

APE: Authenticated Permutation-Based Encryption for Lightweight Cryptography

Elena Andreeva^{1,2}, Begül Bilgin^{1,2,3}, Andrey Bogdanov⁴, Atul Luykx^{1,2},
Bart Mennink^{1,2}, Nicky Mouha^{1,2}, and Kan Yasuda^{1,5}

¹ Department of Electrical Engineering, ESAT/COSIC, KU Leuven, Belgium.
`firstname.lastname@esat.kuleuven.be`

² iMinds, Belgium.

³ Faculty of EEMCS, DIES, UTwente, Netherlands.

⁴ Department of Mathematics, Technical University of Denmark, Denmark.
`anbog@dtu.dk`

⁵ NTT Secure Platform Laboratories, Japan.
`yasuda.kan@lab.ntt.co.jp`

Abstract. The domain of lightweight cryptography focuses on cryptographic algorithms for extremely constrained devices. It is very costly to avoid nonce reuse in such environments, because this requires either a hardware source of randomness, or non-volatile memory to store a counter. At the same time, a lot of cryptographic schemes actually require the nonce assumption for their security. In this paper, we propose APE as the first permutation-based authenticated encryption scheme that is resistant against nonce misuse. We formally prove that APE is secure, based on the security of the underlying permutation. To decrypt, APE processes the ciphertext blocks in reverse order, and uses inverse permutation calls. APE therefore requires a permutation that is both efficient for forward and inverse calls. We instantiate APE with the permutations of three recent lightweight hash function designs: QUARK, PHOTON, and SPONGENT. For any of these permutations, an implementation that supports both encryption and decryption requires less than 1.9 kGE and 2.8 kGE for 80-bit and 128-bit security levels, respectively.

Keywords. APE, Authenticated Encryption, Sponge Function, Online, Deterministic, Permutation-based, Misuse Resistant.

1 Introduction

In constrained environments, conventional solutions to cryptographic problems are prohibitively expensive to implement. Lightweight cryptography deals with cryptographic algorithms within the stringent requirements imposed by devices such as low-cost smart cards, sensor networks, and electronic body implants where energy, power, or hardware area consumption can be heavily restricted.

Although symmetric-key cryptography predominantly makes use of solutions based on block ciphers, recently permutation-based constructions [8] are gaining traction for a wide range of platforms, and on lightweight devices in particular.

Lightweight permutation-based hash functions include GLUON [5], PHOTON [18], QUARK [1, 2], and SPONGENT [10].

Lightweight applications in practice require not only hash functions but also secret-key cryptographic functions, such as authenticated encryption (AE). AE is a cryptographic primitive that guarantees two security goals: privacy and integrity. The prevalent solutions in this direction are block cipher based [23, 27, 30]. Permutation-based AE schemes were only recently proposed, such as the deterministic key-wrap scheme [20] of Khovratovich and SpongeWrap [6, 9] of the KECCAK team.

These two constructions unfortunately have their limitations. With the key-wrap scheme [20], the message length is restricted to one block by design. While sufficient for key wrapping [28], this construction cannot handle arbitrary-length data and is therefore not a full AE scheme. SpongeWrap [6, 9] can encrypt messages of varying lengths but relies on the uniqueness of the nonce value: failure to ensure so makes it possible to reuse the keystream of the encryption. For example, if a pair of plaintexts share a common prefix, the XOR of the first pair of plaintext blocks after this common prefix is leaked.

In Rogaway’s security formalism of nonce-based encryption [25, 26], the nonce is considered to be unique for every evaluation. While this approach has theoretical merits, in practice it is challenging to ensure that a nonce is never reused. This is especially the case in lightweight cryptography, as a nonce is realized either by keeping a state (and correctly updating it) or by providing a hardware source of randomness. Indeed, nonce misuse is a security threat in plenty of practical applications, not necessarily limited to the lightweight setting. Examples include flawed implementations of nonces [12, 14, 21, 22, 31], bad management of nonces by the user, and backup resets or virtual machine clones when the nonce is stored as a counter.

Nonce *misuse resistance* has become an important criterion in the design of AE schemes. The upcoming CAESAR competition [13] considers misuse resistance in detail for their selection of a portfolio of AE algorithms. The problem of nonce misuse has also been addressed by the recent deterministic AE scheme SIV [28], by the online AE scheme McOE [17], and in part by the aforementioned deterministic key-wrap scheme [20]. However, there are currently no permutation-based AE schemes that are resistant to nonce misuse.

Our Contributions

In this work we introduce APE (Authenticated Permutation-based Encryption). APE is the first permutation-based *and* nonce misuse resistant authenticated encryption scheme. APE is inspired by SpongeWrap [6, 9], but differs in several fundamental aspects in order to achieve misuse resistance. Most importantly, in APE the ciphertexts are extracted from the state, whereas SpongeWrap generates a keystream to perform the encryption. APE encryption processes data in an online manner, whereas decryption is done backwards using the inverse of the permutation. APE is formally introduced in Sect. 3. Here, we initially focus on associated data and messages of an integral number of blocks. In Sect. 6, we

show how APE can be generalized at almost no extra cost to handle fractional associated data and message blocks.

We prove that APE achieves privacy and integrity up to about $2^{c/2}$ queries, where c is the capacity parameter of APE. This result is proven in two different models: first, in Sect. 4, we prove security of APE in the ideal permutation model, where the underlying permutation is assumed to behave perfectly random. Next, in Sect. 5, we step away from this ideality assumption and prove that these security results also imply the security of APE in the standard model, up to the same bound.

APE is designed to be well suited for lightweight applications. In Sect. 7, we implement APE in less than 1.9 kGE and 2.8 kGE for 80-bit and 128-bit security respectively with the permutations of QUARK, PHOTON, and SPONGENT. The results also indicate that the by design enforced decryption function leads to a small overhead. Particularly, this overhead is very low compared to the cost of providing a hardware source of randomness to generate nonces.

We conclude our paper in Sect. 8.

2 Notation

Set $\mathbf{R} := \{0, 1\}^r$ and $\mathbf{C} := \{0, 1\}^c$. Given two strings A and B , we use $A\|B$ and AB interchangeably, so for example $A\|B \in \mathbf{R} \times \mathbf{C} \cong \{0, 1\}^{r+c} \ni AB$. Given $X \in \mathbf{R} \times \mathbf{C}$, X_r denotes its rate part and X_c its capacity part. We write $0 \in \mathbf{R}$ for a shorthand for $00 \cdots 0 \in \mathbf{R}$ and $1 \in \mathbf{C}$ for $00 \cdots 01 \in \mathbf{C}$. The symbol \oplus denotes the bitwise XOR operation of two (or more) strings.

An element of \mathbf{R} is called a block. Let \mathbf{R}^* denote the set of strings whose length is a multiple of r , at most $2^{c/2}$ blocks. This explicit bound of $2^{c/2}$ is needed in order to define a random function as being sampled over a finite set of functions. Note that the bounds we prove become trivial for queries of length $2^{c/2}$ blocks. Similarly, let \mathbf{R}^+ denote the set of strings whose length is a positive multiple of r , at most $2^{c/2}$ blocks. Given $M \in \mathbf{R}^+$, we divide it into blocks and write $M[1]M[2] \cdots M[w] \leftarrow M$, where each $M[i]$ is a block and w the block length of the string M .

Let \mathcal{A} be some class of computationally bounded adversaries. For convenience, we use the notation

$$\Delta_{\mathcal{A}}[f, g] := \sup_{A \in \mathcal{A}} |\Pr[A^f = 1] - \Pr[A^g = 1]|$$

to denote the supremum of the distinguishing advantages over all adversaries distinguishing f and g . Providing access to multiple algorithms is denoted with a comma, e.g. $\Delta[f_1, f_2; g_1, g_2]$ denotes distinguishing the combination of f_1 and f_2 from the combination of g_1 and g_2 .

3 APE Authenticated Encryption Mode

We now define our APE mode for the case of plaintexts and associated data of length a multiple of the block size. We refer to Sect. 6 for the generalization of

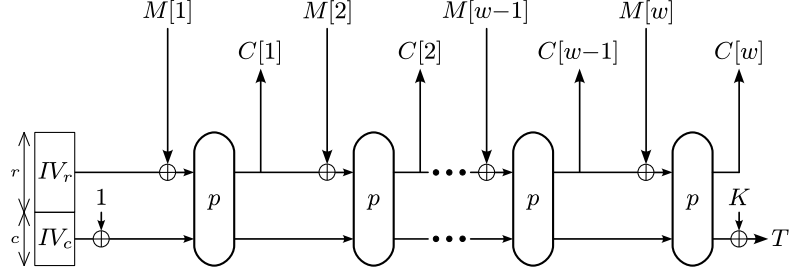


Fig. 1. The APE mode of operation (encryption). If there is no associated data ($A = \emptyset$), we set $IV_r := 0$ and $IV_c := K$. Otherwise, IV is computed from the associated data (see Fig. 2).




Fig. 2. Hashing associated data A in APE: IV used in Fig. 1 is computed as in this figure.

