

Forking a Blockcipher for Authenticated Encryption of Very Short Messages

Elena Andreeva¹, Reza Reyhanitabar², Kerem Varici¹ and Damian Vizár³

¹ imec-COSIC, KU Leuven, Belgium

`elena.andreeva@esat.kuleuven.be`, `kerem.varici@esat.kuleuven.be`

² Elektrobit Automotive GmbH, Germany

`reza.reyhanitabar@elektrobit.com`

³ CSEM, Switzerland

`damian.vizar@csem.ch`

Abstract. Highly efficient encryption and authentication of *short* messages has been identified as an essential requirement for enabling security in constrained computation and communication scenarios such as the CAN FD in automotive systems (with maximum message length of 64 bytes), massive IoT and critical communication domains of 5G, and Narrowband IoT (NB-IoT), to mention some. Accordingly, NIST has specified, as a design requirement in the lightweight cryptography project, that AEAD submissions shall be “optimized to be efficient for short messages (e.g., as short as 8 bytes)”. We propose AEAD schemes that exceed in efficiency over all previous general-purpose modular AEAD designs at processing (very) short inputs. The main ingredient in our solution is a new low-level primitive, called a tweakable *forkcipher*, which we introduce and formalize in this paper. We give an instance of the tweakable forkcipher and dub it ForkAES. It is based on the tweakable blockcipher KIASU, which relies on the round function of AES and uses the TWEAKEY framework to derive round keys from a 128-bit secret key and a 64-bit tweak. Finally, we demonstrate the applicability of a tweakable forkcipher by designing several provably-secure nonce-based AEAD modes of operation, optimized to be efficient for short messages. Considering the AES block size (16 bytes) as a reference, our new AE schemes can beat *all* known schemes for single-block messages while still performing better than majority of the existing schemes for combined message and associated data lengths up to 4 blocks. While ForkAES as a concrete instantiation for a forkcipher is based on KIASU, we note that our solution provides a general recipe for lightweight AEAD for short messages, even for very resource-constrained scenarios in which AES may not be considered a lightweight option. In those environments, our schemes can be instantiated using a forkcipher that is realized based on the best off-the-shelf lightweight blockcipher, following the TWEAKEY framework.

Keywords: Authenticated encryption, short messages, lightweight cryptography, forkcipher, ForkAES.

1 Introduction

Authenticated encryption (AE) aims at achieving two fundamental security goals of symmetric-key cryptography: confidentiality (privacy) and integrity (together with authentication) for data at rest and data in transit.

Historically, these two goals were achieved by generic composition of an encryption scheme (for confidentiality) and a message authentication code (MAC) [12,13]. For instance, *old* versions of major security protocols such as TLS, SSH and IPsec included variants of

generic composition method, namely MAC-then-Encrypt, Encrypt-and-MAC and Encrypt-then-MAC schemes, respectively. But it turned out that this approach is neither the most efficient (as it needs processing the whole message twice) nor the most robust to security and implementation issues [11, 35, 37, 38]; rather it is easy for practitioners to get it wrong even when using the best known method among three, i.e. Encrypt-then-MAC, following standards [35].

The notion of AE as a primitive in its own right—integrating encryption and authentication functionalities and exposing a single abstract interface—put forth by Bellare and Rogaway [15] and independently by Katz and Yung [31] in 2000. The notion was further enhanced by Rogaway [40] to authenticated encryption with associated data (AEAD). Being able to process associated data (AD) is now a default requirement for any authenticated encryption scheme; therefore we use AE and AEAD interchangeably. After attracting near two decades of research and standardization activities, recently fostered by the CAESAR competition (2014–2018) [17], we have now a rich set of *general-purpose* AEAD schemes, some already standardized (e.g. GCM and CCM) and some expected to be adopted by new applications and standards (e.g. the seven finalists of CAESAR).

This progress may lead to the belief that the AEAD problem is “solved”. However, as evidenced by the ECRYPT-CSA report in 2017 [8], several critical ongoing “Challenges in Authenticated Encryption” still need research effort stretching years into the future. Therefore, it is interesting to investigate to what extent CAESAR has resulted in solutions for these problems.

OUR TARGET CHALLENGE. Among the four categories of challenges—security, interface, performance, mistakes and malice—reported by the ECRYPT-CSA [8], we aim at digging into the performance regarding authenticated encryption of *very short messages*. General-purpose AEAD schemes are normally optimized for the cost of handling (moderately) long messages as they often incur some initialization and/or finalization cost that is amortized when the message is long. To quote the provocative statement of the ECRYPT-CSA report: “The performance target is wrong . . . Another increasingly common scenario is that an authenticated cipher is applied to many small messages . . . The challenge here is to minimize overhead.”

Therefore, designing efficient AEAD for short messages is a timely and compelling objective as also evidenced by NIST’s first call for submissions (May 14, 2018) for lightweight cryptography [36], where it is stressed as a *design requirement* that lightweight AEAD submissions shall be “optimized to be efficient for short messages (e.g., as short as 8 bytes)”.

MOTIVATING USE CASES. The need for high-performance and low-latency processing of short messages can be identified as an essential requirement in several security and safety critical use cases in different domains. Examples are Secure Onboard Communication (SecOC) in automotive systems [6], handling of short data bursts in critical communication and massive IoT domains of 5G [1], and Narrowband IoT (NB-IoT) [2, 5] systems. For instance, the new CAN FD standard (ISO 11898-1) for vehicle bus technology [3, 4], which is expected to be implemented in most cars by 2020, allows for a payload up to 64 bytes. In NB-IoT standard [2, 5] the maximum transport block size (TBS) is 680 bits in downlink and 1000 bits in uplink (the minimum TBS size is 16 bits in both cases). In use cases with tight requirements on delay and latency, the typical packet sizes should be small as large packets occupy a link for more time, causing more delays to subsequent packets and increasing latency. In use cases such as smart parking lots the actual data to be sent is just a bit (for “free” or “occupied” status), so a minimum allowed TBS size of 2 bytes (16 bits) would suit the application.

OUR GOAL. Our main objective is to construct secure, modular AEAD schemes that exceed in efficiency over previous modular AEAD constructions at processing very short inputs, while also being able to process longer inputs, albeit somewhat less efficiently. We

insist that our AEAD schemes ought to be able to securely process inputs of arbitrary lengths to be fairly comparable to other general-purpose (long message focused) schemes, and to be qualified as a full-fledged variable-input-length AEAD scheme according to the requirements in NIST’s call for lightweight cryptography primitives.

Towards this goal, we take an approach that can be seen as a parallel to the shift from generic composition to dedicated AEAD designs, but on the level of the primitive. We rethink the way a low level fixed-input-length (FIL) primitive is designed, and how variable-input-length (VIL) AEAD schemes are constructed from such a new primitive.

THE GAP BETWEEN THE PRIMITIVES AND AE. Our first observation is that there is a large gap between the high level security goal to be achieved by the VIL AEAD schemes and the security properties that their underlying FIL primitives can provide. Modular AEAD designs typically confine the AE security to the mode of operation only; the used lower-level primitives, such as (tweakable) block ciphers, cryptographic permutations and compression functions, are never meant to possess any AE-like features, and in particular they are never expanding as required for providing ciphertext integrity in AEAD. Therefore, a VIL AEAD scheme Π designed as a mode of operation for an FIL primitive F plays two roles: not only it extends the domain of its underlying FIL primitive but also it transforms and boosts the security property of the primitive to match to the AEAD security notion. A natural question then arises, whether by explicitly decoupling these two roles of the AEAD mode we can have more efficient designs and more transparent security proofs.

FORKCIIPHER: A NEW PRIMITIVE. The first, most obvious approach to resolving this question is to remove the security gap between the mode and its primitive altogether, i.e., to start from a FIL primitive F which itself is a secure FIL AEAD. This way a VIL AEAD mode will only have one role: a property-preserving domain extender for the primitive F . Property-preserving domain extension is a well-studied and popular design paradigm for other primitives such as hash functions [7, 14, 39]. Informally speaking, the best possible security that a FIL AEAD scheme with a *fixed* ciphertext expansion (stretch) can achieve is to be indistinguishable from a *tweakable random injective* function, i.e., to be a tweakable pseudorandom injection (PRI) [26, 42]. But starting directly with a FIL tweakable PRI, we did not achieve a desirable solution in our quest for the most *efficient AEAD design for short messages*.¹ It seems that, interestingly, narrowing the security gap between the mode and its primitive, but not removing the gap entirely, is what helps us achieve our ultimate goal of efficient AEAD for short messages.

We introduce a different kind of low-level primitive—calling it a tweakable **forkcipher**—that does yield the most efficient AEAD design for short messages. A tweakable forkcipher is *nearly*—but not exactly—a FIL AE primitive; “nearly” because it produces *expanded* ciphertexts with a non-trivial redundancy, and not exactly because it has no integrity-checking mechanisms.² When keyed and tweaked, a forkcipher maps an n bit input block to $2n$ output bits. Intuitively, evaluating a secure tweakable forkcipher on an input X is equivalent to evaluating *two independent* tweakable permutations on X but with an *amortized computational cost*.

We give an instance of the tweakable forkcipher and dub it ForkAES. It is based on the tweakable blockcipher KIASU [28, 30], which relies on the round function of AES and uses the TWEAKEY framework to derive round keys from a 128-bit secret key and a 64-bit tweak. To obtain ForkAES, we apply our newly proposed *iterate-fork-iterate* paradigm: when encrypting a block X , we derive 17 ($3 \cdot 5 + 2$) round keys, and first transform X into X' using 5 AES rounds with 5 round keys. Then, we fork the encryption process by applying two parallel paths each comprising 5 AES rounds with 5 round keys, followed by

¹See Section 5.10 for a brief discussion.

²We can demonstrate that when used in a minimalistic mode of operation, a secure tweakable forkcipher yields a miniature FIL AEAD scheme which achieves tweakable PRI security.

the final whitenings using the remaining two sub-keys. The iterate-fork-iterate paradigm is conceptually easy and we give arguments about its security.

THE NEW AEAD SCHEMES. We demonstrate the applicability of tweakable forkciphers by designing three provably secure nonce-based AEAD modes of operation, all suitable for short messages, but having some different features.

Our first scheme, called **PAEF** (Parallel AE from a forkcipher), is fully parallelizable, but its security completely falls apart with the first nonce repetition. This mode will be useful in applications where longer queries occur more often, and where one of the parties is able to perform parallel computations.

Our second scheme, called **SAEF** (Sequential AE from a forkcipher), has an online computable encryption, but is not parallelizable. While its integrity and confidentiality bounds in the nonce-respecting setting are inferior to those of the parallel mode, we conjecture that a nonce-misusing universal forgery will require data complexity at the birthday bound. In addition, SAEF also lends itself well to low-overhead implementations, as it does not need to store any block counters. In a typical setting where very short queries occur with an overwhelming probability, this mode is our primary recommendation.

Our third scheme, called **fGCM** (forkcipher-based GCM), is inspired by the blockcipher-based standard scheme GCM, but differs from it in some critical details. It is almost fully parallelizable, and it beats the other two modes in efficiency for plaintexts of more than n bits, but this requires an n -bit increase of the memory footprint to store a derived key in order to achieve full efficiency.

PERFORMANCE COMPARISON: NO ONE-SIZE-FITS-ALL MODE. Regarding the state of the art in AE designs, it appears that aiming for a *provably secure* AE scheme that achieves the best performance for both very long and very short message scenarios is a very ambitious goal. Instead we aimed at addressing the specific problem of high-performance AE for very short inputs. All three proposed modes, PAEF, SAEF and fGCM, can be efficiently implemented when instantiated with ForkAES. The ForkAES-based instances of our schemes beat all of the existing blockcipher-based AEAD modes when instantiated with AES, for the shortest queries. The performance comparison (in terms of the number of AES128 calls) of encrypting very short queries (up to 4 blocks) with AES-GCM [23,33], AES-CCM [44], AES-OCB [32], AES-CLOC [27], Deoxys-I [29] and KIASU[≠] [28] is summarized in Figure 1. Our new AE schemes can beat all known schemes for single-block messages while still performing better than most of the existing schemes up to 4-block messages, and then leaving the competition to some of the general-purpose schemes such as AES-OCB.

TO BE MODULAR AND PROVABLY SECURE, OR NOT TO BE. In this work we followed the well-established modular AE design approach for data of arbitrary lengths in the provable security framework. However, we note that there seems to be no consensus in the cryptography community whether AE schemes can claim higher merits for being modular and provably secure or not. For instance, we note that 3 out of 7 CAESAR [17] finalists, namely ACORN, AEGIS and MORUS are novel monolithic designs from scratch and do not follow the provable security paradigms. Nevertheless, we choose as our design paradigm the modular and provable security methodology for its several well-advertised benefits [9,41]. Interestingly, the class of provably secure AE designs includes all currently standardized AE schemes (e.g. GCM and CCM) as well as the majority of CAESAR finalists.

RECIPE FOR LIGHTWEIGHT AEAD FOR SHORT MESSAGES. For very resource constrained IoT devices in which AES could not be considered a lightweight option—which seems to be a driving motivation behind the current NIST lightweight cryptography standardization—our proposed SAEF, PAEF and fGCM modes can be instantiated using a forkcipher

scheme	Cost of encryption in (# of AES rounds)/10									
	a = 0				a = 1				a = 2	
	m=1	m=2	m=3	m=4	m=0	m=1	m=2	m=3	m=1	m=2
GCM	2+2	3+3	4+4	5+5	1+2	2+3	3+4	4+5	2+4	3+5
CCM	4	6	8	10	3	5	7	9	6	8
OCB3	3	4	5	6	3	4	5	6	5	6
CLOC	3	5	7	9	2	4	6	8	5	7
Deoxys-I	2.8	4.2	5.6	7	2.8	4.2	5.6	7	5.6	7
KIASU [≠]	2	3	4	5	2	3	4	5	4	5
PAEF	1.5	3	4.5	6	1.5	2.5	4	5.5	3.5	5
SAEF	1.5	3	4.5	6	1.5	2.5	4	5.5	3.5	5
fGCM	1.5+1	2.5+2	3+3	5+4	1.5+2	1.5+2	2.5+3	3+4	1.5+3	2.5+4

Figure 1: The comparison of performance of several existing single-pass AE schemes and the forkcipher modes presented in Section 5. Here $a = |A|_n$ and $m = |M|_n$, and the cells of the form “ $b + g$ ” mean “ $b \cdot 10$ AES rounds and g multiplications in $\text{GF}(2^{128})$.” GCM, CCM, OCB3 and CLOC are assumed to be instantiated with AES. For schemes that compute a derived key that is the same for all queries (OCB and GCM), the complexity of this step is ignored.

which is based on any off-the-shelf lightweight blockcipher. The crux would be a careful realization of the forkcipher following the TWEAKEY framework.

2 Preliminaries

All strings are binary strings. The set of all strings of length n (for a positive integer n) is denoted $\{0, 1\}^n$. We let $\{0, 1\}^{\leq n}$ denote the set of all strings of length at most n . We denote by $\text{Perm}(n)$ the set of all permutations of $\{0, 1\}^n$. We denote by $\text{Func}(m, n)$ the set of all functions with domain $\{0, 1\}^m$ and range $\{0, 1\}^n$, and we let $\text{Inj}(m, n) \subset \text{Func}(m, n)$ denote the set of all injective functions with the same signature.

For a string X of ℓ bits, we let $X[i]$ denote the i^{th} bit of X for $i = 0, \dots, \ell - 1$ (starting from the left) and $X[i \dots j] = X[i] \| X[i + 1] \| \dots \| X[j]$ for $0 \leq i < j < \ell$. We let $\text{left}_\ell(X) = X[0 \dots (\ell - 1)]$ denote the ℓ leftmost bits of X and $\text{right}_r(X) = X[(|X| - r) \dots (|X| - 1)]$ the r rightmost bits of X , such that $X = \text{left}_\chi(X) \| \text{right}_{|X| - \chi}(X)$ for any $0 \leq \chi \leq |X|$. We let $(L, R) = \text{lsplit}_{X, n}$ denote splitting a string $X \in \{0, 1\}^*$ into two parts such that $L = \text{left}_{\min(|X|, n)}(X)$ and $R = \text{right}_{|X| - |L|}(X)$. In particular, for $n \geq |X|$ we have $(X, \varepsilon) = \text{lsplit}_{X, n}$. We further let $(M', M_*) = \text{msplit}_n(M)$ denote a splitting of a string $M \in \text{bits}^*$ into two parts $M' \| M_* = M$, such that $|M_*| \equiv |M| \pmod{n}$, and $0 \leq |M_*| \leq n$, where $|M_*| = 0$ if and only if $|M| = 0$. We let $(C', C_*, T) = \text{csplit}_n(C)$ splitting a string C of at least n bits into three parts $C' \| C_* \| T = C$, such that $|C_*| = n$, $|T| \equiv |C| \pmod{n}$, and $0 \leq |T| \leq n$, where $|T| = 0$ if and only if $|C| = n$. Finally, we let $C'_1, \dots, C'_m, C_*, T \leftarrow \text{csplit}_{\mathbf{b}_n}(C)$ denote a version of $\text{csplit}_n(C)$, where the string C' further gets partitioned into $|C'|_n$ blocks of n bits, such that $C' = C'_1 \| \dots \| C'_m$.

Given a string X and an integer n , we let $X_1, \dots, X_x, X_* \stackrel{n}{\leftarrow} X$ denote partitioning X into n -bit blocks, such that $|X_i| = n$ for $i = 1, \dots, x$, $0 \leq |X_*| \leq n$ and $X = X_1 \| \dots \| X_x \| X_*$, so $x = \max(0, \lfloor X/n \rfloor - 1)$. We let $|X|_n = \lceil X/n \rceil$. Given a (possibly implicit) positive integer n and an $X \in \{0, 1\}^*$, we let $X \| 10^*$ denote $X \| 10^{n - (|X| \bmod n) - 1}$ for simplicity.

The symbol \perp denotes an error signal, or an undefined value. We denote by $X \leftarrow_{\mathcal{S}}$ sampling an element X from a finite set \mathcal{X} following the uniform distribution. We let $(n)_q$ denote the falling factorial $n \cdot (n - 1) \cdot (n - 2) \cdot \dots \cdot (n - q + 1)$.

3 ForkAES

We design an expanding primitive with fixed input length which we dub ForkAES. It is best described by its name: it is an iterated design with a 128-bit input block, and its internal state is *forked* after half of the rounds to produce *two* redundant 128-bit output blocks. We also add a tweak to facilitate the design of simple modes of operation. ForkAES is obtained by combining two ingredients: the KIASU [28,30] tweakable blockcipher (which is, in turn, a derivative of AES, hence the name), and our newly proposed *iterate-fork-iterate* paradigm.

3.1 Specification

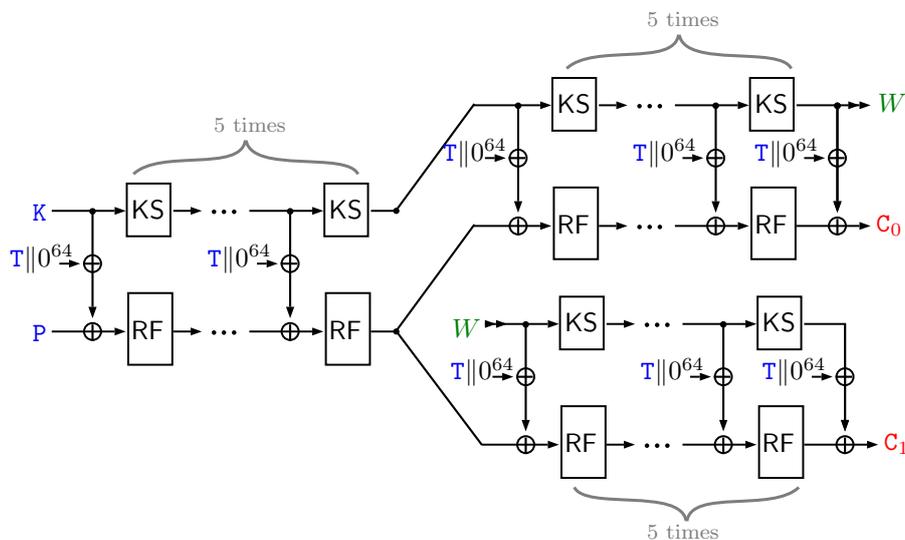


Figure 2: Illustration of an encryption by ForkAES. A 128 bit plaintext P , a 128 bit key K and 64 bit tweak T (all in blue) are used to compute a 256 bit ciphertext $C = C_0 || C_1$ (in red). RF denotes a single iteration of the AES round function and KS denotes a single iteration of the AES keyschedule.

ForkAES is a deterministic cryptographic algorithm which takes a 128-bit plaintext P , a 64-bit tweak T and a 128-bit secret key K as input, and outputs a 256-bit ciphertext C (i.e., $\text{ForkAES}(K, T, P) = C$).

It is based on the tweakable blockcipher KIASU. In KIASU, a round function based on the SubBytes, Shiftrows and Mixcolumn operations of AES is iteratively applied to the plaintext block. Following the TWEAKEY framework [30], the secret key and tweak are used to generate subkeys which are xored to the intermediate internal state before every application of the round function.

Iterate-fork-iterate. How ForkAES differs from both AES and KIASU is that after half of the rounds, the encryption is forked and two copies of the internal states are further processed with different sets of *independent* subkeys. The additional required subkeys are generated by doing the necessary number of extra iterations of the key schedule (beyond what would have been done in the original (tweakable) blockcipher). We call this simple mechanism of turning an iterated (tweakable) blockcipher into a forkcipher the *iterate-fork-iterate* construction paradigm.

Inverse algorithms. Associated to ForkAES are the decryption algorithm ForkAES^{-1} and the *reconstruction* algorithm ForkAES^{ρ} . Because the two output blocks produced by ForkAES are redundant, either one of them is sufficient for decryption. The decryption algorithm thus takes a secret key K , a tweak T , a half-ciphertext C of 128 bits, and a bit b that indicates whether this is the left half or the right half, and inverts the “fork” indicated by b and then the initial common processing. For every $K, P \in \{0, 1\}^{128}$ and every $T \in \{0, 1\}^{64}$ we have

$$P = \text{ForkAES}^{-1}(K, T, \text{left}_n(\text{ForkAES}(K, T, P)), 0) = \text{ForkAES}^{-1}(K, T, \text{right}_n(\text{ForkAES}(K, T, P)), 1).$$

Similarly, the redundancy can be used to recompute one output block from the other which is what the reconstruction algorithm does. It takes a secret key K , a tweak T , a half-ciphertext C of 128 bits, and a bit b that indicates whether this is the left half or the right half, inverts the indicated “fork”, and then recomputes the other one. For every $K, P \in \{0, 1\}^{128}$ and every $T \in \{0, 1\}^{64}$ we have

$$\text{ForkAES}^{\rho}(K, T, \text{left}_n(\text{ForkAES}(K, T, P)), 0) = \text{right}_n(\text{ForkAES}(K, T, P))$$

and

$$\text{ForkAES}^{\rho}(K, T, \text{right}_n(\text{ForkAES}(K, T, P)), 1) = \text{left}_n(\text{ForkAES}(K, T, P))$$

The formal algorithmic description of all three algorithms is given in Figure 3, and the encryption operation is illustrated in Figure 2.

