Areion: Highly-Efficient Permutations and Its Applications (Extended Version)*

Takanori Isobe[1,2], Ryoma Ito[2], Fukang Liu[1], Kazuhiko Minematsu[3], Motoki Nakahashi[1], Kosei Sakamoto[1] and Rentaro Shiba[4]

[1] University of Hyogo, Kobe, Japan.
takanori.isobe@ai.u-hyogo.ac.jp,liufukangs@gmail.com,
motoki.n1998@gmail.com,k.sakamoto0728@gmail.com

[2] National Institute of Information and Communications Technology, Koganei, Japan.
itorym@nict.go.jp

[3] NEC, Kawasaki, Japan.
k-minematsu@nec.com

[4] Mitsubishi Electric Corporation, Kamakura, Japan.
shiba.rentaro@dc.mitsubishielectric.co.jp

Abstract. In real-world applications, the overwhelming majority of cases require (authenticated) encryption or hashing with relatively short input, say up to 2K bytes. Almost all TCP/IP packets are 40 to 1.5K bytes, and the maximum packet lengths of major protocols, *e.g.*, Zigbee, Bluetooth low energy, and Controller Area Network (CAN), are less than 128 bytes. However, existing schemes are not well optimized for short input. To bridge the gap between real-world needs (in the future) and limited performances of state-of-the-art hash functions and authenticated encryptions with associated data (AEADs) for short input, we design a family of wide-block permutations Areion that fully leverages the power of AES instructions, which are widely deployed in many devices. As for its applications, we propose several hash functions and AEADs. Areion significantly outperforms existing schemes for short input and even competitive to relatively long messages. Indeed, our hash function is surprisingly fast, and its performance is less than three cycles/byte in the latest Intel architecture for any message size. It is significantly much faster than existing state-of-the-art schemes for short messages up to around 100 bytes, which are the most widely-used input size in real-world applications, on both the latest CPU architectures (IceLake, Tiger Lake, and Alder Lake) and mobile platforms (Pixel 7, iPhone 14, and iPad Pro with Apple M2).

Keywords: Short message · AES instruction · hash function · authenticated encryption · beyond 5G · IoT

1 Introduction

1.1 Background

In real-world communication environments, the overwhelming majority of cases require (authenticated) encryption or hashing with relatively short input, say up to 2K bytes. It is common knowledge that "real-world" TCP/IP packet length is biased towards short

*This is an updated and extended version from [28]. We added authenticated encryptions with associated data (AEADs) as an new application of our permutations. In addition, we have updated the performances of SHA2-256 and BLAKE3 by using code adapted from SUPERCOP (https://bench.cr.yp.to/supercop.html).

packets [54], as implemented by the standard benchmark method (Internet Mix[1] and the variants) for Internet routers etc. Packet sizes on the Internet generally follow a bimodal distribution, where 44% of packets are between 40 and 100 bytes long, and 37% are between 1400 and 1500 bytes in size. Low-power wireless protocols employ short packets, *e.g.*, the maximum packet length of Zigbee is 127 bytes and 47 bytes for Bluetooth low energy. The next Controller Area Network (CAN) standard, CAN-FD, has a maximum packet size of 64 bytes. In the use of narrow-band IoT [2], even the communication of 1-bit messages (*e.g.*, for device monitoring) is considered one of the target applications. For end-to-end encryption schemes in real-time video conference systems such as Zoom [34] and Webex [67, 57], which rapidly became popular due to the COVID-19 pandemic, the frame size is about 1K bytes. In these applications, an efficient hash function for short messages is essential. Specifically, to maintain the authenticity of the message, particularly against potentially malicious servers, the hash value of each packet/frame should be signed by digital signatures [29]. As such systems require real-time processing, the hash function should be as fast as possible for short messages.

Short inputs are also crucial for the future of mobile communications. So-called "beyond 5G" or 6G mobile communication technology will require short packets to achieve ultra-low latency communication. In comparison, some applications of 6G are expected to require a peak speed of over 100 GBps [44].

The importance of the short-input encryption/hash function has been widely recognized in the cryptographic research community. The NIST report on Lightweight Cryptography (LwC) [47, Sect 2.3.2] explicitly mentioned that lightweight applications typically need a hash function optimized for short messages, such as 256 bits. Also, fast processing for short messages is one of the important criteria for the ongoing NIST LwC standardization project for AEADs (see [1, Sect. 3.4]). Some NIST LwC proposals advertise their performances for short inputs, such as the winner, Ascon [19], a finalist Romulus [30], and second-round candidates ForkAE [3] and Saturnin [12].

The NIST LwC project targeted lightweight AEADs and hash functions, but only a few proposals use AES because the project mainly focuses on devices with low computational resources. On the other hand, the percentage of CPUs that have (the components of) AES as a dedicated instruction is rapidly increasing in the mobile and desktop PC world, represented by Intel AES-NI and ARMv8 AES instructions. Steam Hardware Survey shows that the number of CPUs with AES instructions is as high as 96.65% of the clients as of December 2022[2]. Standardization of AES instructions is also being considered for the RISC-V architecture [45], which is expected to become popular. This trend will also spread to low-end platforms like IoT edge devices.

1.2 Related Work

Short-(Fixed)-Input-Length (SFIL) Hash Functions. Haraka v2 is a SFIL hashing for post-quantum applications such as hash-based signature schemes [41]. However, a recent study by Bao *et al.* [6] reveals that preimage attacks on Haraka-256 and Haraka-512 up to 9 out of 10 rounds and 11 out of 10 rounds are feasible, respectively. That is, Haraka-512 is completely broken by their cryptanalyses, and the security margin of Haraka-256 is only one round. Simpira v2 is a family of permutations [25], and a short-input hashing is one of its applications. Although the security flaw of this application has yet to be found, it needs to be better optimized for recent CPU architectures, especially for a single permutation call, which is required for this application. For example, one round of the 256-bit variant requires two times AES round function calls, and each AES call should be sequentially executed because the second execution requires the output of the first execution. Because

[1] https://en.wikipedia.org/wiki/Internet_Mix
[2] https://store.steampowered.com/hwsurvey/Steam-Hardware-Software-Survey-Welcome-to-Steam

Intel Ice Lake or later processors can pipeline up to 6 AES instructions, it does not take full advantage of the pipeline.

Variable-Input-Length (VIL) Hash Functions. As efficient VIL hash functions, there are KangarooTwelve [10], ParallelHash256 [37], and BLAKE3 [56], in which parallel/tree hash structures allow to leverage pipeline and parallel executions to enhance the performance in software. However, these are effective only for long input, ie, processing short messages (less than 2K bytes) is much less efficient. In addition, KangarooTwelve and BLAKE3 guarantee the 128-bit preimage security. Other standard hash functions, such as SHA2-256 [65] and SHA3-256 [66], are also inefficient for short messages.

SFIL AEADs. As an efficient SFIL AEAD, there is KWF proposed by Khovratovich [39]. It uses a single public permutation as the underlying component to provide an efficient (deterministic) AEAD, also known as a "key wrapping scheme". In this scheme, associated data is processed with an unkeyed cryptographic hash function, and then the hashed associated data and the padded message are processed simultaneously with a single permutation. To achieve 128-bit security, it suggests Keccak-p[800] or Keccak-p[1600] as the permutation and SHA2-256 or SHA3-256 as the unkeyed cryptographic hash function, but these are not efficient for short inputs.

VIL AEADs. Standard authenticated encryption modes, such as OCB3 [42] and GCM [46], can be run extremely fast by using AES instructions for short messages. However, their data security is limited to 64 bits (due to the 128-bit block size of AES); namely, they can be broken with $\mathcal{O}(2^{64})$ encrypted blocks. This security level is not enough for the 6G era. On the other hand, more secure AEADs, such as Deoxys and OPP with BLAKE2b, are not well optimized for short messages, as we will discuss in detail below.

Deoxys is a family of nonce-based AEAD schemes [32, 33]. It is based on a new family of tweakable block ciphers, Deoxys-TBC, which uses the AES round function as a building block; thus, it can greatly benefit from AES instructions. Deoxys is regarded as a very efficient AEAD scheme for short messages, but it can be seen from [33, Tables 9 and 10] that short messages (*e.g.*, 64 bytes) are still processed around three times slower than long messages (*e.g.*, 65K bytes).

OPP is a permutation-based AEAD scheme [23]. It has excellent performance when compared to other permutation-based AEAD schemes, such as CAESAR submissions (*e.g.*, Ascon [18], Keyak [9], and NORX [4]), or the general SpongeWrap schemes [8, 50]. According to [23], OPP instantiated with the reduced-round BLAKE2b permutation achieves a peak speed of 0.55 cycles per byte on an Intel Haswell processor. Moreover, when compared to its competitors AES-GCM [20, 27], OCB3 [42], ChaCha20-Poly1305 [55], and Deoxys-I-128 [32, 33], this instantiation is faster by factors of around 1.87, 1.25, 3.80, and 1.74, respectively. Given that the instantiation of OPP without AES-NI is faster than Deoxys-I-128 with AES-NI, a permutation-based OPP scheme with AES-NI should have further high performance.

1.3 Motivation

Looking into real-world applications, efficient hash functions, and AEADs for short inputs up to 2K bytes is essential and will become increasingly important. However, existing schemes need to be better optimized for short input. Especially, there still needs to be satisfactory VIL hash functions for short messages regarding speed and security. To bridge the gap between the need for real-world applications and the performance of state-of-the-art hash functions and AEADs, we aim to design efficient and secure hash functions and AEADs for short messages.

Specifically, our design goal for hash functions is to be highly efficient for short messages up to 2K bytes, and competitive even for long messages to software-efficient hash functions KangarooTwelve, ParallelHash256, and BLAKE3 on modern desktop and mobile platforms. In addition, our design goal for AEADs is to be more efficient than existing software-efficient schemes, such as Deoxys-I-128 and BLAKE2s-OPP.

1.4 Our Contribution

To achieve our design goals, we first specify a family of efficient permutations Areion that is optimized for the latest CPU architectures, including Intel and ARM, by fully leveraging the power of AES instructions. As its applications, we propose SFIL and VIL hashing in addition to SFIL and VIL AEADs. We then evaluate the security of underlying permutations and their applications and measure software performances in several architectures. Our contributions in this paper are summarized as follows.

Software-Efficient AES-Based Permutations. For environments where AES instructions are available, we design a family of permutations, dubbed Areion, that can be implemented by only AES instructions such as aesenc and aesenclast in AES-NI or vaeseq and vaesmcq in ARMv8 NEON as AES instructions are most efficient cryptographic operation among SIMD operations. As an underlying structure, we propose pipeline-friendly Feistel-type schemes in which additional F functions are appended to Feistel-type schemes to take full advantage of the pipeline executions. We find optimal instantiations of F functions by thoroughly analyzing the security and performance of all possible candidates. As a result, the performance of Areion is significantly faster than existing permutations in the latest CPU architectures. Especially, Areion outperforms other permutations in the encrypt direction. It is an important characteristic of our target applications.

SFIL Hash Function. For an SFIL hashing, we apply Areion to the Davies-Meyer (DM) construction, which consists of a permutation with a feed-forward (applying the XOR operation) of the input as with Simpira v2 [25] and Haraka v2 [41]. Our schemes provide a 256-bit security level against preimage attacks. In addition, these are about 1.4 times faster than the schemes based on Simpira v2.

VIL Hash Function. For a VIL hashing, we design a compression function based on Areionand implant it to the general Merkle-Damgård (MD) construction [51, 16]. Our scheme performs much faster than any other hash functions for the input size up to 1024 bytes and even competitive with other software-efficient hash functions for longer inputs in laptop and mobile environments while ensuring the 256-bit security level of preimage attacks. Its performance is less than three cycles/byte for any message size. It is much faster than existing state-of-the-art schemes for short messages up to around 100 bytes. Such message lengths are typical in real-world applications on the latest CPU architectures (IceLake, Tiger Lake, and Alder Lake) and mobile platforms (Pixel 7, iPhone 14, and iPad Pro with Apple M2).

SFIL AEAD. For an SFIL AEAD scheme, we apply Areion to a (deterministic) AEAD proposed by Khovratovich [39]. Our scheme is about more than 1.2 times faster than the scheme based on Simpira v2 for both encryption and decryption.

VIL AEAD. For VIL AEAD schemes, we apply Areion to the Offset Public Permutation (OPP) mode [23] and the Offset Two-Round (OTR) mode [52]. Consequently, our schemes have better performance than any other target AEADs, such as Deoxys-I-128 and the schemes based on BLAKE2s and Simpira v2.

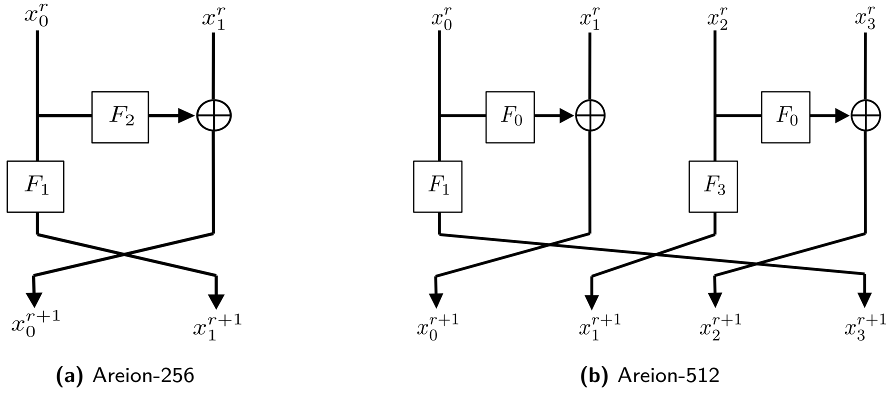

(a) Areion-256 **(b)** Areion-512

Figure 1: The round functions of Areion.

1.5 Paper Organization

In Sect. 2, we describe the specification of Areion. Sect. 3 explains details of our design rationale of Areion and discusses the optimally of our design choices. In Sect. 4, we show several applications of Areion. In Sects. 5 and 6, we give the security and performance evaluations of Areion and its applications, respectively. Sect. 7 concludes the paper.

2 Specification of Permutations

We show the specification of Areion. Areion is based on Simpira v2 but has the structure that allows more AES instructions to be executed in parallel. We provide the following two variants of our permutation: Areion-256 and Areion-512. The former accepts a 256-bit block, and the latter accepts a 512-bit block as input.

To illustrate the specification of each permutation, we denote by F_i $(i \in \{0, 1, 2, 3\})$ the function based on the operations in the AES round function. Let SubBytes, ShiftRows, MixColumns, and AddRoundConstant in the AES round function be SB, SR, MC, and AC, respectively. AC is equivalent to AddRoundKey in ordinal AES, but the constant is added instead of the round key. F_i consists of a combination of SB, SR, MC and AC. For each value of i, F_i is defined as follows:

$$
\begin{aligned}
F_0 &= MC \circ SR \circ SB \\
F_1 &= SR \circ SB \\
F_2 &= MC \circ SR \circ SB \circ AC \circ MC \circ SR \circ SB \\
F_3 &= MC \circ SR \circ SB \circ AC \circ SR \circ SB
\end{aligned}
$$

A combination of AES instructions in AES-NI or NEON can implement these functions. Areion-256 consists of F_1 and F_2, and Areion-512 consists of F_0, F_1, and F_3. The round function of each variant is shown in Fig. 1.

We set the number of rounds of Areion-256 and Areion-512 are 10 and 15, respectively. These are derived from our security evaluation. Sect. 5 describes the details. The round constants are derived from the binary digits of a fraction part of $\pi = 3.1415926\cdots$. Table 1 shows round constants in hexadecimal notation, and round constants are used in little-endian byte order. In the r-th round of Areion, RC_r is added to the state.

Table 1: Round constants.

RC	Round constant
RC_0	0x243f6a8885a308d313198a2e03707344
RC_1	0xa4093822299f31d0082efa98ec4e6c89
RC_2	0x452821e638d01377be5466cf34e90c6c
RC_3	0xc0ac29b7c97c50dd3f84d5b5b5470917
RC_4	0x9216d5d98979fb1bd1310ba698dfb5ac
RC_5	0x2ffd72dbd01adfb7b8e1afed6a267e96
RC_6	0xba7c9045f12c7f9924a19947b3916cf7
RC_7	0x801f2e2858efc16636920d871574e690
RC_8	0xa458fea3f4933d7e0d95748f728eb658
RC_9	0x718bcd5882154aee7b54a41dc25a59b5
RC_{10}	0x9c30d5392af26013c5d1b023286085f0
RC_{11}	0xca417918b8db38ef8e79dcb0603a180e
RC_{12}	0x6c9e0e8bb01e8a3ed71577c1bd314b27
RC_{13}	0x78af2fda55605c60e65525f3aa55ab94
RC_{14}	0x5748986263e8144055ca396a2aab10b6

3 The Design

3.1 AES Instructions and SIMD

SIMD is an abbreviation for Single Instruction Multiple Data and a type of parallel processing. Most modern processors support instructions set for SIMD. SIMD instructions perform operations vector-wise using data stored in dedicated registers, which allows arithmetic/bitwise operations in parallel and advanced operations like data shuffling to be performed with a single instruction.

An example of SIMD instructions that can perform complex operations is an instruction for executing AES, the dominant block cipher. This instruction belongs to AES-NI (AES New Instructions set) in the Intel/AMD processors. AES-NI includes aesenc to perform the round function of the encryption, aesenclast for the final round, instructions for decryption, and instructions to support the round key generation. On the other hand, in the ARMv8 processors, AES instructions are included in the NEON instructions set. AES instructions in NEON include vaeseq for AddRoundKey, SubBytes, and ShiftRows, and vaesmcq for MixColumns. NEON also supports the decryption instructions, while instructions for the round key generation are not supported.

The performance of SIMD instructions can be measured by their latency, throughput, and port usage. Latency means the number of clock cycles that are required for the execution of an instruction. Throughput means the number of clock cycles required before the responsible ports can accept the same instruction again. Dispatched instructions are decomposed into micro-operations and then processed by each execution port.

According to the website by Abel and Reineke *et al.* [58], the latency and throughput of aesenc/aesenclast in Ice Lake are 3 and 0.5, respectively. A throughput of 0.5 means that two execution ports can accept the micro-operation from aesenc/aesenclast and each operation's throughput is 1 [58]. Fig. 2 illustrates the pipelined execution of multiple aesenc on Ice Lake. We can see that up to 6 aesenc can be executed in 5 cycles on Ice Lake using two execution ports, port 0 and port 1.

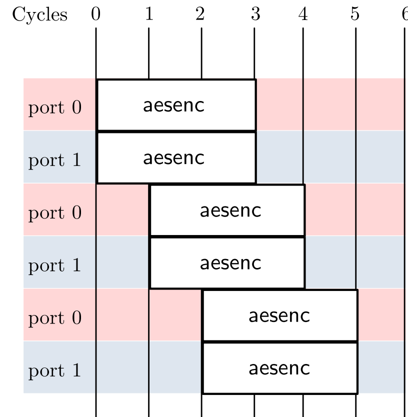

Figure 2: Execution of aesenc on Ice Lake processors.

Table 2: The latency and throughput of aesenc, referred by [58].

<table>
<tr><th rowspan="2">Processor</th><th colspan="2">aesenc</th></tr>
<tr><th>Latency</th><th>Throughput</th></tr>
<tr><td>Skylake</td><td rowspan="4">4</td><td rowspan="3">1</td></tr>
<tr><td>Kaby Lake</td></tr>
<tr><td>Coffee Lake</td></tr>
<tr><td>Cannon Lake</td><td rowspan="4">0.5</td></tr>
<tr><td>Ice Lake</td><td rowspan="3">3</td></tr>
<tr><td>Tiger Lake</td></tr>
<tr><td>Alder Lake</td></tr>
<tr><td>Zen +</td><td rowspan="2">4</td><td rowspan="2">0.5</td></tr>
<tr><td>Zen2</td></tr>
</table>

3.2 General Construction

3.2.1 Permutations Realized by only AES Instructions

To construct optimal permutations in environments where hardware instructions of AES are available, we focus on a class of *permutations that can be implemented solely by AES instructions* such as aesenc and aesenclast in AES-NI or vaeseq and vaesmcq in ARMv8 NEON for the following reasons.

- The latency of AES instructions in AES-NI becomes smaller as the processor's architecture is upgraded. Moreover, Intel 9th generation and later processors have an additional execution port that accepts micro-operations generated from AES instructions, which improves the throughput from 1 to 0.5. The latency and throughput of aesenc in Intel processors from 6 to 11 generation are shown in Table 2.

- Schemes based solely on AES instructions are beneficial in terms of performance and security. Since NIST selected AES as a standard block cipher in 2001, no attack has been published despite considerable cryptanalytic efforts over the past 20 years, and its security is deeply understood in the community of symmetric cryptography. Thus, it is easy to evaluate its security by existing tools convincingly and accumulated cryptanalysis knowledge.