APE to fractional data blocks. APE iterates a fixed permutation $p : \mathbf{R} \times \mathbf{C} \rightarrow \mathbf{R} \times \mathbf{C}$ in a way similar to the sponge construction. The permutation p is the only underlying cryptographic primitive used by APE. A diagram of APE is given in Figs. 1 and 2.

More formally, APE consists of two functionalities: encryption \mathcal{E} and decryption \mathcal{D} , which are formally defined in Fig. 3. These functions employ several subroutines which are given in Fig. 4.

The encryption algorithm \mathcal{E} takes as input a key $K \in \mathcal{K} = \mathbf{C}$, associated data $A \in \mathbf{R}^*$, and a message $M \in \mathbf{R}^+$, and returns a ciphertext $C \in \mathbf{R}^+$ and a tag $T \in \mathbf{C}$, as $(C, T) \leftarrow \mathcal{E}_K(A, M)$. On the other hand, \mathcal{D} takes as input a key $K \in \mathbf{C}$, associated data $A \in \mathbf{R}^*$, a ciphertext $C \in \mathbf{R}^+$, and a tag $T \in \mathbf{C}$, and returns either a message $M \in \mathbf{R}^+$ or the reject symbol \perp , as $M/\perp \leftarrow \mathcal{D}_K(A, C, T)$. The two functionalities are sound, in the sense that whenever we encrypt a message as $(C, T) \leftarrow \mathcal{E}_K(A, M)$, we always get the message back, not \perp , via the decryption process $M \leftarrow \mathcal{D}_K(A, C, T)$.

In APE, the inclusion of a nonce is optional. If a nonce is required, it can be included as part of the associated data.

Algorithm 1: $\mathcal{E}_K(A, M)$	Algorithm 2: $\mathcal{D}_K(A, C, T)$
Input: $K \in \mathbf{C}, A \in \mathbf{R}^*, M \in \mathbf{R}^+$ Output: $C \in \mathbf{R}^+, T \in \mathbf{C}$ 1 if $A = \emptyset$ then 2 $IV \leftarrow (0, K) \in \mathbf{R} \times \mathbf{C}$ 3 else 4 $IV \leftarrow \text{hash-data}_{(0, K)}(A)$ 5 end 6 $(C, \hat{V}_c) \leftarrow \text{enc-message}_{IV}(M)$ 7 $T \leftarrow \hat{V}_c \oplus K$ 8 return (C, T)	Input: $K \in \mathbf{C}, A \in \mathbf{R}^*, C \in \mathbf{R}^+, T \in \mathbf{C}$ Output: $M \in \mathbf{R}^+$ or \perp 1 if $A = \emptyset$ then 2 $IV \leftarrow (0, K) \in \mathbf{R} \times \mathbf{C}$ 3 else 4 $IV \leftarrow \text{hash-data}_{(0, K)}(A)$ 5 end 6 $(M, V_c) \leftarrow \text{dec-ctxt}_{IV_r}(C, T \oplus K)$ 7 if $IV_c = V_c$ then 8 return M 9 else 10 return \perp 11 end

Fig. 3. The encryption $\mathcal{E}_K(A, M)$ and decryption $\mathcal{D}_K(A, C, T)$ algorithms of APE. In Fig. 4, the subroutines $\text{enc-message}_{IV}(M)$, $\text{dec-ctxt}_{IV_r}(C, \hat{V}_c)$ and $\text{hash-data}_V(A)$ are defined.

APE can be used with any permutation that is indistinguishable from a random permutation. However, APE decrypts in inverse direction and requires an efficiently invertible permutation. For lightweight applications, where hardware area size is critical, we provide implementation results for APE based on the permutations of PHOTON, QUARK, and SPONGENT. In Sect. 7, we show that including the inverses of these permutations only leads to a marginal increase of the size of the implementation.

4 Privacy and Integrity of APE

We prove that APE satisfies privacy (CPA) and integrity security up to about $c/2$ bits. Before doing so, in Sect. 4.1 we present the security model, where we formalize the notion of an ideal online function, and where we introduce the CPA and integrity security definitions. Then, privacy is proven in Sect. 4.2 and integrity in Sect. 4.3. The security results in this section are in the ideal permutation model, where we assume that the underlying permutation p is ideal. Later, in Sect. 5, we consider the security of APE in the standard model.

4.1 Security Model

Let $\text{Perm}(n)$ be the set of all permutations on n bits. By \perp , we denote a function that returns \perp on every input. When writing $x \stackrel{\$}{\leftarrow} \mathbf{X}$ for some finite set \mathbf{X} we mean that x is sampled uniformly from \mathbf{X} . To avoid confusion, for $X \in \mathbf{R} \times \mathbf{C}$ we sometimes write $[X]_c := X_c$ to denote the projection of X onto \mathbf{C} .

<hr/> Algorithm 3: $\text{enc-message}_{IV}(M)$ <hr/> Input: $IV \in \mathbf{R} \times \mathbf{C}, M \in \mathbf{R}^+$ Output: $C \in \mathbf{R}^+, \hat{V}_c \in \mathbf{C}$ <pre> 1 $V \leftarrow IV \oplus (0, 1)$ 2 $M[1]M[2] \dots M[w] \leftarrow M$ 3 for $i = 1$ to w do 4 $\hat{V} \leftarrow p(M[i] \oplus V_r, V_c)$ 5 $C[i] \leftarrow \hat{V}_r$ 6 $V \leftarrow \hat{V}$ 7 end 8 return $(C[1]C[2] \dots C[w], \hat{V}_c)$ </pre> <hr/>	<hr/> Algorithm 4: $\text{dec-ctxt}_{IV_r}(C, \hat{V}_c)$ <hr/> Input: $IV_r \in \mathbf{R}, C \in \mathbf{R}^+, \hat{V}_c \in \mathbf{C}$ Output: $M \in \mathbf{R}^+, V_c \in \mathbf{C}$ <pre> 1 $C[1]C[2] \dots C[w] \leftarrow C$ 2 $C[0] \leftarrow IV_r$ 3 for $i = w$ to 1 do 4 $V \leftarrow p^{-1}(C[i], \hat{V}_c)$ 5 $M[i] \leftarrow C[i-1] \oplus V_r$ 6 $\hat{V}_c \leftarrow V_c$ 7 end 8 return $(M[1]M[2] \dots M[w], V_c \oplus 1)$ </pre> <hr/>
<hr/> Algorithm 5: $\text{hash-data}_V(A)$ <hr/> Input: $V \in \mathbf{R} \times \mathbf{C}, A \in \mathbf{R}^+$ Output: $\hat{V} \in \mathbf{R} \times \mathbf{C}$ <pre> 1 $A[1]A[2] \dots A[u] \leftarrow A$ 2 for $i = 1$ to u do 3 $\hat{V} \leftarrow p(A[i] \oplus V_r, V_c)$ 4 $V \leftarrow \hat{V}$ 5 end 6 return \hat{V} </pre> <hr/>	

Fig. 4. The subroutines $\text{enc-message}_{IV}(M)$, $\text{dec-ctxt}_{IV_r}(C, \hat{V}_c)$, and $\text{hash-data}_V(A)$ which are used in Fig. 3 of the APE algorithms.

Online functions were first introduced in [3]. We deviate slightly from their approach by explicitly defining our ideal online function in terms of random functions.

Definition 1 (Ideal Online Function). *Let $g : \mathbf{R}^* \times \mathbf{R}^* \rightarrow \mathbf{R}$ and $g' : \mathbf{R}^* \times \mathbf{R}^* \rightarrow \mathbf{C}$ be random functions. Then, on input of (A, M) with $w = |M|/r$, we define $\$: \mathbf{R}^* \times \mathbf{R}^+ \rightarrow \mathbf{R}^+ \times \mathbf{C}$ as*

$$\$(A, M[1]||M[2]|| \dots ||M[w]) = (C[1]||C[2]|| \dots ||C[w], T),$$

where

$$\begin{aligned} C[j] &= g(A, M[1]|| \dots ||M[j]) \text{ for } j = 1, \dots, w, \\ T &= g'(A, M). \end{aligned}$$

Notice that the above function is actually online: prefixes of outputs remain the same if prefixes of the inputs remain constant. Furthermore, if the associated data in the ideal online function is unique for each invocation of the function,

then we achieve full privacy. This is because the inputs to g and g' will then be unique for each invocation, and since g and g' are random functions, we get outputs that are independent and uniformly distributed. By comparing our scheme to the above ideal online function we will automatically achieve the notions of CPA security and integrity from Rogaway and Zhang [29] and Fleischmann et al. [17].

Definition 2. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ denote an AE scheme. The CPA advantage of a distinguisher D is defined as

$$\mathbf{Adv}_{\Pi}^{\text{cpa}}(D) = \left| \Pr \left[p \xleftarrow{\$} \text{Perm}(r+c), K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K, \perp, p, p^{-1}} = 1 \right] - \Pr \left[p \xleftarrow{\$} \text{Perm}(r+c) : D^{\$, \perp, p, p^{-1}} = 1 \right] \right|.$$

By $\mathbf{Adv}_{\Pi}^{\text{cpa}}(q, m)$ we denote the supremum taken over all distinguishers making q queries of total length m blocks.

