S_i with two 16-byte message blocks m_a and m_b . $S_{i+1} = \text{StateUpdate128L}(S_i, m_a, m_b)$ is illustrated as follows:

$$\begin{aligned}
S_{i+1,0} &= \text{AESRound}(S_{i,7}, S_{i,0} \oplus m_a); \\
S_{i+1,1} &= \text{AESRound}(S_{i,0}, S_{i,1}); \\
S_{i+1,2} &= \text{AESRound}(S_{i,1}, S_{i,2}); \\
S_{i+1,3} &= \text{AESRound}(S_{i,2}, S_{i,3}); \\
S_{i+1,4} &= \text{AESRound}(S_{i,3}, S_{i,4} \oplus m_b); \\
S_{i+1,5} &= \text{AESRound}(S_{i,4}, S_{i,5}); \\
S_{i+1,6} &= \text{AESRound}(S_{i,5}, S_{i,6}); \\
S_{i+1,7} &= \text{AESRound}(S_{i,6}, S_{i,7});
\end{aligned}$$

2.5.2 The initialization of AEGIS-128L

The initialization of AEGIS-128L consists of loading the key and IV into the state, and running the cipher for 10 steps with the key and IV being used as message.

1. Load the key and IV into the state as follows:

$$\begin{aligned}
S_{-10,0} &= K_{128} \oplus IV_{128}; \\
S_{-10,1} &= \text{const}_1; \\
S_{-10,2} &= \text{const}_0; \\
S_{-10,3} &= \text{const}_1; \\
S_{-10,4} &= K_{128} \oplus IV_{128}; \\
S_{-10,5} &= K_{128} \oplus \text{const}_0; \\
S_{-10,6} &= K_{128} \oplus \text{const}_1; \\
S_{-10,7} &= K_{128} \oplus \text{const}_0;
\end{aligned}$$

2. For $i = -10$ to -1 , $S_{i+1} = \text{StateUpdate128L}(S_i, IV_{128}, K_{128})$.

2.5.3 Processing the authenticated data

After the initialization, the associated data AD is used to update the state.

1. If the length of associated data is not a multiple of 256 bits, use 0 bits to pad it to $\lceil \frac{adlen}{256} \rceil \times 256$ bits, and the padded associated data is used to update the state. Note that if $adlen = 0$, the state will not get updated.
2. For $i = 0$ to $\lceil \frac{adlen}{256} \rceil - 1$, we update the state.

$$S_{i+1} = \text{StateUpdate128L}(S_i, AD_{2i}, AD_{2i+1});$$

2.5.4 The encryption of AEGIS-128L

After the initialization, in every step of the encryption, two 16-byte plaintext blocks P_{2i} and P_{2i+1} are used to update the state S_i to obtain the state S_{i+1} , and the plaintext blocks get encrypted.

1. If the size of the message is not a multiple of 256 bits, use 0 bits to pad it to $\lceil \frac{msglen}{256} \rceil \times 256$ bits.
2. Let $u_L = \lceil \frac{adlen}{256} \rceil$, $v_L = \lceil \frac{msglen}{256} \rceil$. For $i = 0$ to $v_L - 1$, we update the state and perform encryption.

$$\begin{aligned}
C_{2i} &= P_{2i} \oplus S_{u_L+i,1} \oplus S_{u_L+i,6} \oplus (S_{u_L+i,2} \& S_{u_L+i,3}) ; \\
C_{2i+1} &= P_{2i+1} \oplus S_{u_L+i,2} \oplus S_{u_L+i,5} \oplus (S_{u_L+i,6} \& S_{u_L+i,7}) ; \\
S_{u_L+i+1} &= \text{StateUpdate128L}(S_{u_L+i}, P_{2i}, P_{2i+1}) ;
\end{aligned}$$

2.5.5 The finalization of AEGIS-128L

After encrypting all the plaintext blocks, we generate the authentication tag using seven more steps. The message being used at this stage is part of the state at the end of the encryption, together with the length of the associated data and the length of the message.