To generate of round keys, we set the secret key as the first round key, iterate the key schedule of AES 16 times, and xor the tweak to the 8 leftmost bytes of each round key. This is exactly what is done in KIASU, except we iterate the key schedule 6 more times. The round key generation algorithm is described in Figure 3.

3.2 Security Evaluation

In this section, we briefly discuss the security of ForkAES against the most important cryptanalytic attacks. We only consider classical black-box attacks, i.e., we do not consider side-channel attacks.

Differential Cryptanalysis. Differential cryptanalysis is one of the most powerful security analysis methods and showing the security of a cipher against it is essential part of the security evaluation. For a cipher based on the Substitution Permutation Network (SPN) the analysis is relatively easy and well-understood and it is based on counting the number of active s-boxes over the cipher rounds. When the active s-boxes reach a certain threshold then the cipher is assumed to be secure against differential cryptanalysis. For example, in the case of AES in the single-key model, one can guarantee at least 25 active s-boxes for a differential path of four rounds due to the careful choice of a permutation layer (which is a diffusion matrix with branching number five). If each active s-box reaches the maximal differential probability of the AES S-box $p_{max} = 2^{-6}$, then the probability of the differential path becomes $2^{-150} < 2^{-128}$. Hence, four AES rounds already provide enough protection. Since our ForkAES design uses the AES round function, we can easily deduce that our design will provide enough security in this setting after four rounds against differential attacks in the single-key model.

Related-TWEAKEY Attacks. The extra freedom provided from key K (in our case tweak T as well) makes the security evaluation of ciphers against related-key (in our case related-tweakey) attacks more challenging. Over the years, many search algorithms [18, 19, 24, 25, 34, 43] were given to compute an upper bound for the related-key differential characteristics. The KIASU designers gave a comprehensive related-key analysis for KIASU by extending the search algorithms to cover the related-tweak option and we summarize their results in Table 1. Our design is based on the KIASU algorithm and its tweakey schedule and thus a closer inspection reveals that the latter results also apply to ForkAES.

<pre> 1: Algorithm ForkAES(K, T, P) 2: $K_0, \dots, K_{16} \leftarrow \text{KeySched}(K, T)$ 3: $S \leftarrow P$ 4: for $i \leftarrow 0$ to 4 do 5: $S \leftarrow S \oplus K_i$ 6: $S \leftarrow \text{AESrnd}(S)$ 7: end for 8: $S_0 \leftarrow S; S_1 \leftarrow S$ 9: for $i \leftarrow 5$ to 9 do 10: $S_0 \leftarrow S_0 \oplus K_i$ 11: $S_0 \leftarrow \text{AESrnd}(S_0)$ 12: end for 13: $C_0 \leftarrow S_0 \oplus K_{10}$ 14: for $i \leftarrow 11$ to 15 do 15: $S_1 \leftarrow S_1 \oplus K_i$ 16: $S_1 \leftarrow \text{AESrnd}(S_1)$ 17: end for 18: $C_1 \leftarrow S_1 \oplus K_{16}$ 19: return $C_0 \parallel C_1$ 20: end Algorithm </pre>	<pre> 6: $S \leftarrow \text{AESrnd}^{-1}(S)$ 7: $S \leftarrow S \oplus K_i$ 8: end for 9: for $i \leftarrow 5 + b' \cdot 6$ to $9 + b' \cdot 6$ do 10: $S \leftarrow S \oplus K_i$ 11: $S \leftarrow \text{AESrnd}(S)$ 12: end for 13: $C' \leftarrow S \oplus K_{10+b' \cdot 6}$ 14: return C' 15: end Algorithm </pre>
<pre> 1: Algorithm ForkAES⁻¹(K, T, C, b) 2: $K_0, \dots, K_{16} \leftarrow \text{KeySched}(K, T)$ 3: $S \leftarrow C \oplus K_{10+b \cdot 6}$ 4: for $i \leftarrow 9 + b \cdot 6$ to $5 + b \cdot 6$ do 5: $S \leftarrow \text{AESrnd}^{-1}(S)$ 6: $S \leftarrow S \oplus K_i$ 7: end for 8: for $i \leftarrow 4$ to 0 do 9: $S \leftarrow \text{AESrnd}^{-1}(S)$ 10: $S \leftarrow S \oplus K_i$ 11: end for 12: return S 13: end Algorithm </pre>	<pre> 1: Algorithm KeySched(K, T) 2: $K_0 \leftarrow K \oplus \mathbb{T} \parallel 0^{64}$ 3: $W_0, \dots, W_3 \leftarrow \frac{32}{K}$ 4: for $i \leftarrow 1$ to 16 do 5: $W'_0 \leftarrow W_0 \oplus \text{SubWord}(W[3]) \oplus \text{Rcon}[i]$ 6: for $j \leftarrow 1$ to 3 do 7: $W'_j \leftarrow W_j \oplus W'_{j-1}$ 8: end for 9: for $j \leftarrow 0$ to 3 do 10: $W_j \leftarrow W'_j$ 11: end for 12: $K_i \leftarrow W_0 \parallel W_1 \parallel W_2 \parallel W_3 \oplus \mathbb{T} \parallel 0^{64}$ 13: end for 14: return K_0, \dots, K_{16} 15: end Algorithm </pre>
<pre> 1: Algorithm ForkAES^ρ(K, T, C, b) 2: $K_0, \dots, K_{16} \leftarrow \text{KeySched}(K, T)$ 3: $b' \leftarrow b \oplus 1$ 4: $S \leftarrow C \oplus K_{10+b \cdot 6}$ 5: for $i \leftarrow 9 + b \cdot 6$ to $5 + b \cdot 6$ do </pre>	<pre> 1: Algorithm AESrnd(S) 2: $S \leftarrow \text{SubBytes}(S)$ 3: $S \leftarrow \text{ShifRows}(S)$ 4: $S \leftarrow \text{MixColumns}(S)$ 5: return S 6: end Algorithm </pre>
<pre> 1: Algorithm AESrnd⁻¹(S) 2: $S \leftarrow \text{iMixColumns}(S)$ 3: $S \leftarrow \text{iShifRows}(S)$ 4: $S \leftarrow \text{iSubBytes}(S)$ 5: return S 6: end Algorithm </pre>	

Figure 3: The algorithms ForkAES, ForkAES⁻¹ and ForkAES^ρ.

Table 1: Upper bounds on probabilities of related-TWEAKEY differential characteristics [28, Table 4.1].

Rounds	Active S-boxes	Probability (upper bound)	Method
1	0	2^0	trivial
2	0	2^0	trivial
3	1	2^{-6}	Matsui's
4	8	2^{-48}	Matsui's
5	≥ 14	2^{-84}	Matsui's
7	≥ 22	2^{-132}	extended split (3R + 4R)

Meet-in-the-Middle Attack. There are numerous meet-in-the-middle attacks performed against AES [20–22] and for all those attacks the key schedule plays an important role. In these attacks partial encryption/decryption is done by guessing keys to prepare pre-computed tables. To reduce the amount of guessed key bytes (and respective attack complexities), the existing linear relations of the AES key schedule are exploited. In our design, we use KIASU as our core encryption operation which in turn replies on the AES cipher with the tweak addition to key schedule. The tweak is a fixed and known constant value T and therefore the existing meet-in-the-middle attacks for AES-128 will apply to both KIASU and our design.

Security Against Other Attacks. Our forkcipher ForkAES is based purposely on the AES block cipher regarding round function and key schedule designs. Moreover, we borrow the KIASU tweak (tweak and key) treatment to support the use of the additional tweak input in our design. Since we do not introduce any novel design complexities, the security of our forkcipher design can be reduced to the security of the AES and KIASU ciphers for further type of attacks.

4 Forkcipher

In this section, we formalize the syntax and security goals of a *forkcipher*; the *kind* of FIL and expanding tweakable primitive that ForkAES is. We recall that in ForkAES, the forward computation starts with a single input X , but forks into two independent branches in the middle, resulting in a double output block $C_0\|C_1$. This makes it possible to compute the original preimage X from either of the two output block-halves C_0 or C_1 , or to reconstruct one half from another. These are the basic properties captured in the concept of a forkcipher.

When keyed and tweaked and given a block X of n bits, a forkcipher computes its image under two permutations $C_0 = \pi_0(X)$ and $C_1 = \pi_1(X)$ *simultaneously*, returning $2n$ bit block $C_0\|C_1$. The “inverse” forkcipher, consists of two algorithms. They each take a binary flag b and an n bit block C as input, and return n bits: one computes the inverse of the input under the permutation selected by b , i.e. $X' = \pi_b^{-1}(C)$ and the other reconstructs the *image* $C' = \pi_{b\oplus 1} \circ \pi_b^{-1}(C)$ of the putative input under the other permutation. This is illustrated in Figure 4.

When used with a random key, each of these permutations would be perfectly random for an ideal forkcipher, i.e., it would implement a *pair* of independent random permutations for every tweak. We thus define a secure forkcipher to be computationally indistinguishable from such an idealized object - a tweak-indexed collection of *pairs* of random permutations.

A trivial forkcipher. It may be clear at this point that the security notion towards which we are headed models two instances of a secure tweakable blockcipher that are used in parallel. In the same spirit, one could instantiate a forkcipher by a tweakable blockcipher used with two independent keys (or a tweak-space separation mechanism).

The main novelty in a forkcipher is that it aims to provide the same security as a pair of tweakable blockciphers at a reduced computational cost. Yet this amortization of computation has nothing to do with the security goals and syntax; these only model what kind of object a forkcipher inevitably is, and which security properties it aspires to achieve.

4.1 Syntax

A forkcipher is a triple of deterministic algorithms, the encryption³ algorithm $F : \{0, 1\}^k \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$, the inversion algorithm $F^{-1} : \{0, 1\}^{2n} \times \mathcal{T} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$

³We again conflate the label for the primitive with the label of the encryption algorithm.

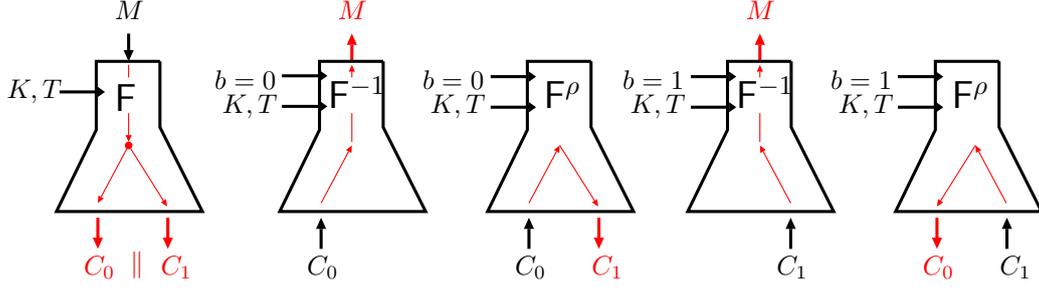


Figure 4: Forward and backward execution of a forkcipher.

and the tag reconstruction algorithm $F^\rho\{0,1\}^k \times \mathcal{T} \times \{0,1\}^n \times \{0,1\} \rightarrow \{0,1\}^n$. We call k, n and \mathcal{T} the keysize, blocksize and tweak space of F , respectively.

A tweakable forkcipher F meets the *correctness condition*, if for every $K, T, M, \beta \in \{0,1\}^k \times \mathcal{T} \times \{0,1\}^n \times \{0,1\}$ we have

$$F^{-1}(K, T, F(K, T, M)[(\beta \cdot n) \dots (\beta \cdot n + n - 1)], \beta) = M$$

and

$$F(K, T, M)[((1-\beta) \cdot n) \dots ((1-\beta) \cdot n + n - 1)] = F^\rho(K, T, F(K, T, M)[(\beta \cdot n) \dots (\beta \cdot n + n - 1)], \beta).$$

We again focus on two specific forms of \mathcal{T} only: when $\mathcal{T} = \{0,1\}^t$ for some positive t , and when \mathcal{T} is an empty set. If $\mathcal{T} = \emptyset$, then we call F simply forkcipher, and we omit \mathcal{T} from the functional notation of both F and F^{-1} , as well as the tweaks from the input arguments of the two algorithms.

4.2 Security Definition

We define the security of forkciphers by indistinguishability from the closest, most natural idealized version of the primitive.

We formalize the security of forkciphers through the notion of a pseudorandom tweakable forked permutation, with help of security games defined in Figure 5. A forked permutation is a pair of oracles, that internally make use of two permutations, s.t. the “left” permutation is used in the usual way (as a permutation), and the “right” permutation is always *forked* from the preimage of the left permutation, no matter if the former is used in the forward or the backward direction.

An adversary \mathcal{A} that aims at breaking a tweakable forkcipher F plays the games **prtfp-real** and **prtfp-ideal** and define the advantage of \mathcal{A} at distinguishing F from a random tweakable injection in a *chosen ciphertext attack* as

$$\text{Adv}_F^{\text{prtfp}}(\mathcal{A}) = \Pr[\mathcal{A}^{\text{prtfp-real}} \Rightarrow 1] - \Pr[\mathcal{A}^{\text{prtfp-ideal}} \Rightarrow 1].$$

4.3 Iterate-Fork-Iterate: A Generic Validation

In order to rule out that the iterate-fork-iterate construction succumbs to a *generic* attack (i.e., attacks that do not use any property of the iterated primitive, but only treat it as a blackbox), we carry out a generic analysis in the spirit of provable security. This result does not, of course, imply the security of ForkAES, or any other instance of forkcipher, because the iterated primitive is not a secure pseudorandom permutation. It does show, however, that the construction itself does not introduce any exploitable weakness.

Game prtfp-real _F	Game prtfp-ideal _F
$K \leftarrow_{\$} \{0, 1\}^k$ $b \leftarrow \mathcal{A}^{\text{ENC,DEC}}$ return b	for $T \in \mathcal{T}$ do $\pi_{T,0}, \pi_{T,1} \leftarrow_{\$} \text{Perm}(n)$ $b \leftarrow \mathcal{A}^{\text{ENC,DEC}}$ return b
Oracle $\text{ENC}(T, M)$ return $F(K, T, M)$	Oracle $\text{ENC}(T, M)$ return $\pi_{T,0}(M) \parallel \pi_{T,1}(M)$
Oracle $\text{DEC}(T, C, \beta)$ return $F^{-1}(K, T, C, \beta)$	Oracle $\text{DEC}(T, C, \beta)$ return $\pi_{T,\beta}^{-1}(C)$

Figure 5: Games **prtfp-real**, and **prtfp-ideal** used to define security of a (strong) forkcipher.

IFI: a generalized ForkAES. We define the IFI construction which combines three tweakable permutations (i.e., three tweak-indexed collections of random permutations) in the same way ForkAES applies the three groups of Tweak-AESround iterations. Fix the block length n and the tweak length t . Formally, for three tweakable random permutations p, p_0, p_1 (i.e. $p = (p_{\mathbf{T}} \leftarrow_{\$} \text{Perm}(n))_{\mathbf{T} \in \{0,1\}^t}$ is a collection of independent uniform elements of $\text{Perm}(n)$ indexed by the elements of $\mathbf{T} \in \{0, 1\}^t$, and similar applies for p_0 and p_1), the forkcipher $F = \text{IFI}[p, p_0, p_1]$ is defined by $F^{\mathbf{T}}(X) = p_{\mathbf{T},0}(p_{\mathbf{T}}(X)) \parallel p_{\mathbf{T},1}(p_{\mathbf{T}}(X))$, $F^{-1\mathbf{T}}(C, b) = p_{\mathbf{T}}^{-1}(p_{\mathbf{T},b}^{-1}(C))$ and $F^{\rho\mathbf{T}}(C, b) = p_{\mathbf{T},b \oplus 1}(p_{\mathbf{T}}^{-1}(C))$. We note that the three tweakable random permutations act as a key for $\text{IFI}[p, p_0, p_1]$ and we omit them for the sake of simplicity.

The analysis. We are going to analyze the indistinguishability of the IFI construction from a forked random permutation. Since both objects an adversary is attempting to distinguish are information-theoretic, we can assume w.l.o.g. that the adversaries are computationally unbounded and deterministic.

Theorem 1. *Fix a blocklength n and a tweaklength t . Then for any adversary \mathcal{A} that makes at most q queries we have that*

$$\text{Adv}_{\text{IFI}[p,p_0,p_1]}^{\text{prtfp}}(\mathcal{A}) = 0$$

where p, p_0, p_1 are random tweakable permutations.

Proof. We use the game Γ from Figure 6 to show that the IFI construction used with a triple of random tweakable permutations yields a perfect forkcipher. For any partially defined permutation π , we let $\mathcal{D}(\pi) \subseteq \{0, 1\}^n$ denote those domain points with a defined image, and we let $\mathcal{R}(\pi) \subseteq \{0, 1\}^n$ denote those range points with a defined preimage. For simplicity, we will denote $\text{IFI}[p, p_0, p_1]$ as F .

We first prove by an induction over adversary's queries that at any point during the execution of game Γ , and for any $\mathbf{T} \in \{0, 1\}^t$ the following properties hold:

1. $\mathcal{D}(\pi_{\mathbf{T},0}) = \mathcal{D}(\pi_{\mathbf{T},1}) = \mathcal{D}(p_{\mathbf{T}})$,
2. $\mathcal{R}(p_{\mathbf{T}}) = \mathcal{D}(p_{\mathbf{T},0}) = \mathcal{D}(p_{\mathbf{T},1})$,

3. $\mathcal{R}(\pi_{\mathbb{T},0}) = \mathcal{R}(p_{\mathbb{T},0})$,
4. $\mathcal{R}(\pi_{\mathbb{T},1}) = \mathcal{R}(p_{\mathbb{T},1})$,
5. $p_{\mathbb{T},0}(p_{\mathbb{T}}(M)) = \pi_{\mathbb{T},0}(M)$ and $p_{\mathbb{T},1}(p_{\mathbb{T}}(M)) = \pi_{\mathbb{T},1}(M)$ for each $M \in \mathcal{D}(p_{\mathbb{T}})$

At the beginning of the game Γ , $p_{\mathbb{T}}, p_{\mathbb{T},0}, p_{\mathbb{T},1}, \pi_{\mathbb{T},0}$ and $\pi_{\mathbb{T},1}$ are undefined for every $\mathbb{T} \in \{0, 1\}^t$, so all four properties are trivially true. Then, assuming that all five properties are true, we examine the effect of \mathcal{A} 's ENC and DEC queries.

When \mathcal{A} makes an ENC(\mathbb{T}, M) query, and $\pi_{\mathbb{T},0}(M) \neq \perp$ (i.e., $M \in \mathcal{D}(\pi_{\mathbb{T},0})$) then none of the partial permutations is extended, and the properties are trivially preserved by the induction assumption.

If $\pi_{\mathbb{T},0}(M) = \perp$, then by property 1 the images of $\pi_{\mathbb{T},1}(M)$ and $p_{\mathbb{T}}(M)$ are undefined as well. We assign a new image to M in each of the three partial permutations, so property 1 is preserved. The value Y is included in $\mathcal{R}(p_{\mathbb{T}}), \mathcal{D}(p_{\mathbb{T},0})$ and $\mathcal{D}(p_{\mathbb{T},1})$, so property 2 is preserved as well. Similarly, $\mathcal{R}(\pi_{\mathbb{T},0})$ and $\mathcal{R}(p_{\mathbb{T},0})$ both get extended by the same value Z_0 , and similarly $\mathcal{R}(\pi_{\mathbb{T},1})$ and $\mathcal{R}(p_{\mathbb{T},1})$ both get extended by Z_1 . Thus properties 3 and 4 are preserved as well. Finally, if property 5 held before the current query then it also holds after it is made, as $p_{\mathbb{T}}(M) = Y$, $p_{\mathbb{T},0}(Y) = Z_0 = \pi_{\mathbb{T},0}(M)$ and $p_{\mathbb{T},1}(Y) = Z_1 = \pi_{\mathbb{T},1}(M)$.

When \mathcal{A} makes a DEC(\mathbb{T}, C, β) query and $\pi_{\mathbb{T},\beta}^{-1}(C) \neq \perp$, no changes are made to the partial permutations and all properties are trivially preserved. Otherwise, the value X extends the domains of $p_{\mathbb{T}}, \pi_{\mathbb{T},0}$ and $\pi_{\mathbb{T},1}$, preserving property 1. The range of $p_{\mathbb{T}}$ is extended by the value Y , as are the domains of $p_{\mathbb{T},0}$ and $p_{\mathbb{T},1}$, preserving property 2. The adversarial input C is added to both $\mathcal{R}(\pi_{\mathbb{T},\beta})$ and $\mathcal{R}(p_{\mathbb{T},\beta})$, and the value $Z_{\beta \oplus 1}$ extends both $\mathcal{R}(\pi_{\mathbb{T},\beta \oplus 1})$ and $\mathcal{R}(p_{\mathbb{T},\beta \oplus 1})$, so the properties 3 and 4 are preserved. Finally, we have $p_{\mathbb{T}}(X) = Y$, $p_{\mathbb{T},\beta}(Y) = C = \pi_{\mathbb{T},\beta}(X)$ and $p_{\mathbb{T},\beta \oplus 1}(Y) = Z_{\beta \oplus 1} = \pi_{\mathbb{T},\beta \oplus 1}(X)$, so property 5 is preserved as well.

It is easy to see, that the games Γ and **prtfp-ideal_F^A** are equivalent. The framed lines in Figure 6 do not affect the outputs of oracle queries; Γ just lazily samples two tweakable random permutations π_0 and π_1 , and uses them to reply the ENC and DEC queries the same way as in **prtfp-ideal_F**. Therefore $\Pr[\mathcal{A}^\Gamma \Rightarrow 1] = \Pr[\mathcal{A}^{\text{prtfp-ideal}_F} \Rightarrow 1]$.

