# PAEQ: Parallelizable Permutation-based Authenticated Encryption <br> (Full Version) 

Alex Biryukov ${ }^{1}$, Dmitry Khovratovich ${ }^{1}$<br>University of Luxembourg<br>alex.biryukov,dmitry.khovratovich@uni.lu


#### Abstract

We propose a new authenticated encryption scheme PAEQ, which employs a fixed public permutation. In contrast to the recent sponge-based proposals, our scheme is fully parallelizable. It also allows flexible key and nonce length, and is one of the few which achieves 128 -bit security for both confidentiality and data authenticity with the same key length. The permutation within PAEQ is a new design called AESQ, which is based on AES and is 512 bits wide. In contrast to similar constructions used in the SHA-3 competition, our permutation fully benefits from the newest Intel AES instructions and runs at 2.5 cycles per byte if used as the counter-mode PRF.


Keywords: Authenticated encryption, SHA-3, permutation-based cryptography, AES.

## 1 Introduction

It has been known for a while that standard blockcipher modes of operation do not provide any integrity/authenticity protection, and hence additional mechanisms are needed to ensure the receiver that the ciphertext has not been modified of generated by the adversary. Whenever two distinct keys are available, the problem can be solved with a simple combination of encryption and MAC generation [4, and the Encrypt-then-MAC paradigm has become an international standard [25]. In contrast to confidentiality-only modes, the authenticated encryption schemes do not provide any decryption but return "invalid" $(\perp)$, if the ciphertext has been created or modified by an adversary.

Since at least the year of 2000, cryptographers have tried to design an authenticated encryption (AE) scheme, which would use a single key and would be at least as efficient as Encrypt-then-MAC. The research went in two directions. The first one deals with new modes of operation which use an arbitrary block cipher. The ISO standards 25] GCM, CCM, and OCB are typical examples. The patented OCB mode runs almost as fast as the counter mode of encryption, which yields the speed below one cycle per byte on modern CPUs if used with AES (19). The second direction is to design a dedicated AE scheme, like Phelix [26] ALE [9, or AEGIS [28].

Modern authenticated encryption schemes are also able to authenticate so called associated data (AD) without encrypting it [24]. A typical application is Internet packets, whose contents are encrypted, whereas headers are not (for routing purposes), while they still should be bound to the encrypted data.

Features of modern authenticated encryption schemes AE schemes provide larger functionality compared to regular encryption modes, and the list of features has been expanding for a while. Let us comment on the most important properties of such schemes:

- Performance (both software and hardware). If the scheme uses AES as underlying block cipher, it can use special instructions on modern CPUs. The resulting encryption speed ranges from 0.6 to 4 cycles per byte [19].
- Security level. Whereas AES supports keys up to 256 bits long, the existing security proofs guarantee little security beyond the half of the block size ( 64 bits for all AES variants). Some dedicated schemes provide higher level of security, but they have not been subject of third-party cryptanalysis or have been broken. Some modes (e.g., GCM) have non-trivial security limitations because of tag length (security drops too fast for small tags). Whereas 64 -bit security seems fine for now, it might not be so for the distant future.
- The existence of security proofs closely relates to the previous feature. Modes of operation usually enjoy reasonable, but lengthy and complicated proofs of security in the standard model assuming that the underlying block cipher is a secure PRP, or proofs in the random oracle model, which assume absence of non-trivial properties of the underlying fixed permutation. Dedicated designs usually have little to no security proofs, but provide some evidence of resisting to certain types of attacks. As lengthy proofs are subject to bugs [16], it is desirable to have them compact and verifiable by the third parties.
- Parallelism. If a lot of resources are available, one may desire to run the encryption, decryption, or verification in parallel. The modes that use chaining (like the ones derived from the CBC encryption or the sponge construction) are difficult to parallelize.
- Tag update. Most of the schemes, with a few exceptions like GCM, require to recompute almost entire ciphertext if a small portion of plaintext is modified. If the ciphertext can be updated quickly, it would allow much faster processing of large amounts of encrypted data, e.g., hard-drive encryption.
- Variable key, nonce, or tag length. All the three parameters are usually restricted by the application that uses an AE scheme. In turn, AE schemes have their own, sometimes incompatible restrictions. The more variability the scheme has, the more applications it suits.
- Processing of associated data (AD). All the modern schemes allow for the authentication of associated data, which is not encrypted. Some of them, however, can not preprocess AD before the plaintext is over, which might be a penalty on the performance.

Blockcipher- vs. permutation-based designs The most generic AE schemes employ blockciphers. However, they are not the only source of good transformations. If the mode is encryptiononly (see above), the transformation does not need to be invertible (cf. the folklore use of hash functions for the CTR mode). Quite recently, the hash function Keccak (7), which employs a 1600 -bit permutation, has been selected as the new SHA-3 standard. We expect that it will be widely deployed in the near future, and hence its building block will be readily accessible to other cryptographic applications. On resource-constrained devices, where the space is limited, it would be very tempting to use a single cryptographic primitive, such as the Keccak permutation, for many purposes. Whenever Keccak or AES are considered too expensive for a device, the lightweight hash functions like Spongent [8] and Quark [3] are also based on a single permutation and may offer it for other schemes. A wide permutation also simplifies the security proofs, as additional inputs such as tweaks and counters can be easily accommodated within the permutation input.

This idea also fits the recent paradigm of the permutation-based cryptography 10 as opposed to the blockcipher-based cryptography. From the practical point of view, it would allow to have a single permutation for all purposes, whereas it would simplify the analysis as a target for a cryptanalyst would be much simpler. The downside of the permutation-based approach is that the security proof has to be devised in the random oracle/permutation model, and does not rely on the PRP assumption.

Our contributions. We offer a new mode, called PPAE, which employs a public fixed permutation (let us denote its bit width by $n$ ). We have tried to make the mode as universal as possible, and to provide the users with almost every capability an authenticated encryption mode might have. We summarize the properties of our scheme in Table 1.

Alongside the mode, we propose an AES-based permutation called $A E S Q$, which is 512-bit wide and has been optimized for recent CPU with AES instructions. It can be used in other permutation-based constructions, e.g. the extended Even-Mansour cipher or the sponge construction. It has a large security margin against the most popular attacks on permutation-based schemes: differential-based collision search, rebound attack, and meet-in-the-middle attack. The authenticated encryption scheme $P A E Q$ is the instantiation of PPAE with AESQ.

The wide permutation allows to get a much higher security level compared to the other AESbased designs (AES-GCM, AES-CBC+HMAC-SHA-256, OCB, COPA [1], COBRA, OTR). These schemes can not deliver a security level higher than 64 bits due to the birthday phenomena at the 128 -bit AES state. In contrast, PAEQ easily brings the security level of 128 bits and higher. We note that this security level assumes a nonce-respecting adversary, who does not repeat nonces in the encryption requests.

PAEQ allows to encrypt and decrypt the data on the arbitrary number of subprocessors with tiny amount of shared memory. Those processors, threads, or other computation units may perform decryption and encryption of incoming blocks in any order.

PAEQ is based on the AES block cipher, and does not use any operations except those of AES. The mode of operation around the AES-based permutation AESQ uses only XORs and counter increments, and aims to be amongst easiest authenticated encryption modes to implement.