1. Let $tmp = S_{u_L+v_L,2} \oplus (adlen \parallel msglen)$, where $adlen$ and $msglen$ are represented as 64-bit integers.
2. For $i = u_L + v_L$ to $u_L + v_L + 6$, we update the state:

$$S_{i+1} = \text{StateUpdate128L}(S_i, tmp, tmp).$$

3. We generate the authentication tag from the state $S_{u_L+v_L+7}$ as follows:

$$T' = \bigoplus_{i=0}^6 S_{u_L+v_L+7,i} .$$

The authentication tag T is the first t bits of T' .

Chapter 3

Security Goals

The security goals of AEGIS are given in Table 3.1. In AEGIS, each key, IV pair is used to protect only one message. If verification fails, the new tag and the decrypted ciphertext should not be given as output.

Note that the authentication security in Table 3.1 includes the integrity security of plaintext, associated data and nonce.

Table 3.1: Security Goals of AEGIS

	Encryption ^a	Authentication ^b
AEGIS-128L	128-bit	128-bit
AEGIS-128	128-bit	128-bit
AEGIS-256	256-bit	128-bit

^aThe encryption security is under the assumption that the attacker could not forge a message through repeated trials.

^bThe authentication security is under the assumption that the secret key is unknown to the attacker, and a 128-bit tag is used.

Chapter 4

Security Analysis

The following requirements should be satisfied in order to use AEGIS securely.

1. Each key should be generated uniformly at random.
2. Each key and IV pair should not be used to protect more than one message; and each key and IV pair should not be used with two different tag sizes.
3. If verification fails, the decrypted plaintext and the wrong authentication tag should not be given as output.

If the above requirements are satisfied, we have the following security claims:

Claim 1. The success rate of a forgery attack is 2^{-t} , where t is the tag size. If the forgery attack is repeated n times, the success rate of a forgery attack is about $n \times 2^{-t}$.

Claim 2. The state and key cannot be recovered faster than exhaustive key search if the forgery attack is not successful. We recommend the use of a 128-bit tag size for AEGIS in order to resist repeated forgery attacks. (Note that with 128-bit tag, the state of AEGIS-256 can be recovered faster than exhaustive key search if a forgery attack is repeated for about 2^{128} times for the same key and IV pair.)

4.1 The security of the initialization

A difference in IV is the main threat to the security of the initialization of AEGIS. A difference in IV would eventually propagate into the ciphertexts, and thus it is possible to apply a differential attack against AEGIS. In AEGIS-128, there are 50 AES round functions (10 steps) in the initialization. If there is a difference in IV , the difference would pass through more than 10 AES round functions. In AEGIS-256, there are 96 AES round functions (16 steps)

in the initialization. If there is a difference in IV , the difference would pass through more than 16 AES round functions. In AEGIS-128L, there are 80 AES round functions (10 steps) in the initialization. If there is a difference in IV , the difference would pass through more than 20 AES round functions. Furthermore, in order to prevent the difference in the state being eliminated completely in the middle of the initialization, we inject the IV difference repeatedly into the state (5, 8 and 10 times into the state of AEGIS-128, AEGIS-256 and AEGIS-128L, respectively). We expect that a differential attack against the initialization would be more expensive than exhaustive key search.

4.2 The security of the encryption process

We emphasize here that AEGIS encryption is a stream cipher with a large state which is updated continuously. The attacks against a block cipher cannot be applied directly to AEGIS. The state update function involves five AES round functions in AEGIS-128, six AES round functions in AEGIS-256, and eight AES round functions in AEGIS-128L. We should ensure that IV is not reused for the same key; otherwise, the states of AEGIS can be recovered easily with either known-plaintext attacks or chosen plaintext attacks. For example, if we re-use an IV and inject a difference into P_i , the difference would propagate into C_{i+2} , and part of the state can be attacked by analyzing the difference pair $(\Delta P_i, \Delta C_{i+2})$. If an authenticated encryption algorithm is secure for re-used IV s, we expect that such an algorithm can only be as fast as a block cipher, as pointed out in [30]. This can be argued as follows: once an IV is re-used, the attacks that are relevant for a block cipher can be applied to attack the state.

