

Counter-in-Tweak: Authenticated Encryption Modes for Tweakable Block Ciphers^{*}

Thomas Peyrin and Yannick Seurin

SPMS, NTU, Singapore
thomas.peyrin@ntu.edu.sg

ANSSI, Paris, France
yannick.seurin@m4x.org

May 27, 2016

Abstract. We propose the Synthetic Counter-in-Tweak (SCT) mode, which turns a tweakable block cipher into a nonce-based authenticated encryption scheme (with associated data). The SCT mode combines in a SIV-like manner a Wegman-Carter MAC inspired from PMAC for the authentication part and a new counter-like mode for the encryption part, with the unusual property that the counter is applied on the tweak input of the underlying tweakable block cipher rather than on the plaintext input. Unlike many previous authenticated encryption modes, SCT enjoys provable security beyond the birthday bound (and even up to roughly 2^n tweakable block cipher calls, where n is the block length, when the tweak length is sufficiently large) in the nonce-respecting scenario where nonces are never repeated. In addition, SCT ensures security up to the birthday bound even when nonces are reused, in the strong nonce-misuse resistance sense (MRAE) of Rogaway and Shrimpton (EUROCRYPT 2006). To the best of our knowledge, this is the first authenticated encryption mode that provides at the same time close-to-optimal security in the nonce-respecting scenario and birthday-bound security for the nonce-misuse scenario. While two passes are necessary to achieve MRAE-security, our mode enjoys a number of desirable features: it is simple, parallelizable, it requires the encryption direction only, it is particularly efficient for small messages compared to other nonce-misuse resistant schemes (no precomputation is required) and it allows incremental update of associated data.

Keywords: authenticated encryption, tweakable block cipher, nonce-misuse resistance, beyond-birthday-bound security, CAESAR competition

^{*} © IACR 2016. This is the full version of the article submitted by the authors to the IACR and to Springer-Verlag in May 2016, which appears in the proceedings of CRYPTO 2016.

1 Introduction

BACKGROUND ON AUTHENTICATED ENCRYPTION. Confidentiality and authenticity of data are the two main security properties that one must ensure when communicating over an insecure channel. In the symmetric key setting, it has long been known how to ensure both of them independently, e.g., starting from a secure block cipher, by using a suitable encryption mode for confidentiality [BJR97] and a block cipher-based MAC for authenticity [BKR00]. However, how exactly to combine both tools has long been left to the practitioners, leading to major security breaches [Kra01, DR11, AP13]. Sometimes, protocol designers even overlooked that authenticity was a necessary requirement besides confidentiality, as exemplified by padding oracle attacks [Vau02]. Even when the combination of the encryption and the MAC schemes is properly done, it might not be the most efficient solution, especially when the two parts rely on two different primitives. For these reasons, interest has shifted towards designing “integrated” Authenticated Encryption (AE) schemes ensuring jointly authenticity and confidentiality of data, which are more efficient and less likely to be incorrectly used. Besides, it has become standard for an AE scheme to have the ability to handle so-called associated data (AD), which are authenticated but not encrypted [Rog02] (such a scheme was for a time called an AEAD scheme, but since this feature is so important in practice, virtually all modern AE schemes provide it; we will only talk of AE in this paper, implicitly meaning AEAD). Even though ad-hoc AE schemes were already used since a long time, the formal treatment of these constructions only started around 2000 [BN00, BR00, KY06]. At about the same time, provably secure AE designs started to appear, such as IAPM [Jut01, Jut08], XCBC [GD01], CCM [WHF02], OCB [RBB03, Rog04a], or GCM [MV04]. The CAESAR competition [CAE] for authenticated encryption, started in 2014, recently put this research topic in the limelight. Various AE schemes were proposed, from purely ad-hoc designs to (tweakable) block cipher operating modes.

NONCE-MISUSE RESISTANCE. Since most symmetric-key primitives (in particular block ciphers) from which AE schemes are built are deterministic, a random IV or a nonce (i.e., a value which must never be repeated for the same secret key) is a necessary ingredient for achieving strong security goals. Failing to ensure the corresponding requirement (high entropy for an IV, non-repetition for a nonce) can have dramatic consequences for security. For example, reusing a nonce just a single time for encrypting two messages in OCB completely breaks confidentiality: an attacker can immediately detect repeated message blocks since the corresponding ciphertext blocks will be equal. The non-repeating requirement on the nonce can be challenging to fulfill in some contexts, for example when encryption is implemented in a stateless device.¹ It is likely (and it has happened before) that some implementations will, e.g., simply generate nonces at random,

¹ Similarly, the high-entropy requirement on the IV is hard to meet when no good randomness source is available.

“hoping” that no collision will occur. For that reason, a recent trend in AE has been to design schemes achieving *nonce-misuse resistance*, which informally means that the impact on security of a nonce repetition should be as limited as possible. This goal was first put forward by Rogaway and Shrimpton [RS06], who formalized the notion of *misuse-resistant AE* (MRAE). For a scheme enjoying this property, authenticity is not harmed by nonce repetitions, while confidentiality is only damaged insofar as the adversary can detect whether the same triple of nonce, AD and message values is repeated. Example of schemes achieving this security notion are EAX [BRW04], SIV [RS06], AEZ [HKR15], or GCM-SIV [GL15]. Because the MRAE notion cannot be achieved for an online scheme (since each bit of the ciphertext must depend on every bit of the plaintext), Fleischmann *et al.* [FFL12] proposed a relaxation of the MRAE notion called *online AE* (OAE), which can be achieved with a single pass on the input. Examples of schemes ensuring this security property are McOE [FFL12], COPA [ABL⁺13], or POET [AFF⁺14]. However, the interest in the OAE notion has been recently reduced by some serious security concerns, notably the so-called chosen-prefix/secret-suffix (CPSS) generic attack [HRRV15], that shares some similarities with the BEAST attack [DR11].

BIRTHDAY AND BEYOND-BIRTHDAY SECURITY. Another important shortcoming of most AE operating modes is that they provide only birthday-bound security with respect to the block length of the underlying primitive. Since virtually all existing block ciphers have block length at most 128 bits (in particular the current block cipher standard AES), this means that security is lost at 2^{64} block cipher calls at best, which is low given modern security requirements (for 64-bit block ciphers, the situation is even more problematic). Moreover, this is rarely a problem with the tightness of the security proof: attacks matching the bound are often known. For example, Ferguson [Fer02] described a simple collision-based attack on OCB that breaks authenticity with 2^{64} blocks of messages. Recently, some AE schemes providing security beyond the birthday bound (BBB) were proposed [Iwa06, Iwa08], but they usually come at an expensive performance price. One could argue that using a double-block-length block cipher would provide the expected security, but this solution comes with an important efficiency penalty (as can be seen in generic double-block-length block cipher constructions) and would be highly problematic for hardware implementations where internal state size is a major contribution to the total area cost.

AE FROM TWEAKABLE BLOCK CIPHERS. Compared with a conventional block cipher, a tweakable block cipher (TBC) \tilde{E} takes an additional input called a *tweak* bringing inherent variability to the primitive (equivalently, a TBC can be seen as a family of block ciphers indexed by the tweak). In the same paper that formalized the corresponding security notion [LRW02], it was pointed out that a TBC was a very convenient starting point for building various schemes. In particular, for AE schemes, two prominent examples are the sibling modes TAE [LRW02] and Θ CB [KR11] (the TBC-based generalization of OCB). They have “perfect” security in the sense that, when used with an ideal TBC, the advantage of any adversary is zero against confidentiality and close to $2^{-\tau}$ against

authenticity, where τ is the tag length. However, as already pointed out, a weakness of both TAE and Θ CB (even when used with an ideal TBC) is that their security completely collapses as soon as a nonce is repeated. As a matter of fact, existing AE schemes built from an ideal TBC either ensure perfect security in the nonce-respecting scenario only (like TAE or Θ CB), or fulfill the weak OAE notion only (e.g. COPA, once recast to use an ideal TBC), or ensure MRAE-security but only up to the birthday bound, even if nonces are not repeated (like AEZ). The PIV construction by Shrimpton and Terashima [ST13] allows to construct a variable-input-length TBC with BBB-security, which in turn allows to construct (via the Encode-then-Encipher method) an AEAD scheme with BBB-security against nonce-respecting adversaries and birthday-bound security against nonce-misusing ones. However, PIV requires as a building block a fixed-input-length TBC with *variable* tweak length (comparable to the maximal input length of the PIV construction), which in turn requires to appeal to universal hash functions with key length comparable to the maximal tweak length. Hence, the resulting AEAD scheme must use very large keys to ensure BBB-security for large messages. As of today, there is no AEAD scheme based on a fixed-tweak-length TBC that ensures both BBB-security in the nonce-respecting scenario and (at least) birthday-bound security in the nonce-misuse scenario. Yet this seems a very desirable goal since such a scheme would at the same time yield very high (BBB) security guarantees in the nominal, nonce-respecting use case and retain acceptable (birthday-bound) security when inadvertently misused.

OUR CONTRIBUTIONS. In this paper, we propose the SCT (Synthetic Counter-in-Tweak) nonce-based AE mode for tweakable block ciphers and prove that it ensures BBB-security in the nonce-respecting scenario, and birthday-bound security in the nonce-misuse scenario (in the strong MRAE sense [RS06]). More precisely, for the nonce-respecting case, when using a ideal TBC with block length n and “effective” tweak length² w , SCT is secure up to roughly 2^n TBC calls when $w \geq n$, and up to roughly $2^{(n+w)/2}$ TBC calls when $w \leq n$, which is always larger than $2^{n/2}$. The SCT mode requires two passes (as is inevitable for MRAE-security), but it is simple, parallelizable, it requires the encryption direction only, it is particularly efficient for small messages compared to other nonce-misuse resistant schemes (no precomputation is required) and it allows incremental update of associated data.

With respect to how authentication and encryption are combined, our design draws inspiration from the SIV generic composition method [RS06]: the nonce N , the associated data A , and the message M are first input to a keyed function F_K , yielding an pseudorandom initial value IV , which will serve as authentication tag. The message is then encrypted, using the generated IV and the nonce N (see Fig. 4). This “recycling” of the nonce in the encryption part of the mode is

² The SCT mode uses 5 tweak prefixes to separate the different usages of the TBC. The “effective” tweak length is what remains once 3 bits have been used to encode the prefix.

what makes our high-level construction (called NSIV) crucially different from SIV and allows to reach BBB-security in the nonce-respecting case.³

It remains to instantiate the two components of the NSIV construction, the keyed function F_K and the encryption scheme. Since we aim at BBB-security in the nonce-respecting case, a natural starting point for F_K is the Wegman-Carter paradigm [WC81, Bra82, Sho96]. Hence, we propose a nonce-based MAC mode called PWC (*Parallel Wegman-Carter*) which combines a xor-universal hash function inspired from PMAC [BR02, Rog04a] applied to the AD and the message, and a simple pseudorandom function applied to the nonce. In order to achieve nonce-misuse resistance (which in general Wegman-Carter MACs do not provide), we add an additional encryption layer, which results in the EPWC (*Encrypted PWC*) mode.

The real challenge lies in designing an encryption scheme which is BBB-secure in the nonce-respecting case. Since on one hand it seems hard to leverage on the non-repeating property of the nonce without actually giving the nonce as input to the encryption mode, and on the other hand we need to make use in some way of the pseudorandom IV computed from F_K ,⁴ it appears that what we need to design is actually a *combined nonce- and IV-based encryption scheme* (*nivE scheme* for short). To the best of our knowledge, this notion has never appeared before, and we introduce it in this paper. The encryption mode that we propose, called CTRT (*CounTeR in Tweak*), is a counter-like mode with the unusual property that the counter is applied on the tweak input of the underlying TBC rather than on the plaintext input, where the nonce comes in. The combination of EPWC and CTRT through the NSIV construction (the IV generated by EPWC being used as initial counter in CTRT) yields the SCT mode, illustrated in Fig. 1.

For completeness, we also describe the CTPWC (CTR-then-PWC) mode, an *online* nonce-based AE scheme which combines in an “encrypt-then-MAC” manner a slight variant of the CTRT encryption mode and the PWC authentication mode. The security guarantees provided by CTPWC are similar to those of Θ CB, but it is roughly twice less efficient, so that we do not claim that it is of particularly high interest. One small advantage that we see for this mode compared with Θ CB is that the nonce length can be as large as the block length of the underlying TBC, whereas for Θ CB the sum of the nonce length and of the maximal length of encrypted messages must be less than the tweak length of the underlying tweakable block cipher, which might be restrictive in some settings (e.g., for KIASU-BC [JNP14c]). It might also escape the patent issues which hindered the adoption of OCB.

INSTANTIATING THE TBC. As just discussed, our new AE modes offer BBB-security (in the nonce-respecting case) when used with an ideal TBC. If one aims at leveraging this security level in the real world, one must instantiate

³ While SIV corresponds to generic composition method A4 in the nomenclature of Namprempre *et al.* [NRS14], NSIV does not fit any of the NRS schemes.

⁴ This excludes for example a simple OCB-like encryption mode since it is only nonce-based, not IV-based.




Fig. 1. The SCT mode, using a TBC \tilde{E} with tweak space $\{1, \dots, 5\} \times \mathcal{T}$ and domain $\mathcal{X} = \{0, 1\}^n$. For each call to \tilde{E}_K , the tweak enters left while the plaintext enters on top. We denote $\tilde{E}_K^i(T, X)$ for $\tilde{E}_K((i, T), X)$ and $\tilde{E}_K^{i/j}$ means that prefix i is used when the input block is complete and unpadded, whereas j is used when the input block is incomplete and padded. Function Inc is a cyclic permutation of \mathcal{T} , and Conv is a regular function from \mathcal{X} to \mathcal{T} (e.g., truncation when $\mathcal{X} = \{0, 1\}^n$ and $\mathcal{T} = \{0, 1\}^w$, $w \leq n$).

the TBC with care. Most existing TBCs are built from conventional block ciphers in a generic way, the prominent example being the XE/XEX construction [Rog04a] which only ensures security up to the birthday bound. Hence, using XE/XEX in our schemes would in a sense waste their nice security promises.⁵ To remedy this problem, one can use either generic TBC constructions with BBB-security [Min09, LST12, LS13, Men15] (but they are often inefficient or provably secure in the ideal cipher model only), or ad-hoc TBC designs without known weaknesses. The second option was chosen for a number of CAESAR candidates [JNP14a, JNP14b, JNP14c, GLS⁺14]. In fact, the SCT mode was explicitly designed as a replacement to the COPA mode used in versions 1.1 and 1.2 of CAESAR candidates Deoxys [JNP14a] (128-bit blocks, 128-bit tweaks) and Joltik [JNP14b] (64-bit blocks, 64-bit tweaks).⁶ We refer to the submission documents of these two candidates for a detailed report on implementation results, which are quite competitive. Other potential candidates for instantiating the SCT mode are Scream [GLS⁺14] and Threefish, the TBC on which the hash function Skein [FLS⁺10] is based. There is currently a shift towards designing dedicated TBCs achieving higher security and efficiency than generic BC-based constructions, and we hope to see more and more TBC proposals that could be used with SCT.

