

Adding more parallelism to the AEGIS authenticated encryption algorithms

Frank Denis

Fastly Inc.
fde@00f.net

Abstract. While the round function of the AEGIS authenticated encryption algorithms is highly parallelizable, their mode of operation is not.

We introduce two new modes to overcome that limitation: AEGIS-128X and AEGIS-256X, that require minimal changes to existing implementations and retain the security properties of AEGIS-128L and AEGIS-256.

1 Introduction

AEGIS [WP14] is a family of three authenticated encryption algorithms, originally designed to leverage the AES-NI instructions set introduced by Intel in 2010 [ADF⁺10]. These instructions perform several compute intensive parts of the AES algorithm, significantly improving the performance of software AES implementations while minimizing the risks of side channel attacks.

However, concurrent AES round instructions are required to fully utilize the AES pipelines. The AEGIS round function was specifically designed with this in mind, and allows up to 8 AES blocks to be updated concurrently. Its design made it the fastest candidate of the CAESAR competition on Intel CPUs with hardware AES acceleration [ARAR16].

Nonetheless, the mode of operation is similar to a duplex: after its initialization, the state is recursively updated. That effectively limits the parallelism of the construction to the parallelism of the round function.

In [BLT15], Bogdanov, Lauridsen, and Tischhauser made a similar observation regarding multiple candidates of the CAESAR competition. They proposed a novel "comb scheduler" to process multiple messages simultaneously.

The modes presented here also encrypt multiple messages simultaneously using the same cipher, but assume that they are fragments of the same message, and share the same key, initialization vector and length.

Given a parallelism degree ν , an input message is split into ν evenly distributed parts, that can be encrypted concurrently. The resulting ciphertexts are then combined, as well as their authentication tags.

The underlying encryption algorithms remain the existing AEGIS algorithms, with a minor addition to their initialization functions.

2 Operations, Variables and Functions

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

2.1 Operations

$ x $: Size of x in bits
$x \oplus y$: Bit-wise exclusive <i>OR</i>
$\mathbb{F}_{128L}(S, m_0, m_1)$: AEGIS-128L state update function
$\mathbb{F}_{256}(S, m)$: AEGIS-256 state update function
$x \ y$: Concatenation of x and y
$Pad(x, \ell)$: Add trailing 0 bits to pad x to ℓ bits
$Enc_{128L}(CTX, K, IV, \bar{A}_i, \bar{P}_i, A , P)$: The AEGIS-128L encryption function with context separation
$Enc_{256}(CTX, K, IV, \bar{A}_i, \bar{P}_i, A , P)$: The AEGIS-256 encryption function with context separation
$Enc_{128X}[\nu](K, IV, A, P)$: The AEGIS-128X $[\nu]$ parallel encryption function
$Enc_{256X}[\nu](K, IV, A, P)$: The AEGIS-256X $[\nu]$ parallel encryption function
$Trunc(x, \ell)$: Truncate x to the first ℓ bits

2.2 Variables and constants

A	: Associated data or $\{\}$ if unspecified
A_i	: 128-bit associated data block
\hat{A}	: A , padded to $r \cdot \nu$ bits
\hat{A}_i	: 128-bit associated data block
\bar{A}_i	: \hat{A} fragment ($\lfloor \frac{ \hat{A} }{\nu} \rfloor$ bits)
$\bar{A}_{i,j}$: 128-bit \hat{A} fragment block
C	: Ciphertext
\hat{C}	: C , padded to $r \cdot \nu$ bits
\hat{C}_i	: A 128-bit ciphertext block
\bar{C}_i	: C fragment ($\lfloor \frac{ \hat{C} }{\nu} \rfloor$ bits)
$\bar{C}_{i,j}$: 128-bit C fragment block
$const_0$: First half of the AEGIS constant (128 bits)
$const_1$: Second half of the AEGIS constant (128 bits)
CTX	: Context separator
K_{128}	: 128-bit encryption key (AEGIS-128, -128L)
K_{256}	: 256-bit encryption key (AEGIS-256)
$K_{256,0}$: First half of a 256-bit key
$K_{256,1}$: Second half of a 256-bit key
P	: Plaintext
\hat{P}	: P , padded to $r \cdot \nu$ bits
\hat{P}_i	: 128-bit plaintext block
\bar{P}	: P fragment ($\lfloor \frac{ \hat{P} }{\nu} \rfloor$ bits)
\bar{P}_i	: 128-bit P fragment block
IV_{128}	: 128-bit initialization vector (AEGIS-128, -128L)
IV_{256}	: 256-bit initialization vector (AEGIS-256)
$IV_{256,0}$: First half of a 256-bit initialization vector
$IV_{256,1}$: Second half of a 256-bit initialization vector
ν	: Parallelism degree (≥ 1)
r	: Absorption rate (128 or 256 bits)
S	: AEGIS state
S_i	: A 128-bit AEGIS state block
T	: Authentication tag for C
\bar{T}_i	: Authentication tag for \bar{C}_i

3 Context separation

From an application perspective, new AEGIS variants should ideally share the same interface as existing variants. Namely, they should accept a single message, optional associated data, a 128 or 256 bit key, and a 128 or 256 bit initialization vector.