At the same time, in a non-trivial ENC(\mathbb{T}, M) query, we lazily sample an image Y of $p_{\mathbb{T}}(M)$, which was previously undefined due to property 1. The lines 3 and 4 do a correct lazy sampling of $p_{\mathbb{T},0}$ and $p_{\mathbb{T},1}$: the images $p_{\mathbb{T},0}(Y)$ and $p_{\mathbb{T},1}(Y)$ were previously undefined due to property 2, and the sampling of the images Z_0 and Z_1 is correct due to properties 3 and 4. Finally, due to property 5, we see that the ENC oracle actually implements the F construction.

Similarly, in a DEC(\mathbb{T}, M, β) query, we sample a preimage Y of previously undefined $p_{\mathbb{T},\beta}^{-1}(C)$ (due to property 3 or 4). Then, the previously unassigned $p_{\mathbb{T}}^{-1}(Y)$ and $p_{\mathbb{T},\beta \oplus 1}(Y)$ (due to property 2) get a correctly sampled preimage X , resp. image $Z_{\beta \oplus 1}$ (and sampling is correct due to property 1 and property 3 or 4). Finally, the assignment is compatible with the F construction (due to property 5). Thus the games Γ and **prtfp-real_F** are equivalent, and $\Pr[\mathcal{A}^\Gamma \Rightarrow 1] = \Pr[\mathcal{A}^{\text{prtfp-real}_F} \Rightarrow 1]$. This concludes the proof. \square

5 Tweakable Forkcipher Modes

We demonstrate the applicability of forkciphers by designing provably secure AE modes of operation for a tweakable forkcipher. The designs are motivated by the following objectives: (1) the resulting AE scheme must be able to process strings of *arbitrary length* but (2) it must be most efficient for encryption queries whose total number of blocks (in AD and message) is very small, e.g. below four.

```

1: proc initialize
2:   bad  $\leftarrow$  false
3:   for  $T \in \{0, 1\}^t$  do
4:      $p_T \leftarrow \perp$ ;  $p_{T,0} \leftarrow \emptyset$ ;  $p_{T,1} \leftarrow \emptyset$ 
5:      $\pi_{T,0} \leftarrow \perp$ ;  $\pi_{T,1} \leftarrow \perp$ 
6:   end for

1: proc ENC( $T, M$ )
2:   if  $\pi_{T,0}(M) = \perp$  then
3:      $Z_0 \leftarrow \{0, 1\}^n \setminus \mathcal{R}(\pi_{T,0})$ 
4:      $Z_1 \leftarrow \{0, 1\}^n \setminus \mathcal{R}(\pi_{T,1})$ 
5:      $\pi_{T,0}(M) \leftarrow Z_0$ 
6:      $\pi_{T,1}(M) \leftarrow Z_1$ 
7:      $Y \leftarrow \{0, 1\}^n \setminus \mathcal{R}(p_T)$ 
8:      $p_T(M) \leftarrow Y$ 
9:      $p_{T,0}(Y) \leftarrow Z_0$ 
10:     $p_{T,1}(Y) \leftarrow Z_1$ 

11:  end if
12:  return  $\pi_{T,0}(M) \parallel \pi_{T,1}(M)$ 

1: proc DEC( $T, C, \beta$ )
2:   if  $\pi_{T,\beta}^{-1}(C) = \perp$  then
3:      $X \leftarrow \{0, 1\}^n \setminus \mathcal{D}(\pi_{T,\beta})$ 
4:      $Z_{\beta \oplus 1} \leftarrow \{0, 1\}^n \setminus \mathcal{R}(\pi_{T,\beta \oplus 1})$ 
5:      $\pi_{T,\beta}^{-1}(C) \leftarrow X$ 
6:      $\pi_{T,\beta \oplus 1}(X) \leftarrow Z_{\beta \oplus 1}$ 
7:      $Y \leftarrow \{0, 1\}^n \setminus \mathcal{D}(p_{T,\beta})$ 
8:      $p_{T,\beta}^{-1}(C) \leftarrow Y$ 
9:      $p_T^{-1}(Y) \leftarrow X$ 
10:     $p_{T,\beta \oplus 1}(Y) \leftarrow Z_{\beta \oplus 1}$ 
11:  end if
12:  return  $\pi_{T,\beta}^{-1}(C)$ 

```

Figure 6: Game Γ used in the proof of security of the $\text{IFI}[p, p_0, p_1]$ construction. The tweakable permutations p, p_0, p_1, π_0 and π_1 are initially undefined.

We define three nonce-based AE modes of operation for a tweakable forkcipher. One of the modes is fully parallelizable, but its security falls apart with the first nonce repetition. This mode will be useful in applications where longer queries occur more often, and where one of the parties has more powerful computational capabilities.

The second mode has an online computable encryption, but is not parallelizable. While its integrity and confidentiality bounds in the nonce-respecting setting are inferior to those of the parallel mode, we conjecture that a nonce-misusing universal forgery will require data complexity at the birthday bound. In addition to the resilience to accidental nonce reuse, the sequential mode also lends itself well to low-overhead implementations, as it does not require to store any message-block counter.

The third mode is inspired by the blockcipher-based scheme GCM [23], but differs from it in some critical details. It is almost fully parallelizable, and it beats the other two modes in efficiency for plaintexts of more than n bits, but this increased efficiency comes at the price of an n -bit larger memory footprint necessary to store a derived authentication key.

While the first two of our AE schemes appear to be computationally less efficient, we stress that all three can be efficiently implemented when instantiated with ForkAES. The ForkAES-based instances will even beat all of the existing AE modes for the shortest queries (see Section 5.8).

A small AE primitive. While a secure forkcipher does not directly capture any integrity, we show in Section 5.9 that a secure forkcipher *can* be used as an AEAD scheme with fixed length messages and AD *in the natural way*, provably delivering robust AE security guarantees.

5.1 Syntax and Security of AE

Our modes target the security of nonce-based AE with associated data, following the syntax proposed by Rogaway [40].

A scheme for authenticated encryption with associated data is a triplet $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$. The key space \mathcal{K} is a finite set endowed with the uniform distribution. The deterministic encryption algorithm $\mathcal{E} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \rightarrow \mathcal{C}$ maps a secret key K , a nonce N , an

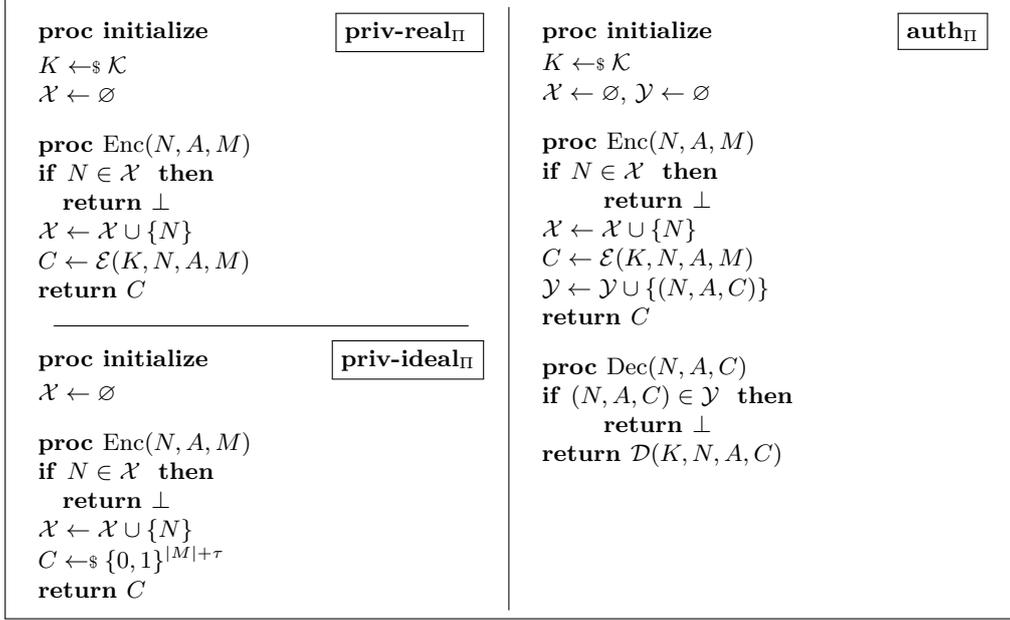


Figure 7: Games for defining the security of a nonce based AE scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ with ciphertext expansion τ .

associated data A and a message M to a ciphertext $C = \mathcal{E}(K, N, A, M)$. The nonce, AD and message domains are all subsets of $\{0, 1\}^*$. The deterministic decryption algorithm $\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{C} \rightarrow \mathcal{M} \cup \{\perp\}$ takes a tuple (K, N, A, C) and either returns a message $M \in \mathcal{M}$, or a distinguished symbol \perp to signalize an authentication error.

We require that for every $M \in \mathcal{M}$, we have $\{0, 1\}^{|M|} \subseteq \mathcal{M}$ (i.e. for any integer m , either all or no strings of length m belong to \mathcal{M}) and that for all $K, N, A, M \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ we have $|\mathcal{E}(K, N, A, M)| = |M| + \tau$ for some non-negative integer τ called the stretch of Π . For correctness of Π , we require that for all $K, N, A, M \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ we have $M = \mathcal{D}(K, N, A, \mathcal{E}(K, N, A, M))$. We let $\mathcal{E}_K(N, A, M) = \mathcal{E}(K, N, A, M)$ and $\mathcal{D}_K(N, A, M) = \mathcal{D}(K, N, A, M)$.

We use the two-requirement definition of AE security. We model a chosen plaintext attack of an adversary \mathcal{A} against the confidentiality of a nonce-based AE scheme Π with the help of the security games **priv-real** and **priv-ideal** in Figure 7. We define the advantage of \mathcal{A} in breaking the confidentiality of Π as

$$\text{Adv}_{\Pi}^{\text{priv}}(\mathcal{A}) = \Pr[\mathcal{A}^{\text{priv-real}\Pi} \Rightarrow 1] - \Pr[\mathcal{A}^{\text{priv-ideal}\Pi} \Rightarrow 1].$$

We model a chosen ciphertext attack against the integrity of Π with help of the game **auth** in Figure 7. We define the advantage of \mathcal{A} in breaking the integrity of Π as

$$\text{Adv}_{\Pi}^{\text{auth}}(\mathcal{A}) = \Pr[\mathcal{A}^{\text{auth}\Pi} \text{ forges}]$$

where “ \mathcal{A} forges” denotes the event that there is a decryption query that returns a value other than \perp .

5.2 Parallel AE from a Forkcipher

We define the nonce based AEAD scheme PAEF (as in “Parallel AE from a Forkcipher”). PAEF is parameterized by a forkcipher F (as defined in Section 4) with $\mathcal{T} = \{0, 1\}^t$ for a positive t . It is further parameterized by a nonce length $0 < \nu \leq t - 4$. An instance

PAEF[F, ν] = ($\mathcal{K}, \mathcal{E}, \mathcal{D}$) has $\mathcal{K} = \{0, 1\}^k$ and the encryption and decryption algorithms as defined in Figure 8. Its nonce space is $\mathcal{N} = \{0, 1\}^\nu$, and its message and AD space are respectively $\mathcal{M} = \{0, 1\}^{\leq n \cdot 2^{(t-\nu-3)}}$ and $\mathcal{A} = \{0, 1\}^{\leq n \cdot 2^{(t-\nu-3)}}$. The ciphertext expansion of PAEF[F, ν] is n . The encryption algorithm is illustrated in Figure 9.

In an encryption query, AD and message are partitioned into blocks of n bits. Each block is processed with exactly one call to F using a tweak that is composed of

1. a three-bit flag $f_0 \| f_1 \| f_2$,
2. the nonce,
3. a $(t - \nu - 3)$ -bit encoding of the block index (indexing is individual for both AD and message),

so the nonce-length is a parameter that allows to make a trade-off between maximal message length and maximal number of queries with the same key. The bit $f_0 = 1$ iff the final block of message is being processed, the bit $f_1 = 1$ iff a block of message is being processed, and $f_2 = 1$ iff the final block of the current input (depending on f_1) is being processed and is incomplete. The ciphertext blocks are the “left” output blocks of F from message blocks, and an xor of all “right” output blocks of F is xored to the rightmost n bits of the call to F that processes the final message block.

The decryption proceeds similarly as the encryption, except that “right” output blocks of the message blocks are reconstructed from ciphertext blocks (using the reconstruction algorithm) to recompute the tag, which is then checked.

5.3 Security of PAEF

We state the formal claim about the nonce-based AE security of PAEF in Theorem 2.

Theorem 2. *Let F be a tweakable forkcipher with $\mathcal{T} = \{0, 1\}^t$, and let $0 < \nu \leq t - 4$. Then for any nonce-respecting adversary \mathcal{A} whose queries lie in the proper domains of the encryption and decryption algorithms and who makes at most q_v decryption queries, we have*

$$\mathbf{Adv}_{\text{PAEF[F,}\nu]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{B})$$

and

$$\mathbf{Adv}_{\text{PAEF[F,}\nu]}^{\text{auth}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{C}) + \frac{q_v \cdot 2^n}{(2^n - 1)^2}$$

for some adversaries \mathcal{B} and \mathcal{C} who make at most twice as many queries in total as is the total number of blocks in all encryption, respectively all encryption and decryption queries made by \mathcal{A} , and who run in time given by the running time of \mathcal{A} plus an overhead that is linear in the total number of blocks in all \mathcal{A} 's queries.

Proof. For both confidentiality and authenticity, we first replace F with a pair of independent random tweakable permutations π_0, π_1 , i.e. $\pi_0 = (\pi_{\mathbb{T},0} \leftarrow \$ \text{Perm}(n))_{\mathbb{T} \in \{0,1\}^t}$ is a collection of independent uniform elements of $\text{Perm}(n)$ indexed by the elements of $\mathbb{T} \in \{0, 1\}^t$ (and similarly $\pi_1 = (\pi_{\mathbb{T},1} \leftarrow \$ \text{Perm}(n))_{\mathbb{T} \in \{0,1\}^t}$). We let $\text{PAEF}[(\pi_0, \pi_1), \nu]$ denote the PAEF mode that uses π_0, π_1 instead of F. We have that

$$\mathbf{Adv}_{\text{PAEF[F,}\nu]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{B}) + \mathbf{Adv}_{\text{PAEF}[(\pi_0, \pi_1), \nu]}^{\text{priv}}(\mathcal{A})$$

because a distinguisher \mathcal{B} for F can perfectly simulate the games $\mathbf{priv-real}_{\text{PAEF[F,}\nu]}$ and $\mathbf{priv-real}_{\text{PAEF}[(\pi_0, \pi_1), \nu]}$ for \mathcal{A} using its own oracles. In place of any F^ρ call, \mathcal{B} has to make a decryption query followed by an encryption query. By copying \mathcal{A} 's output, \mathcal{B} can achieve the same advantage as \mathcal{A} does, with the same data complexity as \mathcal{A} and a very similar

<pre> 1: Algorithm $\mathcal{E}(K, N, A, M)$ 2: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 3: $M_1, \dots, M_m, M_* \xleftarrow{n} M$ 4: $T \leftarrow 0^n$ 5: for $i \leftarrow 1$ to a do 6: $T \leftarrow 000\ N\ \langle i \rangle_{t-\nu-3}$ 7: $T \leftarrow T \oplus \text{right}_n(\mathbf{F}_K^T(A_i))$ 8: end for 9: if $A_* = n$ then 10: $T \leftarrow 001\ N\ \langle a+1 \rangle_{t-\nu-3}$ 11: $T \leftarrow T \oplus \text{right}_n(\mathbf{F}_K^T(A_*))$ 12: else if $A_* > 0$ or $M = 0$ then 13: $T \leftarrow 011\ N\ \langle a+1 \rangle_{t-\nu-3}$ 14: $T \leftarrow T \oplus \text{right}_n(\mathbf{F}_K^T(A_*\ 10^*))$ 15: end if \triangleright Do nothing if $A = \varepsilon, M \neq \varepsilon$ 16: for $i \leftarrow 1$ to m do 17: $T \leftarrow 100\ N\ \langle i \rangle_{t-\nu-3}$ 18: $C_i, T' \leftarrow \text{lsplit}_n(\mathbf{F}_K^T(M_i))$ 19: $T \leftarrow T \oplus T'$ 20: end for 21: if $M_* = n$ then 22: $T \leftarrow 101\ N\ \langle m+1 \rangle_{t-\nu-3}$ 23: $R \leftarrow \mathbf{F}_K^T(M_*)$ 24: else if $M_* > 0$ then 25: $T \leftarrow 111\ N\ \langle m+1 \rangle_{t-\nu-3}$ 26: $R \leftarrow \text{left}_{n+ M_* }(\mathbf{F}_K^T(M_*\ 10^*))$ 27: else 28: $R \leftarrow 0^n$ 29: end if 30: $R \leftarrow (0^{ M_* }\ T) \oplus R$ 31: return $C_1\ \dots\ C_m\ R$ 32: end Algorithm </pre>	<pre> 1: Algorithm $\mathcal{D}(K, N, A, C)$ 2: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 3: $C_1, \dots, C_m, C_*, T \leftarrow \text{csplit-b}_n(C)$ 4: $\bar{T} \leftarrow 0^n$ 5: for $i \leftarrow 1$ to a do 6: $T \leftarrow 000\ N\ \langle i \rangle_{t-\nu-3}$ 7: $\bar{T} \leftarrow \bar{T} \oplus \text{right}_n(\mathbf{F}_K^T(A_i))$ 8: end for 9: if $A_* = n$ then 10: $T \leftarrow 001\ N\ \langle a+1 \rangle_{t-\nu-3}$ 11: $\bar{T} \leftarrow \bar{T} \oplus \text{right}_n(\mathbf{F}_K^T(A_*))$ 12: else if $A_* > 0$ or $T = 0$ then 13: $T \leftarrow 011\ N\ \langle a+1 \rangle_{t-\nu-3}$ 14: $\bar{T} \leftarrow \bar{T} \oplus \text{right}_n(\mathbf{F}_K^T(A_*\ 10^*))$ 15: end if \triangleright Do nothing if $A = \varepsilon, M \neq \varepsilon$ 16: for $i \leftarrow 1$ to m do 17: $T \leftarrow 100\ N\ \langle i \rangle_{t-\nu-3}$ 18: $M_i \leftarrow \mathbf{F}_K^{-1}(C_i, 0)$ 19: $\bar{T} \leftarrow \bar{T} \oplus \mathbf{F}_K^{\rho}(C_i, 0)$ 20: end for 21: $C_*, T \leftarrow \text{lsplit}_n((C_*\ T) \oplus (0^{ T }\ \bar{T}))$ 22: if $T = n$ then 23: $T' \leftarrow T$ 24: else if $T > 0$ then 25: $T \leftarrow 111\ N\ \langle m+1 \rangle_{t-\nu-3}$ 26: $T' \leftarrow 10^{n- T -1}\ T$ 27: else 28: if $C_* \neq 0^n$ then return \perp 29: return ε 30: end if 31: $\bar{T} \leftarrow \text{left}_{ T }(\mathbf{F}_K^{\rho}(C_*, 0))$ 32: $M_* \leftarrow \mathbf{F}_K^{-1}(C_*, 0)$ 33: if $T' \neq \text{right}_{n- T }(M_*)\ \bar{T}$ then 34: return \perp 35: end if 36: $M_* \leftarrow \text{left}_{ T }(M_*)$ 37: return $M_1\ \dots\ M_m\ M_*$ 38: end Algorithm </pre>
--	--

Figure 8: The PAEF[F, ν] AEAD scheme. Here $\langle i \rangle_\ell$ is the canonical encoding of an integer i as an ℓ -bit string.

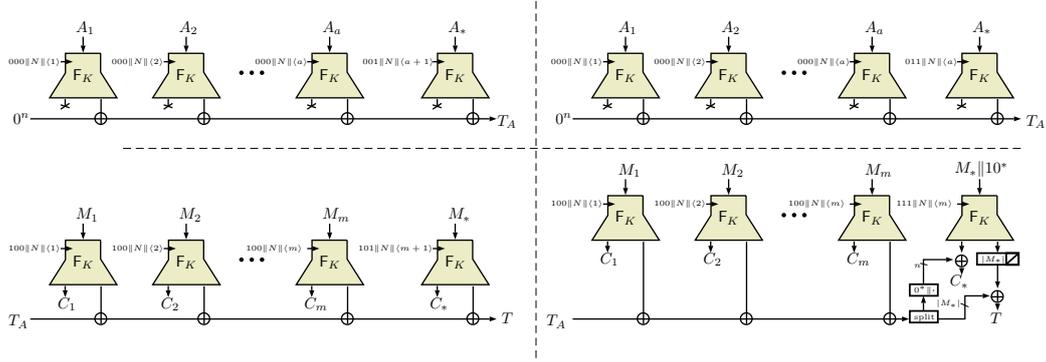


Figure 9: The encryption algorithm of PAEF[F] mode. The picture illustrates the processing of AD when length of AD is a multiple of n (**top left**) and when the length of AD is not a multiple of n (**top right**), and the processing of the message when length of the message is a multiple of n (**bottom left**) and when the length of message is not a multiple of n (**bottom right**)

running time. This implies that the gap between these games is bounded by $\mathbf{Adv}_F^{\text{prtfp}}(\mathcal{B})$. By a similar argument, we have that

$$\mathbf{Adv}_{\text{PAEF}[F, \nu]}^{\text{auth}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{prtfp}}(\mathcal{C}) + \mathbf{Adv}_{\text{PAEF}[(\pi_0, \pi_1), \nu]}^{\text{priv}}(\mathcal{A}).$$

For confidentiality, it is easy to see that in a nonce-respecting attack, every message block is processed with a unique tweak. Every ciphertext block is produced as the only image under and independent random permutation, and thus uniformly distributed. The final $|M_*| + n$ bits of every ciphertext are produced as xor-sum that always contains (a part of) the concatenated output of π_0 and π_1 applied to the last message block with a unique tweak, such that this part is not used for any ciphertext block. Since all ciphertexts are uniformly distributed we get perfect privacy and hence our privacy result.

For authenticity, we analyse the probability of forgery for an adversary that makes a single decryption query against $\text{PAEF}[(\pi_0, \pi_1), \nu]$ and then use a result of Bellare [10] to extend our result to multiple queries (still against $\text{PAEF}[(\pi_0, \pi_1), \nu]$).

We will denote the encryption queries of \mathcal{A} and the corresponding replies as (N^i, A^i, M^i) and C^i for $i = 1, \dots, q$, where q is the number of encryption queries made by \mathcal{A} . For each i we let $C_1^i, \dots, C_m^i, C_*^i, T = \text{csplit-b}_n(C^i)$. We let (N, A, C) denote the only decryption query of \mathcal{A} and we let $C_1, \dots, C_m, C_*, T = \text{csplit-b}_n(C)$. When the forgery (N, A, C) is made, we have two base cases. If the nonce N is fresh, then the forgery attempt is equivalent to guessing the value of a uniform string of n bits, thus succeeds with probability 2^{-n} . This holds even if $|T| < 0$, because the rightmost $(n - |T|)$ bits of the final image under the inverse of π_0 must have a specific value.