OPEN PROBLEMS AND FUTURE WORK. The CTRT encryption scheme has the notable feature that its security degrades gracefully with the maximal number of repetitions of nonces: when the nonce repetitions are limited, security remains close to the security bound in the nonce-respecting case. In contrast, the security of the EPWC authentication mode (and more generally of any encrypted Wegman-Carter MAC) falls back to birthday bound as soon as the adversary can repeat one single nonce twice (see Remark 2 in Section 5). It remains a pending question to modify EPWC so that it provides graceful security degradation with the maximal number of nonce repetitions as well. This would make the resulting AE scheme a good candidate for high security in both nonce-respecting and nonce-misuse models for most practical scenarios. Another challenging open problem would be to construct an AE scheme which remains BBB-secure even when nonces are arbitrarily repeated. The main difficulty is to build a deterministic, stateless, BBB-secure MAC, which is known to be notably hard [Yas11, DS11]. Another possible direction for future work would be to design a mode similar to SCT using only one pass and achieving online nonce-misuse resistance in the OAE sense [FFL12]. Such a feature would allow users to smoothly choose the best security achievable, depending on whether two passes can be tolerated or not by the application. Finally, it would be interesting to analyze how to strengthen SCT against other misuse scenarios such as release of unverified plaintext [ABL⁺14], and to study how its security is affected by tag truncation [HKR15].

⁵ Similarly, the only reason why OCB is secure up to the birthday bound whereas Θ CB is “perfectly” secure is because it relies on XE/XEX for instantiating the TBC.

⁶ The tweak prefixes used in this paper were chosen for ease of exposition and are slightly different from the ones used in Deoxys and Joltik v1.3, which were chosen mainly for efficiency reasons.

ORGANIZATION. In Section 2 we provide a high-level description of the various possibilities that we considered for constructing a BBB-secure nivE scheme. After introducing the notation and standard security notions in Section 3, we describe the CTRT encryption scheme and prove its security in Section 4, while we describe the PWC and EPWC nonce-based MAC schemes and prove their security in Section 5. Finally, we explain how to combine CTRT and EPWC using the NSIV construction to build the nonce-based AE mode SCT and prove its security in Section 6. In Appendix D, we describe the CTPWC mode, an online nonce-based AE mode derived from CTRT and PWC.

2 Counter-in-Tweak for Beyond-Birthday Security

As motivated in introduction, our goal is to design a simple TBC-based AE scheme that provides BBB-security in the nonce-respecting setting and (at least) birthday-bound MRAE-security. As already mentioned, an encrypted Wegman-Carter MAC solves the problem for the authentication part, so that we focus here on encryption. Hence, our problem is as follows: given a nonce and a synthetic IV generated pseudorandomly from the nonce, the AD and the message, how do we use them to encrypt the message with BBB-security? We give a quick overview of the various constructions that we considered and why, except the CTRT encryption mode we propose, they fail or are unsatisfactory.

A natural direction to explore is to start from a scheme providing BBB-security for non-repeating nonces, such as TAE [LRW02] or Θ CB [KR11], which are similar with regard to encryption: it simply consists in a “tweakable” codebook mode, the tweak holding the nonce and a message block counter. This is obviously not nonce-misuse resistant: repeating the nonce a single time will lead to a complete break of confidentiality since a constant message block leads to a constant ciphertext block. One could incorporate the IV by simply concatenating it to the nonce and the counter to form the tweak. However, this would require a TBC with larger tweak, which is usually very costly to achieve.⁷ Rather than concatenating the nonce and the IV, one could try to combine them into a single shorter string S , but this would presumably result in birthday security even in the nonce-respecting scenario (since a collision on S would directly break confidentiality). Hence, a codebook encryption mode does not seem to be a very convenient starting point.

For this reason, we preferred to consider a counter mode (note that this was the encryption mode favored by Rogaway and Shrimpton to instantiate the SIV composition method [RS06]). The question now is: how do we feed the nonce, the IV, and the i -th counter to the TBC in order to create the mask that will be xored to the i -th message block? We considered several possibilities (we do not claim this to be exhaustive):

- (a) One can put the nonce in the tweak input, and the sum of the IV and the counter in the plaintext input. The problem is that confidentiality caps at

⁷ For example, for TBCs following the TWEAKEY approach [JNP14a, JNP14b, JNP14d], there is a large gap in the number of rounds needed to make the TBC secure as the tweak length increases.

birthday bound even in the nonce-respecting scenario: the adversary can query the encryption of a single message with $2^{n/2}$ equal blocks, and observe that no collision occurs in the corresponding ciphertext blocks (since the nonce is fixed and all TBC calls use the same tweak), which will distinguish the ciphertext from a random string (for which a collision would be expected).

- (b) One can concatenate the nonce and the counter in the tweak input, and use the IV for the plaintext input. Since the tweak is different for each message block position, this solves the issue of the previous solution. We conjecture that this mode meets our security objectives, but an important drawback is that a larger tweak length is required and, as mentioned before, this is very costly.
- (c) One can put the sum of the nonce and the counter in the tweak input (instead of concatenating them) and the IV in the plaintext input. This mode might meet our security objectives, however the adversary can very easily provoke collisions on tweak inputs even in the nonce-respecting scenario, which might complicate the proof of BBB-security. Another drawback is that in the nonce-misuse scenario, a collision on the IV immediately breaks confidentiality, which dashes any hope for BBB nonce-misuse resistance. Note that one could imagine variants where the nonce and the IV are first encrypted before being used, but it is not clear if this would prevent the issues just mentioned and this would presumably make the security proof quite complex.
- (d) Finally, one can put the sum of the IV and the counter in the tweak input and the nonce in the plaintext input, which is exactly the CTRT mode. We will prove in Section 4 that it meets our security goal. The first idea is that the counter in the tweak input ensures that all the calls to the internal TBC will use different tweaks for one single message query, so that the ciphertext looks uniformly random in that case. Thus, the adversary has to query several messages with different nonce values and hope that many collisions will occur between tweak inputs in order to observe a divergence from uniformity in the ciphertexts. However, these collisions are hard to control since they depend on the pseudorandom IV (in contrast with other modes discussed above, where the tweak input can be easily controlled by the adversary). We will show in Section 4, using a “balls-into-bins” analysis, that the number of tweak collisions remains small, so that the distribution of the ciphertexts remains close to uniform. Moreover, the nonce-misuse scenario helps the adversary only if it can repeat the same nonce a very high number of times until a collision happens on the tweaks, so that the security of the CTRT mode degrades gracefully with the maximal number of nonce repetitions.

3 Preliminaries

NOTATION. Given a string $X \in \{0, 1\}^*$, $|X|$ denotes its length. If X and Y are respectively n -bit and m -bit strings, $n < m$, then $X \oplus Y$ denotes the n -bit string obtained by xoring X with the n leftmost bits of Y . Given some implicit length n and a bit-string X of length $1 \leq |X| < n$, we denote $X10^*$ the string obtained by

appending a single 1 and $(n - |X| - 1)$ 0's to X . Given two sets \mathcal{X} and \mathcal{Y} , the set of all functions from \mathcal{X} to \mathcal{Y} is denoted $\text{Func}(\mathcal{X}, \mathcal{Y})$. A function $F \in \text{Func}(\mathcal{X}, \mathcal{Y})$ is said *regular* if all $Y \in \mathcal{Y}$ have the same number of preimages by F (this obviously requires $|\mathcal{X}|$ to be a multiple of $|\mathcal{Y}|$).

TWEAKABLE BLOCK CIPHERS. A *tweakable block cipher* (TBC) with key space \mathcal{K} , tweak space \mathcal{T} , and domain \mathcal{X} is a mapping $\tilde{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X}$ such that for any key $K \in \mathcal{K}$ and any tweak $T \in \mathcal{T}$, $X \mapsto \tilde{E}(K, T, X)$ is a permutation of \mathcal{X} . We often write $\tilde{E}_K(T, X)$ or $\tilde{E}_K^T(X)$ in place of $\tilde{E}(K, T, X)$. We denote $\text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{X})$ the set of all tweakable block ciphers with key space \mathcal{K} , tweak space \mathcal{T} , and domain \mathcal{X} . A *tweakable permutation* with tweak space \mathcal{T} and domain \mathcal{X} is a mapping $\tilde{P} : \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{X}$ such that for any tweak $T \in \mathcal{T}$, $X \mapsto \tilde{P}(T, X)$ is a permutation of \mathcal{X} . We often write $\tilde{P}^T(X)$ in place of $\tilde{P}(T, X)$. We denote $\text{TP}(\mathcal{T}, \mathcal{X})$ the set of all tweakable permutations with tweak space \mathcal{T} and domain \mathcal{X} . The security of a TBC is defined as follows.

Definition 1 (TPRP security). Let $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{X})$ and A be an adversary with oracle access to a tweakable permutation with tweak space \mathcal{T} and domain \mathcal{X} . The advantage of A in breaking the TPRP-security of \tilde{E} is defined as

$$\text{Adv}_{\tilde{E}}^{\text{TPRP}}(A) = \left| \Pr \left[K \leftarrow_{\S} \mathcal{K} : A^{\tilde{E}_K} = 1 \right] - \Pr \left[\tilde{P} \leftarrow_{\S} \text{TP}(\mathcal{T}, \mathcal{X}) : A^{\tilde{P}} = 1 \right] \right|.$$

Note that we do not need the strongest “two-sided” version of TPRP-security (where the adversary also has access to a decryption oracle) since all constructions considered in this paper only use the forward (encryption) direction of the underlying TBC.

TWEAK SEPARATION. Let \tilde{E} be a TBC with tweak space of the form $\mathcal{T}' = \mathcal{I} \times \mathcal{T}$ for some subset $\mathcal{I} \subset \mathbb{N}$ and some set \mathcal{T} . We call \mathcal{T} the *effective* tweak space of \tilde{E} . Then, for $i \in \mathcal{I}$, we denote \tilde{E}^i the tweakable block cipher with the same key and message spaces as \tilde{E} and tweak space \mathcal{T} defined by

$$\tilde{E}^i(K, T, X) = \tilde{E}(K, (i, T), X).$$

By the same convention as before, we write $\tilde{E}_K^i(T, X)$ or $\tilde{E}_K^{i,T}(X)$ for $\tilde{E}^i(K, T, X)$. Clearly, when \tilde{E} is an ideal TBC drawn uniformly at random from $\text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{M})$, then each \tilde{E}^i is an independent ideal TBC drawn uniformly at random from $\text{TBC}(\mathcal{K}, \mathcal{T}, \mathcal{M})$. Given a bit-string X of length $1 \leq |X| \leq n$, we compactly write

$$\tilde{E}^{i/j}(K, T, X(10^*)), \quad \tilde{E}_K^{i/j}(T, X(10^*)), \quad \text{or} \quad \tilde{E}_K^{i/j,T}(X(10^*))$$

to mean $\tilde{E}^i(K, T, X)$ when $|X| = n$ and $\tilde{E}^j(K, T, X10^*)$ when $|X| < n$.

STANDARD SECURITY NOTIONS. We give the security definitions of a nonce-based PRF, a nonce-based MAC, and a nonce-based Authenticated Encryption scheme. All these are standard, except maybe the nonce-based PRF notion which is a

straightforward adaptation of the classical definition of a PRF to the nonce-based setting. Our definition of the security of a MAC is indistinguishability-based (which will be more convenient later), but it is easy to see that it is equivalent to the more conventional unforgeability-based definition. In the following, a nonce-based keyed function is a function $F : \mathcal{K} \times \mathcal{N} \times \mathcal{D} \rightarrow \mathcal{Y}$, where \mathcal{K} is the key space, \mathcal{N} the nonce space, \mathcal{D} the domain and \mathcal{Y} the range.

Definition 2 (Nonce-Based PRF). *Let $F : \mathcal{K} \times \mathcal{N} \times \mathcal{D} \rightarrow \mathcal{Y}$ be a nonce-based keyed function, and let us write $F_K(N, D)$ for $F(K, N, D)$. Let A be an adversary with oracle access to a function from $\mathcal{N} \times \mathcal{D}$ to \mathcal{Y} . The advantage of A against the PRF-security of F is defined as*

$$\text{Adv}_F^{\text{PRF}}(A) = \left| \Pr [K \leftarrow_{\S} \mathcal{K} : A^{F_K} = 1] - \Pr [R \leftarrow_{\S} \text{Func}(\mathcal{N} \times \mathcal{D}, \mathcal{Y}) : A^R = 1] \right|.$$

The adversary is said nonce-respecting if it never repeats a nonce $N \in \mathcal{N}$ in its oracle queries. In that case, we denote its advantage $\text{Adv}_F^{\text{nPRF}}(A)$.

Definition 3 (Nonce-Based MAC). *Let F be as in Definition 2. Let B be an adversary with oracle access to two oracles, the first oracle being a function from $\mathcal{N} \times \mathcal{D}$ to \mathcal{Y} , the second oracle with inputs in $\mathcal{N} \times \mathcal{D} \times \mathcal{Y}$ and outputs in $\{1, \perp\}$. The advantage of B against the MAC-security of F is defined as*

$$\text{Adv}_F^{\text{MAC}}(B) = \left| \Pr [K \leftarrow_{\S} \mathcal{K} : B^{F_K, \text{Ver}_K} = 1] - \Pr [K \leftarrow_{\S} \mathcal{K} : B^{F_K, \text{Rej}} = 1] \right|,$$

where Ver_K is an oracle which takes as input a triple $(N, D, \text{tag}) \in \mathcal{N} \times \mathcal{D} \times \mathcal{Y}$ and returns 1 if $F_K(N, D) = \text{tag}$, and \perp otherwise, and Rej is an oracle which always returns \perp . The adversary is not allowed to ask a verification query (N, D, tag) if a previous query (N, D) to F_K returned tag . The adversary is said nonce-respecting if it never repeats a nonce $N \in \mathcal{N}$ in its queries to the first oracle F_K . In that case, we denote its advantage $\text{Adv}_F^{\text{nMAC}}(B)$.

Note that in the general case where the adversary is allowed to repeat nonces, F can be seen as a standard (i.e., not nonce-based) keyed function with domain $\mathcal{N} \times \mathcal{D}$, in which case one recovers the standard definitions of a PRF and a MAC (hence our notation of the advantage when the adversary is unrestricted w.r.t. nonces). While it is a well-known fact that if F is a secure PRF, then it is a secure MAC [BKR00, GGM86], we stress that this is not true for the nonce-based variants of the two notions, which are in fact incomparable.⁸

We then give the definition of a nonce-based Authenticated Encryption (nAE) scheme (with associated data), for which we first recall the syntax. Let \mathcal{K} , \mathcal{N} , \mathcal{A} , and \mathcal{M} be non-empty sets. A nAE scheme is a tuple $\Pi = (\mathcal{K}, \mathcal{N}, \mathcal{A}, \mathcal{M}, \text{Enc}, \text{Dec})$,

⁸ E.g., an nPRF-secure function F might depend only on the nonce, in which case it is trivial to forge and break nMAC-security, while an nMAC-secure function F might have all its outputs starting with a 0 bit, which allows to trivially break nPRF-security.

where Enc and Dec are deterministic algorithms. The encryption algorithm Enc takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \mathcal{A}$, and a message $M \in \mathcal{M}$, and outputs a binary string $C \in \{0, 1\}^*$ (we assume that Enc returns \perp if one of the inputs is not in the intended set). The decryption algorithm Dec takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, associated data $A \in \mathcal{A}$, and a binary string $C \in \{0, 1\}^*$, and outputs either a message $M \in \mathcal{M}$, or a special symbol \perp . We require that $\text{Dec}(K, N, A, \text{Enc}(K, N, A, M)) = M$ for all tuples $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$. We write $\text{Enc}_K(N, A, M)$ for $\text{Enc}(K, N, A, M)$ and $\text{Dec}_K(N, A, C)$ for $\text{Dec}(K, N, A, C)$.