However, AEGIS is meant to be used in a nonce-respecting scenario [VV18]. Clearly, reusing the same key and IV to encrypt different parts of a message would violate that contract.

In order to avoid universal forgery and decryption attacks, we augment the AEGIS initialization functions with a context, meant to provide domain separation. That is, for two different values $CTX_0 \neq CTX_1$, $Enc.[\cdot](CTX_0, \cdot)$ and $Enc.[\cdot](CTX_1, \cdot)$ act as two different functions of $\{K, IV, A, M\}$.

3.1 Augmenting AEGIS-128L for context separation

AEGIS-128L defines the initial state S as a vector of eight AES blocks $\{S_0, S_1, \dots, S_7\}$ set to:

Block	Initial value
S_0	$K_{128} \oplus IV_{128}$
S_1	$const_1$
S_2	$const_0$
S_3	$const_1$
S_4	$K_{128} \oplus IV_{128}$
S_5	$K_{128} \oplus const_0$
S_6	$K_{128} \oplus const_1$
S_7	$K_{128} \oplus const_0$

From this state, the original AEGIS-128L initialization function performs 10 updates as described in algorithm 1.

Algorithm 1 Contextless AEGIS-128L initialization

```

function INITIALIZE(K, IV)
   $S \leftarrow \{K_{128} \oplus IV_{128}, const_1, const_0, const_1, K_{128} \oplus IV_{128}, K_{128} \oplus const_0, K_{128} \oplus const_1, K_{128} \oplus const_0\}$ 
   $i \leftarrow 0$ 
  while  $i < 10$  do
     $S \leftarrow \mathbb{F}_{128L}(S, IV_{128}, K_{128})$ 
     $i \leftarrow i + 1$ 
  end while
end function

```

We augment this function to accept a context parameter CTX . Before each update, the context is added to the blocks at indices 3 and 7 of the state, as described in algorithm 2

Note that when $CTX = 0$, the resulting state is exactly the same as AEGIS-128L, as originally specified, without a context.

Algorithm 2 AEGIS-128L initialization with context

```

function INITIALIZE(CTX, K, IV)
   $S \leftarrow \{K_{128} \oplus IV_{128}, const_1, const_0, const_1, K_{128} \oplus IV_{128}, K_{128} \oplus const_0, K_{128} \oplus const_1, K_{128} \oplus const_0\}$ 
   $i \leftarrow 0$ 
  while  $i < 10$  do
     $S_3 \leftarrow S_3 \oplus Pad(CTX, 128)$ 
     $S_7 \leftarrow S_7 \oplus Pad(CTX, 128)$ 
     $S \leftarrow \mathbb{F}_{128L}(S, IV_{128}, K_{128})$ 
     $i \leftarrow i + 1$ 
  end while
end function

```

3.2 Augmenting AEGIS-256 for context separation

AEGIS-256 accepts a 256-bit key K_{256} made of two AES blocks $\{K_{256,0}, K_{256,1}\}$, as well as 256-bit initialization vector IV_{256} made of two AES blocks $\{IV_{256,0}, IV_{256,1}\}$.

The initial state S is a vector of six AES blocks $\{S_0, S_1, \dots, S_5\}$ set to:

Block	Initial value
S_0	$K_{256,0} \oplus IV_{256,0}$
S_1	$K_{256,1} \oplus IV_{256,1}$
S_2	$const_0$
S_3	$const_1$
S_4	$K_{256,0} \oplus const_1$
S_5	$K_{256,1} \oplus const_0$

From this state, the original AEGIS-256 initialization function performs 16 updates as described in algorithm 3.

Algorithm 3 Contextless AEGIS-256 initialization

```

function INITIALIZE(K, IV)
   $S \leftarrow \{K_{256,0} \oplus IV_{256,0}, K_{256,1} \oplus IV_{256,1}, const_0, const_1, K_{256,0} \oplus const_1, K_{256,1} \oplus const_0\}$ 
   $i \leftarrow 0$ 
  while  $i < 4$  do
     $S \leftarrow \mathbb{F}_{256}(S, K_{256,0})$ 
     $S \leftarrow \mathbb{F}_{256}(S, K_{256,1})$ 
     $S \leftarrow \mathbb{F}_{256}(S, IV_{256,0})$ 
     $S \leftarrow \mathbb{F}_{256}(S, IV_{256,1})$ 
     $i \leftarrow i + 1$ 
  end while
end function

```

We augment this function to accept a context parameter CTX . Before each update, the context is added to the blocks at indices 3 and 5 of the state, as described in algorithm 4

Note that when $CTX = 0$, the resulting state is exactly the same as AEGIS-256, as originally specified, without a context.