Definition 3. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ denote an AE scheme. The integrity advantage of a distinguisher D is defined as

$$\mathbf{Adv}_{\Pi}^{\text{int}}(D) = \left| \Pr \left[p \xleftarrow{\$} \text{Perm}(r+c), K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K, \mathcal{D}_K, p, p^{-1}} = 1 \right] - \Pr \left[p \xleftarrow{\$} \text{Perm}(r+c), K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K, \perp, p, p^{-1}} = 1 \right] \right|.$$

We assume that the distinguisher does not make a decryption query (A, C, T) if it ever obtained $(C, T) \leftarrow \mathcal{E}_K(A, M)$ for some M . By $\mathbf{Adv}_{\Pi}^{\text{int}}(q, m)$ we denote the supremum taken over all distinguishers making q queries of total length m blocks.

4.2 Privacy

In this section, we present a privacy security proof for APE.

Theorem 1. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be the APE construction. Then,

$$\mathbf{Adv}_{\Pi}^{\text{cpa}}(q, m) \leq \frac{m^2}{2^{r+c}} + \frac{m(m+1)}{2^c}.$$

Proof. We consider the strongest possible type of distinguishers: let D be any information-theoretic distinguisher which has unbounded computational power and whose complexity is measured solely by the number of queries it makes to its oracles. Without loss of generality, we restrict ourselves to distinguishers that do not ask “trivial” queries, queries to which it knows the answer in advance.

As a first step, we make a PRP-PRF switch [4]: we move from random permutation (p, p^{-1}) to a primitive (f, f^{-1}) defined as follows. This primitive maintains an initially empty list of responses, \mathcal{F} , and we denote its domain/range by $\text{dom}(\mathcal{F})/\text{rng}(\mathcal{F})$. Now, on a non-trivial forward query $f(x)$, the response y is

randomly drawn from $\mathbf{R} \times \mathbf{C}$. The primitive aborts if y happens to be in $\text{rng}(\mathcal{F})$ already; otherwise, the fresh tuple (x, y) is added to \mathcal{F} . Similarly for inverse queries to f^{-1} . Clearly, (p, p^{-1}) and (f, f^{-1}) behave identically as long as the latter does not abort. Given that the distinguisher makes at most q queries of total length m blocks (each block corresponds to a new (f, f^{-1}) -query), such an abort happens with probability at most $\binom{m}{2}/2^{r+c} \leq m^2/2^{r+c+1}$. We apply this PRP-PRF switch to both the ideal and the real world, and hence we find

$$\Delta_D(\mathcal{E}_K, \perp, p, p^{-1}; \$, \perp, p, p^{-1}) \leq \frac{m^2}{2^{r+c}} + \Delta_D(\mathcal{E}_K, \perp, f, f^{-1}; \$, \perp, f, f^{-1}). \quad (1)$$

In the remainder, we consider D to have oracle access to one of the two worlds: (F, f, f^{-1}) , where $F \in \{\mathcal{E}_K, \$\}$ (without loss of generality we can drop the \perp).

If f is called by D then we call this a direct f -query, and similar for direct f^{-1} -queries. A call of f by \mathcal{E}_K (as a result of D calling \mathcal{E}_K) is called an indirect f -query. When we do not specify whether an f -query is indirect or direct, we mean that it could be either. Note that indirect queries do not occur in the random world $(\$, f, f^{-1})$. Every indirect f -query has a sequence of associated data blocks and message blocks leading up to it (from the \mathcal{E}_K -query calling it); we call this sequence the message chain associated to the indirect f -query.

Let \mathcal{Q}_i denote the set of all prefixes of all queries made by D to its F -oracle before the i th (f, f^{-1}) -query, where a query (A, M) results in prefixes $\{A[1], A[1] \parallel A[2], \dots, A \parallel M\}$. Regarding all *direct* queries before the i th query, we denote by X_i^{dir} the set of all capacity values input to f -queries or output of f^{-1} -queries. For example, a direct forward query $y \leftarrow f(x)$ adds $[x]_c$ to X_i^{dir} and a direct inverse query $x \leftarrow f^{-1}(y)$ adds $[x]_c$ to X_i^{dir} . Similarly, by X_i^{ind} we denote the set of all capacity values input to *indirect* f -queries before the i th f -query. We write $X_i = X_i^{\text{dir}} \cup X_i^{\text{ind}}$, and initialize $X_0^{\text{ind}} = \{K\}$.

We define event $\mathbf{E}_i = \mathbf{E}_i^{\text{dir-}X} \cup \mathbf{E}_i^{\text{ind-}X}$, where

$$\begin{aligned} \mathbf{E}_i^{\text{dir-}X} &: \text{direct query } y \leftarrow f(x) \text{ or } x \leftarrow f^{-1}(y) \text{ satisfies } [x]_c \in X_i^{\text{ind}} \cup X_i^{\text{ind}} \oplus 1, \\ \mathbf{E}_i^{\text{ind-}X} &: \text{indirect query } f(x) \text{ with message chain } (A, M) \notin \mathcal{Q}_i \text{ satisfies} \\ &\quad [f(x)]_c \in X_i \cup X_i \oplus 1. \end{aligned}$$

We furthermore define

$$\hat{\mathbf{E}}_i := \mathbf{E}_i \cap \bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j, \text{ and } \mathbf{E} := \bigcup_{i=1}^m \hat{\mathbf{E}}_i, \quad (2)$$

where $\overline{\mathbf{E}}_j$ is the complement of \mathbf{E}_j .

Now, the remainder of the proof is divided as follows. In Lem. 1 we will prove that $(\mathcal{E}_K, f, f^{-1})$ and $(\$, f, f^{-1})$ are indistinguishable as long as \mathbf{E} does not occur. From (1) and the fundamental lemma of game playing [4] we find

$$\text{Adv}_\Pi^{\text{cpa}}(q, m) \leq \frac{m^2}{2^{r+c}} + \Pr[D^{\mathcal{E}_K, f, f^{-1}} \text{ sets } \mathbf{E}].$$

Then, in Lem. 2, we will prove that $\Pr[D^{\mathcal{E}_K, f, f^{-1}} \text{ sets } \mathbf{E}] \leq \frac{m(m+1)}{2^c}$, which completes the proof. \square

Lemma 1. *Given that \mathbf{E} does not occur, $(\mathcal{E}_K, f, f^{-1})$ and $(\$, f, f^{-1})$ are indistinguishable.*

Proof. Note that in the ideal world, each direct f -query is new, and is answered with a uniformly randomly drawn response. Now, consider a direct query $f(x)$ in the real world. As the distinguisher does not make trivial queries, it does not coincide with any previous direct query. Additionally, if $[x]_c \in X_i^{\text{ind}} \cup X_i^{\text{ind}} \oplus 1$, where $f(x)$ is the i th f -query, then this would trigger $\mathbf{E}_i^{\text{dir-}X}$, hence we can assume $[x]_c \notin X_i^{\text{ind}} \cup X_i^{\text{ind}} \oplus 1$. This means that the query $f(x)$ is truly new, and its value is independently and uniformly distributed. The same reasoning applies to f^{-1} -queries. Therefore, we only need to consider queries to the big oracle $F \in \{\mathcal{E}_K, \$\}$. Let (A, M) be a query made by the distinguisher. Denote by w the number of blocks of M . Denote the corresponding ciphertext and tag by (C, T) .

First consider the case $(A, M[1] \parallel \dots \parallel M[j]) \in \mathcal{Q}_i$ for some $j \in \{0, \dots, w\}$ and assume j is maximal (we will come back to the case of $(A, *) \notin \mathcal{Q}_i$ later in the proof). Let (A', M') be the corresponding earlier query, so $M[1] \parallel \dots \parallel M[j] = M'[1] \parallel \dots \parallel M'[j]$, and denote its ciphertext and tag by (C', T') and block length by w' . Clearly, in the ideal world $(\$, f, f^{-1})$, we have $C[i] = C'[i]$ for $i = 1, \dots, j$, but $C[i]$ for $i = j+1, \dots, w$ and T are uniformly randomly drawn. We will consider how these values are distributed in the real world $(\mathcal{E}_K, f, f^{-1})$. We first consider the general case $j < w$, the case $j = w$ is discussed afterwards.

1. $C[1], \dots, C[j]$. Also in the real world, these values equal $C'[1], \dots, C'[j]$, which follows clearly from the specification of \mathcal{E}_K . Note that in particular, the state value V equals V' after the j th round.
2. $C[j+1]$. We make a distinction between $j > 0$ and $j = 0$, and start with the former case.