- Haraka v2 [41] is a family of permutations. It is an SPN-type scheme based on AES instructions and word shuffle operations, such as unpack instructions. Shiba

et al. show that the structure of Haraka v2 is optimal among SPN-type schemes based solely on AES instructions and shuffle operations [64]. Thus, presenting a new SPN-type scheme with better performance than Haraka v2should be challenging.

- Word shuffle operations provide only simple linear transformations. In contrast, AES instructions include not only more complex linear operations (*i.e.*, MixColumns and ShiftRows) but also nonlinear operations (*i.e.*, 16 parallel executions of 8-bit S-box) by only a single instruction call. In addition, the latency of word shuffle operations requires one, even on the latest CPU architectures. Thus, arguably AES instructions are the most efficient and cryptographically-strong operations in all SIMD instructions.

- Haraka v2 does not provide a sufficient level of security as a hash function according to the recent study by Bao *et al.* [6]. They present preimage attacks on Haraka-256 and Haraka-512 up to 9 out of 10 rounds and 11 out of 10 rounds, respectively. In addition, designers of Haraka v2 did not claim any security as a public permutation. According to these facts, Haraka v2 should require roughly 1.2 to 1.5 times of recommended rounds by the designers, *i.e.*, about 12 to 15 rounds to ensure the security as public permutations of Haraka v2 and hash functions. These additional rounds degrade the performance of Haraka v2 significantly. We remark that, due to the structure of Haraka v2, increasing the number of rounds requires not only more AES instructions but also more word shuffle operations. Thus, it significantly impacts the overall performance of the tweaked versions of Haraka v2 compared to the Feistel-type scheme such as Simpira v2, which is a class of *permutations that can be implemented solely by AES instructions.*

For the above reasons, we choose the Feistel-type scheme to design new 256- and 512-bit permutations from 128-bit AES instructions.

3.2.2 Feistel-type Scheme for Leveraging the Pipeline

Limitations of Simpira v2. For the 256- and 512-bit variants of Simpira v2 (hereafrer, we will refer to each variant as Simpira-256 and Simpira-512, respectively), there is still room for improvement in their design, considering the characteristic of AES instructions in modern processors, especially for applications that require sequential executions of underlying permutations, *e.g.*, SFIL and VIL hash functions.

Specifically, the one-block encryption of Simpira-256 requires two times of executions, and each AES call should be sequential because the second execution requires the output of the first execution. On the other hand, one-block encryption of Simpira-512 is capable of pipelining up to two 2-round AES executions. However, since Intel Ice Lake or later processors can pipeline up to 6 AES instructions, the structure of Simpira-256 and Simpira-512 does not take full advantage of the pipeline.

Pipeline-Friendly Feistel-type Schemes. To take advantage of the pipeline as possible, we design pipeline-friendly Feistel-type schemes in which F functions are added in the left branch for the 256-bit version and first and third branches for the 512-bit version to Feistel-type scheme, respectively, as shown in Fig. 3. These allow for pipelined execution of two and four AES instructions, respectively.

As another possible scheme, we can add F functions in the right branch for the 256-bit version and the second and fourth branches for the 512-bit version to the above schemes before XOR operations, respectively. However, our initial evaluation confirmed that these additional instructions do not improve the performance because they cannot significantly reduce the required number of rounds to ensure the security of structural attacks on Feistel, such as impossible differential and integral attacks. Besides, the critical path in

Table 3: Instructions per cycle (IPC) of each permutation.

Algorithm	#Round	IPC
Areion-256	10	**0.66**
Simpira-256	15	0.46
Areion-512	15	**0.92**
Simpira-512	15	0.52

the decryption of this scheme becomes three times longer than that of the encryption. From these facts, we conclude that *the schemes in Fig. 1 are optimal for 2- and 4-line Feistel-type schemes for high performance.*

Comparison. In order to compare the degree of utilization of the pipeline, we checked instructions per cycle (IPC) of each variant of Areion and Simpira v2 by static code analysis using LLVM machine code analyzer (llvm-mca). Table 3 shows the results. For both variants, the results show the IPC of Areion is larger than that of Simpira v2. Based on this fact, the construction of Areion can utilizes the pipeline more effectively.

3.3 Finding Optimal Constructions

Possible Candidates of F Functions. Recall that our permutations are realized solely by AES instructions. As already discussed in [41, 25, 64], F functions consisting of one or two AES round functions are optimal in Feistel- and SPN-type schemes. In this work, to find further efficient constructions, we also consider last-round instructions such as aesenclast in AES-NI or vaesmcq in ARMv8 NEON, respectively, as underlying instructions. Thus, F functions should be realized by one or two combinations of aesenc and aesenclast in AES-NI or vaeseq and vaesmcq in ARMv8 NEON, respectively.

There are six possible candidates of F_i ($i \in \{0, 1, 2, 3, 4, 5\}$), where F_0, F_1, F_2, F_3 are defined in Sect. 2 and F_4 and F_5 are as follows.

$$F_4 = SR \circ SB \circ AC \circ MC \circ SR \circ SB$$
$$F_5 = SR \circ SB \circ AC \circ SR \circ SB$$

For AES-NI, F_0, F_1, F_2, F_3, F_4 and F_5 are implemented by aesenc, aesenclast, aesenc $\rightarrow$ aesenc, aesenclast $\rightarrow$ aesenc, aesenc $\rightarrow$ aesenclast, and aesenclast $\rightarrow$ aesenclast, respectively. Note that XOR operations in the Feistel-type scheme are executed by the operation of AddRoundKey, which is the last operation of aesenc and aesenclast, respectively. This feature of AddRoundKey is the reason why AC is absent in the last of these equations.

For ARMv8 NEON, F_0, F_1, F_2, F_3, F_4 and F_5 are implemented by vaeseq $\rightarrow$ vaesmcq, vaeseq, vaeseq $\rightarrow$ vaesmcq $\rightarrow$ vaeseq $\rightarrow$ vaesmcq, vaeseq $\rightarrow$ vaeseq $\rightarrow$ vaesmcq, vaeseq $\rightarrow$ vaesmcq $\rightarrow$ vaeseq, vaeseq $\rightarrow$ vaeseq, respectively. As vaeseq performs AddRoundKey before SubBytes, the AddRoundKey operation of the first vaeseq in each function is used to realize the XOR operation of the previous round for Feistel-type schemes. This observation implies that our schemes can be implemented solely by vaeseq and vaesmcqin NEON, except for the XOR operation in the last round.

How to Find F functions. To find optimal combinations of functions F_i ($i \in \{0, 1, 2, 3, 4, 5\}$) in Fig. 3, we first evaluate the security against differential/linear, impossible differential, and integral attacks using Mixed-Integer Linear Programming (MILP) for all combinations. Let R_1, R_2, and R_3 be the number of rounds where the following three conditions are satisfied, respectively.

- R_1: The number of rounds where the minimum number of active S-boxes is enough to ensure security against differential/linear attacks.

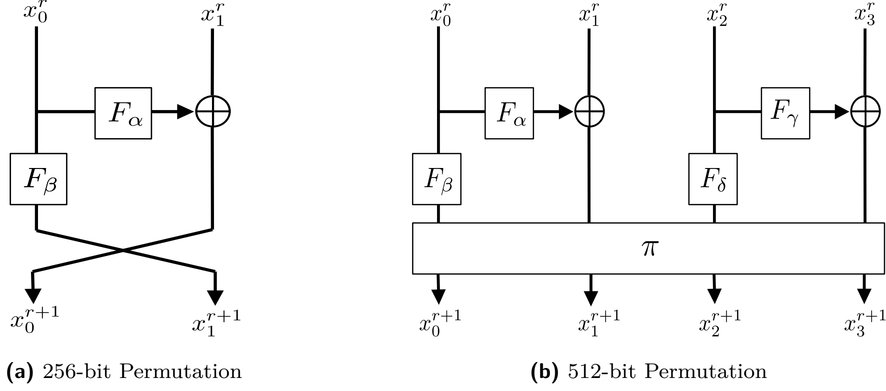

(a) 256-bit Permutation **(b)** 512-bit Permutation

Figure 3: Target constructions of our permutations.

R_2: The number of rounds with no byte-truncated impossible differential characteristic.

R_3: The number of rounds with no byte-wise integral distinguisher.

Besides, we define $\max\{R_1, R_2, R_3\}$ as $R_{\max}$. After obtaining $R_{\max}$, we will look into characterizes for the performance in $R_{\max}$ to find the most efficient ones. The details are explained in the following.

3.3.1 On 256-bit Permutations

Let a 256-bit permutation with F_α and F_β functions be (α, β)-perm, where $\alpha, \beta \in \{0, 1, 2, 3, 4, 5\}$, as illustrated in Fig. 3. As a 256-bit permutation has two functions with six possible candidates, the total number of combinations is 36 $(= 6 \times 6)$. Among them, we look for combinations implemented by the lowest number of AES instructions in $R_{\max}$, *i.e.*, we choose the ones that can achieve the required security level with the lowest number of AES instructions.

Table 4 shows R_1, R_2, R_3, $R_{\max}$ and the number of AES instructions in $R_{\max}$ of all 36 candidates. According to this table, the lowest one is $(2, 1)$-perm for which, R_1, R_2 and R_3 are estimated as 5, 5, and 4, respectively, namely, $R_{\max} = 5$, and #AES instructions in 5 rounds is only 15. From this result, we select $(2, 1)$-perm as underlying one for Areion-256.

3.3.2 On 512-bit Permutations

Let a 512-bit permutation with F_α, F_β, F_γ and F_δ functions using π block shuffle layer be $(\alpha, \beta, \gamma, \delta, \pi)$-perm, where $\alpha, \beta, \gamma, \delta \in \{0, 1, 2, 3, 4, 5\}$, as illustrated in Fig. 3. As a 512-bit permutation has four F functions in which there are six possible candidates and π block shuffle has 24 $(= 4!)$ patterns, the total number of combinations is estimated as 31104 $(= 6 \times 6 \times 6 \times 6 \times 4!)$.

We thoroughly analyze security and performance using the following procedures to find the most efficient combination among them.

Step 1: Limiting the Number of AES Instructions in $R_{\max}$. As with the 256-bit case, we focus on combinations implemented by the lowest number of AES instructions in $R_{\max}$. As a result of our search, we find 30 candidates in which the lowest number of AES instructions in $R_{\max}(= 9)$ is 45, as shown in Table 5.

Step 2: Eliminating the Equivalent Candidates. Twenty-eight candidates out of the remaining 30 can be classified into 14 equivalent classes, *i.e.*, each two candidates of them is mapped to one equivalent class. Based on this fact, we can eliminate 14 equivalent classes, and then reduce to 16 $(= 30 - 14)$ candidates.

Table 4: Search results on 256-bit permutations.

(α, β)	R_1	R_2	R_3	R_{max}	#AES instructions
$(0,0)$	23	7	5	23	46
$(0,1)$	8	9	5	9	18
$(0,2)$	6	5	4	6	18
$(0,3)$	16	7	4	16	48
$(0,4)$	6	6	4	6	18
$(0,5)$	6	8	5	8	24
$(1,0)$	8	9	6	9	18
$(1,1)$	33	-	-	-	-
$(1,2)$	6	5	4	6	18
$(1,3)$	6	9	6	9	27
$(1,4)$	6	9	5	9	27
$(1,5)$	23	-	-	-	-
$(2,0)$	6	5	4	6	18
$(2,1)$	**5**	**4**	**5**	**5**	**15**
$(2,2)$	6	4	3	6	24
$(2,3)$	6	5	4	6	24
$(2,4)$	4	5	3	5	20
$(2,5)$	5	5	4	5	20
$(3,0)$	16	7	4	16	48
$(3,1)$	7	9	5	9	27
$(3,2)$	6	5	4	6	24
$(3,3)$	12	-	4	-	-
$(3,4)$	4	6	4	6	24
$(3,5)$	7	-	5	-	-
$(4,0)$	6	7	4	7	21
$(4,1)$	6	9	5	9	27
$(4,2)$	4	5	3	5	20
$(4,3)$	4	6	4	6	24
$(4,4)$	7	-	3	-	-
$(4,5)$	6	-	5	-	-
$(5,0)$	7	8	6	8	24
$(5,1)$	23	-	-	-	-
$(5,2)$	4	5	4	5	20
$(5,3)$	7	-	6	-	-
$(5,4)$	6	-	5	-	-
$(5,5)$	17	-	-	-	-

Step 3: Considering Efficiency in NEON Instructions. The remaining 16 combinations can be classified into three different classes. Specifically, each different class has the following different π:

$$\pi_1 : x_0^r||x_1^r||x_2^r||x_3^r \mapsto x_1^{r+1}||x_2^{r+1}||x_3^{r+1}||x_0^{r+1}$$
$$\pi_2 : x_0^r||x_1^r||x_2^r||x_3^r \mapsto x_3^{r+1}||x_0^{r+1}||x_1^{r+1}||x_2^{r+1}$$
$$\pi_3 : x_0^r||x_1^r||x_2^r||x_3^r \mapsto x_1^{r+1}||x_3^{r+1}||x_0^{r+1}||x_2^{r+1}$$

The two constructions in π_3 are unsuitable for implementations using NEON instructions in ARMv8. This is because the implementation of these two constructions by NEON requires successive XORs, which hampers the implementation with only `vaeseq` and `vaesmcq` while maintaining the compatibility of the implementation on ARM and Intel. Based on this fact, we eliminate these two constructions using π_3

Table 5: Search results of 512-bit permutations.

$(\alpha, \beta, \gamma, \delta, \pi)$	R_1	R_2	R_3	R_{max}	#AES instructions
$(0, 0, 0, 4, \pi_1)$	8	9	5	9	45
$(0, 0, 1, 2, \pi_1)$	9	9	6	9	45
$(0, 0, 2, 1, \pi_1)$	9	9	6	9	45
$(0, 0, 4, 0, \pi_1)$	8	9	6	9	45
$(0, 0, 5, 0, \pi_1)$	9	9	7	9	45
$(0, 1, 0, 2, \pi_1)$	9	9	6	9	45
$(0, 1, 0, 3, \pi_1)$	9	9	6	9	45
$(0, 1, 0, 4, \pi_1)$	8	9	6	9	45
$(0, 1, 2, 0, \pi_1)$	9	9	6	9	45
$(0, 1, 2, 1, \pi_1)$	8	9	6	9	45
$(0, 2, 0, 1, \pi_1)$	9	9	6	9	45
$(0, 2, 1, 0, \pi_1)$	9	9	6	9	45
$(0, 3, 0, 1, \pi_1)$	9	9	6	9	45
$(0, 4, 0, 0, \pi_1)$	8	9	5	9	45
$(0, 4, 0, 1, \pi_1)$	8	9	6	9	45
$(1, 0, 0, 2, \pi_1)$	9	9	6	9	45
$(1, 0, 2, 0, \pi_1)$	9	9	7	9	45
$(1, 2, 0, 0, \pi_1)$	9	9	6	9	45
$(2, 0, 0, 1, \pi_1)$	9	9	6	9	45
$(2, 0, 1, 0, \pi_1)$	9	9	7	9	45
$(2, 1, 0, 0, \pi_1)$	9	9	6	9	45
$(2, 1, 0, 1, \pi_1)$	8	9	6	9	45
$(4, 0, 0, 0, \pi_1)$	8	9	6	9	45
$(5, 0, 0, 0, \pi_1)$	9	9	7	9	45
$(0, 1, 0, 2, \pi_2)$	8	9	5	9	45
$(0, 2, 0, 1, \pi_2)$	8	9	5	9	45
$(1, 0, 2, 0, \pi_2)$	9	9	5	9	45
$(2, 0, 1, 0, \pi_2)$	9	9	5	9	45
$(0, 0, 1, 2, \pi_3)$	9	9	6	9	45
$(0, 0, 1, 4, \pi_3)$	9	9	6	9	45

from the candidates. As a result, we obtain 14 constructions.

Step 4: Estimating Theoretical Number of Cycles. For the remaining 14 candidates, we use a performance analysis tool llvm-mca to estimate the theoretical number of cycles in Ice Lake or later architecture. Table 6 shows theoretical values of total cycles in 15-round encryption, calculated by llvm-mca. According to this result, we reduce to 6 candidates with the lowest number of cycles to perform the encryption.

Step 5: Performing Experimental Evaluations. We measure the performance of the remaining six candidates on several platforms. Table 6 shows the results on Ice Lake architecture (Intel(R) Core(TM) i7-1068NG7 CPU @ 2.30GHz). From these results, we selected $(0, 1, 0, 3, \pi_1)$-perm as the optimal combination for Areion-512.

Table 6: Theoretical and experimental value of total cycles at 15-round encryption.

$(\alpha, \beta, \gamma, \delta, \pi)$	total cycle (by llvm-mca)	cpb (by experiments)
$(0,0,0,4,\pi_1)$	**8899**	1.09016
$(0,0,1,2,\pi_1)$	**8899**	1.08927
$(0,0,2,1,\pi_1)$	11528	-
$(0,0,4,0,\pi_1)$	11528	-
$(0,0,5,0,\pi_1)$	11528	-
$(0,1,0,2,\pi_1)$	**8899**	1.08918
$(0,1,0,3,\pi_1)$	**8899**	**1.08882**
$(0,1,0,4,\pi_1)$	**8899**	1.08989
$(0,1,2,0,\pi_1)$	11528	-
$(0,1,2,1,\pi_1)$	11528	-
$(1,0,0,2,\pi_1)$	**8899**	1.09017
$(1,0,2,0,\pi_1)$	11528	-
$(0,1,0,2,\pi_2)$	9109	-
$(1,0,2,0,\pi_2)$	9919	-

4 Applications

4.1 Permutation-based Hash Functions

4.1.1 SFIL Hash Function.

For an SFIL hashing, we apply Areion to the Davies-Meyer (DM) construction, which consists of a permutation with a feed-forward (applying the XOR operation) of the input. The use of DM for SFIL hashing has already been discussed in [25, 41]. In particular, Haraka v2 implemented two SFIL hash functions, Haraka256-DM : $\mathbb{F}_2^{256} \to \mathbb{F}_2^{256}$ and Haraka512-DM : $\mathbb{F}_2^{512} \to \mathbb{F}_2^{256}$, defined as follows:

$$\mathsf{Haraka256\text{-}DM}(x) = \pi_{256}(x) \oplus x, \quad (1)$$

$$\mathsf{Haraka512\text{-}DM}(x) = \mathsf{trunc}(\pi_{512}(x) \oplus x), \quad (2)$$

where π_{256} and π_{512} are the 256- and 512-bit permutations of Haraka v2, respectively; and trunc : $\mathbb{F}_2^{512} \to \mathbb{F}_2^{256}$ is a truncation function defined as follows:

$$\mathsf{trunc}(x_0||\cdots||x_{15}) = x_2||x_3||x_6||x_7||x_8||x_9||x_{12}||x_{13}, \quad (3)$$

where $x = x_0||\cdots||x_{15} \in \mathbb{F}_2^{512}$. Our SFIL hash functions, Areion256-DM and Areion512-DM, use Areion-256 and Areion-512 instead of Haraka v2's ones. The DM construction uses only the forward direction of the permutation, and the overhead beyond the permutation is negligible. Thus, the performances of Areion256-DM and Areion512-DM are effectively the same as those of the forward direction of underlying permutations.

The designers of Simpira v2 suggested its application to SFIL hash functions [25]. Then, for performance comparison, we define DM construction instantiations of Simpira v2 in the same way as above and refer to them as Simpira256-DM and Simpira512-DM.

4.1.2 VIL Hash Function.

For a VIL hashing, we apply Areion-512 to the Merkle-Damgård (MD) construction, a classical method of building a cryptographic hash function from a compression function [51, 16].

Our VIL hash function, Areion512-MD, is an MD construction instantiated with Areion512-DM. Other design details of Areion512-MD follow SHA2-256 [65]. SHA2-256 has

two phases, preprocessing and hash computation phases. The former is further divided into three steps: padding the message, parsing the message into message blocks, and setting the initial hash value. For Areion512-MD, padding, and message parsing are executed in the same procedure as SHA2-256. However, the length of the padded message should be adjusted to be a multiple of 256 bits instead of a multiple of 512 bits; and the size of the parsed message block is 256 bits (see [65, Section 5] for more details). Areion512-MD uses the same initial hash value H of SHA2-256, and it consists of the following two 128-bit words:

$$H_0 = \texttt{0x6a09e667bb67ae853c6ef372a54ff53a}, \tag{4}$$

$$H_1 = \texttt{0x510e527f9b05688c1f83d9ab5be0cd19}. \tag{5}$$

Then, Areion512-DM is used for the hash computation phase. The parsed message block is inserted into x_0 and x_1 of the input word positions in Areion-512, and the initial hash value and chaining values (that is, the output value of each compression function) are set into x_2 and x_3 of the input word positions in Areion-512 (see Fig. 1b). Finally, the output value of the last DM compression function becomes a 256-bit message digest.