Algorithm 4 AEGIS-256 initialization with context

```

function INITIALIZE(CTX, K, IV)
   $S \leftarrow \{K_{256,0} \oplus IV_{256,0}, K_{256,1} \oplus IV_{256,1}, const_0, const_1, K_{256,0} \oplus const_1, K_{256,1} \oplus const_0\}$ 
   $i \leftarrow 0$ 
  while  $i < 4$  do
     $S_3 \leftarrow S_3 \oplus Pad(CTX, 128)$ 
     $S_5 \leftarrow S_5 \oplus Pad(CTX, 128)$ 
     $S \leftarrow \mathbb{F}_{256}(S, K_{256,0})$ 
     $S_3 \leftarrow S_3 \oplus Pad(CTX, 128)$ 
     $S_5 \leftarrow S_5 \oplus Pad(CTX, 128)$ 
     $S \leftarrow \mathbb{F}_{256}(S, K_{256,1})$ 
     $S_3 \leftarrow S_3 \oplus Pad(CTX, 128)$ 
     $S_5 \leftarrow S_5 \oplus Pad(CTX, 128)$ 
     $S \leftarrow \mathbb{F}_{256}(S, IV_{256,0})$ 
     $S_3 \leftarrow S_3 \oplus Pad(CTX, 128)$ 
     $S_5 \leftarrow S_5 \oplus Pad(CTX, 128)$ 
     $S \leftarrow \mathbb{F}_{256}(S, IV_{256,1})$ 
     $i \leftarrow i + 1$ 
  end while
end function

```

3.3 Explicit lengths for finalization

The finalization functions of AEGIS-128L and AEGIS-256 absorb the lengths of the associated data and message before performing seven state updates (algorithms 5 and 6).

The Enc_{128L} and Enc_{256} functions defined here require explicit lengths for $|A|$ and $|P|$, to be used in their respective finalization step.

This is required because in the parallel variants of AEGIS, individual encryption functions process distinct fragments of A and P simultaneously. However, in the finalization phase of AEGIS-128X $[\nu]$ and AEGIS-256X $[\nu]$, these functions use the same values for $|A|$ and $|P|$ instead of the individual fragment lengths.

Algorithm 5 The AEGIS-128L finalization function

```

function FINALIZE( $|A|$ ,  $|P|$ ,  $|T|$ )
   $T \leftarrow Pad(|A|, 128) || Pad(|P|, 128)$ 
   $i \leftarrow 0$ 
  while  $i < 7$  do  $\mathbb{F}_{128L}(S, T, T)$ 
     $i \leftarrow i + 1$ 
  end while
  if  $|T| = 256$  then
    return  $(S_0 \oplus S_1 \oplus S_2 \oplus S_3) || (\oplus S_4 \oplus S_5 \oplus S_6 \oplus S_7)$ 
  else
    return  $S_0 \oplus S_1 \oplus S_2 \oplus S_3 \oplus S_4 \oplus S_5 \oplus S_6$ 
  end if
end function

```

Algorithm 6 The AEGIS-256 finalization function

```

function FINALIZE( $|A|$ ,  $|P|$ ,  $|T|$ )
   $T \leftarrow Pad(|A|, 128) || Pad(|P|, 128)$ 
   $i \leftarrow 0$ 
  while  $i < 7$  do  $\mathbb{F}_{256}(S, T)$ 
     $i \leftarrow i + 1$ 
  end while
  if  $|T| = 256$  then
    return  $(S_0 \oplus S_1 \oplus S_2) || (\oplus S_3 \oplus S_4 \oplus S_5)$ 
  else
    return  $S_0 \oplus S_1 \oplus S_2 \oplus S_3 \oplus S_4 \oplus S_5$ 
  end if
end function

```

4 The AEGIS-128X $[\nu]$ and AEGIS-256X $[\nu]$ modes

AEGIS absorbs the associated data and message with a rate r with $r = 256$ for AEGIS-128L and $r = 128$ for AEGIS-128 and AEGIS-256.

We define two new modes: AEGIS-128X $[\nu]$ and AEGIS-256X $[\nu]$, that absorb $r \cdot \nu$ bits per state update, spread over ν concurrent instances of the AEGIS-128L or AEGIS-256 encryption functions.

The associated data A and plaintext P are split into interleaved blocks with a stride of $r \cdot \nu$ bits as they arrive. The last blocks are padded if necessary.

We first pad A and P by adding trailing zero bits until they match the stride length:

$$\hat{A} = Pad(A, r \cdot \nu)$$

$$\hat{P} = Pad(P, r \cdot \nu)$$

\hat{A} is split into 128-bit blocks $\{\hat{A}_0, \hat{A}_1, \dots, \hat{A}_{\frac{|\hat{A}|}{128}-1}\}$.