The paper is structured as follows. We provide a formal syntax of authenticated encryption modes and describe PPAE and PAEQ in this context in Section 2. We discuss its performance in Section 3 and give the design rationale in Section 4. The security aspects are discussed in Sections 5 and 6.

| Key/nonce/tag length | We allow keys, nonces, and tags of arbitrary length, as long as they fit into the <br> permutation, fulfill some minimal requirements, and constitute the integer number of <br> bytes. |
| :--- | :--- |
| Performance | Depending on the key size, the encryption speed is about 6 cycles per byte on modern <br> CPUs with permutation AESQ. |
| Security level | Depending on the key length and the permutation width, we support a range of <br> security levels from 64 to 256 bits. For the permutation of width $n$ bits we can use <br> a key of about $(n / 3-6)$ bits and get the same security level for both confidentiality <br> and ciphertext integrity. Hence a permutation of 400 bits width already delivers a <br> security level of 128 bits. |
| Security proof | Our mode is provably secure in the random permutation model, whereas the security <br> proof is short and verifiable by the third parties. |
| Parallelism | Our scheme is fully parallelizable: all blocks of plaintext and associated data can be <br> processed in parallel; only the last call of the permutation needs all the operations to <br> finish. |
| Online processing | Our scheme is fully online, being able to process plaintext blocks or blocks of associ- <br> ated data as soon as they are ready without knowing the final length. |
| Patents | It is not patented, and we are not aware of any patent covering any part of the <br> submission. |
| Tag update | If the tag is not truncated, then the last permutation call can be inverted given the <br> key, and only two extra permutation calls are needed to encrypt and authenticate a <br> new plaintext with one new block. |
| Inverse | The permutation is used in the forward direction only, with a sole exception of tag <br> update, if this feature is needed. |
| Nonce misuse | If the nonce is reused, then the integrity is still provided. Additionally, a user may <br> generate a nonce out of the key, the plaintext, and the associated data with a dedicated <br> routine. |

Table 1. PPAE/PAEQ features.

## 2 PAEQ as Authenticated Encryption Scheme

### 2.1 Notation

The authenticated encryption scheme is denoted by $\Pi$ and is defined as a pair of functions $\mathcal{E}$ and $\mathcal{D}$, which provide encryption and decryption, respectively. The inputs to $\mathcal{E}$ are a plaintext $P \in \mathcal{P}$, associated data $A \in \mathcal{A}$, a nonce $N \in \mathcal{N}=\{0,1\}^{r}$, a key $K \in \mathcal{K}=\{0,1\}^{k}$, and a tag length $\tau . \mathcal{P}$ is the set of byte strings with length between 1 and $2^{96}$, and $A$ is the set of byte strings with length between 0 and $2^{96}$. The encryption function outputs ciphertext $C \in \mathcal{C}$ and $\operatorname{tag} T \in \mathcal{T}=\{0,1\}^{\tau}$ :

$$
\mathcal{E}: \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{P} \rightarrow \mathcal{C} \times \mathcal{T}
$$

For fixed ( $N, A, K$ ) the encryption function is injective and hence defines the decryption function. The decryption function takes a key, a nonce, associated data, the ciphertext, and the tag as input and returns either a plaintext or the "invalid" message $\perp$ :

$$
\mathcal{D}: \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C} \times \mathcal{T} \rightarrow \mathcal{P} \cup\{\perp\} .
$$

For brevity we will write $\mathcal{E}_{K}^{N, A}(\cdot)$ and $\mathcal{D}_{K}^{N, A}(\cdot)$.
To accomodate nonce-misuse cases, the user may choose to generate $N$ as a keyed function of plaintext, key, and associated data. It is then called extra nonce and denoted by $N_{e}$. The user is then supposed to communicate $N_{e}$ to the receiver, with ordinary nonce transmission rules applied.

Let $X$ be an internal $s$-bit variable. Then we refer to its bits as $X[1], X[2], \ldots, X[s]$ and denote the subblock with bits from $s_{1}$ to $s_{2}$ as $X\left[s_{1} \ldots s_{2}\right]$. In pictures and formulas the least' significant bits and bytes are at the left, and the concatenation of multi-bit variables is defined as follows:

$$
X \| Y=X[1] X[2] \ldots X[t] Y[1] Y[2] \ldots Y[s] .
$$

The counter values have their least significant bits as the least significant bits of corresponding variables.

PAEQ is a concrete instantiation of generic mode PPAE, which takes permutation $\mathcal{F}$ of width $n$.

### 2.2 PPAE mode of operation

For the domain separation the following two-byte constants are used:

$$
D_{i}=(k,(r+i) \quad(\bmod 256)), i=1,2,3,4,5,6,
$$

where the second value is taken modulo 256 .
We refer to Figure 1 for a graphical illustration of the PAEQ functionality.

## Encryption.

1. During the first stage the plaintext is split into blocks $P_{i}$ of $(n-k-16)$ bits $(n / 8-k / 8-2$ bytes) $P_{1}, P_{2}, \ldots, P_{l}$. We encrypt block $P_{i}$ as follows:

$$
\begin{aligned}
V_{i} & \leftarrow D_{0}\left\|R_{i}\right\| N \| K ; \\
W_{i} & \leftarrow \mathcal{F}\left(V_{i}\right) ; \\
C_{i} & \leftarrow W_{i}[17 . .(n-k)] \oplus P_{i} .
\end{aligned}
$$

Here $V_{i}, W_{i}$ are $n$-bit intermediate variables, and $C_{i}$ is a ( $n-k-16$ )-bit block. The counter $R_{i}=i$ occupies the ( $n-k-r-16$ )-bit block. If last plaintext block is incomplete and has length $t^{\prime}$, then $D_{0}$ is replaced with $D_{1}$.


Fig. 1. Encryption and authentication with PAEQ.
2. During the second stage we compute intermediate variables for authentication:

$$
\begin{aligned}
X_{i} & \leftarrow D_{2}\left\|C_{i}\right\| W_{i}[(n-k+1) . . n] ; \\
Y_{i} & \leftarrow\left(\mathcal{F}\left(X_{i}\right)\right)[17 . .(n-k)] .
\end{aligned}
$$

If the last plaintext block $P_{t}$ is incomplete, then we define $P_{t}^{\prime}=P_{t}\|a\| a\|\ldots\| a$, where $a$ is one-byte variable with value equal to $t^{\prime}$, the byte length of $P_{t}$, and define

$$
\begin{aligned}
& C_{t}^{\prime} \leftarrow W_{t}[17 . .(n-k)] \oplus P_{t}^{\prime} ; \\
& X_{t} \leftarrow D_{3}\left\|C_{t}^{\prime}\right\| W_{i}[(n-k+1) . . n] .
\end{aligned}
$$

Finally,

$$
Y \leftarrow \bigoplus_{i} Y_{i}
$$

3. During the third stage we compute an intermediate variable for authenticating associated data. The AD is splitted into blocks $A_{1}, A_{2}, \ldots, A_{s}$ of length $n-2 k-16$. Then we compute

$$
\begin{aligned}
X_{i}^{\prime} & \leftarrow D_{4}\left\|R_{i}\right\| A_{i} \| K ; \\
Y_{i}^{\prime} & \leftarrow\left(\mathcal{F}\left(X_{i}^{\prime}\right)\right)[17 . .(n-k)],
\end{aligned}
$$

where $R_{i}$ is a $k$-bit counter starting with 1 . If the last AD block is incomplete (has $t^{\prime \prime}$ bytes), it is padded with bytes whose value is $t^{\prime \prime}$, and the constant $D_{4}$ is changed to $D_{5}$. Finally,

$$
Y^{\prime} \leftarrow \bigoplus_{i} Y_{i}^{\prime}
$$

4. In the final stage we compute the tag. First, we define

$$
Z \leftarrow Y \bigoplus Y^{\prime}
$$

and then,

$$
T \leftarrow \mathcal{F}\left(D_{6}\|Z\| K\right) \oplus\left(0^{n-k} \| K\right),
$$

where $0^{n-k}$ stands for $(n-k)$ zero bits.
The $\operatorname{tag} T$ is truncated to $T[1 \ldots \tau]$.
The encryption and authentication process is illustrated in Figure 1
Decryption. The decryption process repeats the encryption process with minor corrections. We decrypt as

$$
P_{i} \leftarrow W_{i}[17 . .(n-k)] \oplus C_{i}
$$

with appropriate corrections for the incomplete block if needed. $C_{i}$ is used as given when composing $X_{i}$. Finally, the tag $T^{\prime}$ is computed and matched with submitted $T$. If $T \neq T^{\prime}$ (including length mismatch), the decryption process returns invalid.