The designers of Simpira v2 and Haraka v2 did not mention its application to VIL hash functions [25, 41]. However, for performance comparison, we define an MD construction instantiation of Simpira512-DM and Haraka512-DM in the same way as above and refer to them as Simpira512-MD and Haraka512-MD, respectively.

4.2 Permutation-based AEAD Schemes

4.2.1 SFIL AEAD.

To implement an SFIL AEAD scheme, we apply Areion to a (so-called) key wrapping scheme, a class of deterministic authenticated encryption (DAE) defined by Rogaway and Shrimpton [60]. More precisely, key wrapping schemes allow us to wrap (encrypt) a new short-term secret key with a long-term master key shared between a sender and a receiver. We present Areion512-KWF, which is a key wrapping scheme instantiated with Areion-512 and Areion512-MD. Our design is similar to KWF [39]. It is a variant of encode-then-encipher scheme [7] and needs a single permutation call to process the whole input. The specification of Areion512-KWF is the same as the original KWF, and Areion-512 and Areion512-MD are used as a fixed 512-bit permutation $\mathcal{F}$ and a collision-resistant hash function $\mathcal{G}$ for an associated data, respectively (see [39, Section 3.2] for more details). According to [39], to achieve the security level of s bits, the recommended parameters are defined as follows:

$$k \geq s, \quad \ell \geq 2s, \quad n \geq 2s + k, \quad m = n - k - \ell, \tag{6}$$

where k, ℓ, n, and m denote the key size, the hash size for associated data, the block size of the permutation, and the size of a padded message, respectively. For example, to achieve 128-bit security for $m = 128$, we need $k \geq 128$, $\ell \geq 256$, and $n \geq 512$. This observation suggests that Areion-512 is the natural choice to implement (a variant of) KWF.

In [39], it is proved that KWF is provably secure as a DAE. The security proof assumes that the underlying permutation is a public random permutation and that the hash function is collision-resistant. The former assumption can be translated into the situation of no structural weaknesses in the permutation. Our design goals for Areion-512 and Areion512-MD cover both assumptions. It is known that, by making a nonce as a part of associated data, a DAE also implements a misuse-resistant AE (MRAE), which provides the maximum protection against misuse (repeat) of a nonce [60]. Standard AEAD schemes, such as GCM [46], OCB3 [42], and ChaCha20-Poly1305 [55], are not MRAE hence lack this property. In this sense, Areion512-KWF provides stronger security than these schemes.

The designers of Simpira v2 provided its application to AEAD schemes [25]. For performance comparison, we design KWF instantiated with Simpira-512 and Simpira512-MD in the same way as above, and we refer to this variant as Simpira512-KWF.

4.2.2 VIL AEADs.

To build VIL AEADs, we rely on Offset Public Permutation (OPP) [23] and Offset Two-Round (OTR) modes [52].

OPP. The OPP mode, designed by Granger *et al.* at EUROCRYPT 2016 [23], is a fully parallelizable nonce-based AEAD scheme based on the Masked Even-Mansour (MEM) tweakable block cipher construction. The MEM construction improves the efficiency of the conventional Tweakbale Even-Mansour (TEM) construction [63, 13, 14, 49] by using an efficient word-oriented LFSR- and powering-up-based masking function. We can expect a highly-efficient AEAD scheme by combining the MEM construction with a highly-efficient permutation such as Areion. We present Areion256-OPP, an OPP mode instantiated with Areion-256. The specification of Areion256-OPP is almost the same as the original OPP mode (see [23, Sects. 2–4] and the publicly available source code[3] for more details), excluding an LFSR update function in the masking function and a method of setting the secret key K and nonce N, and Areion-256 is used as the underlying permutation in the MEM construction.

The masking function: The masking function in the MEM consists of word-oriented LFSRs. The LFSR update function depends on the permutation size in the MEM construction and the word size of the LFSR. [23, Table 1] describes several examples of the LFSR update function, and we choose the following:

$$\varphi : (x_0, x_1, x_2, x_3) \mapsto (x_1, x_2, x_3, (x_0 \lll 3) \oplus (x_3 \gg 5)), \tag{7}$$

where the state size b is 256 bits, the word size w is 64 bits, and the number of words n is 4 (see [23, Section 3.4] for more details).

Input formatting: [23] did not specify the recommended parameters, such as the key and nonce sizes. We define 128- or 256-bit key and 128-bit nonce as our recommendations. In OPP, the underlying permutation in MEM takes the concatenation of key and nonce, *i.e.*, $N||K$, for the initialization of the masking function. When using a 128-bit key, this method works with Areion-256. However, when using a 256-bit key, the above method does not work as the input size exceeds the block size of Areion-256. In this case, we recommend setting the following value as the input word positions x_0 and x_1 in Areion-256 (see Fig. 1a) with a 256-bit key:

$$x_0||x_1 = (N||0^{128}) \oplus K, \tag{8}$$

where 0^{128} is the 128-bit zero padding value[4]. Another solution would be to use Areion-512 instead of Areion-256. However, Areion-256 can achieve a 256-bit security with the method of setting the key and nonce following Eq. (8), and is superior to Areion-512 in terms of its efficiency (see Sect. 6 for more details); therefore, it is better to use Areion-256 as the underlying permutation in the OPP mode.

For performance comparison, we instantiate OPP with Simpira-256 similarly and refer to this variant as Simpira256-OPP.

[3] https://github.com/MEM-AEAD/mem-aead

[4] We do not claim the security against related-key attacks because it is evident that a tuple (N, N', K, K') satisfying $(N||0^{128}) \oplus K = (N'||0^{128}) \oplus K'$ breaks the scheme.

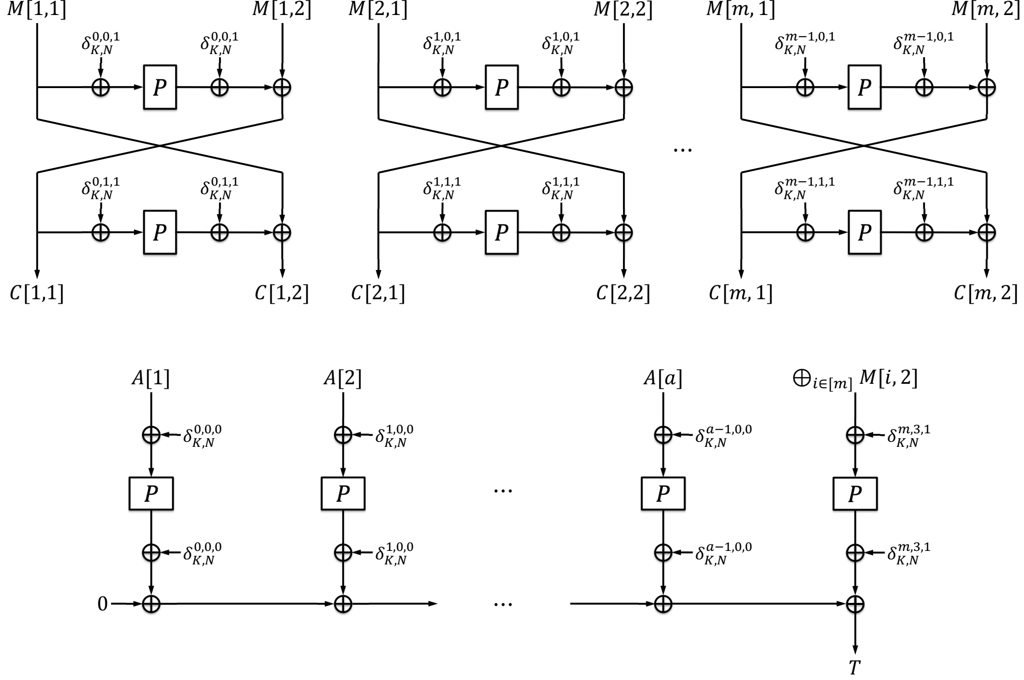

Figure 4: The overall structure of the permutation-based OTR mode. K, N, M, C, A, T, δ, and P denote the secret key, nonce, message, ciphertext, associated data, tag, MEM masking function, and the underlying permutation.

OTR. The OTR mode, proposed by Minematsu at EUROCRYPT 2014 [52], is a block cipher mode of operation to realize an AEAD scheme. The OTR mode is a rate-1 (*i.e.* one block cipher call per one input block) parallel AEAD scheme. It has provable security up to the birthday bound. In other words, it is secure up to $O(2^{n/2})$ processed blocks. Moreover, OTR only needs the block cipher function's forward direction (inverse-freeness) for its encryption and decryption. Effectively, these features show that OTR is as efficient and secure as OCB [59], while removing the necessity of block cipher inverse. We define Areion256-OTR, a permutation-based OTR mode based on the MEM instantiated with Areion-256. Areion256-OTR has a similar structure to Prøst-OTR [36]. Dobraunig *et al.* presented a related-key forgery attack on Prøst-OTR [17]. Since Areion256-OTR adopted MEM to realize the internal permutation-based tweakable block cipher, this attack cannot work against Areion256-OTR. The designers of the MEM claimed that MEM can be proven the mixed tweakable pseudorandom permutation (MTPRP) security; thus, the permutation-based OTR mode based on the MEM can be regarded to have provable security if there are no weaknesses in the underlying permutation, *i.e.*, Areion-256. In addition, OTR is inverse-free; thus, the permutation-based OTR mode using the MEM construction should have an advantage in terms of its efficiency in both encryption and decryption because Areion-256 can realize highly-efficient encryption.

The overall structure of the permutation-based OTR mode is shown in Fig. 4. The specification of Areion256-OTR is based on the original OPP and OTR modes. Namely, (1) we replace the combination of the masking part and the block cipher encryption E_K part in the original specification [52] with MEM; (2) use the same padding function as [23] for the last associated data block and the last two message blocks; and (3) use the tag generating function, LFSR update function, and input formatting (for the key and nonce) as in the case of Areion256-OPP.

For performance comparison, we also specify a Simpira v2-based OTR mode, Simpira256-

OTR, in the same way as above.

5 Security Evaluation

5.1 Security for Underlying Permutations

We evaluate the security of Areion-256 and Areion-512 as public permutations against differential, linear, impossible differential, and integral attacks.

Claimed Security for Underlying Permutations. We claim 128-bit security for both Areion-256 and Areion-512 as with Simpira v2, *i.e.*, we consider the attacks up to 2^{128} complexity. There is no rigorous definition of a distinguisher for a public permutation. In the literature, there is, however, a very related concept called the known-key distinguisher [40] or the correlation intractability [11]. Note that once the key is known for a block cipher, the block cipher becomes a public permutation. Roughly speaking, a known-key distinguisher is that for a block cipher, a relation exists such that given the key, it is easy to find plaintext-ciphertext pairs satisfying this relation. However, it is difficult to find them for a random permutation [40]. Moreover, if the relationship is simply the description of the block cipher itself, this should be meaningless for the following reasons. First, every block cipher will be vulnerable to this attack with only 1 query. Second, the relationship is not interesting at all from the designers' perspective [40]. In [21], a more formal definition of the known-key distinguisher for a block cipher was given, which is a rigorous description of the above statement. In both the known-key distinguishers on AES [40, 21], they indeed are the extensions of the well-known integral attack on round-reduced AES, where the attackers start from a middle round and aim to find an input-output set such that the sum of some bytes in the input and output are all zero, respectively. We will rely on similar start-from-the-middle techniques to construct zero-sum distinguishers for our proposed public permutations. Moreover, our zero-sum distinguishers also resemble the known-key distinguishers on AES [40, 21] because we similarly find distinguishers based on the well-known integral attack on AES.

Differential/Linear Attacks. We estimate the security against differential/linear attacks by obtaining the lower bound for the number of differentially/linearly active S-boxes with an MILP-based method proposed by Mouha *et al.* [53]. AS_D and AS_L denote the lower bound for the number of differentially and linearly active S-boxes, respectively.

Since the maximal differential and linear probability of the S-box of AES are both 2^{-6}, $AS_{D/L}$ of ≥ 22 ($2^{-6\times 22} < 2^{-128}$) is sufficient to ensure 128-bit security against differential/linear attacks. Table 7 shows the lower bound for the differentially/linearly active S-boxes for Areion-256 and Areion-512. In our evaluation, Areion-256/Areion-512 achieves both AS_D and AS_L of ≥ 22 at 4/6 rounds, and both AS_D and AS_L at 12 rounds for both permutations outnumber well over 22. Therefore, we expect full rounds of Areion-256 and Areion-512 can resist differential and linear attacks.

Impossible Differential Attacks. The miss-in-the-middle approach is known as an efficient way to find the longest impossible differences, which can be implemented by an MILP with a small change from an MILP model for counting the number of differentially active S-boxes [62, 15]. In our evaluation, we search a class of impossible differential characteristics where input and output differences activate only one byte to find the longest impossible differences efficiently.

By this approach, we find the impossible differences at 4/8 rounds of Areion-256/Areion-512, which are the longest ones we can find. Since there is still enough margin to full

Table 7: The lower bound for the number of differentially/linearly active S-boxes for Areion-256 and Areion-512. Here, AS_D and AS_L denote the number of differentially and linearly active S-boxes, respectively.

Primitives \ Rounds		1	2	3	4	5	6	7	8	9	10	11	12
Areion-256	AS_D	0	6	12	38	46	53	60	86	92	99	108	135
Areion-512		0	2	5	8	20	36	62	73	90	106	119	135
Areion-256	AS_L	0	1	9	24	42	48	65	79	91	102	114	128
Areion-512		0	1	4	8	21	35	50	68	89	103	120	131

rounds for both permutations, we expect that full rounds of Areion-256 and Areion-512 can resist impossible differential attacks.

Integral Attacks. To find the integral distinguisher, we evaluate the byte-wise division property with a MILP-based method proposed by Xiang *et al.* [68]. We search the input space where only one byte is constant, and the remaining bytes are active, *i.e.*, the data/time complexity of the integral distinguishers are 2^{248} and 2^{504} for Areion-256 and Areion-512, respectively.

As a result, we find the 3- and 5-round integral distinguisher on Areion-256 and Areion-512, respectively. It should be emphasized that the required data/time complexities for these distinguishers exceed our security claim. Hence, the longest integral distinguishers with up to 2^{128} data/time complexity, which are in our security claim, are expected to exist on fewer rounds than that of these distinguishers. Thus, we expect full rounds of Areion-256 and Areion-512 can resist integral attacks.

Zero-sum Distinguishers. The zero-sum distinguisher [5] is a popular attack on public permutations. The overall attack procedure is straightforward. Specifically, the attackers first choose a particular set of intermediate state values and then propagate this set of values backward and forwards, respectively. If, in the corresponding set of inputs and outputs, the sum of some input bits and output bits are zero, respectively, a zero-sum distinguisher is found. We have evaluated the resistance against this attack based on the well-known 4-round integral distinguisher for AES. It is found that there are zero-sum distinguishers for 5-round Areion-256 and 10-round Areion-512, respectively. The data and time complexities of the two zero-sum distinguishers are the same, which are both 2^{32}. We give the details below.

The distinguisher for 5-round Areion-256. First, we explain the zero-sum distinguisher for 5-round Areion-256, as shown in Fig. 5.

Specifically, we choose 4 bytes of x_1^2 which traverses all the 2^{32} possible values. For x_0^2, it is assigned to a random constant value. According to the round function, we have

$$\begin{aligned}
x_0^4 &= G_1 \circ G_0(x_0^3) \oplus x_1^3 = G_1 \circ G_0(G_1 \circ G_0(x_0^2) \oplus x_1^2) \oplus G_0(x_0^2),\\
x_1^4 &= G_0(x_0^3) = G_0(G_1 \circ G_0(x_0^2) \oplus x_1^2),\\
x_0^5 &= G_1 \circ G_0(x_0^4) \oplus x_1^4,\\
x_1^5 &= G_0(x_0^4),
\end{aligned}$$

For the term $G_1 \circ G_0(x_0^4)$ in x_0^5, with our input form for (x_0^2, x_1^2), it is equivalent to that x_1^2 passes 4 AES rounds. The term x_1^4 in x_0^5 is equivalent to that x_1^2 passes 2 AES rounds. Hence, we need to use a data set of size 2^{32}, and all the bytes in x_0^5 will be balanced.

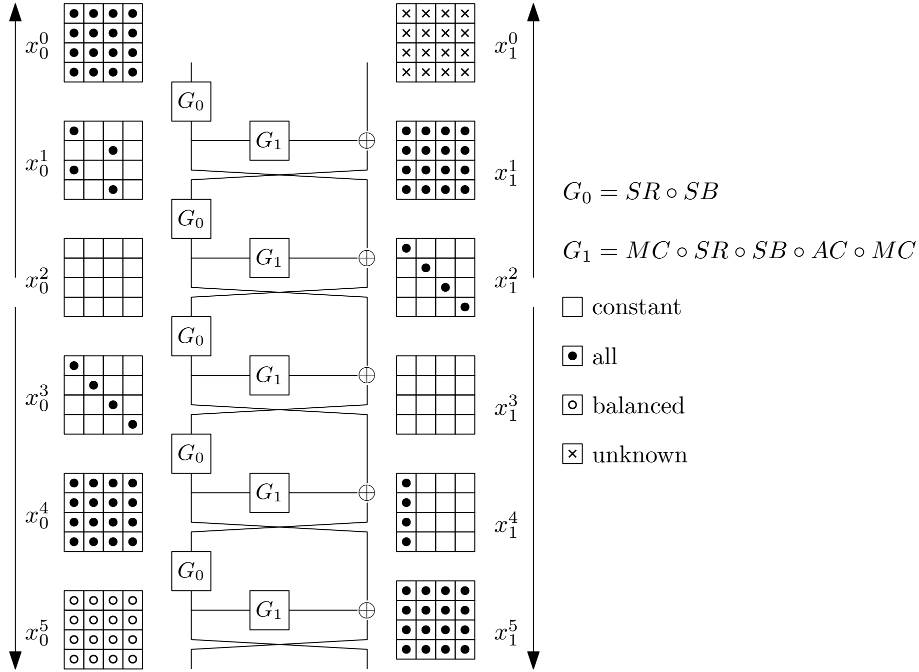

Figure 5: The zero-sum distinguisher for 5-round Areion-256.

For x_1^5, as x_0^4 can be viewed as applying 2 AES rounds to x_1^2, each byte of x_1^5 will also be balanced. The above observation also explains why the automatic method based on the division property could only detect a 3-round integral distinguisher in the forward direction, *i.e.*, we at least need to consider 4 AES rounds.

In the backward direction, we have

$$\begin{aligned} x_0^1 &= G_0^{-1}(x_1^2), \\ x_1^1 &= G_0 \circ G_1(x_0^1) \oplus x_0^2, \\ x_0^0 &= G_0^{-1}(x_1^1). \end{aligned}$$

Therefore, all the bytes in x_0^0 will be balanced. To better understand this, one can first consider the case when only one byte of x_1^2 traverses all the 2^8 possible values. For such a case, it can be easily checked that each byte in x_0^0 will also traverse all the 2^8 possible values. Hence, if one diagonal of x_1^2 takes all the possible 2^{32} values, all bytes in x_0^0 are also balanced.

The distinguisher for 10-round Areion-512. Next, we explain the zero-sum distinguisher for 10-round Areion-512, as shown in Fig. 6. We start from the state $(x_0^4, x_1^4, x_2^4, x_3^4)$ after 4 rounds of permutation. For the input form, we restrict that 4 bytes of x_0^4 will traverse all the 2^{32} possible values, as shown in Fig. 6. Then, we randomly choose a 128-bit constant $\mathcal{C}$ such that $F_0(x_0^4) \oplus x_1^4 = \mathcal{C}$ always holds. In other words, the value of x_1^4 is conditioned, and it is dynamically chosen according to x_0^4. For x_2^4, we assign a random constant value to it. For x_3^4, we also assign a random constant value $\mathcal{C}'$ to it but we require that the first column of $x_0^3 = F_1^{-1}(\mathcal{C}')$ is all 0. Note that (F_0, F_1, F_2, F_3) are defined in Sect. 2.

For such an input state, in the forward direction, we can trivially deduce that (x_0^6, x_1^6, x_3^6) are constants and one diagonal of x_2^6 will take all the 2^{32} possible values. Therefore, we can also deduce that (x_0^7, x_3^7, x_3^8) are all constants.