These blocks are interleaved to produce ν independent $\frac{|\hat{A}|}{\nu}$ bit messages $\{\bar{A}_0, \bar{A}_1, \dots, \bar{A}_{\nu-1}\}$.

$$\begin{aligned}\bar{A}_0 &= \hat{A}_0 \| \hat{A}_\nu \| \hat{A}_{2\nu} \| \hat{A}_{3\nu} \| \dots \\ \bar{A}_1 &= \hat{A}_1 \| \hat{A}_{\nu+1} \| \hat{A}_{2\nu+1} \| \hat{A}_{3\nu+1} \| \dots \\ \bar{A}_2 &= \hat{A}_2 \| \hat{A}_{\nu+2} \| \hat{A}_{2\nu+2} \| \hat{A}_{3\nu+2} \| \dots \\ &\vdots \\ \bar{A}_{\nu-1} &= \hat{A}_{\nu-1} \| \hat{A}_{\nu+(\nu-1)} \| \hat{A}_{2\nu+(\nu-1)} \| \hat{A}_{3\nu+(\nu-1)} \| \dots\end{aligned}$$

Similarly, \hat{P} is split into ν independent $\frac{|\hat{P}|}{\nu}$ bit messages $\{\bar{P}_0, \bar{P}_1, \dots, \bar{P}_{\nu-1}\}$.

4.1 AEGIS-128X

AEGIS-128X $[\nu]$ first encrypts and authenticates the plaintext and associated data fragments independently, producing ν ciphertexts $\{\bar{C}_0, \bar{C}_1, \dots, \bar{C}_{\nu-1}\}$ and authentication tags $\{\bar{T}_0, \bar{T}_1, \dots, \bar{T}_{\nu-1}\}$.

$$\begin{aligned}\{\bar{C}_0, \bar{T}_0\} &= Enc_{128L}(CTX \leftarrow 0, K, IV, \bar{A}_0, \bar{M}_0, |A|, |P|) \\ \{\bar{C}_1, \bar{T}_1\} &= Enc_{128L}(CTX \leftarrow 1, K, IV, \bar{A}_1, \bar{M}_1, |A|, |P|) \\ \{\bar{C}_2, \bar{T}_2\} &= Enc_{128L}(CTX \leftarrow 2, K, IV, \bar{A}_2, \bar{M}_2, |A|, |P|) \\ &\vdots \\ \{\bar{C}_{\nu-1}, \bar{T}_{\nu-1}\} &= Enc_{128L}(CTX \leftarrow \nu - 1, K, IV, \bar{A}_{\nu-1}, \bar{M}_{\nu-1}, |A|, |P|)\end{aligned}$$

$\{\bar{C}_0, \bar{C}_1, \dots, \bar{C}_{\nu-1}\}$ are deinterleaved to produce the final ciphertext:

$$\begin{aligned}\hat{C} &= \\ &\bar{C}_{0,0} \| \bar{C}_{1,0} \| \bar{C}_{2,0} \| \dots \| \bar{C}_{(\nu-1),0} \| \\ &\bar{C}_{0,1} \| \bar{C}_{1,1} \| \bar{C}_{2,1} \| \dots \| \bar{C}_{(\nu-1),1} \| \\ &\bar{C}_{0,2} \| \bar{C}_{1,2} \| \bar{C}_{2,2} \| \dots \| \bar{C}_{(\nu-1),2} \| \dots\end{aligned}$$

$$C = Trunc(\hat{C}, |P|)$$

Finally, the AEGIS-128X $[\nu]$ authentication tag is the bit-wise exclusive *OR* of the AEGIS-128L authentication tags:

$$T = \bar{T}_0 \oplus \bar{T}_1 \oplus \dots \oplus \bar{T}_{\nu-1}$$

Note that AEGIS-128L and AEGIS-128X[1] are equivalent.

4.2 AEGIS-256X

AEGIS-256X $[\nu]$ uses the exact same interleaving technique as AEGIS-128X $[\nu]$ in order to process $r \cdot \nu$ bits per state update.

The only difference being that fragments are encrypted and authenticated using the AEGIS-256 encryption function instead of the AEGIS-128L one.

$$\begin{aligned}\{\bar{C}_0, \bar{T}_0\} &= Enc_{256}(CTX \leftarrow 0, K, IV, \bar{A}_0, \bar{M}_0, |A|, |P|) \\ \{\bar{C}_1, \bar{T}_1\} &= Enc_{256}(CTX \leftarrow 1, K, IV, \bar{A}_1, \bar{M}_1, |A|, |P|) \\ \{\bar{C}_2, \bar{T}_2\} &= Enc_{256}(CTX \leftarrow 2, K, IV, \bar{A}_2, \bar{M}_2, |A|, |P|) \\ &\vdots \\ \{\bar{C}_{\nu-1}, \bar{T}_{\nu-1}\} &= Enc_{256}(CTX \leftarrow \nu - 1, K, IV, \bar{A}_{\nu-1}, \bar{M}_{\nu-1}, |A|, |P|)\end{aligned}$$