Write the indirect query corresponding to the j th round as $f(x)$. The input of the $(j+1)$ th query will be $f(x) \oplus (M[j+1], 0)$. Note that $(A, M' \parallel M[j+1]) \notin \mathcal{Q}$, as this would contradict the fact that j is maximal. Now, assume this $(j+1)$ th query has already been made before, i.e. $[f(x)]_c \in X^{\text{dir}} \cup X^{\text{ind}}$. This may be the case (it may even date from before the evaluation of (A, M')), but at this particular time the capacity part $[f(x)]_c$ did not hit any element from $X^{\text{dir}} \cup X^{\text{ind}}$ (otherwise it would have triggered $\mathbf{E}^{\text{ind-}X}$). After this query has been made, there has not been any newer indirect query or any newer direct query whose capacity part hit $[f(x)]_c$ (both cases would have triggered $\mathbf{E}^{\text{dir-}X} \cup \mathbf{E}^{\text{ind-}X}$). Thus, the query corresponding to the $(j+1)$ th round is generated independently and uniformly at random.

Now, in the case $j = 0$, denote the state coming from $\text{hash-data}_{0,K}(A)$ by IV (if $A = \emptyset$, $IV = (0, K)$). The same story as before applies with the difference that now the input to the $(j+1)$ th query is $IV \oplus (M[j+1], 1)$. Here we use that by $\bar{\mathbf{E}}$, no other query hit $X_j^{\text{ind}} \oplus 1$ (for direct queries) or $X_j \oplus 1$ (for indirect queries) in the meanwhile, and that X^{ind} is initialized with $\{K\}$.

3. $C[j+2], \dots, C[w]$. By the above argument, the indirect query made in the $(j+1)$ th round of (A, M) , say $f(x)$ for the sake of presentation, is responded

with a uniformly random answer. This query would have triggered $\mathbf{E}^{\text{ind-}X}$ if $[f(x)]_c \in X_i$. Therefore, we know that also the $(j+2)$ th query is truly random and so is $C[j+2]$. The same reasoning applies up to $C[w]$.

4. T . The same reasoning applies: the previous query is responded with a truly random answer $f(x)$. Consequently $T = [f(x)]_c \oplus K$ is random too.

A special treatment is needed for $j = w$. In this case, $C[1], \dots, C[w]$ equals $C'[1], \dots, C'[w]$ by construction, but the query producing T is not new. Yet, the distinguisher never made that query itself by virtue of $\mathbf{E}^{\text{dir-}X}$, so it never learnt $T \oplus K$. Besides, due to the absence of indirect capacity collisions, $\mathbf{E}^{\text{ind-}X}$, every f -query will produce a tag at most once. This means that T will look uniformly random to the distinguisher, as it would look if it were produced by $\$$.

Finally, we consider the case $(A, *) \notin \mathcal{Q}_i$, hence this is the first time a query for this particular associated data A is made. Then, the above reasoning carries over for $j = 0$ with the simplification that if $A \neq \emptyset$, the value IV_c can be considered new. \square

Lemma 2. $\Pr[D^{\mathcal{E}_K, f, f^{-1}} \text{ sets } \mathbf{E}] \leq \frac{m(m+1)}{2^c}.$

Proof. Inspired by (2), we start bounding $\Pr[\mathbf{E}_i \cap \bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j]$ for $i \in \{1, \dots, m\}$. Clearly,

$$\Pr[\mathbf{E}_i \cap \bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j] \leq \Pr[\mathbf{E}_i \mid \bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j].$$

Therefore, we assume $\bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j$ and consider the probability the i th query triggers \mathbf{E}_i .

If the i th query is a direct (forward or inverse) query, it triggers $\mathbf{E}_i^{\text{dir-}X}$ if the distinguisher guesses (in case of forward) or hits (in case of inverse) a capacity part in $X_i^{\text{ind}} \cup X_i^{\text{ind}} \oplus 1$, which happens with probability at most $2|X_i^{\text{ind}}|/2^c$. On the other hand, if the i th query is a new indirect query (i.e. for which $(A, M) \notin \mathcal{Q}_i$) it triggers $\mathbf{E}_i^{\text{ind-}X}$ if $[f(x)]_c \in X_i \cup X_i \oplus 1$. This occurs with probability at most $2|X_i|/2^c$.

As the query is either direct or indirect, we could take the maximum of both values. Given that $|X_i^{\text{ind}}| \leq |X_i| \leq i$, we find:

$$\Pr[\mathbf{E}_i \mid \bigcap_{j=1}^{i-1} \overline{\mathbf{E}}_j] \leq \frac{2i}{2^c}.$$

The result is now obtained by summing over $i = 1, \dots, m$ (as in (2)). \square

4.3 Integrity

In this section, we present an integrity security proof for APE.

Theorem 2. *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be the APE construction. Then,*

$$\mathbf{Adv}_{\Pi}^{\text{int}}(q, m) \leq \frac{m^2}{2^{r+c}} + \frac{2m(m+1)}{2^c}.$$

Proof. The basic idea of the proof is the same as for Thm. 1. Again, let D be any information-theoretic distinguisher. By the PRP-PRF switch, we find

$$\Delta_D(\mathcal{E}_K, \mathcal{D}_K, p, p^{-1}; \mathcal{E}_K, \perp, p, p^{-1}) \leq \frac{m^2}{2^{r+c}} + \Delta_D(\mathcal{E}_K, \mathcal{D}_K, f, f^{-1}; \mathcal{E}_K, \perp, f, f^{-1}). \quad (3)$$

We consider D to have oracle access to one of the two worlds: $(\mathcal{E}_K, F, f, f^{-1})$, where $F \in \{\mathcal{D}_K, \perp\}$.

We use the same notation as in Thm. 1, but slightly more involved definitions are required and we re-introduce them. If f is called by D then we call this a direct f -query, and similar for direct f^{-1} -queries. A call of f by \mathcal{E}_K or \mathcal{D}_K (as a result of D calling them) is called an indirect f -query, and similar for indirect f^{-1} -queries (via \mathcal{D}_K). Every indirect f -query has a sequence of associated data blocks and/or message blocks leading up to it (from the \mathcal{E}_K - or \mathcal{D}_K -query calling it); we call this sequence the message chain associated to the indirect f -query. Every indirect f^{-1} -query has a tag and a sequence of ciphertext blocks leading up to it, and we call this sequence the associated ciphertext chain.

Let \mathcal{Q}_i denote the set of all prefixes of all queries made by D to its \mathcal{E}_K -oracle before the i th (f, f^{-1}) -query, where an \mathcal{E}_K -query (A, M) results in prefixes $\{A[1], A[1]\|A[2], \dots, A\|M\}$. In this set, we also include $\{A[1], \dots, A\}$ for an F -query (A, C, T) . Let \mathcal{Q}_i^{-1} denote the set of all suffixes of all queries made by D to its F -oracle before the i th query, where an F -query (A, C, T) results in suffixes $\{C[w]\|T, C[w-1]\|C[w]\|T, \dots, C\|T\}$. (The tag value T is included here for technical reasons.) Regarding all *direct* queries before the i th query, we denote by X_i^{dir} the set of all capacity values input to f -queries or output of f^{-1} -queries, and by Y_i^{dir} the set of all capacity values input to f^{-1} -queries or output of f -queries. For example, a direct forward query $y \leftarrow f(x)$ adds $[x]_c$ to X_i^{dir} and $[y]_c$ to Y_i^{dir} . The sets X_i^{ind} and Y_i^{ind} are defined similarly. We write $X_i = X_i^{\text{dir}} \cup X_i^{\text{ind}}$ and $Y_i = Y_i^{\text{dir}} \cup Y_i^{\text{ind}}$, and initialize $X_0^{\text{ind}} = Y_0^{\text{ind}} = \{K\}$.

We define event $\mathbf{E}_i = \mathbf{E}_i^{\text{dir-}X} \cup \mathbf{E}_i^{\text{ind-}X} \cup \mathbf{E}_i^{\text{dir-}Y} \cup \mathbf{E}_i^{\text{ind-}Y}$, where $\mathbf{E}_i^{\text{dir-}X}$ and $\mathbf{E}_i^{\text{ind-}X}$ are as in the proof of Thm. 1 with the renewed definitions of the sets, and where

$\mathbf{E}_i^{\text{dir-}Y}$: direct query $y \leftarrow f(x)$ or $x \leftarrow f^{-1}(y)$ satisfies $[y]_c \in Y_i^{\text{ind}} \cup Y_i^{\text{ind}} \oplus 1$,

$\mathbf{E}_i^{\text{ind-}Y}$: indirect query $f^{-1}(y)$ with ciphertext chain $(C, T) \notin \mathcal{Q}_i^{-1}$ satisfies

$$[f^{-1}(y)]_c \in Y_i \cup Y_i \oplus 1 \text{ or } [y]_c \in Y_i^{\text{dir}} \oplus K.$$

Definitions $\hat{\mathbf{E}}_i$ and \mathbf{E} are as before. The latter condition of $\mathbf{E}_i^{\text{ind-}Y}$, $[y]_c \in Y_i^{\text{dir}} \oplus K$, covers the case the distinguisher obtains the key by making a direct inverse query and a \mathcal{D}_K -query.