Since

$$\begin{aligned} x_2^{10} &= F_0(x_2^9) \oplus x_3^9, \\ x_2^9 &= F_0(x_2^8) \oplus x_3^8, \\ x_2^8 &= F_0(x_2^7) \oplus x_3^7, \\ x_3^9 &= F_1(x_0^8) = F_1(F_0(x_0^7) \oplus x_1^7), \end{aligned}$$

we can rewrite x_2^{10} as follows where $\mathcal{C}_i$ are 128-bit constants:

$$\begin{aligned} x_2^{10} &= F_0(F_0(F_0(x_2^7) \oplus \mathcal{C}_0) \oplus \mathcal{C}_1) \oplus F_1(x_1^7 \oplus \mathcal{C}_2) \\ &= F_0(F_0(F_0(F_0(x_2^6) \oplus \mathcal{C}_3) \oplus \mathcal{C}_0) \oplus \mathcal{C}_1) \oplus F_1(F_3(x_2^6) \oplus \mathcal{C}_2). \end{aligned}$$

Since $F_0 = MC \circ SR \circ SB$, $F_1 = SR \circ SB$ and $F_3 = MC \circ SR \circ SB \circ AC \circ SR \circ SB$, the term $F_0(F_0(F_0(F_0(x_2^6) \oplus \mathcal{C}_3) \oplus \mathcal{C}_0) \oplus \mathcal{C}_1)$ is equivalent to applying 4 AES rounds to x_2^6. The term $F_1(F_3(x_2^6) \oplus \mathcal{C}_2)$ is equivalent to applying 2.5 AES rounds to x_2^6. The above observation implies that we need to use a data set of size 2^{32} to detect an integral property at x_2^{10}. Since one diagonal of x_2^6 takes all the 2^{32} possible values, each byte in x_2^{10} is balanced. For $(x_0^{10}, x_1^{10}, x_3^{10})$, we will lose the zero-sum property, and this can be deduced similarly. In other words, we can obtain a 6-round integral distinguisher with data complexity 2^{32} in the forward direction, which is one more round than the result obtained with the automatic method based on the division property. The main reason is that we dynamically choose values for x_1^4 such that $F_0(x_0^4) \oplus x_1^4$ is always a constant when x_0^4 varies.

In the backward direction, we consider a subset of $(x_0^4, x_1^4, x_2^4, x_3^4)$. Specifically, we consider the case when the first byte x_0^4 takes all the 2^8 possible values. In this case, the value of the first column of x_1^4 is dynamically chosen such that $F_0(x_0^4) \oplus x_1^4$ is a constant $\mathcal{C}$, as shown in Fig. 6. Then, we have 2^{24} such subsets in total.

Since

$$\begin{aligned} x_1^4 &= F_3(x_2^3) = F_0 \circ AC \circ SR \circ SB(x_2^3), \\ \mathcal{C} &= F_0(x_0^4) \oplus x_1^4, \end{aligned}$$

we have

$$AC \circ SR \circ SB(x_2^3) = F_0^{-1}(F_0(x_0^4) \oplus \mathcal{C}).$$

Since $F_0 = MC \circ SR \circ SB$, the above formula implies that the first byte of $AC \circ SR \circ SB(x_2^3)$ will traverse all the 2^8 possible values. Hence, only the first byte of x_2^3 will traverse all the 2^8 possible values. Therefore, we obtain the form of $(x_0^3, x_1^3, x_2^3, x_3^3)$ shown in Fig. 6.

Deducing $(x_0^2, x_1^2, x_2^2, x_3^2)$ from $(x_0^3, x_1^3, x_2^3, x_3^3)$ is trivial and we omit the details. Deducing (x_0^1, x_1^1) is also trivial based on $(x_0^2, x_1^2, x_2^2, x_3^2)$. Next, we mainly focus on (x_2^1, x_3^1). Similarly, we have

$$\begin{aligned} x_1^2 &= F_3(x_2^1) = F_0 \circ AC \circ SR \circ SB(x_2^1), \\ x_0^3 &= F_0(x_0^2) \oplus x_1^2, \\ AC \circ SR \circ SB(x_2^1) &= F_0^{-1}(F_0(x_0^2) \oplus x_0^3). \end{aligned}$$

As the first column of x_0^3 is zero, the first diagonal of x_2^1 will equal the first diagonal of $AC \circ SR \circ SB(x_2^1)$. Hence, we obtain the form of (x_2^1, x_3^1) as shown in Fig. 6. Based on $(x_0^1, x_1^1, x_2^1, x_3^1)$, deducing $(x_0^0, x_1^0, x_2^0, x_3^0)$ is trivial and we omit the details.

In a word, we can construct a zero-sum distinguisher for 10-round Areion-512.

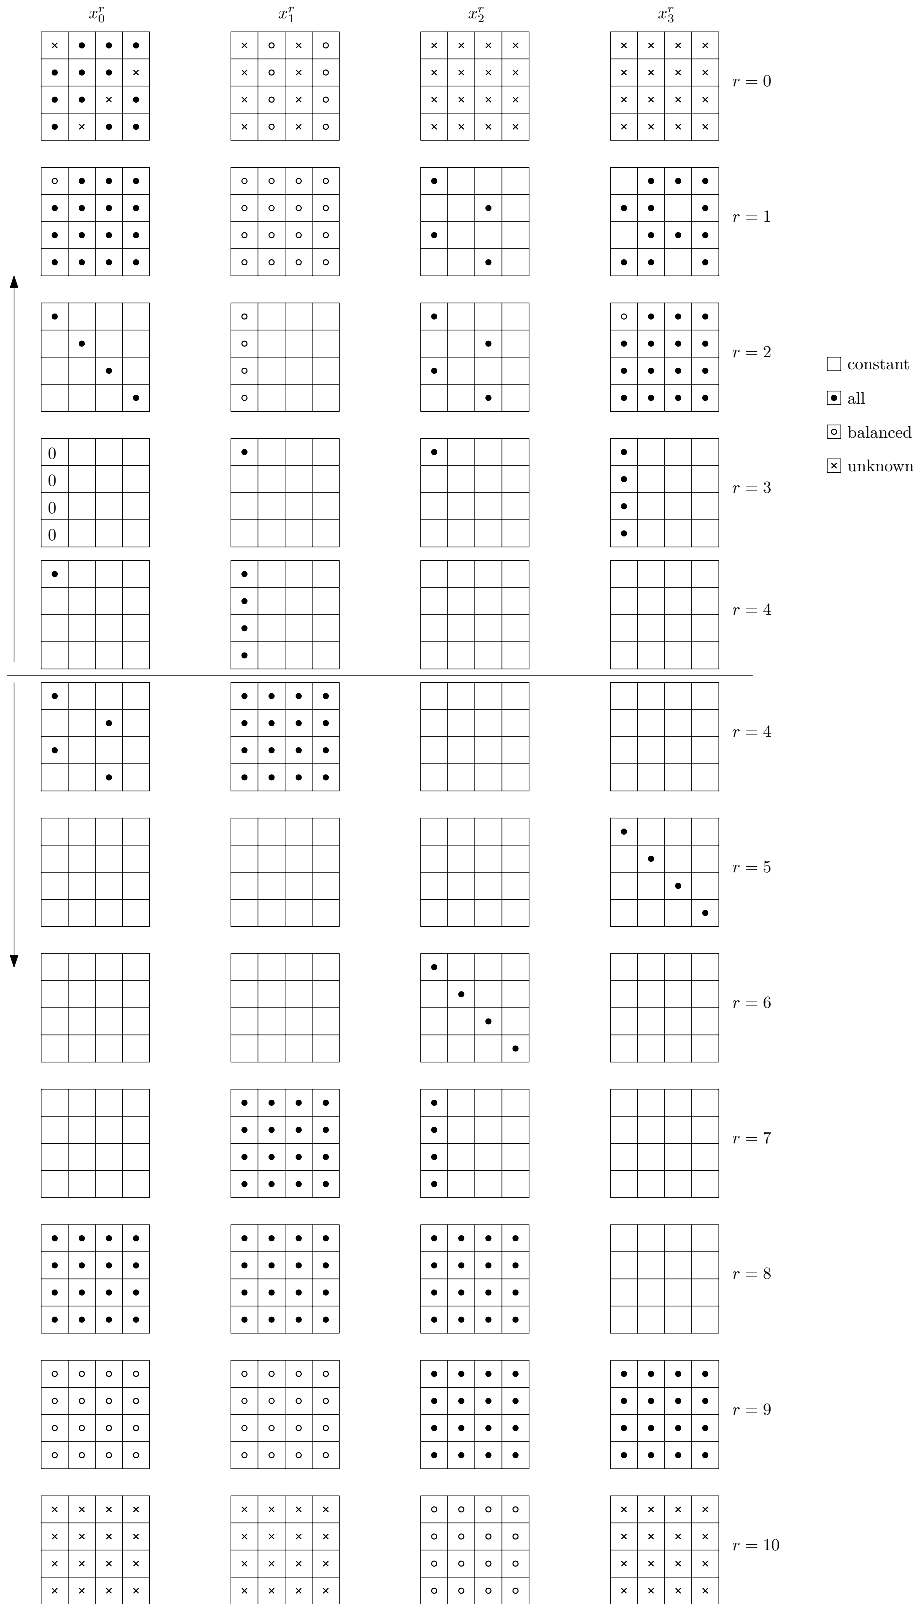

Figure 6: The zero-sum distinguisher for 10-round Areion-512.

5.2 Security for Hash Functions

Claimed Security for Hash Functions. We claim 256-bit security against the preimage attack for both Areion256-DM and Areion512-DM. However, as in Haraka v2 and the SFIL hash function built on Simpira v2, we do not claim their resistances against the collision attack since it is unnecessary for their applications.

For the MD-based hash function, we claim 256-bit security against the preimage attack and 128-bit security against the collision attack, the same as SHA2-256. Due to the generic second-preimage attack on the MD construction [38], our MD-based hash scheme could only provide about 193-bit security for second-preimage attacks. This limitation is because the maximal number of allowed message blocks is 2^{64} and $193 = 256 - 64 + 1$, which is the same security level as SHA2-256.

Meet-in-the-Middle Preimage Attack. For the DM-based SFIL hash functions by using Areion-256 and Areion-512 as the underlying permutations, respectively, it is necessary to take into account Sasaki's meet-in-the-middle (MITM) preimage attack [61]. This attack is the most powerful preimage attack on such hash functions. Indeed, the designers of Haraka v2 have evaluated its resistance against this attack in a dedicated way. We also performed a careful analysis to understand the security of our constructions better. We found preimage attacks on 5-round Areion256-DM and 10-round Areion512-DM, respectively. Therefore, there is still a sufficiently large security margin. We detail our analysis below.

To save space, we only describe the general procedure of Sasaki's meet-in-the-middle preimage attack, as shown below:

Step 1: Identify the bytes fixed to constants and assign proper values to them.

Step 2: Identify the bytes that are to be exhausted. Classify them into backward neutral bytes and forward neutral bytes.

Step 3: In the forward direction, we assume that the backward neutral bytes are unknown and compute the internal state values based on the constant bytes and the forward neutral bytes. In other words, we only compute the bytes that can be computed from the knowledge of the constant bytes and the forward neutral bytes. This step is repeated for all the possible values of the forward neutral bytes, and we store the corresponding computed information.

Step 4: In the backward direction, we assume that the forward neutral bytes and unknown, and we only compute the bytes that can be computed from the knowledge of the constant bytes and the backward neutral bytes. This step is repeated for all the possible values of the backward neutral bytes, and we store the corresponding computed information[5].

Step 5: Find matches between the store information obtained at Step 3 and Step 4. Suppose the matching probability is 2^{-p} and there are 2^{b_f} and 2^{b_b} possible values for the forward neutral bytes and the backward neutral bytes, respectively. Moreover, for each obtained state information at Step 3, if it is possible to identify the matched information obtained at Step 4 with time complexity 1, or vice versa, we can say that we find $2^{b_f+b_b-p}$ possible pairs among the $2^{b_f+b_b}$ pairs with time complexity $max(2^{b_f}, 2^{b_b})$ where usually $b_f + b_b - p \leq 0$. In other words, we exhaust $2^{b_f+b_b}$ possible candidates only with time complexity $max(2^{b_f}, 2^{b_b})$. Hence, the MITM preimage attack is $min(2^{b_b}, 2^{b_f})$ times faster than the brute force.

[5]Note that in the actual implementations, we only need to store either the information obtained at Step 3 or Step 4. For simple explanations, we assume both are stored.

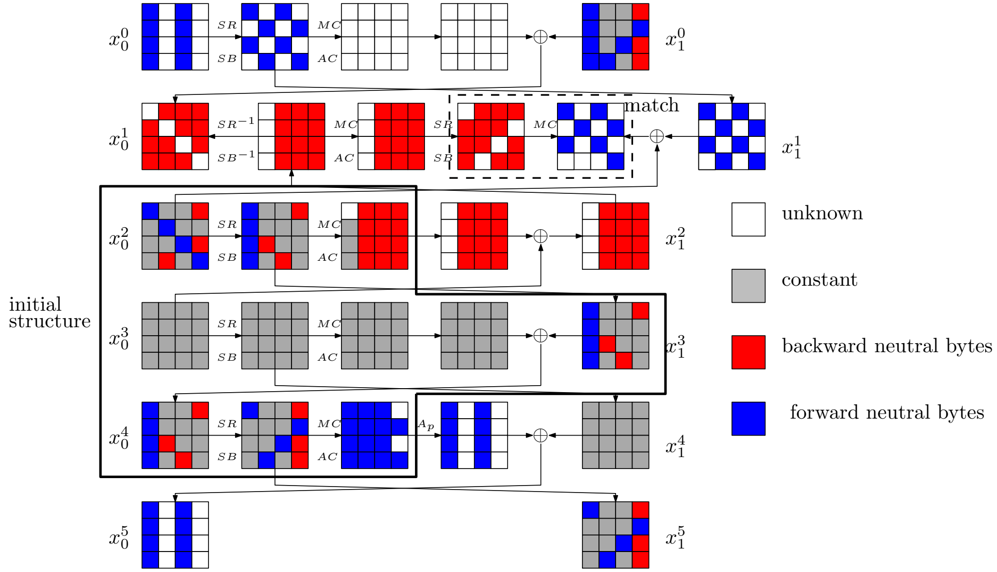

Figure 7: The preimage attack on 5-round Areion256-DM.

Hence, this attack aims to identify the forward and backward neutral bytes as well as an efficient matching method. We performed careful analysis for the two short-input hash schemes and found preimage attacks on 5-round Areion256-DM and 10-round Areion512-DM, respectively. In the two attacks, $b_b = b_f = 8$ and the matching phase can be efficiently finished with time complexity 1. Hence, both the preimage attacks are 2^8 times faster than the brute force. The corresponding illustration of the two preimage attacks can be referred to Figs. 7 and 8, respectively.

Collision Attacks. The most powerful collision attack on AES-based hash functions is the rebound attack [48], especially when built on the DM construction, as the attacker can fully control the whole internal state. However, as already mentioned in Haraka v2 and the SFIL hash function based on Simpira v2, the collision resistance of SFIL hash schemes is not necessary when they are used in the signature scheme, which is also the case of our SFIL hash functions.

Security of MD Construction. For our hash scheme built on the MD construction, the attacker will soon lose the capability to fully control the internal state since each message block is only 256 bits, *i.e.*, half of the state size. However, by using $j > 1$ message blocks, Sasaki's MITM attack can still be applied in the same way as in the attack on the DM constructions. Specifically, although the 256-bit initial value set at (x_2, x_3) in the first input state is fixed, the attackers can view the 256-bit chaining variable (CV) in the last input state as a controllable part. Then, Sasaki's MITM attack is applied, and we aim to find 2^i solutions of the last input state to match the given hash value in less than 2^{256} time. This way, 2^i candidates of CV in the last input state can be obtained. Finally, we randomly pick values for the first $j-1$ message blocks to compute the corresponding CV for the last input state and expect one such CV to match one of the 2^i candidates obtained by the MITM attack. Hence, we need to try 2^{256-i} different values for the first $j-1$ message blocks, and the time complexity is below 2^{256}.

For the collision resistance, we consider the rebound attack, the most efficient technique for AES-based hash functions. In particular, the most powerful rebound attack is always based on the Super-Sbox technique [22, 43]. For such a technique, the attacker can control

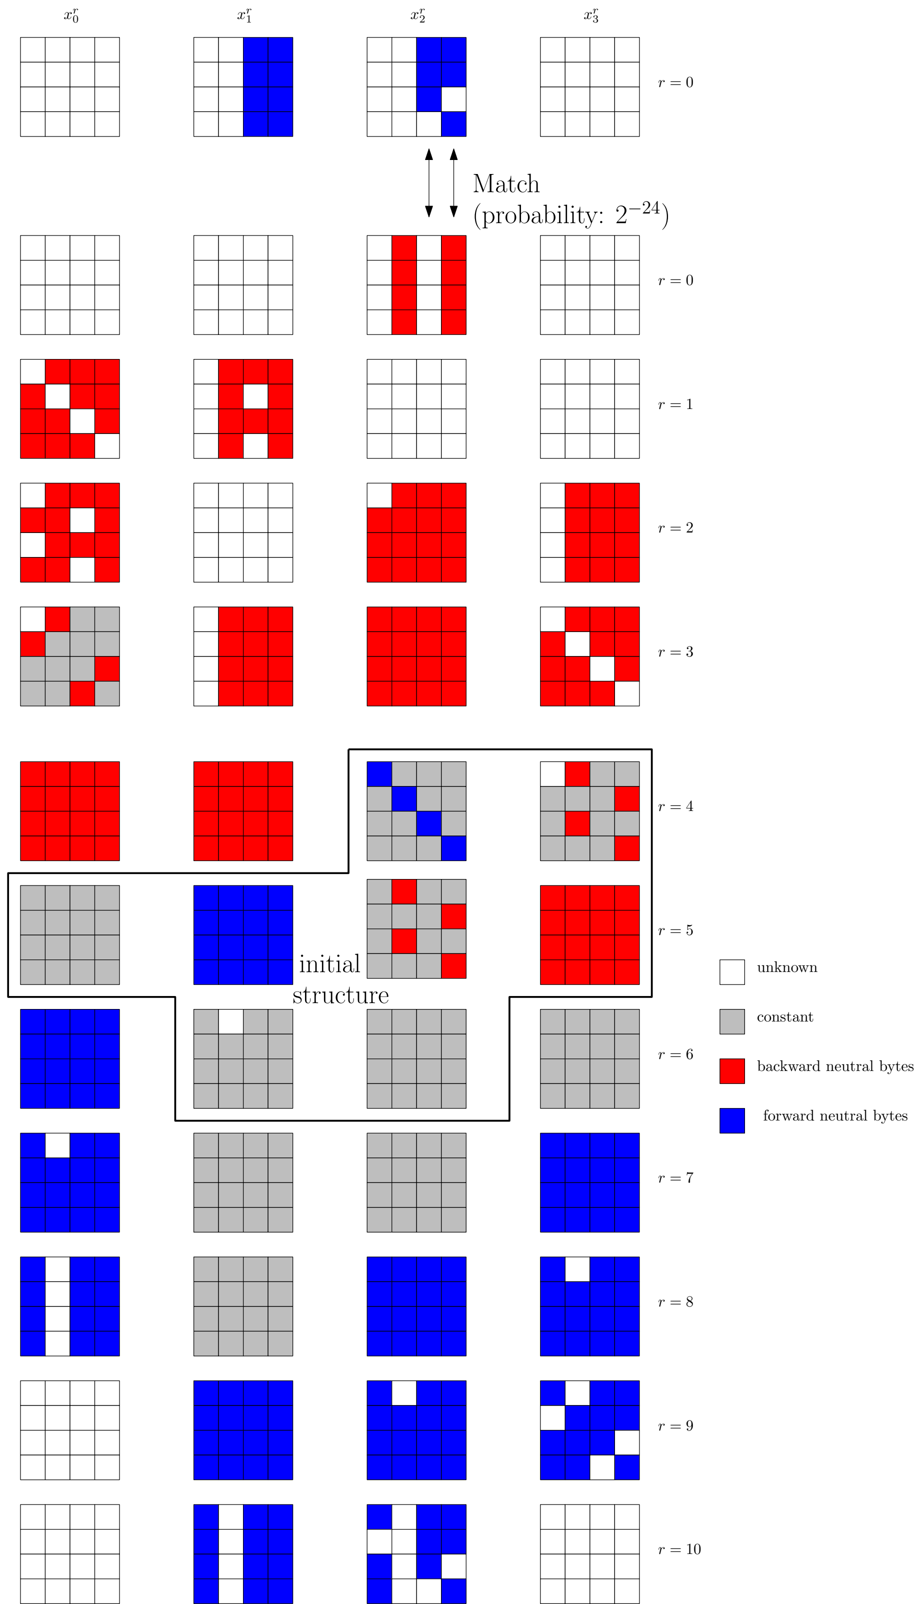

Figure 8: The preimage attack on 10-round Areion512-DM.

the difference transitions over two consecutive AES rounds with a pre-computation phase called the inbound phase, as shown in Fig. 9. Combined with the feature of the rebound

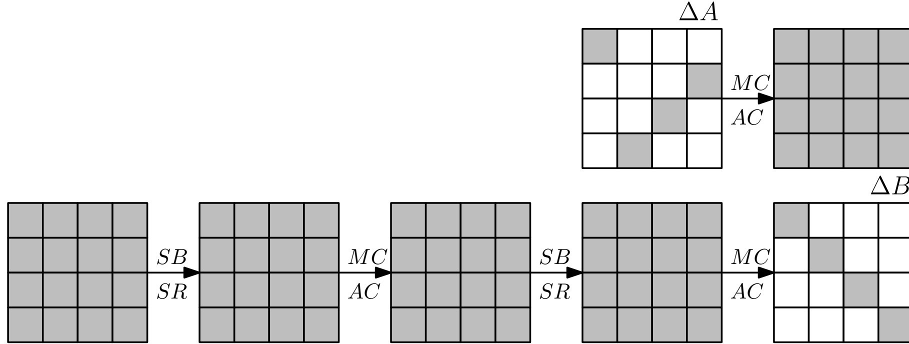

Figure 9: The inbound phase: precomputing the pairs $(A, \Delta A)$ such that $\Delta A \rightarrow \Delta B$ holds with probability 1.

attack, this technique allows the attacker to ignore the influence of $4 + 16 + 16 + 4 = 38$ active S-boxes by using 128 free bits. Since the size of one message block in our VIL hash function is 256 bits, we expect that the attacker can ignore $38 \times 2 = 76$ active S-boxes with the Super-Sbox technique. However, we emphasize that it does not necessarily imply that the attacker can always ignore 76 active S-boxes in the actual attack because the rebound attack is also a start-from-the-middle-style attack, and one should be careful of the consistency in the CV.