$\{\bar{C}_0, \bar{C}_1, \dots, \bar{C}_{\nu-1}\}$ are deinterleaved to produce the AEGIS-256X $[\nu]$ ciphertext:

$$\hat{C} = \begin{array}{l} \overline{C}_{0,0} \| \overline{C}_{1,0} \| \overline{C}_{2,0} \| \dots \| \overline{C}_{(\nu-1),0} \| \\ \overline{C}_{0,1} \| \overline{C}_{1,1} \| \overline{C}_{2,1} \| \dots \| \overline{C}_{(\nu-1),1} \| \\ \overline{C}_{0,2} \| \overline{C}_{1,2} \| \overline{C}_{2,2} \| \dots \| \overline{C}_{(\nu-1),2} \| \dots \end{array}$$

$$C = \text{Trunc}(\hat{C}, |P|)$$

Finally, the AEGIS-256X[ν] authentication tag is the bit-wise exclusive *OR* of the AEGIS-256 authentication tags:

$$T = \overline{T}_0 \oplus \overline{T}_1 \oplus \dots \oplus \overline{T}_{\nu-1}$$

Note that AEGIS-256 and AEGIS-256X[1] are equivalent.

5 Rationale

The AEGIS security claims have the following requirements:

- Each key should be generated uniformly at random.
- 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.
- If verification fails, the decrypted plaintext and the wrong authentication tag should not be given as output.

AEGIS-128X[ν] and AEGIS-256X[ν] have the same requirements.

In the finalization step, the same total sizes are absorbed by all the instances of the encryption function. This ensures that the parallel variants of AEGIS can only be used within the same limits as the non-parallel variants. It also improves randomness, as any change to the input lengths will affect every authentication tag T_i , regardless of the rate.

We'd like to emphasize that AEGIS-128X[ν] and AEGIS-256 are not new algorithms. They are modes, built on top of AEGIS-128L and AEGIS-256. designed to preserve the same security guarantees and requirements. Keys must be generated uniformly at random, key and *IV* pairs must not be reused, and for both variants the success rate of a forgery attack remains $2^{-|T|}$.

We did not include a parallel variant of AEGIS-128, as opposed to AEGIS-128L. AEGIS-128 trades performance for a smaller state size. But given that performance is the main justification for AEGIS-128X, AEGIS-128L feels like a more natural choice.

5.1 Implications of the AEGIS-128L context addition

$Enc_{128X}[\nu](K, IV, A, P)$ can be seen as ν evaluations of AEGIS-128L, on ν independent messages of the same length.

One way to satisfy the AEGIS-128L contract while reusing the key is to use distinct initialization vectors for each of the message fragments.

The parallelism degree ν , and thus the bounds of *CTX*, are limited by the hardware, and guaranteed to be small.

We could limit the AEGIS-128X[ν] *IV* size to $128 - \log_2(\nu)$ bits (instead of 128 for AEGIS-128L), encoding the context in the remaining bits to create the *IV* used by the underlying AEGIS-128L function. That would be effectively AEGIS-128L, evaluated with independent messages, and distinct (K, IV) pairs. However, from an application perspective, a $128 - \log_2(\nu)$ bit initialization vector would be unusual, and at odds with AEGIS-128L.

Ideally, we'd like AEGIS-128L to internally support $128 + \log_2(\nu)$ -bit initialization vectors: AEGIS-128X $[\nu]$ applications would use 128 bit initialization vectors, but the context could still be encoded to separate the parallel AEGIS-128L instances. To put it differently, we need to introduce a context with the same differential properties as the initialization vector.

In the proposed tweak to the initialization function, the context is added to the constants in blocks 3 and 7 AEGIS-128L of the initial state. The purpose of the constants (simply derived from the Fibonacci sequence) is to resist attacks exploiting the symmetry of the AES round function and of the overall AEGIS state.

Given its limited range, adding ν cannot turn them into weak constants, and doesn't alter any of the AEGIS-128L properties. Also note that ν is expected to be a hyperparameter, that an adversary cannot have control of.

Differential attacks could be a concern with the same (K, IV) pair used in different contexts. But in AEGIS-128L, there are 80 AES round functions (10 update steps) in the initialization function. In [STSI23], Shiraya et al. showed that the initialization phase of AEGIS-128L is secure against differential attacks after 3 update steps.

Furthermore, in order to prevent the difference in the state being eliminated completely in the middle of the initialization, the context difference is repeatedly injected into the state. This is consistent with how 128-bit initialization vectors are absorbed in AEGIS-128L.

The addition of a short context is thus unlikely to invalidate any of the current AEGIS-128L security claims.

The above security claims require a key and IV pair not to be used with different tag sizes. The AEGIS-128X $[\nu]$ construction guarantees that internal AEGIS-128L evaluations will always share the same tag size.

Note that the addition of a context to the AEGIS-128L initialization function could also be used to create a different initial state for different tag sizes, effectively increasing misuse resistance.

We do not encourage implementations to expose the context parameter in their public APIs. However, it can be used for constructions based on the standard AEGIS AEADs that extend beyond the ones proposed in this paper.

5.2 Implications of the AEGIS-256 context addition

The same observations apply to AEGIS-256X $[\nu]$ and its $Enc_{256X}[\nu]$ encryption function.

AEGIS-256 has a large initialization vector (256 bits). Exposing a shorter IV to applications while leveraging the remaining space for context separation would be reasonable. That would still allow applications to use random nonces with a negligible collision probability.