If N is reused, i.e. if $N = N^i$ for some $N^i \in \{N^1, \dots, N^q\}$, then we perform a case analysis. Note that we can disregard all encryption queries except the i^{th} , because their ciphertexts are computed using independent random permutations. Every case assumes the negation of all previous case-conditions.

Case 1, $|C|_n \neq |C^i|_n$: We have several subcases.

- If $|C| = n$, then C is a xor-sum of $\pi_{T,1}$ images from the associated data (denoted as T_A in Figure 9), such that we can possibly have $A^i = A$. However, due to the assumption in this case, we must have $|M^i| > 0$, so the xor-sum T_{A^i} computed

in the i^{th} encryption query is xor-masked with uniform bits produced by the processing of M_*^i . Therefore T_{A^i} is statistically independent of C^i , and the adversary has no information when trying to guess the value of the T_A sum. The probability of a successful forgery is 2^{-n} .

- When $|C| > n$, regardless if C has more or less blocks than C^i , the successful forgery is equivalent to guessing the value of an image under π_1 (respectively the value of n out of $2n$ bits produced by $\pi_{\top,0}^{-1}(\text{left}_n(T_*))$ and $\pi_{\top,1}(\pi_{\top,0}^{-1}(\text{left}_n(T_*)))$) such that the tweak $\top = 110\|N\|(m+1)_{t-\nu-3}$ (respectively $\top = 111\|N\|(m+1)_{t-\nu-3}$) was not used before. The probability of this event is 2^{-n} .

The probability of a successful forgery in **Case 1** is at most 2^{-n} . In the following cases, $|C|_n = |C^i|_n$.

Case 2, $|A|_n \neq |A^i|_n$: Again, we have a few subcases to consider.

- If $|A|_n > |A^i|_n$, a successful forgery is equivalent to guessing an output value of $\pi_{\top,1}$ with a previously unused tweak ($\top = 0b1\|N\|(a+1)_{t-\nu-3}$ for $b \in \{0,1\}$) thanks to $a > a_i$, succeeding with probability of 2^{-n} .
- If $0 < |A|_n < |A^i|_n$ but more than 0, then a successful forgery is still equivalent to guessing an output value of $\pi_{\top,1}$ with a previously unused tweak ($\top = 10b1\|N\|(a+1)_{t-\nu-3}$ for $b \in \{0,1\}$), thanks to the three-bit domain-separation flag (which was set to 000 in the i^{th} encryption query). This succeeds with probability 2^{-n} .
- Finally if $|A| = 0$, then $|A|_n \neq |A^i|_n$ implies that $|A^i|_n > 0$. Forging in this case is equivalent to guessing the image $\pi_{(011\|N\|1),1}(10^{n-1})$, such that the random permutation $\pi_{(011\|N\|1),1}$ was evaluated on no more than a single other input $A_*^i\|10^* \neq 10^{n-1}$ in the whole game. This succeeds with probability at most $1/(2^n - 1)$.

Thus the probability of a successful forgery in this case is at most $1/(2^n - 1)$. In the remaining cases, we have $|C|_n = |C^i|_n > 1$ and $|A|_n = |A^i|_n > 0$.

Case 3, $|C| \neq |C^i|$ and $|T| = n$ or $|T^i| = n$: In this case, the forgery verification will use $\pi_{\top,1}$ with a fresh tweak \top because the “incomplete-block” bit of the three-bit flag will have different values in the processing of the decryption query, and in the processing of the i^{th} encryption query. The forgery succeeds with probability 2^{-n} .

Case 4, $|A| \neq |A^i|$ and $|A_*| = n$ or $|A_*^i| = n$: This is analogous with the previous case; the probability of forgery is 2^{-n} . In the remaining cases, we have $|C|_n = |C^i|_n > 1$, $|A|_n = |A^i|_n > 0$ and $|T| > 0$, $|T^i| > 0$, $|A_*| > 0$, $|A_*^i| > 0$.

Case 5, $|C| \neq |C^i|$ and $|T| < n$ and $|T^i| < n$: In this case, both the encryption query and the decryption query use the same tweak \top to process M_*^i and C_*, T , respectively. There are two conditions for the forgery to succeed. First, the preimage $X = \pi_{\top,0}^{-1}(C_* \oplus (0^{n-|T|}\text{left}_{|T|}(\bar{T})))$ (as per line 21 in Figure 8) must be equal to $W\|10^{n-|T|-1} \neq M_*^i\|10^{n-|T^i|-1}$ (noting that the case condition implies $|T| \neq |T^i|$) for some $W \in \{0,1\}^{|T|}$. This is no easier than finding a fresh value whose preimage falls into a set of size $2^{|T|}$. With a single image of $\pi_{\top,0}^{-1}$ already used, this succeeds with probability bounded by $(2^{|T|})/(2^n - 1)$. *Secondly*, the image $Y = \pi_{\top,1}(X)$ must be equal to $T\|Z$ for some $Z \in \{0,1\}^{n-|T|}$, *conditioned on X having the correct format*. This is equivalent to guessing a fresh image under $\pi_{\top,1}$ with $(n - |T|)$ free bits. As a single image of $\pi_{\top,1}$ has been used already, this happens with probability at most $(2^{n-|T|})/(2^n - 1)$. The probability of a successful forgery in this case is therefore bounded by $(2^{|T|})/(2^n - 1) \cdot (2^{n-|T|})/(2^n - 1) = 2^n/(2^n - 1)^2$.

Case 6, $|A| \neq |A^i|$ and $|A_a| < n$ and $|A_a^i| < n$: In this case, the final blocks A_* and A_*^i are processed by the same random permutation $\pi_{\tau,1}$, but as $A_* \parallel 10^{n-|A_*|} \neq A_*^i \parallel 10^{n-|A_*^i|}$, successfully forging in this case is equivalent to guessing the yet unsampled image $\pi_{\tau,1}(A_* \parallel 10^{n-|A_*|})$. With a single image of $\pi_{\tau,1}$ used before, this happens with probability at most $1/(2^n - 1)$.

Case 7, $|C| = |C^i|$ and $|A| = |A^i|$: In this case, there must be at least a single block of either AD or ciphertext where the two queries differ. We investigate the following subcases.

- If the forgery N, A, C differs from N, A^i, C^i only in $C_* \parallel T$, then, if we ran the decryption algorithm on N, A^i, C^i and N, A, C in parallel, the values \bar{T}^i and \bar{T} used on the line 21 of the decryption algorithm in Figure 8 would be the same, and thus necessarily $(C_* \parallel T) \oplus (0^{n-|T|} \parallel \bar{T}) \neq (C_*^i \parallel T^i) \oplus (0^{n-|T|} \parallel \bar{T}^i)$. The probability of a successful forgery is at most $2^n/(2^n - 1)^2$ by a similar argument as in **Case 5**.
- If $A, C \parallel T$ and $A^i, C^i \parallel T^i$ differ in a single block, such that $C_* \parallel T = C_*^i \parallel T^i$, a forgery is impossible (because $\pi_{\tau,0}$ and $\pi_{\tau,1}$ are all permutations).
- If there are at least two blocks in $A_1, \dots, A_a, A_*, C_1, \dots, C_m, C_*, T$ that differ from the corresponding blocks in $A_1^i, \dots, A_a^i, A_*^i, C_1^i, \dots, C_m^i, C_*^i, T^i$, then the forgery can succeed in two ways. The first is if $(C_* \parallel T) \oplus (0^{n-|T|} \parallel \bar{T}) = (C_*^i \parallel T^i) \oplus (0^{n-|T|} \parallel \bar{T}^i)$. This happens with probability at most $1/(2^n - 1)$, as there will be at least one index j for which $A_j \neq A_j^i$ (or $C_j \neq C_j^i$), and for which $\pi_{\tau,1}(A_j) \oplus \pi_{\tau,1}(A_j^i)$ (respectively $\pi_{\tau,1}(\pi_{\tau,0}^{-1}(C_j)) \oplus \pi_{\tau,1}(\pi_{\tau,0}^{-1}(C_j^i))$) would have to take a particular value. The probability follows from the fact that whatever τ , the random permutations $\pi_{\tau,1}$ and $\pi_{\tau,1}^{-1}$ we sampled only once. The second way is if $(C_* \parallel T) \oplus (0^{n-|T|} \parallel \bar{T}) \neq (C_*^i \parallel T^i) \oplus (0^{n-|T|} \parallel \bar{T}^i)$ but the verification still succeeds. This is analogous to **Case 5**.

The probability of a successful forgery in this case is therefore bounded by $2^n/(2^n - 1)^2$.

Thus a single forgery succeeds with probability no greater than $2^n/(2^n - 1)^2$. By applying the result of Bellare [10], we can bound the probability of a successful forgery among q_v decryption queries as $(q_v \cdot 2^n)/(2^n - 1)^2$. \square

5.4 Sequential AE from a Forkcipher

We now define SAEF (as in “Sequential AE from a Forkcipher,” pronounce as “safe”), a nonce-based AEAD scheme parameterized by a tweakable forkcipher F (as defined in Section 4) with $\mathcal{T} = \{0, 1\}^t$ for a positive $t \leq n$. An instance $\text{SAEF}[F] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ has a key space $\mathcal{K} = \{0, 1\}^k$, nonce space $\mathcal{N} = \{0, 1\}^{t-3}$, and the AD and message spaces are both $\{0, 1\}^*$ (although the maximal message length influences the security). The ciphertext expansion of $\text{SAEF}[F]$ is n . The encryption and decryption algorithms are defined in Figure 11 and the encryption algorithm is illustrated in Figure 12.

In an encryption query, first AD and then message processed in blocks of n bits. Each block is processed with exactly one call to F , using a tweak that is composed of

1. the nonce, and
2. a three-bit flag f .

The flag f takes different values for processing of different types of data blocks in the encryption algorithm. These are given in Figure 10.

The “right” output of every F call is used as a whitening mask for the following F call, masking either the input (in AD processing) or both the input and the output (in message

value of f	type of block processed by the encryption algorithm
000	processing non-final AD block
010	processing final complete AD block
011	processing final incomplete AD block
110	processing final complete AD block to produce tag
111	processing final incomplete AD block to produce tag
001	processing non-final message block
100	processing final complete message block
101	processing final incomplete message block

Figure 10: The binary flag values used in the SAEF mode for forkcipher.

processing) of this subsequent call. The initial F call in the query is unmasked. The tag is the last “right” output of F produced in the query.

The decryption proceeds similarly as the encryption, except that the “right” outputs of F in the message processing part are reconstructed from ciphertext blocks (using the F reconstruction algorithm). Only if the verification succeeds will the actual decryption take place.

5.5 Security of SAEF

For the security analysis of SAEF, we will work with adversarial resources that are more fine-grained than it is usual. The reason for this is that the security in a typical setting, most of the messages will be very short. We will show that this can greatly improve security of SAEF, if we account for adversarial resources per query length.

For an adversary \mathcal{A} , we will watch the total number of encryption and decryption queries, q and q_v respectively. For data complexity, we will work with a vector $\mathbf{q} \in \mathbb{N}^L$ with $L = \{\ell \mid \exists A \in \mathcal{A}, M \in \mathcal{M} \text{ s.t. } |A|_n = \ell \text{ or } |M|_n = \ell\}$. In other words, \mathbf{q} is a vector of non-negative integers whose elements are indexed by a subset of those integers that can be the number of either message or AD blocks in a query to SAEF[F]. The vector \mathbf{q} associated to \mathcal{A} will then contain for each ℓ the maximal number \mathbf{q}_ℓ of queries (N, A, M) made by \mathcal{A} with no more than $\ell = \max(|A|_n, |M|_n)$ blocks of either the message or AD. We further let $\ell_{\max} = \max\{|A|_n + |X|_n \mid \mathcal{A} \text{ queries } (N, A, X) \text{ or } (N, A, X||T)\}$ denote the maximal number of blocks of AD and message *together* in a single query. We finally let σ denote the maximal total data complexity of \mathcal{A} , measured in the total number of blocks of AD, messages, and possibly ciphertexts in all \mathcal{A} 's queries. It is then easy to see that $\sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \ell \cdot \mathbf{q}_\ell \leq 2\sigma$ and that $q + q_v \leq \sum_{\ell \in L} \mathbf{q}_\ell$.

We state the formal claim about the nonce-based AE security of SAEF in Theorem 3.

Theorem 3. *Let F be a tweakable forkcipher with $\mathcal{T} = \{0, 1\}^t$. Then for any nonce-respecting adversary \mathcal{A} whose resources are bounded by $q, q_v, \sigma, \ell_{\max}$ and \mathbf{q} (in the sense we just defined) such that $\ell_{\max} \leq 2^n/4$, we have*

$$\begin{aligned} \text{Adv}_{\text{SAEF[F]}}^{\text{priv}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{B}) + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} 3 \cdot \frac{\mathbf{q}_\ell \cdot \ell \cdot (\ell - 1)}{2^n}, \\ \text{Adv}_{\text{SAEF[F]}}^{\text{auth}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{C}) \\ &\quad + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \frac{\mathbf{q}_\ell \cdot 3\ell \cdot (\ell - 1)}{2^n} + \frac{q_v \cdot (2\ell_{\max} + 1)}{2^n} + \frac{q_v}{(2^n - 1)} \end{aligned}$$

for some adversaries \mathcal{B} and \mathcal{C} who make at most 2σ queries, and who run in time given by the running time of \mathcal{A} plus $\gamma \cdot \sigma$ for some constant γ .

<pre> 1: Algorithm $\mathcal{E}(K, N, A, M)$ 2: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 3: $M_1, \dots, M_m, M_* \xleftarrow{n} M$ 4: $\text{noM} \leftarrow 0$ 5: if $M = 0$ then $\text{noM} \leftarrow 1$ 6: $\Delta \leftarrow 0^n$; $\text{T} \leftarrow 000\ N$ 7: for $i \leftarrow 1$ to a do 8: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}(A_i \oplus \Delta))$ 9: end for 10: if $A_* = n$ then 11: $\text{T} \leftarrow \text{noM}\ 10\ N$ 12: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}(A_* \oplus \Delta))$ 13: else if $A_* > 0$ or $M = 0$ then 14: $\text{T} \leftarrow \text{noM}\ 11\ N$ 15: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}((A_*\ 10^*) \oplus \Delta))$ 16: end if \triangleright Do nothing if $A = \varepsilon, M \neq \varepsilon$ 17: $\text{T} \leftarrow 001\ N$ 18: for $i \leftarrow 1$ to m do 19: $S \leftarrow \text{F}_K^{\text{T}}(M_i \oplus \Delta) \oplus (\Delta\ 0^n)$ 20: $C_i, \Delta \leftarrow \text{lsplit}_n(S)$ 21: end for 22: if $M_* = n$ then 23: $\text{T} \leftarrow 100\ N$ 24: $R \leftarrow \text{F}_K^{\text{T}}(M_* \oplus \Delta) \oplus (\Delta\ 0^n)$ 25: else if $M_* > 0$ then 26: $\text{T} \leftarrow 101\ N$ 27: $R \leftarrow \text{F}_K^{\text{T}}((M_*\ 10^*) \oplus \Delta) \oplus (\Delta\ 0^n)$ 28: $R \leftarrow \text{left}_{n+ M_* }(R)$ 29: else 30: $R \leftarrow \Delta$ 31: end if 32: return $C_1\ \dots\ C_m\ R$ 33: end Algorithm </pre>	<pre> 1: Algorithm $\mathcal{D}(K, N, A, C)$ 2: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 3: $C_1, \dots, C_m, C_*, \text{Tcsplit-b}_n C$ 4: $\text{noM} \leftarrow 0$ 5: if $C = n$ then $\text{noM} \leftarrow 1$ 6: $\Delta \leftarrow 0^n$; $\text{T} \leftarrow 000\ N$ 7: for $i \leftarrow 1$ to a do 8: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}(A_i \oplus \Delta))$ 9: end for 10: if $A_* = n$ then 11: $\text{T} \leftarrow \text{noM}\ 10\ N$ 12: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}(A_* \oplus \Delta))$ 13: else if $A_* > 0$ or $T = 0$ then 14: $\text{T} \leftarrow \text{noM}\ 11\ N$ 15: $\Delta \leftarrow \text{right}_n(\text{F}_K^{\text{T}}((A_*\ 10^*) \oplus \Delta))$ 16: end if \triangleright Do nothing if $A = \varepsilon, M \neq \varepsilon$ 17: $\text{T} \leftarrow 001\ N$ 18: for $i \leftarrow 1$ to m do 19: $M_i \leftarrow \text{F}^{-1}_{\text{T}}(C_i \oplus \Delta, 0) \oplus \Delta$ 20: $\Delta \leftarrow \text{F}^{\rho}_{\text{T}}(C_i \oplus \Delta, 0)$ 21: end for 22: if $T = n$ then 23: $\text{T} \leftarrow 100\ N$ 24: $\bar{T} \leftarrow T$ 25: else if $T > 0$ then 26: $\text{T} \leftarrow 101\ N$ 27: $\bar{T} \leftarrow 10^{n- T -1}\ T$ 28: else 29: if $\Delta \neq C_*$ then return \perp 30: return ε 31: end if 32: $\bar{T} \leftarrow \text{left}_{ T }(\text{F}^{\rho}_{\text{T}}(C_* \oplus \Delta, 0))$ 33: $M_* \leftarrow \text{F}^{-1}_{\text{T}}(C_* \oplus \Delta, 0) \oplus \Delta$ 34: if $\bar{T} \neq \text{right}_{ T }(M_*)\ \bar{T}$ then 35: return \perp 36: end if 37: $M_* \leftarrow \text{left}_{ T }(M_*)$ 38: return $M_1\ \dots\ M_m\ M_*$ 39: end Algorithm </pre>
--	---

Figure 11: The SAEF[F] AEAD scheme.

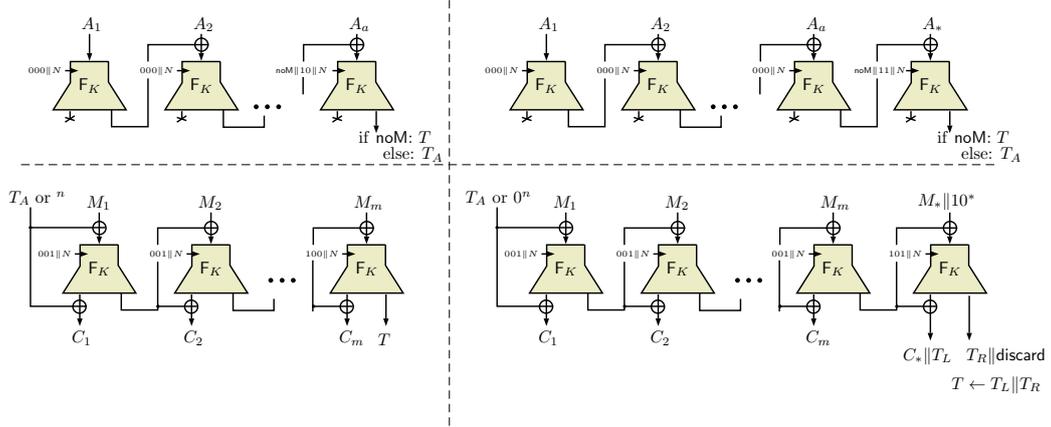


Figure 12: The encryption algorithm of SAEF[F] mode. The bit $\text{noM} = 1$ iff $|M| = 0$. The picture illustrates the processing of AD when length of AD is a multiple of n (**top left**) and when the length of AD is not a multiple of n (**top right**), and the processing of the message when length of the message is a multiple of n (**bottom left**) and when the length of message is not a multiple of n (**bottom right**)

Corollary 1. From $\sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \ell \cdot q_{\ell} \leq 2\sigma$ it immediately follows that

$$\begin{aligned} \text{Adv}_{\text{SAEF[F]}}^{\text{priv}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{B}) + 12 \cdot \frac{\sigma^2}{2^n}, \\ \text{Adv}_{\text{SAEF[F]}}^{\text{auth}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{C}) + 12 \cdot \frac{\sigma^2}{2^n} + 4 \cdot \frac{q_v \cdot (2^{\ell_{\max}} + 1)}{2^n} + \frac{q_v}{(2^n - 1)}. \end{aligned}$$

Remark 1 (Interpretation of the bounds in Theorem 3 and Corollary 1). Corollary 1 clearly shows that in general, the security of SAEF is birthday bounded, which is unusual for a mode of operation for a tweakable primitive. However, SAEF was designed especially for applications where the encryption queries are very *short*. The security guarantees given by the fine-grained bound in Theorem 3 become very competitive in this case.

For example, in an (IoT-like) application where q_v is reasonably small, and where we know for certain that the vast majority of queries consists of a single block of AD and a single message block, and a tiny portion $\gamma \ll q$ of messages is expected to contain a small number of message blocks β , the contribution of the short queries vanishes and the security bounds of SAEF become near optimal:

$$\begin{aligned} \text{Adv}_{\text{SAEF[F]}}^{\text{priv}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{B}) + 3 \cdot \frac{\gamma \cdot \beta \cdot (\beta - 1)}{2^n} \\ \text{Adv}_{\text{SAEF[F]}}^{\text{auth}}(\mathcal{A}) &\leq \text{Adv}_{\text{F}}^{\text{prtfp}}(\mathcal{C}) + 3 \cdot \frac{\gamma \cdot \beta \cdot (\beta - 1)}{2^n} + \frac{q_v \cdot (\beta^2 + 1)}{2^n} + \frac{q_v}{(2^n - 1)} \end{aligned}$$

Remark 2 (Sequential encryption vs forgeries.). The sequential structure of SAEF causes a birthday-type deterioration of the security bounds in the nonce-respecting setting. However, we conjecture that it also lends a degree for resistance to forgeries up to the birthday bound in the nonce-misuse setting to SAEF. Right now this is an unproven claim, and we intend to prove it for the full version of the paper.

Here is an intuition behind this claim. In the case of nonce-reuse, issuing queries with no AD allows the adversary to learn direct input-“left” output pairs for F . However, the adversary never learns the “right” outputs of F for the non-final blocks of message. Therefore, finding a collision on the input to F that would allow a forgery should still take a query complexity at the birthday bound.

Proof of Theorem 3. The security analysis of SAEF is slightly more involved than in the case of PAEF. We first tackle confidentiality and then integrity.