Statistical Attacks. If the IV is used only once for each key, it is impossible to apply a differential attack to the encryption process. It is extremely difficult to apply a linear attack (or correlation attack) to recover the secret state since the state of AEGIS is updated in a nonlinear way. In general, it would be difficult to apply any statistical attack to recover the secret state due to the nonlinear state update function (the statistical correlation between any two states vanishes quickly as the distance between them increases).

LEX [4, 5] is an AES-based stream cipher that generates keystream from part of the state. We would like to mention here that AEGIS is not vulnerable to the attack against LEX [8]. There is a fundamental reason why LEX is vulnerable to a statistical attack while AEGIS is not: the round keys used in LEX are fixed, while the whole state of AEGIS is updated continuously in a nonlinear way.

Randomness of Keystream. Brice Minaud analyzed the bias of the keystream of AEGIS in [19]. It shows that the biases of the keystreams of AEGIS-128 and AEGIS-256 are 2^{-77} and 2^{-89} , respectively. It requires 2^{140} and 2^{188} data to distinguish the keystreams of AEGIS-128 and AEGIS-256, respectively.

4.3 The security of message authentication

There are two main approaches to attack a MAC algorithm. One approach is to recover the secret key or secret state, another approach is to introduce/detect an internal state collision. Besides these two approaches, when we analyze the security of message authentication, we need to consider that the AEGIS encryption may affect the security of message authentication.

4.3.1 Recovering key or state.

From Sect. 4.1, we expect that the secret key cannot be recovered faster than exhaustive search by attacking the initialization. From Sect. 4.2, we expect that the state cannot be recovered faster than exhaustive search by attacking the encryption process if the IV is used only once. Similarly, we expect that the state cannot be recovered faster than exhaustive search by attacking the tag generation process if IV is not reused.

An attacker can still inject a difference into the state in the tag verification process and obtain the decrypted plaintext if the forgery attack is allowed to be repeated for multiple times for the same key and IV pair. In a forgery attack, the decrypted plaintext is known to the attacker with probability 2^{-t} (if the verification is successful). It becomes possible to recover the state if the forgery attack is repeated many times. We recommend the use of 128-bit tag so that recovering the state requires at least 2^{128} forgery attempts.

The security level of the AEGIS-256 state is only 128 bits with a 128-bit tag (if we consider that a forgery attack becomes successful). However, we believe that repeating the forgery attack for around 2^{128} times to recover a state is impractical.

4.3.2 Internal collisions.

A powerful attack against MAC is to introduce and detect internal collisions. A general approach based on the birthday attack was given by Preneel and van Oorschot [26]: an internal collision can be detected after a key is used to generate the authentication tags of about $2^{n/2}$ chosen messages, where n is the state size and tag size in bits. The internal collision can be exploited to forge the tags of new messages. The birthday attack was later applied to other MAC algorithms [32]. AEGIS resists this type of attacks due to its large state size. Another approach to introduce internal collision is through differential cryptanalysis. Suppose that the difference cancellation in the state occurs with probability 2^{-a} ; then we can detect an internal collision after a secret key is used to generate the tags of those 2^a message pairs. The resulting internal collision can be used to forge the tags of new messages.

An attacker can inject a difference into the state in the decryption and tag verification process by modifying the ciphertext. However, AEGIS provides a large security margin against this type of attack since differences are introduced into a large state. The security of AEGIS against forgery attack is stronger than