The plaintext is returned only if the tags match.

### 2.3 The AESQ permutation

The $\mathcal{F}$ permutation in PAEQ is the AESQ permutation, which is defined below. AESQ operates on 512 -bit inputs, which are viewed as four 128 -bit registers. The state undergoes 20 identical rounds. The rounds use standard AES operations: SubBytes, ShiftRows, and MixColumns, which are applied to the 128 -bit registers exactly as in AES (Figure 5).

```
Input: 128-bit states \(A, B, C, D\), round constants \(Q_{i, j, k}\)
for \(0 \leq i<R=10\) do
    for \(0 \leq j<2\) do
        \(A \leftarrow\) MixColumns ○ ShiftRows ○ SubBytes \((A)\);
        \(A_{0} \leftarrow A_{0} \oplus Q_{i, j, 1} ;\)
        \(B \leftarrow\) MixColumns o ShiftRows o SubBytes \((B)\);
        \(B_{0} \leftarrow B_{0} \oplus Q_{i, j, 2} ;\)
        \(C \leftarrow\) MixColumns o ShiftRows o SubBytes \((C)\);
        \(C_{0} \leftarrow C_{0} \oplus Q_{i, j, 3}\);
        \(D \leftarrow\) MixColumns o ShiftRows o SubBytes \((D)\);
        \(D_{0} \leftarrow D_{0} \oplus Q_{i, j, 4} ;\)
    end
    \((A, B, C, D) \leftarrow \operatorname{Shuffle}(A, B, C, D)\)
end
```

Algorithm 1: Pseudocode for the AESQ permutation with $2 R$ rounds

The round constants are chosen as follows in the matrix register view:

$$
Q_{i, j, k}=\left(\begin{array}{cccc}
8 i+4 j+k & 8 i+4 j+k & 8 i+4 j+k & 8 i+4 j+k \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0
\end{array}\right)
$$

Here is the Shuffle mapping that permutes columns of the internal states:

|  | $A$ |  |  |  | $C$ |  |  | $C$ |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| From | $A[0]$ | $A[1]$ | $A[2]$ | $A[3]$ | $B[0]$ | $B[1]$ | $B[2]$ | $B[3]$ | $C[0]$ | $C[1]$ | $C[2]$ | $C[3]$ | $D[0]$ | $D[1]$ | $D[2]$ | $D[3]$ |
| To | $A[3]$ | $D[3]$ | $C[2]$ | $B[2]$ | $A[1]$ | $D[1]$ | $C[0]$ | $B[0]$ | $A[2]$ | $D[2]$ | $C[3]$ | $B[3]$ | $A[0]$ | $D[0]$ | $C[1]$ | $B[1]$ |

### 2.4 Extra nonce

The extra nonce $N_{e}$ is a function of the key, the plaintext, and the associated data. It is not online, i.e. it needs to know the length of all inputs and the output. It is the sponge hash function:

$$
T_{1}=\mathcal{F}\left(Q_{1} \| 0^{2 k}\right), T_{2}=\mathcal{F}\left(T_{1} \oplus\left(Q_{2} \| 0^{2 k}\right)\right), \ldots, T_{m}=\mathcal{F}\left(T_{m-1} \oplus\left(Q_{m} \| 0^{2 k}\right)\right),
$$

and $N_{e}$ is the truncation of $T_{m}$ to $r$ bits. The injection blocks $Q_{i}$ come from the string $Q$, which is composed as follows:

$$
Q=|P|\||A|\| k_{b}\left\|r_{b}\right\| K\|P\| A \| 10^{*} 1,
$$

where $|P|$ is the plaintext length in bytes, $|A|$ is the associated data length in bytes, $k_{b}$ is the key length in bytes, $r_{b}$ is the nonce length in bytes, and $10^{*} 1$ is the sponge padding: one byte with value 1 , then as many zero bytes as needed to fill all but one bytes in the injection block, and then the byte with value 1 .

## 3 Performance

### 3.1 Benchmarks for PAEQ

Now we provide speed benchmarks for PAEQ. On the recent Haswell CPU family we obtain the speed of AESQ as 5 cycles per byte (cpb) for the 64 -bit key, 6 cpb for the 128 -bit key, and 9 cpb for the 256 -bit key.

| Security level / Key length | PAEQ (20 rounds, cycles per byte) |
| :---: | :---: |
| 64 | 4.9 |
| 80 | 5.1 |
| 128 | 5.8 |
| 256 | 8.9 |

### 3.2 Implementing AESQ

Our permutation is best suited for the last Intel and AMD processors equipped with special AES instructions. Of the AES-NI instruction set, we use only AESENC instruction that performs a single round of encryption. More precisely, $\operatorname{AESENC}(S, K)$ applies ShiftRows, SubBytes, and MixColumnsto $S$ and then XORs the subkey $K$. In our scheme the subkey is a round constant. For two rounds of AESQ, we use the following instructions:

- 8 aesenc instructions for AES round calls;
- 8 vpunpckhdq instructions to permute the state columns between registers;
- 8 vpaddq instructions to update round constants.

This gives a total of 24 instructions per 64 bytes of the state in two rounds. This means that the full $2 R$-round AESQ permutation needs $24 R$ instructions, and theoretically may run in $3 R / 8$ cycles per byte if properly pipelined. This speed may be achieved in practice, since the throughput of AES instructions is 1 cycle starting from the Sandy Bridge architecture (2010), and the vpaddq instruction may be even faster. The following strategy gives us the best performance on the newest Haswell architecture:

- Process two 512 -bit states in parallel in order to mitigate the latency of 8 cycles of the AESENC instruction. Store each state into $4 \times \mathrm{mm}$ registers.
- Store round constants in $4 \times \mathrm{mm}$ registers and use them in the two parallel computations.
- Use one $x m m$ register to store the constant 1 to update the round constants, and one temporary register for the Shuffle operation.

We note that on the earlier Westmere architecture the AESENC instruction has latency 6 and throughput 2, hence the AES round calls should be interleaved with mixing instructions conducted on several parallel states.

We have made our own experiments and concluded that the 16 -round version of AESQ ( $R=8$ ) runs at 2.4 cycles per byte on a Haswell-family CPU, whereas the 20 -round version runs at 3 cycles per byte. This can be compared to the speed of the Keccak-1600 permutation. As eBASH reports, on a Haswell CPU the Keccak hash function with rate 1088 runs at 10.6 cpb , which implies that the full 1600 -bit permutation runs at approximately $\frac{10.6 \cdot 1088}{1600}=7.2 \mathrm{cpb}$. Therefore, the 20 -round AESQ is 2.5 times as fast as Keccak-1600.

## 4 Design rationale

### 4.1 Design of PPAE

When creating the new scheme, we pursued the following goals:

- Offer high security level, up to 128 bits, ideally equal to the key length,
- Make the mode of operation simple enough to yield compact and reliable security proofs.
- Deliver as many features as possible.