Definition 4 (Nonce-Based AE). Let $\Pi = (\mathcal{K}, \mathcal{N}, \mathcal{A}, \mathcal{M}, \text{Enc}, \text{Dec})$ be a nAE scheme. The advantage of an adversary \mathbf{A} in breaking Π is defined as

$$\text{Adv}_{\Pi}^{\text{AE}}(\mathbf{A}) = \left| \Pr \left[K \leftarrow_{\S} \mathcal{K} : \mathbf{A}^{\Pi.\text{Enc}_K(\cdot, \cdot, \cdot), \Pi.\text{Dec}_K(\cdot, \cdot, \cdot)} = 1 \right] - \Pr \left[\mathbf{A}^{\text{Rand}(\cdot, \cdot, \cdot), \text{Rej}(\cdot, \cdot, \cdot)} = 1 \right] \right|,$$

where Rand is an oracle which on input $(N, A, M) \in \mathcal{N} \times \mathcal{A} \times \mathcal{M}$ outputs a random⁹ string of length $|\Pi.\text{Enc}_K(N, A, M)|$ and Rej is an oracle which always outputs \perp . The adversary is not allowed to make a decryption query (N, A, C) if a previous encryption query (N, A, M) returned C . The adversary is said nonce-respecting if it never repeats a nonce $N \in \mathcal{N}$ in its encryption queries, in which case we denote its advantage $\text{Adv}_{\Pi}^{\text{nAE}}(\mathbf{A})$.

Note that in the general case where the adversary is allowed to repeat nonces, Π can be seen as a deterministic AE scheme [RS06] with header space (in the terms of [RS06]) $\mathcal{N} \times \mathcal{A}$, so that one exactly recovers the definition of the MRAE notion of Rogaway and Shrimpton [RS06] (which we simply abbreviate to AE here).

ADVERSARY CHARACTERISTICS. In all the paper, given some implicit parameter n , a (q, m, ℓ, σ, t) -adversary against a nonce-based scheme is an adversary:

- which makes at most q oracle queries; when the adversary has access to two oracles (i.e., when attacking the MAC-security of a keyed function or a nAE scheme), this means q queries in total to both oracles;
- which uses any nonce at most m times throughout its queries ($m = 1$ for a nonce-respecting adversary); when the adversary has access to two oracles, this only applies to queries to its first oracle (MAC or encryption oracle);
- such that the length of any of its queries (nonce excluded) is at most ℓ blocks of n bits; for a keyed function with domain $\mathcal{D} = \mathcal{A} \times \mathcal{M}$ or a nAE scheme, this means that *both* the AD length and the message length of any query is at most ℓ blocks of n bits;
- such that the total length of all its queries (nonce excluded) is at most σ blocks of n bits; for a keyed function with domain $\mathcal{D} = \mathcal{A} \times \mathcal{M}$ or a nAE scheme, this means the sum of the AD and the message length over all queries;
- which runs in time at most t .

⁹ We assume that Rand returns the same output if a query is repeated.

4 The CTRT Encryption Mode

4.1 Syntax and Security of nivE Schemes

Most existing encryption schemes are either nonce-based [Rog04b] or IV-based [BDJR97], i.e., they employ an externally provided value which either should not repeat (nonce), or should be selected uniformly at random (IV). (See also [NRS14].) Here, we introduce the notion of combined nonce- and IV-based encryption scheme (*nivE* for short).

Syntactically, a nivE scheme is a tuple $\Pi = (\mathcal{K}, \mathcal{N}, \mathcal{IV}, \mathcal{M}, \text{Enc}, \text{Dec})$ where \mathcal{K} , \mathcal{N} , \mathcal{IV} and \mathcal{M} are non-empty sets and Enc and Dec are deterministic algorithms. The encryption algorithm Enc takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, an initial value $IV \in \mathcal{IV}$, and a message $M \in \mathcal{M}$, and outputs a binary string $C \in \{0, 1\}^*$ (we assume that Enc returns \perp if one of the inputs is not in the intended set). The decryption algorithm Dec takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, an initial value $IV \in \mathcal{IV}$, and a binary string $C \in \{0, 1\}^*$, and outputs either a message $M \in \mathcal{M}$, or a special symbol \perp . We require that

$$\text{Dec}(K, N, IV, \text{Enc}(K, N, IV, M)) = M$$

for all tuples $(K, N, IV, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{IV} \times \mathcal{M}$.

We denote Enc^{\S} the probabilistic algorithm which takes as input $(K, N, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{M}$, internally generates a uniformly random $IV \leftarrow_{\S} \mathcal{IV}$, computes $C = \text{Enc}(K, N, IV, M)$, and outputs $(IV, C) \in \mathcal{IV} \times \{0, 1\}^*$. We write $\text{Enc}_K(N, IV, M)$ for $\text{Enc}(K, N, IV, M)$ and $\text{Enc}_K^{\S}(N, M)$ for $\text{Enc}^{\S}(K, N, M)$. The security of a nivE scheme is defined as follows.

Definition 5 (Security of a nivE scheme). *Let $\Pi = (\mathcal{K}, \mathcal{N}, \mathcal{IV}, \mathcal{M}, \text{Enc}, \text{Dec})$ be a nivE scheme. The advantage of an adversary \mathbf{A} in breaking Π is defined as*

$$\text{Adv}_{\Pi}^{\text{ivE}}(\mathbf{A}) = \left| \Pr \left[K \leftarrow_{\S} \mathcal{K} : \mathbf{A}^{\Pi, \text{Enc}_K^{\S}(\cdot, \cdot)} = 1 \right] - \Pr \left[\mathbf{A}^{\text{Rand}(\cdot, \cdot)} = 1 \right] \right|,$$

where Rand is an oracle which on input $(N, M) \in \mathcal{N} \times \mathcal{M}$ outputs a random string of length $|\Pi.\text{Enc}_K^{\S}(N, M)|$. The adversary is said nonce-respecting if it never repeats a nonce $N \in \mathcal{N}$ in its oracle queries, in which case we denote its advantage $\text{Adv}_{\Pi}^{\text{nivE}}(\mathbf{A})$.

Note that when the adversary is allowed to repeat nonces, Π can be seen as a family of purely IV-based encryption (ivE) schemes [NRS14] indexed by the nonce space \mathcal{N} , hence our notation of the advantage in that case.

4.2 Definition and Analysis of the CTRT Mode

We now define the CTRT (*CounTeR in Tweak*) mode, turning a tweakable block cipher into a nivE scheme. Let \mathcal{K} and \mathcal{T} be non-empty sets, and let

```

1 algorithm CTRT[ $\tilde{E}$ ].Enc $_K(N, IV, M)$ 
2    $\ell := |M|/n$ 
3   parse  $M$  as  $M_1 \parallel \dots \parallel M_\ell$ , with  $|M_1| = \dots = |M_{\ell-1}| = n$  and  $1 \leq |M_\ell| \leq n$ 
4   for  $i = 1$  to  $\ell$  do
5      $C_i := M_i \oplus \tilde{E}_K^{1, IV}(N)$ 
6      $IV := \text{Inc}(IV)$ 
7   return  $C_1 \parallel C_2 \parallel \dots \parallel C_\ell$ 

```

Fig. 2. Definition of the CTRT mode, using a TBC $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{T}' = \{1\} \times \mathcal{T}$ and $\mathcal{X} = \{0, 1\}^n$.

$\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ be a tweakable block cipher with key space \mathcal{K} , tweak space¹⁰ $\mathcal{T}' = \{1\} \times \mathcal{T}$, and domain $\mathcal{X} = \{0, 1\}^n$. Let Inc be a cyclic permutation of \mathcal{T} . We construct from \tilde{E} a nivE scheme $\text{CTRT}[\tilde{E}]$ with key space \mathcal{K} , nonce space $\mathcal{N} = \mathcal{X} = \{0, 1\}^n$, IV space $\mathcal{IV} = \mathcal{T}$, and message space $\mathcal{M} = \{0, 1\}^*$ as defined in Fig. 2 and illustrated on bottom of Fig. 1.

The security of CTRT is captured by Theorem 1 below. Logarithms are in base 2 and $t_{\text{CTRT}}(\sigma)$ is an upper bound on the time needed for computing $\text{CTRT}[\tilde{E}].\text{Enc}_K$ on inputs of total message length at most σ blocks of n bits when calls to \tilde{E}_K cost unit time.

Theorem 1 (Security of CTRT). *Let $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{X} = \{0, 1\}^n$, $\mathcal{T}' = \{1\} \times \mathcal{T}$, and $|\mathcal{T}| \geq 8$. Let \mathbf{A} be a (q, m, ℓ, σ, t) -adversary against $\text{CTRT}[\tilde{E}]$ with $\ell \leq |\mathcal{T}|$. Then there exists an adversary \mathbf{A}' against the TPRP-security of \tilde{E} , making at most σ oracle queries and running in time at most $t + t_{\text{CTRT}}(\sigma)$, such that*

$$\text{Adv}_{\text{CTRT}[\tilde{E}]}^{\text{ivE}}(\mathbf{A}) \leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{A}') + \frac{2(m-1)\sigma + 1}{|\mathcal{T}|} + f(\sigma),$$

where

$$\begin{aligned} f(\sigma) &= \frac{2\sigma \log^2 \sigma}{|\mathcal{X}|} && \text{when } 8 \leq \sigma \leq |\mathcal{T}|, \\ &= \frac{2\sigma^2 \log^2 |\mathcal{T}|}{|\mathcal{X}||\mathcal{T}|} && \text{when } \sigma \geq |\mathcal{T}|. \end{aligned}$$

In particular, if \mathbf{A} is nonce-respecting ($m = 1$), one has

$$\text{Adv}_{\text{CTRT}[\tilde{E}]}^{\text{nivE}}(\mathbf{A}) \leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{A}') + \frac{1}{|\mathcal{T}|} + f(\sigma).$$

¹⁰ The CTRT mode does not need tweak separation per se. We use a single 1 prefix in order to conveniently combine CTRT with EPWC later.

Before proceeding to the proof, we comment the security bound of this theorem. Consider first the case of a nonce-respecting adversary A . Assuming $|\mathcal{T}| \geq |\mathcal{X}|$, then A must make queries of total message length σ blocks of n bits with σ close to $|\mathcal{X}| = 2^n$ (neglecting logarithmic factors) before being able to distinguish the outputs of $\text{CTRTR}[\tilde{E}]$ from random.¹¹ On the other hand, if $|\mathcal{T}| = 2^w < |\mathcal{X}|$, then $\text{CTRTR}[\tilde{E}]$ is secure up to roughly $2^{(n+w)/2}$ TBC calls (again, neglecting logarithmic factors), which is always larger than $2^{n/2}$. In particular, if $w = n/2$ (as e.g. for KIASU-BC [JNP14c]), then security is ensured up to roughly $2^{3n/4}$ TBC calls. In the nonce-misuse scenario, note that the additional term $2(m-1)\sigma/|\mathcal{T}|$ remains small as long as nonces are not repeated too many times (e.g. $m \leq 100$) and $\sigma \ll |\mathcal{T}|$, and turns into a birthday-like term only in the extreme case where a few nonces are repeated close to σ times. This means that a few nonce repetitions will not hurt and that nonces must be “seriously” mishandled before security goes down to birthday bound.

PROOF OF THEOREM 1. Fix a $(q', m, |\mathcal{T}|, \sigma, t)$ -adversary A against $\text{CTRTR}[\tilde{E}]$ (we denote q' the maximal number of adversarial queries and will later use q for the actual number of queries in a specific attack). The first part of the proof is standard, and consists in introducing an intermediate game where all calls to \tilde{E}_K in the CTRTR construction are replaced by calls to a random tweakable permutation \tilde{P} . Consider the following adversary A' against the TPRP-security of \tilde{E} . Let $G \in \{\tilde{E}_K, \tilde{P}\}$ be the oracle to which A' has access. Adversary A' runs A , answers its encryption queries (N, M) by drawing a random IV and executing the code in Fig. 2 on input (N, IV, M) , replacing calls to \tilde{E}_K by oracle calls to G , and finally outputs the same bit as A . Clearly, A' makes at most σ oracle queries and runs in time at most $t + t_{\text{CTRTR}}(\sigma)$. Moreover, it is easy to see that

$$\text{Adv}_{\text{CTRTR}[\tilde{E}]}^{\text{ivE}}(A) \leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(A') + \delta, \quad (1)$$

where

$$\delta = \left| \Pr \left[\tilde{P} \leftarrow_{\S} \text{TP}(\mathcal{T}, \mathcal{X}) : A^{\text{CTRTR}[\tilde{P}].\text{Enc}^{\S}} = 1 \right] - \Pr \left[A^{\text{Rand}} = 1 \right] \right| \quad (2)$$

and $\text{CTRTR}[\tilde{P}]$ is a slight abuse of notation for the CTRTR construction based on an arbitrary tweakable permutation \tilde{P} .

Upper bounding δ is now a purely information-theoretic problem, so that we allow A to be computationally unbounded, and hence, *wlog*, deterministic. The adversary is now trying to distinguish between $\text{CTRTR}[\tilde{P}]$ for a random \tilde{P} (thereafter called the “real world”) and Rand (thereafter called the “ideal world”). We assume *wlog* that A always makes queries of length a multiple of the block length n , and of total length σ blocks (if not, we pad all queries whose final block is incomplete with zeros for free, which can only increase the adversary’s advantage).

¹¹ Note that in that case the size of the tweak space \mathcal{T} only impacts the maximal message length ℓ , not the security bound.

Following the H-coefficients method [Pat08b, CS14], we summarize the interaction of \mathbf{A} with its oracle in the so-called *transcript* of the attack

$$\tau = ((N_1, M_1, IV_1, C_1), \dots, (N_q, M_q, IV_q, C_q)),$$

where (N_i, M_i) denotes the i -th query of the attacker and (IV_i, C_i) the corresponding answer of the oracle. Furthermore, we denote $M_i = M_{i,1} \parallel \dots \parallel M_{i,\ell_i}$, $C_i = C_{i,1} \parallel \dots \parallel C_{i,\ell_i}$, where ℓ_i is the number of blocks of the i -th message, and $IV_{i,j} = \text{Inc}^{j-1}(IV_i)$ the j -th counter for the i -th message, $j = 1, \dots, \ell_i$. Let Θ_{re} , resp. Θ_{id} , denote the distribution of the transcript in the real world, resp. ideal world. We say that a transcript τ is \mathbf{A} -attainable (or simply attainable) if the probability to obtain τ in the ideal world is non-zero. Note that the number of queries q and the lengths of the queries ℓ_1, \dots, ℓ_q are themselves random variables (they can vary for distinct \mathbf{A} -attainable transcripts), yet by the assumption that the attacker always asks the maximal number of allowed blocks throughout its queries, one always has $\sum_{i=1}^q \ell_i = \sigma$.

From τ we define for each possible tweak $T \in \mathcal{T}$ the “load” of the tweak as

$$L(T) = |\{(i, j) : IV_{i,j} = T\}|.$$

In words, $L(T)$ is the number of times the tweak T appears as a counter when encrypting the queries of the adversary. Clearly, one has

$$\sum_{T \in \mathcal{T}} L(T) = \sigma. \quad (3)$$

The proof relies on the fundamental lemma of the H-coefficients technique (see e.g. [CS14] for the proof).

Lemma 1. *Assume that the set of \mathbf{A} -attainable transcripts is partitioned into two disjoint sets $\text{Good}\mathcal{T}$ and $\text{Bad}\mathcal{T}$, and that there exists ε_1 and ε_2 such that for any $\tau \in \text{Good}\mathcal{T}$, one has*

$$\frac{\Pr[\Theta_{\text{re}} = \tau]}{\Pr[\Theta_{\text{id}} = \tau]} \geq 1 - \varepsilon_1,$$

and $\Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}] \leq \varepsilon_2$. Then $\delta \leq \varepsilon_1 + \varepsilon_2$, with δ as defined by (2).