Now, the remainder of the proof is divided as follows. In Lem. 3 we will prove that $(\mathcal{E}_K, \mathcal{D}_K, f, f^{-1})$ and $(\mathcal{E}_K, \perp, f, f^{-1})$ are indistinguishable as long as \mathbf{E} does not occur. From (3) and the fundamental lemma of game playing [4] we find

$$\mathbf{Adv}_H^{\text{cpa}}(q, m) \leq \frac{m^2}{2^{r+c}} + \Pr[D^{\mathcal{E}_K, \mathcal{D}_K, f, f^{-1}} \text{ sets } \mathbf{E}].$$

Then, in Lem. 4, we will prove that $\Pr[D^{\mathcal{E}_K, \mathcal{D}_K, f, f^{-1}} \text{ sets } \mathbf{E}] \leq \frac{2m(m+1)}{2^c}$, which completes the proof. \square

Lemma 3. *Given that \mathbf{E} does not occur, $(\mathcal{E}_K, \mathcal{D}_K, f, f^{-1})$ and $(\mathcal{E}_K, \perp, f, f^{-1})$ are indistinguishable.*

Proof. For direct (f, f^{-1}) -queries, the analysis of Lem. 1 carries over with the difference that $[y]_c \in Y_i^{\text{ind}} \cup Y_i^{\text{ind}} \oplus 1$ would trigger event $\mathbf{E}^{\text{dir-}Y}$. Also, queries made to \mathcal{E}_K in the real and ideal world are handled the same. Here, we use that in the real world, indirect f -queries coming from \mathcal{D}_K (corresponding to the evaluation of `hash-data0,K`), do not ruin the distribution of the responses from \mathcal{E}_K by $\overline{\mathbf{E}}^{\text{ind-}X}$, where we now deal with a larger set X_i^{ind} . Therefore, we only need to consider queries to the big oracle $F \in \{\mathcal{D}_K, \perp\}$. Let (A, C, T) be a query made by the distinguisher. Denote by w the number of blocks of M . Denote the state coming from `hash-data0,K`(A) by IV , and the corresponding message and capacity value by (M, V_c) . Here, $V_c \oplus 1$ equals the capacity part of the call to f^{-1} corresponding to $C[1]$.

Clearly, in the ideal world $(\mathcal{E}_K, \perp, f, f^{-1})$, the query is responded with \perp . Therefore, the distinguisher has no advantage in the real world unless $IV_c = V_c$. We distinguish two cases:

- $(C, T) \notin \mathcal{Q}_i^{-1}$. A similar reasoning as for the value T in Lem. 1 subcase “ $i < w$ ” results in the observation that the indirect f^{-1} -query corresponding to the ciphertext block $C[1]$ is new and the response is uniformly randomly drawn. The only fundamental difference lies in the fact that we now use events $\mathbf{E}^{\text{dir-}Y}$ and $\mathbf{E}^{\text{ind-}Y}$, and rely on the fact the first indirect query never matches a direct query $\oplus K$ (by $\overline{\mathbf{E}}_i^{\text{ind-}Y}$) or vice versa (by $\overline{\mathbf{E}}_i^{\text{dir-}Y}$). We skip the details. Now, D succeeds if the capacity part of this value, say $[f^{-1}(y)]_c$, equals $IV_c \oplus 1$, but then this indirect query would have triggered $\mathbf{E}^{\text{ind-}Y}$. We note that this value IV_c may be an older value, e.g., if $A \in \mathcal{Q}_i$, but this does not invalidate the analysis.
- $(C, T) \in \mathcal{Q}_i^{-1}$. The further reasoning depends on whether $(A, M') \in \mathcal{Q}_i$ for some M' .
 - $(A, *) \notin \mathcal{Q}_i$. If $A = \emptyset$, the distinguisher succeeding would mean that $V_c = K \oplus 1$. But this means that at some earlier point in time an indirect f^{-1} -query has hit $K \oplus 1$ and thus invalidated $\mathbf{E}^{\text{dir-}Y} \cup \mathbf{E}^{\text{ind-}Y}$. Now, if $A \neq \emptyset$, the analysis of Lem. 1 for the same case “ $(A, *) \notin \mathcal{Q}_i$ ” carries over: the value IV_c can be considered new and if it hits $V_c \oplus 1$, the query would trigger $\mathbf{E}^{\text{ind-}X}$.
 - $(A, M') \in \mathcal{Q}_i$ for some M' . By assumption, (A, M') and (C, T) cannot correspond to one and the same query. Without loss of generality, (A, M') corresponds to an earlier query than (C, T) . Now, the distinguisher’s query is valid if $IV_c = V_c \oplus 1$. If this is the case, the (C, T) query must have triggered $\mathbf{E}^{\text{ind-}Y}$. Also, in case $A = \perp$, this equation could not hold as $V_c = K \oplus 1$ would have triggered $\mathbf{E}^{\text{ind-}Y}$.

This completes the proof of Lem. 3. \square

Lemma 4. $\Pr[D^{\mathcal{E}_K, \mathcal{D}_K, f, f^{-1}} \text{ sets } \mathbb{E}] \leq \frac{2m(m+1)}{2^c}.$

Proof. The analysis is fairly similar to the proof of Lem. 2, with the difference that events $\mathbb{E}^{\text{dir-}Y}$ and $\mathbb{E}^{\text{ind-}Y}$ have to be considered too.

If the i th query is a direct (forward or inverse) query, it triggers $\mathbb{E}_i^{\text{dir-}X}$ or $\mathbb{E}_i^{\text{dir-}Y}$ with probability at most $\frac{4i}{2^c}$. If it is an indirect f -query, the analysis of Lem. 2 carries over, which means that $\mathbb{E}_i^{\text{ind-}X}$ is satisfied with probability at most $\frac{2i}{2^c}$. On the other hand, if it is an indirect f^{-1} -query, it sets $\mathbb{E}_i^{\text{ind-}Y}$ with probability at most $\frac{3i}{2^c}$ (note that $\mathbb{E}_i^{\text{ind-}Y}$ contains an additional condition, compared with $\mathbb{E}_i^{\text{ind-}X}$). As the query is either direct or indirect, and either forward or inverse, we could take the maximum of the three values, and our bound is obtained by summing over $i = 1, \dots, m$. \square

5 Standard Model Security of APE

As is conventionally done for existing permutation-based designs, our proof for APE assumes that the underlying permutation is ideal. Without making any changes to the specification of APE, we now provide a standard model security argument for our scheme. Inspired by [15], we note that APE can also be described as a block cipher based design: we drop the key additions at the beginning and end, and replace the permutations with a keyed block cipher E_K defined by $E_K := KpK := \oplus_{0\parallel K} \circ p \circ \oplus_{0\parallel K}$. (One can view E_K as the Even-Mansour [16] block cipher with partial key addition.) We remark that this is, indeed, an equivalent description of APE. In our notation we denote APE as described and based on some block cipher E by $\Pi' = (\mathcal{K}, \mathcal{E}^E, \mathcal{D}^E)$. We first give the privacy and integrity definitions in the standard model and then show that our results of Thm. 1 and Thm. 2 easily translate to a standard model security of Π' .

Definition 4. Let E be a block cipher, and let $\Pi' = (\mathcal{K}, \mathcal{E}^E, \mathcal{D}^E)$ denote an AE scheme. The CPA advantage of a distinguisher D is defined as

$$\text{Adv}_{\Pi'}^{\text{cpa}}(D) = \left| \Pr \left[K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K^E} = 1 \right] - \Pr \left[D^{\$} = 1 \right] \right|.$$

By $\text{Adv}_{\Pi'}^{\text{cpa}}(t, q, m)$ we denote the supremum taken over all distinguishers running in time t and making q queries of total length m blocks. Alternatively, we write $\text{Adv}_{\Pi'}^{\text{cpa}}(t, q, m) = \Delta_{q, m}^t(\mathcal{E}_K^E; \$)$ as in Def. 2 with the inclusion of t .

Definition 5. Let E be a block cipher, and let $\Pi' = (\mathcal{K}, \mathcal{E}^E, \mathcal{D}^E)$ denote an AE scheme. The integrity advantage of a distinguisher D is defined as

$$\text{Adv}_{\Pi'}^{\text{int}}(D) = \left| \Pr \left[K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K^E, \mathcal{D}_K^E} = 1 \right] - \Pr \left[K \xleftarrow{\$} \mathcal{K} : D^{\mathcal{E}_K^E, \perp} = 1 \right] \right|.$$

By $\text{Adv}_{\Pi'}^{\text{int}}(t, q, m)$ we denote the supremum taken over all distinguishers running in time t and making q queries of total length m blocks. Alternatively, we write $\text{Adv}_{\Pi'}^{\text{int}}(t, q, m) = \Delta_{q, m}^t(\mathcal{E}_K^E, \mathcal{D}_K^E; \mathcal{E}_K^E, \perp)$ as in Def. 3 with the inclusion of t .

In both definitions we refer to the rate of Π' , the number of block cipher calls per message block, as ρ . Furthermore, we need the notion of strong pseudorandom permutation, or $\text{prp}\pm 1$, security of E .

