Reconsidering the Security Bound of AES-GCM-SIV

Tetsu Iwata[1] and Yannick Seurin[2]

[1] Nagoya University, Japan tetsu.iwata@nagoya-u.jp
[2] ANSSI, Paris, France yannick.seurin@m4x.org

Abstract. We make a number of remarks about the AES-GCM-SIV nonce-misuse resistant authenticated encryption scheme currently considered for standardization by the Crypto Forum Research Group (CFRG). First, we point out that the security analysis proposed in the ePrint report 2017/168 is incorrect, leading to overly optimistic security claims. We correct the bound and re-assess the security guarantees offered by the scheme for various parameters. Second, we suggest a simple modification to the key derivation function which would improve the security of the scheme with virtually no efficiency penalty.

Keywords: authenticated encryption · AEAD · GCM-SIV · AES-GCM-SIV · CAESAR competition

1 Introduction

Authenticated Encryption. An authenticated encryption scheme aims at providing both confidentiality and authenticity when communicating over an insecure channel. The recent CAESAR competition [CAE] has spawned a lot of candidate schemes as well as more theoretical works on the subject.

One of the most widely deployed AEAD schemes today is GCM [MV04], which combines, in the "encrypt-then-MAC" fashion [BN00], a Wegman-Carter MAC [WC81, Sho96] based on a polynomial hash function called GHASH, and the counter encryption mode [BDJR97]. GCM is nonce-based [Rog04], i.e., for each encryption the sender must provide a non-repeating value N. Unfortunately, the security of GCM becomes very brittle in case the same nonce N is reused (something called nonce-misuse), in particular a simple attack allows to completely break authenticity [Jou06, BZD+16] (damages to confidentiality are to some extent less dramatic [ADL17]).

AES-GCM-SIV. In order to remedy the nonce-misuse problem faced by GCM, Gueron and Lindell [GL15] proposed the GCM-SIV mode. It is based on the same components as GCM (and as such it can benefit from dedicated CPU instructions that were developed to accelerate GCM) but it combines them through the SIV composition method [RS06] which endows the resulting scheme with "nonce-misuse resistance", meaning that repeating a nonce does not affect authenticity of the scheme and only allows an adversary to detect if the same message was already encrypted along with the same nonce before.

Some time later, Gueron, Langley, and Lindell [GLL16] proposed a variant of GCM-SIV called AES-GCM-SIV[1] as a candidate for standardization to the Crypto Forum Research

[1] The name AES-GCM-SIV is somehow a misnomer: ingenuously, one would think that this designates the AEAD *scheme* resulting from instantiating the GCM-SIV *mode of operation* [GL15] with the AES *block cipher*. It is unclear whether the designers of AES-GCM-SIV think of it as a pure mode of operation, as the AEAD scheme resulting from instantiating this mode with AES, or both.

Group (CFRG) of IETF. In order to overcome some limitations of GCM-SIV, this mode slightly differs from the latter in essentially three ways: (i) it uses a variant of GHASH called POLYVAL, (ii) it uses a key derivation function to derive a hashing key and an encryption key from the nonce and the master key, whereas GCM-SIV uses hashing and encryption keys that are independent from the nonce, and (iii) the initial counter consists of the entire pseudorandom tag (except for its most significant bit), whereas in GCM-SIV the 32 least significant bits of the counter are initialized with zeros. The last modification was also suggested under the name GCM-SIV1 by Iwata and Minematsu [IM16]. A security analysis of AES-GCM-SIV was proposed by the designers [GLL17], which covers versions 3 to 5 of the CFRG specification; weaknesses were spotted in version 2 of the specification by the NSA [NSA17], leading to minor changes in the scheme.

Our Findings. We uncover a number of flaws in the security proofs presented in [GL15] and [GLL17] which are serious enough to make the final security bound derived for AES-GCM-SIV in [GLL17] essentially unusable. We give a simple attack that contradicts this security bound, thus making the question of the provable security of AES-GCM-SIV open. In order to fix the situation and correctly gauge AES-GCM-SIV's security, we present a corrected security proof and then turn to the task of interpreting this bound for concrete parameters.

Based on their result, Gueron *et al.* [GLL17] claimed that the security bound of AES-GCM-SIV is dominated by

$$\frac{QR^2}{2^{n-k-2}}, \tag{1}$$

where n is the block length of the underlying block cipher, Q is the number of distinct nonces used throughout encryption queries, R is the maximal number of repetitions of any nonce in encryption queries, and the maximum message length is $2^k - 1$ blocks. This term essentially captures the probability that two counters used for encryption collide, resulting in an immediate break of confidentiality. We show that the corrected bound is actually dominated by

$$\frac{QR^2}{2^{n-2k+1}}, \tag{2}$$

which is roughly 2^k times larger than term (1) and which captures the adversary's advantage in distinguishing the outputs of the underlying block cipher from random (in other words, this is a classical "PRP-PRF switching" term, albeit in the so-called "multi-user" setting). We stress that this bound is tight and matched by a simple PRP-PRF distinguishing attack. For large values of k (which can be up to 32), the difference between (1) and (2) is significant, and many parameters deemed secure in [GLL17] are in fact not secure at all (see Table 1 in Section 3.3). All details can be found in Section 3.

One might be tempted to argue that attacks against the counter encryption mode based on distinguishing the underlying block cipher from a random function through (the absence of) collisions in outputs is much less dangerous than collisions in counters which immediately reveal the xor of two plaintext blocks. However, this is a very dubious and dangerous reasoning, as shown by the following textbook example [Jou09, Sect. 6.1.1.2]. Assume that the adversary knows that a sender will encrypt one out of two possible "unrelated" plaintexts M_0 and M_1 of the same (sufficiently large) length (in blocks) ℓ, and that it intercepts the corresponding ciphertext C. Then it can simply compute $C \oplus M_0$ and $C \oplus M_1$ and look for collisions among blocks of the resulting strings: no collision can occur for the correct plaintext, whereas a collision will occur with probability roughly $\ell^2/2^n$ for the incorrect plaintext. See also [McG12], which shows that this kind of attacks can have a real impact in practice.

Improving the Key Derivation Function. As a secondary contribution, we point out that the key derivation function used in AES-GCM-SIV can be replaced with the "sum of PRPs" construction [Luc00] or a variant of CENC [Iwa06] to improve the security bound without harming efficiency. Details can be found in Section 5.

Recommendations. It is claimed in the abstract of [GLL17] (as well as in the CFRG draft [GLL16, Sect. 9]) that AES-GCM-SIV "allows for encrypting up to 2^{50} messages with the same key" without any precision on the maximal message length. In light of our results, we think that it is necessary to revise the recommended parameters and usage limitations for AES-GCM-SIV. In particular, whether AES-GCM-SIV can securely encrypt more than 2^{32} messages with the same key, which was the limit for both GCM and GCM-SIV and the main reason for designing AES-GCM-SIV in the first place, becomes questionable. If we follow NIST recommendations for GCM [Dwo07, Sect. 9] and pose that the adversary's advantage, dominated by Equation (2), should not exceed 2^{-32}, we see that for $n = 128$, parameters Q, R and k must satisfy

$$QR^2 \leq 2^{97-2k}.$$

Hence, for $k = 32$ and tolerating up to $R = 2^8$ repetitions of any nonce (as suggested in [GLL16, Sect. 9]), the total number Q of distinct nonces in encryptions must be at most 2^{17} (which implies that the total number of encrypted messages cannot be more than $QR = 2^{25}$). In the case where nonces are drawn uniformly and independently at random, the dominating term of the security bound becomes

$$\frac{9N_E}{2^{n-2k+1}},$$

where N_E is the total number of encryptions. Hence, when $n = 128$ and $k = 32$, N_E must be less than approximately 2^{30} for this term to be less than 2^{-32}.

In conclusion, the suitability of the adoption of AES-GCM-SIV as a CFRG standard or its deployment in large-scale protocols such as QUIC [QUI] should be reconsidered based on the new security analysis presented in this paper.