that of Pelican MAC when the message or the tag gets modified. In Pelican MAC, four AES round functions are used to process each 16-byte message block; while in AEGIS, at least four AES round functions are used. Furthermore, the state size of AEGIS-128 is at least 5 times that of Pelican MAC, and it becomes much more difficult to eliminate the difference in the large state. A simple description of our analysis is given below. We notice that the first difference being injected into ciphertext would pass through five round functions without being affected by another ciphertext difference in AEGIS-128, and there are at least 26 active Sboxes being involved. If we consider only a single differential path, the probability of the difference cancellation in the state is less than $2^{-6 \times 26} = 2^{-156}$. Thus generating a state collision in the verification process requires at least 2^{156} modifications to the ciphertext. Note that the differential attack here is slightly different from that against block cipher since the AEGIS verification process would guarantee that each forgery attack generates only one useful difference pair (the failed forgery attacks would not give outputs). It shows that AEGIS-128 is strong against forgery attack when the ciphertext or tag gets modified. Multiple differential paths would not have a significant effect on the forgery attack here, since each differential path has to cancel its own differences being left in the state. Attacking AEGIS-256 is more difficult since it involves a larger state and more AES round functions. The security of AEGIS-128L against forgery attack is slightly weaker than AEGIS-128 since a difference passes through at least four AES rounds. However, a forgery attack against AEGIS-128L still requires at least 2^{150} modifications to the ciphertext. Note that our analysis above is very conservative since when a difference passes through five AES round functions, the difference would be injected into each 16-byte element in the state.

We now analyze whether the noninvertible AEGIS state update function affects the security of the authentication of AEGIS. In AEGIS, a difference in the state could be eliminated even if there is no difference being introduced to cancel it. However, it would only happen if the difference in every 16-byte element is able to eliminate the difference in the next element after passing through an AES round function. It means that at least 26 active Sboxes are involved in this difference elimination process in AEGIS-128, and generating these particular differences in the state involves more than 26 additional active Sboxes. We consider that this type of weak state difference has a negligible effect on the security of the authentication of AEGIS.

The analysis given above shows that the authentication of AEGIS is very strong.

4.4 Other attacks

There are weak states in AEGIS. In one type of weak states, all the 16-byte elements in a state are equal: consequently all the 16-byte elements in the next state would be equal (if the message block is 0). However, there are only 2^{128} such states, so this type of weak state appears with probabilities 2^{-512} ,

2^{-640} and 2^{-896} for AEGIS-128, AEGIS-256 and AEGIS-128L, respectively. In another type of weak states, the four columns in each 16-byte element are equal and every 16-byte element has such a property: in this case, the same property would appear in the next state (if the message block also has such a property). However, there are only $2^{32 \times 5} = 2^{160}$ such states in AEGIS-128, $2^{32 \times 6} = 2^{192}$ such states in AEGIS-256, $2^{32 \times 8} = 2^{256}$, so we expect that this type of weak state appears with probabilities 2^{-480} , 2^{-608} and 2^{-768} for AEGIS-128, AEGIS-256 and AEGIS-128L, respectively.

Chapter 5

Features

- Efficient. On the latest Intel Haswell microprocessors the speed of AEGIS-128L is more than twice that of AES-GCM.
 - The computational cost of AEGIS is less than half that of AES-GCM.
 - Authentication is achieved almost for free.
 - Encryption/decryption share the same algorithm.
 - Parallel AES round functions at each step, suitable for fast software implementation using AES-NI, and suitable for fast hardware implementation.
- Secure. AEGIS provides 128-bit authentication security, stronger than AES-GCM.

Chapter 6

The Performance of AEGIS

6.1 Software Performance

To process a 16-byte message block, AEGIS-128L, AEGIS-128 and AEGIS-256 use four, five and six AES round functions, respectively. In AEGIS, the critical path for processing a 16-byte message block is about one AES round. The computational cost of AEGIS is about half that of AES for each message block, thus the speed of AEGIS is about twice that of AES when they are implemented using table lookups. For implementations based on bit-slicing techniques (e.g. Käsper and Schwabe [15]), the difference is smaller as AEGIS-128 and AEGIS-256 allow for 5 and 6 parallel AES operations rather than 8; but the speed of AEGIS-128L is still about twice that of AES-128. AEGIS is very efficient when it is implemented using the AES new instructions (AES-NI) which are available on some x86 processors since 2010.

The following software performance data of AEGIS on the Intel Skylake micro-processor (Core i5-6600) is from Supercop-2016-08-06 [9].

Table 6.1: The speed comparison (in cycles per byte) for different plaintext length on Intel Skylake. EA means encryption-authentication; DV means decryption-verification. The length of associated data is zero.