Definition 6. Let E be a block cipher. The $\text{prp}\pm 1$ advantage of a distinguisher \mathcal{D} is defined as

$$\mathbf{Adv}_E^{\text{prp}\pm 1}(\mathcal{D}) = \left| \Pr \left[K \xleftarrow{\$} \mathcal{K} : D^{E_K, E_K^{-1}} = 1 \right] - \Pr \left[\pi \xleftarrow{\$} \text{Perm}(r+c) : D^{\pi, \pi^{-1}} = 1 \right] \right|.$$

By $\mathbf{Adv}_E^{\text{prp}\pm 1}(t, q)$ we denote the maximum advantage taken over all distinguishers that run in time t and make q queries.

We demonstrate that the standard model security of APE is implied by the results of Sect. 4. To this end we introduce two propositions, one with respect to the integrity and one with respect to the privacy of Π' .

Proposition 1. Let E be a block cipher.

$$\mathbf{Adv}_{\Pi'}^{\text{int}}(t, q, m) \leq \frac{m^2}{2^{r+c}} + \frac{3m(m+1)}{2^c} + 2\mathbf{Adv}_E^{\text{prp}\pm 1}(t', \rho m).$$

Proof. Let $K \xleftarrow{\$} \mathcal{K}$. Let E be a publicly available block cipher and $\pi, p \xleftarrow{\$} \text{Perm}(r+c)$ be random permutations. We first switch from E to random π :

$$\begin{aligned} \Delta_{q,m}^t(\mathcal{E}_K^E, \mathcal{D}_K^E; \mathcal{E}_K^E, \perp) &\leq \Delta_{q,m}^t(\mathcal{E}_K^E, \mathcal{D}_K^E; \mathcal{E}_K^\pi, \mathcal{D}_K^\pi) + \Delta_{q,m}^t(\mathcal{E}_K^\pi, \mathcal{D}_K^\pi; \mathcal{E}_K^\pi, \perp) + \\ &\quad \Delta_{q,m}^t(\mathcal{E}_K^\pi, \perp; \mathcal{E}_K^E, \perp) \\ &\leq \Delta_{q,m}^t(\mathcal{E}_K^\pi, \mathcal{D}_K^\pi; \mathcal{E}_K^\pi, \perp) + 2\mathbf{Adv}_E^{\text{prp}\pm 1}(t', \rho m), \end{aligned}$$

where $t' \approx t$. As π is a random permutation, we could give the distinguisher unlimited time (effectively considering information-theoretic distinguishers), and the bound simplifies to:

$$\Delta_{q,m}^t(\mathcal{E}_K^E, \mathcal{D}_K^E; \mathcal{E}_K^E, \perp) \leq \Delta_{q,m}(\mathcal{E}_K^\pi, \mathcal{D}_K^\pi; \mathcal{E}_K^\pi, \perp) + 2\mathbf{Adv}_E^{\text{prp}\pm 1}(t', \rho m).$$

For the remaining Δ -term:

$$\begin{aligned} \Delta_{q,m}(\mathcal{E}_K^\pi, \mathcal{D}_K^\pi; \mathcal{E}_K^\pi, \perp) &\leq \Delta_{q,m}(\mathcal{E}_K^\pi, \mathcal{D}_K^\pi; \mathcal{E}_K^{KpK}, \mathcal{D}_K^{KpK}) + \\ &\quad \Delta_{q,m}(\mathcal{E}_K^{KpK}, \mathcal{D}_K^{KpK}; \mathcal{E}_K^{KpK}, \perp) + \Delta_{q,m}(\mathcal{E}_K^{KpK}, \perp; \mathcal{E}_K^\pi, \perp) \\ &\leq 0 + \Delta_{q,m}(\mathcal{E}_K^{KpK}, \mathcal{D}_K^{KpK}; \mathcal{E}_K^{KpK}, \perp) + 0, \end{aligned}$$

where we use that π and KpK are identically distributed as π and p are random permutations and K is random and unknown. The middle term equals $\mathbf{Adv}_{\Pi'}^{\text{int}}(q, m)$ with the difference that the distinguisher cannot access p . A distinguisher would only benefit from such additional access, thus:

$$\Delta_{q,m}(\mathcal{E}_K^{KpK}, \mathcal{D}_K^{KpK}; \mathcal{E}_K^{KpK}, \perp) \leq \mathbf{Adv}_{\Pi'}^{\text{int}}(q, m),$$

which is bounded in Thm. 2. This completes the proof. \square

Proposition 2. *Let E be a block cipher.*

$$\mathbf{Adv}_{\Pi'}^{\overline{\text{cpa}}}(t, q, m) \leq \frac{m^2}{2^{r+c}} + \frac{2m(m+1)}{2^c} + \mathbf{Adv}_E^{\text{prpt1}}(t', \rho m).$$

Proof. The proof is a straightforward simplification of the proof of Prop. 1, and therefore omitted. \square

6 APE for Fractional Data

The APE description of Sect. 3 should only be used when the application can guarantee that the length of the plaintext and the associated data is always a multiple of the block size r . In this section, we explain how to adjust APE to handle fractional plaintext and associated data. This is done by applying ‘10*’-padding to all plaintext and associated data (fractional or not).

The extension of $\text{hash-data}_{0,K}$ to fractional associated data is given in Fig. 5, and for fractional messages the extension of enc-message_{IV} is given in Fig. 6. We elaborate on the extension for fractional messages (the extension for fractional associated data being similar). Consider a message M and denote its last block by $M[w]$, where $|M[w]| = |M| \bmod r$. We distinguish among three cases:

- $|M[w]| \leq r - 1$ and $w = 1$. The procedure can be seen in the top part of Fig. 6. Note that the corresponding ciphertext will be r bits. This is required for decryption to be possible;
- $|M[w]| \leq r - 1$ and $w \geq 2$. The procedure is depicted in the bottom part of Fig. 6. Note that the ciphertext $C[w - 1]$ is of size equal to $M[w]$. The reason we opt for this design property is the following: despite $M[w]$ being smaller than r bits, we require its corresponding ciphertext to be r bits for decryption to be possible. As a toll, the extended APE generates ciphertext $C[w - 1]$ to be of size equal to $M[w]$;
- $|M[w]| = r$. In this special case where M consists of integral message blocks, we nevertheless need a padding. However, instead of occupying an extra message block for this, the ‘10*’-padding spills over into the capacity. This can be seen as an XOR of $10 \cdots 00$ into the capacity part of the state. We recall the reader of the fact that the $\oplus 1$ in the beginning of the function is a shorthand notation for $\oplus 00 \cdots 01$, and hence, these values do not cancel each other out.

The adjustments have no influence on the decryption algorithm \mathcal{D} , except if $|M| \leq r$ for which a slightly more elaborate function is needed. Note that the spilling of the padding in case $|M[w]| = r$ causes security to degrade by half a bit: intuitively, APE is left with a capacity of $c' = c - 1$ bits. We have opted for this degrading over an efficiency loss due to an additional round.

For the purpose of the security analysis, we define the padding functions pad_{data} and $\text{pad}_{\text{message}}$ corresponding to the associated data and message. They output blocks of $r' = r + 1$ bits as suggested by Fig. 5 and Fig. 6, respectively. In more detail, $\text{pad}_{\text{message}}$ is defined as follows: (i) if $|M| \leq r$ then $\text{pad}_{\text{message}}(M) =$




Fig. 5. A generalization of APE that can handle fractional associated data blocks.

$M || 10^{r-|M|}$, (ii) if $|M| > r$, we derive the r -bit message blocks $M[1] \cdots M[w-1]$ and the $|M| \bmod r$ -bit message block $M[w]$, and set $\text{pad}_{\text{message}}$ to output the w message blocks $M[1] || 0, \dots, M[w-1] || 0$, and $M[w] || 10^{r-|M[w]|}$. The definition of pad_{data} is analogous. Using these injective padding functions, the privacy and integrity proofs of Thm. 1 and Thm. 2 carry over with minor modifications, we only highlight the differences.

Theorem 3. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be the APE construction for fractional data. Then,

$$\text{Adv}_{\Pi}^{\text{cpa}}(q, m) \leq \frac{m^2}{2^{r+c}} + \frac{m(m+1)}{2^{c-1}}.$$

Proof. Throughout this proof, for associated data A and message M , denote $pA = pA[1] || pA[2] || \cdots || pA[u] \leftarrow \text{pad}_{\text{data}}(A)$ and $pM = pM[1] || pM[2] || \cdots || pM[w] \leftarrow \text{pad}_{\text{message}}(M)$. Note that these are blocks of $r' = r+1$ bits, meaning that we have capacity $c' = c-1$. The proof of Thm. 1 now carries over using the new notation, with the difference that Lem. 1, bounding the distance between $(\mathcal{E}_K, f, f^{-1})$ and $(\$, f, f^{-1})$ as long as \mathbf{E} does not occur, needs closer attention as we need to consider different cases. However, ciphertext blocks are always of length r bits, except $C[w-1]$ if $w > 1$ and $|M[w]| \neq r$. But in the current setting \mathcal{D} has no access to the decryption oracle. Therefore, for simplicity we assume that all blocks $C[j]$ are of length r bits. This does not invalidate the argument. This makes the case $|M| \leq r$ a special case of $|M| > r$, and the analysis of Lem. 1 carries over directly, using pA and pM instead. \square