Timeline. This paper is based on ePrint version 20160310:063701 of [GL15] and ePrint version 20170223:140759 of [GLL17]. These two ePrint reports were updated after we sent a preliminary version of this paper to the authors on July 7, 2017.

2 Preliminaries

2.1 Notation and Security Definitions

General Notation. We let $\{0,1\}^n$, $\{0,1\}^*$, and $(\{0,1\}^n)^+$ denote respectively the set of bit strings of length n, the set of all bit strings (including the empty string of length 0), and the set of non-empty tuples of n-bit strings. The length of a bit string x is denoted $|x|$. Given two bit strings x and y, $x||y$ denotes their concatenation. Given a bit string x of length at least m, $\mathsf{Trunc}_m(x)$ denotes the m rightmost bits of x. Given an integer $i \leq 2^a - 1$, we let $[i]_a$ denote its a-bit binary representation. Given a finite non-empty set $\mathcal{X}$, $x \leftarrow_\$ \mathcal{X}$ denotes the sampling of x uniformly at random in $\mathcal{X}$. Given non-empty sets $\mathcal{X}$ and $\mathcal{Y}$, the set of all functions from $\mathcal{X}$ to $\mathcal{Y}$ is denoted $\mathsf{Func}(\mathcal{X}, \mathcal{Y})$, and the set of all permutations of $\mathcal{X}$ is denoted $\mathsf{Perm}(\mathcal{X})$.

PRFs and Block Ciphers. A keyed function with key space $\mathcal{K}$, domain $\mathcal{X}$, and range $\mathcal{Y}$ is a function $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$. We denote $F_K(X)$ for $F(K, X)$. A (q, t)-adversary against F is an algorithm A with oracle access to a function from $\mathcal{X}$ to $\mathcal{Y}$, making at most q oracle queries, running in time at most t, and outputting a single bit. The advantage of A in breaking the PRF-security of F is defined as

$$\mathbf{Adv}^{\mathrm{prf}}_F(\mathsf{A}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathsf{A}^{F_K} = 1\right] - \Pr\left[R \leftarrow_{\$} \mathsf{Func}(\mathcal{X}, \mathcal{Y}) : \mathsf{A}^{R} = 1\right]\right|.$$

A block cipher with key space $\mathcal{K}$ and domain $\mathcal{X}$ is a mapping $E : \mathcal{K} \times \mathcal{X} \to \mathcal{X}$ such that for any key $K \in \mathcal{K}$, $X \mapsto E(K, X)$ is a permutation of $\mathcal{X}$. We denote $E_K(X)$ for $E(K, X)$. A (q, t)-adversary against E is an algorithm A with oracle access to a permutation of $\mathcal{X}$, making at most q oracle queries, running in time at most t, and outputting a single bit. The advantage of A in breaking the PRP-security of E is defined as

$$\mathbf{Adv}^{\mathrm{prp}}_E(\mathsf{A}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathsf{A}^{E_K} = 1\right] - \Pr\left[P \leftarrow_{\$} \mathsf{Perm}(\mathcal{X}) : \mathsf{A}^{P} = 1\right]\right|.$$

IV-based Encryption Schemes. Syntactically, an IV-based encryption (ivE) scheme is a tuple $\Pi = (\mathcal{K}, \mathsf{Enc}, \mathsf{Dec})$ where $\mathcal{K}$ is a non-empty key set and Enc and Dec are deterministic algorithms. The encryption algorithm Enc takes as input a key $K \in \mathcal{K}$, an initial value $IV \in \{0,1\}^{\mathtt{ivl}}$, where $\mathtt{ivl}$ is the IV length, and a message $M \in \{0,1\}^*$, and outputs a ciphertext $C \in \{0,1\}^*$. The decryption algorithm Dec takes as input a key $K \in \mathcal{K}$, an initial value $IV \in \{0,1\}^{\mathtt{ivl}}$, and a ciphertext $C \in \{0,1\}^*$, and outputs a message $M \in \{0,1\}^*$. We require that

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

for all tuples $(K, IV, M) \in \mathcal{K} \times \{0,1\}^{\mathtt{ivl}} \times \{0,1\}^*$.

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

Definition 1 (Security of an ivE scheme)**.** Let $\Pi = (\mathcal{K}, \mathsf{Enc}, \mathsf{Dec})$ be an ivE scheme. The advantage of an adversary A in breaking Π is defined as

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

where $\$(\cdot)$ is an oracle which on input $M \in \{0,1\}^*$ outputs a random string of length $|\Pi.\mathsf{Enc}^{\$}_K(M)|$.

Authenticated Encryption. A nonce-based Authenticated Encryption with Associated Data (AEAD) scheme is a tuple $\Pi = (\mathcal{K}, \mathsf{Enc}, \mathsf{Dec})$ where $\mathcal{K}$ is a non-empty key set and Enc and Dec are deterministic algorithms. The encryption algorithm Enc takes as input a key $K \in \mathcal{K}$, a nonce $N \in \{0,1\}^{\mathtt{nl}}$ where $\mathtt{nl}$ is the nonce-length, associated data (AD) $A \in \{0,1\}^*$, and a message $M \in \{0,1\}^*$, and returns a string $Y \in \{0,1\}^*$. The decryption algorithm Dec takes as input a key $K \in \mathcal{K}$, a nonce $N \in \{0,1\}^{\mathtt{nl}}$, associated data $A \in \{0,1\}^*$, and a string $Y \in \{0,1\}^*$, and returns either a message $M \in \{0,1\}^*$ or a special value $\perp$ indicating that inputs are invalid. We write $\mathsf{Enc}_K(N, A, M)$ for $\mathsf{Enc}(K, N, A, M)$ and $\mathsf{Dec}_K(N, A, Y)$ for $\mathsf{Dec}(K, N, A, Y)$.

For many AEAD schemes (in particular for AES-GCM-SIV), any non-$\perp$ output Y of the encryption algorithm consists of the concatenation of a ciphertext C of the same size as the message M and a tag $T \in \{0,1\}^{\mathtt{tl}}$ where $\mathtt{tl}$ is the tag-length.

We use the following Misuse Resistant Authenticated Encryption (MRAE) security notion [RS06].

Definition 2 (MRAE-security)**.** Let $\Pi = (\mathcal{K}, \mathsf{Enc}, \mathsf{Dec})$ be a nonce-based AEAD scheme. The advantage of an adversary A in breaking the MRAE-security of Π is defined as

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathsf{A}^{\Pi.\mathsf{Enc}_K(\cdot,\cdot,\cdot),\, \Pi.\mathsf{Dec}_K(\cdot,\cdot,\cdot)} = 1\right] - \Pr\left[\mathsf{A}^{\$(\cdot,\cdot,\cdot),\perp(\cdot,\cdot,\cdot)} = 1\right]\right|,$$

where $\$(\cdot,\cdot,\cdot)$ is an oracle which on input (N, A, M) outputs a random string of length $|\Pi.\mathsf{Enc}_K(N, A, M)|$ and $\perp(\cdot,\cdot,\cdot)$ is an oracle which always outputs $\perp$. The adversary is not allowed to repeat a query or to make a decryption query (N, A, Y) if a previous encryption query (N, A, M) returned Y. The adversary is said to be nonce-respecting if it never repeats a nonce N in its encryption queries.

Multi-User Security. All security definitions above can be formulated in the *multi-user* setting [ML15, BT16, HT16, LMP17]. All oracles to which the adversary has access take an additional "identifier" input $i \in \mathcal{I} \subset \mathbb{N}$. For each identifier $i \in \mathcal{I}$, a key K is drawn independently at random in the "real world", whereas an independent ideal oracle is implemented in the "ideal world". We denote the corresponding security notion with prefix "mu-". For example, the mu-PRF security of a keyed function $F : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is defined as

$$\mathbf{Adv}^{\text{mu-prf}}_{F}(\mathsf{A}) = \Big|\Pr\left[K_i \leftarrow_{\$} \mathcal{K}, i \in \mathcal{I} : \mathsf{A}^{(i,x)\mapsto F_{K_i}(x)} = 1\right] - \Pr\left[R_i \leftarrow_{\$} \mathsf{Func}(\mathcal{X}, \mathcal{Y}), i \in \mathcal{I} : \mathsf{A}^{(i,x)\mapsto R_i(x)} = 1\right]\Big|.$$

2.2 Description of AES-GCM-SIV

The AES-GCM-SIV AEAD mode for a block cipher $E : \mathcal{K}_2 \times \{0,1\}^n \to \{0,1\}^n$ combines through the SIV composition method [RS06] a variable input-length PRF built from a polynomial hash function called POLYVAL, a variant of GHASH [MV04], with the counter encryption mode. The high-level structure of the mode is depicted in Figure 1. We describe each component in details below.