According to Table 7, the minimal number of active S-boxes in 11-round Areion-512 is 119. By ignoring 76 active S-boxes, there are still $119 - 76 = 43$ active S-boxes left. In the outbound phase, we usually need to cancel the truncated differences. In the best case, we only need to consider half of the left active S-boxes, *i.e.*, we know the propagation of the truncated differences, and we only add conditions on the sum of the two truncated differences, as shown in Fig. 10. Even if we only consider $43/2 \approx 21$ active S-boxes, they still correspond to a very low uncontrolled probability of $2^{-21 \times 8} = 2^{-168}$. Note that we have not yet taken into account the extra conditions on the truncated input and output differences to generate a collision. If they are considered, the truncated differential may be worse (*i.e.*, there are more active S-boxes), and the uncontrolled probability may further decrease. These analyses suggest that the VIL hash function based on the 15-round Areion-512 is secure against the collision attack.

We also note that there is a variant method [31] of the 2-round Super-Sbox technique that can cover three consecutive AES rounds, which can allow the attackers to ignore

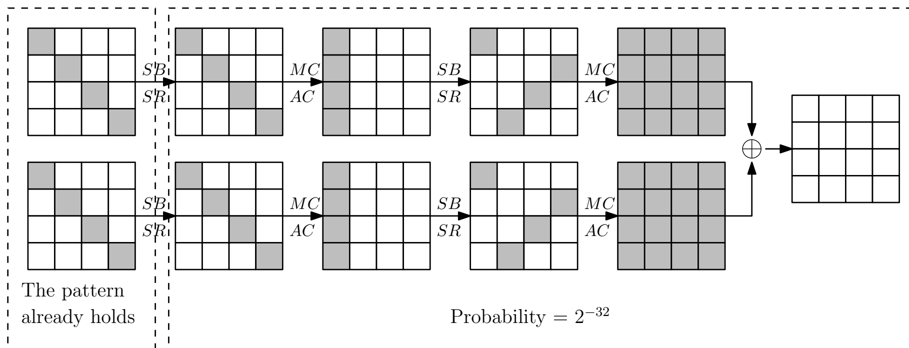

Figure 10: Cancel the truncated differences.

the influence of $4 + 16 + 16 + 16 + 4 = 54$ active S-boxes. However, this technique does not come for free. Specifically, different from the 2-round Super-Sbox technique to satisfy $4 + 16 + 16 + 4 = 38$ active S-boxes where lots of degrees of freedom are left after this phase, there is no degrees of freedom left after performing such a 3-round Super-Sbox technique and finding a solution to satisfy these 54 active S-boxes succeeds with probability 2^{-64}. In other words, it is like the 2-round Super-Sbox technique with satisfying extra 16 active S-boxes with a probability of only 2^{-64}, which is a considerable improvement over the 2-round Super-Sbox technique. We also note that it is almost equivalent to our conservative estimation that we only need to consider half of the remaining active S-boxes at the outbound phase when using the 2-round Super-Sbox technique for the inbound phase.

6 Performance Evaluation

In this section, we evaluate the performance of both Areion and its applications to the permutation-based hash functions and AEAD schemes described in Sect. 4. To this end, we used the available source code at GitHub[6] to evaluate the cycle counters, *i.e.*, cycles per byte (cpb), in the target primitive. All our evaluations were performed on the following widely deployed platforms: the Ice Lake, Tiger Lake, and Alder Lake platforms. The Ice Lake platform has an Intel(R) Core(TM) i7-1068NG7 CPU @ 2.30GHz. The Tiger Lake platform has an Intel(R) Core(TM) i7-1165G7 CPU @ 2.80GHz. The Alder Lake platform has an Intel(R) Core(TM) i9-12900K CPU @ 3.20GHz on a performance-core (P-core) and 2.40GHz on an efficient-core (E-core). Turbo Boost technology has been switched off for all our evaluations. We note here that the P-core has been specified for our evaluations on the Alder Lake platform because there is almost no difference in the benchmarks between using either the P-core or E-core.

Besides, we also evaluate the performance of NEON implementations of permutation-based hash functions proposed in Sect. 4 in several mobile environments. The NEON implementations of Areion is shown in Appendix A.3.

6.1 Underlying Permutations

We first evaluate the performance of the underlying permutations, *i.e.*, Areion-256 and Areion-512. These implementations are given in Appendix A.1. For comparison, we used the underlying permutations of Simpira v2, Haraka v2, and the 512-bit permutation BLAKE2s. We can find the source codes of Haraka v2 and BLAKE2s available at GitHub[7,8], but we could not find the available source code for Simpira v2. For this reason, we implemented it as described in Appendix A.2.

According to [25, 41], Simpira v2 and Haraka v2 are supposed to operate on multiple message blocks, not just a single message block, to get the highest performance. Based on this concept, we also evaluate the performance when operating on eight message blocks in parallel and a single message block.

Tables 8 and 9 show benchmarks for single and parallel encryption/decryption on our platforms. From these tables, Haraka v2 appears to be the fastest encryption, but it cannot be regarded to have a security margin sufficiently, as discussed in Sect. 3.2.1. For this reason, we consider there is no problem even if Haraka v2 is excluded from our comparison. Instead of the original Haraka v2, we select the 12/15-round variants of Haraka v2, Haraka-256 (x1.2/x1.5) and Haraka-512 (x1.2/x1.5), for our comparison. This selection is because DM-based instantiations of the tweaked variants, Haraka256-DM (x1.2/x1.5) and

[6] https://github.com/seb-m/cycles
[7] https://github.com/kste/haraka
[8] https://github.com/BLAKE2/BLAKE2

Table 8: Benchmarks for single block encryption/decryption on the Ice Lake, Tiger Lake, and Alder Lake platforms. All values are given as cpb.

	Ice Lake		Tiger Lake		Alder Lake	
Primitive	Enc	Dec	Enc	Dec	Enc	Dec
Areion-256	**1.92**	**2.84**	**1.91**	**2.83**	**1.93**	**2.81**
Simpira-256	2.94	2.94	2.92	2.92	2.94	2.94
Haraka-256	1.58	4.08	1.58	4.08	1.55	4.00
Haraka-256 (x1.2)	1.90	5.28	1.90	5.28	1.86	4.80
Haraka-256 (x1.5)	2.37	6.12	2.37	6.12	2.32	6.00
Areion-512	**1.09**	**2.52**	**1.09**	**2.52**	**1.09**	**2.52**
Simpira-512	1.47	1.47	1.46	1.46	1.47	1.47
Haraka-512	1.06	2.58	1.06	2.58	1.09	2.58
Haraka-512 (x1.2)	1.27	3.10	1.27	3.10	1.31	3.10
Haraka-512 (x1.5)	1.59	3.87	1.59	3.87	1.63	3.87

Haraka512-DM (x1.2/x1.5), can be regarded to have a similar security level as Areion256-DM and Areion512-DM. Indeed, the security margins against MITM preimage attacks of Areion256-DM, Haraka256-DM (x1.2), and Haraka256-DM (x1.5) are 5, 3, and 6, respectively. Similarly, the security margins of Areion512-DM, Haraka512-DM (x1.2), and Haraka512-DM (x1.5) are 5, 1, and 4, respectively. We summarize the performance comparison for the underlying permutations as follows:

- Areion-256 realizes the fastest encryption among the target permutations, excluding Haraka-256 (x1.2) for single block encryption (although there are almost no differences in performance). Specifically, Areion-256 performs at least 1.52 and 1.20 times faster than Simpira-256 and Haraka-256 (x1.5) for single block encryption, respectively, and at least 1.12 and 1.03 times faster than Simpira-256 and Haraka-256 (x1.2) for parallel block encryption, respectively. On the other hand, for single and parallel block decryptions, Areion-256 performs faster than Haraka-256 (x1.2/x1.5), but there are almost no differences in performance between Areion-256 and Simpira-256.

- Areion-512 realizes the fastest encryption among the target permutations, excluding Simpira-512 for parallel block encryption (although there are almost no differences in performance). Specifically, Areion-512 performs at least 1.34 and 1.16 times faster than Simpira-512 and Haraka-512 (x1.2) for single block encryption, respectively, and at least 1.26 times faster than Haraka-512 (x1.2) for parallel block encryption. On the other hand, Areion-512 performs faster than Haraka-256 (x1.2/x1.5) and BLAKE2s, especially for parallel block decryption, but it performs at least 2.00 times slower than Simpira-512.

Given that the Areion-512 decryption function is not used for the proposed applications of Areion described in Sect. 4, we consider that there is no problem even if Areion-512 performs slower than Simpira-512 for decryption. Therefore, Areion has the strongest advantage of performing faster than any other target permutations, especially in terms of encryption direction.

Regarding the advantage of Areion-256 over Areion-512, Table 9 suggests that Areion-256 is consistently faster than Areion-512 for parallel processing and even the fastest among all the selected 256-/512-bit permutations in many cases. In addition, it has a balanced performance for encryption and decryption thanks to its Feistel-like structure, unlike Haraka-256, and faster than the Feistel-based Simpira-256. That is, it should work more efficiently with the existing parallelizable permutation-based authenticated encryption modes, *e.g.*, OPP [23] and a permutation-based counterpart of OTR [52] than other permutations. The latter would be similar to Prøst-OTR [35] adopting the masking scheme

Table 9: Benchmarks for parallel block encryption/decryption on the Ice Lake, Tiger Lake, and Alder Lake platforms. All values are given as cpb.

	Ice Lake		Tiger Lake		Alder Lake	
Primitive	Enc	Dec	Enc	Dec	Enc	Dec
Areion-256	**0.55**	**0.66**	**0.55**	**0.66**	**0.51**	**0.56**
Simpira-256	0.69	0.69	0.68	0.68	0.57	0.57
Haraka-256	0.53	1.75	0.54	1.74	0.44	1.52
Haraka-256 (x1.2)	0.64	2.10	0.65	2.09	0.53	1.83
Haraka-256 (x1.5)	0.79	2.62	0.81	2.61	0.66	2.28
Areion-512	**0.64**	**1.24**	**0.63**	**1.25**	**0.61**	**1.13**
Simpira-512	0.63	0.62	0.62	0.61	0.53	0.53
Haraka-512	0.67	2.06	0.66	2.04	0.64	1.83
Haraka-512 (x1.2)	0.81	2.48	0.80	2.45	0.77	2.20
Haraka-512 (x1.5)	1.00	3.09	0.99	3.06	0.96	2.74

of OPP for provable security and for avoiding the attack specific to (the masking scheme of) Prøst-OTR [17]. Its applications to the parallel authenticated encryption modes are left as our future work. On the other hand, Areion-512 is the fastest among the selected 256-/512-bit permutations for single block encryption direction (Table 8). That is, it should work more efficiently with the existing permutation-based compression functions, such as DM construction, and the existing sequential hash functions, such as MD construction. These are the target applications for this study.

6.2 Permutation-based Hash Functions

Next, we evaluate the performance of the permutation-based hash functions, *i.e.*, the SFIL and VIL hash functions (DM and MD constructions). These instantiations of Areion are implemented based on the source codes of Areion-256 and Areion-512 described in Appendix A.1. For comparison regarding the SFIL hash functions, we used DM constructions instantiated with Simpira v2 and Haraka v2. On the other hand, for comparison regarding the VIL hash functions, we used AES-based VIL hash functions, such as Simpira512-MD, Haraka512-MD, and double-block-length hash functions proposed by Hirose at FSE 2006 [26]. We refer to Hirose's hash function as Hirose-DBL. These instantiations are also implemented similarly to those of Areion. In addition, we used SHA2-256, SHA3-256, ParallelHash256, KangarooTwelve, and BLAKE3. We can find these source codes available at SUPERCOP[9] and GitHub[10,11,12]; then, we modified these source codes to use for our comparison.

Tables 10 and 11 show benchmarks for the SFIL and VIL hash functions on our platforms. From Table 10, Haraka512-DM appears to be the fastest SFIL hash function, but Haraka v2 cannot be regarded to have the security margin sufficiently; thus, we use Haraka256-DM (x1.2/x1.5) and Haraka512-DM (x1.2/x1.5) for our comparison regarding the SFIL hash functions, as discussed in Sect. 6.1. Similarly, we use Haraka512-MD (x1.2/x1.5) to compare the VIL hash functions. We summarize the performance comparison for the SFIL hash functions as follows:

- Areion256-DM realizes the fastest SFIL hashing among the target DM constructions with the 256-bit permutation, excluding Haraka256-DM (x1.2) (although there are almost no differences in performance). Specifically, Areion256-DM performs at least 1.41 and 1.21 times faster than Simpira256-DM and Haraka256-DM (x1.5), respectively.

[9] https://bench.cr.yp.to/supercop.html
[10] https://github.com/wereHamster/sha256-sse
[11] https://github.com/XKCP/XKCP
[12] https://github.com/BLAKE3-team/BLAKE3

Table 10: Benchmarks for SFIL hash functions on the Ice Lake, Tiger Lake, and Alder Lake platforms. All values are given as cpb.

Primitive	Ice Lake	Tiger Lake	Alder Lake
Areion256-DM	**2.01**	**2.01**	**1.99**
Simpira256-DM	2.84	2.83	2.81
Haraka256-DM	1.64	1.63	1.61
Haraka256-DM (x1.2)	1.97	1.96	1.94
Haraka256-DM (x1.5)	2.46	2.44	2.41
Areion512-DM	**1.05**	**1.05**	**1.04**
Simpira512-DM	1.41	1.41	1.40
Haraka512-DM	1.12	1.10	1.13
Haraka512-DM (x1.2)	1.35	1.32	1.36
Haraka512-DM (x1.5)	1.68	1.65	1.69

Table 11: Benchmarks for VIL hash functions on the Ice Lake, Tiger Lake, and Alder Lake platforms. All values are given as cpb.

Platform	Primitive	Security level†	Impl.	Input sizes (bytes) 64	128	256	512	1024	2048	4096
Ice Lake	**Areion512-MD**	256	AES-NI	**1.99**	**2.35**	**2.54**	**2.60**	**2.63**	**2.65**	**2.66**
	Simpira512-MD	256	AES-NI	2.03	2.65	3.04	3.25	3.34	3.38	3.40
	Haraka512-MD	256	AES-NI	2.12	2.47	2.57	2.66	2.69	2.71	2.72
	Haraka512-MD (x1.2)	256	AES-NI	2.55	2.96	3.09	3.19	3.23	3.25	3.26
	Haraka512-MD (x1.5)	256	AES-NI	3.18	3.70	3.86	3.99	4.03	4.07	4.08
	Hirose-DBL	256	AES-NI	12.29	12.21	12.21	12.34	12.27	12.28	12.27
	SHA256	256	AVX2	5.81	4.27	3.48	3.09	2.90	2.79	2.75
	ParallelHash256	256	AVX2	61.09	30.65	19.88	14.43	11.71	10.33	9.37
	ParallelHash256	256	AVX512	42.61	21.39	13.89	9.90	7.91	6.92	6.24
	BLAKE3	128	SSE	4.70	3.63	3.22	3.02	2.92	1.53	0.84
	KangarooTwelve	128	AVX2	11.59	5.96	5.27	4.86	4.15	3.79	3.61
	KangarooTwelve	128	AVX512	8.16	4.17	3.78	3.34	2.79	2.52	2.38
Tiger Lake	**Areion512-MD**	256	AES-NI	**1.89**	**2.31**	**2.50**	**2.57**	**2.61**	**2.63**	**2.64**
	Simpira512-MD	256	AES-NI	1.98	2.60	3.01	3.22	3.32	3.37	3.40
	Haraka512-MD	256	AES-NI	2.09	2.44	2.57	2.64	2.68	2.71	2.72
	Haraka512-MD (x1.2)	256	AES-NI	2.51	2.92	3.08	3.16	3.22	3.26	3.27
	Haraka512-MD (x1.5)	256	AES-NI	3.14	3.65	3.85	3.95	4.02	4.07	4.08
	Hirose-DBL	256	AES-NI	12.30	12.21	12.24	12.34	12.24	12.24	12.24
	SHA256	256	AVX2	5.48	3.89	3.11	2.72	2.53	2.43	2.38
	ParallelHash256	256	AVX2	61.04	30.62	19.86	14.40	11.67	10.35	9.36
	ParallelHash256	256	AVX512	42.40	21.77	13.83	9.98	7.86	6.87	6.20
	BLAKE3	128	SSE	5.03	3.76	3.28	3.05	2.94	1.53	0.84
	KangarooTwelve	128	AVX2	11.60	5.93	5.30	4.87	4.16	3.78	3.61
	KangarooTwelve	128	AVX512	8.22	4.20	3.78	3.33	2.78	2.50	2.36
Alder Lake‡	**Areion512-MD**	256	AES-NI	**1.60**	**2.16**	**2.42**	**2.60**	**2.66**	**2.68**	**2.70**
	Simpira512-MD	256	AES-NI	1.65	2.30	2.87	3.19	3.32	3.39	3.42
	Haraka512-MD	256	AES-NI	1.68	2.15	2.41	2.55	2.62	2.65	2.67
	Haraka512-MD (x1.2)	256	AES-NI	2.02	2.58	2.90	3.05	3.14	3.18	3.21
	Haraka512-MD (x1.5)	256	AES-NI	2.52	3.23	3.62	3.82	3.92	3.97	4.01
	Hirose-DBL	256	AES-NI	12.67	12.61	12.58	12.59	12.61	12.61	12.61
	SHA256	256	AVX2	4.45	3.26	2.63	2.35	2.20	2.12	2.06
	ParallelHash256	256	AVX2	59.38	29.47	19.48	14.25	11.62	10.30	9.36
	ParallelHash256	256	AVX512	–	–	–	–	–	–	–
	BLAKE3	128	SSE	5.20	4.05	3.77	3.64	3.57	1.85	1.26
	KangarooTwelve	128	AVX2	10.65	5.40	4.99	4.73	4.06	3.73	3.57
	KangarooTwelve	128	AVX512	–	–	–	–	–	–	–

† The security level is against preimage attacks.

‡ Our Alder Lake platform does not support the AVX512 instruction set.

- Areion512-DM realizes the fastest SFIL hashing among the target DM constructions with the 512-bit permutation. Specifically, Areion512-DM performs at least 1.34 and 1.25 times faster than Simpira256-DM and Haraka256-DM (x1.2), respectively.

Table 12: Benchmarks for hash functions on the Pixel 5, Pixel 6, Pixel 7, iPhone13, iPhone14, and iPadPro. All values are given as Gbps.