However, the large initialization vector has practical benefits, such as the ability to include a *protocol-specific* context with no overhead. In order to retain the same convenience as AEGIS-256 and for consistency with AEGIS-128X, our own context is added in a similar fashion.

In the AEGIS-256 initialization function, there are 96 round functions (16 update steps). According to the MILP-based evaluation from [STSI23], AEGIS-256 is secure against differential attacks after 6 update steps.

6 Implementation notes

Implementing AEGIS-128X $[\nu]$ and AEGIS-256X $[\nu]$ only requires trivial modifications to existing AEGIS-128L and AEGIS-256 implementations.

They apply the exact same operations as AEGIS-128L and AEGIS-256, to vectors of ν AES blocks instead of single blocks.

For example, with 256-bit registers, two AEGIS-128L states S and S' can be stored as:

$$\{S_0, S'_0\}, \{S_1, S'_1\}, \{S_2, S'_2\}, \dots \{S_7, S'_7\}$$

This perfectly matches the AEGIS-128X[2] interleaved representation. In addition to the forward AES permutation, updating an AEGIS state only requires the bit-wise OR and AND operations. Even with wide registers, such operations are very efficient on any CPU with vector instructions.

On CPUs that don't implement vectorized versions of the AES permutation, AEGIS-128X[ν] and AEGIS-256X[ν] can be implemented in different ways:

- by emulating AES vector instructions. This is the easiest option, keeping the code close to hardware-accelerated versions.
- by evaluating $\{\overline{A}_0, \overline{A}_1, \overline{A}_2, \dots, \overline{A}_{\nu-1}\}$ and $\{\overline{C}_0, \overline{C}_1, \overline{C}_2, \dots, \overline{C}_{\nu-1}\}$ sequentially, with a periodic synchronization, for example after every memory page. This reduces cache-locality but also register pressure.

7 Performance evaluation

We implemented AEGIS-128X[ν] and AEGIS-256X[ν] using the Zig programming language. The code [Den23] is nearly identical to the reference implementations from the AEGIS specification [DSL23], with the AesBlock type extended to 2 AES blocks.

On Intel, AMD and ARM CPUs with vector registers but without AES vector instructions, AEGIS-128L and AEGIS-256 have better or comparable performance.

And unsurprisingly, on CPUs without vector registers, AEGIS-128L and AEGIS-256 are consistently faster than their parallel counterparts.

However, on CPUs with the VAES instruction set such as an Intel Raptor Lake processor (table 1) or an AMD EPYC CPU (table 2, plot 7), AEGIS-128X[2] is almost twice as fast as AEGIS-128L with medium to large inputs.

With short inputs, the overhead of the parallel versions is small (plot 7). AEGIS-128X[2] is faster than AEGIS-128L as soon as the input size reaches 256 bytes, and AEGIS-256X[2] beats AEGIS-256 for inputs of size 128 or more.