The variable-input-length PRF underlying AES-GCM-SIV follows a slight variant of the standard "hash-then-encrypt" (a.k.a. "UHF-then-PRF") construction. It will be convenient for the proof of Theorem 1 below to describe this function, that we denote $\mathsf{HtE}[H, E]$, in the following modular way. It relies on E and an additional keyed function $H : \mathcal{K}_1 \times \{0,1\}^* \times \{0,1\}^* \to \{0,1\}^n$. It takes as input a key $(K_1, K_2) \in \mathcal{K}_1 \times \mathcal{K}_2$, a nonce $N \in \{0,1\}^{\mathtt{nl}}$ with $\mathtt{nl} < n$, associated data $A \in \{0,1\}^*$, and a message $M \in \{0,1\}^*$, and outputs a tag $T \in \{0,1\}^n$ defined as

$$\mathsf{HtE}[H, E]_{K_1,K_2}(N, A, M) = E_{K_2}(H_{K_1}(A, M) \oplus N), \tag{3}$$

where N is left-padded with zeros (AES-GCM-SIV sets $n = 128$ and restricts the nonce-length to 96 bits, but it could potentially be larger).

The specific hash function H used in AES-GCM-SIV is defined as follows. It uses a polynomial hash function POLYVAL taking as input a "hashing" key $K_1 \in \{0,1\}^n$ and a tuple in $(\{0,1\}^n)^+$ and returning a string in $\{0,1\}^n$ (the exact specification of POLYVAL, which is only defined for $n = 128$, is not needed in this paper). It also uses an encoding function Encode taking as input associated data $A \in \{0,1\}^*$ and a message $M \in \{0,1\}^*$ and returning a unique encoding of (A, M) in $(\{0,1\}^n)^+$ by padding A and M with zeros and appending an n-bit block encoding of the length of A and M. Then

$$H_{K_1}(A, M) = 0 \| \mathsf{Trunc}_{n-1}\Big(\mathsf{POLYVAL}\big(K_1, \mathsf{Encode}(A, M)\big)\Big). \tag{4}$$

The specific counter mode CTR used in AES-GCM-SIV uses $(n-1)$-bit counters with a random initial value. More formally, on input a non-empty message $M \in \{0,1\}^*$ parsed

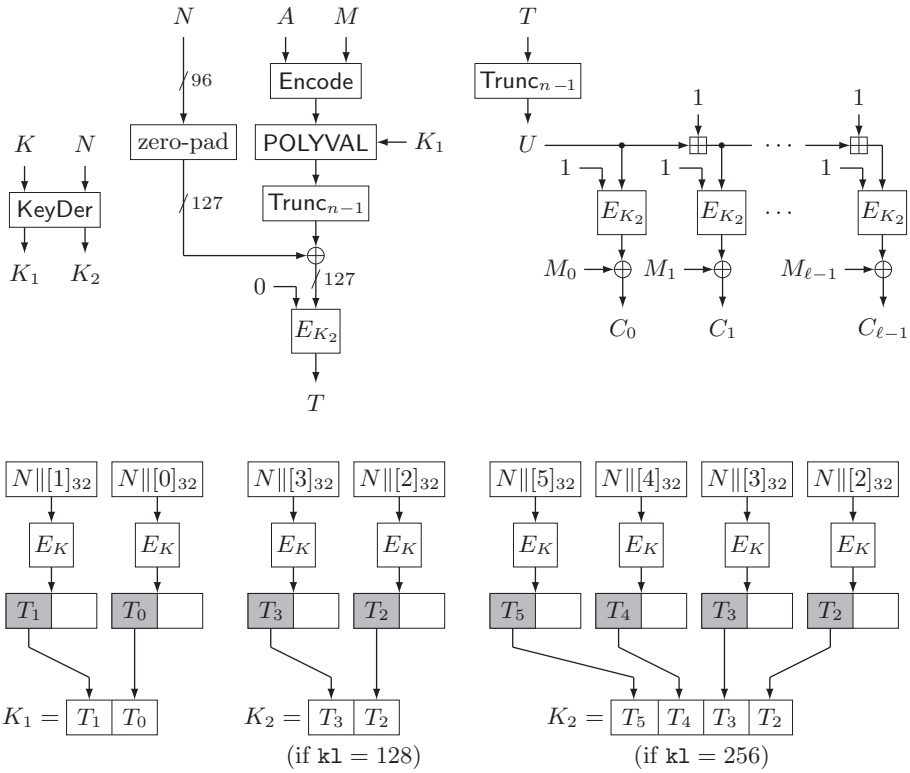

Figure 1: The AES-GCM-SIV mode (top) and the key derivation function (bottom).

in n-bit blocks $(M_0, \ldots, M_{\ell-1})$ where $\ell = \lceil |M|/n \rceil$, $|M_i| = n$ for $i \in \{0, \ldots, \ell-2\}$, and $|M_{\ell-1}| \leq n$, an initial counter U is drawn uniformly at random in $\{0,1\}^{n-1}$ and the i-th ciphertext block is

$$C_i = M_i \oplus E_{K_2}(1 \| (U \boxplus i)),$$

where $U \boxplus i$ denotes addition modulo 2^{32} of the 32 least significant bits of U and i. We denote $\mathsf{CTR}[E]_{K_2}(U, M)$ the result of encrypting message $M \in \{0,1\}^*$ under key $K_2 \in \mathcal{K}_2$ with initial counter $U \in \{0,1\}^{n-1}$.

From the components above, Gueron *et al.* [GLL17] define as an intermediate layer of abstraction the GCM-SIV$^+$ AEAD mode for E as follows. It has key space $\mathcal{K}_1 \times \mathcal{K}_2$, and on input a key pair (K_1, K_2) and a triple (N, A, M) it returns $C \| T$ where

$$T = \mathsf{HtE}[H, E]_{K_1, K_2}(N, A, M),$$
$$C = \mathsf{CTR}[E]_{K_2}(\mathsf{Trunc}_{n-1}(T), M).$$

Finally, the AES-GCM-SIV mode adds key derivation on top of GCM-SIV$^+$. More specifically, keys K_1 and K_2 are derived from a master key $K \in \mathcal{K}_2$ and the nonce N through a key derivation function $\mathsf{KeyDer}[E] : \mathcal{K}_2 \times \{0,1\}^{\mathtt{nl}} \to \mathcal{K}_1 \times \mathcal{K}_2$ (constructed from E, see below). On input a key $K \in \mathcal{K}_2$ and a triple (N, A, M), it returns $C \| T$ defined as

$$(K_1, K_2) = \mathsf{KeyDer}(K, N),$$
$$C \| T = \mathsf{GCM\text{-}SIV}^+_{K_1, K_2}(N, A, M).$$

Remark 1. We note that both GCM-SIV$^+$ and AES-GCM-SIV slightly depart from the "pure" SIV composition method: GCM-SIV$^+$ uses the same key K_2 both in the PRF and in the

encryption scheme, and AES-GCM-SIV uses a nonce-based key derivation function. Note however that once E_{K_2} has been replaced by a uniformly random function F^*, GCM-SIV$^+$ becomes a strict instantiation of SIV: indeed, in the tag generation part, F^* is only called on inputs whose most significant bit is 0, while in the encryption part, it is only called on inputs whose most significant bit is 1, which amounts to having independent uniformly random functions in each part.