We say that an attainable transcript τ is bad if one of the two following conditions are met:

- (C-1) there exists $(i, j) \neq (i', j')$ such that $IV_{i,j} = IV_{i',j'}$ and $N_i = N_{i'}$;
- (C-2) there exists $(i, j) \neq (i', j')$ such that $IV_{i,j} = IV_{i',j'}$, $N_i \neq N_{i'}$, and $M_{i,j} \oplus C_{i,j} = M_{i',j'} \oplus C_{i',j'}$.

Note that condition (C-1) can only be satisfied for a nonce-misuse adversary, since (by the assumption that the length of each query is at most $|\mathcal{T}|$ blocks of n bits so that counters do not loop) $IV_{i,j} = IV_{i',j'}$ requires $i \neq i'$, which implies that $N_i \neq N_{i'}$ for a nonce-respecting adversary. Note also that the

condition (C-2) can only be satisfied in the ideal world. Indeed, in the real world, $M_{i,j} \oplus C_{i,j} = \tilde{P}(IV_{i,j}, N_i)$ and $M_{i',j'} \oplus C_{i',j'} = \tilde{P}(IV_{i',j'}, N_{i'})$, so that if $IV_{i,j} = IV_{i',j'}$ and $N_i \neq N_{i'}$ one necessarily has $M_{i,j} \oplus C_{i,j} \neq M_{i',j'} \oplus C_{i',j'}$.

We let **BadT** be the set of bad transcripts, and **GoodT** be the set of attainable transcripts which are not bad, henceforth called good transcripts. We first consider good transcripts.

Lemma 2. *Let $\tau \in \text{GoodT}$ be a good transcript. Then*

$$\frac{\Pr[\Theta_{\text{re}} = \tau]}{\Pr[\Theta_{\text{id}} = \tau]} \geq 1.$$

Proof. Note that a good transcript has the property that for each $(i, j) \neq (i', j')$ such that $IV_{i,j} = IV_{i',j'}$, one has $N_i \neq N_{i'}$ and $M_{i,j} \oplus C_{i,j} \neq M_{i',j'} \oplus C_{i',j'}$. In other words, the transcript encodes a partial tweakable permutation, where for each tweak $T \in \mathcal{T}$ there are exactly $L(T)$ distinct values N_i mapped to some value $M_{i,j} \oplus C_{i,j}$. The probability to obtain a good transcript τ in the ideal and real worlds can now be easily computed. In the ideal world, since the IV_i 's and the C_i 's are uniformly random, one has

$$\Pr[\Theta_{\text{id}} = \tau] = \frac{1}{|\mathcal{IV}|^q \cdot |\mathcal{X}|^\sigma}.$$

In the real world, the IV_i 's are random as well, but now the probability to obtain the C_i 's can easily be seen to be the probability that the random tweakable permutation \tilde{P} is compatible with the partial tweakable permutation encoded by τ [CS14]. Hence, one has

$$\Pr[\Theta_{\text{re}} = \tau] = \frac{1}{|\mathcal{IV}|^q} \cdot \prod_{T \in \mathcal{T}} \frac{1}{(|\mathcal{X}|)_{L(T)}},$$

where $(a)_b$ denotes the falling factorial $a(a-1)\cdots(a-b+1)$, with $(a)_0 = 1$ by convention. From this, we deduce that

$$\frac{\Pr[\Theta_{\text{re}} = \tau]}{\Pr[\Theta_{\text{id}} = \tau]} = \frac{|\mathcal{X}|^\sigma}{\prod_{T \in \mathcal{T}} (|\mathcal{X}|)_{L(T)}} \stackrel{(3)}{=} \prod_{T \in \mathcal{T}} \frac{|\mathcal{X}|^{L(T)}}{(|\mathcal{X}|)_{L(T)}} \geq 1. \quad \square$$

It remains to upper bound the probability to obtain a bad transcript in the ideal world. For $i \in \{1, 2\}$, let BadT_i be the set of attainable transcripts satisfying condition (C- i). We first consider condition (C-1).

Lemma 3. *One has*

$$\Pr[\Theta_{\text{id}} \in \text{BadT}_1] \leq \frac{2(m-1)\sigma}{|\mathcal{T}|}.$$

Proof. Consider two distinct queries (N_i, M_i) and $(N_{i'}, M_{i'})$. If the nonces are the same ($N_i = N_{i'}$), then the probability, over the random draw of IV_i and $IV_{i'}$

in \mathcal{T} , that there exists j and j' such that $IV_{i,j} = IV_{i',j'}$, is $(\ell_i + \ell_{i'} - 1)/|\mathcal{T}|$. If the nonces are distinct, then clearly condition (C-1) cannot be satisfied for i and i' . Hence, summing over all possible nonces, we have

$$\Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}_1] \leq \sum_{N \in \mathcal{N}} \sum_{\substack{1 \leq i < i' \leq q \\ N_i = N_{i'} = N}} \frac{\ell_i + \ell_{i'} - 1}{|\mathcal{T}|}.$$

Fix some nonce N , and assume for notational simplicity that the first q' queries use nonce N , with $q' \leq m$ by assumption. Then the probability that condition (C-1) is met for nonce N is at most

$$\begin{aligned} \sum_{i=1}^{q'-1} \sum_{i'=i+1}^{q'} \frac{\ell_i + \ell_{i'} - 1}{|\mathcal{T}|} &\leq \sum_{i=1}^{q'-1} \frac{(q' - 1)\ell_i + \ell(N)}{|\mathcal{T}|} \\ &\leq \frac{2(q' - 1)\ell(N)}{|\mathcal{T}|} \\ &\leq \frac{2(m - 1)\ell(N)}{|\mathcal{T}|}, \end{aligned}$$

where $\ell(N)$ is the total length of queries using nonce N . The result follows by summing over all possible nonces, using $\sum_{N \in \mathcal{N}} \ell(N) = \sigma$. \square

We handle condition (C-2) in the following lemma.

Lemma 4. *One has*

$$\Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}_2] \leq \frac{1}{|\mathcal{T}|} + \min\{\sigma, |\mathcal{T}|\} \cdot \frac{(L_{\max})^2}{2|\mathcal{X}|},$$

where $L_{\max} = 2 \log \sigma$ when $8 \leq \sigma \leq |\mathcal{T}|$ and $L_{\max} = \frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|}$ when $\sigma \geq |\mathcal{T}|$.

Proof. Let $\text{Bad}\mathcal{T}_3$ be the set of transcripts satisfying the following condition (L_{\max} being defined as in the statement of the lemma):

(C-3) there exists $T \in \mathcal{T}$ such that $L(T) \geq L_{\max}$.

Then

$$\Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}_2] \leq \Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}_2 \mid \Theta_{\text{id}} \notin \text{Bad}\mathcal{T}_3] + \Pr[\Theta_{\text{id}} \in \text{Bad}\mathcal{T}_3].$$

Note that in the ideal world, the values $L(T)$ only depend on the random draw of the IV_i 's, and that once the $L(T)$'s are fixed, condition (C-2) only depends on the random draw of the $C_{i,j}$'s. In particular, since in the ideal world the $C_{i,j}$'s

are uniformly random, one has

$$\begin{aligned}
& \Pr[\Theta_{\text{id}} \in \text{BadT}_2 \mid \Theta_{\text{id}} \notin \text{BadT}_3] \\
& \leq \sum_{T \in \mathcal{T}} \sum_{\substack{(i,j),(i',j') \\ IV_{i,j} = IV_{i',j'} = T}} \Pr[M_{i,j} \oplus C_{i,j} = M_{i',j'} \oplus C_{i',j'}] \\
& \leq \sum_{T \in \mathcal{T}} \frac{L(T)(L(T) - 1)}{2|\mathcal{X}|} \\
& \leq \min\{\sigma, |\mathcal{T}|\} \cdot \frac{(L_{\max})^2}{2|\mathcal{X}|},
\end{aligned}$$

where for the third inequality we used that there are at most $\min\{\sigma, |\mathcal{T}|\}$ tweaks T such that $L(T) \geq 1$.

It remains to upper bound the probability that condition (C-3) is satisfied, which can be recast as a “balls-into-bins” problem. Thinking of each tweak T as a bin, each random IV_i determines a sequence of ℓ_i consecutive¹² bins where a ball is thrown. If the attacker only made queries of length one block, then this would be a standard “balls-into-bins” problem, where each ball is thrown independently in a bin chosen uniformly at random, and we could use classical results about the maximal occupancy of any bin directly. However, the attacker can choose the length of each message at will and we need to take this into account.¹³ Intuitively, for some fixed total number σ of balls, using messages of length $\ell_i > 1$ should *lower* the maximal occupancy since balls thrown in consecutive bins cannot end in the same bin. We formalize this intuition in a separate lemma below, which implies that

$$\Pr[\Theta_{\text{id}} \in \text{BadT}_3] \leq \frac{1}{|\mathcal{T}|}.$$

The result follows. \square

The lemma below is a simple variant on the standard balls-into-bins problem. A similar result was proved in [ANSS16] (and potentially in many other papers).

Lemma 5. *Consider a set of $|\mathcal{T}| \geq 8$ bins and $\sigma \geq 8$ balls. Fix an integer $q \leq \sigma$ and a sequence of integers (ℓ_1, \dots, ℓ_q) with $1 \leq \ell_i \leq |\mathcal{T}|$ and $\sum_{i=1}^q \ell_i = \sigma$. Consider the following random process: for $i = 1, \dots, q$, a chain of ℓ_i balls is thrown in consecutive bins, the initial bin being chosen independently and uniformly at random. Then the probability that, at the end of the process, any bin contains L_{\max} balls or more, is less than $1/|\mathcal{T}|$, where*

- (a) $L_{\max} = 2 \log \sigma$ when $\sigma \leq |\mathcal{T}|$;
- (b) $L_{\max} = \frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|}$ when $\sigma \geq |\mathcal{T}|$.

Proof. Deferred to Appendix A. \square

¹² The successor of the tweak T is $\text{Inc}(T)$.

¹³ Note that the adversary must commit to the length ℓ_i of the chain *before* knowing the initial bin IV_i since it first makes the query (N_i, M_i) and only then receives the answer (IV_i, C_i) .

COMPLETING THE PROOF OF THEOREM 1. From Lemmas 3 and 4, we obtain by the union bound that

$$\Pr[\Theta_{\text{id}} \in \text{BadT}] \leq \frac{2(m-1)\sigma + 1}{|\mathcal{T}|} + f(\sigma), \quad (4)$$

with $f(\sigma)$ as in the statement of Theorem 1. Combining (4) with Lemmas 1 and 2 (taking $\varepsilon_1 = 0$), we obtain the same upper bound for δ (defined by (2)) as for $\Pr[\Theta_{\text{id}} \in \text{BadT}]$. Finally, Eq. (1) yields the result.

VARIANTS. In Appendix C, we describe two variants of CTRT, a purely nonce-based one and a purely IV-based one.

5 The PWC and EPWC Message Authentication Codes

In this section, we describe two related modes for message authentication, PWC (*Parallel Wegman-Carter*) and EPWC (*Encrypted PWC*). Let \mathcal{K} and \mathcal{T} be two sets, and let \tilde{E} be a tweakable block cipher with key space \mathcal{K} , tweak space¹⁴ $\mathcal{T}' = \{2, \dots, 5\} \times \mathcal{T}$, and domain $\mathcal{X} = \{0, 1\}^n$. Let Inc be a cyclic permutation of \mathcal{T} . From \tilde{E} , we construct two nonce-based keyed functions, $\text{PWC}[\tilde{E}]$ and $\text{EPWC}[\tilde{E}]$, both with key space \mathcal{K} , nonce-space $\mathcal{N} = \mathcal{X} = \{0, 1\}^n$, domain $\mathcal{D} = \mathcal{A} \times \mathcal{M}$, where $\mathcal{A} = \mathcal{M} = \{0, 1\}^*$,¹⁵ and range $\mathcal{Y} = \mathcal{X} = \{0, 1\}^n$, as defined in Fig. 3 and illustrated on top of Fig. 1 (for PWC, just omit the final call to $\tilde{E}_K^{4,0}$). We will prove that both $\text{PWC}[\tilde{E}]$ and $\text{EPWC}[\tilde{E}]$ are 2^n -secure as nonce-based MAC and nonce-based PRF, and that $\text{EPWC}[\tilde{E}]$ is moreover a birthday bound-secure PRF in the nonce-misuse scenario.

The PWC construction follows the Wegman-Carter paradigm [WC81, Bra82, Sho96] by combining a xor-universal hash function H inspired from PMAC [BR02, Rog04a] applied to (A, M) , and a pseudorandom function F applied to the nonce N . This pseudorandom function is constructed from \tilde{E} by summing two independent pseudorandom permutations in order to obtain security beyond the birthday bound [Luc00]. The EPWC construction is simply PWC with an additional layer of encryption to provide nonce-misuse resistance.

Before stating and proving the security results for (E)PWC, we focus on how to obtain the BBB-secure pseudorandom function F from \tilde{E} . A straightforward way would be to “put the nonce in the tweak”, e.g.,

$$F'_K(N) = \tilde{E}_K^{6,N}(0).$$

This would result in a uniformly random value for each new nonce, but this is only possible when the intended nonce space is smaller than the effective tweak

¹⁴ We use a set of prefixes which is disjoint from the set used for the CTRT mode in order to later combine the two modes smoothly.

¹⁵ When constructing an AE scheme, it is more convenient to directly define a vector-input MAC, rather than a string-input MAC that must later be transformed to handle vectors of strings, as required for an AE scheme.

space \mathcal{T} of \tilde{E} . In order to allow the nonce length to be as large as the block length of \tilde{E} , we use instead the “sum-of-PRPs” construction by defining (the exact tweak prefixes are unimportant)

$$F_K(N) = \tilde{E}_K^{2,0}(N) \oplus \tilde{E}_K^{2,1}(N). \quad (5)$$

The pseudorandomness of this construction has been well studied. Assuming that $\tilde{E}_K^{2,0}$ and $\tilde{E}_K^{2,1}$ are perfectly random and independent permutations, Lucks [Luc00, Theorem 5] showed that an information-theoretic adversary trying to distinguish F_K from a random function $\rho : \{0,1\}^n \rightarrow \{0,1\}^n$ within q queries has an advantage upper bounded by $q^3/2^{2n-1}$ (see also [CLP14]). Better bounds were proposed in three different papers: Bellare and Impagliazzo [BI99] proved that the advantage is upper bounded by $\mathcal{O}(n)(q/2^n)^{1.5}$, while Patarin proved in two different ways [Pat08a, Pat13] an upper bound $\mathcal{O}(q/2^n)$. However, in all three cases the exact $\mathcal{O}(\cdot)$ function was left unspecified and the upper bound was not explicitly worked out. For the sake of concreteness, we propose the following optimistic conjecture.

Conjecture 1. There is an absolute constant C such the advantage of any adversary trying to distinguish the sum of two independent random permutations of \mathcal{X} from a random function from \mathcal{X} to \mathcal{X} within q queries is at most $Cq/|\mathcal{X}|$.

The security of PWC and EPWC is captured by Theorems 2 and 3 below. We denote by $t_{\text{PWC}}(\sigma)$, resp. $t_{\text{EPWC}}(\sigma)$, an upper bound on the time needed to compute $\text{PWC}[\tilde{E}]$, resp. $\text{EPWC}[\tilde{E}]$ on inputs of total (AD + message) length at most σ blocks of n bits when calls to \tilde{E}_K cost unit time.