Confidentiality of SAEF. We first replace the forkcipher F with a pair of tweakable permutations π_0 and π_1 . I.e. $\pi_0 = (\pi_{T,0} \leftarrow \$ \text{Perm}(n))_{T \in \{0,1\}^t}$ is a collection of independent uniform elements of $\text{Perm}(n)$ indexed by the elements of $T \in \{0,1\}^t$ (and similarly for $\pi_1 = (\pi_{T,1} \leftarrow \$ \text{Perm}(n))_{T \in \{0,1\}^t}$). We let $\text{SAEF}[\pi_0, \pi_1]$ denote the SAEF mode that uses π_0, π_1 instead of F . This replacement implies the following inequality:

$$\mathbf{Adv}_{\text{SAEF}[F]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{prtfp}}(\mathcal{B}) + \mathbf{Adv}_{\text{SAEF}[\pi_0, \pi_1]}^{\text{priv}}(\mathcal{A})$$

by a similar argument as in the proof of Theorem 2.

We now further replace the two families of random permutations π_0 and π_1 with families of *random functions* f_0 and f_1 with the same signature. I.e. $f_b = (f_{T,b} \leftarrow \$ \text{Func}(n))_{T \in \{0,1\}^t}$ for $b \in \{0,1\}$. Denoting the SAEF mode using these random functions by $\text{SAEF}[f_0, f_1]$, we have that

$$\mathbf{Adv}_{\text{SAEF}[\pi_0, \pi_1]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{SAEF}[f_0, f_1]}^{\text{priv}}(\mathcal{A}) + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} 2 \cdot \frac{q_\ell \cdot \ell \cdot (\ell - 1)}{2^n}$$

because each query uses a distinct nonce, and thus for each N , the functions $f_{T,0}$ and $f_{T,1}$ with $T = b_0 b_1 b_2 \| N$ and $b_0, b_1, b_2 \in \{0,1\}$ are always used only in a single query. Moreover, for $b_0 b_1 b_2 \in \{0,1\}^3 \setminus \{000, 001\}$, the associated random functions are used at most once, so their distribution is indistinguishable from that of the corresponding random permutations. For $b_0 b_1 b_2 \in \{000, 001\}$, there will be exactly q_ℓ values of N for which both $f_{T,0}$ and $f_{T,1}$ will be used at most ℓ times. Replacing such a $\pi_{b_0 b_1 b_2 \| N, b}$ by $f_{b_0 b_1 b_2 \| N, b}$ augments the bound by at most $\ell \cdot (\ell - 1) \cdot 2^{-n-1}$ by the RP-RF switching lemma [16] and a standard hybrid argument. By summing over ℓ , $b_0 b_1 b_2 \in \{000, 001\}$ and $b \in \{0,1\}$, we obtain the bound.

We now bound $\mathbf{Adv}_{\text{SAEF}[f_0, f_1]}^{\text{priv}}(\mathcal{A})$. For this, we use the games G_0 and G_1 defined in Figure 13. In both games, the set \mathcal{D}_T collects the domain points, on which the functions $f_{T,0}$ and $f_{T,1}$ were already evaluated. It is easy to verify that G_0 actually implements **priv-real** $_{\text{SAEF}[f_0, f_1]}$, as the flag **bad** and the sets \mathcal{D}_T have no influence on the outputs of Enc . It is also possible to verify that $\Pr[\mathcal{A}^{\text{priv-ideal}}_{\text{SAEF}[f_0, f_1]} \Rightarrow 1] = \Pr[\mathcal{A}^{G_1} \Rightarrow 1]$: unless **bad** is set, every ciphertext block C_i is an image of a distinct input to $f_{T,0}$. The final blocks C_* and T are an xor of the sum of previous “right” blocks, and of bits produced by two random functions $f_{10b \| N, 0}$ and $f_{10b \| N, 1}$ that are each used exactly once (due to the non-repetition of the nonces and the domain separation flags $10b$ for $b \in \{0,1\}$). Thus, all the output bits of Enc are uniform. Once **bad** is set, all the ciphertext blocks and each value of Δ is replaced by a uniform string, so the simulation is perfect. Thus we have $\mathbf{Adv}_{\text{SAEF}[f_0, f_1]}^{\text{priv}}(\mathcal{A}) \leq \Pr[\mathcal{A}^{G_0} \Rightarrow 1] - \Pr[\mathcal{A}^{G_1} \Rightarrow 1]$.

We also have that G_0 and G_1 are identical until **bad**, so by the Fundamental lemma of gameplaying [16] we have that $\mathbf{Adv}_{\text{SAEF}[f_0, f_1]}^{\text{priv}}(\mathcal{A}) \leq \Pr[\mathcal{A}^{G_0} \text{ sets bad}]$, where \mathcal{A}^{G_0} sets **bad** denotes the event that **bad** = **true** when \mathcal{A} issues its final output. We bound $\Pr[\mathcal{A}^{G_0} \text{ sets bad}]$ by union bound, iterating over the probability that the i^{th} query sets **bad**, given that **bad** was not set before.

For an encryption query (N, A, M) with $\ell = |M|_n$, all the sets $\mathcal{D}_{b_0 b_1 b_2 \| N}$ used in the query are initially empty, due to the uniqueness of N . The flag **bad** can only be set due to a collision on $\mathcal{D}_{001 \| N}$. We use the tweak $001 \| N$ when processing the message. With the first message block, the adversary can possibly choose the input to $f_{001 \| N, 1}$, but this cannot cause a collision because $\mathcal{D}_{001 \| N}$ is empty. Then for every next block M_i , the input to $f_{001 \| N, 1}$ is masked by the current value Δ . If **bad** has not been set

before, Δ is a fresh, uniform value, so the probability that $(M_i \oplus \Delta) \in \mathcal{D}_{001\|N}$ is at most $|\mathcal{D}_{001\|N}| \cdot 2^{-n} = (i-1) \cdot 2^{-n}$. The probability that **bad** gets set due to AD processing is then bounded by $\ell \cdot (\ell-1) \cdot 2^{-n-1}$ by summing over $i = 2, \dots, \ell$.

We know that for each ℓ , there are exactly \mathbf{q}_ℓ queries (N, A, M) with $\ell = \max(|A|_n, |M|_n)$ and for $\ell > \ell_{\max}$, $\mathbf{q}_\ell = 0$. So, by applying the union bound over all queries, we get $\Pr[\mathcal{A}^{G_0} \text{ sets } \mathbf{bad}] \leq \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \mathbf{q}_\ell \cdot \ell \cdot (\ell-1) \cdot 2^{-n}$. Combined with the previously derived inequalities, this completes the proof of the confidentiality bound.

<pre> 1: proc initialize 2: for $\mathsf{T} \in \{0, 1\}^t$ do 3: $f_{\mathsf{T},0} \leftarrow \text{\\$} \text{Func}(n)$ 4: $f_{\mathsf{T},1} \leftarrow \text{\\$} \text{Func}(n)$ 5: $\mathcal{D}_{\mathsf{T}} \leftarrow \emptyset$ 6: end for 7: bad \leftarrow false 1: proc Enc(N, A, M) 2: $A_1, \dots, A_a, A_* \xleftarrow{\\$} A$ 3: $M_1, \dots, M_m, M_* \xleftarrow{\\$} M$ 4: noM \leftarrow 0 5: if $M = 0$ then noM \leftarrow 1 6: $\Delta \leftarrow 0^n$; $\mathsf{T} \leftarrow 000\ N$ 7: for $i \leftarrow 1$ to a do 8: $\Delta \leftarrow f_{\mathsf{T},1}(A_i \oplus \Delta)$ 9: end for 10: if $A_* = n$ then 11: $\mathsf{T} \leftarrow \text{noM}\ 10\ N$ 12: $\Delta \leftarrow f_{\mathsf{T},1}(A_* \oplus \Delta)$ 13: else if $A_* > 0$ or $M = 0$ then 14: $\mathsf{T} \leftarrow \text{noM}\ 11\ N$ 15: $\Delta \leftarrow f_{\mathsf{T},1}((A_* \ 10^*) \oplus \Delta)$ 16: end if 17: $\mathsf{T} \leftarrow 001\ N$ </pre>	<pre> 18: for $i \leftarrow 1$ to m do 19: if $M_i \oplus \Delta \in \mathcal{D}_{\mathsf{T}}$ then 20: bad \leftarrow true 21: end if 22: $\mathcal{D}_{\mathsf{T}} \leftarrow \mathcal{D}_{\mathsf{T}} \cup (M_i \oplus \Delta)$ 23: $C_i \leftarrow f_{\mathsf{T},0}(M_i \oplus \Delta) \oplus \Delta$ 24: $\Delta \leftarrow f_{\mathsf{T},1}(M_i \oplus \Delta)$ 25: if bad = true then 26: $C_i \ \Delta \leftarrow \text{\\$} \{0, 1\}^{2n}$ 27: end if 28: end for 29: if $M_* = n$ then 30: $\mathsf{T} \leftarrow 100\ N$ 31: $C_* \leftarrow f_{\mathsf{T},0}(M_* \oplus \Delta) \oplus \Delta$ 32: $T \leftarrow f_{\mathsf{T},1}(M_* \oplus \Delta)$ 33: else if $M_* > 0$ then 34: $\mathsf{T} \leftarrow 101\ N$ 35: $C_* \leftarrow f_{\mathsf{T},0}((M_* \ 10^*) \oplus \Delta) \oplus \Delta$ 36: $T \leftarrow f_{\mathsf{T},1}((M_* \ 10^*) \oplus \Delta)$ 37: $T \leftarrow \text{left}_{ M_* }(T)$ 38: else 39: $T \leftarrow \Delta$ 40: end if 41: return $C_1 \ \dots \ C_m \ C_* \ T$ </pre>
--	---

Figure 13: The games G_0 and G_1 for bounding $\text{Adv}_{\text{SAEF}[f_0, f_1]}^{\text{priv}}$. The game G_0 does *not* contain the boxed statement, while G_1 does.

Integrity of SAEF. We again replace the forkcipher F with a pair of tweakable permutations $\pi_0 = (\pi_{\mathsf{T},0} \leftarrow \text{\$} \text{Perm}(n))_{\mathsf{T} \in \{0,1\}^t}$ and $\pi_1 = (\pi_{\mathsf{T},1} \leftarrow \text{\$} \text{Perm}(n))_{\mathsf{T} \in \{0,1\}^t}$, such that we have

$$\text{Adv}_{\text{SAEF}[F]}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prtfp}}(\mathcal{C}) + \text{Adv}_{\text{SAEF}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A})$$

by a similar argument as in the proof of Theorem 2.

We additionally replace the tweakable permutation π_1 by a tweakable function f_1 with the same signature, yielding

$$\text{Adv}_{\text{SAEF}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_{\text{SAEF}[\pi_0, f_1]}^{\text{auth}}(\mathcal{A}) + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \frac{\mathbf{q}_\ell \cdot \ell \cdot (\ell-1)}{2^n}$$

by a similar argument as in the proof of SAEF's confidentiality.

To bound $\text{Adv}_{\text{SAEF}[\pi_0, f_1]}^{\text{auth}}(\mathcal{A})$, we consider the games G_2 and G_3 in Figure 14. It is easy to see that the game G_2 actually implements the game $\text{auth}_{\text{SAEF}[\pi_0, f_1]}$, because the sets

\mathcal{D}_T for $T \in \{0, 1\}^t$ and the flag **bad** have no effect on the outputs of the game. Moreover, unless **bad** is set to true, the games G_2 and G_3 execute the same code. Thus, by the Fundamental lemma of gameplaying [16], we have that $\Pr[\mathcal{A}^{G_2} \text{ forges}] - \Pr[\mathcal{A}^{G_3} \text{ forges}] \leq \Pr[\mathcal{A}^{G_2} \text{ sets bad}]$ and consequently $\mathbf{Adv}_{\text{SAEF}[\pi_0, f_1]}^{\text{auth}}(\mathcal{A}) \leq \Pr[\mathcal{A}^{G_2} \text{ sets bad}] + \Pr[\mathcal{A}^{G_3} \text{ forges}]$.

Transition from G_2 to G_3 . The flag **bad** being set means that for some $T \in \{0, 1\}^t$, the permutation $\pi_{T,0}$ and the function $f_{T,1}$ were used twice on the same input in an encryption queries. Informally speaking, this event may allow the adversary to forge trivially by simply truncating the ciphertext, or the associated data used in an encryption query with such a collision. We disallow this kind of victory in the game G_3 .

Similarly as in the proof of confidentiality bound, we bound $\Pr[\mathcal{A}^{G_2} \text{ sets bad}]$ by the union bound, iterating over the probability that the i^{th} query sets **bad**, given that **bad** was not set before. In an encryption query (N, A, M) with $\ell = \max(|A|_n, |M|_n)$, the flag **bad** can only be set by a collision on set $\mathcal{D}_{b\|N}$ with $b \in \{000, 001\}$. For the first block of AD (and the first block of message) there can be no collision because $\mathcal{D}_{000\|N}$ is initially empty (as is $\mathcal{D}_{001\|N}$). For each subsequent block A_i with $2 \leq i < \ell$, A_i is xored with a mask Δ . Given that **bad** has not been set before, Δ is an image of a fresh input under $f_{(000\|N),1}$. Therefore, the probability that A_i sets **bad** is bounded by $|\mathcal{D}_{000\|N}| \cdot 2^{-n} \leq (i-1) \cdot 2^{n-1}$ due to the fact $i < \ell \leq \ell_{\max}$ and due to the assumption $\ell_{\max} \leq 2^{n-1}$. Summing over i , we get that AD processing cannot set **bad** with a probability bigger than $\ell \cdot (\ell-1) \cdot 2^{-n}$. The probability that **bad** is set while processing message blocks is bounded by $\ell \cdot (\ell-1) \cdot 2^{-n}$ by a similar argument. Summing over all queries and knowing that for each ℓ , there are exactly \mathbf{q}_ℓ queries with $\ell = \max(|A|_n, |M|_n)$, we get that $\Pr[\mathcal{A}^{G_2} \text{ sets bad}] \leq \sum_{\ell \in \mathbb{L} \cap \{0, \dots, \ell_{\max}\}} \mathbf{q}_\ell \cdot 2\ell \cdot (\ell-1) \cdot 2^{-n}$.

Forgery in G_3 . We proceed to bounding $\Pr[\mathcal{A}^{G_3} \text{ forges}]$. We carry out the analysis for an adversary \mathcal{A}' that makes a single verification query, and then obtain $\Pr[\mathcal{A}^{G_3} \text{ forges}] \leq q_v \cdot \Pr[\mathcal{A}'^{G_3} \text{ forges}]$, referring to a result by Bellare to support the claim [10].

In what follows, we let $(N^i, A^i, M^i), C^i$ denote the i^{th} encryption query made by \mathcal{A}' , and (N, A, C) denote the only decryption query. For each i , we let $C_1^i, \dots, C_m^i, C_*^i, T^i \leftarrow \text{csplit-b}_n(C^i)$ and we let $C_1, \dots, C_m, C_*, T \leftarrow \text{csplit-b}_n(C)$. Additionally, we will refer to the values of the Δ variable. We will indicate by $\Delta_{A,j}$ the j^{th} value that the variable Δ takes when processing the j^{th} block of A from the decryption query (N, A, C) , and by $\Delta_{M,j}$ the j^{th} value that the variable Δ takes when processing the j^{th} block of the ciphertext C . We note that we can have $j = *$ and that $\Delta_{A,1} = 0^n$. We define $\Delta_{A,j}^i$ and $\Delta_{M,j}^i$ in a similar way for (N^i, A^i, M^i) .

We bound $\Pr[\mathcal{A}^{G_3} \text{ forges}]$ by the following case analysis.

Case 1, $N \neq N^i$ for all $1 \leq i \leq q$: In this case, T is either compared to an output of $f_{(b\|N),1}$ with $b \in \{100, 110, 111\}$, or the rightmost $n - |T|$ bits of the preimage $\pi_{(101\|N),0}^{-1}(C_*)$ must have a specific value, and T is compared to $|T|$ bits of an image under $f_{(101\|N),1}$. Because all these random permutations (respectively functions) have not yet been sampled, the probability that the forgery succeeds is at most 2^{-n} .

In all the remaining cases, we assume that $\exists 1 \leq i \leq q : N = N^i$, i.e. the nonce N is reused from the i^{th} encryption query.

Case 2, $|T| = n$ and $|T^i| < n$, or $|T| < n$ and $|T^i| = n$: If $|T| = n$, then it must be equal to an image under a so far unused random function $f_{(100\|N),1}$ (because of the three-bit flag 100). Thus forging is equivalent to guessing the value of n random bits and succeeds with probability 2^{-n} .

If on the other hand $|T| < n$, then either the rightmost $n - |T|$ bits of the preimage $\pi_{(101\|N),0}^{-1}(C_*)$ must have a specific value and T is compared to $|T|$ bits of an image under

$f_{(101\|N),1}$, or C_* is compared to the output of $f_{(11b\|N),1}$ for $b \in \{0,1\}$ (if $|T| = 0$). With each of these three permutations unused in the experiment until now, the forgery succeeds with probability 2^{-n} .

Case 3, $|T| = n \wedge |T^i| = n$: The tag T is compared to an image under the random function $f_{(100\|N),1}$, which was also used to produce T^i . The adversary can forge in two ways, either by making sure that $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$ and reusing $T = T^i$, or by trying to guess the tag $T \neq T^i$ without this collision. Abusing the notation, we therefore have $\Pr[\mathcal{A}^{G_3} \text{ forges} | \mathbf{Case 3}] \leq \Pr[\text{coll} | \mathbf{Case 3}] + \Pr[\mathcal{A}^{G_3} \text{ forges} | \neg \text{coll} \wedge \mathbf{Case 3}]$ where coll denotes the event $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$ conditioned on the event that $\text{bad} \neq \text{true}$ in game G_3 .

We have that $\Pr[\mathcal{A}^{G_3} \text{ forges} | \neg \text{coll} \wedge \mathbf{Case 3}] \leq 2^{-n}$, because forging is equivalent to guessing the value of a fresh image under a random function. To bound $\Pr[\text{coll} | \mathbf{Case 3}]$, we investigate four subcases:

Case 3.1, $|A_*^i| = n$ and $|A_*| < n$ or $|A_*^i| < n$ and $|A_*| = n$: For the event coll to occur, there either are two indices $1 \leq j \leq m^i$ and $1 \leq j' \leq m$ for which $M_j^i \oplus \Delta_{M,j}^i = M_{j'}^i \oplus \Delta_{M,j'}^i$, or there are no such two indices and yet $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$ (in the first case, \mathcal{A}' may simply reuse a part of C^i). Omitting the condition “**Case 3.1**”, the probability of coll is bounded by

$$\Pr[\text{coll} | \mathbf{Case 3.1}] \leq \sum_{j,j'} \Pr[\text{coll}_{j,j'}] + \Pr[\text{coll}_*]$$

where $\text{coll}_{j,j'}$ denotes the event $M_j^i \oplus \Delta_{M,j}^i = M_{j'}^i \oplus \Delta_{M,j'}^i$ conditioned on $\# \bar{j} : M_j^i \oplus \Delta_{M,\bar{j}}^i = M_{j'-1}^i \oplus \Delta_{M,j'-1}^i$ (if $j' > 1$), and where coll_* denotes the event $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$ conditioned on $\# \bar{j}', \bar{j} : M_{\bar{j}'}^i \oplus \Delta_{M,\bar{j}}^i = M_{\bar{j}'}^i \oplus \Delta_{M,\bar{j}'}^i$.

For $j' = 1$, \mathcal{A} may force $\Delta_{M,1} = 0^n$ by setting $A = \varepsilon$, making $\text{coll}_{j,1}$ equivalent to $\Delta_{M,j}^i = C_1 \oplus C_1^i$. However, all the masks $\Delta_{M,j}^i$ were generated uniformly, either due to the $\text{bad} \neq \text{true}$ condition for $j > 1$ or because $A^i \neq \varepsilon$ for $j = 1$, so the probability that any of them is equal to $C_1 \oplus C_1^i$ is 2^{-n} . If $A \neq \varepsilon$, then $\Delta_{M,1}$ is produced by a fresh random function, and therefore uniformly distributed, and the probability of $\text{coll}_{j,1}$ is thus 2^{-n} .

For $j' > 1$, the condition $\# \bar{j} : M_j^i \oplus \Delta_{M,\bar{j}}^i = M_{j'-1}^i \oplus \Delta_{M,j'-1}^i$ implies that $\Delta_{M,j'}$ is uniformly distributed, and so the probability of $\text{coll}_{j,j'}$ is 2^{-n} for any j .

Finally, the event coll_* is equivalent to guessing the value of random n -bit string, and occurs with probability 2^{-n} .

As we have $j \leq m^i \leq \ell_{\max}$ and $j' \leq m \leq \ell_{\max}$, the probability of coll in this subcase is at most $(\ell_{\max}^2 + 1) \cdot 2^{-n}$, and a forgery succeeds with probability at most $(\ell_{\max}^2 + 2) \cdot 2^{-n}$.

Case 3.2, $|A_*^i| = n$ and $|A| = n$: The analysis is very similar as in **Case 3.1**, except that now both $\Delta_{M,1}$ and $\Delta_{M,1}^i$ are produced by the same random function $f_{010\|N,1}$. We therefore need to consider the possibility, that \mathcal{A}' forces $\Delta_{M,1} = \Delta_{M,1}^i$ by triggering a non-trivial collision $\Delta_{A,*} \oplus A_* = \Delta_{A,1}^i \oplus A_*^i$. If no such collision happens, the analysis proceeds exactly as in **Case 3.1**.

The analysis of the collision in AD processing is almost identical to the analysis of the collision in the message processing, except that we trivially have $\Delta_{A,1} = \Delta_{A,1}^i = 0^n$. This does not help the adversary in achieving a *non-trivial* collision, because such a collision can only occur past the trivial blockwise common prefix of A and A^i .

Denoting by coll_A the event $\Delta_{A,*} \oplus A_* = \Delta_{A,1}^i \oplus A_*^i$, we have

$$\Pr[\text{coll}_A | \mathbf{Case 3.2}] \leq \sum_{j,j'} \Pr[\text{coll}_{j,j'}] + \Pr[\text{coll}_*]$$

where $\text{coll}_{j,j'}$ denotes the event $A_j^i \oplus \Delta_{A,j}^i = A_{j'}^i \oplus \Delta_{A,j'}^i$ conditioned on $\bar{j} : A_j^i \oplus \Delta_{A,\bar{j}}^i = A_{j'-1} \oplus \Delta_{A,j'-1}$ (if $j' > 1$), and where coll_* denotes the event $A_* \oplus \Delta_{A,*} = A_*^i \oplus \Delta_{A,*}^i$ conditioned on $\bar{j}, \bar{j} : A_{\bar{j}}^i \oplus \Delta_{A,\bar{j}}^i = A_{\bar{j}} \oplus \Delta_{A,\bar{j}}$, such that we iterate over $\min\{\bar{j} | A_{\bar{j}}^i \neq A_{\bar{j}}\} \leq j' \leq m$ and $1 \leq j \leq m^i$.