Platform	Primitive	Input sizes (bytes)							
		32	64	128	256	512	1024	2048	4096
Pixel 5 (Snapdragon765G)	Areion512-MD	**3.99**	**4.38**	**4.82**	**5.07**	**5.22**	**5.29**	**5.33**	**5.33**
	SHA2-256	0.33	0.63	1.23	2.21	3.69	5.57	7.42	8.96
	SHA3-256	0.20	0.40	0.80	1.17	1.51	1.76	1.93	2.09
Pixel 6 (Google Tensor)	Areion512-MD	**6.07**	**7.28**	**7.20**	**7.19**	**7.18**	**7.18**	**7.18**	**7.18**
	SHA2-256	0.45	0.86	1.62	2.87	4.71	6.94	9.07	10.72
	SHA3-256	0.28	0.56	1.12	1.66	2.19	2.59	2.86	3.11
Pixel 7 (Google Tensor2)	Areion512-MD	**5.81**	**7.40**	**7.33**	**7.33**	**7.27**	**7.31**	**7.28**	**7.29**
	SHA2-256	0.44	0.86	1.61	2.87	4.75	6.96	9.02	10.89
	SHA3-256	0.28	0.57	1.13	1.67	2.21	2.63	2.88	3.14
iPhone 13 (A15)	Areion512-MD	**8.39**	**14.71**	**13.84**	**11.15**	**10.19**	**9.75**	**9.56**	**9.46**
	SHA2-256	0.96	1.81	3.44	6.00	9.08	12.19	14.70	16.42
	SHA3-256	0.47	0.96	1.97	2.78	3.51	3.98	4.33	4.67
iPhone 14 (A15)	Areion512-MD	**8.03**	**14.84**	**13.95**	**11.20**	**10.21**	**9.79**	**9.58**	**9.48**
	SHA2-256	0.94	1.77	3.29	5.77	8.98	12.19	14.63	16.36
	SHA3-256	0.50	0.97	1.90	2.69	3.39	3.87	4.17	4.49
iPad Pro (Apple M1)	Areion512-MD	**8.39**	**14.98**	**14.38**	**11.74**	**10.75**	**10.31**	**10.11**	**10.00**
	SHA2-256	0.49	1.81	3.40	5.92	8.68	12.03	14.48	16.19
	SHA3-256	0.47	0.95	1.98	2.76	3.49	3.92	4.33	4.56
iPad Pro (Apple M2)	Areion512-MD	**8.67**	**15.59**	**15.08**	**12.09**	**11.04**	**10.58**	**10.36**	**10.24**
	SHA2-256	1.02	1.95	3.65	6.32	9.55	12.84	15.92	17.76
	SHA3-256	0.51	1.03	2.06	2.91	3.64	4.11	4.45	4.78

Consequently, It can be considered that Areion256-DM and Areion512-DM are the fastest SFIL hash functions. On the other hand, we summarize the performance comparison for the VIL hash functions as follows:

- Areion512-MD realizes the fastest VIL hashing among the target hash functions with a 256-bit security level for input sizes up to around 4K bytes. Specifically, its performance is less than 3 cpb for any message size. Moreover, it is much faster than existing state-of-the-art schemes (*e.g.*, SHA2-256, SHA3-256, and ParallelHash256) for short messages up to around 100 bytes, a widely-used input size in real-world applications.

Considering the need for cryptographic primitives resistant to symmetric-key cryptanalysis based on quantum algorithms (*e.g.*, Grover's algorithm [24]), hash functions with a 256-bit security level must be required for the future. For this reason, we consider that there is no problem even if Areion512-MD performs slower than KangarooTwelve when the input size is 2K bytes or more. In addition, according to the current study on packet sizes on the Internet [54], it is known that around 44% of packets are between 40 and 100 bytes long and 37% are between 1400 and 1500 bytes in size. Given that most of the packet sizes on the Internet are 1.5K bytes or less, Areion512-MD has the strongest advantage of performing faster than any other target VIL hash functions with a 256-bit security level.

Tables 12 shows benchmarks for the VIL hash functions using NEON implementations in Appendix A.3 on mobile environments. We compare with existing schemes of SHA2-256 and SHA3-256 which are available for optimized implementations in OpenSSL. Areion512-MD achieves outstanding performance for short messages, especially up to 128 bytes.

6.3 Permutation-based AEADs

Finally, we evaluate the performance of the permutation-based AEADs, *i.e.*, SFIL and VIL AEADs (the KWF, OPP, and OTR modes). These instantiations of Areion are implemented based on the source codes of Areion-256 and Areion-512 described in Appendix A.1. For comparison regarding the SFIL AEAD, we used Simpira512-KWF. It is implemented in the same way as Areion512-KWF. We note here that the key wrapping functions instantiated

Table 13: Benchmarks for SFIL AEADs on the Ice Lake, Tiger Lake, and Alder Lake platforms. The sizes of associated data and plaintext are 256 and 16 bytes, respectively. All values are given as cpb.

	Ice Lake		Tiger Lake		Alder Lake	
Primitive	Enc	Dec	Enc	Dec	Enc	Dec
Areion512-KWF	**2.38**	**2.38**	**2.37**	**2.37**	**2.30**	**2.30**
Simpira512-KWF	2.87	2.87	2.86	2.86	2.69	2.69

Table 14: Benchmarks for VIL AEADs on the Ice Lake, Tiger Lake, and Alder Lake platforms. The size of the associated data is 128 bytes. All values are given as cpb.

		Input sizes (bytes)					
		64	128	256	512	1024	2048
Platform	Primitive	Enc / Dec	Enc / Dec	Enc / Dec	Enc / Dec	Enc / Dec	Enc / Dec
Ice Lake	**Areion256-OPP**	**1.95 / 2.10**	**1.52 / 2.03**	**1.22 / 1.46**	**1.03 / 1.32**	**0.88 / 1.20**	**0.80 / 1.14**
	Areion256-OTR	**2.23 / 2.23**	**1.82 / 1.82**	**1.29 / 1.45**	**1.06 / 1.26**	**0.91 / 1.15**	**0.82 / 1.08**
	Simpira256-OPP	2.39 / 2.39	1.83 / 1.87	1.52 / 1.49	1.29 / 1.27	1.13 / 1.10	1.04 / 1.02
	Simpira256-OTR	2.75 / 2.75	2.29 / 2.29	1.62 / 1.64	1.35 / 1.40	1.18 / 1.23	1.07 / 1.13
	Deoxys-I-128	2.72 / 3.99	2.20 / 3.17	1.69 / 2.50	1.27 / 1.85	1.02 / 1.39	0.87 / 1.13
	BLAKE2s-OPP	5.28 / 4.84	4.55 / 4.03	3.91 / 3.25	3.44 / 2.72	3.12 / 2.33	2.93 / 2.10
	BLAKE2s-OTR	6.67 / 6.40	4.81 / 4.80	3.93 / 3.96	3.61 / 3.59	3.32 / 3.29	3.15 / 3.14
Tiger Lake	**Areion256-OPP**	**1.96 / 2.12**	**1.53 / 2.04**	**1.22 / 1.45**	**1.03 / 1.29**	**0.87 / 1.18**	**0.79 / 1.13**
	Areion256-OTR	**2.22 / 2.23**	**1.79 / 1.80**	**1.30 / 1.45**	**1.06 / 1.27**	**0.91 / 1.15**	**0.82 / 1.07**
	Simpira256-OPP	2.35 / 2.36	1.85 / 1.89	1.53 / 1.50	1.29 / 1.26	1.13 / 1.10	1.04 / 1.01
	Simpira256-OTR	2.76 / 2.73	2.33 / 2.32	1.61 / 1.63	1.36 / 1.40	1.18 / 1.22	1.07 / 1.12
	Deoxys-I-128	2.76 / 3.96	2.28 / 3.15	1.76 / 2.49	1.31 / 1.84	1.03 / 1.38	0.88 / 1.12
	BLAKE2s-OPP	5.27 / 4.83	4.54 / 4.02	3.91 / 3.25	3.45 / 2.71	3.12 / 2.32	2.92 / 2.10
	BLAKE2s-OTR	6.64 / 6.40	4.81 / 4.80	3.95 / 3.93	3.61 / 3.58	3.32 / 3.29	3.14 / 3.13
Alder Lake	**Areion256-OPP**	**1.94 / 2.06**	**1.52 / 1.67**	**1.15 / 1.26**	**0.95 / 1.09**	**0.81 / 0.97**	**0.73 / 0.89**
	Areion256-OTR	**2.00 / 2.00**	**1.74 / 1.77**	**1.22 / 1.26**	**1.00 / 1.06**	**0.86 / 0.91**	**0.77 / 0.83**
	Simpira256-OPP	2.33 / 2.31	1.78 / 1.84	1.34 / 1.37	1.12 / 1.16	0.96 / 1.01	0.86 / 0.91
	Simpira256-OTR	2.46 / 2.46	2.24 / 2.25	1.53 / 1.58	1.27 / 1.32	1.10 / 1.16	1.00 / 1.06
	Deoxys-I-128	2.29 / 3.22	1.71 / 2.33	1.36 / 1.82	1.09 / 1.42	0.90 / 1.15	0.79 / 0.99
	BLAKE2s-OPP	4.99 / 4.55	4.30 / 3.81	3.68 / 3.12	3.21 / 2.60	2.86 / 2.23	2.68 / 2.00
	BLAKE2s-OTR	6.47 / 6.15	4.64 / 4.62	3.70 / 3.71	3.33 / 3.33	3.04 / 3.02	2.87 / 2.87

with Haraka-512 and Haraka-512 (x1.2/x1.5) are excluded from our comparison because Haraka-512 cannot be regarded to have the security margin sufficiently, and the underlying permutation of Haraka-512 (x1.2/x1.5) are clearly slower than that of Areion-512, as explained Sect. 6.1. On the other hand, for comparison regarding the VIL AEADs, we used Deoxys-I-128 and the OPP and OTR modes instantiated with BLAKE2s called BLAKE2s-OPP and BLAKE2s-OTR. We can find the source codes of Deoxys-I-128 and BLAKE2s-OPP at SUPERCOP[13] and GitHub[14], respectively; then, we modified these source codes to use for our performance comparison. In addition, BLAKE2s-OTR is implemented in the same way as Areion256-OTR.

Our platforms' benchmarks for the SFIL and VIL AEADs are shown in Tables 13 and 14. From Table 13, Areion512-KWF performs at least 1.17 times faster than Simpira512-KWF for encryption and decryption; thus, Areion512-KWF is the fastest SFIL AEAD. On the other hand, we summarize the performance comparison regarding the VIL AEADs as follows:

- Areion256-OPP realizes the fastest encryption among the target AEADs. Specifically, Areion256-OPP is the only AEAD to achieve less than two cpb for 64-bit messages on all our platforms. In addition, Areion256-OPP provides the fastest encryption/decryption among the target AEADs, excluding Areion256-OTR, on the latest CPU architecture Alder Lake.

- Areion256-OTR realizes the fastest encryption/decryption among the target AEADs, excluding Areion256-OPP and Simpira256-OPP. Specifically, compared to its main

[13] https://bench.cr.yp.to/supercop.html
[14] https://github.com/MEM-AEAD/mem-aead

competitor Simpira256-OTR, Areion256-OTR provides faster encryption/decryption than Simpira256-OTR for all messages up to 2K bytes. In addition, when compared to Areion256-OPP, Areion256-OTR is balanced in encryption and decryption for long messages of 256 bits or more.

From the above viewpoints, Areion256-OPP and Areion256-OTR have better performance than any other target AEADs.

7 Conclusion

We proposed a family of wide-block permutations Areion that fully leverages the power of AES instructions and show its applications of hash functions and AEADs. Our schemes significantly outperform existing schemes for short input and are competitive for relatively-long messages. Among them, our hash function is surprisingly fast. Its performance is less than three cycles/byte in the latest Intel architectures for any message size. It is about ten times faster than existing schemes for short messages up to around 100 bytes, which are the most widely-used input size in real-world applications, on both of on latest CPU architectures (IceLake, Tiger Lake, and Alder Lake) and mobile environments (Pixel 7, iPhone 14, and iPad Pro with Apple M2).

Acknowledgments

Takanori Isobe is supported by JST, PRESTO Grant Number JPMJPR2031. These research results were also obtained from the commissioned research (No.05801) by National Institute of Information and Communications Technology (NICT), Japan. Fukang Liu is supported by Grant-in-Aid for Research Activity Start-up (Grant No. 22K21282). Kosei Sakamoto is supported by Grant-in-Aid for JSPS Fellows (KAKENHI 20J23526) for Japan Society for the Promotion of Science.

We would like to thank Samuel Neves for informing us about more efficient implementations of SHA2-256 and BLAKE3. We also would like to thank Frank Denis for pointing out the issues of the test vectors and suggesting improvements on the code.

References

[1] Submission Requirements and Evaluation Criteria for the Lightweight Cryptography Standardization Process (2018), https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/final-lwc-submission-requirements-august2018.pdf, National Institute of Standards and Technology

[2] 3GPP TS 36.213: Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures. https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2427 (2019)

[3] Andreeva, E., Lallemand, V., Purnal, A., Reyhanitabar, R., Roy, A., Vizár, D.: Forkcipher: A New Primitive for Authenticated Encryption of Very Short Messages. In: Galbraith, S.D., Moriai, S. (eds.) Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part II. Lecture Notes in Computer Science, vol. 11922, pp. 153–182. Springer (2019). https://doi.org/10.1007/978-3-030-34621-8_6, https://doi.org/10.1007/978-3-030-34621-8_6

[4] Aumasson, J.P., Jovanovic, P., Neves, S.: Norx v3.0. Submission to CAESAR competition (2016), https://competitions.cr.yp.to/round3/norxv30.pdf

[5] Aumasson, J.P., Meier, W.: Zero-sum Distinguishers for Reduced Keccak-f and for the Core Functions of Luffa and Hamsi (2009), https://131002.net/data/papers/AM09.pdf

[6] Bao, Z., Dong, X., Guo, J., Li, Z., Shi, D., Sun, S., Wang, X.: Automatic Search of Meet-in-the-Middle Preimage Attacks on AES-like Hashing. In: Canteaut, A., Standaert, F. (eds.) Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12696, pp. 771–804. Springer (2021). https://doi.org/10.1007/978-3-030-77870-5_27, https://doi.org/10.1007/978-3-030-77870-5_27

[7] Bellare, M., Rogaway, P.: Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In: Okamoto, T. (ed.) Advances in Cryptology - ASIACRYPT 2000, 6th International Conference on the Theory and Application of Cryptology and Information Security, Kyoto, Japan, December 3-7, 2000, Proceedings. Lecture Notes in Computer Science, vol. 1976, pp. 317–330. Springer (2000). https://doi.org/10.1007/3-540-44448-3_24, https://doi.org/10.1007/3-540-44448-3_24

[8] Bertoni, G., Daemen, J., Peeters, M., Assche, G.V.: Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications. In: Miri, A., Vaudenay, S. (eds.) Selected Areas in Cryptography - 18th International Workshop, SAC 2011, Toronto, ON, Canada, August 11-12, 2011, Revised Selected Papers. Lecture Notes in Computer Science, vol. 7118, pp. 320–337. Springer (2011). https://doi.org/10.1007/978-3-642-28496-0_19, https://doi.org/10.1007/978-3-642-28496-0_19

[9] Bertoni, G., Daemen, J., Peeters, M., Assche, G.V., Keer, R.V.: Keyak v2.2. Submission to CAESAR competition (2016), https://competitions.cr.yp.to/round3/keyakv22.pdf

[10] Bertoni, G., Daemen, J., Peeters, M., Assche, G.V., Keer, R.V., Viguier, B.: KangarooTwelve: Fast Hashing Based on Keccak-p. In: Preneel, B., Vercauteren, F. (eds.) Applied Cryptography and Network Security - 16th International Conference, ACNS 2018, Leuven, Belgium, July 2-4, 2018, Proceedings. Lecture Notes in Computer Science, vol. 10892, pp. 400–418. Springer (2018). https://doi.org/10.1007/978-3-319-93387-0_21, https://doi.org/10.1007/978-3-319-93387-0_21

[11] Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. J. ACM **51**(4), 557–594 (2004)

[12] Canteaut, A., Duval, S., Leurent, G., Naya-Plasencia, M., Perrin, L., Pornin, T., Schrottenloher, A.: Saturnin: a suite of lightweight symmetric algorithms for post-quantum security. IACR Trans. Symmetric Cryptol. **2020**(S1), 160–207 (2020). https://doi.org/10.13154/tosc.v2020.iS1.160-207, https://doi.org/10.13154/tosc.v2020.iS1.160-207

[13] Cogliati, B., Lampe, R., Seurin, Y.: Tweaking Even-Mansour Ciphers. In: Gennaro, R., Robshaw, M. (eds.) Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9215, pp. 189–208. Springer (2015). https://doi.org/10.1007/978-3-662-47989-6_9, https://doi.org/10.1007/978-3-662-47989-6_9

[14] Cogliati, B., Seurin, Y.: Beyond-Birthday-Bound Security for Tweakable Even-Mansour Ciphers with Linear Tweak and Key Mixing. In: Iwata, T., Cheon, J.H. (eds.) Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9453, pp. 134–158. Springer (2015). https://doi.org/10.1007/978-3-662-48800-3_6, https://doi.org/10.1007/978-3-662-48800-3_6

[15] Cui, T., Jia, K., Fu, K., Chen, S., Wang, M.: New Automatic Search Tool for Impossible Differentials and Zero-Correlation Linear Approximations. IACR Cryptol. ePrint Arch. p. 689 (2016), http://eprint.iacr.org/2016/689

[16] Damgård, I.: A Design Principle for Hash Functions. In: Brassard, G. (ed.) Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings. Lecture Notes in Computer Science, vol. 435, pp. 416–427. Springer (1989). https://doi.org/10.1007/0-387-34805-0_39, https://doi.org/10.1007/0-387-34805-0_39

[17] Dobraunig, C., Eichlseder, M., Mendel, F.: Related-Key Forgeries for Prøst-OTR. In: Leander, G. (ed.) Fast Software Encryption - 22nd International Workshop, FSE 2015, Istanbul, Turkey, March 8-11, 2015, Revised Selected Papers. Lecture Notes in Computer Science, vol. 9054, pp. 282–296. Springer (2015). https://doi.org/10.1007/978-3-662-48116-5_14, https://doi.org/10.1007/978-3-662-48116-5_14

[18] Dobraunig, C., Eichlseder, M., Mendel, F., Schläffer, M.: Ascon v1.2. Submission to CAESAR competition (2016), https://competitions.cr.yp.to/round3/asconv12.pdf

[19] Dobraunig, C., Eichlseder, M., Mendel, F., Schläffer, M.: Ascon v1.2. Submission to Round 1 of the NIST Lightweight Cryptography project (2019), https://csrc.nist.gov/CSRC/media/Projects/Lightweight-Cryptography/documents/round-1/spec-doc/ascon-spec.pdf

[20] Dworkin, M.: NIST SP 800-38D, Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC (2007), U.S. Department of Commerce/NIST

[21] Gilbert, H.: A simplified representation of AES. In: ASIACRYPT (1). Lecture Notes in Computer Science, vol. 8873, pp. 200–222. Springer (2014)

[22] Gilbert, H., Peyrin, T.: Super-sbox cryptanalysis: Improved attacks for aes-like permutations. In: FSE. Lecture Notes in Computer Science, vol. 6147, pp. 365–383. Springer (2010)

[23] Granger, R., Jovanovic, P., Mennink, B., Neves, S.: Improved Masking for Tweakable Blockciphers with Applications to Authenticated Encryption. In: Fischlin, M., Coron, J. (eds.) Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9665, pp. 263–293. Springer (2016). https://doi.org/10.1007/978-3-662-49890-3_11, https://doi.org/10.1007/978-3-662-49890-3_11

[24] Grover, L.K.: A Fast Quantum Mechanical Algorithm for Database Search. In: Miller, G.L. (ed.) Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996. pp. 212–219. ACM (1996). https://doi.org/10.1145/237814.237866, https://doi.org/10.1145/237814.237866

[25] Gueron, S., Mouha, N.: Simpira v2: A Family of Efficient Permutations Using the AES Round Function. In: Cheon, J.H., Takagi, T. (eds.) Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10031, pp. 95–125 (2016). https://doi.org/10.1007/978-3-662-53887-6_4, https://doi.org/10.1007/978-3-662-53887-6_4

[26] Hirose, S.: Some Plausible Constructions of Double-Block-Length Hash Functions. In: Robshaw, M.J.B. (ed.) Fast Software Encryption, 13th International Workshop, FSE 2006, Graz, Austria, March 15-17, 2006, Revised Selected Papers. Lecture Notes in Computer Science, vol. 4047, pp. 210–225. Springer (2006). https://doi.org/10.1007/11799313_14, https://doi.org/10.1007/11799313_14

[27] Housley, R.: Using AES-CCM and AES-GCM Authenticated Encryption in the Cryptographic Message Syntax (CMS). RFC **5084**, 1–11 (2007). https://doi.org/10.17487/RFC5084, https://doi.org/10.17487/RFC5084

[28] Isobe, T., Ito, R., Liu, F., Minematsu, K., Nakahashi, M., Sakamoto, K., Shiba, R.: Areion: Highly-efficient permutations and its applications to hash functions for short input. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2023**(2), 115–154 (2023). https://doi.org/10.46586/tches.v2023.i2.115-154, https://doi.org/10.46586/tches.v2023.i2.115-154

[29] Isobe, T., Ito, R., Minematsu, K.: Security Analysis of SFrame. In: Bertino, E., Shulman, H., Waidner, M. (eds.) Computer Security - ESORICS 2021 - 26th European Symposium on Research in Computer Security, Darmstadt, Germany, October 4-8, 2021, Proceedings, Part II. Lecture Notes in Computer Science, vol. 12973, pp. 127–146. Springer (2021). https://doi.org/10.1007/978-3-030-88428-4_7, https://doi.org/10.1007/978-3-030-88428-4_7

[30] Iwata, T., Khairallah, M., Minematsu, K., Peyrin, T.: Duel of the Titans: The Romulus and Remus Families of Lightweight AEAD Algorithms. IACR Trans. Symmetric Cryptol. **2020**(1), 43–120 (2020). https://doi.org/10.13154/tosc.v2020.i1.43-120, https://doi.org/10.13154/tosc.v2020.i1.43-120