The Key Derivation Function. To be complete, it remains to describe the key derivation routine $\mathsf{KeyDer}[E] : (K, N) \mapsto (K_1, K_2)$. It is specified for $E \in \{\mathsf{AES}_{128}, \mathsf{AES}_{256}\}$ (hence the block length is $n = 128$ and $\mathcal{K}_2 = \{0,1\}^{\mathtt{kl}}$ where the key-length $\mathtt{kl}$ is 128 or 256) and nonce-length $\mathtt{nl} = 96$. For $i \in \{0, \ldots, 5\}$, let

$$T_i = \mathsf{Trunc}_{64}\big(\mathsf{AES}_K(N \| [i]_{32})\big).$$

Then $K_1 = T_1 \| T_0$ and

$$\begin{aligned} K_2 &= T_3 \| T_2 && \text{if } \mathtt{kl} = 128, \\ &= T_5 \| T_4 \| T_3 \| T_2 && \text{if } \mathtt{kl} = 256. \end{aligned}$$

In more abstract words, it relies on a PRP-to-PRF conversion method [BKR98, HWKS98] which consists in concatenating truncated outputs of AES applied to the input N and distinct indices. We will see in Section 5 that an equally efficient but more secure PRP-to-PRF conversion method could have been used.

Remark 2. The decryption algorithm is not defined in [GLL17]. When discussing the security of AES-GCM-SIV in the following sections, we will assume the following natural decryption algorithm. On input a master key $K \in \mathcal{K}$ and a triple $(N, A, C\|T)$, we let

$$\begin{aligned} (K_1, K_2) &= \mathsf{KeyDer}(K, N), \\ M &= \mathsf{CTR}[E]_{K_2}(\mathsf{Trunc}_{n-1}(T), C), \\ T' &= \mathsf{HtE}[H, E]_{K_1, K_2}(N, A, M), \end{aligned}$$

and the decryption algorithm returns M if $T = T'$, and $\perp$ otherwise. We note that $M = \mathsf{CTR}[E]_{K_2}(\mathsf{Trunc}_{n-1}(T), C)$ corresponds to counter mode decryption.

3 About the Security Bound

3.1 Problems in GLL's Security Bound

The following security bound for AES-GCM-SIV was claimed in [GLL17]. We omit the running time of adversaries since, unlike queries, they are irrelevant for our discussion.

Theorem ([GLL17], Theorem 6)**.** *Let* A *be an adversary against the MRAE-security of* $\Pi = \mathsf{AES\text{-}GCM\text{-}SIV}[E]$ *where* $E = \mathsf{AES}$*. Assume that* A*:*

- *makes encryption queries of length (in* 128*-bit blocks) at most* $2^k - 1$*,*
- *uses at most* Q *distinct nonces in encryption queries,*
- *repeats any nonce at most* R *times in encryption queries,*
- *makes at most* q_D *decryption queries of total length at most* L *bits for each distinct nonce.*

Then there exists an adversary A' *against the PRF-security of* AES *making at most* $Q(2R + 2q_D + L/128)$ *oracle queries and an adversary* A'' *against the PRP-security of* AES *making at most* $6Q$ *oracle queries such that*

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'') + \min\left\{\frac{36Q^2}{2^{129}}, \frac{6Q}{2^{96}}\right\} + Q\left(2\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') + \frac{R^2}{2^{126-k}} + \frac{R^2 + 2q_D}{2^{127}}\right).$$

We note that it is very unusual and confusing for the security bound of a mode of operation to contain both the PRP- and the PRF-insecurity of the underlying block cipher. Block ciphers are designed to be good PRPs, not PRFs. The remaining of this section will show that confusing the two can be misleading.

More concernedly, this security bound is flawed in at least two respects. First, the authors use an hybrid "multi-user" argument to infer the security bound for the variant of AES-GCM-SIV denoted Π' where a uniformly random function is used to derive keys for each nonce. However, such an argument must take into account *all* keys derived in the security experiment by encryption *and decryption queries*. Since nonces in decryption queries are arbitrary (in particular, they can be completely different from the ones used in encryption queries), this means that the number of hybrid experiments must be $Q+q_D$ (and hence the multiplicative factor in front of the security bound given by [GLL17, Theorem 4] should be $Q + q_D$ as well).[2] To formally disprove the bound as stated, simply consider an attacker against AES-GCM-SIV with AES replaced by a uniformly random permutation, which makes no encryption queries ($Q = 0$) and simply attempts to forge a valid ciphertext within q_D random decryption queries: the bound of [GLL17, Theorem 6] indicates that the advantage of this adversary should be zero, whereas it is clearly not (it is roughly $q_D/2^n$, the probability that a random tag be valid). We note that this problem also affects the term $\min\{36Q^2/2^{129}, 6Q/2^{96}\}$ which accounts for the adversary's advantage in distinguishing the key derivation function from random and which should also consider decryption queries, so that Q should be replaced by $Q + q_D$ in this term as well.

Second, the number of queries made by A', which is claimed to be at most $Q(2R + 2q_D + L/128)$, is incorrect. Details are not given, but the factor Q seems to come from the Q hybrid "multi-user" security experiments, while the term $2R + 2q_D + L/128$ comes from [GLL17, Theorem 4]. However, Q should only multiply the advantage, not the number of queries of the adversary constructed in each hybrid experiment. Besides, the number of queries claimed by [GLL17, Theorem 4] for each hybrid experiment is also erroneous, as we explain in more details in Section 3.2. As we will see later, a correct upper bound for the number of queries made by the PRF-adversary A' against AES is essential for accurately analyzing AES-GCM-SIV's security.

Remark 3. Upper bounding the number of queries made by reduction A' is actually quite straightforward. A source of mistake in [GLL17, Theorem 6] was to derive this upper bound by *composing* reductions, which, in addition to usually resulting in looser bounds, is also quite error-prone. When analyzing the security of a high-level mode of operation which consists of the combination of several components (as is the case for AES-GCM-SIV), one should rather begin with replacing the underlying primitive with its uniformly random counterpart in the high-level mode, and only then analyze the security of each component in the information-theoretic setting.

[2] The discussion here depends on the exact decryption algorithm. We assume that it is defined as in Remark 2.

3.2 Correcting the Security Bound

In order to remedy the situation, we prove a corrected version of the bound. We start with a corrected version of the security bound for the GCM-SIV$^+$ mode. For the sake of clarity, our notation slightly departs from the one used in [GLL17] (we will revert to the original notation when comparing the bounds). We also specify an upper bound on the AD length, which is needed to upper bound the maximal differential probability of POLYVAL. The *message length* of an encryption query (N, A, M), resp. decryption query $(N, A, C\|T)$ is the length of M, resp. C. All lengths below are measured in n-bit blocks.

Theorem 1 (GCM-SIV$^+$ security bound)**.** *Let E be a block cipher with n-bit blocks. Let A be an adversary against the MRAE-security of $\Pi = \mathsf{GCM\text{-}SIV}^+[E]$ making at most q_E encryption queries and q_D decryption queries, such that*

- *the message length in any encryption or decryption query is at most ℓ_m,*
- *the AD length in any encryption or decryption query is at most ℓ_a,*
- *the total message length in encryption queries is at most σ_E,*
- *the total message length in decryption queries is at most σ_D.*

Then there exists an adversary A' against the PRF-security of E making at most $q_E + q_D + \sigma_E + \sigma_D$ oracle queries such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') + \frac{q_E^2 \ell_m}{2^{n-1}} + \frac{(q_E + q_D)^2(\ell_m + \ell_a + 1)}{2^n} + \frac{q_D}{2^n}.$$

Proof. The adversary A has access to either the real encryption and decryption oracles of (with a slight abuse of notation) $\mathsf{GCM\text{-}SIV}^+[H_{K_1}, E_{K_2}]$ for uniformly random keys K_1 and K_2, or $(\$, \bot)$. First, we replace the block cipher E_{K_2} in the real encryption and decryption oracles with a uniformly random function F^*, and denote Π^* the resulting scheme. Consider the adversary A' against the PRF-security of E, having access to an oracle O (which is either E_{K_2} for a random key K_2 or a uniformly random function F^* from $\{0,1\}^n$ to $\{0,1\}^n$), which simply runs A, draws a uniformly random hashing key K_1, and answers all encryption/decryption queries made by A using K_1 and its oracle O. Then A' makes at most $q_E + q_D + \sigma_E + \sigma_D$ oracle queries (in details, for each encryption or decryption query it makes exactly one oracle query when computing the tag and as many oracle queries as the length (in blocks) of the message when encrypting or decrypting it). Moreover, one has