The probability of each $\text{coll}_{j,j'}$ is at most 2^{-n} by a similar argument as in the analysis of **Case 3.1** for all pairs $j'j$ except for $j = j' = \min\{\bar{j} | A_{\bar{j}}^i \neq A_{\bar{j}}\}$; for $j' = \min\{\bar{j} | A_{\bar{j}}^i \neq A_{\bar{j}}\}$ the event $\Delta_{A,j'} \oplus A_{j'} = \Delta_{A,j'}^i \oplus A_{j'}^i$ is impossible. There are never more than ℓ_{\max} possible values for j' .

We deduce that the probability of forger in this case is bounded by $(2\ell_{\max}^2 + 2) \cdot 2^{-n}$.

Case 3.3, $0 < |A_*^i| < n$ and $0 < |A| < n$: Because of the use of injective padding in AD processing, the analysis of this case is almost identical with the analysis of **Case 3.2**, and a forgery succeeds with probability bounded by $(2\ell_{\max}^2 + 1) \cdot 2^{-n}$.

Case 3.4, $A = A^i = \varepsilon$: In this special case, we have $\Delta_{M,1} = \Delta_{M,1}^i = 0^n$. The analysis in this case is the same as in case **Case 3.1**, except that we iterate over $\min\{\bar{j} | C_{\bar{j}}^i \neq C_{\bar{j}}\} \leq j' \leq m$, and deduce that a forgery succeeds with a probability bounded by $(\ell_{\max}^2 + 2) \cdot 2^{-n}$.

Summing up, the probability of forgery in **Case 3** is bounded by $(2\ell_{\max}^2 + 2) \cdot 2^{-n}$.

Case 4, $|T| = 0$ and $|T^i| = 0$: In this case, both messages are empty. We have the following subcases:

Case 4.1, $|A_*| = n$ and $|A_*^i| < n$, or $|A_*| < n$ and $|A_*^i| = n$: In this case, the value C_* is compared to an image of a random function that has not been sampled in the experiment. A forgery succeeds with probability 2^{-n} in this case.

Case 4.2, $|A_*| = n$ and $|A_*^i| = n$: In this case, \mathcal{A}' can manage to forge either by triggering the collision $\Delta_{A,*} \oplus A_* = \Delta_{A,1}^i \oplus A_*^i$, or by guessing the tag otherwise. By a similar argument as **Case 3.2**, the probability of forgery is no more than $(\ell_{\max}^2 + 2) \cdot 2^{-n}$ in this case.

Case 4.3, $|A_*| < n$ and $|A_*^i| < n$: Due to the use of the injective 10^* padding, this case is symmetric to **Case 4.2** and the probability of forgery is no more than $(\ell_{\max}^2 + 2) \cdot 2^{-n}$.

The probability of forgery in this case is no more than $(\ell_{\max}^2 + 2) \cdot 2^{-n}$.

Case 5, $0 < |T| < n$ and $0 < |T^i| < 0$: In this case, we have the following subcases:

Case 5.1, $|T| = |T^i|$: In this case, the adversary may forge either by forcing $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$ and reusing the tag, or by guessing a fresh valid pair (C_*, T) . The analysis is very similar as in **Case 3**, except that the probability of forgery conditioned on $C_* \oplus \Delta_{M,*} \neq C_*^i \oplus \Delta_{M,*}^i$ changes.

In particular, the freshly sampled preimage $X = \pi_{101\|N,0}^{-1}(C_* \oplus \Delta_{M,*})$ will need to be of the form $X = Z\|10^*$ for some $Z \in \{0, 1\}^{|T|}$ and simultaneously, the first $|T|$ bits of the freshly sampled image $Y = f_{101\|N,1}(X)$ will need to be equal to T . This happens with probability no greater than $((2^{|T|} - 1) \cdot (2^{n-|T|})) / ((2^n - 1) \cdot 2^n) \leq 1 / (2^n - 1)$.

The total probability of forgery is no greater than $(2\ell_{\max} + 1) \cdot 2^{-n} + (2^n - 1)^{-1}$

Case 5.2, $|T| \neq |T^i|$: As in the previous case, the adversary may attempt to trigger $C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i$. If this does not succeed, the probability of forgery is bounded by $(2^n - 1)/2^{2n}$ as in **Case 5.1**.

If this does succeed, the forgery is not as easily granted as before, as \mathcal{A}' also needs $C_* \oplus \Delta_{M,*} = Z\|10^*$ for some $Z \in \{0, 1\}^{|T|}$. Because clearly

$$\Pr[C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i \text{ and } C_* \oplus \Delta_{M,*} = Z\|10^*] \leq \Pr[C_* \oplus \Delta_{M,*} = C_*^i \oplus \Delta_{M,*}^i],$$

We conclude that the probability of forgery in this case is also bounded by $(2^{\ell_{\max}} + 1) \cdot 2^{-n} + (2^n - 1)^{-1}$.

The probability of forgery in this case is no more than $(2^{\ell_{\max}} + 1) \cdot 2^{-n} + (2^n - 1)^{-1}$.

By taking the maximum over all cases, the probability that a single-decryption-query adversary \mathcal{A}' forgers in the game G_3 is at most $(2^{\ell_{\max}} + 1) \cdot 2^{-n} + (2^n - 1)^{-1}$. The adversary \mathcal{A} making q_v decryption queries thus forges with probability bounded by $q_v \cdot (2^{\ell_{\max}} + 1) \cdot 2^{-n} + q_v \cdot (2^n - 1)^{-1}$. By back-substituting all the previous equalities, we obtain the claimed result. \square

5.6 Modifying GCM for a Forkcipher

We finally define fGCM (as “forkcipher-based GCM,”), a nonce-based AEAD scheme parameterized by a tweakable forkcipher F (as defined in Section 4) with $\mathcal{T} = \{0, 1\}^t$ for a positive t . An instance $\text{fGCM}[F] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ has a key space $\mathcal{K} = \{0, 1\}^k$, nonce space $\mathcal{N} = \{0, 1\}^{t-2}$, and the AD and message spaces are both $\{0, 1\}^{\leq 2^{n/2}}$ (although the maximal message length influences the security). The ciphertext expansion of $\text{fGCM}[F]$ is n . The encryption and decryption algorithms are defined in Figure 15 and the encryption algorithm is illustrated in Figure 16.

The scheme fGCM is inspired by the original Galois-Counter Mode [23] in that it processes all-but-last message block using counter mode, and computes a polynomial hash of thus-computed ciphertext blocks and the AD. Unlike the original GCM, the output of the GHASH is used as a masking offset for a primitive call, that encrypts the final ciphertext block and simultaneously produces the authentication tag.

In an encryption query, the message processed in blocks of n bits. The first $\max(0, \lceil |M|/n \rceil - 1)$ message blocks are encrypted with counter mode, using $\lceil \max(0, \lceil |M|/n \rceil - 1) / 2 \rceil$ calls to F (because each F call produces $2n$ pseudorandom bits). Then, AD and the string formed by concatenating the just-computed ciphertext blocks is processed by GHASH to obtain a masking value Δ . The key for GHASH is derived as $\text{left}_n(F(K, 0^t, 0^n))$. The final message block is padded to n bits if needed, and then processed by a single F call, whose input and “left” output are xor-masked by Δ . The first $|M| \bmod n + n$ bits of the output of this call form the final ciphertext block and the tag.

Each call to F made during the encryption is using a tweak that is composed of

1. the nonce, and
2. a two-bit flag f .

The flag f takes different values for processing of different types of data blocks in the encryption algorithm:

$f = 01$ when in counter mode,

$f = 10$ in the final call if $|M| \equiv 0 \pmod{n}$,

$f = 11$ in the final call if $|M| \not\equiv 0 \pmod{n}$.

The decryption recomputes the value Δ , with it decrypts the final ciphertext block and and verifies the integrity of the inputs. If the verification succeeds, the rest of the ciphertext gets decrypted, otherwise the error symbol is returned.

<pre> 1: proc initialize 2: for $T \in \{0, 1\}^t$ do 3: $\pi_{T,0} \leftarrow \text{Perm}(n)$ 4: $f_{T,1} \leftarrow \text{Func}(n)$ 5: $\mathcal{D}_T \leftarrow \emptyset$ 6: end for 7: bad \leftarrow false 1: proc Enc(N, A, M) 2: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 3: $M_1, \dots, M_m, M_* \xleftarrow{n} M$ 4: noM \leftarrow 0 5: if $M = 0$ then noM \leftarrow 1 6: $\Delta \leftarrow 0^n$; $T \leftarrow 000\ N$ 7: for $i \leftarrow 1$ to a do 8: if $A_i \oplus \Delta \in \mathcal{D}_T$ then 9: bad \leftarrow true 10: end if 11: $\mathcal{D}_T \leftarrow \mathcal{D}_T \cup (A_i \oplus \Delta)$ 12: $\Delta \leftarrow f_{T,1}(A_i \oplus \Delta)$ 13: end for 14: if $A_* = n$ then 15: $T \leftarrow \text{noM}\ 10\ N$ 16: $\Delta \leftarrow f_{T,1}(A_* \oplus \Delta)$ 17: else if $A_* > 0$ or $M = 0$ then 18: $T \leftarrow \text{noM}\ 11\ N$ 19: $\Delta \leftarrow f_{T,1}((A_*\ 10^*) \oplus \Delta)$ 20: end if 21: $T \leftarrow 001\ N$ 22: for $i \leftarrow 1$ to m do 23: if $M_i \oplus \Delta \in \mathcal{D}_T$ then 24: bad \leftarrow true 25: end if 26: $\mathcal{D}_T \leftarrow \mathcal{D}_T \cup (M_i \oplus \Delta)$ 27: $C_i \leftarrow \pi_{T,0}(M_i \oplus \Delta) \oplus \Delta$ 28: $\Delta \leftarrow f_{T,1}(M_i \oplus \Delta)$ 29: end for 30: if $M_* = n$ then 31: $T \leftarrow 100\ N$ 32: $C_* \leftarrow \pi_{T,0}(M_* \oplus \Delta) \oplus \Delta$ 33: $T \leftarrow f_{T,1}(M_* \oplus \Delta)$ 34: else if $M_* > 0$ then 35: $T \leftarrow 101\ N$ 36: $C_* \leftarrow \pi_{T,0}((M_*\ 10^*) \oplus \Delta) \oplus \Delta$ 37: $T \leftarrow f_{T,1}((M_*\ 10^*) \oplus \Delta)$ 38: $T \leftarrow \text{left}_{ M_* }(T)$ 39: else </pre>	<pre> 40: $T \leftarrow \Delta$ 41: end if 42: $T \leftarrow \Delta$ 43: return $C_1\ \dots\ C_m\ C_*\ T$ 1: proc Dec(N, A, C) 2: if bad = true then 3: return \perp 4: end if 5: $A_1, \dots, A_a, A_* \xleftarrow{n} A$ 6: $C_1, \dots, C_m, C_*, T \leftarrow \text{csplit-}b_n(C)$ 7: noM \leftarrow 0 8: if $C = n$ then noM \leftarrow 1 9: $\Delta \leftarrow 0^n$; $T \leftarrow 000\ N$ 10: for $i \leftarrow 1$ to a do 11: $\Delta \leftarrow f_{T,1}(A_i \oplus \Delta)$ 12: end for 13: if $A_* = n$ then 14: $T \leftarrow \text{noM}\ 10\ N$ 15: $\Delta \leftarrow f_{T,1}(A_* \oplus \Delta)$ 16: end if 17: if $A_* > 0$ or $T = 0$ then 18: $T \leftarrow \text{noM}\ 11\ N$ 19: $\Delta \leftarrow f_{T,1}((A_*\ 10^*) \oplus \Delta)$ 20: end if 21: $T \leftarrow 001\ N$ 22: for $i \leftarrow 1$ to m do 23: $M_i \leftarrow \pi_{T,0}^{-1}(C_i \oplus \Delta, 0) \oplus \Delta$ 24: $\Delta \leftarrow f_{T,1}(\pi_{T,0}^{-1}(C_i \oplus \Delta, 0))$ 25: end for 26: if $T = n$ then 27: $T \leftarrow 100\ N$ 28: $\bar{T} \leftarrow T$ 29: else if $T > 0$ then 30: $T \leftarrow 100\ N$ 31: $\bar{T} \leftarrow 10^{n- T -1}\ T$ 32: else 33: if $\Delta \neq C_*$ then return \perp 34: return \perp 35: end if 36: $M_* \leftarrow \pi_{T,0}^{-1}(C_* \oplus \Delta) \oplus \Delta$ 37: $\bar{T} \leftarrow \text{left}_{ T }(f_{T,1}(M_* \oplus \Delta))$ 38: if $\bar{T} \neq \text{right}_{ T }(M_*)\ \bar{T}$ then 39: return \perp 40: end if 41: $M_* \leftarrow \text{left}_{ T }(M_*)$ 42: return $M_1\ \dots\ M_m\ M_*$ </pre>
---	--

Figure 14: The games G_2 and G_3 for bounding $\text{Adv}_{\text{SAEF}[\pi_0, f_1]}^{\text{auth}}$. The game G_2 does not contain the boxed statements, while G_3 does.

<pre> 1: Algorithm $\mathcal{E}(K, N, A, M)$ 2: $L \leftarrow \text{left}_n(\mathbb{F}_K^{0^t}(0^n))$ 3: $M', M_* \leftarrow \text{msplit}_n(M)$ 4: $S \leftarrow \varepsilon$ 5: for $i \leftarrow 1$ to $\lceil M' /2n \rceil$ do 6: $S \leftarrow S \parallel \mathbb{F}_K^{01 \parallel N}(\langle i \rangle_n)$ 7: end for 8: $C' \leftarrow M' \oplus \text{left}_{ M' }(S)$ 9: $\Delta \leftarrow \text{GHASH}(L, A, C')$ 10: if $M_* = n$ then 11: $T \leftarrow 10 \parallel N$ 12: else 13: $M_* \leftarrow M_* \parallel 10^*$ 14: $T \leftarrow 11 \parallel N$ 15: end if 16: $R \leftarrow \mathbb{F}_K^T(M_* \oplus \Delta) \oplus (\Delta \parallel 0^n)$ 17: $C \leftarrow C' \parallel \text{left}_{ M_* +n}(R)$ 18: return C 19: end Algorithm 1: Algorithm $\text{GHASH}(L, A, C)$ 2: $\alpha \leftarrow n - (A \bmod n)$ 3: $\gamma \leftarrow n - (C \bmod n)$ 4: $X \leftarrow A \parallel 0^\alpha \parallel C \parallel 0^\gamma \parallel \langle A \rangle_{n/2} \parallel \langle C \rangle_{n/2}$ 5: $X_1, \dots, X_x \xleftarrow{n} X$ 6: $Y \leftarrow 0^n$ 7: for $j \leftarrow 1$ to x do 8: $L \cdot (Y \oplus X_j) \quad \triangleright \text{in } \text{GF}(2^n)$ 9: end for 10: return Y 11: end Algorithm </pre>	<pre> 1: Algorithm $\mathcal{D}(K, N, A, C)$ 2: $L \leftarrow \text{left}_n(\mathbb{F}_K^{0^t}(0^n))$ 3: $C', C_*, T \leftarrow \text{csplit}_n(C)$ 4: $\Delta \leftarrow \text{GHASH}(L, A, C')$ 5: if $T = n$ then 6: $T \leftarrow 10 \parallel N$ 7: $\bar{T} \leftarrow T$ 8: else 9: $T \leftarrow 11 \parallel N$ 10: $\bar{T} \leftarrow 10^{n- T -1} \parallel T$ 11: end if 12: $\tilde{T} \leftarrow \text{left}_{ T }(\mathbb{F}_K^T(C_* \oplus \Delta, 0))$ 13: $M_* \leftarrow \mathbb{F}_K^{-1 \parallel T}(C_* \oplus \Delta, 0) \oplus \Delta$ 14: if $\bar{T} \neq \text{right}_{n- T }(M_*) \parallel \tilde{T}$ then 15: return \perp 16: end if 17: $M_* \leftarrow \text{left}_{ T }(M_*)$ 18: $S \leftarrow \varepsilon$ 19: for $i \leftarrow 1$ to $\lceil C' /2n \rceil$ do 20: $S \leftarrow S \parallel \mathbb{F}_K^{01 \parallel N}(\langle i \rangle_n)$ 21: end for 22: $M' \leftarrow C' \oplus \text{left}_{ C' }(S)$ 23: return $M' \parallel M_*$ 24: end Algorithm </pre>
--	--

Figure 15: The fGCM[F] AEAD scheme. Refer to Section 2 for the definitions of string manipulation notation.

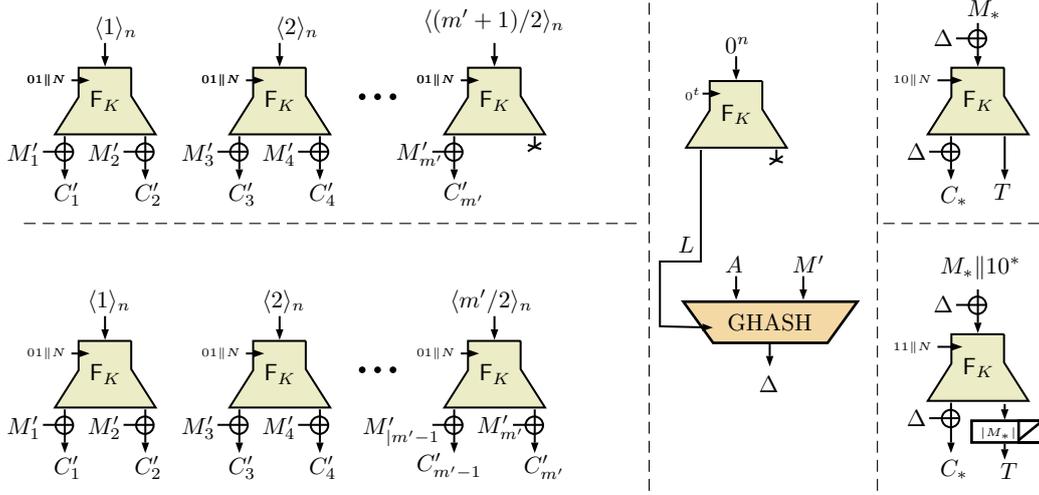


Figure 16: The encryption algorithm of fGCM[F] mode (see Figure 15 for the definition of M'). We have $m' = |M'|_n$. The picture illustrates the processing of M' when length of M' is an odd multiple of n (**top left**) and when the length of M' is an even multiple of n (**bottom left**), the processing of the final message block M_* when length of the message M is a multiple of n (**top right**) and when the length of message M is either zero or not a multiple of n (**bottom right**), and the computation of the value Δ (**middle**).

5.7 Security of fGCM

For the security analysis of fGCM, we will again work with the previously defined fine-grained adversarial resources. We state the formal claim about the nonce-based AE security of fGCM in Theorem 4.

Theorem 4. *Let F be a tweakable forkcipher with $\mathcal{T} = \{0, 1\}^t$. Then for any nonce-respecting adversary \mathcal{A} whose resources are bounded by $q, q_v, \sigma, \ell_{\max}$ and \mathbf{q} (in the sense we just defined) such that $\ell_{\max} \leq 2^n/4$, we have*

$$\begin{aligned} \text{Adv}_{\text{fGCM[F]}}^{\text{priv}}(\mathcal{A}) &\leq \text{Adv}_F^{\text{prtfp}}(\mathcal{B}) + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \frac{\mathbf{q}_\ell \cdot \ell \cdot (\ell - 1)}{2^n}, \\ \text{Adv}_{\text{fGCM[F]}}^{\text{auth}}(\mathcal{A}) &\leq \text{Adv}_F^{\text{prtfp}}(\mathcal{C}) + \frac{q_v \cdot \ell_{\max}}{2^n} + \frac{q_v \cdot 2^n}{(2^n - 1)^2} \end{aligned}$$

for some adversaries \mathcal{B} and \mathcal{C} who make at most 2σ queries, and who run in time given by the running time of \mathcal{A} plus $\gamma \cdot \sigma$ for some constant γ .

Corollary 2. *From $\sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} \ell \cdot \mathbf{q}_\ell \leq 2\sigma$ it immediately follows that*

$$\text{Adv}_{\text{fGCM[F]}}^{\text{priv}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prtfp}}(\mathcal{B}) + 4 \cdot \frac{\sigma^2}{2^n}$$

Remark 3 (Interpretation of the bounds in Theorem 4 and Corollary 2). According to Corollary 2, the security of fGCM is birthday bounded in the general case. However, when most of the encryption queries are very *short*, the security guarantees given by the fine-grained bound in Theorem 4 become very competitive.

Revisiting the example where q_v is reasonably small, and where the vast majority of queries consists of a single block of AD and a single message block, and a tiny portion $\gamma \ll q$ of messages is expected to contain a small number of message blocks β , the security

bounds of fGCM become near optimal:

$$\begin{aligned}\mathbf{Adv}_{\text{fGCM}[\mathbb{F}]}^{\text{priv}}(\mathcal{A}) &\leq \mathbf{Adv}_{\mathbb{F}}^{\text{prtfp}}(\mathcal{B}) + \frac{\gamma \cdot \beta \cdot (\beta - 1)}{2^n} \\ \mathbf{Adv}_{\text{fGCM}[\mathbb{F}]}^{\text{auth}}(\mathcal{A}) &\leq \mathbf{Adv}_{\mathbb{F}}^{\text{prtfp}}(\mathcal{C}) + \frac{q_v \cdot \beta}{2^n} + \frac{q_v \cdot 2^n}{(2^n - 1)^2}.\end{aligned}$$

The proof of Theorem 4 relies on the AXU property of GHASH stated in Lemma 1.

Lemma 1 (GHASH [23]). *For any positive integers n, ℓ_{\max} such that $\ell_{\max} \cdot n < 2^{n/2}$, for any pairs of strings $(A, C'), (\tilde{A}, \tilde{C}') \in (\{0, 1\}^{\leq \ell_{\max} \cdot n})^2$ such that $(A, C') \neq (\tilde{A}, \tilde{C}')$, and for any $c \in \{0, 1\}^n$ we have that*

$$\Pr[L \leftarrow_{\$} \{0, 1\}^n : \text{GHASH}(L, A, C') \oplus \text{GHASH}(L, \tilde{A}, \tilde{C}') = c] \leq \frac{\ell_{\max} + 1}{2^n}$$

with GHASH defined in Figure 15.