To achieve these goals, we decided to trade performance for clarity and verifiability.
Existing block ciphers were poor choice for these goals. They commonly have a 128 -bit block, which almost inevitably results in the loss of security at the level of $2^{64}$ cipher calls. The 256 -bit cipher Threefish could have been used, but the lack of cryptanalysis in the single-key model makes it a risky candidate. Our mode of operation would be also restricted to a single cipher.

Instead, we constructed a permutation-based mode, which takes a permutation of any width if it is at least twice as large as the key. This choice makes the scheme much more flexible, allows for variable key and nonce length, and simplifies the proof. The key update also becomes very easy. The downside of the permutation-based approach is that the security proof has to be devised in the random oracle/permutation model, and does not rely on the PRP assumption.

This is inevitable, but the success of the sponge-based constructions tells that it is not necessarily a drawback.

For the encryption stage we have chosen an analogy with the CTR mode, so that we do not have to use the permutation inverse. It also allows us to truncate some parts of the intermediate variables. For the authentication stage we use a parallel permutation-based construction. It takes the yet unused secret input from the encryption stage, which provides pseudo-randomness.

It remains to choose a permutation. Initially we thought of using a family of permutations with different widths. Examples could be Keccak [7], Spongent [8], or Quark [3] permutations. However, the performance loss would be too high given two invocations of a permutation per plaintext block. Instead, we designed our own permutation which shows the best performance on modern CPUs. It can be used in other permutation-based constructions, e.g. the extended Even-Mansour cipher or the sponge construction.

### 4.2 Design of AESQ

When designing AESQ, we needed a permutation wide enough to accomodate 128 -bit keys and nonces. The AES permutation would be too short, while AES-based permutations used in the SHA-3 context would be too large or not well optimized for AES instructions on modern CPUs.

We decided to run 4 AES states in parallel and regularly shuffle the state bytes. Since two AES rounds provide full diffusion, the shuffle should occur every two rounds. The shuffle operation should make each state to affect all four states, resembling the ShiftRows transformation in AES. The recent Intel processors, along with dedicated AES instructions, provide a set of instructions that interleave the 32 -bit subwords of 256 -bit registers. Those subwords are columns of the AES state, so we shuffled the columns. The shuffle function in this submission is one of the permutations that provide full diffusion and needs the minimal 8 number of processor instructions.

## 5 Security of PPAE

In this section we provide the security analysis of the PPAE mode of operation, which takes a permutation $\mathcal{F}$ of width $n$. Though PAEQ fixes the permutation to AESQ and its width to 512 bits, the following proof is useful when defining other scheme within the PPAE mode.

The security of a AE scheme is defined as the inability to distinguish between the two worlds, where an adversary has access to some oracles and a permutation. One world consists of the encryption oracle $\mathcal{E}_{K}(\cdot, \cdot, \cdot)$ and decryption oracle $\mathcal{D}_{K}(\cdot, \cdot, \cdot)$, where the secret key is randomly chosen and shared. The second world consists of the "random-bits" oracle $\$(\cdot, \cdot, \cdot, \cdot)$ and the "always-invalid" oracle $\perp(\cdot, \cdot, \cdot$,$) . In addition, the adversary and all the oracles have an oracle$ access to the permutation $\mathcal{F}$. The encryption requests must be nonce-respecting, i.e. all the nonces in those requests must be distinct. A decryption request ( $N, A, C, T$ ) shall not contain the ciphertext previously obtained with $(N, A)$, but it may repeat nonces with other ciphertexts.

We give a security proof for a fixed key length $k$, nonce length $r$, and random permutation $\mathcal{F}$ of width $n$. The variable-length security proof is not given in this paper, but can be mounted thanks to the $D_{i}$ constants. The proof also does not take into account the calls to $\mathcal{F}$ made in the extra nonce computation.

### 5.1 Confidentiality

We prove confidentiality as indistinguishability of the pair (ciphertext, tag) from a random string for the tag length $n$. The result for truncated tags is a trivial corollary. Let $\Pi[\mathcal{F}]$ be the PPAE mode instantiated with permutation $\mathcal{F}$. We do not consider the decryption oracle here, as we later show that with overwhelming probability it always returns $\perp$.
 be the total number of queries to $\mathcal{F}$ made during the calls to the $\Pi$ oracle, and $\sigma_{\mathcal{F}}$ be the total number of queries to $\mathcal{F}$ and $\mathcal{F}^{-1}$ oracles together. Then his advantage of distinguishing the oracle $\Pi[\mathcal{F}]$ from the random-bits oracle $\$$ is upper bounded as follows:

$$
\begin{equation*}
\operatorname{Adv}_{\Pi}^{\operatorname{conf}}(\mathcal{A}) \leq \frac{\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2}}{2^{n}}+\frac{2 \sigma_{\Pi}^{2}}{2^{n-k-16}}+\frac{2 \sigma_{\mathcal{F}}}{2^{k}}+\frac{2 \sigma_{\mathcal{F}} \sigma_{\Pi}}{2^{n-16}} \tag{1}
\end{equation*}
$$

Proof. We denote $i$-th ciphertext by $C^{(i)}$ and its $j$-th block by $C_{j}^{(i)}$. The same notation is used for intermediate variables $X, Y, Z$.

In the games $G_{i}$ below, $p_{i}$ is the probability that $\mathcal{A}$ outputs 1 . The following transitions between games are made to prove the high probability of the following properties: first, that $\Pi$ and the adversary ask non-intersecting sets of queries to $\mathcal{F}$ or $\mathcal{F}^{-1}$, and secondly, that the ciphertexts and the tags are composed of random and independent blocks and hence are indistinguishable from random strings.

Game $G_{0}$. The permutation $\mathcal{F}$ is randomly chosen. The adversary interacts with $\Pi, \mathcal{F}$, and $\mathcal{F}^{-1}$ 。

Game $G_{1}$. The permutation $\mathcal{F}$ is still randomly chosen but is defined in a lazy manner: we maintain a table $T_{\mathcal{F}}=\{(x, \mathcal{F}(x))\}$ which is initially empty, and fill the table values whenever $\mathcal{F}$ or $\mathcal{F}^{-1}$ are queried alone or within $\Pi$.

Game $G_{2}$. We modify oracles $\mathcal{F}$ or $\mathcal{F}^{-1}$ so that when a new value has to be added to the table, it is chosen completely at random from $\{0,1\}^{n}$. If this choice creates a collision, i.e. either $\mathcal{F}$ or $\mathcal{F}^{-1}$ is no longer a permutation, then we explicitly abort. The probability of such an abort is upper bounded by $\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2} / 2^{n}$.

Game $G_{3}$. We now require $\Pi$ to explicitly abort if there is a collision in the set $\left\{X_{j}^{(i)}[17 \ldots(n-\right.$ $k)]\}_{i, j}$ or in the set $\left\{T[1 \ldots(n-k)]^{(i)}\right\}$. The probability of this event is upper bounded by $\sigma_{\Pi}^{2} / 2^{n-k-16}$.

Game $G_{4}$. We now require that the adversary does not ask the oracle $\mathcal{F}$ with a query ending with $K$. Since the outputs of $\mathcal{F}^{-1}$ are completely random, the adversary has no choice but to guess $K$. Hence the probability of this event is upper bounded by $\sigma_{\mathcal{F}} / 2^{k}$.