	Very Long	1536B	64B
AEGIS-128L(EA)	0.25	0.34	2.50
AEGIS-128L(DV)	0.25	0.37	3.16
AEGIS-128(EA)	0.43	0.51	2.22
AEGIS-128(DV)	0.41	0.49	2.41
AEGIS-256(EA)	0.47	0.59	3.19
AEGIS-256(DV)	0.46	0.57	3.31

6.2 Hardware Performance

We implemented AEGIS-128L in VHDL. Our implementation targets high throughput, so eight AES round functions are implemented. On the FPGA Xilinx Virtex 7, AEGIS-128L is implemented with 2424 slices (8815 LUTs), and it runs at the speed of 78.3 Gbits/second. More implementations of AEGIS will be available in the third round.

Debjyoti Bhattacharjee and Anupam Chattopadhyay applied various throughput-area improvement techniques to implement AEGIS-128 on 65 nm ASIC [3]. Their results show that AEGIS-128 reaches the high throughput of 121.07 Gbits/second with 172.72 KGE; AEGIS-128 reaches the low-area of 18.72 KGE at the speed of 1.32 Gbits/second.

Chapter 7

Design Rationale

The goal of AEGIS is to achieve high performance and strong security. To achieve high performance, we use the AES round function which is now implemented on the latest Intel and AMD microprocessors as Intel AES New Instructions (AES-NI). AES-NI is very efficient for achieving diffusion and confusion on a modern microprocessor. In the design of AEGIS, we use several parallel AES round functions in each step so as to use most of the pipeline stages in AES instruction. AES instructions are implemented on Intel Westmere (06_25H, 06_2CH, 06_2FH) microprocessors with a three-stage pipeline (6 clock cycles), and are implemented on Intel Sandy Bridge (06_2AH) microprocessors with an eight-stage pipeline (8 clock cycles) [12]. Using several parallel AES round functions in AEGIS significantly improves its performance by utilizing the pipeline of AES-NI.

To achieve strong encryption security, we ensure that the IV difference is randomized at the initialization stage, and the state cannot be recovered from the ciphertext. There are at least 10 steps in the initialization of AEGIS, so we expect that the initialization of AEGIS is strong. To ensure that the state cannot be recovered from the ciphertext faster than brute force attack, we ensure that at least 20, 30 and 24 AES round functions are involved in the state recovery attack against AEGIS-128, AEGIS-256 and AEGIS-128L, respectively.

To achieve strong authentication security, we ensure that any difference being injected results in a particular difference with sufficiently small probability so that it is difficult to launch a forgery attack. Our design is partly motivated by the design of Pelican MAC [7]. In Pelican MAC, a difference would pass through 4 AES round functions before meeting with another difference, so at least 25 active Sboxes are involved. The security proof against differential forgery attack is very simple for Pelican MAC (however, there is a birthday type attack against Pelican MAC due to its 128-bit size [32]). In AEGIS, the first difference in the state would pass through at least 4 AES round functions before being affected by another difference. In addition, when a difference passes through AES round functions, the differences are injected into at least four elements in the state, so it becomes more difficult to eliminate the state difference.

7.1 Reply to the Comments of CAESAR Committee Members

The CAESAR committee members gave very helpful comments on AEGIS. In the following, we provide explanation to some comments.

1. The purpose of using constants in the initialization is to resist the attacks which exploit the symmetric structure of AEGIS. There are two symmetric structures in AEGIS: the symmetric property in the AES round function (for example, we may rotate the four columns, i.e., the i th column is shifted to the $((i + 1) \bmod 4)$ th column); the symmetric property of the overall state (for example, in AEGIS-128L, we can shift the i th 128-bit word in the state to the $((i + 1) \bmod 8)$ th word).
2. We did not use 256-bit key directly in AEGIS-128L due to the concern of guess-and-determine attack on AEGIS.
3. In the finalization, it is a good idea to generate the tag in a similar way as generating keystream (as suggested by the committee member). In AEGIS, we generate the tag in a linear way from the state (different from the keystream generation) is to get a more random tag. We designed AEGIS mainly targeting the microprocessors with AES instructions. In software implementation, the performance of these two approaches are almost the same.