Proof of Theorem 4. We first tackle confidentiality and then integrity of fGCM.

Confidentiality of fGCM. As before, we first replace the forkcipher \mathbb{F} with a pair of tweakable permutations π_0 and π_1 (i.e., $\pi_0 = (\pi_{\mathbb{T},0} \leftarrow_{\$} \text{Perm}(n))_{\mathbb{T} \in \{0,1\}^t}$ is a collection of independent uniform elements of $\text{Perm}(n)$ indexed by the elements of $\mathbb{T} \in \{0, 1\}^t$, and similarly for $\pi_1 = (\pi_{\mathbb{T},1} \leftarrow_{\$} \text{Perm}(n))_{\mathbb{T} \in \{0,1\}^t}$). We let $\text{fGCM}[\pi_0, \pi_1]$ denote the fGCM mode that uses π_0, π_1 instead of \mathbb{F} , and deduce the following inequality:

$$\mathbf{Adv}_{\text{fGCM}[\mathbb{F}]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathbb{F}}^{\text{prtfp}}(\mathcal{B}) + \mathbf{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{priv}}(\mathcal{A})$$

by a similar argument as in the proof of Theorem 2.

We further replace the two families of random permutations π_0 and π_1 with families of *random functions* f_0 and f_1 with the same signature (i.e., $f_b = (f_{\mathbb{T},b} \leftarrow_{\$} \text{Func}(n))_{\mathbb{T} \in \{0,1\}^t}$ for $b \in \{0, 1\}$). Denoting the fGCM mode using these random functions by $\text{fGCM}[f_0, f_1]$, we have that

$$\mathbf{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{fGCM}[f_0, f_1]}^{\text{priv}}(\mathcal{A}) + \sum_{\ell \in L \cap \{0, \dots, \ell_{\max}\}} 2 \cdot \frac{q_{\ell} \cdot \ell \cdot (\ell - 1)}{2^n}$$

because each query uses a distinct nonce, and thus for each N , the functions $f_{\mathbb{T},0}$ and $f_{\mathbb{T},1}$ with $\mathbb{T} = b_0 b_1 \| N$ and $b_0 b_1 \in \{0, 1\}^2$ are always used only in a single query. For $b_0 b_1 \in \{00, 10, 11\}$, the associated random functions are used at most once, so their distribution is indistinguishable from that of the corresponding random permutations. For $b_0 b_1 = 01$, there will be exactly q_{ℓ} values of N for which both $f_{\mathbb{T},0}$ and $f_{\mathbb{T},1}$ will be used at most ℓ times. Replacing such a $\pi_{01 \| N, b}$ by $f_{01 \| N, b}$ augments the bound by at most $\ell \cdot (\ell - 1) \cdot 2^{-n-1}$ by the RP-RF switching lemma [16] and a standard hybrid argument. By summing over ℓ , and $b \in \{0, 1\}$, we obtain the bound.

Bounding $\mathbf{Adv}_{\text{fGCM}[f_0, f_1]}^{\text{priv}}(\mathcal{A})$ is relatively straightforward, because in each ciphertext C that \mathcal{A} sees, one part of C is computed as an xor of message bits and images of the random functions $f_{01 \| N, 0}$ and $f_{01 \| N, 1}$. Since none of the functions is ever evaluated twice on any input, all these bits are uniformly distributed. The remainder of C consists of a concatenation of the images of two random functions, that are only evaluated once during the whole experiment. Every ciphertext is thus distributed uniformly, and we have $\mathbf{Adv}_{\text{fGCM}[f_0, f_1]}^{\text{priv}}(\mathcal{A}) = 0$.

Integrity of fGCM. We again replace the forkcipher F with a pair of tweakable permutations $\pi_0 = (\pi_{T,0} \leftarrow \$ \text{Perm}(n))_{T \in \{0,1\}^t}$ and $\pi_1 = (\pi_{T,1} \leftarrow \$ \text{Perm}(n))_{T \in \{0,1\}^t}$, and have

$$\text{Adv}_{\text{fGCM}[F]}^{\text{auth}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prtfp}}(C) + \text{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A})$$

by a similar argument as in the proof of Theorem 2.

To bound $\text{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A})$, we first use the result of Bellare [10] that states that

$$\text{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A}) \leq q_v \cdot \text{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A}')$$

where \mathcal{A}' is an adversary with identical resources to those of \mathcal{A} , except that \mathcal{A}' only makes a single decryption query.

What remains is to bound $\text{Adv}_{\text{fGCM}[\pi_0, \pi_1]}^{\text{auth}}(\mathcal{A}')$. The adversary \mathcal{A}' makes q encryption queries N^i, A^i, M^i and receives ciphertexts C^i in response for $i = 1, \dots, q$. For each i we define $M_*^i, M_*^i = \text{msplit}_n(M^i)$ and $C_*^i, C_*^i, T^i = \text{csplit}_n(C^i)$, and we let Δ^i denote the output of GHASH in the i^{th} query. \mathcal{A}' also makes a single decryption query (N, A, C) , for which we define $C', C_*, T = \text{csplit}_n(C)$, and let Δ denote the output of GHASH evaluation in the decryption query. We proceed with the following case analysis.

Case 1, $N \neq N^i$ for all $1 \leq i \leq q$: T is either compared to an output of $\pi_{(b_0 b_1 \| N), 1}$ with $b_0 b_1 \in \{10, 11\}$, or the rightmost $n - |T|$ bits of the preimage $\pi_{(11 \| N), 0}^{-1}(C_*)$ must have a specific value, and T is compared to $|T|$ bits of an image under $\pi_{(11 \| N), 1}$. Because all these random permutations have not yet been sampled, the probability that the forgery succeeds is at most 2^{-n} .

In all the remaining cases, we assume that $\exists 1 \leq i \leq q : N = N^i$, i.e., the nonce N is reused from the i^{th} encryption query.

Case 2, $|T| = n \wedge |T^i| < n$ or $|T| < n \wedge |T^i| = n$: If $|T| = n$, the tag T is compared to an image under the random permutation $\pi_{(10 \| N), 1}$. However, no image under this permutation was used in the game (because we only used $\pi_{(11 \| N), 1}$). Thus forging is equivalent to guessing the value of n random bits and succeeds with probability 2^{-n} .

If on the other hand $|T| < n$, the rightmost $n - |T|$ bits of the preimage $\pi_{(11 \| N), 0}^{-1}(C_*)$ must have a specific value, and T is compared to $|T|$ bits of an image under $\pi_{(11 \| N), 1}$. These two permutation being unused until now, the forgery succeeds with probability 2^{-n} .

Case 3, $|T| = n \wedge |T^i| = n$: We will investigate two subcases:

Case 3.1, $(A, C') \neq (A^i, C'^i)$: The forgery can succeed in two ways in this case. If the adversary reuses the tag, i.e. $T = T^i$, then the forgery succeeds only if $\pi_{(10 \| N), 1}(\pi_{(10 \| N), 0}^{-1}(C_* \oplus \Delta)) = \pi_{(10 \| N), 1}(\pi_{(10 \| N), 0}^{-1}(C_*^i \oplus \Delta^i))$, which is equivalent with the event $\Delta \oplus \Delta_i = C_* \oplus C_*^i$. As $(A, C') \neq (A^i, C'^i)$, the probability of this event is no more than $\ell_{\max}/2^n$ due to the AXU property of GHASH (note that ℓ_{\max} is the maximal number of blocks in M , but C' always has 1 block less than M).

Otherwise, if $T \neq T^i$, the forgery succeeds if and only if $C_* \oplus \Delta \neq C_*^i \oplus \Delta_i \wedge \pi_{(10 \| N), 1}(\pi_{(10 \| N), 0}^{-1}(C_* \oplus \Delta)) = T$. Because $\pi_{(10 \| N), 1} \circ \pi_{(10 \| N), 0}^{-1}$ is a random permutation, and it was only evaluated on a single image, the probability of this event is at most $1/(2^n - 1)$.

Because the two ways of forging in this case are mutually exclusive, the probability of forgery is bounded by $\max(\ell_{\max}/2^n, 1/(2^n - 1))$.

Case 3.2, $(A, C') = (A^i, C'^i)$: Here, we must have $(C_*, T) \neq (C_*^i, T^i)$ for the forgery to be valid, but we also have $\Delta = \Delta^i$. This implies that $C_* \neq C_*^i$ and $T \neq T^i$ (because the tags are produced as images of permutations). A successful forgery implies that $\pi_{(10 \| N), 1}(\pi_{(10 \| N), 0}^{-1}(C_* \oplus \Delta)) = T$, such that the random permutation $\pi_{(10 \| N), 1}$

$\pi_{(10\|N),0}^{-1}$ was only evaluated on $C_*^i \oplus \Delta^i$ before, which occurs with probability at most $1/(2^n - 1)$.

The probability of forgery in this case is no more than $\max(\ell_{\max}/2^n, 1/(2^n - 1))$.

Case 4, $|T| < n \wedge |T^i| < n$: Again, we have two subcases:

Case 4.1, $(A, C') \neq (A^i, C'^i)$: Similarly to **Case 3**, \mathcal{A}' can succeed in two ways in this case. The adversary may succeed in creating the collision $C_* \oplus \Delta = C_*^i \oplus \Delta_i$ and forge by setting $T = T^i$, which happens with probability bounded by $\ell_{\max}/2^n$ due to the AXU property of GHASH.

Otherwise, if $C_* \oplus \Delta \neq C_*^i \oplus \Delta_i$, the forgery succeeds if and only if $\pi_{(10\|N),0}^{-1}(C_* \oplus \Delta) \oplus \Delta = Y\|10^{n-|T|-1} \wedge \pi_{(10\|N),1}(\pi_{(10\|N),0}^{-1}(C_* \oplus \Delta)) = T\|Z$ for arbitrary $Y \in \{0,1\}^{|T|}$ and $Z \in \{0,1\}^{n-|T|}$. The probability of the event $\pi_{(10\|N),0}^{-1}(C_* \oplus \Delta) \oplus \Delta = Y\|10^{n-|T|-1}$ is at most $\{0,1\}^{|T|}/(2^n - 1)$ because $\pi_{(10\|N),0}^{-1}$ has only been evaluated on a single point $C_*^i \oplus \Delta_i \neq C_* \oplus \Delta$ before. The probability of the event $\pi_{(10\|N),1}(Y\|10^{n-|T|-1}) = T\|Z$ conditioned on $\pi_{(10\|N),0}^{-1}(C_* \oplus \Delta) = Y\|10^{n-|T|-1}$ is then at most $\{0,1\}^{n-|T|}/(2^n - 1)$, because $\pi_{(10\|N),1}$ was also only evaluated on a single input different from $Y\|10^{n-|T|-1}$. Multiplying the two terms yields $2^n/(2^n - 1)^2$.

Because the two ways of forging in this case are not clearly mutually exclusive (both allow $T = T^i$), the probability of forgery is bounded by $\ell_{\max}/2^n + 2^n/(2^n - 1)^2$.

Case 4.2, $(A, C') = (A^i, C'^i)$: We must have $(C_*, T) \neq (C_*^i, T^i)$ for the forgery to be valid, and we also have $\Delta = \Delta^i$. This implies that $C_* \neq C_*^i$ (because forging with $C_* = C_*^i$ and $T \neq T^i$ is impossible), which in turn implies $C_* \oplus \Delta \neq C_*^i \oplus \Delta_i$. By a similar argument as in the second part of **Case 4.1**, the probability of forgery in this case is bounded by $2^n/(2^n - 1)^2$.

The probability of forgery in this case is thus no more than $\ell_{\max}/2^n + 2^n/(2^n - 1)^2$.

By taking the maximum over all cases, the probability that a single-decryption-query adversary \mathcal{A}' manages to forge $\ell_{\max} \cdot 2^{-n} + 2^n \cdot (2^n - 1)^{-2}$. The adversary \mathcal{A} making q_v decryption queries thus forges with probability bounded by $q_v \cdot \ell_{\max} \cdot 2^{-n} + q_v \cdot 2^n \cdot (2^n - 1)^{-2}$. By back-substituting all the previous inequalities, we obtain the claimed result. \square

5.8 Efficient Instantiation

Judging by the generic definitions of PAEF and SAEF in Figures 8 and 11, these two modes look anything but efficient. First, half of the output of every forkcipher call is being wasted when processing the associated data, and the decryption algorithms have to call both the inversion and the reconstruction algorithm of the forkcipher for every ciphertext block.

However, both PAEF and SAEF were designed with our proposed instance for tweakable forkcipher, the ForkAES, in mind. Addressing the two issues just raised, and referring to Section 3, with ForkAES

- The computational cost of "wasting" half of the output corresponds to 10 AES rounds; we evaluate the first five rounds to compute state from which the computation is forked, and then we only need to compute the five more rounds to compute the "right" output. Thus in the AD processing, PAEF and SAEF perform as any other rate-1 mode.

scheme	Cost of encryption in (# of AES rounds)/10									
	$a = 0$				$a = 1$				$a = 2$	
	$m=1$	$m=2$	$m=3$	$m=4$	$m=0$	$m=1$	$m=2$	$m=3$	$m=1$	$m=2$
GCM	2+2	3+3	4+4	5+5	1+2	2+3	3+4	4+5	2+4	3+5
CCM	4	6	8	10	3	5	7	9	6	8
OCB3	3	4	5	6	3	4	5	6	5	6
CLOC	3	5	7	9	2	4	6	8	5	7
Deoxys-I	2.8	4.2	5.6	7	2.8	4.2	5.6	7	5.6	7
KIASU [≠]	2	3	4	5	2	3	4	5	4	5
PAEF	1.5	3	4.5	6	1.5	2.5	4	5.5	3.5	5
SAEF	1.5	3	4.5	6	1.5	2.5	4	5.5	3.5	5
fGCM	1.5+1	2.5+2	3+3	5+4	1.5+2	1.5+2	2.5+3	3+4	1.5+3	2.5+4

Figure 17: The comparison of performance of GCM [33], CCM [44], OCB3 [32], CLOC [27], Deoxys-I [29], KIASU [28] and the AE schemes presented in Section 5. Here $a = |A|_n$ and $m = |M|_n$, and the cells of the form “ $b + g$ ” mean “ $b \cdot 10$ AES rounds and g multiplications in $\text{GF}(2^{128})$.” GCM, CCM, OCB3 and CLOC are assumed to be instantiated with AES. For schemes that compute a derived key that is the same for all queries (OCB and GCM), the complexity of this step is ignored.

- The computational cost of the inversion and the reconstruction algorithm can be factored when evaluated on *the same input*. We first compute the 5 inverse rounds that are common for both the reconstruction and the inversion algorithms to reach the forking state S , and then compute five more backwards round to compute the plaintext, and five other *forward* rounds to compute the other output half. The computational cost of this is exactly the same as the complete forward evaluation of ForkAES: 15 AE rounds.

In an optimized implementation, the actual interface of ForkAEs would be slightly different as well. The general rule for the implementations could be characterized as “implement only what you need.” For example for the PAEF and SAEF modes, one would have both the usual forward interface $\text{ForkAES}(\cdot, \cdot, \cdot)$ with $2n$ output bits, but also a dedicated interface $\text{ForkAES-r}(\cdot, \cdot, \cdot)$ that evaluates only the right branch of the forward ForkAES, for the AD processing. On the other hand, one would have a dedicated “composite” inverse interface $\text{ForkAES}(\cdot, \cdot, \cdot)$ that would take the key, the tweak, and *the left output block*, and directly output $2n$ bits—the plaintext block, and the reconstructed right output block. Thus we would factor the inversion and reconstruction, and dispense with the indicator bit input.

In this light, the performance of $\text{SAEF}[\text{ForkAES}]$ and $\text{PAEF}[\text{ForkAES}, \nu]$ looks much more promising, with the cost of both encryption and decryption algorithm being equivalent to about 1 AES evaluation per block of AD, and 1.5 AES evaluations per block of message.

PAEF and SAEF versus other modes. Of course, the factor of 1.5 is restrictive for long plaintexts, but SAEF and PAEF were designed to be efficient for *short messages*. To illustrate their efficiency, we compare the computational cost (in terms of the number of AES128 calls) of encrypting very short queries by PAEF and SAEF with that of GCM, CCM, OCB, CLOC and Deoxys-I in Figure 17.

Note also that in the plausible application scenarios (as those mentioned in Section 1), it is unlikely that a pipelined hardware accelerator for AES would be available. Therefore, the metric used in Table 17 yields an accurate comparison of actual performance of the mentioned AE modes.

5.9 Deterministic MiniAE

In the introduction, we stated that a forkcipher is nearly, but not exactly, an AE primitive: we clarify this statement in this section. It is easy to see that the syntax and security goals of a forkcipher, as proposed in Section 4, capture neither AE functionality nor AE security goals. Yet, *constructing* a secure PRI (with the same signature) from the forkcipher is trivial. In this section, we demonstrate that when used in a minimalistic “mode” of operation, a secure forkcipher yields a miniature AE scheme for fixed-size messages, which achieves PRI security [42].

PRI security of an AEAD scheme. Informally speaking, the best possible security that an AEAD scheme *with a fixed stretch* can achieve is to be (computationally) indistinguishable from a random injection from $\mathcal{N} \times \mathcal{A} \times \mathcal{M}$ to \mathcal{C} , because any AE scheme that is correct, must also be injective. This intuition is formalized as follows. The advantage of an adversary \mathcal{A} in distinguishing an AEAD scheme Π with ciphertext expansion τ from a random injection with the same signature is defined as

$$\text{Adv}_{\Pi}^{\text{pri}}(\mathcal{A}) = \Pr[\mathcal{A}^{\text{pri-real}_{\Pi}} \Rightarrow 1] - \Pr[\mathcal{A}^{\text{pri-ideal}_{\Pi}} \Rightarrow 1]$$

with the games **pri-real_Π** and **pri-ideal_Π** defined in Figure 18.

<pre> proc initialize $K \leftarrow \mathcal{K}$ proc Enc(N, A, M) return $\mathcal{E}(K, N, A, M)$ proc Dec(N, A, C) return $\mathcal{D}(K, N, A, C)$ </pre>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">pri-real_Π</div>	<pre> proc initialize for $N, A \in \mathcal{N} \times \mathcal{A}$ do $f_{N,A} \leftarrow \text{Inj}(\tau)$ proc Enc(N, A, M) return $f_{N,A}(M)$ proc Dec(N, A, C) if $\exists M \in \mathcal{M}$ s.t. $f_{N,A}(M) = C$ then return M else return \perp </pre>	<div style="border: 1px solid black; padding: 2px; display: inline-block;">pri-ideal_Π</div>
---	--	--	---

Figure 18: Pseudo-random injection (PRI) security games for a scheme $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ with ciphertext expansion τ .

Given a tweakable forkcipher F with $\mathcal{T} = \{0, 1\}^t$ and a $1 \leq \nu < t$, we define the AEAD scheme $\text{MAE}[F, \nu] = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ (as in “mini AE”) with $\mathcal{K} = \{0, 1\}^k$. The message space $\mathcal{M} = \{0, 1\}^n$ is given by the block-size of F , the nonce space $\mathcal{N} = \{0, 1\}^\nu$ and the AD space $\mathcal{A} = \{0, 1\}^\alpha$ with $\alpha = t - \nu$, so the parameter ν allows to make a trade-off between the nonce and AD sizes. The ciphertext expansion is n . The encryption and the decryption algorithm are defined in Figure 19.

The MAE mode captures the immediate intuition behind the “AE-potential” of a forkcipher: just use the redundancy contained in the right output block as a “tag”.

Security of MAE. We have the following statement about the security of MAE.

Theorem 5. *Let F be a tweakable forkcipher with $\mathcal{T} = \{0, 1\}^t$, let $1 \leq \nu < t$ and let $1 \leq \tau \leq n$. Then for adversary \mathcal{A} whose queries lie in the proper domains of the encryption and decryption algorithms and who makes q encryption queries and q_v decryption queries such that $q + q_v \leq 2^{n-1}$, we have*

$$\text{Adv}_{\text{MAE}[F, \nu, \tau]}^{\text{pri}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prtfp}}(\mathcal{B}) + \frac{(q + q_v)^2}{2^n}$$

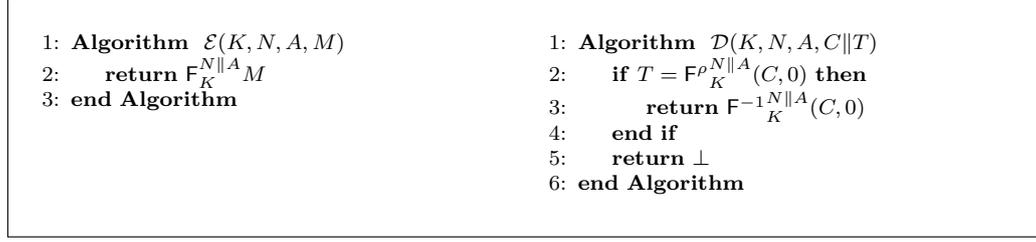


Figure 19: The MAE[F, ν] AEAD scheme.

for some adversary \mathcal{B} who makes at most twice as many queries in total as \mathcal{A} , and who runs in time given by the running time of \mathcal{A} plus an overhead that is linear in the total number \mathcal{A} 's queries.

Proof (sketch). We first replace F by a pair of tweakable random permutations $\pi_0 = (\pi_{\tau,0} \leftarrow \$ \text{Perm}(n))_{\tau \in \{0,1\}^t}$ and $\pi_1 = (\pi_{\tau,1} \leftarrow \$ \text{Perm}(n))_{\tau \in \{0,1\}^t}$. Letting MAE[π_0, π_1] denote the MAE mode that uses π_0, π_1 instead of F , we have

$$\text{Adv}_{\text{MAE}[F]}^{\text{pri}}(\mathcal{A}) \leq \text{Adv}_F^{\text{prtfp}}(\mathcal{B}) + \text{Adv}_{\text{MAE}[\pi_0, \pi_1]}^{\text{pri}}(\mathcal{A})$$

by a similar argument as in the proof of Theorem 2. In the rest of the analysis, we will refer to MAE[π_0, π_1] simply by Π .