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') + \mathbf{Adv}^{\mathrm{mrae}}_{\Pi^*}(\mathsf{A}).$$

It remains to upper bound A's advantage against Π^*. By Remark 1, Π^* is an instantiation of the generic SIV construction (also called composition method A4 in [NRS14]), so that we can apply Theorem 2 of [RS06] or the result of [NRS14, Appendix A.3] to obtain

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi^*}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{ivE}}_{\mathsf{CTR}[F^*]}(\mathsf{B}) + \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{HtE}[H,F^*]}(\mathsf{B}') + \frac{q_D}{2^n},$$

where B is an adversary against the ivE-security of the counter mode $\mathsf{CTR}[F^*]$ making at most q_E queries, each of length (in n-bit blocks) at most ℓ_m and B' is an adversary against the PRF-security of $\mathsf{HtE}[H, F^*]$, where H is defined as in Equation (4), making at most[3] $q_E + q_D$ queries. (Note that the last term in the bound of [RS06, Theorem 2] is

[3]We note that [NRS14, Lemma 3] charges B' with a higher number of queries, namely $2(q_E + q_D)$, which has been propagated through Theorems 2.2, 3.5, 4.2, and 4.3 of [GL15] up to Theorem 4 of [GLL17]. This is actually due to the fact that Lemma 3 of [NRS14] applies to other composition methods than A4 that require this higher number of queries. It is in fact easy to see (e.g., from the proof of Theorem 2 in [RS06]) that the factor 2 is superfluous.

$(q_E + q_D)/2^n$, but it is easy to see from its proof that this can be improved to $q_D/2^n$ as stated here or in [NRS14, Appendix A.3].)

Clearly, the outputs of $\mathsf{CTR}[F^*]$ are perfectly indistinguishable from random unless two counters collide. This cannot happen for two counters used in the same encryption query. Consider now two distinct encryption queries. Since they are both of length at most ℓ_m, the set of counters used for these two encryptions will overlap *iff* the initial counter of the second encryption falls in a set of size at most $2\ell_m - 1$, which happens with probability at most $(2\ell_m - 1)/2^{n-1}$ (recall that counters are $(n-1)$-bit long). Summing over all pairs of encryption queries, one has

$$\mathbf{Adv}^{\mathrm{ivE}}_{\mathsf{CTR}[F^*]} \leq \frac{q_E(q_E-1)}{2} \cdot \frac{2\ell_m - 1}{2^{n-1}} \leq \frac{q_E^2 \ell_m}{2^{n-1}}. \tag{5}$$

The PRF-security of the UHF-then-PRF construction is standard (xoring the nonce to the output of the hash function implies that one needs the hash function to be AXU rather than AU). Assuming H is ε-AXU, since B' makes at most $q_E + q_D$ queries, one has [GL15, Lemma 3.3]

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{HtE}[H,F^*]}(\mathsf{B}') \leq \frac{(q_E+q_D)^2\varepsilon}{2}.$$

It remains to upper bound ε, which depends on the maximal length of inputs to $\mathsf{POLYVAL}$. Since Encode appends exactly one block to the concatenation of M and A, all inputs to H have length at most $\ell_m + \ell_a + 1$. By [GLL17, Lemma 2], we have that H is ε-AXU for

$$\varepsilon = \frac{\ell_m + \ell_a + 1}{2^{n-1}}. \tag{6}$$

(Note that the denominator is 2^{n-1} rather than 2^n because of the truncation of the most significant bit of $\mathsf{POLYVAL}$.) Combining all equations above yields the result. □

Comparison with GLL's Bound. Theorem 4 in [GLL17] states that for any adversary A making at most q_E encryption queries of maximal message length (in n-bit blocks) $2^k - 1$ and at most q_D decryption queries of overall message length at most L bits,[4] there exists an adversary A' making at most $2q_E + 2q_D + L/n$ oracle queries such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq 2\mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') + \frac{q_E^2}{2^{n-k-2}} + \frac{q_E^2 + q_D}{2^{n-1}}.$$

Restating Theorem 1 using notation of [GLL17, Theorem 4] for the sake of comparison, we obtain that for any adversary A making at most q_E encryption queries of maximal message length (in n-bit blocks) $2^k - 1$ and at most q_D decryption queries of overall message length at most L bits, and such that the AD length (in n-bit blocks) in any query is at most ℓ_a, there exists an adversary A' making at most $q_E + q_D + q_E(2^k - 1) + L$ oracle queries[5] such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') + \frac{q_E^2}{2^{n-k-1}} + \frac{(q_E+q_D)^2(2^k+\ell_a)}{2^n} + \frac{q_D}{2^n}.$$

Putting aside factors 2 or so coming from overlooked optimizations when defining A' and computing a few probabilities, various mistakes throughout the proofs of Theorems 3.5, 4.2, and 4.3 in [GL15] explain the differences between [GLL17, Theorem 4] (whose proof relies on the aforementioned theorems) and Theorem 1 above:

[4] As far as we can tell, no upper bound is put on the AD length in the statement of [GLL17, Theorem 4].
[5] Note that if all messages in decryption queries are 1-bit long, then $\sigma_D = L$.

- The number of queries made by A' does not take into account A's encryption queries. The problem seems to have slipped in when going from [GL15, Theorem 4.3], where L denotes the overall length of all (encryption and decryption) queries to [GLL17, Theorem 4], where L denotes the overall length of decryption queries only.
- The number of queries of adversary B' used for deriving the final security bound in [GL15, Theorem 3.5] is not correct. This adversary, denoted $\mathcal{B}(\mathcal{A})$ in the proof of [GL15, Theorem 3.5], is said to make at most $2(q_E + q_D)$ queries in the middle of the proof (which is correct up to the remark of Footnote 3), but this becomes q_E in Equation (3) at the end of the proof.
- The length of the AD is not taken into account when computing the maximal differential probability ε of H in [GL15, Theorem 4.3]. This is correctly taken into account in the bound of [GL15, Theorem 4.2] through the term $(\lceil L/n \rceil + 1)q_E^2/2^n$ (up to correcting q_E to $q_E + q_D$, see the point above) where L denotes the overall length of all queries (presumably including the AD length). However, in the proof of [GL15, Theorem 4.2], the authors use the inequality $\lceil L/n \rceil \leq 2^k$, where 2^k has been defined in Sect. 2 as the maximum message length (in n-bit blocks), which, as far as we understand, excludes the AD length. The mistake is then propagated to [GL15, Theorem 4.3].

Based on Theorem 1, we can now state the following corrected security bound for AES-GCM-SIV used with a block cipher with 128-bit blocks such as AES for the particular (and simpler) case where $q_D = 0$ (i.e., we consider a CPA-adversary against the privacy of the scheme), which will be sufficient to make our point. We give a more general bound in Section 4.