[31] Jean, J., Naya-Plasencia, M., Peyrin, T.: Improved rebound attack on the finalist grøstl. In: Canteaut, A. (ed.) Fast Software Encryption - 19th International Workshop, FSE 2012, Washington, DC, USA, March 19-21, 2012. Revised Selected Papers. Lecture Notes in Computer Science, vol. 7549, pp. 110–126. Springer (2012). https://doi.org/10.1007/978-3-642-34047-5_7, https://doi.org/10.1007/978-3-642-34047-5_7

[32] Jean, J., Nikolic, I., Peyrin, T.: Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In: Sarkar, P., Iwata, T. (eds.) Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II. Lecture Notes in Computer Science, vol. 8874, pp. 274–288. Springer (2014). https://doi.org/10.1007/978-3-662-45608-8_15, https://doi.org/10.1007/978-3-662-45608-8_15

[33] Jean, J., Nikolic, I., Peyrin, T., Seurin, Y.: The Deoxys AEAD Family. J. Cryptol. **34**(3), 31 (2021). https://doi.org/10.1007/s00145-021-09397-w, https://doi.org/10.1007/s00145-021-09397-w

[34] Josh Blum and Simon Booth and Oded Gal and Maxwell Krohn and Julia Len and Karan Lyons and Antonio Marcedone and Mike Maxim and Merry Ember Mou and Jack O'Connor and Miles Steele and Matthew Green and Lea Kissner and Alex Stamos: E2E Encryption for Zoom Meetings – Version 3.2 (October 2021), https://github.com/zoom/zoom-e2e-whitepaper

[35] Kavun, E.B., Lauridsen, M.M., Leander, G., Rechberger, C., Schwabe, P., Yalçın, T.: Prøst. CAESAR Proposal (2014), http://proest.compute.dtu.dk

[36] Kavun, E.B., Lauridsen, M.M., Leander, G., Rechberger, C., Schwabe, P., Yalçın, T.: Prøst v1.1. Submission to CAESAR competition (2014), https://competitions.cr.yp.to/round1/proestv11.pdf

[37] Kelsey, J., jen Chang, S., Perlner, R.: SHA-3 Derived Functions: cSHAKE, KMAC, TupleHash and ParallelHash. NIST special publication **800**, 185 (2016)

[38] Kelsey, J., Schneier, B.: Second preimages on n-bit hash functions for much less than 2^n work. In: EUROCRYPT. Lecture Notes in Computer Science, vol. 3494, pp. 474–490. Springer (2005)

[39] Khovratovich, D.: Key Wrapping with a Fixed Permutation. In: Benaloh, J. (ed.) Topics in Cryptology - CT-RSA 2014 - The Cryptographer's Track at the RSA Conference 2014, San Francisco, CA, USA, February 25-28, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8366, pp. 481–499. Springer (2014). https://doi.org/10.1007/978-3-319-04852-9_25, https://doi.org/10.1007/978-3-319-04852-9_25

[40] Knudsen, L.R., Rijmen, V.: Known-key distinguishers for some block ciphers. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 4833, pp. 315–324. Springer (2007)

[41] Kölbl, S., Lauridsen, M.M., Mendel, F., Rechberger, C.: Haraka v2 - Efficient Short-Input Hashing for Post-Quantum Applications. IACR Trans. Symmetric Cryptol. **2016**(2), 1–29 (2016). https://doi.org/10.13154/tosc.v2016.i2.1-29, https://doi.org/10.13154/tosc.v2016.i2.1-29

[42] Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) Fast Software Encryption - 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers. Lecture Notes in Computer Science, vol. 6733, pp. 306–327. Springer (2011). https://doi.org/10.1007/978-3-642-21702-9_18, https://doi.org/10.1007/978-3-642-21702-9_18

[43] Lamberger, M., Mendel, F., Rechberger, C., Rijmen, V., Schläffer, M.: Rebound distinguishers: Results on the full whirlpool compression function. In: ASIACRYPT. Lecture Notes in Computer Science, vol. 5912, pp. 126–143. Springer (2009)

[44] Latva-aho, M., Leppänen, K.: Key drivers and research challenges for 6g ubiquitous wireless intelligence. 6G Research Visions 1 (2019)

[45] Marshall, B., Newell, G.R., Page, D., Saarinen, M.O., Wolf, C.: The design of scalar AES Instruction Set Extensions for RISC-V. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2021**(1), 109–136 (2021). https://doi.org/10.46586/tches.v2021.i1.109-136, https://doi.org/10.46586/tches.v2021.i1.109-136

[46] McGrew, D.A., Viega, J.: The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In: Canteaut, A., Viswanathan, K. (eds.) Progress in Cryptology - INDOCRYPT 2004, 5th International Conference on Cryptology in India, Chennai, India, December 20-22, 2004, Proceedings. Lecture Notes in Computer Science, vol. 3348, pp. 343–355. Springer (2004). https://doi.org/10.1007/978-3-540-30556-9_27, https://doi.org/10.1007/978-3-540-30556-9_27

[47] McKay, K.A., Bassham, L., Turan, M.S., Mouha, N.: Report on Lightweight Cryptography (2017), https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8114.pdf, National Institute of Standards and Technology IR 8114

[48] Mendel, F., Rechberger, C., Schläffer, M., Thomsen, S.S.: The rebound attack: Cryptanalysis of reduced whirlpool and grøstl. In: Dunkelman, O. (ed.) Fast Software Encryption, 16th International Workshop, FSE 2009, Leuven, Belgium, February 22-25, 2009, Revised Selected Papers. Lecture Notes in Computer Science, vol. 5665, pp. 260–276. Springer (2009). https://doi.org/10.1007/978-3-642-03317-9_16, https://doi.org/10.1007/978-3-642-03317-9_16

[49] Mennink, B.: XPX: Generalized Tweakable Even-Mansour with Improved Security Guarantees. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I. Lecture Notes in Computer Science, vol. 9814, pp. 64–94. Springer (2016). https://doi.org/10.1007/978-3-662-53018-4_3, https://doi.org/10.1007/978-3-662-53018-4_3

[50] Mennink, B., Reyhanitabar, R., Vizár, D.: Security of Full-State Keyed Sponge and Duplex: Applications to Authenticated Encryption. In: Iwata, T., Cheon, J.H. (eds.) Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9453, pp. 465–489. Springer (2015). https://doi.org/10.1007/978-3-662-48800-3_19, https://doi.org/10.1007/978-3-662-48800-3_19

[51] Merkle, R.C.: One Way Hash Functions and DES. In: Brassard, G. (ed.) Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings. Lecture Notes in Computer Science, vol. 435, pp. 428–446. Springer (1989). https://doi.org/10.1007/0-387-34805-0_40, https://doi.org/10.1007/0-387-34805-0_40

[52] Minematsu, K.: Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In: Nguyen, P.Q., Oswald, E. (eds.) Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8441, pp. 275–292. Springer (2014). https://doi.org/10.1007/978-3-642-55220-5_16, https://doi.org/10.1007/978-3-642-55220-5_16

[53] Mouha, N., Wang, Q., Gu, D., Preneel, B.: Differential and Linear Cryptanalysis Using Mixed-Integer Linear Programming. In: Wu, C., Yung, M., Lin, D. (eds.) Information Security and Cryptology - 7th International Conference, Inscrypt 2011, Beijing, China, November 30 - December 3, 2011. Revised Selected Papers. Lecture Notes in Computer Science, vol. 7537, pp. 57–76. Springer (2011). https://doi.org/10.1007/978-3-642-34704-7_5, https://doi.org/10.1007/978-3-642-34704-7_5

[54] Murray, D., Koziniec, T., Zander, S., Dixon, M., Koutsakis, P.: An Analysis of Changing Enterprise Network Traffic Characteristics. In: 23rd Asia-Pacific Conference

on Communications, APCC 2017, Perth, Australia, December 11-13, 2017. pp. 1–6. IEEE (2017). https://doi.org/10.23919/APCC.2017.8303960, https://doi.org/10.23919/APCC.2017.8303960

[55] Nir, Y., Langley, A.: ChaCha20 and Poly1305 for IETF Protocols. RFC **8439**, 1–46 (2018). https://doi.org/10.17487/RFC8439, https://doi.org/10.17487/RFC8439

[56] O'Connor, J., Aumasson, J.P., Neves, S., Wilcox-O'Hearn, Z.: BLAKE3: One Function, Fast Everywhere (2020), https://github.com/BLAKE3-team/BLAKE3-specs/blob/master/blake3.pdf

[57] Omara, E., Uberti, J., Gouaillard, A., Murillo, S.G.: Secure Frame (SFrame). https://tools.ietf.org/html/draft-omara-sframe-03 (August 2021)

[58] Real-Time, Lab, E.S.: uops.info. Official webpage, https://www.uops.info/ (2021)

[59] Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: a block-cipher mode of operation for efficient authenticated encryption. In: Reiter, M.K., Samarati, P. (eds.) CCS 2001, Proceedings of the 8th ACM Conference on Computer and Communications Security, Philadelphia, Pennsylvania, USA, November 6-8, 2001. pp. 196–205. ACM (2001). https://doi.org/10.1145/501983.502011, https://doi.org/10.1145/501983.502011

[60] Rogaway, P., Shrimpton, T.: A Provable-Security Treatment of the Key-Wrap Problem. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4004, pp. 373–390. Springer (2006). https://doi.org/10.1007/11761679_23, https://doi.org/10.1007/11761679_23

[61] Sasaki, Y.: Meet-in-the-middle preimage attacks on AES hashing modes and an application to whirlpool. In: Joux, A. (ed.) Fast Software Encryption - 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers. Lecture Notes in Computer Science, vol. 6733, pp. 378–396. Springer (2011). https://doi.org/10.1007/978-3-642-21702-9_22, https://doi.org/10.1007/978-3-642-21702-9_22

[62] Sasaki, Y., Todo, Y.: New Impossible Differential Search Tool from Design and Cryptanalysis Aspects - Revealing Structural Properties of Several Ciphers. In: Coron, J., Nielsen, J.B. (eds.) Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part III. Lecture Notes in Computer Science, vol. 10212, pp. 185–215 (2017). https://doi.org/10.1007/978-3-319-56617-7_7, https://doi.org/10.1007/978-3-319-56617-7_7

[63] Sasaki, Y., Todo, Y., Aoki, K., Naito, Y., Sugawara, T., Murakami, Y., Matsui, M., Hirose, S.: Minalpher v1.1. Submission to CAESAR competition (2015), https://competitions.cr.yp.to/round2/minalpherv11.pdf

[64] Shiba, R., Sakamoto, K., Isobe, T.: Efficient constructions for large-state block ciphers based on AES New Instructions. IET Information Security (2021). https://doi.org/https://doi.org/10.1049/ise2.12053, https://ietresearch.onlinelibrary.wiley.com/doi/abs/10.1049/ise2.12053

[65] of Standards, N.I., Technology: Secure Hash Standard (SHS). Federal Information Processing Standards Publication. FIPS PUB 180-4 (August 2015), https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.180-4.pdf

[66] of Standards, N.I., Technology: SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions. Federal Information Processing Standards Publication. FIPS PUB 202 (August 2015), https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.202.pdf

[67] Systems, C.: Zero-Trust Security for Cisco Webex (2020), https://www.cisco.com/c/en/us/solutions/collateral/collaboration/white-paper-c11-744553.html

[68] Xiang, Z., Zhang, W., Bao, Z., Lin, D.: Applying MILP Method to Searching Integral Distinguishers Based on Division Property for 6 Lightweight Block Ciphers. In: Cheon, J.H., Takagi, T. (eds.) Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I. Lecture Notes in Computer Science, vol. 10031, pp. 648–678 (2016). https://doi.org/10.1007/978-3-662-53887-6_24, https://doi.org/10.1007/978-3-662-53887-6_24

A Reference Implementations

A.1 Areion-256 and Areion-512

```
#include <stdint.h>
#include <immintrin.h>

/* Round Constants */
const uint32_t RC[15*4] = {
   0x243f6a88, 0x85a308d3, 0x13198a2e, 0x03707344,
   0xa4093822, 0x299f31d0, 0x082efa98, 0xec4e6c89,
   0x452821e6, 0x38d01377, 0xbe5466cf, 0x34e90c6c,
   0xc0ac29b7, 0xc97c50dd, 0x3f84d5b5, 0xb5470917,
   0x9216d5d9, 0x8979fb1b, 0xd1310ba6, 0x98dfb5ac,
   0x2ffd72db, 0xd01adfb7, 0xb8e1afed, 0x6a267e96,
   0xba7c9045, 0xf12c7f99, 0x24a19947, 0xb3916cf7,
   0x801f2e28, 0x58efc166, 0x36920d87, 0x1574e690,
   0xa458fea3, 0xf4933d7e, 0x0d95748f, 0x728eb658,
   0x718bcd58, 0x82154aee, 0x7b54a41d, 0xc25a59b5,
   0x9c30d539, 0x2af26013, 0xc5d1b023, 0x286085f0,
   0xca417918, 0xb8db38ef, 0x8e79dcb0, 0x603a180e,
   0x6c9e0e8b, 0xb01e8a3e, 0xd71577c1, 0xbd314b27,
   0x78af2fda, 0x55605c60, 0xe65525f3, 0xaa55ab94,
   0x57489862, 0x63e81440, 0x55ca396a, 0x2aab10b6
};

/* Round constants are used in little-endian byte order. */
#define RC0(i) _mm_setr_epi32(RC[(i)*4+3], RC[(i)*4+2], RC[(i)*4+1], RC[(i)*4+0])
#define RC1(i) _mm_setr_epi32(0, 0, 0, 0)

/* Round Function for the 256-bit permutation */
#define Round_Function_256(x0, x1, i) do { \
   x1 = _mm_aesenc_si128(_mm_aesenc_si128(x0, RC0(i)), x1); \
   x0 = _mm_aesenclast_si128(x0, RC1(i)); \
} while(0)

/* 256-bit permutation */
#define perm256(x0, x1) do { \
   Round_Function_256(x0, x1, 0); \
   Round_Function_256(x1, x0, 1); \
   Round_Function_256(x0, x1, 2); \
   Round_Function_256(x1, x0, 3); \
   Round_Function_256(x0, x1, 4); \
   Round_Function_256(x1, x0, 5); \
   Round_Function_256(x0, x1, 6); \
   Round_Function_256(x1, x0, 7); \
   Round_Function_256(x0, x1, 8); \
   Round_Function_256(x1, x0, 9); \
} while(0)

/* Inversed Round Function for the 256-bit permutation */
#define Inv_Round_Function_256(x0, x1, i) do { \
   x0 = _mm_aesdeclast_si128(x0, RC1(i)); \
   x1 = _mm_aesenc_si128(_mm_aesenc_si128(x0, RC0(i)), x1); \
} while(0)

/* Inversed 256-bit permutation */
#define Inv_perm256(x0, x1) do { \
   Inv_Round_Function_256(x1, x0, 9); \
   Inv_Round_Function_256(x0, x1, 8); \
   Inv_Round_Function_256(x1, x0, 7); \
   Inv_Round_Function_256(x0, x1, 6); \
   Inv_Round_Function_256(x1, x0, 5); \
   Inv_Round_Function_256(x0, x1, 4); \
   Inv_Round_Function_256(x1, x0, 3); \
   Inv_Round_Function_256(x0, x1, 2); \
```

```
   Inv_Round_Function_256(x1, x0, 1); \
   Inv_Round_Function_256(x0, x1, 0); \
} while(0)

/* Round Function for the 512-bit permutation */
#define Round_Function_512(x0, x1, x2, x3, i) do { \
   x1 = _mm_aesenc_si128(x0, x1); \
   x3 = _mm_aesenc_si128(x2, x3); \
   x0 = _mm_aesenclast_si128(x0, RC1(i)); \
   x2 = _mm_aesenc_si128(_mm_aesenclast_si128(x2, RC0(i)), RC1(i)); \
} while (0)

/* 512-bit permutation */
#define perm512(x0, x1, x2, x3) do { \
   Round_Function_512(x0, x1, x2, x3, 0); \
   Round_Function_512(x1, x2, x3, x0, 1); \
   Round_Function_512(x2, x3, x0, x1, 2); \
   Round_Function_512(x3, x0, x1, x2, 3); \
   Round_Function_512(x0, x1, x2, x3, 4); \
   Round_Function_512(x1, x2, x3, x0, 5); \
   Round_Function_512(x2, x3, x0, x1, 6); \
   Round_Function_512(x3, x0, x1, x2, 7); \
   Round_Function_512(x0, x1, x2, x3, 8); \
   Round_Function_512(x1, x2, x3, x0, 9); \
   Round_Function_512(x2, x3, x0, x1, 10); \
   Round_Function_512(x3, x0, x1, x2, 11); \
   Round_Function_512(x0, x1, x2, x3, 12); \
   Round_Function_512(x1, x2, x3, x0, 13); \
   Round_Function_512(x2, x3, x0, x1, 14); \
} while(0)

/* Inversed Round Function for the 512-bit permutation */
#define Inv_Round_Function_512(x0, x1, x2, x3, i) do { \
   x0 = _mm_aesdeclast_si128(x0, RC1(i)); \
   x2 = _mm_aesdeclast_si128(_mm_aesimc_si128(x2), RC0(i)); \
   x2 = _mm_aesdeclast_si128(x2, RC1(i)); \
   x1 = _mm_aesenc_si128(x0, x1); \
   x3 = _mm_aesenc_si128(x2, x3); \
} while (0)

/* Inversed 512-bit permutation */
#define Inv_perm512(x0, x1, x2, x3) do { \
   Inv_Round_Function_512(x3, x0, x1, x2, 14); \
   Inv_Round_Function_512(x2, x3, x0, x1, 13); \
   Inv_Round_Function_512(x1, x2, x3, x0, 12); \
   Inv_Round_Function_512(x0, x1, x2, x3, 11); \
   Inv_Round_Function_512(x3, x0, x1, x2, 10); \
   Inv_Round_Function_512(x2, x3, x0, x1, 9); \
   Inv_Round_Function_512(x1, x2, x3, x0, 8); \
   Inv_Round_Function_512(x0, x1, x2, x3, 7); \
   Inv_Round_Function_512(x3, x0, x1, x2, 6); \
   Inv_Round_Function_512(x2, x3, x0, x1, 5); \
   Inv_Round_Function_512(x1, x2, x3, x0, 4); \
   Inv_Round_Function_512(x0, x1, x2, x3, 3); \
   Inv_Round_Function_512(x3, x0, x1, x2, 2); \
   Inv_Round_Function_512(x2, x3, x0, x1, 1); \
   Inv_Round_Function_512(x1, x2, x3, x0, 0); \
} while(0)

/* Areion-256 */
void permute_areion_256(__m128i dst[2], const __m128i src[2])
{
    __m128i x0 = src[0];
    __m128i x1 = src[1];
    perm256(x0, x1);
    dst[0] = x0;
```

```
    dst[1] = x1;
}

/* Invresed Areion-256 */
void inverse_areion_256(__m128i dst[2], const __m128i src[2])
{
    __m128i x0 = src[0];
    __m128i x1 = src[1];
    Inv_perm256(x0, x1);
    dst[0] = x0;
    dst[1] = x1;
}

/* Areion-512 */
void permute_areion_512(__m128i dst[4], const __m128i src[4])
{
    __m128i x0 = src[0];
    __m128i x1 = src[1];
    __m128i x2 = src[2];
    __m128i x3 = src[3];
    perm512(x0, x1, x2, x3);
    dst[0] = x3;
    dst[1] = x0;
    dst[2] = x1;
    dst[3] = x2;
}

/* Invresed Areion-512 */
void inverse_areion_512(__m128i dst[4], const __m128i src[4])
{
    __m128i x0 = src[0];
    __m128i x1 = src[1];
    __m128i x2 = src[2];
    __m128i x3 = src[3];
    Inv_perm512(x0, x1, x2, x3);
    dst[0] = x1;
    dst[1] = x2;
    dst[2] = x3;
    dst[3] = x0;
}
```