Theorem 2 (PRF-Security of PWC and EPWC). *Let $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{X} = \{0,1\}^n$ and $\mathcal{T}' = \{2, \dots, 5\} \times \mathcal{T}$, and assume Conjecture 1. Let A be a (q, m, ℓ, σ, t) -adversary against the PRF-security of $\text{PWC}[\tilde{E}]$, resp. $\text{EPWC}[\tilde{E}]$, with $\ell \leq |\mathcal{T}'| - 2$. Then there exists an absolute constant C and an adversary A' , resp. A'' , against the TPRP-security of \tilde{E} , making at most $\sigma + 2q$, resp. $\sigma + 3q$ oracle queries and running it time at most $t + t_{\text{PWC}}(\sigma)$, resp. $t + t_{\text{EPWC}}(\sigma)$, such that*

(a) if A is nonce-respecting ($m = 1$), then

$$\begin{aligned} \text{Adv}_{\text{PWC}[\tilde{E}]}^{\text{nPRF}}(A) &\leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(A') + \frac{Cq}{|\mathcal{X}|}; \\ \text{Adv}_{\text{EPWC}[\tilde{E}]}^{\text{nPRF}}(A) &\leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(A'') + \frac{Cq}{|\mathcal{X}|}; \end{aligned}$$

(b) if A is allowed to repeat nonces ($m > 1$), then

$$\text{Adv}_{\text{EPWC}[\tilde{E}]}^{\text{PRF}}(A) \leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(A'') + \frac{q^2}{|\mathcal{X}|}.$$

```

1 algorithm PWC[ $\tilde{E}$ ] $_K(N, A, M)$  EPWC[ $\tilde{E}$ ] $_K(N, A, M)$ 
2    $\ell_a := |A|/n$ 
3    $\ell_m := |M|/n$ 
4   parse  $A$  as  $A_1 \parallel \dots \parallel A_{\ell_a}$ , with  $|A_1| = \dots = |A_{\ell_a-1}| = n$ ,  $1 \leq |A_{\ell_a}| \leq n$ 
5   parse  $M$  as  $M_1 \parallel \dots \parallel M_{\ell_m}$ , with  $|M_1| = \dots = |M_{\ell_m-1}| = n$ ,  $1 \leq |M_{\ell_m}| \leq n$ 
6    $\text{auth} := \tilde{E}_K^{2,0}(N) \oplus \tilde{E}_K^{2,1}(N)$ 
7   if  $\ell_a > 0$  then
8     for  $i = 1$  to  $\ell_a - 1$  do
9        $\text{auth} := \text{auth} \oplus \tilde{E}_K^{2,i+1}(A_i)$ 
10      if  $|A_{\ell_a}| = n$  then
11         $\text{auth} := \text{auth} \oplus \tilde{E}_K^{2,\ell_a+1}(A_{\ell_a})$ 
12      else
13         $\text{auth} := \text{auth} \oplus \tilde{E}_K^{3,\ell_a+1}(A_{\ell_a} 10^*)$ 
14  if  $\ell_m > 0$  then
15    for  $i = 1$  to  $\ell_m - 1$  do
16       $\text{auth} := \text{auth} \oplus \tilde{E}_K^{4,i}(M_i)$ 
17      if  $|M_{\ell_m}| = n$  then
18         $\text{auth} := \text{auth} \oplus \tilde{E}_K^{4,\ell_m}(M_{\ell_m})$ 
19      else
20         $\text{auth} := \text{auth} \oplus \tilde{E}_K^{5,\ell_m}(M_{\ell_m} 10^*)$ 
21   $\text{tag} := \tilde{E}_K^{4,0}(\text{auth})$ 
22  return tag

```

Fig. 3. Definition of the PWC and EPWC modes, using a TBC $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{T}' = \{2, \dots, 5\} \times \mathcal{T}$ and $\mathcal{X} = \{0, 1\}^n$. The boxed statement only applies to EPWC. For notational simplicity, we identify \mathcal{T} with $\{0, \dots, |\mathcal{T}| - 1\}$ and $\text{Inc}^i(0)$ with i .

Proof. Fix a $(q, m, |\mathcal{T}| - 2, \sigma, t)$ -adversary A against the PRF-security of PWC[\tilde{E}] or EPWC[\tilde{E}], trying to distinguish the construction from a random function $R \leftarrow_{\S} \text{Func}(\mathcal{N} \times \mathcal{D}, \mathcal{Y})$, where $\mathcal{D} = \mathcal{A} \times \mathcal{M}$. We start by proving (a), assuming A is nonce-respecting ($m = 1$). First, slightly abusing the notation, let us see (E)PWC as a construction based on an arbitrary tweakable permutation, identifying (E)PWC[\tilde{E}] $_K$ with (E)PWC[\tilde{E}_K]. We start by replacing \tilde{E}_K in the security experiment by a uniformly random tweakable permutation \tilde{P} . One can see $A^{\text{PWC}[\cdot]}$, resp. $A^{\text{EPWC}[\cdot]}$, as an adversary A' , resp. A'' against the TPRP-security of \tilde{E} , making at most $\sigma + 2q$, resp. $\sigma + 3q$ queries to its oracle (since a query of ℓ_i blocks to PWC, resp. EPWC, costs $\ell_i + 2$, resp. $\ell_i + 3$ calls to \tilde{E}) and running in time at most $t' = t + t_{\text{PWC}}(\sigma)$, resp. $t' = t + t_{\text{EPWC}}(\sigma)$, so that

$$\text{Adv}_{\text{PWC}[\tilde{E}]}^{\text{nPRF}}(A) \leq \text{Adv}_{\tilde{E}}^{\text{TPRP}}(A') + \delta, \quad (6)$$

$$\mathbf{Adv}_{\text{EPWC}[\tilde{E}]}^{\text{nPRF}}(\mathbf{A}) \leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{A}'') + \delta, \quad (7)$$

where

$$\delta = \left| \Pr \left[\tilde{P} \leftarrow_{\S} \text{TP}(\mathcal{T}', \mathcal{X}) : \mathbf{A}^{(\text{E})\text{PWC}[\tilde{P}]} = 1 \right] - \Pr \left[R \leftarrow_{\S} \text{Func}(\mathcal{N} \times \mathcal{D}, \mathcal{Y}) : \mathbf{A}^R = 1 \right] \right|. \quad (8)$$

In order to upper bound δ , we abstract the high-level structure of $(\text{E})\text{PWC}[\tilde{P}]$ as follows. First, we see how A and M are handled as applying a keyed hash function (the key being \tilde{P}) to the pair (A, M) , viz.

$$H_{\tilde{P}}(A, M) \stackrel{\text{def}}{=} \left(\bigoplus_{i=1}^{\ell_a-1} \tilde{P}^{2,i+1}(A_i) \right) \oplus \tilde{P}^{2/3, \ell_a+1}(A_{\ell_a}(10^*)) \\ \oplus \left(\bigoplus_{i=1}^{\ell_m-1} \tilde{P}^{4,i}(M_i) \right) \oplus \tilde{P}^{4/5, \ell_m}(M_{\ell_m}(10^*)). \quad (9)$$

We also define a pseudorandom function (again with key \tilde{P}) as

$$F_{\tilde{P}}(N) \stackrel{\text{def}}{=} \tilde{P}^{2,0}(N) \oplus \tilde{P}^{2,1}(N). \quad (10)$$

Then, $(\text{E})\text{PWC}[\tilde{P}]$ can be written

$$\text{PWC}[\tilde{P}](N, A, M) = F_{\tilde{P}}(N) \oplus H_{\tilde{P}}(A, M), \quad (11)$$

$$\text{EPWC}[\tilde{P}](N, A, M) = \tilde{P}^{4,0}(F_{\tilde{P}}(N) \oplus H_{\tilde{P}}(A, M)), \quad (12)$$

which should make it clear that the PWC construction follows the Wegman-Carter paradigm [WC81] with an additional layer of encryption for EPWC (note that the three sets of tweaks used in F , H , and for the final encryption call are disjoint, so that these three building blocks are independent).

We start by showing that the hash function family $(H_{\tilde{P}})$, with $\tilde{P} \in \text{TP}(\mathcal{T}', \mathcal{X})$, is xor-universal, i.e., for any two distinct inputs (A, M) , (A', M') , and any $X \in \mathcal{X} = \{0, 1\}^n$, the probability, over the random draw of $\tilde{P} \leftarrow_{\S} \text{TP}(\mathcal{T}', \mathcal{X})$, that

$$H_{\tilde{P}}(A, M) \oplus H_{\tilde{P}}(A', M') = X, \quad (13)$$

is less than $1/|\mathcal{X}|$. Assume that $A \neq A'$ (the reasoning is similar if $A = A'$ and $M \neq M'$), let $\ell_a = |A|/n$ and $\ell'_a = |A'|/n$, and assume *wlog* that $\ell_a \geq \ell'_a$. Denote $A = A_1 \parallel \dots \parallel A_{\ell_a}$ and $A' = A'_1 \parallel \dots \parallel A'_{\ell'_a}$. Assume first that $\ell_a > \ell'_a$. Then (13) is equivalent to

$$\tilde{P}^{2/3, \ell_a+1}(A_{\ell_a}(10^*)) = Z,$$

where Z is independent of permutations \tilde{P}^{2, ℓ_a+1} and \tilde{P}^{3, ℓ_a+1} , hence the probability is exactly $1/|\mathcal{X}|$. Assume now that $\ell_a = \ell'_a$. There is necessarily an index $i \leq \ell_a$ such that $A_i \neq A'_i$. If $i < \ell_a$, then (13) is equivalent to

$$\tilde{P}^{2, i+1}(A_i) \oplus \tilde{P}^{2, i+1}(A'_i) = Z,$$

where Z is a value potentially depending on \tilde{P} for tweaks different from $(2, i + 1)$. Since the probability of this equality (over the random draw of $\tilde{P}^{2, i+1}$) is either 0 when $Z = 0$ or exactly $1/|\mathcal{X}|$ when $Z \neq 0$, it follows that the condition is met with probability at most $1/|\mathcal{X}|$ in that case. Similarly, if $i = \ell_a$, then (13) is equivalent to

$$\tilde{P}^{2/3, i+1}(A_i(10^*)) \oplus \tilde{P}^{2/3, i+1}(A'_i(10^*)) = Z,$$

where Z is a value potentially depending on \tilde{P} for tweaks different from $(2, i + 1)$ and $(3, i + 1)$. Again, the condition is met with probability at most $1/|\mathcal{X}|$ in that case. This concludes the proof that H is xor-universal.

As a second step, we replace $F_{\tilde{P}}$ by a uniformly random function ρ from $\mathcal{N} = \{0, 1\}^n$ to $\mathcal{X} = \{0, 1\}^n$. Let $\text{PWC}'[\rho, \tilde{P}]$, resp. $\text{EPWC}'[\rho, \tilde{P}]$ be defined as in (11), resp. (12), except that $F_{\tilde{P}}$ is replaced by a call to ρ . Since A is nonce-respecting, then both $\text{PWC}'[\rho, \tilde{P}]$ and $\text{EPWC}'[\rho, \tilde{P}]$ are perfectly indistinguishable from Rand (this is obvious for PWC' , while for EPWC' this follows from the fact that applying any fixed permutation to uniformly random values yields uniformly random values). Hence, it remains to upper bound A 's advantage in distinguishing $(\text{E})\text{PWC}[\tilde{P}]$ from $(\text{E})\text{PWC}'[\rho, \tilde{P}]$. By a straightforward hybrid argument, this is exactly the advantage of an adversary A''' simulating H (and $\tilde{P}^{4,0}$ for EPWC) in distinguishing $F_{\tilde{P}}$ from ρ within at most q queries (since each query to the construction translates in exactly one query to the function applied to the nonce). Using Conjecture 1, this advantage is upper bounded by $Cq/|\mathcal{X}|$. Combining this with (6), resp. (7), we obtain the result.

We then prove (b), assuming A is allowed to repeat nonces ($m > 1$). Exactly as before, one has

$$\mathbf{Adv}_{\text{EPWC}[\tilde{P}]}^{\text{PRF}}(A) \leq \mathbf{Adv}_E^{\text{TPRP}}(A'') + \delta,$$

with δ defined as in (8). We now see $\text{EPWC}[\tilde{P}]$ as a construction based on a universal hash function applied to (N, A, M) followed by a PRF. More precisely, let

$$H'_{\tilde{P}}(N, A, M) = F_{\tilde{P}}(N) \oplus H_{\tilde{P}}(A, M),$$

with H and F as defined in resp. (9) and (10). Then

$$\text{EPWC}[\tilde{P}](N, A, M) = \tilde{P}^{4,0}(H'_{\tilde{P}}(N, A, M)).$$

It is easy to adapt the proof that H is xor-universal to show that H' is also xor-universal (hence, in particular, universal, which is all we need here). The remaining of the proof is now standard [Sho04], and we only sketch it. We first replace $\tilde{P}^{4,0}$ in $\text{EPWC}[\tilde{P}]$ by a uniformly random function $\rho : \mathcal{X} \rightarrow \mathcal{X}$, and denote $\text{EPWC}''[\rho, \tilde{P}]$ the resulting construction. By the PRP-PRF switching lemma, A can distinguish $\text{EPWC}[\tilde{P}]$ from $\text{EPWC}''[\rho, \tilde{P}]$ with advantage at most $q^2/(2|\mathcal{X}|)$, and because H' is universal, it can distinguish $\text{EPWC}''[\rho, \tilde{P}]$ from Rand with advantage at most $q^2/(2|\mathcal{X}|)$. The result follows. \square

Theorem 3 (MAC-Security of PWC and EPWC). *Let $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{X} = \{0, 1\}^n$ and $\mathcal{T}' = \{2, \dots, 5\} \times \mathcal{T}$, and assume Conjecture 1. Let \mathbf{B} be a nonce-respecting $(q, 1, \ell, \sigma, t)$ -adversary against the MAC-security of $\text{PWC}[\tilde{E}]$, resp. $\text{EPWC}[\tilde{E}]$, with $\ell \leq |\mathcal{T}'| - 2$. Then there exists an absolute constant C and an adversary \mathbf{B}' , resp. \mathbf{B}'' , against the TPRP-security of \tilde{E} , making at most $\sigma + 2q$, resp. $\sigma + 3q$ oracle queries and running it time at most $t + t_{\text{PWC}}(\sigma)$, resp. $t + t_{\text{EPWC}}(\sigma)$, such that*

$$\begin{aligned} \mathbf{Adv}_{\text{PWC}[\tilde{E}]}^{\text{nMAC}}(\mathbf{B}) &\leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{B}') + \frac{(C+1)q}{|\mathcal{X}|}; \\ \mathbf{Adv}_{\text{EPWC}[\tilde{E}]}^{\text{nMAC}}(\mathbf{B}) &\leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{B}'') + \frac{(C+1)q}{|\mathcal{X}|}. \end{aligned}$$