Chapter 8

Changes in the Third Round Submission

There is no tweak to AEGIS.

We made the following changes to the document:

1. We cited an reference [19] in Section 4.2. The randomness of the keystream of AEGIS is analyzed in [19].
2. We updated the software performance on Intel microprocessor Skylake according to the Supercop measurement data in Section 6.1.
3. We added the hardware performance data in Section 6.2.
4. We provided some explanation to some comments given by the CAESAR committee members.

Chapter 9

No Hidden Weakness

We state here that the designer/designers have not hidden any weaknesses in this cipher.

Chapter 10

Intellectual property

We state that AEGIS is not patented and it is freely available for all applications.

If any of this information changes, the submitter will promptly (and within at most one month) announce these changes on the `crypto-competitions` mailing list.

Chapter 11

Consent

The submitter hereby consents to all decisions of the CAESAR selection committee regarding the selection or non-selection of this submission as a second-round candidate, a third-round candidate, a finalist, a member of the final portfolio, or any other designation provided by the committee. The submitter understands that the committee will not comment on the algorithms, except that for each selected algorithm the committee will simply cite the previously published analyses that led to the selection of the algorithm. The submitter understands that the selection of some algorithms is not a negative comment regarding other algorithms, and that an excellent algorithm might fail to be selected simply because not enough analysis was available at the time of the committee decision. The submitter acknowledges that the committee decisions reflect the collective expert judgments of the committee members and are not subject to appeal. The submitter understands that if he disagrees with published analyses then he is expected to promptly and publicly respond to those analyses, not to wait for subsequent committee decisions. The submitter understands that this statement is required as a condition of consideration of this submission by the CAESAR selection committee.

Chapter 12

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments, especially the idea of fully utilizing the 8-stage pipeline of AES-NI on the Sandy Bridge processor to achieve higher performance by increasing the state size. We would also like to thank the anonymous CAESAR committee members for their helpful comments. We would like to thank Tao Huang for implementing AEGIS in VHDL. And we would like to thank Ivica Nikolic for analyzing the security of the finalization part of AEGIS-128L. The first author has been funded by the NAP grant of the Nanyang Technological University. The second author has been funded in part by the Research Council KU Leuven (GOA TENSE) and the FWO Flanders.

Bibliography

- [1] M. Ågren, M. Hell, T. Johansson, W. Meier. Grain-128a: A New Version of Grain-128 with Optional Authentication. *International Journal of Wireless and Mobile Computing 2011*, Vol. 5, No. 1 pp. 48–59.
- [2] M. Bellare and C. Namprempre. Authenticated Encryption: Relations among notions and analysis of the generic composition paradigm. *Advances in Cryptology – Asiacrypt 2000*, LNCS 1976, pp. 531–545.
- [3] D. Bhattacharjee and A. Chattopadhyay. Efficient Hardware Accelerator for AEGIS-128 Authenticated Encryption. *International Conference on Information Security and Cryptology – ICISC 2015*, pp. 385–402.
- [4] A. Biryukov, The Design of a Stream Cipher LEX, *Selected Areas in Cryptography – SAC 2006*, LNCS 4356, pp. 67–75.
- [5] A. Biryukov. The Tweak for LEX-128, LEX-192, LEX-256. ECRYPT stream cipher project report 2006/037. Available at <http://www.ecrypt.eu.org/stream>.
- [6] A. Bogdanov, F. Mendel, F. Regazzoni, V. Rijmen, and E. Tischhauser. ALE: AES-Based Lightweight Authenticated Encryption. *Fast Software Encryption – FSE 2013*.
- [7] J. Daemen, V. Rijmen. The Pelican MAC Function. IACR Cryptology ePrint Archive 2005: 88 (2005).
- [8] O. Dunkelman, N. Keller. A New Attack on the LEX Stream Cipher. *Advances in Cryptology – Asiacrypt 2008*, LNCS 5350, pp. 539–556.
- [9] eBAEAD: ECRYPT Benchmarking of Authenticated Ciphers. Available at <https://bench.cr.yp.to/results-caesar.html>.
- [10] N. Ferguson, D. Whiting, B. Schneier, J. Kelsey, S. Lucks and T. Kohno. Helix, Fast Encryption and Authentication in a Single Cryptographic Primitive. *Fast Software Encryption – FSE 2003*, LNCS 2887, pp. 330–346.
- [11] V. Gligor and P. Donescu. Fast encryption and authentication: XCBC encryption and XECB authentication modes. *Fast Software Encryption – FSE 2001*, LNCS 2355, pp. 92–108.