Game $G_{5}$. We now require the oracle $\mathcal{F}^{-1}$ to explicitly abort if its reply ends with $K$. The adversary can guess the unknown parts of $X$ or $T \oplus K$, respectively (he can not hope for collisions in $X$ or $T$ due to the transition in game $G_{3}$ ). The probability of this event for all cases is upper bounded by $\sigma_{\mathcal{F}} / 2^{k}$.

Game $G_{6}$. We require $\Pi$ to abort if it produces such $X_{j}^{(i)}$ that it is already in $T_{\mathcal{F}}$. The probability of this event is upper bounded by $\sigma_{\mathcal{F}} \sigma_{\Pi} / 2^{n-16}$. We note that the adversary can not manipulate $X_{j}$ with $M_{j}$ since the choice is made before the $X_{j}$ is generated.

Game $G_{7}$. In the same manner, we require $\mathcal{F}$ to abort if it is queried with an already generated $X_{j}^{(i)}$ and $\mathcal{F}^{-1}$ to abort if it is queried with an already generated $X_{j}^{(i)}$. The probability of this event is upper bounded by $\sigma_{\mathcal{F}} \sigma_{\Pi} / 2^{n-16}$.

Game $G_{8}$. Finally, we require $\Pi$ to abort if there is a collision in the set $\left\{Z^{(i)}\right\}_{i}$. Since every $Z$ is the XOR of uniformly generated $Y_{i}$ and $Y_{i}^{\prime}$, the probability of this event is upper bounded by $\sigma_{I}^{2} / 2^{n-k-16}$.

Let us now look at game $G_{8}$. The oracle $\Pi$, in fact, never asks the oracle $\mathcal{F}$ with a query that corresponds to an existing entry in $T_{\mathcal{F}}$. We ruled out such an event in games $G_{4}, G_{5}$ for the first layer of encryption and the layer of the AD processing. We ruled out the event in games $G_{6}, G_{7}$ for the second layer of encryption. Finally, we ensure this for the authentication level in game $G_{8}$. The game $G_{3}$ is auxiliary.

The indistinguishability of the pair (ciphertext,tag) from a random string follows from the following facts:

- all the $V_{j}^{(i)}$ inputs are distinct (nonce-misuse requirement);
- hence, all the values $W_{j}^{(i)}$ are chosen uniformly in Games $G_{2}-G_{8}$;
- all the $Z^{(i)}$ values are distinct (Game $G_{8}$ );
- all the tags $T$ are chosen uniformly (and independently of ciphertexts) in Game $G_{8}$.

Therefore, Game $G_{8}$ is indistinguishable from the interaction with the random-bit oracle. The adversary's advantage in Game $G_{0}$ is upper bounded by the sum of transition probabilities in Games $G_{1}-G_{8}$, which yields Equation (1).

This concludes the proof.

### 5.2 Ciphertext integrity

Theorem 2. Let an adversary be in the setting of Theorem 1. Suppose he makes $q$ decryption requests to $\Pi$. Then his advantage of distinguishing $\Pi^{-1}$ from the "always-invalid" oracle $\perp$ is upper bounded as follows:

$$
\begin{equation*}
\operatorname{Adv}_{\Pi}^{\mathrm{int}}(\mathcal{A}) \leq \operatorname{Adv}_{I I}^{\text {conf }}(\mathcal{A})+\frac{\sigma_{\Pi}^{2}}{2^{n-k-16}}+\frac{q}{2^{\tau}}+\frac{\sigma_{\Pi} q}{2^{n-k-16}}+\frac{q}{2^{k}} \tag{2}
\end{equation*}
$$

Proof. We first need to strengthen our setting. With probability $\operatorname{Adv}_{I I}^{\text {conf }}(\mathcal{A})$ we are in Game $G_{8}$. Let us make additional computations in Game $G_{9}$.

Let $M$ be a plaintext to encrypt, and $A$ be the associated data, $Y(N, M)$ and $Y^{\prime}(A)$ be intermediate variables. Then for each encryption request $(N, M, A)$ we also compute $Y\left(N, M^{\prime}\right)$ and $Y^{\prime}\left(A^{\prime}\right)$, for every full-block prefix $M^{\prime}$ of $M$ and every full-block prefix $A^{\prime}$ of $A$. We store all the resulting $Y$ and $Y^{\prime}$ in special tables $Y_{\text {ext }}$ and $Y_{\text {ext }}^{\prime}$. Let us note that the size of both tables is upper bounded by $\sigma_{\Pi}$.

Game $G_{9}$. We additionally require $\Pi$ to explicitly abort if there are collisions in the set $\left\{Z_{\text {ext }}=\right.$ $\left.Y \oplus Y^{\prime} \mid Y \in Y_{e x t}, Y^{\prime} \in Y_{e x t}^{\prime}\right\}$. The probability of this event is upper bounded by $\sigma_{I}^{2} / 2^{n-k-16}$.

Lemma 1. Suppose we are in Game $G_{9}$. A valid decryption ( $N, A, C, T$ ) request does not generate $\perp$ with probability upper bounded by

$$
2^{-\tau}+\sigma_{\Pi} / 2^{n-k-16}+1 / 2^{k} .
$$

The proof can be found in Appendix A.
Now we are ready to prove Theorem 2. The requirements of Game $G_{9}$ are fulfilled with probability $\operatorname{Adv}_{I I}^{\text {conf }}(\mathcal{A})+\frac{\sigma_{I}^{2}}{2^{n-k-16}}$. The first decryption request returns $\perp$ with probability lower bounded in Lemma 1. Then the adversary learns nothing with his request, and we can ignore filling the table $T_{\mathcal{F}}$. All the other requests have the same lower bound to be invalid, so we merely multiply this probability by $q$.

This concludes the proof of Theorem 22


Fig. 2. Rebound attack on AESQ.

### 5.3 Basic robustness to nonce misuse

In this section we investigate the level of robustness of PAEQ to nonce misuse, i.e. to the multiple use of the same nonce(s) for encryption.

First, we note that confidentiality is almost destroyed in the case of nonce reuse. In fact, it is easy for any person to compute a new ciphertext $C^{\prime}$ for message $M$ under the same nonce $N$, as

$$
C^{\prime}=M \oplus M^{\prime} \oplus C
$$

However, with high probability the tag $T$ is still indistinguishable from the random string, which allows us to prove the integrity property in the case of nonce reuse, though with a worse bound.

Lemma 2. Suppose adversary $\mathcal{A}$ has access to $\Pi[\mathcal{F}], \mathcal{F}$, and $\mathcal{F}^{-1}$, and let $K \stackrel{\$ \mathcal{K} \text {. Let } \sigma_{\Pi}, ~\left(\frac{1}{\leftarrow}\right)}{ }$ be the total number of queries to $\mathcal{F}$ made during the calls to the $\Pi$ oracle, and $\sigma_{\mathcal{F}}$ be the total number of queries to $\mathcal{F}$ and $\mathcal{F}^{-1}$ oracles together. The adversary is allowed to repeat nonces, but not to repeat plaintexts.