Proof. The proof is standard and only sketched. Let \mathbf{B} be a nonce-respecting $(q, 1, |\mathcal{T}'| - 2, \sigma, t)$ -adversary against the MAC-security of $\text{PWC}[\tilde{E}]$ or $\text{EPWC}[\tilde{E}]$. Let $W_1 = ((\text{E})\text{PWC}[\tilde{E}]_K, \text{Ver}_K)$ and $W_4 = ((\text{E})\text{PWC}[\tilde{E}]_K, \text{Rej})$ be the two worlds that \mathbf{B} tries to distinguish, with Ver_K and Rej as in Definition 3. Let Win be the event that \mathbf{B} asks a right (verification) query which returns 1. Clearly, \mathbf{B} 's distinguishing advantage is equal to the probability of event Win when \mathbf{B} interacts with W_1 , which we now upper bound. As in the proof of Theorem 2, we first replace all calls to \tilde{E}_K made by the two oracles in W_1 by a call to a uniformly random tweakable permutation \tilde{P} , and let (with obvious simplification to the notation) $W_2 = ((\text{E})\text{PWC}[\tilde{P}], \text{Ver}_{\tilde{P}})$. In W_3 , we further replace the pseudorandom function $F_{\tilde{P}}$ as defined by (10) by a uniformly random function $\rho : \mathcal{N} \rightarrow \mathcal{X}$. Then, depending on whether we are considering PWC or EPWC, and using Conjecture 1, one has

$$\begin{aligned} \Pr[\mathbf{B}^{W_1} : \text{Win}] &\leq |\Pr[\mathbf{B}^{W_1} : \text{Win}] - \Pr[\mathbf{B}^{W_2} : \text{Win}]| \\ &\quad + |\Pr[\mathbf{B}^{W_2} : \text{Win}] - \Pr[\mathbf{B}^{W_3} : \text{Win}]| + \Pr[\mathbf{B}^{W_3} : \text{Win}] \\ &\leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(\mathbf{B}'/\mathbf{B}'') + \frac{Cq}{|\mathcal{X}|} + \Pr[\mathbf{B}^{W_3} : \text{Win}], \end{aligned}$$

with \mathbf{B}' and \mathbf{B}'' as in the statement of the theorem. In W_3 , \mathbf{B} is trying to forge a MAC for

$$\text{PWC}'[\rho, \tilde{P}](N, A, M) = \rho(N) \oplus H_{\tilde{P}}(A, M), \quad (14)$$

$$\text{or EPWC}'[\rho, \tilde{P}](N, A, M) = \tilde{P}^{4,0}(\rho(N) \oplus H_{\tilde{P}}(A, M)), \quad (15)$$

where H is defined by (9) and was shown to be xor-universal in the proof of Theorem 2. The classical result on Wegman-Carter MACs [Sho96] applies (the additional encryption layer for EPWC does not modify it), and gives

$$\Pr[\mathbf{B}^{W_3} : \text{Win}] \leq \frac{q}{|\mathcal{X}|}.$$

This concludes the proof. \square

Remark 1. While it is in principle possible to save one encryption call in the EPWC construction by keeping the final AD or message block unencrypted as in the standard PMAC construction [BR02, Rog04a], we avoid this to ensure that static AD always gets treated the same, independently of the message. Indeed, applying this optimization would result in a construction where the final block of AD should be treated differently depending on whether the message is empty or not. Handling the AD independently of the message allows to precompute

$$\text{auth}' = \bigoplus_{i=1}^{\ell_a-1} \tilde{P}^{2,i+1}(A_i) \oplus \tilde{P}^{2/3,\ell_a+1}(A_{\ell_a}(10^*))$$

and to process the nonce and the message later (in particular, when the AD is static, auth' need not be recomputed each time).

Remark 2. In the nonce-misuse scenario, there is a simple birthday attack against EPWC $[\tilde{E}]$ as soon as the adversary can repeat a single nonce twice. The attack proceeds as follows: simply query EPWC $[\tilde{E}]_K$ for roughly $2^{n/2}$ pairs (N_i, A, M) with distinct nonces and the same AD and message until a collision occurs on the outputs for two nonces N_1 and N_2 . Clearly, a collision on the MACs implies that $F_K(N_1) = F_K(N_2)$ (where F_K is given by (5)). Hence, the adversary can now query $Y = \text{EPWC}[\tilde{E}]_K(N_1, A', M')$ for a new pair $(A', M') \neq (A, M)$. Then Y is a valid forgery for (N_2, A', M') . It remains an open problem to design a nonce-based MAC scheme ensuring graceful degradation of security with the maximal number of nonce repetitions.


6 The SCT Mode

6.1 The NSIV Construction

In this section, we present the nAE mode SCT and analyze its security. We first describe a generic composition method named NSIV, which defines a nAE scheme from a nonce-based keyed function and an nivE scheme. The NSIV construction results from a small (but important from a security viewpoint) modification to the (generic) SIV construction [RS06]. While in SIV the encryption part is purely IV-based, NSIV relies on a combined nonce- and IV-based encryption (nivE) scheme, the nonce being used as input *both* to the keyed function and the nivE scheme. This is the only difference with SIV, where the nonce is only given as input to the keyed function.

More formally, let F be a nonce-based keyed function with key-space \mathcal{K}_1 , nonce space \mathcal{N} , domain $\mathcal{D} = \mathcal{A} \times \mathcal{M}$, and range \mathcal{Y} , and $\Pi = (\mathcal{K}_2, \mathcal{N}, \mathcal{IV}, \mathcal{M}, \text{Enc}, \text{Dec})$ be a nivE scheme. Fix a regular function $\text{Conv} : \mathcal{Y} \rightarrow \mathcal{IV}$. We define the nAE scheme¹⁶ NSIV $[F, \Pi] = (\mathcal{K}, \mathcal{N}, \mathcal{A}, \mathcal{M}, \text{Enc}, \text{Dec})$ with key-space $\mathcal{K} = \mathcal{K}_1 \times \mathcal{K}_2$ as specified on Fig. 4.

¹⁶ Our formalization of an nAE scheme in Definition 4 assumes that the ciphertext is a binary string, whereas in our description, NSIV $[F, \Pi].\text{Enc}$ returns a pair (C, tag) . We assume some implicit encoding of this pair into a single binary string.



```

1 algorithm NSIV[F, \Pi].Enc_{K_1, K_2}(N, A, M)
2   tag := F_{K_1}(N, A, M)
3   IV := Conv(tag)
4   C := \Pi.Enc_{K_2}(N, IV, M)
5   return (C, tag)
6
7 algorithm NSIV[F, \Pi].Dec_{K_1, K_2}(N, A, C, tag)
8   IV := Conv(tag)
9   M := \Pi.Dec_{K_2}(N, IV, C)
10  tag' := F_{K_1}(N, A, M)
11  if tag' = tag then return M else return \perp

```

Fig. 4. The NSIV construction, defining a nAE scheme from a nonce-based keyed function $F : \mathcal{K}_1 \times \mathcal{N} \times \mathcal{D} \rightarrow \mathcal{Y}$ where $\mathcal{D} = \mathcal{A} \times \mathcal{M}$ and a nivE scheme $\Pi = (\mathcal{K}_2, \mathcal{N}, \mathcal{IV}, \mathcal{M}, \text{Enc}, \text{Dec})$. Function Conv is a regular function from \mathcal{Y} to \mathcal{IV} .

The security of NSIV[F, \Pi] is given by Theorem 4 below. We assume that $\mathcal{A} = \mathcal{M} = \{0, 1\}^*$ for convenience, but this restriction can be lifted easily. We denote by $t_\Pi(\sigma)$ an upper bound on the time needed for computing $\Pi.\text{Enc}$ or $\Pi.\text{Dec}$ on inputs of total message length at most σ blocks of n bits, and we assume that computing $\text{Conv}(\text{tag})$ or sampling uniformly from $\text{Conv}^{-1}(\text{IV})$ takes negligible time for any $\text{tag} \in \mathcal{Y}$ and $\text{IV} \in \mathcal{IV}$. The proof of this theorem is similar to the security proof of SIV, and deferred to Appendix B.

Theorem 4 (Security of NSIV). *Let $F : \mathcal{K}_1 \times \mathcal{N} \times \mathcal{D} \rightarrow \mathcal{Y}$, where $\mathcal{D} = \mathcal{A} \times \mathcal{M}$, be a nonce-based keyed function, $\Pi = (\mathcal{K}_2, \mathcal{N}, \mathcal{IV}, \mathcal{M}, \text{Enc}, \text{Dec})$ be a nivE scheme, and $\text{Conv} : \mathcal{Y} \rightarrow \mathcal{IV}$ be a regular function. Let \mathbf{A} be a (q, m, ℓ, σ, t) -adversary against NSIV[F, \Pi]. Then, letting $t' = t + t_\Pi(\sigma)$, the following holds:*

- (a) *if \mathbf{A} is allowed to repeat nonces ($m > 1$), then there exists a (q, m, ℓ, σ, t') -adversary \mathbf{A}' against Π and a (q, q, ℓ, σ, t') -adversary \mathbf{A}'' against the PRF-security of F such that*

$$\text{Adv}_{\text{NSIV}[F, \Pi]}^{\text{AE}}(\mathbf{A}) \leq \text{Adv}_{\Pi}^{\text{ivE}}(\mathbf{A}') + \text{Adv}_F^{\text{PRF}}(\mathbf{A}'') + \frac{q}{|\mathcal{Y}|};$$

- (b) if A is nonce-respecting ($m = 1$), then there exists a $(q, 1, \ell, \sigma, t')$ -adversary A' against Π and $(q, 1, \ell, \sigma, t')$ -adversaries A'' and A''' against respectively the PRF- and the MAC-security of F , all nonce-respecting, such that

$$\mathbf{Adv}_{\text{NSIV}[F, \Pi]}^{\text{nAE}}(A) \leq \mathbf{Adv}_{\Pi}^{\text{nivE}}(A') + \mathbf{Adv}_F^{\text{nPRF}}(A'') + \mathbf{Adv}_F^{\text{nMAC}}(A''').$$

6.2 From NSIV to SCT

The $\text{SCT}[\tilde{E}]$ mode is simply $\text{NSIV}[F, \Pi]$ where F is instantiated with $\text{EPWC}[\tilde{E}]$ and Π is instantiated with $\text{CTRTR}[\tilde{E}]$. Additionally, in order to be able to use the same key for calls to \tilde{E} both in EPWC and in CTRTR, we use tweak separation to ensure that all calls to \tilde{E} in EPWC and in CTRTR are independent. The resulting construction is illustrated in Fig. 1. Combining Theorem 4 with Theorems 1, 2, and 3, we finally obtain the following result for the security of SCT.¹⁷ We denote by $t_{\text{SCT}}(\sigma)$ an upper bound on the time needed for computing $\text{SCT}[\tilde{E}].\text{Enc}_K$ or $\text{SCT}[\tilde{E}].\text{Dec}_K$ on inputs of total (AD + message) length at most σ blocks of n bits when calls to \tilde{E}_K cost unit time.

Theorem 5 (Security of SCT). *Let $\tilde{E} \in \text{TBC}(\mathcal{K}, \mathcal{T}', \mathcal{X})$ with $\mathcal{X} = \{0, 1\}^n$, $\mathcal{T} = \{1, \dots, 5\} \times \mathcal{T}'$, and $|\mathcal{T}| \geq 8$. Let Conv be a regular¹⁸ function from \mathcal{X} to \mathcal{T} . Assume Conjecture 1 and let $f(\sigma)$ be defined as in Theorem 1. Let A be a (q, m, ℓ, σ, t) -adversary against $\text{SCT}[\tilde{E}]$ with $\ell \leq |\mathcal{T}| - 2$. Then there exists an absolute constant C and an adversary A' against the TPRP-security of \tilde{E} , making at most $\sigma + 3q$ oracle queries and running in time at most $t + t_{\text{SCT}}(\sigma)$, such that*

- (a) if A is allowed to repeat nonces in encryption queries ($m > 1$), then

$$\mathbf{Adv}_{\text{SCT}[\tilde{E}]}^{\text{AE}}(A) \leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(A') + \frac{2(m-1)\sigma + 1}{|\mathcal{T}|} + f(\sigma) + \frac{q^2 + q}{|\mathcal{X}|};$$

- (b) if A is nonce-respecting ($m = 1$), then

$$\mathbf{Adv}_{\text{SCT}[\tilde{E}]}^{\text{nAE}}(A) \leq \mathbf{Adv}_{\tilde{E}}^{\text{TPRP}}(A') + \frac{1}{|\mathcal{T}|} + f(\sigma) + \frac{(2C+1)q}{|\mathcal{X}|}.$$

Acknowledgements

The authors would like to thank Jérémy Jean and Ivica Nikolic for their remarks on early designs. The first author is supported by the Singapore National Research Foundation Fellowship 2012 (NRF-NRFF2012-06).

¹⁷ In more details, it is more convenient to prove Theorem 5 by first replacing \tilde{E}_K by a uniformly random tweakable permutation, and then applying Theorems 1, 2, and 3 for a perfect TBC.

¹⁸ Note that this regularity condition imposes $|\mathcal{T}| \leq |\mathcal{X}|$. However, when $|\mathcal{T}| > |\mathcal{X}|$, the security bounds of CTRTR and EPWC do not depend on the tweak length (only the maximal message length does). Hence, one can always use a subset of tweaks of size $|\mathcal{X}|$ in case $|\mathcal{T}| > |\mathcal{X}|$.

References

- [ABL⁺13] Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. Parallelizable and Authenticated Online Ciphers. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 (Proceedings, Part I)*, volume 8269 of *LNCS*, pages 424–443. Springer, 2013.
- [ABL⁺14] Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. How to Securely Release Unverified Plaintext in Authenticated Encryption. In *Advances in Cryptology - ASIACRYPT 2014 (Proceedings, Part I)*, volume 8873 of *LNCS*, pages 105–125. Springer, 2014.
- [AFF⁺14] Farzaneh Abed, Scott R. Fluhrer, Christian Forler, Eik List, Stefan Lucks, David A. McGrew, and Jakob Wenzel. Pipelineable On-line Encryption. In Carlos Cid and Christian Rechberger, editors, *Fast Software Encryption - FSE 2014*, volume 8540 of *LNCS*, pages 205–223. Springer, 2014.
- [ANSS16] Gilad Asharov, Moni Naor, Gil Segev, and Ido Shahaf. Searchable Symmetric Encryption: Optimal Locality in Linear Space via Two-Dimensional Balanced Allocations. IACR Cryptology ePrint Archive, Report 2016/251, 2016. To appear at STOC 2016. Preprint available at <http://eprint.iacr.org/2016/251>.
- [AP13] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky Thirteen: Breaking the TLS and DTLS Record Protocols. In *2013 IEEE Symposium on Security and Privacy - SP 2013*, pages 526–540. IEEE Computer Society, 2013.
- [BDJR97] Mihir Bellare, Anand Desai, E. Jorjipii, and Phillip Rogaway. A Concrete Security Treatment of Symmetric Encryption. In *Symposium on Foundations of Computer Science - FOCS '97*, pages 394–403. IEEE Computer Society, 1997.
- [BI99] Mihir Bellare and Russell Impagliazzo. A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. IACR Cryptology ePrint Archive, Report 1999/024, 1999. Available at <http://eprint.iacr.org/1999/024>.
- [BKR00] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The Security of the Cipher Block Chaining Message Authentication Code. *Journal of Computer and System Sciences*, 61(3):362–399, 2000.
- [BN00] Mihir Bellare and Chanathip Namprempre. Authenticated Encryption: Relations among Notions and Analysis of the Generic Composition Paradigm. In Tatsuaki Okamoto, editor, *Advances in Cryptology - ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 531–545. Springer, 2000.
- [BR00] Mihir Bellare and Phillip Rogaway. Encode-Then-Encipher Encryption: How to Exploit Nonces or Redundancy in Plaintexts for Efficient Cryptography. In Tatsuaki Okamoto, editor, *Advances in Cryptology - ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 317–330. Springer, 2000.
- [BR02] John Black and Phillip Rogaway. A Block-Cipher Mode of Operation for Parallelizable Message Authentication. In Lars R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 384–397. Springer, 2002.
- [Bra82] Gilles Brassard. On Computationally Secure Authentication Tags Requiring Short Secret Shared Keys. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology - CRYPTO '82*, pages 79–86. Plenum Press, New York, 1982.