- [12] Intel. Intel 64 and IA-32 Architectures Optimization Reference Manual. Available at <http://www.intel.com/content/dam/doc/manual/64-ia-32-architectures-optimization-manual.pdf>
- [13] G. Jakimoski and S. Khajuria. ASC-1: An Authenticated Encryption Stream Cipher. *Selected Area in Cryptography – SAC 2011*, LNCS 7118, pp. 356–372.
- [14] C. Jutla, Encryption modes with almost free message integrity. *Advances in Cryptology – EUROCRYPT 2001*, LNCS 2045, pp. 529–544.
- [15] E. Käsper and P. Schwabe. Faster and Timing-Attack Resistant AES-GCM. *Cryptographic Hardware and Embedded Systems – CHES 2009*, LNCS 5747, pp. 1–17.
- [16] J. Katz and M. Yung. Unforgeable encryption and adaptively secure modes of operation. *Fast Software Encryption–FSE 2000*, LNCS 1978, pp. 284–299.
- [17] T. Krovetz, P. Rogaway. The Software Performance of Authenticated-Encryption Modes. *Fast Software Encryption – FSE 2011*, LNCS 6733, pp. 306–327.
- [18] D. McGrew and J. Viega. The security and performance of the Galois/Counter Mode (GCM) of operation. *Progress in Cryptology – INDOCRYPT 2004*, LNCS 3348, pp. 343–355.
- [19] B. Minaud. Linear Biases in the AEGIS Keystream. *Selected Area in Cryptography – SAC 2014*, pp. 290–305.
- [20] National Institute of Standards and Technology. Advanced Encryption Standard. FIPS 197.
- [21] National Institute of Standards and Technology. Recommendation for Block Cipher Modes of Operation. NIST special publication 800-38A, 2001 Edition.
- [22] National Institute of Standards and Technology. The Keyed-Hash Message Authentication Code (HMAC). FIPS PUB 198.
- [23] National Institute of Standards and Technology. Recommendations for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality. NIST special publication 800-38C, May 2004.
- [24] National Institute of Standards and Technology. Recommendations for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST special publication 800-38D, November 2007.
- [25] National Institute of Standards and Technology. Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication. NIST special publication 800-38B.

- [26] B. Preneel, P. C. van Oorschot. On the Security of Iterated Message Authentication Codes. *IEEE Transactions on Information Theory* 45(1), 188–199 (1999).
- [27] P. Rogaway, M. Bellare, and J. Black. OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Trans. on Information and System Security*, 6(3), pp. 365–403, 2003. Earlier version, with T. Krovetz, in *CCS* 2001.
- [28] P. Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. *Advances in Cryptology – ASIACRYPT 2004*, LNCS 3329, pp. 16–31.
- [29] D. Whiting, B. Schneier, S. Lucks and F. Muller. Phelix: Fast Encryption and Authentication in a Single Cryptographic Primitive. eSTREAM, ECRYPT Stream Cipher Project Report 2005/027.
- [30] H. Wu, B. Preneel. Differential-Linear Attacks Against the Stream Cipher Phelix. *Fast Software Encryption – FSE 2007*, LNCS 4593, pp. 87–100.
- [31] H. Wu, B. Preneel. AEGIS: A Fast Authenticated Encryption Algorithm. *Selected Area in Cryptography – SAC 2013*.
- [32] Z. Yuan, W. Wang, K. Jia, G. Xu, X. Wang. New Birthday Attacks on Some MACs Based on Block Ciphers. *Advances in Cryptology – CRYPTO 2009*, LNCS 5677, pp. 209–230.