For the rest of the analysis, we use the game Γ_0 and Γ_1 defined in Figure 20. We claim that $\Pr[\mathcal{A}^{\text{pri-real}\Pi} \rightarrow 1] = \Pr[\mathcal{A}^{\Gamma_1} \rightarrow 1]$ and that $\Pr[\mathcal{A}^{\text{pri-ideal}\Pi} \rightarrow 1] = \Pr[\mathcal{A}^{\Gamma_0} \rightarrow 1]$, which yields

$$\text{Adv}_{\text{MAE}[\pi_0, \pi_1]}^{\text{pri}}(\mathcal{A}) \leq \Pr[\mathcal{A}^{\Gamma_1} \rightarrow 1] - \Pr[\mathcal{A}^{\Gamma_0} \rightarrow 1].$$

It is easy to verify the latter equality; with the boxed statements removed, the code in Figure 20 implements a family of random injections indexed by (N, A) by lazy sampling. In particular, note that the probability that non-trivial decryption query succeeds in Γ_0 is $\Pr[b' = 1 \wedge b'' = 1] = 1 \cdot (2^n - |f_{N,A}|)^2 / (2^n - |f_{N,A}|) \cdot (2^{2n} - |f_{N,A}|)$ which is equal to the probability of finding a preimage of a random injection for which $|f_{N,A}|$ range points with known images (or known to have no preimages).

The former equality holds, because the framed lines in game Γ_1 make sure that f does in fact implement MAE based on a family of *pairs* of random permutations indexed by (N, A) . First, there is an implicit bijection between $\{0, 1\}^t$ and $\mathcal{N} \times \mathcal{A}$, so they are interchangeable. Then, the conditions of lines 7 and 7 make sure that the functions $\pi_{(N,A),0}$ and $\pi_{(N,A),1}$ defined by $\pi_{(N,A),0}(M) = \text{left}_n(f(m))$ and $\pi_{(N,A),1}(M) = \text{right}_n(f(m))$. The framed lines following the line 8 make sure that the distribution of ciphertext is the same as when produced by a pair of random permutations. The boxed statement after line 10 rejects ciphertexts that can never be produced by MAE[π_0, π_1]. The boxed statements after line 20 make sure that the probability that a non-trivial decryption query succeeds is the same as for MAE[π_0, π_1].

In addition, the games Γ_0 and Γ_1 are identical until **bad**, so we have $\Pr[\mathcal{A}^{\Gamma_1} \rightarrow 1] - \Pr[\mathcal{A}^{\Gamma_0} \rightarrow 1] \leq \Pr[\mathcal{A}^{\Gamma_1} \text{ sets } \mathbf{bad}]$ by the Fundamental lemma of gameplaying [16]. We define \mathbf{bad}^i for $i = 1, \dots, q + q_v$ to be the event **bad** is set to **true** in the i^{th} query made by the adversary. We further let \mathbf{bad}_1^i denote the event that \mathbf{bad}^i is true due to line 8, \mathbf{bad}_2^i denote the event that \mathbf{bad}^i is true due to line 10 and \mathbf{bad}_3^i denote the event that \mathbf{bad}^i is true due to line 20. Then we have that $\Pr[\mathcal{A}^{\Gamma_1} \text{ sets } \mathbf{bad}] \leq \sum_{i=1}^3 \sum_{j=1}^{q+q_v} \Pr[\mathbf{bad}_j^i]$.

We have that

$$\Pr[\mathbf{bad}_1^i] \leq (i-1) \cdot \frac{2 \cdot 2^n - 1}{2^{2n} - i + 1} \leq (i-1) \cdot \frac{2^{n+1}}{2^n(2^n - 1)} \leq \frac{2 \cdot (i-1)}{2^n - 1}$$

because if `bad` was not set previously, there are at most $i - 1$ elements in both $\mathcal{R}_l(f_{N,A})$ and $\mathcal{R}_r(f_{N,A})$ for any (N, A) , and for each X element of either $\mathcal{R}_l(f_{N,A})$ or $\mathcal{R}_r(f_{N,A})$, there are at most $2^n - 1$ elements of $\{0, 1\}^{2^n} \setminus \mathcal{R}_\perp(f_{N,A})$ that collide with X on their n leftmost, or respectively rightmost bits. The rest follows from the assumption $(q + q_v) \leq 2^n$ implied by $q + q_v \leq 2^{n-1}$. Summing over i , we get that $\sum_{i=1}^{q+q_v} \Pr[\text{bad}_1^i] \leq 2 \cdot (q + q_v)^2 / 2 \cdot (2^n - 1)$.

Then, we have that

$$\Pr[\text{bad}_2^i] \leq \frac{2^n}{2^{2n-i} - 1} \leq \frac{2^n}{2^n(2^n - 1)} \leq \frac{1}{2^n - 1}$$

because in the i^{th} query, we have $0 \leq |f_{N,A}| \leq i - 1$ for any (N, A) , and this determines the parameter of the Bernoulli variable which can set `bad`. The inequality follows using the assumption $(q + q_v) \leq 2^n$ implied by $(q + q_v) \leq 2^{n-1}$. Summing over i , we get that $\sum_{i=1}^{q+q_v} \Pr[\text{bad}_2^i] \leq (q + q_v) / (2^n - 1)$.

Finally, we have that

$$\Pr[\text{bad}_3^i] = \frac{1}{2^n - |f_{N,A}|} \cdot \left(1 - \frac{(2^n - |f_{N,A}|)^2}{2^{2n} - |f_{N,A}|}\right) \leq \frac{1}{2^n - |f_{N,A}|} \leq \frac{1}{2^n - (i - 1)} \leq \frac{1}{2^{n-1}}$$

because `badji` occurs in the i^{th} query if and only if $b' = b'' = 1$. The final inequality then follows from the assumption $(q + q_v) \leq 2^{n-1}$. Summing over i , we get that $\sum_{i=1}^{q+q_v} \Pr[\text{bad}_3^i] \leq (q + q_v) / (2^{n-1})$.

The claimed bound is obtained by adding up the sums $\sum_{j=1}^{q+q_v} \Pr[\text{bad}_j^i]$ for $j = 1, 2, 3$. \square

5.10 Discussion

We presented three modes of operation for a tweakable forkcipher. Each of the three modes reduces to a single call to the used forkcipher F in the case that the input only consists of a single block of data (either AD or message, but not both). This, together with an appropriate instantiation of F yields concrete AE schemes that excel in short-message encryption and decryption performance.

However, all three modes also share a certain deficiency: the technique we use to process arbitrary-length messages requires the ciphertext expansion to be exactly n bits. This can be limiting in many applications, in which bandwidth is limited. We leave solving this issue as an interesting open question.

Starting from a FIL PRI. As mentioned in Section 1, we also considered designing the VIL AEAD schemes as modes of operation of a FIL (tweakable) PRI, but we encountered several setbacks when heading in this direction.

Firstly, the FIL PRI does not seem suitable for designing modes that process *arbitrary* length messages. Out of the three modes we proposed, PAEF and SAEF cannot be recast as modes of a PRI, because they effectively discard one half of every call to the used primitive F . If F were a PRI, this would make decryption impossible. The parallel-permutation structure is crucial for these two modes. The fGCM, on the other hand, could be recast as a mode of a FIL PRI, as it does not require the computation of the inverse of F , except for the final message block. The problem here is that when used with a FIL PRI, fGCM could not have a constant ciphertext expansion. When instantiated with a n -to- $2n$ bit PRI, an encryption of a single-bit message would result in a ciphertext with $2n - 1$ bits of expansion; an attempt to truncate the ciphertext would render decryption impossible. Thus we were unable to find a mode of operation for a FIL-PRI that would process arbitrary length messages with *constant* stretch. Whether such a mode of operation for forkciphers exists is another interesting open question.

```

1: proc initialize
2:   for  $N, A \in \mathcal{N} \times \mathcal{A}$  do
3:      $f_{N,A} = \emptyset$ 
4:   end for
5:    $\text{bad} \leftarrow \text{false}$ 

1: proc Enc( $N, A, M$ )
2:   if  $\exists C$  s.t.  $(M, C) \in f_{N,A}$  then
3:     return  $C$ 
4:   end if
5:    $C \leftarrow \{0, 1\}^{2^n} \setminus \mathcal{R}_\perp(f_{N,A})$ 
6:    $C_l \leftarrow \text{left}_n(C); C_r \leftarrow \text{right}_n(C)$ 
7:   if  $C_l \in \mathcal{R}_l(f_{N,A})$  or  $C_r \in \mathcal{R}_r(f_{N,A})$ 
then
8:      $\text{bad} \leftarrow \text{true}$ 
9:     if  $C_l \in \mathcal{R}_l(f_{N,A})$  then
10:       $X \leftarrow \{0, 1\}^n \setminus \mathcal{R}_l(f_{N,A})$ 
11:       $C \leftarrow X \| C_r$ 
12:    end if
13:    if  $C_r \in \mathcal{R}_r(f_{N,A})$  then
14:       $X \leftarrow \{0, 1\}^n \setminus \mathcal{R}_r(f_{N,A})$ 
15:       $C \leftarrow C_l \| X$ 
16:    end if
17:  end if
18:   $f_{N,A} \leftarrow f_{N,A} \cup \{(M, C)\}$ 
19:  return  $C$ 

1: proc Dec( $N, A, C \| T$ )
2:   if  $\exists M$  s.t.  $(M, C) \in f_{N,A}$  then
3:     return  $M$ 

4:   end if
5:    $M \leftarrow \{0, 1\}^n \setminus \mathcal{D}(f_{N,A})$ 
6:    $C_l \leftarrow \text{left}_n(C); C_r \leftarrow \text{right}_n(C)$ 
7:   if  $C_l \in \mathcal{R}_l(f_{N,A})$  or  $C_r \in \mathcal{R}_r(f_{N,A})$ 
then
8:      $b \leftarrow \text{Be} \left( \frac{2^n - |f_{N,A}|}{2^{2^n} - |f_{N,A}|} \right)$ 
9:     if  $b = 1$  then
10:       $\text{bad} \leftarrow \text{true}$ 
11:      return  $\perp$ 
12:       $f_{N,A} \leftarrow f_{N,A} \cup \{(M, C)\}$ 
13:      return  $M$ 
14:    end if
15:    else
16:       $b' \leftarrow \text{Be} \left( \frac{1}{2^{2^n} - |f_{N,A}|} \right)$ 
17:       $b'' \leftarrow \text{Be} \left( \frac{(2^n - |f_{N,A}|)^2}{2^{2^n} - |f_{N,A}|} \right)$ 
18:      if  $b' = 1$  then
19:        if  $b'' = 0$  then
20:           $\text{bad} \leftarrow \text{true}$ 
21:           $f_{N,A} \leftarrow f_{N,A} \cup \{(M, C)\}$ 
22:          return  $M$ 
23:        else
24:           $f_{N,A} \leftarrow f_{N,A} \cup \{(M, C)\}$ 
25:          return  $M$ 
26:        end if
27:      end if
28:    end if
29:     $f_{N,A} \leftarrow f_{N,A} \cup \{(\perp, C)\}$ 
30:    return  $\perp$ 

```

Figure 20: The games Γ_0 and Γ_1 for bounding $\text{Adv}_{\text{MAE}[\pi_0, \pi_1]}^{\text{pri}}$. The game Γ_1 does *not* contain the boxed statements, while Γ_0 does. The games implement the (partially defined) injective functors $f_{N,A} : \{0, 1\}^n \rightarrow \{0, 1\}^{2^n}$ as initially-empty sets of preimage-image pairs; a pair (\perp, C) signifies that C has no preimage under the given function. We define the domain, range, and the “left” and “right” range of any $f_{N,A}$ as $\mathcal{D}(f_{N,A}) = \{M \in \{0, 1\}^n \mid \exists (M, C) \in f_{N,A}\}$, $\mathcal{R}(f_{N,A}) = \{C \in \{0, 1\}^{2^n} \mid \exists (M, C) \in f_{N,A} \text{ s.t. } M \neq \perp\}$, $\mathcal{R}_l(f_{N,A}) = \{L \in \{0, 1\}^n \mid \exists \text{ some } L \| X \in \mathcal{R}(f_{N,A})\}$ and $\mathcal{R}_r(f_{N,A}) = \{R \in \{0, 1\}^n \mid \exists \text{ some } X \| R \in \mathcal{R}(f_{N,A})\}$. We additionally define the extended range $\mathcal{R}_\perp(f_{N,A}) = \{C \in \{0, 1\}^{2^n} \mid \exists (M, C) \in f_{N,A}\}$. $\text{Be}(p)$ denotes a random variable with Bernoulli distribution with $\Pr[\text{Be}(p) = 1] = p$.

Constructing a FIL PRI. The second question is if we can find better instances of a FIL PRI. Our preferred choice was AES-based and thus the ForkAES. Yet, as evidenced by the result in Section 5.9, there is an unavoidable birthday-type quantitative gap between the PRI security, and the kind of security that ForkAES inherently possesses. A direct instance of a *true* FIL tweakable PRI is a question we leave open.

References

- [1] 3GPP TS 22.261: Service requirements for next generation new services and markets. <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3107>
- [2] 3GPP TS 36.213: Evolved Universal Terrestrial Radio Access (E-UTRA); Physical layer procedures. <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=2427>
- [3] CAN FD Standards and Recommendations. <https://www.can-cia.org/news/cia-in-action/view/can-fd-standards-and-recommendations/2016/9/30/>
- [4] ISO 11898-1:2015: Road vehicles – Controller area network (CAN) – Part 1: Data link layer and physical signalling. <https://www.iso.org/standard/63648.html>
- [5] NB-IoT: Enabling New Business Opportunities. http://www.huawei.com/minisite/iot/img/nb_iot_whitepaper_en.pdf
- [6] Specification of Secure Onboard Communication. https://www.autosar.org/fileadmin/user_upload/standards/classic/4-3/AUTOSAR_SWS_SecureOnboardCommunication.pdf
- [7] Andreeva, E., Neven, G., Preneel, B., Shrimpton, T.: Seven-Property-Preserving Iterated Hashing: ROX. In: Kurosawa, K. (ed.) ASIACRYPT 2007. LNCS, vol. 4833, pp. 130–146. Springer (2007)
- [8] Aumasson, J.P., Babbage, S., Bernstein, D.J., Cid, C., Daemen, J., Gaj, O.D.K., Gueron, S., Junod, P., Langley, A., McGrew, D., Paterson, K., Preneel, B., Rechberger, C., Rijmen, V., Robshaw, M., Sarkar, P., Schaumont, P., Shamir, A., Verbauwhede, I.: CHAE: Challenges in Authenticated Encryption. ECRYPT-CSA D1.1, Revision 1.05, 1 March 2017
- [9] Bellare, M.: Practice-oriented provable-security. In: Okamoto, E., Davida, G.I., Mambo, M. (eds.) ISW '97, Tatsunokuchi. LNCS, vol. 1396, pp. 221–231. Springer (1998)
- [10] Bellare, M., Goldreich, O., Mityagin, A.: The Power of Verification Queries in Message Authentication and Authenticated Encryption. IACR Cryptology ePrint Archive 2004, 309 (2004)
- [11] Bellare, M., Kohno, T., Namprempre, C.: Breaking and provably repairing the SSH authenticated encryption scheme: A case study of the encode-then-encrypt-and-mac paradigm. ACM Trans. Inf. Syst. Secur. 7(2), 206–241 (2004)
- [12] Bellare, M., Namprempre, C.: Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer (2000)

-
- [13] Bellare, M., Namprempe, C.: Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. *J. Cryptology* 21(4), 469–491 (2008)
- [14] Bellare, M., Ristenpart, T.: Multi-property-preserving hash domain extension and the EMD transform. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 299–314. Springer (2006)
- [15] Bellare, M., Rogaway, P.: Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 317–330. Springer (2000)
- [16] Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings. Lecture Notes in Computer Science, vol. 4004, pp. 409–426. Springer (2006), https://doi.org/10.1007/11761679_25
- [17] Bernstein, D.J.: Cryptographic competitions: CAESAR. <http://competitions.cr.yp.to>
- [18] Biryukov, A., Nikolic, I.: Automatic search for related-key differential characteristics in byte-oriented block ciphers: Application to aes, camellia, khazad and others. In: Gilbert, H. (ed.) Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6110, pp. 322–344. Springer (2010), https://doi.org/10.1007/978-3-642-13190-5_17
- [19] Biryukov, A., Nikolic, I.: Search for related-key differential characteristics in DES-Like ciphers. In: Joux, A. (ed.) Fast Software Encryption - 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers. Lecture Notes in Computer Science, vol. 6733, pp. 18–34. Springer (2011), https://doi.org/10.1007/978-3-642-21702-9_2
- [20] Demirci, H., Selçuk, A.A.: A meet-in-the-middle attack on 8-round AES. In: Nyberg, K. (ed.) Fast Software Encryption, 15th International Workshop, FSE 2008, Lausanne, Switzerland, February 10-13, 2008, Revised Selected Papers. Lecture Notes in Computer Science, vol. 5086, pp. 116–126. Springer (2008), https://doi.org/10.1007/978-3-540-71039-4_7
- [21] Derbez, P., Fouque, P., Jean, J.: Improved key recovery attacks on reduced-round AES in the single-key setting. In: Johansson, T., Nguyen, P.Q. (eds.) Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7881, pp. 371–387. Springer (2013), https://doi.org/10.1007/978-3-642-38348-9_23
- [22] Dunkelman, O., Keller, N., Shamir, A.: Improved single-key attacks on 8-round AES-192 and AES-256. In: Abe, M. (ed.) Advances in Cryptology - ASIACRYPT 2010 - 16th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 5-9, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6477, pp. 158–176. Springer (2010), https://doi.org/10.1007/978-3-642-17373-8_10

-
- [23] Dworkin, M.: Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D (Nov 2007)
- [24] Emami, S., Ling, S., Nikolić, I., Pieprzyk, J., Wang, H.: The resistance of present-80 against related-key differential attacks. *Cryptography Commun.* 6(3), 171–187 (Sep 2014), <http://dx.doi.org/10.1007/s12095-013-0096-8>
- [25] Fouque, P., Jean, J., Peyrin, T.: Structural evaluation of AES and chosen-key distinguisher of 9-round AES-128. In: Canetti, R., Garay, J.A. (eds.) *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I. Lecture Notes in Computer Science*, vol. 8042, pp. 183–203. Springer (2013), https://doi.org/10.1007/978-3-642-40041-4_11
- [26] Hoang, V.T., Krovetz, T., Rogaway, P.: Robust Authenticated-Encryption AEZ and the Problem That It Solves. In: Oswald, E., Fischlin, M. (eds.) *EUROCRYPT 2015. LNCS*, vol. 9056, pp. 15–44. Springer (2015)
- [27] Iwata, T., Minematsu, K., Guo, J., Morioka, S.: CLOC: authenticated encryption for short input. In: Cid, C., Rechberger, C. (eds.) *Fast Software Encryption 2014. LNCS*, vol. 8540, pp. 149–167. Springer (2015)
- [28] J. Jean, I. Nikolić, T.P.: KIASU v1 (2014), "<https://competitions.cr.yp.to/round1/kiasuv1.pdf>"
- [29] J. Jean, I. Nikolić, T.P.Y.S.: Deoxys v1.41 v1 (2016), "<https://competitions.cr.yp.to/round3/deoxysv141.pdf>"
- [30] Jean, J., Nikolic, I., Peyrin, T.: Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In: Sarkar, P., Iwata, T. (eds.) *ASIACRYPT 2014, Part II. LNCS*, vol. 8874, pp. 274–288. Springer (2014)
- [31] Katz, J., Yung, M.: Unforgeable encryption and chosen ciphertext secure modes of operation. In: Schneier, B. (ed.) *FSE 2000. LNCS*, vol. 1978, pp. 284–299. Springer (2001)
- [32] Krovetz, T., Rogaway, P.: The Software Performance of Authenticated-Encryption Modes. In: Joux, A. (ed.) *FSE 2011. LNCS*, vol. 6733, pp. 306–327. Springer (2011)
- [33] McGrew, D.A., Viega, J.: The security and performance of the galois/counter mode (GCM) of operation. In: Canteaut, A., Viswanathan, K. (eds.) *Progress in Cryptology - INDOCRYPT 2004, 5th International Conference on Cryptology in India, Chennai, India, December 20-22, 2004, Proceedings. Lecture Notes in Computer Science*, vol. 3348, pp. 343–355. Springer (2004), https://doi.org/10.1007/978-3-540-30556-9_27
- [34] Mouha, N., Wang, Q., Gu, D., Preneel, B.: Differential and linear cryptanalysis using mixed-integer linear programming. In: Wu, C., Yung, M., Lin, D. (eds.) *Information Security and Cryptology - 7th International Conference, Inscrypt 2011, Beijing, China, November 30 - December 3, 2011. Revised Selected Papers. Lecture Notes in Computer Science*, vol. 7537, pp. 57–76. Springer (2011), https://doi.org/10.1007/978-3-642-34704-7_5
- [35] Namprempe, C., Rogaway, P., Shrimpton, T.: Reconsidering Generic Composition. In: Nguyen, P.Q., Oswald, E. (eds.) *EUROCRYPT 2014. Lecture Notes in Computer Science*, vol. 8441, pp. 257–274. Springer (2014)
- [36] NIST: DRAFT Submission Requirements and Evaluation Criteria for the Lightweight Cryptography Standardization Process. <https://csrc.nist.gov/Projects/Lightweight-Cryptography> (2018)

-
- [37] Paterson, K.G., Yau, A.K.L.: Cryptography in theory and practice: The case of encryption in ipsec. IACR Cryptology ePrint Archive 2005, 416 (2005), <http://eprint.iacr.org/2005/416>
- [38] Paterson, K.G., Yau, A.K.L.: Cryptography in Theory and Practice: The Case of Encryption in IPsec. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 12–29. Springer (2006)
- [39] Reyhanitabar, M.R., Susilo, W., Mu, Y.: Analysis of Property-Preservation Capabilities of the ROX and ESh Hash Domain Extenders. In: Boyd, C., Nieto, J.M.G. (eds.) ACISP 2009. LNCS, vol. 5594, pp. 153–170. Springer (2009)
- [40] Rogaway, P.: Authenticated-Encryption with Associated-Data. In: ACM CCS 2002. pp. 98–107 (2002)
- [41] Rogaway, P.: Practice-oriented provable security and the social construction of cryptography. IEEE Security & Privacy 14(6), 10–17 (2016)
- [42] Rogaway, P., Shrimpton, T.: A Provable-Security Treatment of the Key-Wrap Problem. In: EUROCRYPT 2006. pp. 373–390 (2006)
- [43] Siwei Sun, Lei Hu, P.W.K.Q.X.M.L.S.: Automatic security evaluation and (related-key) differential characteristic search: Application to simon, present, lblock, des(1) and other bit-oriented block ciphers. Cryptology ePrint Archive, Report 2013/676 (2013), <https://eprint.iacr.org/2013/676>
- [44] Whiting, D., Housley, R., Ferguson, N.: Counter with CBC-MAC (CCM). IETF RFC 3610 (Informational) (Sep 2003), <http://www.ietf.org/rfc/rfc3610.txt>