- [BRW04] Mihir Bellare, Phillip Rogaway, and David Wagner. The EAX Mode of Operation. In Bimal K. Roy and Willi Meier, editors, *Fast Software Encryption - FSE 2004*, volume 3017 of *LNCS*, pages 389–407. Springer, 2004.
- [CAE] CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. See <http://competitions.cr.yp.to/caesar.html>.
- [CLP14] Benoît Cogliati, Rodolphe Lampe, and Jacques Patarin. The Indistinguishability of the XOR of k Permutations. In Carlos Cid and Christian Rechberger, editors, *Fast Software Encryption - FSE 2014*, volume 8540 of *LNCS*, pages 285–302. Springer, 2014.
- [CS14] Shan Chen and John Steinberger. Tight Security Bounds for Key-Alternating Ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, 2014. Full version available at <http://eprint.iacr.org/2013/222>.
- [DR11] Thai Duong and Juliano Rizzo. Here Come The \oplus Ninjas. Unpublished manuscript, 2011. Available at <https://bug665814.bugzilla.mozilla.org/attachment.cgi?id=540839>.
- [DS11] Yevgeniy Dodis and John P. Steinberger. Domain Extension for MACs Beyond the Birthday Barrier. In Kenneth G. Paterson, editor, *Advances in Cryptology - EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 323–342. Springer, 2011.
- [Fer02] Niels Ferguson. Collision Attacks on OCB. Unpublished manuscript, 2002. Available at <http://www.cs.ucdavis.edu/~rogaway/ocb/fe02.pdf>.
- [FFL12] Ewan Fleischmann, Christian Forler, and Stefan Lucks. McOE: A Family of Almost Foolproof On-Line Authenticated Encryption Schemes. In Anne Canteaut, editor, *Fast Software Encryption - FSE 2012*, volume 7549 of *LNCS*, pages 196–215. Springer, 2012.
- [FLS⁺10] Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, and Jesse Walker. The Skein Hash Function Family. SHA3 Submission to NIST (Round 3), 2010.
- [GD01] Virgil D. Gligor and Pompiliu Donescu. Fast Encryption and Authentication: XCBC Encryption and XECB Authentication Modes. In Mitsuru Matsui, editor, *Fast Software Encryption - FSE 2001*, volume 2355 of *LNCS*, pages 92–108. Springer, 2001.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GL15] Shay Gueron and Yehuda Lindell. GCM-SIV: Full Nonce Misuse-Resistant Authenticated Encryption at Under One Cycle per Byte. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM Conference on Computer and Communications Security - CCS 2015*, pages 109–119. ACM, 2015.
- [GLS⁺14] Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, Kerem Varici, François Durvaux, Lubos Gaspar, and Stéphanie Kerckhof. SCREAM and iSCREAM. Submitted to the CAESAR competition, 2014.
- [HKR15] Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. Robust Authenticated-Encryption: AEZ and the Problem That It Solves. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 (Proceedings, Part I)*, volume 9056 of *LNCS*, pages 15–44. Springer, 2015.

- [HRRV15] Viet Tung Hoang, Reza Reyhanitabar, Phillip Rogaway, and Damian Vizár. Online Authenticated-Encryption and its Nonce-Reuse Misuse-Resistance. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 (Proceedings, Part I)*, volume 9215 of *LNCS*, pages 493–517. Springer, 2015.
- [Iwa06] Tetsu Iwata. New Blockcipher Modes of Operation with Beyond the Birthday Bound Security. In Matthew J. B. Robshaw, editor, *Fast Software Encryption - FSE 2006*, volume 4047 of *LNCS*, pages 310–327. Springer, 2006.
- [Iwa08] Tetsu Iwata. Authenticated Encryption Mode for Beyond the Birthday Bound Security. In Serge Vaudenay, editor, *Progress in Cryptology - AFRICACRYPT 2008*, volume 5023 of *LNCS*, pages 125–142. Springer, 2008.
- [JNP14a] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Deoxys v1. Submitted to the CAESAR competition, 2014.
- [JNP14b] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Joltik v1. Submitted to the CAESAR competition, 2014.
- [JNP14c] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. KIASU v1. Submitted to the CAESAR competition, 2014.
- [JNP14d] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Tweaks and Keys for Block Ciphers: The TWEAKEY Framework. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 (Proceedings, Part II)*, volume 8874 of *LNCS*, pages 274–288. Springer, 2014.
- [Jut01] Charanjit S. Jutla. Encryption Modes with Almost Free Message Integrity. In Birgit Pfitzmann, editor, *Advances in Cryptology - EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 529–544. Springer, 2001.
- [Jut08] Charanjit S. Jutla. Encryption Modes with Almost Free Message Integrity. *J. Cryptology*, 21(4):547–578, 2008. Earlier version at EUROCRYPT 2001.
- [KR11] Ted Krovetz and Phillip Rogaway. The Software Performance of Authenticated-Encryption Modes. In Antoine Joux, editor, *Fast Software Encryption - FSE 2011*, volume 6733 of *LNCS*, pages 306–327. Springer, 2011.
- [Kra01] Hugo Krawczyk. The Order of Encryption and Authentication for Protecting Communications (or: How Secure Is SSL?). In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001*, volume 2139 of *LNCS*, pages 310–331. Springer, 2001.
- [KY06] Jonathan Katz and Moti Yung. Characterization of Security Notions for Probabilistic Private-Key Encryption. *J. Cryptology*, 19(1):67–95, 2006. Preliminary version in *Proceedings of STOC 2000*.
- [LRW02] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable Block Ciphers. In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002*, volume 2442 of *LNCS*, pages 31–46. Springer, 2002.
- [LS13] Rodolphe Lampe and Yannick Seurin. Tweakable Blockciphers with Asymptotically Optimal Security. In Shiho Moriai, editor, *Fast Software Encryption - FSE 2013*, volume 8424 of *LNCS*, pages 133–151. Springer, 2013.
- [LST12] Will Landecker, Thomas Shrimpton, and R. Seth Terashima. Tweakable Blockciphers with Beyond Birthday-Bound Security. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012*, volume 7417 of *LNCS*, pages 14–30. Springer, 2012. Full version available at <http://eprint.iacr.org/2012/450>.

- [Luc00] Stefan Lucks. The Sum of PRPs Is a Secure PRF. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 470–484. Springer, 2000.
- [Men15] Bart Mennink. Optimally Secure Tweakable Blockciphers. In Gregor Leander, editor, *Fast Software Encryption - FSE 2015*, volume 9054 of *LNCS*, pages 428–448. Springer, 2015. Full version available at <http://eprint.iacr.org/2015/363>.
- [Min09] Kazuhiko Minematsu. Beyond-Birthday-Bound Security Based on Tweakable Block Cipher. In Orr Dunkelman, editor, *Fast Software Encryption - FSE 2009*, volume 5665 of *LNCS*, pages 308–326. Springer, 2009.
- [MV04] David A. McGrew and John Viega. The Security and Performance of the Galois/Counter Mode (GCM) of Operation. In Anne Canteaut and Kapalee Viswanathan, editors, *Progress in Cryptology - INDOCRYPT 2004*, volume 3348 of *LNCS*, pages 343–355. Springer, 2004.
- [NRS14] Chanathip Namprempre, Phillip Rogaway, and Thomas Shrimpton. Reconsidering Generic Composition. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, 2014.
- [Pat08a] Jacques Patarin. A Proof of Security in $O(2^n)$ for the Xor of Two Random Permutations. In Reihaneh Safavi-Naini, editor, *Information Theoretic Security - ICITS 2008*, volume 5155 of *LNCS*, pages 232–248. Springer, 2008. Full version available at <http://eprint.iacr.org/2008/010>.
- [Pat08b] Jacques Patarin. The “Coefficients H” Technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography - SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, 2008.
- [Pat13] Jacques Patarin. Security in $O(2^n)$ for the Xor of Two Random Permutations: Proof with the Standard H Technique. IACR Cryptology ePrint Archive, Report 2013/368, 2013. Available at <http://eprint.iacr.org/2013/368>.
- [RBB03] Phillip Rogaway, Mihir Bellare, and John Black. OCB: A block-cipher mode of operation for efficient authenticated encryption. *ACM Trans. Inf. Syst. Secur.*, 6(3):365–403, 2003.
- [Rog02] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM Conference on Computer and Communications Security - CCS 2002*, pages 98–107. ACM, 2002.
- [Rog04a] Phillip Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In Pil Joong Lee, editor, *Advances in Cryptology - ASIACRYPT 2004*, volume 3329 of *LNCS*, pages 16–31. Springer, 2004.
- [Rog04b] Phillip Rogaway. Nonce-Based Symmetric Encryption. In Bimal K. Roy and Willi Meier, editors, *Fast Software Encryption - FSE 2004*, volume 3017 of *LNCS*, pages 348–359. Springer, 2004.
- [RS06] Phillip Rogaway and Thomas Shrimpton. A Provable-Security Treatment of the Key-Wrap Problem. In Serge Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 373–390. Springer, 2006.
- [Sho96] Victor Shoup. On Fast and Provably Secure Message Authentication Based on Universal Hashing. In Neal Koblitz, editor, *Advances in Cryptology - CRYPTO '96*, volume 1109 of *LNCS*, pages 313–328. Springer, 1996.

- [Sho04] Victor Shoup. Sequences of Games: A Tool for Taming Complexity in Security Proofs. IACR Cryptology ePrint Archive, Report 2004/332, 2004. Available at <http://eprint.iacr.org/2004/332.pdf>.
- [ST13] Thomas Shrimpton and R. Seth Terashima. A Modular Framework for Building Variable-Input-Length Tweakable Ciphers. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013 (Proceedings, Part I)*, volume 8269 of *LNCS*, pages 405–423. Springer, 2013.
- [Vau02] Serge Vaudenay. Security Flaws Induced by CBC Padding - Applications to SSL, IPSEC, WTLS ... In Lars R. Knudsen, editor, *Advances in Cryptology - EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 534–546. Springer, 2002.
- [WC81] Mark N. Wegman and Larry Carter. New Hash Functions and Their Use in Authentication and Set Equality. *J. Comput. Syst. Sci.*, 22(3):265–279, 1981.
- [WHF02] Doug Whiting, Russ Housley, and Niels Ferguson. Counter with CBC-MAC (CCM). Submission to NIST, 2002. Available at <http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/proposedmodes/ccm/ccm.pdf>.
- [Yas11] Kan Yasuda. A New Variant of PMAC: Beyond the Birthday Bound. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *LNCS*, pages 596–609. Springer, 2011.

A Proof of Lemma 5

We restate and prove Lemma 5.

Lemma. Consider a set of $|\mathcal{T}| \geq 8$ bins and $\sigma \geq 8$ balls. Fix an integer $q \leq \sigma$ and a sequence of integers (ℓ_1, \dots, ℓ_q) with $1 \leq \ell_i \leq |\mathcal{T}|$ and $\sum_{i=1}^q \ell_i = \sigma$. Consider the following random process: for $i = 1, \dots, q$, a chain of ℓ_i balls is thrown in consecutive bins, the initial bin being chosen independently uniformly at random. Then the probability that, at the end of the process, any bin contains L_{\max} balls or more, is less than $1/|\mathcal{T}|$, where

- (a) $L_{\max} = 2 \log \sigma$ when $\sigma \leq |\mathcal{T}|$;
- (b) $L_{\max} = \frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|}$ when $\sigma \geq |\mathcal{T}|$.

Proof. For each bin $T \in \mathcal{T}$, let $L(T)$ be the load of bin T , i.e., the number of balls it contains at the end of the process, and let us write $L(T)$ as the sum of independent variables $L_i(T)$, $1 \leq i \leq q$, where $L_i(T) = 1$ if bin T contains a ball from the i -th chain and $L_i(T) = 0$ otherwise. The probability that bin T contains one of the balls of the i -th chain is exactly $\ell_i/|\mathcal{T}|$. Hence, the expected value of $L(T)$ is

$$\mu = \sum_{i=1}^q \frac{\ell_i}{|\mathcal{T}|} = \frac{\sigma}{|\mathcal{T}|}.$$

Moreover, $L(T)$ being the sum of independent random variables, the multiplicative Chernoff bound gives us that for any $\delta > 0$,

$$\Pr [L(T) \geq (1 + \delta)\mu] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu,$$

or equivalently, letting $\alpha = 1 + \delta$,

$$\Pr [L(T) \geq \alpha\mu] \leq \left(\frac{e^{\alpha-1}}{\alpha^\alpha}\right)^\mu < \left(\frac{e}{\alpha}\right)^{\alpha\mu}.$$

We first prove part (a) of the lemma, i.e., assuming $\sigma \leq |\mathcal{T}|$. Taking $\alpha = 2|\mathcal{T}|\log\sigma/\sigma$ yields

$$\begin{aligned} \Pr [L(T) \geq 2\log\sigma] &\leq \left(\frac{e\sigma}{2|\mathcal{T}|\log\sigma}\right)^{2\log\sigma} \\ &= \frac{\sigma^2}{|\mathcal{T}|^2} \left(\frac{e}{2\log\sigma}\right)^{2\log\sigma} \left(\frac{\sigma}{|\mathcal{T}|}\right)^{2\log\sigma-2} \\ &\leq \frac{\sigma^2}{|\mathcal{T}|^2} \left(\frac{1}{2}\right)^{2\log\sigma} \left(\frac{\sigma}{|\mathcal{T}|}\right)^{2\log\sigma-2} && (\sigma \geq 8) \\ &= \frac{1}{|\mathcal{T}|^2} \left(\frac{\sigma}{|\mathcal{T}|}\right)^{2\log\sigma-2} \\ &\leq \frac{1}{|\mathcal{T}|^2} && (\sigma \leq |\mathcal{T}|). \end{aligned}$$

The result follows by a union bound over $T \in \mathcal{T}$.