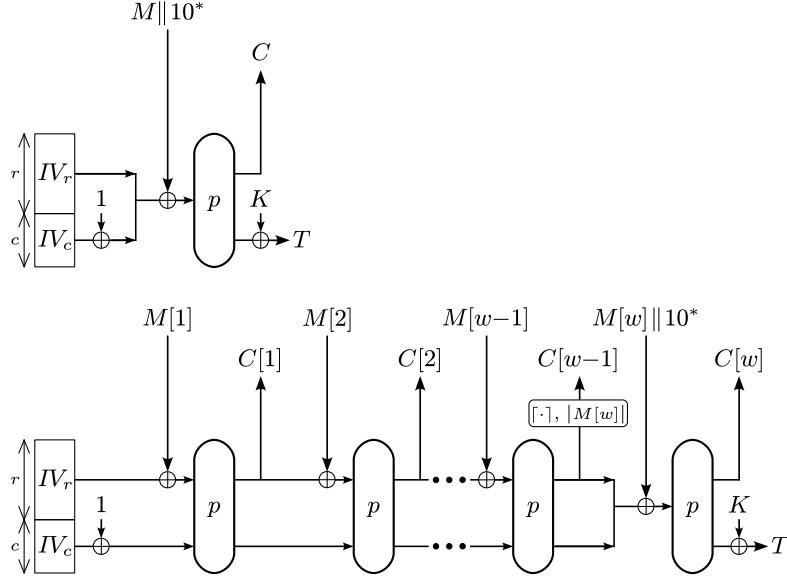


Fig. 6. A generalization of APE that can handle fractional message blocks.

Theorem 4. Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be the APE construction for fractional data. Then,

$$\mathbf{Adv}_{\Pi}^{\text{int}}(q, m) \leq \frac{m^2}{2^{r+c}} + \frac{2m(m+1)}{2^{c-1}}.$$

Proof. Denote $pA = pA[1] \| pA[2] \| \dots \| pA[u]$ and $pM = pM[1] \| pM[2] \| \dots \| pM[w]$ as in Thm. 3. Note that these are blocks of $r' = r + 1$ bits, meaning that we have capacity $c' = c - 1$. The proof of Thm. 2 now carries over using the new notation, with the difference that Lem. 3, bounding the distance between $(\mathcal{E}_K, \mathcal{D}_K, f, f^{-1})$ and $(\mathcal{E}_K, \perp, f, f^{-1})$ as long as \mathbf{E} does not occur, needs closer attention as we need to consider different cases. As in Lem. 3, it suffices to consider queries to the big oracle $F \in \{\mathcal{D}_K, \perp\}$. Let (A, C, T) be a query made by the distinguisher.

First consider $|C| > r$. The ciphertext blocks are always c bits, except for block $C[w-1]$ which may be shorter. For a ciphertext C denote $pC = pC[1] \| pC[2] \| \dots \| pC[w] \leftarrow \text{pad}_{\text{cipher}}(C)$, defined as follows: $pC[j] = C[j]$ for $j \in \{1, \dots, w-2, w\}$ and $pC[w-1] = C[w-1] \| 0^{r-|C[w-1]|}$. Using this notation, it is easy to see that the analysis of Lem. 3 carries over directly.

Next, assume $|C| = r$, hence the corresponding message block M is of size at most r bits. Clearly, in the ideal world $(\mathcal{E}_K, \perp, f, f^{-1})$, the query is responded with \perp . Therefore, the distinguisher has no advantage in the real world unless its query (A, C, T) is correct. Denote the state coming from $\text{hash-data}_{0,K}(A)$ by IV . For a state value x , denote its bottom $c' = c - 1$ bits by $[x]_{c'}$. The decryption query is successful if $[IV]_{c'} \oplus 1 = [f^{-1}(C, T \oplus K)]_{c'}$. We distinguish two cases:

- $(C, T) \notin \mathcal{Q}_i^{-1}$. Consider the indirect f^{-1} -query. As $(C, T) \notin \mathcal{Q}_i^{-1}$ this means that it has never been made as an indirect query before. Also, it does not match an older direct query $\oplus K$ (by $\overline{\mathbf{E}_i^{\text{ind-}Y}}$). This implies the indirect query $f^{-1}(C, T \oplus K)$ is fresh. If its bottom c' bits equal $[IV]_{c'} \oplus 1$, then this indirect query would have triggered $\mathbf{E}^{\text{ind-}Y}$.
- $(C, T) \in \mathcal{Q}_i^{-1}$. The analysis of Lem. 3 carries over.

This completes the proof. \square

7 Hardware Implementation

We implement APE with the permutations of PHOTON [18], QUARK [2], and SPONGENT [10]. The results are given in Table 1. We use these permutations without any modifications to investigate the hardware performance of APE. As the designs of PHOTON, QUARK and SPONGENT follow the hermetic sponge strategy [7], the underlying permutations are assumed to be indistinguishable from random permutations. This assumption is necessary in order to achieve the claimed privacy (Thm. 1) and integrity (Thm. 2) security bounds. Since APE is designed for constrained devices, we focus on a security level of 80 and 128 bits, which correspond to a capacity of 160 or 256 bits, respectively. One exception is APE based on QUARK: since QUARK is not equipped with a version for 128 bits of security we resort to a permutation that offers 112 bits of security. The versions of PHOTON and SPONGENT with 80 bits of security are implemented with a 4-bit serialization, which means that we implement one 4-bit S-box. For the versions with higher security, we use an 8-bit serialization which requires two 4-bit S-boxes for SPONGENT and one 8-bit S-box for PHOTON. Unlike PHOTON and SPONGENT, the round permutation of QUARK is based on Feedback Shift Registers (FSRs). Hence it is possible to update one bit per clock cycle, and in our implementation we choose to do so for area efficiency.

As APE decrypts in reverse order and requires the inverse permutation, for each of the algorithms (PHOTON, QUARK, and SPONGENT) we have provided both an encryption-only implementation and an implementation with encryption and decryption. In brief, we have implemented APE as follows. The initial state is XORed with the first data inserted nibble by nibble (or byte by byte, or bit by bit). After each permutation evaluation, the resulting ciphertext is output as the new data is inserted in the same clock cycle. At the end of the iteration, the entire state is output and the capacity part is XORed with the key to generate the tag. Similarly, for decryption, the first state corresponds to the ciphertext concatenated with an XOR of the key and the tag, and at the end authenticity is verified.

For the hardware implementation results in Table 1, we used ModelSim to verify the functionality of the designs and Synopsys Design Vision D-2010.03-SP4 for synthesis. We used Faraday Standard Cell Library based on UMC 0.18 μm and open-cell 45nm NANGATE [24] library. As main observations, we see that APE with an encryption and decryption mode can be implemented with less than 1.9 kGE and 2.8 kGE for 80-bit and 128-bit security respectively.

Table 1. APE is implemented using the PHOTON, QUARK, and SPONGENT permutations. For each algorithm, we provide an encryption-only implementation, as well as one that does both encryption and decryption (denoted as “e/d”). Our overview also includes lightweight implementations of the authenticated encryption schemes ALE [11], ASC-1 [19], and AES-CCM [30]. We remark that the clock frequency of the APE implementations is 100 kHz, compared to 20 MHz for the other ciphers.

APE on UMC 180 nm CMOS process @ 100 kHz					
Design	Security (bits)	Rate (bits)	Latency (cycles)	Throughput (kbps)	Area (GE)
PHOTON-196	80	36	1248	2.9	1398
PHOTON-196 e/d	80	36	1297	2.8	1634
QUARK-176	80	16	880	1.81	1694
QUARK-176 e/d	80	16	880	1.81	1871
SPONGENT-176	80	16	4050	0.4	1423
SPONGENT-176 e/d	80	16	4094	0.4	1868
PHOTON-288	128	32	924	3.45	2154
PHOTON-288 e/d	128	32	960	3.33	2449
QUARK-256	112	32	1270	2.51	2286
QUARK-256 e/d	112	32	1270	2.51	2470
SPONGENT-272	128	16	4480	0.4	2105
SPONGENT-272 e/d	128	16	4652	0.3	2781

APE on NANGATE 45 nm CMOS process @ 100 kHz					
Design	Security (bits)	Rate (bits)	Latency (cycles)	Throughput (kbps)	Area (GE)
PHOTON-196	80	36	1248	2.9	1309
PHOTON-196 e/d	80	36	1297	2.8	1536
QUARK-176	80	16	880	1.81	1606
QUARK-176 e/d	80	16	880	1.81	1773
SPONGENT-176	80	16	4050	0.4	1598
SPONGENT-176 e/d	80	16	4094	0.4	1838
PHOTON-288	128	32	924	3.45	2104
PHOTON-288 e/d	128	32	960	3.33	2327
QUARK-256	112	32	1270	2.51	2228
QUARK-256 e/d	112	32	1270	2.51	2331
SPONGENT-272	128	16	4480	0.4	2378
SPONGENT-272 e/d	128	16	4652	0.3	2661

Other AE schemes on ST 65 nm CMOS LP-HVT process @ 20 MHz [11]					
Design	Security (bits)	Rate (bits)	Latency (cycles)	Throughput (kbps)	Area (GE)
ALE	128	N/A	105	121.9	2579
ALE e/d	128	N/A	105	121.9	2700
ASC-1 A	128	N/A	370	34.59	4793
ASC-1 A e/d	128	N/A	370	34.59	4964
ASC-1 B	128	N/A	235	54.47	5517
ASC-1 B e/d	128	N/A	235	54.47	5632
AES-CCM	128	N/A	452	28.32	3472
AES-CCM e/d	128	N/A	452	28.32	3765

When implemented with the same permutation, encryption-only implementations of SpongeWrap and APE will have similar implementation figures. This is because in both constructions, the processing of every message block requires one XOR with the rate part and one permutation function call. We recall that the crucial difference between SpongeWrap and APE is that APE provides nonce misuse resistance. For the decryption operation, the cost of misuse resistance for APE is that the backwards permutation must be implemented as well.