Theorem 2 (AES-GCM-SIV privacy bound)**.** *Let* A *be an adversary against the MRAE-security of* $\Pi = \mathsf{AES\text{-}GCM\text{-}SIV}[E]$ *with* $E = \mathsf{AES}$*. Assume that* A*:*

- *makes encryption queries of maximal message length (in* 128*-bit blocks)* $2^k - 1$ *and maximal AD length (in* 128*-bit blocks)* ℓ_a*,*
- *uses at most* Q *distinct nonces in encryption queries,*
- *repeats any nonce at most* R *times in encryption queries,*
- *makes no decryption queries.*

Then there exists an adversary A' *against the PRF-security of* AES *making at most* $R \cdot 2^k$ *queries and an adversary* A'' *against the PRP-security of* AES *making at most* $6Q$ *queries such that*

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'') + \min\left\{\frac{36Q^2}{2^{129}}, \frac{6Q}{2^{96}}\right\} + Q\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') + \frac{QR^2}{2^{126-k}} + \frac{QR^2\ell_a}{2^{128}}.$$

Proof. The proof is similar to the one of [GLL17, Theorem 6]. We first replace the function $\mathsf{KeyDer}[E](K, \cdot)$ by a uniformly random function from $\{0,1\}^{\mathtt{nl}}$ to $\mathcal{K}_1 \times \mathcal{K}_2$, and let Π' denote the resulting AEAD scheme. By [GLL17, Lemma 5], there exists an adversary A' against the PRP-security of AES making at most $6Q$ queries such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'') + \min\left\{\frac{36Q^2}{2^{129}}, \frac{6Q}{2^{96}}\right\} + \mathbf{Adv}^{\mathrm{mrae}}_{\Pi'}(\mathsf{A}).$$

Note that Π' is exactly "$\mathsf{GCM\text{-}SIV}^+$ in the multi-user setting", where a fresh pair of keys (K_1, K_2) is drawn uniformly at random for each nonce.[6] Hence, by a straightforward

[6] This is not entirely true since an adversary against $\mathsf{GCM\text{-}SIV}^+$ in the multi-user setting would be able to freely choose nonces in encryption queries for each user, whereas in the security experiment against Π' the nonce is fixed according to the key pair, but this can only *lower* the adversary's advantage.

multi-user hybrid argument and Theorem 1 with $q_E = R$, $\ell_m = 2^k - 1$, $\sigma_E = R \cdot (2^k - 1)$, and $q_D = \sigma_D = 0$, there is an adversary A' against the PRF-security of AES making at most $R \cdot 2^k$ queries such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi'}(\mathsf{A}) \leq Q \left(\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') + \frac{R^2}{2^{126-k}} + \frac{R^2 \ell_a}{2^{128}} \right),$$

which concludes the proof. □

3.3 Analyzing the Security Bound

We assume to begin with that $\ell_a = 0$ and will come back to the impact of the AD length at the end of this section.

It is claimed in [GLL17] that the security bound of AES-GCM-SIV is dominated by the term $QR^2/2^{126-k}$, which captures both the probability that two counters collide for two encryption queries using the same nonce and the probability that two outputs of POLYVAL (with empty AD input) collide for two queries with the same nonce. We disprove this claim by showing that for virtually all parameters, the bound is actually dominated by the term $Q\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}')$. Indeed, since AES is a pseudorandom *permutation*, for any $q \leq 2^{129/2}$, there exists an adversary B making q queries such that

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{B}) \geq 0.316 \frac{q(q-1)}{2^{128}}.$$

(Adversary B simply checks whether there is a collision among answers received from its oracle. See e.g. [BKR00, Proposition 2.4].) Since A' makes up to $R \cdot 2^k$ queries, this means in particular that the best upper bound we can hope for the term $Q\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}')$ is roughly

$$0.316 \frac{QR^2}{2^{128-2k}}.$$

Stated otherwise, the PRF-advantage term in the security bound should be replaced by a PRP-advantage term at the cost of the PRP-PRF switching lemma (see [BKR00, Proposition 2.5]), i.e.,

$$Q\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') \leq Q\mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}') + \frac{QR^2}{2^{129-2k}},$$

after which the security bound of Theorem 2 becomes (assuming $\ell_a = 0$)

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'') + \min\left\{ \frac{36Q^2}{2^{129}}, \frac{6Q}{2^{96}} \right\} + Q\mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}') + \frac{QR^2}{2^{129-2k}} + \frac{QR^2}{2^{126-k}}.$$

Note that $QR^2/2^{129-2k}$ is larger than $QR^2/2^{126-k}$ as soon as $k \geq 4$ and larger than $6Q/2^{96}$ when $R \geq 2^{18-k}$. Hence, assuming that AES is a good PRP, under these two mild conditions, the bound is dominated by the term $QR^2/2^{129-2k}$ corresponding to the "multi-user" PRF-advantage of A' against AES.

A Matching Attack. We stress that the term $QR^2/2^{129-2k}$ in the bound above is actually tight up to some small constant. For any fixed parameters (Q, R, k), consider the following "multi-user PRP-PRF" distinguisher: for $i \in \{1, \ldots, Q\}$ and $j \in \{1, \ldots, R\}$, it queries $(N_i, \emptyset, M_{i,j})$ to the encryption oracle for arbitrary distinct nonces $N_1, \ldots, N_Q$ and arbitrary distinct messages $M_{i,j}$ of length 2^k (and $\emptyset$ denotes empty AD), and returns 1 *iff* for some i, a collision occurs among the $R \cdot 2^k$ blocks of $(C_{i,1} \oplus M_{i,1}, \ldots, C_{i,R} \oplus M_{i,R})$, where $C_{i,j}$ is the answer of the encryption oracle to the query $(N_i, \emptyset, M_{i,j})$.

This adversary returns 1 with probability at most $QR^2/2^{126-k}$ when interacting with the real encryption oracle (see below), while it returns 1 when interacting with $\$(\cdot,\cdot,\cdot)$ with probability

$$\begin{aligned}1-\left(\prod_{i=1}^{R\cdot 2^k-1}\left(1-\frac{i}{2^{128}}\right)\right)^Q &\geq 1-\left(\prod_{i=1}^{R\cdot 2^k-1} e^{-i/2^{128}}\right)^Q\\ &\geq 1-e^{-QR\cdot 2^k(R\cdot 2^k-1)/2\cdot 2^{128}}\\ &\geq 0.316\frac{QR^2}{2^{129-2k}},\end{aligned}$$

implying that the adversary's advantage is lower bounded by

$$0.316\frac{QR^2}{2^{129-2k}}-\frac{QR^2}{2^{126-k}}=\left(0.316-\frac{1}{2^{k-3}}\right)\cdot\frac{QR^2}{2^{129-2k}}.$$

To see that the adversary returns 1 with probability at most $QR^2/2^{126-k}$ when interacting with the real encryption oracle, we closely follow the proof of Theorem 1. The difference is that we treat E_{K_2} as a uniformly random permutation rather than a uniformly random function.

We first fix $i\in\{1,\ldots,Q\}$, and focus on $(C_{i,1}\oplus M_{i,1},\ldots,C_{i,R}\oplus M_{i,R})$. By construction, for any $j\in\{1,\ldots,R\}$, we never have a collision among the 2^k blocks of $C_{i,j}\oplus M_{i,j}$. We then fix distinct $j,j'\in\{1,\ldots,R\}$ and consider the probability of having a collision between one of the 2^k blocks of $C_{i,j}\oplus M_{i,j}$ and one of the 2^k blocks of $C_{i,j'}\oplus M_{i,j'}$. There are two cases that make the collision occur:

Case 1. $H_{K_1}(M_{i,j})=H_{K_1}(M_{i,j'})$.

Case 2. $U_{i,j}\boxplus\ell=U_{i,j'}\boxplus\ell'$ for some $\ell,\ell'\in\{0,1,\ldots,2^k-1\}$, where $U_{i,j}=\mathsf{Trunc}_{n-1}(T_{i,j})$ and $T_{i,j}$ is the tag for the query $(N_i,\emptyset,M_{i,j})$.

Note that Case 2 refers to the event

$$\{U_{i,j},U_{i,j}\boxplus 1,\ldots,U_{i,j}\boxplus(2^k-1)\}\cap\{U_{i,j'},U_{i,j'}\boxplus 1,\ldots,U_{i,j'}\boxplus(2^k-1)\}\neq\emptyset,$$

i.e., the event that we have a collision among these $2\cdot 2^k$ counters.