Then his advantage of distinguishing the tag $T$ produced by the oracle $\Pi[\mathcal{F}]$ from the randombits oracle $\$$ is upper bounded as follows:

$$
\begin{equation*}
\operatorname{Adv}_{\Pi}^{\mathrm{conf}}(\mathcal{A}) \leq 2 \frac{\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2}}{2^{k}} \tag{3}
\end{equation*}
$$

The proof can be found in Appendix B, Now we can prove the same bounds for integrity.
Theorem 3. Let an adversary be in the setting of Lemma 2. Suppose he makes $q$ decryption requests to $\Pi$. Then his advantage of distinguishing $\Pi^{-1}$ from the "always-invalid" oracle $\perp$ is upper bounded as follows:

$$
\begin{equation*}
\operatorname{Adv}_{\Pi}^{\operatorname{int}}(\mathcal{A}) \leq 2 \frac{\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2}}{2^{k}} \tag{4}
\end{equation*}
$$

Proof. The proof directly follows from Lemma 2. Since the adversary can not distinguish a tag from a random string under repeating nonces, she can not create a forgery. Indeed, suppose that an adversary $B$ has advantage $\alpha$ of distinguishing $\Pi^{-1}$ from the "always-invalid" oracle for a single query, i.e. she can guess the tag with probability $\alpha$. Then $B$ can distinguish $\Pi$ from the random-bit oracle with the same probability, which is in turn bounded in Lemma 2.

### 5.4 Extra nonce

The extra nonce $N_{e}$ is a function of $P, K$, and $A$. It invokes a sponge construction, which is proved indifferentiable from random oracle up to $2^{c / 2}$ queries [6], where the parameter $c$ is equal to $2 k$ in PAEQ (and PPAE). The extra nonce function is thus pseudo-random and generates unique nonces unless it reaches the birthday bound of $2^{r / 2}$ queries (this is smaller than the $2^{k}$ indifferentiability bound in this paper). Therefore, the following statement holds.

Proposition 1. Let $A$ be the adversary that makes $q$ encryption queries $(P, A)$ to the PPAE
 all the pairs $(P, A)$ are distinct, then the advantage of the adversary to violate confidentiality or data integrity are bounded by values given in Theorems 1 and 2 or $\max \left(\frac{q^{2}}{2^{2 k}}, \frac{q^{2}}{2^{r}}\right)$, whatever is larger.

## 6 Security of AESQ

In this section we discuss the security of the AESQ permutation. We expect that the reader is familiar with properties of the AES internal operations, and refer to 12 in case of questions.

### 6.1 Structure and decomposition

First we recall that one round of AES does not provide full diffusion as it mixes together only 32 bits. For instance, consider a column in the AES round that undergoes the MixColumns transformation. Before that its bytes have been permuted by ShiftRows and substituted by SubBytes; afterwards they are xored with a subkey and again go through SubBytes and ShiftRows. These operations can be grouped into a single 32-bit so-called "SuperSBox' parametrized with a 32-bit subkey. As a result, two AES rounds can be viewed as a layer SuperSubBytes of four 32-bit SuperSBoxes followed by ShiftRows and MixColumns, so that we can view AES-128 as a 5 -round cipher with larger S-boxes.

The same strategy applies to the AESQ permutation. The two-round groups that process the registers $A, B, C, D$ can be viewed as a single round with $16=4 \cdot 4$ parallel


Fig. 3. SuperSBox in AES. SuperSBoxes and a large linear transformation. This yields an $R$-round SPN permutation with 32 -bit S-boxes out of the original $2 R$-round one.

We can go further and view as many as four rounds of AESQ as a single round with MegaSBoxes (cf. MegaBoxes in [11). Indeed, refer to Figure 4 and Algorithm 1. Consider a 128 -bit register in round 2 , for instance register $A$. Let us first look at its input. Going into backward direction, the columns spread to all the registers and then each column undergoes MixColumns without any influence from the other parts of the register. Then the values are shuffled by ShiftRows, and finally updated byte-wise by SubBytes. Hence we can recompute the 16 bytes at the beginning of round 1, even though they are located in different registers. Let us now look at the register $A$ at the end of round 3. Its columns again spread to all the registers, and again undergo SubBytes and ShiftRows independently of the other register bytes. As a result, we can view as many as four rounds as a layer MegaSubBytes of four 128-bit MegaSBoxes followed by a MixColumns-based linear transformation, which we call MegaMixColumns (cf. the analysis in [17]). It has branch number 5, as it is exactly a set of MixColumns transformations with reordered inputs and outputs. Note that this decomposition must start with an odd round, and does not work for even rounds.


Fig. 4. MegaSBox in AESQ.

### 6.2 Analysis of permutations in the attack context

Only a few permutations as a single and secure object have been designed for the use in practical constructions. The most well-known is the Keccak 1600-bit permutation, which is used in the Keccak/SHA-3 hashing algorithm; the others are used in the SHA-3 competitors: CubeHash [5], Grostl [14], JH [27]. It is worth noticing that a permutation per se can not be formally defined "secure". The best we can make is an informal statement like the $2^{l}$ "flat sponge" claim [7], which basically states that no attack with complexity below $2^{l}$ and specific for the particular permutation exists. The parameter $l$ is used in defining the capacity parameters in sponge functions and in fact measures the designers' confidence.

In our case we claim $l=256$ or the 256 -bit security of AESQ against all attacks. In order to support our claim, we look at the existing attacks on permutation-based designs and check if they apply to AESQ.

Collision attacks. We first consider collision attacks on sponge-based hash functions. The collision attacks on the reduced Keccak [13] strongly rely on high-probability differential trails [22], and only add a couple of rounds over their length with the help of message-modification techniques. The so-called internal-differential attack [23], while exploiting similarities within the internal state, is also limited by the propagation of difference generated by the round constants. Hence to prevent these attacks we have to demonstrate the absence of high-probability differential trails for a high number of rounds.

Let us now consider compression functions based on permutations. For example, Grostl uses functions

$$
P(x \oplus y) \oplus Q(y) \oplus x \quad \text { and } \quad x \oplus P(x),
$$

where $P$ and $Q$ are AES-based permutations. The main strategy in collision attacks on the AES-based designs [15,20 is the construction of a truncated differential trail with low input and output Hamming weight. Then the conforming inputs are found with the rebound attack and are tested for a collision.

Preimage attacks. The preimage attacks on sponge-based hash functions have been also based on the differential properties of the permutation. As long as a differential generated by message difference $\Delta M$ has high probability in some output bits, it can be used to speed up the preimage search [21]. There are also generic methods that can save a factor of several bits by exploiting incomplete diffusion in the final rounds, but we note that their complexity can not be reduced much. The invariant attacks [2] do not apply because of round constants.

Preimage attacks on Grostl are based on the meet-in-the-middle framework. Whereas it is difficult to formalize the necessary conditions for these attacks to work, we notice that the number of rounds attacked with meet-in-the-middle/bicliques is smaller compared to the rebound attack (even though the attack goals are distinct).

Other attacks. A more generic set of attacks are given by the CICO and multi-CICO problems [7], which require the attacker to find one or more (input, output) pair conforming to certain bit conditions. There is no comprehensive treatment of these attacks, but they seem to be limited by twice the number of rounds needed for full diffusion.

Consider an 8 -round version $F$ of AESQ, for instance, that $F(X)=Y$, and we know the last 384 bits of both $X$ and $Y$. How difficult is it to restore $X$ and $Y$ ? The last 384 bits mean three registers $B, C, D$. Hence, we can compute registers $B, C, D$ through rounds 0,1 and 6,7 . Then we know 96 bits of each register in the beginning of round 2 and in the end of round 5 . By a simple meet-in-the-middle attack we match between rounds 3 and 4 and recover $X$ and $Y$ with complexity $2^{32}$. However, we did not manage to extend this technique to more rounds.

### 6.3 Possible attacks

Differential analysis of $A E S Q$. Let us evaluate the differential properties of AESQ. We are backed up with the analysis of AES by Daemen and Rijmen. They prove the following lemmar

Lemma 3. Any differential trail over 4 AES rounds has at least 25 active $S$-boxes.
We have demonstrated in Section 6.1 that four rounds of AESQ can be viewed as a layer of 128 -bit S-boxes and a large linear transformation, which has branch number 5 . The 128 -bit MegaSBoxes are exactly four-rounds of AES without the key additions and the last MixColumns. Hence each active MegaSBox has at least 25 active S-boxes.