A.2 Simpira-256 and Simpira-512

```
#include <stdint.h>
#include <immintrin.h>

/* Round Constant */
#define RC0(i) _mm_setr_epi32(0x00^(i)^(2), 0x10^(i)^(2), 0x20^(i)^(2), 0x30^(i)^(2))
#define RC1(i) _mm_setr_epi32(0x00^(2*(i)+1)^(4), 0x10^(2*(i)+1)^(4), 0x20^(2*(i)+1)^(4), 0
     x30^(2*(i)+1)^(4))
#define RC2(i) _mm_setr_epi32(0x00^(2*(i)+2)^(4), 0x10^(2*(i)+2)^(4), 0x20^(2*(i)+2)^(4), 0
     x30^(2*(i)+2)^(4))

/* Round Function for the 256-bit permutation */
#define Round_Function_256(x0, x1, i) do { \
   x1 = _mm_aesenc_si128(_mm_aesenc_si128(x0, RC0(i)), x1); \
} while(0)

/* 256-bit permutation */
#define perm256(x0, x1) do { \
   Round_Function_256(x0, x1, 0); \
   Round_Function_256(x1, x0, 1); \
   Round_Function_256(x0, x1, 2); \
   Round_Function_256(x1, x0, 3); \
   Round_Function_256(x0, x1, 4); \
```

```
    Round_Function_256(x1, x0, 5); \
    Round_Function_256(x0, x1, 6); \
    Round_Function_256(x1, x0, 7); \
    Round_Function_256(x0, x1, 8); \
    Round_Function_256(x1, x0, 9); \
    Round_Function_256(x0, x1, 10); \
    Round_Function_256(x1, x0, 11); \
    Round_Function_256(x0, x1, 12); \
    Round_Function_256(x1, x0, 13); \
    Round_Function_256(x0, x1, 14); \
} while(0)

/* Inversed 256-bit permutation */
#define Inv_perm256(x0, x1) do { \
    Round_Function_256(x0, x1, 14); \
    Round_Function_256(x1, x0, 13); \
    Round_Function_256(x0, x1, 12); \
    Round_Function_256(x1, x0, 11); \
    Round_Function_256(x0, x1, 10); \
    Round_Function_256(x1, x0, 9); \
    Round_Function_256(x0, x1, 8); \
    Round_Function_256(x1, x0, 7); \
    Round_Function_256(x0, x1, 6); \
    Round_Function_256(x1, x0, 5); \
    Round_Function_256(x0, x1, 4); \
    Round_Function_256(x1, x0, 3); \
    Round_Function_256(x0, x1, 2); \
    Round_Function_256(x1, x0, 1); \
    Round_Function_256(x0, x1, 0); \
} while(0)

/* Round Function for the 512-bit permutation */
#define Round_Function_512(x0, x1, x2, x3, i) do { \
    x1 = _mm_aesenc_si128(_mm_aesenc_si128(x0, RC1(i)), x1); \
    x3 = _mm_aesenc_si128(_mm_aesenc_si128(x2, RC2(i)), x3); \
} while (0)

/* 512-bit permutation */
#define perm512(x0, x1, x2, x3) do { \
    Round_Function_512(x0, x1, x2, x3, 0); \
    Round_Function_512(x1, x2, x3, x0, 1); \
    Round_Function_512(x2, x3, x0, x1, 2); \
    Round_Function_512(x3, x0, x1, x2, 3); \
    Round_Function_512(x0, x1, x2, x3, 4); \
    Round_Function_512(x1, x2, x3, x0, 5); \
    Round_Function_512(x2, x3, x0, x1, 6); \
    Round_Function_512(x3, x0, x1, x2, 7); \
    Round_Function_512(x0, x1, x2, x3, 8); \
    Round_Function_512(x1, x2, x3, x0, 9); \
    Round_Function_512(x2, x3, x0, x1, 10); \
    Round_Function_512(x3, x0, x1, x2, 11); \
    Round_Function_512(x0, x1, x2, x3, 12); \
    Round_Function_512(x1, x2, x3, x0, 13); \
    Round_Function_512(x2, x3, x0, x1, 14); \
} while(0)

/* Inversed 512-bit permutation */
#define Inv_perm512(x0, x1, x2, x3) do { \
    Round_Function_512(x2, x3, x0, x1, 14); \
    Round_Function_512(x1, x2, x3, x0, 13); \
    Round_Function_512(x0, x1, x2, x3, 12); \
    Round_Function_512(x3, x0, x1, x2, 11); \
    Round_Function_512(x2, x3, x0, x1, 10); \
    Round_Function_512(x1, x2, x3, x0, 9); \
    Round_Function_512(x0, x1, x2, x3, 8); \
    Round_Function_512(x3, x0, x1, x2, 7); \
```

```
   Round_Function_512(x2, x3, x0, x1, 6); \
   Round_Function_512(x1, x2, x3, x0, 5); \
   Round_Function_512(x0, x1, x2, x3, 4); \
   Round_Function_512(x3, x0, x1, x2, 3); \
   Round_Function_512(x2, x3, x0, x1, 2); \
   Round_Function_512(x1, x2, x3, x0, 1); \
   Round_Function_512(x0, x1, x2, x3, 0); \
} while(0)
```

A.3 NEON Implementations of Areion-256 and Areion-512

```
#include<stdint.h>
#include<arm_neon.h>

/* Round Constant aligned for little endian */
const uint32_t RC[][4] = {
    {0x03707344, 0x13198a2e, 0x85a308d3, 0x243f6a88},
    {0xec4e6c89, 0x082efa98, 0x299f31d0, 0xa4093822},
    {0x34e90c6c, 0xbe5466cf, 0x38d01377, 0x452821e6},
    {0xb5470917, 0x3f84d5b5, 0xc97c50dd, 0xc0ac29b7},
    {0x98dfb5ac, 0xd1310ba6, 0x8979fb1b, 0x9216d5d9},
    {0x6a267e96, 0xb8e1afed, 0xd01adfb7, 0x2ffd72db},
    {0xb3916cf7, 0x24a19947, 0xf12c7f99, 0xba7c9045},
    {0x1574e690, 0x36920d87, 0x58efc166, 0x801f2e28},
    {0x728eb658, 0x0d95748f, 0xf4933d7e, 0xa458fea3},
    {0xc25a59b5, 0x7b54a41d, 0x82154aee, 0x718bcd58},
    {0x286085f0, 0xc5d1b023, 0x2af26013, 0x9c30d539},
    {0x603a180e, 0x8e79dcb0, 0xb8db38ef, 0xca417918},
    {0xbd314b27, 0xd71577c1, 0xb01e8a3e, 0x6c9e0e8b},
    {0xaa55ab94, 0xe65525f3, 0x55605c60, 0x78af2fda},
    {0x2aab10b6, 0x55ca396a, 0x63e81440, 0x57489862},
    {0x7c72e993, 0xa15486af, 0x1141e8ce, 0xb4cc5c34},
    {0x741831f6, 0x2ba9c55d, 0x636fbc2a, 0xb3ee1411},
    {0x6c24cf5c, 0xafd6ba33, 0x9b87931e, 0xce5c3e16},
    {0x6b4bb9af, 0x3b8f4898, 0x28958677, 0x7a325381},
    {0xfb21a991, 0x61d809cc, 0x66282193, 0xc4bfe81b},
    {0xe98575b1, 0xef845d5d, 0x5dec8032, 0x487cac60},
    {0xd396acc5, 0x23893e81, 0xeb651b88, 0xdc262302},
    {0x48420040, 0xe0b4482a, 0x3f442392, 0xf6d6ff38},
    {0xf6e96c9a, 0x21c66842, 0x9e1f9b5e, 0x69c8f04a}
};

#define RC0 vmovq_n_u8(0)
#define RC1(i) vreinterpretq_u8_u32(vld1q_u32(RC[i]))

/* Operations for the round function */
#define A1(X, K) vaesmcq_u8((vaeseq_u8(X, K)))
#define A2(X, K) vaeseq_u8(X, K)
#define A3(X) vaesmcq_u8(X)
#define A4(X, K) vaesdq_u8(X, K)
#define XOR(X, Y) veorq_u8(X, Y)

/* Round Function for the 256-bit permutation */
#define R_FIRST(x0, x1, i) \
    do { \
        x1 = A2(A1(A1(x0, RC0), RC1(i)), x1); \
        x0 = A2(x0, RC0); \
    } while (0)

#define R_MIDDLE(x0, x1, i) \
    do { \
        x1 = A2(A1(A1(x0, RC0), RC1(i)), x1); \
    } while (0)

#define R_LAST(x0, x1, i) \
```

```
    do { \
        x1 = XOR(A1(A1(x0, RC0), RC1(i)), x1); \
        x0 = A2(x0, RC0); \
    } while (0)

/* 256-bit permutation */
#define perm256(x0, x1) \
    do { \
        R_FIRST(x0, x1, 0); \
        R_MIDDLE(x1, x0, 1); \
        R_MIDDLE(x0, x1, 2); \
        R_MIDDLE(x1, x0, 3); \
        R_MIDDLE(x0, x1, 4); \
        R_MIDDLE(x1, x0, 5); \
        R_MIDDLE(x0, x1, 6); \
        R_MIDDLE(x1, x0, 7); \
        R_MIDDLE(x0, x1, 8); \
        R_LAST(x1, x0, 9); \
    } while (0)


 /* Inversed Round Function for the 256-bit permutation */
#define Inv_R_FIRST(x0, x1, i) \
    do { \
        x0 = A4(x0, RC0); \
        x1 = A4(A1(A1(x0, RC0), RC1(i)), x1); \
    } while (0)

#define Inv_R_MIDDLE(x0, x1, i) \
    do { \
        x1 = A4(A1(A1(x0, RC0), RC1(i)), x1); \
    } while (0)

#define Inv_R_LAST(x0, x1, i) \
    do { \
        x1 = XOR(A1(A1(x0, RC0), RC1(i)), x1); \
    } while (0)

/* Inversed 256-bit permutation */
#define Inv_perm256(x0, x1) \
    do { \
        Inv_R_FIRST(x1, x0, 9); \
        Inv_R_MIDDLE(x0, x1, 8); \
        Inv_R_MIDDLE(x1, x0, 7); \
        Inv_R_MIDDLE(x0, x1, 6); \
        Inv_R_MIDDLE(x1, x0, 5); \
        Inv_R_MIDDLE(x0, x1, 4); \
        Inv_R_MIDDLE(x1, x0, 3); \
        Inv_R_MIDDLE(x0, x1, 2); \
        Inv_R_MIDDLE(x1, x0, 1); \
        Inv_R_LAST(x0, x1, 0); \
    } while (0)

 /* Round Function for the 512-bit permutation */
#define R_FIRST(x0, x1, x2, x3, i) \
    do { \
        x1 = A2(A1(x0, RC0), x1); \
        x3 = A2(A1(x2, RC0), x3); \
        x0 = A2(x0, RC0); \
        x2 = A1(A2(x2, RC0), RC1(i)); \
    } while (0)

#define R_MIDDLE(x0, x1, x2, x3, i) \
    do { \
        x1 = A2(A1((x0), x1); \
        x3 = A2(A3(x2), x3); \
```

```
        x2 = A1(x2, RC1(i)); \
    } while (0)

#define R_LAST(x0, x1, x2, x3, i) \
    do { \
        x1 = XOR(A3(x0), x1); \
        x3 = XOR(A3(x2), x3); \
        x2 = A1(x2, RC1(i)); \
    } while (0)

/* 512-bit permutation */
#define perm512(x0, x1, x2, x3) \
    do { \
        R_FIRST(x0, x1, x2, x3, 0); \
        R_MIDDLE(x3, x0, x1, x2, 1); \
        R_MIDDLE(x2, x3, x0, x1, 2); \
        R_MIDDLE(x1, x2, x3, x0, 3); \
        R_MIDDLE(x0, x1, x2, x3, 4); \
        R_MIDDLE(x3, x0, x1, x2, 5); \
        R_MIDDLE(x2, x3, x0, x1, 6); \
        R_MIDDLE(x1, x2, x3, x0, 7); \
        R_MIDDLE(x0, x1, x2, x3, 8); \
        R_MIDDLE(x3, x0, x1, x2, 9); \
        R_MIDDLE(x2, x3, x0, x1, 10); \
        R_MIDDLE(x1, x2, x3, x0, 11); \
        R_MIDDLE(x0, x1, x2, x3, 12); \
        R_MIDDLE(x3, x0, x1, x2, 13); \
        R_LAST(x2, x3, x0, x1, 14); \
    } while (0)
```

B Test Vectors

B.1 Areion-256

```
/* test vector #1 */
Input:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Output:
    28 12 a7 24 65 b2 6e 9f ca 75 83 f6 e4 12 3a a1
    49 0e 35 e7 d5 20 3e 4b a2 e9 27 b0 48 2f 4d b8

/* test vector #2 */
Input:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

Output:
    68 84 5f 13 2e e4 61 60 66 c7 02 d9 42 a3 b2 c3
    a3 77 f6 5b 13 bb 05 c7 cd 1f b2 9c 89 af a1 85
```

B.2 Areion-512

```
/* test vector #1 */
Input:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Output:
    b2 ad b0 4f a9 1f 90 15 59 36 71 22 cb 3c 96 a9
    78 cf 3e e4 b7 3c 6a 54 3f e6 dc 85 77 91 02 e7
    e3 f5 50 10 16 ce ed 1d d2 c4 8d 0b c2 12 fb 07
    ad 16 87 94 bd 96 cf f3 59 09 cd d8 e2 27 49 28

/* test vector #2 */
Input:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f

Output:
    b6 90 b8 82 97 ec 47 0b 07 dd a9 2b 91 95 9c ff
    13 5e 9a c5 fc 3d c9 b6 47 a4 3f 4d aa 8d a7 a4
    e0 af bd d8 e6 e2 55 c2 45 27 73 6b 29 8b d6 1d
    e4 60 ba b9 ea 79 15 c6 d6 dd be 05 fe 8d de 40
```

B.3 Areion256-DM

```
/* test vector #1 */
Input:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Output:
    28 12 a7 24 65 b2 6e 9f ca 75 83 f6 e4 12 3a a1
    49 0e 35 e7 d5 20 3e 4b a2 e9 27 b0 48 2f 4d b8

/* test vector #2 */
```

```
Input:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

Output:
    68 85 5d 10 2a e1 67 67 6e ce 08 d2 4e ae bc cc
    b3 66 e4 48 07 ae 13 d0 d5 06 a8 87 95 b2 bf 9a
```

B.4 Areion512-DM

```
/* test vector #1 */
Input:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Output:
    59 36 71 22 cb 3c 96 a9 3f e6 dc 85 77 91 02 e7
    e3 f5 50 10 16 ce ed 1d ad 16 87 94 bd 96 cf f3

/* test vector #2 */
Input:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f

Output:
    0f d4 a3 20 9d 98 92 f0 5f bd 25 56 b6 90 b9 bb
    c0 8e 9f fb c2 c7 73 e5 d4 51 88 8a de 4c 23 f1
```

B.5 Areion512-MD

```
/* test vector #1 */
Input:
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
    00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

Output:
    7f 22 34 44 5f 3a 72 00 65 93 79 42 01 53 6c 94
    09 5d ab d3 fd b5 84 67 48 d3 59 55 5c 52 e6 51

/* test vector #2 */
Input:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f
    40 41 42 43 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f
    50 51 52 53 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f
    60 61 62 63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f
    70 71 72 73 74 75 76 77 78 79 7a 7b 7c 7d 7e 7f

Output:
    3e 4d 31 0f be 21 d0 7b b9 00 46 88 a1 50 36 b7
    ab d9 ae 2f e9 e6 0c 9a ca 2a cc 36 98 5e 60 0b
```

B.6 Areion512-KWF

```
/* test vector #1 */
key:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

associated data:
   3f 3e 3d 3c 3b 3a 39 38 37 36 35 34 33 32 31 30
   2f 2e 2d 2c 2b 2a 29 28 27 26 25 24 23 22 21 20
   1f 1e 1d 1c 1b 1a 19 18 17 16 15 14 13 12 11 10
   0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

message:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

ciphertext:
   01 94 4a 6a f7 14 72 20 81 30 f7 5d 27 ed 10 15
   12 71 dc 0e fd 80 e3 88 71 29 0d 7d fb e0 6e ed
   3c aa 13 ec 37 1a c1 df 29 70 6f c8 6b 90 60 ee
   75 2e 21 50 4f 43 5a 7f 59 ee 5b b3 9c af 63 d5

/* test vector #2 */
key:
   0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

associated data:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

message:
   00 01 02 03 04 05 06 07 08 09 0a 0b

ciphertext:
   a4 27 e4 10 65 cd ab 1b 53 b5 f9 7b c9 ac ad 4f
   e1 f0 5b 4e 0d a4 87 30 e4 6c 55 81 cc c0 30 d5
   ad 86 9c 07 be de 72 31 50 31 ec b7 18 5b f2 f0
   bb a2 46 72 cf bb a4 fc 7e 25 48 28 85 30 fb 11
```

B.7 Areion256-OPP

```
/* test vector #1 */
key:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

nonce:
   0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

associated data:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

plaintext:
   00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
   10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
   20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
   30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f

ciphertext:
   a4 69 c0 ab 00 bf b6 8e 1f f3 74 54 b8 3d da 59
   ef 61 1b 32 30 c0 a7 f0 a7 36 7c ab 36 c8 8a 59
   d4 dc e1 ec 7e cb 9b ad b4 77 16 93 24 b9 22 b4
   ef 04 17 8a 46 58 85 10 c2 44 ae 7b 7c bc 05 a0

tag:
   76 12 8b 16 b6 cd 68 21 e3 7b df 58 69 27 61 a5
   05 dd 89 f4 cc 81 b7 c9 28 96 53 d6 83 a7 a8 a7
```

```
/* test vector #2 */
key:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

nonce:
    0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

associated data:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

plaintext:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f
    40 41 42 43 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f
    50 51 52 53 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f
    60 61 62 63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f
    70 71 72 73 74 75 76 77 78 79 7a 7b 7c 7d 7e 7f

ciphertext:
    16 d7 b2 7a 50 0a a0 3e a1 d1 79 f3 26 63 b3 b9
    e3 f0 41 b9 ba dd 0e 4d 59 f1 bf 87 82 5b 2a 30
    f9 00 11 96 fd 45 30 6d 59 86 d7 a2 57 0c 6c 8a
    df 68 8e 7e a2 0a 27 1b 61 e0 67 39 4f a2 85 5d
    e8 71 76 5c ce 79 5b 4d 81 6c 7e b3 74 b1 66 6f
    dc a1 de c1 af 22 8b bb eb 76 74 86 b8 52 08 c1
    26 f2 b2 79 87 94 0b 03 00 f6 23 27 86 55 ba 5d
    c9 db 3e bc 56 55 69 a0 f2 16 22 9d a4 a6 63 d8

tag:
    25 d9 b9 09 41 45 e6 1f f0 f5 49 be 6d fe 81 a2
    ec 7c e7 8c 8f c0 ba b0 d7 72 1b 9d 80 d4 76 f7
```

B.8 Areion256-OTR

```
/* test vector #1 */
key:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

nonce:
    0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

associated data:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f

plaintext:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f

ciphertext:
    84 48 e2 88 24 9a 90 a9 37 da 5e 7f 94 10 f4 76
    df 50 29 01 04 f5 91 c7 9f 0f 46 90 0a f5 b4 66
    41 1d fc 8b 9b 44 a3 08 7e 60 2b 50 77 d9 2c 31
    d2 22 d3 f0 dd 33 df 9f a0 0d e6 12 cb 54 89 09

tag:
    f2 4f 05 07 54 d1 aa d2 04 5f c5 39 14 ad 2b 5c
    ac 31 6c d8 04 28 91 d7 42 b0 59 fa d7 ca be ef
```

```
/* test vector #2 */
key:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

nonce:
    0f 0e 0d 0c 0b 0a 09 08 07 06 05 04 03 02 01 00

associated data:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f

plaintext:
    00 01 02 03 04 05 06 07 08 09 0a 0b 0c 0d 0e 0f
    10 11 12 13 14 15 16 17 18 19 1a 1b 1c 1d 1e 1f
    20 21 22 23 24 25 26 27 28 29 2a 2b 2c 2d 2e 2f
    30 31 32 33 34 35 36 37 38 39 3a 3b 3c 3d 3e 3f
    40 41 42 43 44 45 46 47 48 49 4a 4b 4c 4d 4e 4f
    50 51 52 53 54 55 56 57 58 59 5a 5b 5c 5d 5e 5f
    60 61 62 63 64 65 66 67 68 69 6a 6b 6c 6d 6e 6f
    70 71 72 73 74 75 76 77 78 79 7a 7b 7c 7d 7e 7f

ciphertext:
    36 f6 90 59 74 2f 86 19 89 f8 53 d8 0a 4e 9d 96
    d3 c1 73 8a 8e e8 38 7a 61 c8 85 bc be 66 14 0f
    74 45 c3 d5 4a dc de 89 6f ae c3 46 c8 99 69 b8
    45 8c 19 51 68 54 fe f6 fa 7b b6 84 02 1a 8b 2c
    21 38 98 01 2d 93 7c 50 81 18 ed 1d 8e a4 a6 ce
    3c 70 73 5c ac 9d c2 09 6a 39 62 54 72 ab 25 7d
    a8 23 9f c2 24 6f ac 0b 13 7c c7 a9 75 cf 49 ee
    0b 4f 74 e7 86 02 ff 6e 8a 7c ce d2 9b f4 e5 48

tag:
    ce 51 cc 06 29 22 f2 10 37 e1 07 ea 84 44 19 77
    e7 f1 73 73 39 94 6c a5 fb f8 3b 92 fa c6 0c 5b
```