The probability of Case 1 is at most $\varepsilon=(2^k+1)/2^{127}$ from [GLL17, Lemma 2]. Assuming that we have $H_{K_1}(M_{i,j})\neq H_{K_1}(M_{i,j'})$, the probability of Case 2 is at most $2^{k+1}/(2^{128}-1)$, since for any fixed $T_{i,j}$, there are $2^{128}-1$ possible values for $T_{i,j'}$ that are different from $T_{i,j}$, and the condition $U_{i,j}\boxplus\ell=U_{i,j'}\boxplus\ell'$ requires that the most significant $127-k$ bits of $U_{i,j'}$ be the same as those of $U_{i,j}$, implying that we have at most 2^{k+1} possibilities for $T_{i,j'}$ that meet the condition. Therefore, for any fixed $i\in\{1,\ldots,Q\}$, we have a collision among $C_{i,j}\oplus M_{i,j}$ and $C_{i,j'}\oplus M_{i,j'}$ with probability at most

$$\frac{2^k+1}{2^{127}}+\frac{2^k}{2^{128}-1}\leq\frac{2^k}{2^{125}},$$

and the claim follows from a union bound, as we have at most $R^2/2$ possible choices for distinct $j,j'\in\{1,\ldots,R\}$ and Q choices for $i\in\{1,\ldots,Q\}$.

Random Nonces. A security bound for AES-GCM-SIV when N is a value drawn at random (a "random IV") rather than a non-repeating nonce was also given in [GLL17, Corollary 8], but since it was inferred from [GLL17, Theorem 6] it is flawed as well.[7]

[7] Additionally, it is argued in the proof sketch of [GLL17, Corollary 8] that since the maximal number of repetitions of any nonce is 3 except with small probability, the number of distinct nonces Q resulting from N_E encryption queries is at most $N_E/3$ except with the same small probability; but actually only a very small number of nonces will repeat three times, so that Q is in fact close to N_E.

Table 1: Security bound for AES-GCM-SIV revised according to the leading term of Theorem 2 and Corollary 1 compared with claims in [GLL17, Fig. 4]. We highlight in gray parameters for which the security bound is above 2^{-32} and should be considered insecure according to NIST recommendations for GCM. For the nonce-based version, the total number of encryptions N_E is set to QR.

Scheme	N_E	Q	R	k	bound	[GLL17] claim
AES-GCM-SIV	2^{32}	2^{32}	1	32	2^{-33}	2^{-61}
(nonce based)	2^{64}	2^{64}	1	32	2^{-1}	2^{-29}
	2^{31}	1	2^{31}	32	2^{-3}	2^{-32}
	2^{31}	1	2^{31}	16	2^{-35}	2^{-48}
	2^{39}	1	2^{39}	16	2^{-19}	2^{-32}
	2^{42}	1	2^{42}	10	2^{-25}	2^{-32}
	2^{50}	2^{42}	2^{8}	32	2^{-7}	2^{-36}
	2^{50}	2^{42}	2^{8}	16	2^{-39}	2^{-51}
	2^{50}	2^{46}	2^{4}	32	2^{-11}	2^{-40}
AES-GCM-SIV	2^{48}	—	—	32	2^{-14}	2^{-44}
(random IV)	2^{63}	—	—	16	2^{-31}	2^{-32}

Hence, we also state a corrected bound for this case. As in [GLL17, Sect. 5.2], the proof follows from the fact that the probability that any value repeats four or more times when drawing N_E 96-bit values uniformly at random is at most $(N_E)^4/(24 \cdot 2^{288})$, and applying Theorem 2 with $Q = N_E$ and $R = 3$.

Corollary 1. *Let* A *be an adversary against the MRAE-security of the random IV-based variant of* $\Pi = \mathsf{AES\text{-}GCM\text{-}SIV}[E]$ *where* $E = \mathsf{AES}$*. Assume that* A *makes at most* N_E *encryption queries of maximal message length (in n-bit blocks)* $2^k - 1$ *with empty AD and no decryption queries. Then there exists an adversary* A' *against the PRF-security of* AES *making at most* $3 \cdot 2^k$ *queries and an adversary* A'' *against the PRP-security of* AES *making at most* $6N_E$ *queries such that*

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \leq \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'') + \min\left\{\frac{36(N_E)^2}{2^{129}}, \frac{6N_E}{2^{96}}\right\} + N_E\left(\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') + \frac{9}{2^{126-k}}\right) + \frac{(N_E)^4}{24 \cdot 2^{288}}.$$

Again, for a large range of parameters, this security bound is dominated by the term

$$N_E\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}') \simeq 9N_E/2^{129-2k}.$$

Parameters Examples. For concreteness, we give in Table 1 a revised version of the claims made in [GLL17, Fig. 4] (which did not take the PRF-security of AES into account) based on the corrected security bound of Theorem 2 and Corollary 1. We see that the security guarantees provided by AES-GCM-SIV are significantly weaker than claimed in [GLL17, Fig. 4].

Apart from the numerical examples in [GLL17, Fig. 4], the authors presented three more concrete examples.

1. For AES-GCM-SIV with a nonce, when $Q = 2^{40}, R = 2^8$, and $k = 10$, it is stated in [GLL17, Sect. 5.1] that the security bound is $\varepsilon'' + 2^{57}\varepsilon' + 2^{-53}$, where $\varepsilon' = \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}')$ for A' making 2^{57} queries and $\varepsilon'' = \mathbf{Adv}^{\mathrm{prp}}_{\mathsf{AES}}(\mathsf{A}'')$ for A'' making $6 \cdot 2^{40}$ queries. Even though the number of queries made by A' is erroneous (it should be $R \cdot 2^k = 2^{18}$), the authors should have concluded that the bound was vacuous since for an adversary A' making as much as 2^{57} queries, $\varepsilon' \simeq 2^{-15}$ and $2^{57}\varepsilon'$ is much larger than 1.[8]

2. For AES-GCM-SIV with a random IV, when $N_E = 2^{32}$ and $k = 32$, it is stated in [GLL17, Sect. 5.2] that the adversary's advantage is at most 2^{-60}, whereas our corrected security bound shows that it is only upper bounded by 2^{-30}.

3. The same paragraph states that when $N_E = 2^{64}$ and $k = 32$, the distinguishing probability is at most 2^{-28}, whereas our corrected security bound becomes void (for the good reason that a variant of the matching attack described above succeeds with advantage close to 1).

We note that [GLL17, Sect. 5.3] acknowledges that $\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{AES}}(\mathsf{A}')$ can be large, but this crucial observation is not taken into consideration in [GLL17, Fig. 4] nor in the surrounding discussion.

Impact of the AD Length. It might seem surprising that the AD length shows up in the privacy bound, since the AD is not supposed to be secret. This is in fact a consequence of the definition of privacy of an AEAD scheme, which demands that the tag be indistinguishable from random. This in turn depends on the maximal differential probability of POLYVAL and hence on the AD length. The term $QR^2\ell_a/2^{128}$ in the security bound of Theorem 2 is actually matched by a simple distinguishing attack: for Q distinct nonces $N_1, \ldots, N_Q$, query the encryption oracle with R triplets $(N_i, A_{i,j}, \emptyset)$, where $A_{i,j}$ are arbitrary ADs of length ℓ_a and $\emptyset$ denotes the empty message, receiving a tag $T_{i,j}$ in response; then there will be a collision between two tags $T_{i,j}$ and $T_{i,j'}$ with probability roughly $QR^2\ell_a/2^{128}$, whereas for truly random tags such a collision should happen with probability approximately $QR^2/2^{128}$.

We note that no maximal length for the AD is given in [GLL17], while the CFRG specification draft [GLL16] sets a maximal length of $2^{61} - 1$ bytes.[9] Even though there is little reason in practice for the AD to be that large, such an upper bound implies that the two terms $QR^2\ell_a/2^{128}$ and $QR^2/2^{129-2k}$ are of similar magnitude when both the message length and the AD length are maximal (i.e., $k = 32$ and $\ell_a \simeq 2^{57}$ blocks).

4 A General Security Bound for AES-GCM-SIV

In this section, we provide a general security bound for AES-GCM-SIV. All lengths below are measured in n-bit blocks.