The branch number of MegaMixColumns implies that eight AESQ rounds starting from an odd round have at least 5 active MegaSBoxes, and hence 125 active S-boxes. Therefore, nine rounds starting from an even round (as required by the permutation) have at least 126 active S-boxes, as the first round must have at least one. More generally, $8 R+1$ rounds must have at least $125 R+1$ active S-boxes. In turn, the last round has at maximum 64 active $S$-boxes, so the lower bound for $8 R$ rounds is $125 R-63$ active S-boxes. The other bounds are given in Table 2 ,

| Rounds | Active S-boxes | Rounds | Active S-boxes |
| :---: | :---: | :---: | :---: |
| 2 | 5 | $8 \mathrm{R}+2$ | $125 \mathrm{R}+2$ |
| 4 | 10 | $8 \mathrm{R}+4$ | $125 \mathrm{R}+7$ |
| 6 | 27 | $8 \mathrm{R}+6$ | $125 \mathrm{R}+27$ |
| 8 R | $125 \mathrm{R}-63$ | $8 \mathrm{R}+8$ | $125 \mathrm{R}+68$ |

Table 2. Minimal number of active S-boxes in the AESQ permutation.

The table implies that a 20 -round trail has at least 257 active S-boxes. Given that the maximal differential probability of each S-box is $2^{-6}$, the vast majority of these trails have

[^0]differential probability close to zero. The effect of clustering these trails into a single differential and the resulting probability is not completely studied, but the bound of maximum expected differential probability for 4 rounds of AES of $2^{-113}$ [18] indicates that these values should be very low for 16 rounds of AESQ and more unless the round constants accidentally admit a large number of conforming pair for a differential (the so called height parameter).

If we consider slide-like attacks, where differences between registers or states shifted by several rounds are considered, we obtain that the round constants generate one-byte difference in every round, so we can not hope for high probability differentials in this case.

Rebound attacks. Rebound attacks aim to construct a conforming input for a truncated differential trail with a little amortized cost. The attack consists of two phases: inbound phase, where a conforming input is constructed for the low-probability part in the meet-in-the-middle manner, and the outbound phase, where the differences are traced through the rest of the primitive without the control of the cryptanalyst.

Again, given the decomposition of AESQ into MegaSBoxes, we can generate a full difference distribution table for MegaSBoxes with the complexity $2^{256}$. Consider MegaSBoxes in rounds $5-8$ of the permutation. Let us select an active $S$-box in rounds 4 and 9 , and select a one-byte difference in them. Then we apply MegaMixColumns and obtain input/output differences for MegaSBoxes. We obtain actual state values and then recompute the difference in the outer directions. Active S-boxes in round 4 activates one MegaSBox in rounds 1-4, and then activates all the $S$-boxes of round 0 . At the other end, the active S-box in round 9 activates a MegaSBox in rounds $9-12$, which activates all S-boxes in round 13 . This yields a clear distinguisher for 12 rounds of AESQ with complexity $2^{256}$, and possibly a distinguisher for 14 rounds.

Final choice. We take the 20 -round AESQ permutation within PAEQ, which should provide a 256 -bit security against all attacks. After the third-party cryptanalysis, if little improvement over our analysis will be shown, we might consider reducing the number of rounds to 16 for 80 and 128 -bit security.

## 7 Acknowledgement

We are grateful to Anonymous for help with AESQ implementation.

## 8 Conclusion

We have proposed a new mode PPAE and its instantiation PAEQ, that delivers virtually all features of a modern authenticated encryption scheme. To the best of our knowledge, no other scheme combines parallelism, 128 -bit security, online processing and a quick tag update. The scheme has a compact security proof in the random permutation model.

Unfortunately, these advantages come at some cost. Our scheme employs $(2+\epsilon)$ calls to a permutation of width $n$ to encrypt $n$ bits of the message. To mitigate this effect, we suggested to use a reduced Keccak permutation or a new AES-like permutation AESQ of width 512 bits. We conduct the basic security analysis of AESQ and concluded that the most relevant attacks leave a significant security margin. The resulting scheme runs at 6-7 cycles per byte on modern processors.

For the future research, we plan to extend our security proof to capture variable-length keys and nonces. We also plan to optimize the permutation AESQ so that it provides the same security with reduced number of operations.

## References

1. Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. Parallelizable and authenticated online ciphers. In $A S I A C R Y P T ' 13$, volume 8269 of Lecture Notes in Computer Science. Springer, 2013.
2. Jean-Philippe Aumasson, Eric Brier, Willi Meier, María Naya-Plasencia, and Thomas Peyrin. Inside the hypercube. In ACISP, volume 5594 of Lecture Notes in Computer Science, pages 202-213. Springer, 2009.
3. Jean-Philippe Aumasson, Luca Henzen, Willi Meier, and María Naya-Plasencia. Quark: A lightweight hash. In CHES'10, volume 6225 of Lecture Notes in Computer Science, pages 1-15. Springer, 2010. https:// 131002.net/quark/quark_full.pdf
4. Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In ASIACRYPT'00, volume 1976 of Lecture Notes in Computer Science, pages 531-545. Springer, 2000.
5. Daniel J. Bernstein. CubeHash specification (2.b.1). Submission to NIST (Round 2), 2009.
6. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. On the indifferentiability of the sponge construction. In EUROCRYPT'08, volume 4965 of Lecture Notes in Computer Science, pages 181197. Springer, 2008.
7. Guido Bertoni, Joan Daemen, Michael Peeters, and Gilles Van Assche. The Keccak reference, version 3.0, 2011. http://keccak.noekeon.org/Keccak-reference-3.0.pdf
8. Andrey Bogdanov, Miroslav Knezevic, Gregor Leander, Deniz Toz, Kerem Varici, and Ingrid Verbauwhede. Spongent: A lightweight hash function. In CHES'11, volume 6917 of Lecture Notes in Computer Science, pages 312-325. Springer, 2011.
9. Andrey Bogdanov, Florian Mendel, Francesco Regazzoni, Vincent Rijmen, and Elmar Tischhauser. ALE: AES-based lightweight authenticated encryption. In FSE'13, to appear, 2013. available at http://www2. compute.dtu.dk/~anbog/fse13-ale.pdf
10. Joan Daemen. Permutation-based symmetric cryptography and Keccak. Technical report, Ecrypt II, Crypto for 2020 Invited Talk, 2013. https://www.cosic.esat.kuleuven.be/ecrypt/cryptofor2020/ slides/KeccakEcryptTenerife.pdf
11. Joan Daemen, Mario Lamberger, Norbert Pramstaller, Vincent Rijmen, and Frederik Vercauteren. Computational aspects of the expected differential probability of 4 -round aes and aes-like ciphers. Computing, 85(1-2):85-104, 2009.
12. Joan Daemen and Vincent Rijmen. The Design of Rijndael. AES - the Advanced Encryption Standard. Springer, 2002.
13. Itai Dinur, Orr Dunkelman, and Adi Shamir. New attacks on Keccak-224 and Keccak-256. In FSE'12, volume 7549 of Lecture Notes in Computer Science, pages 442-461. Springer, 2012.
14. Praveen Gauravaram, Lars R. Knudsen, Krystian Matusiewicz, Florian Mendel, Christian Rechberger, Martin Schlffer, and Sren S. Thomsen. Grøstl - a SHA-3 candidate. Submission to NIST (Round 3), 2011.
15. Kota Ideguchi, Elmar Tischhauser, and Bart Preneel. Improved collision attacks on the reduced-round Grøstl hash function. In ISC, volume 6531 of $L N C S$, pages 1-16. Springer, 2010.
16. Tetsu Iwata, Keisuke Ohashi, and Kazuhiko Minematsu. Breaking and repairing GCM security proofs. In CRYPTO'12, volume 7417 of Lecture Notes in Computer Science, pages 31-49. Springer, 2012.
17. Jérémy Jean, María Naya-Plasencia, and Martin Schläffer. Improved analysis of ECHO-256. In Selected Areas in Cryptography'11, volume 7118 of Lecture Notes in Computer Science, pages 19-36. Springer, 2011.
18. Liam Keliher and Jiayuan Sui. Exact maximum expected differential and linear probability for 2-round Advanced Encryption Standard (AES). IACR Cryptology ePrint Archive, 2005:321, 2005.
19. Ted Krovetz and Phillip Rogaway. The software performance of authenticated-encryption modes. In FSE'11, volume 6733 of Lecture Notes in Computer Science, pages 306-327. Springer, 2011.
20. Florian Mendel, Christian Rechberger, Martin Schlffer, and Sren S. Thomsen. Rebound attacks on the reduced Grøstl hash function. In $C T-R S A$, volume 5985 of $L N C S$, pages 350-365. Springer, 2010.
21. Pawel Morawiecki, Josef Pieprzyk, Marian Srebrny, and Michal Straus. Preimage attacks on the roundreduced Keccak with the aid of differential cryptanalysis. Cryptology ePrint Archive, Report 2013/561, 2013. http://eprint.iacr.org/
22. María Naya-Plasencia, Andrea Röck, and Willi Meier. Practical analysis of reduced-round keccak. In INDOCRYPT'11, volume 7107 of Lecture Notes in Computer Science, pages 236-254. Springer, 2011.
23. Thomas Peyrin. Improved differential attacks for ECHO and Grøstl. In CRYPTO'10, volume 6223 of Lecture Notes in Computer Science, pages 370-392. Springer, 2010.
24. Phillip Rogaway. Authenticated-encryption with associated-data. In ACM Conference on Computer and Communications Security'02, pages 98-107, 2002.
25. ISO/IEC 19772 JTC 1 SC 27. Information technology - Security techniques - Authenticated encryption, 2009.
26. D. Whiting, B. Schneier, S. Lucks, and F. Muller. Phelix: Fast encryption and authentication in a single cryptographic primitive. Technical report.
27. Hongjun Wu. The hash function JH. Submission to NIST (round 3), 2011.
28. Hongjun Wu and Bart Preneel. AEGIS: A fast authenticated encryption algorithm. Cryptology ePrint Archive, Report 2013/695, 2013. http://eprint.iacr.org/2013/695