We then prove part (b) of the result, assuming $\sigma \geq |\mathcal{T}|$. This time, taking $\alpha = 2\log|\mathcal{T}|$ yields

$$\begin{aligned} \Pr \left[L(T) \leq \frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|} \right] &\leq \left(\frac{e}{2\log|\mathcal{T}|}\right)^{\frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|}} \\ &\leq \left(\frac{1}{2}\right)^{\frac{2\sigma \log |\mathcal{T}|}{|\mathcal{T}|}} && (|\mathcal{T}| \geq 8) \\ &\leq \left(\frac{1}{2}\right)^{2\log|\mathcal{T}|} && (|\mathcal{T}| \leq \sigma) \\ &= \frac{1}{|\mathcal{T}|^2}. \end{aligned}$$

Again, the result follows by a union bound over $T \in \mathcal{T}$. □

B Proof of Theorem 4

Let $\tilde{\Pi} = \text{NSIV}[F, \Pi]$ and let \mathbf{A} be a (q, m, ℓ, σ, t) -adversary against $\tilde{\Pi}$ (we do not specify whether \mathbf{A} is nonce-respecting or not yet). We denote

$$\begin{aligned} W_1 &= (\tilde{\Pi}.\text{Enc}_{K_1, K_2}, \tilde{\Pi}.\text{Dec}_{K_1, K_2}), \\ W_4 &= (\text{Rand}, \text{Rej}) \end{aligned}$$

the two worlds that \mathbf{A} must tell apart, where $\text{Rand}(N, A, M)$ returns a random string of length $|\tilde{\Pi}.\text{Enc}_{K_1, K_2}(N, A, M)|$ and Rej always returns \perp . We introduce

intermediate worlds W_2^{AE} , W_2^{nAE} , and W_3 . As the names suggest, W_2^{AE} , resp. W_2^{nAE} , will be used for nonce-misuse, resp. nonce-respecting adversaries, while W_3 will be used in both cases. In W_2^{AE} , we replace F_{K_1} by a uniformly random function $\rho : \mathcal{N} \times \mathcal{D} \rightarrow \mathcal{Y}$ in both the encryption and decryption algorithms. Formally, letting $\tilde{\Pi}' = \text{NSIV}[\text{Func}(\mathcal{N} \times \mathcal{D}, \mathcal{Y}), \Pi]$, we define

$$W_2^{\text{AE}} = (\tilde{\Pi}'.\text{Enc}_{\rho, K_2}, \tilde{\Pi}'.\text{Dec}_{\rho, K_2}).$$

In W_2^{nAE} , on the other hand, we replace the decryption algorithm of W_1 by an always-rejecting oracle, i.e.,

$$W_2^{\text{nAE}} = (\tilde{\Pi}.\text{Enc}_{K_1, K_2}, \text{Rej}).$$

Finally, world W_3 is defined as

$$W_3 = (\tilde{\Pi}'.\text{Enc}_{\rho, K_2}, \text{Rej}).$$

We first consider A 's advantage in distinguishing W_3 from W_4 . Let

$$\delta_{3,4} = |\Pr[A^{W_3} = 1] - \Pr[A^{W_4} = 1]|.$$

Consider the adversary A' against the nivE-security of Π working as follows. Let $G \in \{\Pi.\text{Enc}_{K_2}^{\S}, \text{Rand}'\}$ be the oracle to which A' has access, where $\text{Rand}'(N, M)$ returns a random string of length $|\Pi.\text{Enc}_{K_2}^{\S}(N, M)|$. Adversary A' runs A . When A asks an encryption query (N, A, M) , A' queries $G(N, M)$, thereby obtaining an answer (IV, C) . It samples uniformly at random tag in the set of preimages $\text{Conv}^{-1}(IV)$ and returns (C, tag) to A . When A asks a decryption query, A' simply answers \perp . When A halts and outputs some bit, A' outputs the same bit. Note that independently of whether G is $\Pi.\text{Enc}_{K_2}^{\S}$ or Rand' , the first element IV in its answers is uniformly random, and hence tag is uniformly random as well since Conv is regular. It follows that A' perfectly simulates W_3 when $G = \Pi.\text{Enc}_{K_2}^{\S}$, and perfectly simulates W_4 when $G = \text{Rand}'$. Moreover, it is easy to see that A' is a (q, m, ℓ, σ, t') -adversary. In particular, A' is nonce-respecting iff A is nonce-respecting. Hence,

$$\delta_{3,4} \leq \mathbf{Adv}_{\Pi}^{\text{nivE}}(A') \quad \text{or} \quad \delta_{3,4} \leq \mathbf{Adv}_{\Pi}^{\text{ivE}}(A'),$$

depending on whether A is nonce-respecting or not.

We now prove part (a) of the theorem, assuming A is allowed to repeat nonces.

Let

$$\delta_{1,2}^{\text{AE}} = \left| \Pr[A^{W_1} = 1] - \Pr[A^{W_2^{\text{AE}}} = 1] \right|.$$

Let A'' be the following adversary against the PRF-security of F . Denote $G \in \{F_{K_1}, \rho\}$ the oracle to which A'' has access. Adversary A'' picks a random $K_2 \leftarrow_{\S} \mathcal{K}_2$ and runs A . When A makes an encryption query (N, A, M) , or a decryption query (N, A, C, tag) , A'' simply executes the corresponding code of Fig. 4, replacing all calls to F_{K_1} by a call to its oracle G . When A halts and outputs some bit, A'' outputs the same bit. Then A'' perfectly simulates W_1

when $G = F_{K_1}$ for some random key K_1 , and perfectly simulates world W_2 when $G = \rho$ for some random function ρ , so that

$$\delta_{1,2}^{\text{AE}} \leq \mathbf{Adv}_F^{\text{PRF}}(\mathbf{A}'').$$

Moreover, it is easy to see that \mathbf{A}'' is a (q, q, ℓ, σ, t') -adversary.¹⁹

Next, we upper bound

$$\delta_{2,3}^{\text{AE}} = \left| \Pr \left[\mathbf{A}^{W_2^{\text{AE}}} = 1 \right] - \Pr \left[\mathbf{A}^{W_3} = 1 \right] \right|.$$

For this, we simply apply the H-coefficients technique. See W_2^{AE} as the ideal world and W_3 as the real world in the language of Lemma 1, and define a bad transcript as a transcript which contains a decryption query which returned $M \neq \perp$. Consequently, an attainable transcript τ is good if it only contains decryption queries that were answered by \perp . Since the encryption oracles are exactly the same in worlds W_2^{AE} and W_3 , and since the right (decryption) oracle always rejects in world W_3 , it follows that for any good transcript τ , $\Pr[\Theta_{\text{re}} = \tau] \geq \Pr[\Theta_{\text{id}} = \tau]$. It remains to upper bound $\Pr[\Theta_{\text{id}} \in \text{BadT}]$, i.e., the probability that the adversary makes a decryption query that returns $M \neq \perp$ in world W_2^{AE} . Consider a decryption query (N, A, C, tag) , and let $IV = \text{Conv}(\text{tag})$ and $M = \Pi.\text{Dec}_{K_2}(N, IV, C)$. Assume first that there was a previous encryption query (N, A, M) that returned (C, tag') . By our assumption that the adversary never asks a decryption query (N, A, C, tag') if a previous encryption query (N, A, M) returned (C, tag') , one necessarily has $\text{tag} \neq \text{tag}' = \rho(N, A, M)$, so that the decryption oracle necessarily returns \perp . Hence, the adversary is essentially trying to guess a random value $\rho(N, A, M)$ within at most $q(N, A, M)$ trials, where $q(N, A, M)$ is the number of decryption queries of the form (N, A, C, tag) with $C = \Pi.\text{Enc}_{K_2}(N, \text{Conv}(\text{tag}), M)$, something it can achieve with probability at most $q(N, A, M)/|\mathcal{Y}|$. Summing over all possible values of (N, A, M) , we obtain that $\Pr[\Theta_{\text{id}} \in \text{BadT}] \leq q/|\mathcal{Y}|$. Hence, by Lemma 1, we have

$$\delta_{2,3}^{\text{AE}} \leq \frac{q}{|\mathcal{Y}|}.$$

The result follows since $\mathbf{Adv}_{\Pi}^{\text{AE}}(\mathbf{A}) \leq \delta_{1,2}^{\text{AE}} + \delta_{2,3}^{\text{AE}} + \delta_{3,4}$.

Let us now prove part (b) of the theorem, assuming \mathbf{A} is nonce-respecting.

Let

$$\delta_{2,3}^{\text{nAE}} = \left| \Pr \left[\mathbf{A}^{W_2^{\text{nAE}}} = 1 \right] - \Pr \left[\mathbf{A}^{W_3} = 1 \right] \right|.$$

Let \mathbf{A}'' be the following adversary against the nPRF-security of F . Denote $G \in \{F_{K_1}, \rho\}$ the oracle to which \mathbf{A}'' has access. Adversary \mathbf{A}'' samples $K_2 \leftarrow_{\S} \mathcal{K}_2$ and runs \mathbf{A} . When \mathbf{A} makes an encryption query (N, A, M) , \mathbf{A}'' asks query (N, A, M) to its G oracle to get the corresponding tag value, computes $IV = \text{Conv}(\text{tag})$ and

¹⁹ Note that \mathbf{A}'' may repeat nonces in its oracle queries even when \mathbf{A} is nonce-respecting because a nonce-respecting \mathbf{A} is allowed to repeat nonces in *decryption* queries. This is why this transition from W_1 to W_2^{AE} cannot be used for a nonce-respecting adversary.

$C = H.\text{Enc}_{K_2}(N, IV, M)$, and returns (C, tag) . When A asks a decryption query, A'' simply returns \perp . When A halts and returns some output bit, A'' outputs the same bit. It is easy to see that A'' perfectly simulates W_2^{nAE} when $G = F_{K_1}$, while it perfectly simulates W_3 when $G = \rho$. Moreover, it is easy to see that A'' is a $(q, 1, \ell, \sigma, t')$ -adversary and in particular A'' is nonce-respecting because A is nonce-respecting. Hence,

$$\delta_{2,3}^{\text{nAE}} \leq \mathbf{Adv}_F^{\text{nPRF}}(A'').$$

It remains to upper bound

$$\delta_{1,2}^{\text{nAE}} = \left| \Pr[A^{W_1} = 1] - \Pr[A^{W_2^{\text{nAE}}} = 1] \right|.$$

Let A''' be the following adversary against the nMAC-security of F . Denote (F_{K_1}, G) the two oracles to which A''' has access, with $G \in \{\text{Ver}_{K_1}, \text{Rej}\}$. Adversary A''' samples $K_2 \leftarrow_{\S} \mathcal{K}_2$ and runs A . When A makes an encryption query (N, A, M) , A''' asks query (N, A, M) to its F_{K_1} oracle to get the corresponding tag value, computes $IV = \text{Conv}(\text{tag})$ and $C = H.\text{Enc}_{K_2}(N, IV, M)$, and returns (C, tag) . When A asks a decryption query (N, A, C, tag) , A''' computes $IV = \text{Conv}(\text{tag})$ and $M = H.\text{Dec}_{K_2}(N, IV, C)$, asks query (N, A, M, tag) to its oracle G , and returns M if $G(N, A, M, \text{tag}) = 1$ or \perp if $G(N, A, M, \text{tag}) = \perp$. When A halts and returns some output bit, A''' outputs the same bit. It is easy to see that A''' perfectly simulates W_1 when $G = \text{Ver}_{K_1}$, while it perfectly simulates W_2^{nAE} when $G = \text{Rej}$. Moreover, it is easy to see that A''' is a $(q, 1, \ell, \sigma, t')$ -adversary and in particular A''' is nonce-respecting because A is nonce-respecting. Hence,

$$\delta_{1,2}^{\text{nAE}} \leq \mathbf{Adv}_F^{\text{nMAC}}(A''').$$

The result follows since $\mathbf{Adv}_H^{\text{nAE}}(A) \leq \delta_{1,2}^{\text{nAE}} + \delta_{2,3}^{\text{nAE}} + \delta_{3,4}$.

C Variants of the CTRT Mode

A PURE NONCE-BASED VARIANT. From CTRT, we can define a purely nonce-based encryption scheme by simply generating the IV in a pseudorandom way from the nonce N . Namely, if \tilde{E} has tweak space $\mathcal{T}' = \{0, 1\} \times \mathcal{T}$, then let

$$\begin{aligned} IV &= \text{Conv}(\tilde{E}_K^{(0,0)}(N)), \\ C &= \text{CTR}[\tilde{E}].\text{Enc}_K(N, IV, M), \end{aligned}$$

where Conv is a regular function from \mathcal{X} to $\mathcal{IV} = \mathcal{T}$. We denote the resulting mode NCTR, and it is illustrated in Fig. 5 (gray box). This mode will be useful when defining the online scheme CTPWC in Section D.

Clearly, the NCTR mode offers no security against an adversary allowed to repeat nonces. Against nonce-respecting adversaries however, we expect this mode to enjoy a security bound very close to the one of Theorem 1 (b). The only differences are an increase in the number of oracle queries made by adversary A'

(since there is an additional call to \tilde{E}_K for each encryption query), and the fact that the IV's are not uniformly random any more since they are generated by applying Conv to values $\tilde{E}_K^{(0,0)}(N)$ drawn (up to replacing $\tilde{E}_K^{(0,0)}$ by a random permutation) uniformly at random *without replacement*.²⁰ We leave the extension of Lemma 5 to take into account this non-uniform distribution of the IV's to future work.

Remark 3. A more straightforward way to generate a uniformly random IV for each new nonce would be to put the nonce in the tweak, by letting for example $IV = \text{Conv}(\tilde{E}_K^{(0,N)}(0))$. However, this is only possible when the intended nonce space is smaller than the effective tweak space \mathcal{T} of \tilde{E} . Note that, starting from a block cipher with t -bit tweaks, the nonce in that case can be of length at most $t - 1$ bits because of tweak separation.

A PURE IV-BASED VARIANT. We can also define a purely IV-based variant of CTRT. For this, rename IV and N in CTRT respectively IV_1 and IV_2 , and assume that they are chosen uniformly at random at each encryption call. Call the resulting mode IVCTRT. The security of this variant can be easily derived from Theorem 1. Indeed, consider an adversary A making at most q queries of total message length at most σ blocks of n bits. Since the “nonces” (IV_2) are now uniformly random for each encryption query, it is easy to obtain an upper bound on the maximal number m of repetitions of any “nonce” that holds with high probability. For this, one can simply apply Lemma 5, replacing $|\mathcal{T}|$ with $|\mathcal{X}|$, and considering the case where $\sigma = q$ and $\ell_1 = \dots = \ell_q = 1$ (this is just the “standard” balls-into-bins setting). Hence, the probability that, by the end of the q encryption queries of the adversary, any value IV_2 has been repeated more than m times, is less than $1/|\mathcal{X}|$, where

- (a) $m = 2 \log q$ when $q \leq |\mathcal{X}|$;
- (b) $m = \frac{2q \log |\mathcal{X}|}{|\mathcal{X}|}$ when $q \geq |\mathcal{X}|$.

Combining this with Theorem 1 yields the security bound for IVCTRT, depending on the range of q and σ . Namely, when \tilde{E} is ideal, the advantage of A is at most $1/|\mathcal{X}| + 1/|\mathcal{T}| + f(\sigma) + g(q, \sigma)$, where f is defined as in Theorem 1 and

$$g(q, \sigma) = \begin{cases} \frac{4\sigma \log q}{|\mathcal{T}|} & \text{when } 8 \leq q \leq |\mathcal{X}| \\ \frac{4q\sigma \log |\mathcal{X}|}{|\mathcal{X}||\mathcal{T}|} & \text{when } q \geq |\mathcal{X}|. \end{cases}$$

D The Online CTPWC Mode

In this section, we propose a nonce-based, *online* AEAD mode called CTPWC (*Counter-in-Tweak then Parallel Wegman-Carter*). It combines in an “encrypt-

²⁰ One could, of course, replace $\tilde{E}_K^{(0,0)}$ by a random function in the security proof, but this would require the PRP-PRF switching lemma and one would bump into the birthday bound.

then-MAC” manner the NCTRT nonce-based encryption mode described in Appendix C and the nonce-based keyed function PWC of Section 5. More precisely, according to the NRS nomenclature [NRS14], it uses the N2 generic composition method: on input (N, A, M) , compute $C = \text{NCTRT}[\tilde{E}]_K(N, M)$, $\text{tag} = \text{PWC}[\tilde{E}]_K(N, A, C)$, and return (C, tag) . This is illustrated in Fig. 5. The nAE-security of this mode is tightly related to the security (as a nonce-based encryption scheme [NRS14]) of NCTRT (which, as explained in Section C, we conjecture close to the security of CTRT) and the nPRF- and nMAC-security of PWC given in Theorems 2 and 3. We omit the details.




Fig. 5. The CTPWC mode, using a TBC \tilde{E} with tweak space $\{0, \dots, 5\} \times \mathcal{T}$ and domain $\mathcal{X} = \{0, 1\}^n$. It is obtained by applying the “encrypt-then-MAC” composition method to NCTRT and PWC. We recall that $\tilde{E}_K^i(T, X) = \tilde{E}_K((i, T), X)$ and that $\tilde{E}_K^{i/j}$ means that prefix i is used when the input block is complete and unpadded, whereas j is used when the input block is incomplete and padded. Function Inc is a cyclic permutation of \mathcal{T} , and Conv is a regular function from \mathcal{X} to \mathcal{T} (e.g., truncation when $\mathcal{X} = \{0, 1\}^n$ and $\mathcal{T} = \{0, 1\}^w$ with $w \leq n$).