As shown in Table 1, the overhead of decryption for APE ranges from 4 to 17 percent, which corresponds to at most 0.7 kGE. For devices without non-volatile memory, this overhead is very low compared to the cost of providing a hardware source of randomness to generate nonces.

Note that the permutation-based schemes are implemented on 180 nm and 45 nm CMOS, whereas for the block cipher based schemes, lightweight implementations on 65 nm CMOS are provided. Therefore, we cannot compare these implementations directly. Also note that the clock frequencies of the implementations differ, which lead to different throughput figures. However, it seems that APE and ALE have similar performance figures and APE is smaller than ASC-1 A, ASC-1 B and AES-CCM.

8 Conclusions

In this paper, we introduced APE, the first misuse resistant permutation-based AE scheme. We proved that APE provides security and integrity up to the birthday bound of the capacity, in the ideal permutation model *as well as* in the standard model. This not only ensures security of APE when its underlying primitive is considered an ideal permutation, but also allows to employ it with any secure block cipher of specific form. To achieve misuse resistance, the decryption of APE as a permutation-based construction uses the inverse permutation to decrypt in a backwards manner. The advantage of having backwards decryption is that if the tag or last ciphertext block is missing, then decryption is impossible. Our hardware implementations of APE show that it is well-suited for lightweight applications. In fact, using any of the permutations of QUARK, PHOTON, and SPONGENT, less than 1.9 kGE (80-bit security) and less than 2.8 kGE (128-bit security) is required for an implementation of APE that supports both encryption and decryption. Due to its resistance against nonce reuse and its low area requirements in hardware, APE is suitable for environments where it is prohibitively expensive to require non-volatile memory or a hardware source of randomness.

ACKNOWLEDGMENTS. We would like to thank the various anonymous reviewers for providing useful comments. This work was supported in part by the Research Council KU Leuven: GOA TENSE (GOA/11/007). In addition, this work was supported by the Research Fund KU Leuven, OT/13/071. Elena Andreeva and Nicky Mouha are supported by Postdoctoral Fellowships from the

Flemish Research Foundation (FWO-Vlaanderen). Atul Luykx and Bart Menink are supported by Ph.D. Fellowships from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

References

1. Aumasson, J.P., Henzen, L., Meier, W., Naya-Plasencia, M.: Quark: A Lightweight Hash. In: Mangard, S., Standaert, F.X. (eds.) CHES. Lecture Notes in Computer Science, vol. 6225, pp. 1–15. Springer (2010)
2. Aumasson, J.P., Henzen, L., Meier, W., Naya-Plasencia, M.: Quark: A Lightweight Hash. *J. Cryptology* 26(2), 313–339 (2013)
3. Bellare, M., Boldyreva, A., Knudsen, L.R., Namprempre, C.: On-line Ciphers and the Hash-CBC Constructions. *J. Cryptology* 25(4), 640–679 (2012)
4. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT. Lecture Notes in Computer Science, vol. 4004, pp. 409–426. Springer (2006)
5. Berger, T.P., D’Hayer, J., Marquet, K., Minier, M., Thomas, G.: The GLUON Family: A Lightweight Hash Function Family Based on FCSRs. In: Mitrokotsa, A., Vaudenay, S. (eds.) AFRICACRYPT. Lecture Notes in Computer Science, vol. 7374, pp. 306–323. Springer (2012)
6. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Permutation-Based Encryption, Authentication and Authenticated Encryption. Directions in Authenticated Ciphers (July 2012)
7. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Cryptographic Sponge Functions, available at <http://sponge.noekeon.org/CSF-0.1.pdf>
8. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Sponge Functions. ECRYPT Hash Function Workshop (May 2007)
9. Bertoni, G., Daemen, J., Peeters, M., Van Assche, G.: Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications. In: Miri, A., Vaudenay, S. (eds.) Selected Areas in Cryptography 2011. Lecture Notes in Computer Science, vol. 7118, pp. 320–337. Springer (2012)
10. Bogdanov, A., Knezevic, M., Leander, G., Toz, D., Varici, K., Verbauwhede, I.: SPONGENT: A Lightweight Hash Function. In: Preneel, B., Takagi, T. (eds.) CHES. Lecture Notes in Computer Science, vol. 6917, pp. 312–325. Springer (2011)
11. Bogdanov, A., Mendel, F., Regazzoni, F., Rijmen, V., Tischhauser, E.: ALE: AES-based lightweight authenticated encryption. In: Moriai, S. (ed.) FSE. Lecture Notes in Computer Science, Springer (2013)
12. Borisov, N., Goldberg, I., Wagner, D.: Intercepting mobile communications: the insecurity of 802.11. In: Rose, C. (ed.) MOBICOM. pp. 180–189. ACM (2001)
13. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness (April 2013)
14. Cantero, H.M., Peter, S., Bushing, Segher: Console Hacking 2010 – PS3 Epic Fail. 27th Chaos Communication Congress (December 2010)
15. Chang, D., Dworkin, M., Hong, S., Kelsey, J., Nandi, M.: A Keyed Sponge Construction with Pseudorandomness in the Standard Model. The Third SHA-3 Candidate Conference (March 2012)
16. Even, S., Mansour, Y.: A Construction of a Cipher from a Single Pseudorandom Permutation. *J. Cryptology* 10(3), 151–162 (1997)

17. Fleischmann, E., Forler, C., Lucks, S.: McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes. In: Canteaut, A. (ed.) FSE. Lecture Notes in Computer Science, vol. 7549, pp. 196–215. Springer (2012)
18. Guo, J., Peyrin, T., Poschmann, A.: The PHOTON Family of Lightweight Hash Functions. In: Rogaway, P. (ed.) CRYPTO. Lecture Notes in Computer Science, vol. 6841, pp. 222–239. Springer (2011)
19. Jakimoski, G., Khajuria, S.: ASC-1: An Authenticated Encryption Stream Cipher. In: Miri, A., Vaudenay, S. (eds.) Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 7118, pp. 356–372. Springer (2011)
20. Khovratovich, D.: Key Wrapping with a Fixed Permutation. Cryptology ePrint Archive, Report 2013/145 (2013)
21. Kohno, T.: Attacking and Repairing the WinZip Encryption Scheme. In: Atluri, V., Pfizmann, B., McDaniel, P.D. (eds.) ACM Conference on Computer and Communications Security. pp. 72–81. ACM (2004)
22. Lenstra, A.K., Hughes, J.P., Augier, M., Bos, J.W., Kleinjung, T., Wachter, C.: Public Keys. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO. Lecture Notes in Computer Science, vol. 7417, pp. 626–642. Springer (2012)
23. McGrew, D.A., Viega, J.: The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT. Lecture Notes in Computer Science, vol. 3348, pp. 343–355. Springer (2004)
24. NANGATE: The NanGate 45nm Open Cell Library, available at <http://www.nangate.com>
25. Rogaway, P.: Authenticated-Encryption with Associated-Data. In: Atluri, V. (ed.) ACM Conference on Computer and Communications Security 2002. pp. 98–107. ACM (2002)
26. Rogaway, P.: Nonce-Based Symmetric Encryption. In: Roy, B.K., Meier, W. (eds.) FSE 2004. Lecture Notes in Computer Science, vol. 3017, pp. 348–359. Springer (2004)
27. Rogaway, P., Bellare, M., Black, J.: OCB: A Block-Cipher Mode of Operation for Efficient Authenticated Encryption. ACM Trans. Inf. Syst. Secur. 6(3), 365–403 (2003)
28. Rogaway, P., Shrimpton, T.: A Provable-Security Treatment of the Key-Wrap Problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. Lecture Notes in Computer Science, vol. 4004, pp. 373–390. Springer (2006)
29. Rogaway, P., Zhang, H.: Online Ciphers from Tweakable Blockciphers. In: Kiayias, A. (ed.) CT-RSA. Lecture Notes in Computer Science, vol. 6558, pp. 237–249. Springer (2011)
30. Whiting, D., Housley, R., Ferguson, N.: Counter with CBC-MAC (CCM). Request For Comments 3610 (2003)
31. Wu, H.: The Misuse of RC4 in Microsoft Word and Excel. Cryptology ePrint Archive, Report 2005/007 (2005)