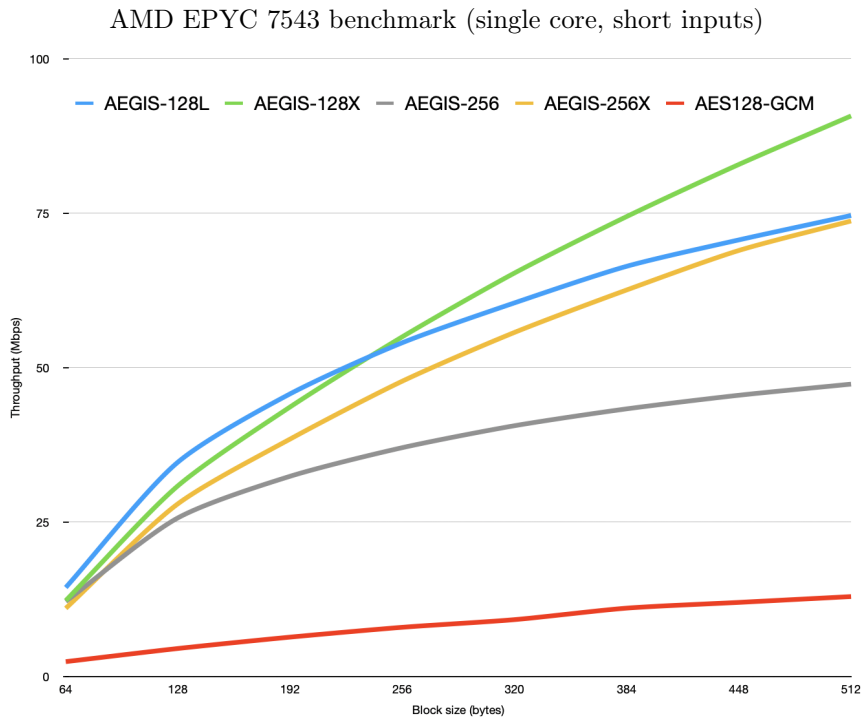
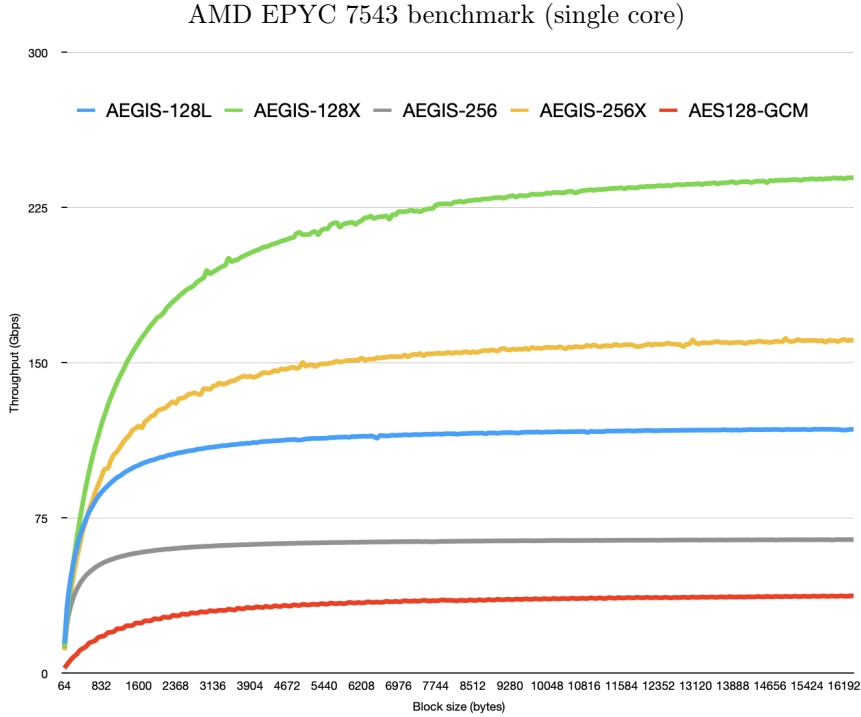
Algorithm	Throughput
AEGIS-128X[2]	318 Gbps
AEGIS-128L	190 Gbps
AES128-GCM	81 Gbps

Table 1. Intel Core i9-13900K benchmark numbers (16 KB messages, single core)

Algorithm	Throughput
AEGIS-128X[2]	239 Gbps
AEGIS-128L	118 Gbps
AES128-GCM	37 Gbps

Table 2. AMD EPYC 7543 benchmark numbers (16 KB messages, single core)

The benchmarked AES128-GCM implementation is the one from OpenSSL 3.1.0, while the AEGIS implementations are the reference Zig code of the AEGIS specification, as well as our modified version to support the 128X and 256X variants.



The AMD Zen4 family of CPUs features the AVX-512 instruction set, as well as 4-way AES instructions.

As OpenSSL 3.1.1 includes an AES-GCM implementation specifically optimized for these CPUs, we also benchmarked AEGIS on an AMD Ryzen 7 7700 CPU. The results are summarized in tables 3 and 4.

We observe that even generic, non-parallel variants of AEGIS outperform the fastest known implementations of AES-GCM on that hardware.

And the 2-way variants of AEGIS immediately double the throughput.

Algorithm	Throughput
AEGIS-128X[4]	348 Gbps
AEGIS-128X[2]	310 Gbps
AEGIS-128L	152 Gbps
AES128-GCM	94 Gbps

Table 3. AMD Ryzen 7 7700 benchmark numbers - 128 bit security (16 KB messages, single core)

Algorithm	Throughput
AEGIS-256X[4]	277 Gbps
AEGIS-256X[2]	184 Gbps
AEGIS-256	93 Gbps
AES256-GCM	81 Gbps

Table 4. AMD Ryzen 7 7700 benchmark numbers - 256 bit security (16 KB messages, single core)

The performance gains of the 4-way variants over the 2-way variants appear to be less significant. However, especially with AEGIS-128X4, we may be limited by other factors, such as the memory bandwidth.

Independent measurements reported a more significant gap between the 2X and 4X variants on Intel Xeon CPUs.

8 Conclusion

We propose two additions to the AEGIS family of authenticated encryption algorithms: AEGIS-128X[ν] and AEGIS-256[ν].

They take advantage of new vector instructions to compute multiple instances of AEGIS-128L and AEGIS-256 in parallel. In order to do so, inputs are simply split into ν pieces that can be processed independently, in a way that directly matches how CPU instructions load, store and use vector registers.

As a result, they are very similar to the non-parallel variant, share the same security guarantees, yet significantly improve the performance of these ciphers on CPUs supporting these vector instructions.