## A Proof of Lemma 1

There are four possible options:

1. Fresh nonce;
2. Old nonce, fresh ciphertext (not obtained with that nonce);
3. Old nonce, old ciphertext, fresh associated data;
4. Old nonce, old ciphertext, old associated data, fresh tag.

The last option clearly never succeeds.
Fresh nonce. Since $N$ is fresh, the decryption request turns out to be a valid encryption request. Let $M^{\prime}$ be the plaintext of length $|C|$ that consists of zero bits, and let $C^{\prime}$ be its ciphertext under nonce $N$. Then the decryption request $(N, A, C, T)$ results in the encryption request $\left(N, A, C \oplus C^{\prime}\right)$, which generates the tag uniformly at random according to Game $G_{8}$. Therefore, the probability to forge the tag is upper bounded by $2^{-\tau}$.

Old nonce $N$, fresh ciphertext $C$. Let $C^{\prime}$ be the ciphertext previously obtained with $N$. Consider two cases.

1. $C^{\prime}$ is a prefix of $C$. Then $Z$ is fresh due to requirements of Game $G_{9}$ and is chosen uniformly at random.
2. $C^{\prime}$ is not a prefix of $C$. Let $j$ be any block index where $C^{\prime}$ differs from $C$. By definition, the variables $X_{j}^{\prime}$ and $X_{j}$ are distinct. There are two cases:

- The block $X_{j}^{\prime}[17 \ldots(n-k)]$ has been previously obtained with nonce $N^{\prime \prime}$ as block $X_{s}^{\prime \prime}[17 \ldots(n-k)]$. Only one such nonce could exist because of the requirements of Game $G_{3}$. Since all $X_{i}$ generated during encryption requests are sampled uniformly at random, the probability of $X_{s}^{\prime \prime}=X_{j}^{\prime}$ is upper bounded by $1 / 2^{k}$.
- The block $X_{j}^{\prime}[17 \ldots(n-k)]$ has never been obtained before. Then $Y\left(X_{j}^{\prime}\right)$ is sampled uniformly at random, and with probability $\sigma_{\Pi} / 2^{n-k-16}$ generates a not-fresh $Z$.
Therefore, $Z$ (and hence the tag) is chosen uniformly at random in all cases except one that has probability $1 / 2^{k}$. Thus probability to forge the tag in this case is upper bounded by

$$
2^{-\tau}+\sigma_{\Pi} / 2^{n-k-16}+1 / 2^{k}
$$

Old nonce $N$, old ciphertext $C$, fresh associated data $A$. Let $Y^{\prime}(A)$ be the value obtained during the processing of the associated data $A$. There are two options:

- $Y^{\prime}(A)=Y^{\prime}\left(A^{0}\right)$ for some previously queried $A^{0}$. Due to requirements of Game $G_{9}$, the value $Y(N, C) \oplus Y^{\prime}\left(A_{0}\right)=Z(N, C, A)$ is fresh.
$-Y^{\prime}(A)$ is fresh. Then it generates a not-fresh $Z$ with probability upper bounded by $\sigma_{\Pi} / 2^{n-k-16}$. Therefore, the probability to forge the tag in this case is upper bounded by

$$
2^{-\tau}+\sigma_{\Pi} / 2^{n-k-16}
$$

This concludes the proof of lemma.

## B Proof of Lemma 2

To prove this lemma, we follow the same games as in Theorems 1 and 2, The transitions from Game $G_{0}$ to $G_{1}$ and $G_{2}$ hold with the same probability.

The Game $G_{3}$ is replaced with Game $G_{3}^{\prime}$, where we require $\Pi$ to abort if there are collisions in $\left\{X_{j}^{(i)}[(n-k+1) \ldots k]\right\}$ for all used nonces and counter values. The probability of this event is upper bounded by $\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2} / 2^{k}$.

Games $G_{4}-G_{7}$ remain unmodified.
Finally, we unite Games $G_{8}$ and $G_{9}$ by requiring no collisions in $Z_{\text {ext }}$. Since all the pairs (plaintext, associated data) are distinct and there is no collision in $X[(n-k+1) \ldots k]$ according to $G_{3}^{\prime}$, every $Z$ has a fresh and uniformly chosen $Y$ or $Y^{\prime}$ in its XOR. Therefore, the upper bound on the collision event is $\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2} / 2^{k}$.

As long as all $Z$ are unique, the tags are uniformly generated. Therefore, the adversary's advantage is upper bounded by $2\left(\sigma_{\mathcal{F}}+\sigma_{\Pi}\right)^{2} / 2^{k}$.


Fig. 5. 8 rounds of the AESQ permutation.


[^0]:    ${ }^{1}$ Actually, they prove a more general theorem for SPN ciphers but the AESQ permutation does not satisfy its assumptions.