References

- ADF⁺10. K. D. Akdemir, M. G. Dixon, W. K. Feghali, P. G. Fay, V. Gopal, J. Guilford, E. Ozturk, G. Wolrich, and R. Zohar. Breakthrough AES Performance with Intel [®] AES New Instructions. 2010.
- ARAR16. Ankele, Ralph, Ankele, and Robin. Software Benchmarking of the 2nd round CAESAR Candidates. Cryptology ePrint Archive, Report 2016/740, 2016. <https://eprint.iacr.org/2016/740>.
- BLT15. A. Bogdanov, M. M. Lauridsen, and E. Tischhauser. Comb to Pipeline: Fast Software Encryption Revisited. In *FSE 2015, LNCS 9054*, pages 150–171. Springer, Heidelberg, March 2015.
- Den23. F. Denis. AEGIS-128X and AEGIS-256X implementations. GitHub, 2023. <https://github.com/jedisct1/aegis-128x>.
- DSL23. F. Denis, F. E. R. Scotoni, and S. Lucas. The AEGIS family of authenticated encryption algorithms. Internet-Draft draft-irtf-cfrg-aegis-aead-02, Internet Engineering Task Force, 2023. Work in Progress.
- STSI23. T. Shiraya, N. Takeuchi, K. Sakamoto, and T. Isobe. MILP-based security evaluation for AEGIS/Tiaoxin-346/Rocca. *IET Information Security*, 2023.
- VV18. S. Vaudenay and D. Vizár. Can Caesar Beat Galois? - Robustness of CAESAR Candidates Against Nonce Reusing and High Data Complexity Attacks. In *ACNS 18, LNCS 10892*, pages 476–494. Springer, Heidelberg, July 2018.
- WP14. H. Wu and B. Preneel. AEGIS: A Fast Authenticated Encryption Algorithm. In *SAC 2013, LNCS 8282*, pages 185–201. Springer, Heidelberg, August 2014.

A Test vectors

Test vectors (in hexadecimal format) of AEGIS-128X[ν] and AEGIS-256X[ν] are given below.

A.1 AEGIS-128X[2]

Key: 000102030405060708090 a0b0c0d0e0f
 IV: 101112131415161718191 a1b1c1d1e1f
 AD: (empty)
 Plaintext: (empty)
 Ciphertext: (empty)
 128-bit tag: df5b76dbe5f2acd08523a284f3b90219

Key: 000102030405060708090 a0b0c0d0e0f
 IV: 101112131415161718191 a1b1c1d1e1f
 AD: 0102030401020304
 Plaintext: 050607080506070805060708
 Ciphertext: 985bac76fe1cefaa0b24d4d2
 128-bit tag: 7f683fe3c1ddeef54759ea8f890a64ef

A.2 AEGIS-256X[2]

Key: 000102030405060708090 a0b0c0d0e0f
 101112131415161718191 a1b1c1d1e1f
 IV: 101112131415161718191 a1b1c1d1e1f
 202122232425262728292 a2b2c2d2e2f
 AD: (empty)
 Plaintext: (empty)
 Ciphertext: (empty)
 128-bit tag: ddbc7695b01623f82eee962f160bf267

Key: 000102030405060708090 a0b0c0d0e0f
 101112131415161718191 a1b1c1d1e1f
 IV: 101112131415161718191 a1b1c1d1e1f
 202122232425262728292 a2b2c2d2e2f
 AD: 0102030401020304
 Plaintext: 050607080506070805060708
 Ciphertext: a0b3f5b6b93db779c9d1b9de
 128-bit tag: fd2e93a6eb0b74dc30eb984fbec1d657

A.3 AEGIS-128X[4]

Key: 000102030405060708090 a0b0c0d0e0f
 IV: 101112131415161718191 a1b1c1d1e1f
 AD: (empty)
 Plaintext: (empty)
 Ciphertext: (empty)
 256-bit tag: e054b894f597ad7c0468793f67e7dc2c
 7c864c07177ae13b4f03ad4e00f9bab1

Key: 000102030405060708090 a0b0c0d0e0f
 IV: 101112131415161718191 a1b1c1d1e1f
 AD: 0102030401020304
 Plaintext: 050607080506070805060708
 Ciphertext: 985bac76fe1cefaa0b24d4d2
 256-bit tag: 8bd8f84f25ae7f5dfe24a761ddf559ab
 70798d72bb94ff11c0a2cddb20445d97

A.4 AEGIS-256X[4]

Key: 000102030405060708090 a0b0c0d0e0f
 101112131415161718191 a1b1c1d1e1f
 IV: 101112131415161718191 a1b1c1d1e1f
 202122232425262728292 a2b2c2d2e2f
 AD: (empty)
 Plaintext: (empty)
 Ciphertext: (empty)
 256-bit tag: 1263af1648b25cfa9fd94204bcd4d8ca
 391629d6f6924775bdd124b105e0850c

Key: 000102030405060708090 a0b0c0d0e0f
 101112131415161718191 a1b1c1d1e1f
 IV: 101112131415161718191 a1b1c1d1e1f
 202122232425262728292 a2b2c2d2e2f
 AD: 0102030401020304
 Plaintext: 050607080506070805060708
 Ciphertext: a0b3f5b6b93db779c9d1b9de
 256-bit tag: 4ccc371413e5f4b7227e8c6d98a1b5ba
 67c9de0fa46352fc6378cf3c3bd410f5