Theorem 3 (AES-GCM-SIV MRAE-security bound)**.** *Let E be a block cipher with n-bit blocks and key space $\mathcal{K}_2 = \{0,1\}^{\mathtt{kl}}$, where* `kl` *is the key length. Let* A *be an adversary against the MRAE-security of* $\Pi = \mathsf{AES\text{-}GCM\text{-}SIV}[E]$ *such that*

- A *uses at most Q distinct nonces in encryption queries,*
- A *repeats any nonce at most R times in encryption queries,*

[8] Curiously, the term 2^{-53} is correct according to our bound, but according to the bound of [GLL17, Theorem 4] it should have been 2^{-60}.

[9] No rationale is given for this choice, which was presumably made according to the GCM specification [Dwo07].

- A *makes at most* q_D *decryption queries,*
- *the message length in any encryption or decryption query is at most* $2^k - 1$*,*
- *the AD length in any encryption or decryption query is at most* ℓ_a*,*
- *the total message length in decryption queries is at most* σ_D*.*

Then there exists an adversary A' *against the mu-PRP-security of* E *making at most* $R \cdot 2^k$ *queries for at most* Q *distinct users and distributing at most* $q_D + \sigma_D$ *additional queries as it wishes among users and an adversary* A'' *against the PRP-security of* E *making at most* $6(Q + q_D)$ *queries such that*

$$\begin{aligned}\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \le \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathsf{A}'') + \min\left\{\frac{36(Q+q_D)^2}{2^{n+1}}, \frac{6(Q+q_D)}{2^{3n/4}}\right\} \\ + \mathbf{Adv}^{\text{mu-prp}}_{E}(\mathsf{A}') + \frac{QR^2}{2^{n-2k}} + \frac{(q_D+\sigma_D)(R\cdot 2^k + q_D + \sigma_D)}{2^n} \\ + \frac{(Q+q_D)^2}{2^{\mathtt{kl}+1}} + \frac{QR^2(2^k+\ell_a)}{2^n} + \frac{Rq_D(2^k+\ell_a)}{2^{n-1}} + \frac{q_D}{2^n} + \frac{QR^2}{2^{n-k-1}}.\end{aligned}$$

The proof is deferred to Appendix A.

Let us briefly comment on the bound that would be obtained with a straightforward multi-user hybrid argument similar to the one used in the proof of [GLL17, Theorem 6] and Theorem 2. As in the proof of Theorem 2, we first replace $\mathsf{KeyDer}[E](K,\cdot)$ by a uniformly random function from $\{0,1\}^{\mathtt{nl}}$ to $\mathcal{K}_1 \times \mathcal{K}_2$, and let Π' denote the resulting AEAD scheme. By [GLL17, Lemma 5], there exists an adversary A'' against the PRP-security of E making at most $6(Q + q_D)$ queries such that

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \le \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathsf{A}'') + \min\left\{\frac{36(Q+q_D)^2}{2^{n+1}}, \frac{6(Q+q_D)}{2^{3n/4}}\right\} + \mathbf{Adv}^{\mathrm{mrae}}_{\Pi'}(\mathsf{A}).$$

Then, we use a multi-user hybrid argument (with $Q + q_D$ users) combined with Theorem 1 where $q_E = R$, $\ell_m = 2^k - 1$, and $\sigma_E = R(2^k - 1)$. This yields the bound

$$\mathbf{Adv}^{\mathrm{mrae}}_{\Pi'}(\mathsf{A}) \le (Q+q_D)\left(\mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') + \frac{R^2}{2^{n-k-1}} + \frac{(R+q_D)^2(2^k+\ell_a)}{2^n} + \frac{q_D}{2^n}\right)$$

for an adversary A' making at most $R \cdot 2^k + q_D + \sigma_D$ oracle queries. (Note that in such a basic multi-user hybrid argument, we have no other choice than assuming that *each* hybrid adversary makes at most q_D decryption queries of total message length σ_D.) By the PRP-PRF switching lemma, we have

$$\mathbf{Adv}^{\mathrm{prf}}_{E}(\mathsf{A}') \le \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathsf{A}') + \frac{(R\cdot 2^k + q_D + \sigma_D)^2}{2^{n+1}}.$$

Combining the three equations above, one obtains

$$\begin{aligned}\mathbf{Adv}^{\mathrm{mrae}}_{\Pi}(\mathsf{A}) \le \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathsf{A}'') + \min\left\{\frac{36(Q+q_D)^2}{2^{n+1}}, \frac{6(Q+q_D)}{2^{3n/4}}\right\} \\ + (Q+q_D)\mathbf{Adv}^{\mathrm{prp}}_{E}(\mathsf{A}') + \frac{(Q+q_D)(R\cdot 2^k + q_D + \sigma_D)^2}{2^{n+1}} \\ + \frac{(Q+q_D)R^2}{2^{n-k-1}} + \frac{(Q+q_D)(R+q_D)^2(2^k+\ell_a)}{2^n} + \frac{(Q+q_D)q_D}{2^n},\end{aligned}$$

which is a very crude bound (note in particular that it contains terms that are cubic in q_D). Instead, we set to prove a better bound with a more careful multi-user argument.

5 About the Key Derivation Function

In this section, we point out that the key derivation function specified in AES-GCM-SIV is sub-optimal w.r.t. security. In short, the designers could have used the "sum of PRPs" construction [BKR98, BI99, Luc00] rather than truncation.

More precisely, consider the key deriving function $\mathsf{KeyDer}'[E]$ which maps (K, N) to (K_1, K_2) where

$$\begin{aligned} K_1 &= E_K(N\|[1]_{32}) \oplus E_K(N\|[0]_{32}) \\ K_2 &= E_K(N\|[3]_{32}) \oplus E_K(N\|[2]_{32}) && \text{if } \mathtt{kl} = n, \\ &= E_K(N\|[5]_{32}) \oplus E_K(N\|[4]_{32}) \big\| E_K(N\|[3]_{32}) \oplus E_K(N\|[2]_{32}) && \text{if } \mathtt{kl} = 2n. \end{aligned}$$

This key derivation function makes exactly the same number of calls to E as the original one. However, its PRF-security is much better and essentially optimal. It has been studied in numerous papers [Luc00, BI99, Pat08a, Pat10, Pat13, DHT17]. In particular, using [DHT17, Theorem 1], the PRF-advantage of any adversary A making at most Q oracle queries against KeyDer' is upper bounded by

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{KeyDer}'}(\mathsf{A}) \leq 3 \cdot \frac{1.5Q + 3\sqrt{Q}}{2^n} \leq \frac{15Q}{2^n}.$$

(Note that in the particular case we are considering, $n = 128$ and $Q \leq 2^{96} = 2^{128-32}$ since nonces are 96 bits long, so that the hypothesis $Q \leq 2^{n-5}$ of [DHT17, Theorem 1] is always met.) Even an adversary which is able to query all 2^{96} possible nonces to KeyDer' has a distinguishing advantage of at most 2^{-28}, whereas it has advantage close to 1 against the original truncation-based key derivation function KeyDer.

Alternatively, one can use of a variant of CENC [Iwa06] to derive the keys as

$$\begin{aligned} K_1 &= E_K(N\|[1]_{32}) \oplus E_K(N\|[0]_{32}) \\ K_2 &= E_K(N\|[2]_{32}) \oplus E_K(N\|[0]_{32}) && \text{if } \mathtt{kl} = n, \\ &= E_K(N\|[3]_{32}) \oplus E_K(N\|[0]_{32}) \big\| E_K(N\|[2]_{32}) \oplus E_K(N\|[0]_{32}) && \text{if } \mathtt{kl} = 2n. \end{aligned}$$

This saves one call to E if $\mathtt{kl} = n$ and two calls if $\mathtt{kl} = 2n$, and the security is comparable to the "sum of PRPs" construction [Pat05, IMV16].

Acknowledgements

Tetsu Iwata was supported by JSPS KAKENHI, Grant-in-Aid for Scientific Research (B), Grant Number 26280045. Yannick Seurin was partially supported by the French Agence Nationale de la Recherche through the BRUTUS project under Contract ANR-14-CE28-0015.

References

[ADL17] Tomer Ashur, Orr Dunkelman, and Atul Luykx. Boosting Authenticated Encryption Robustness With Minimal Modifications. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 (Proceedings, Part III)*, volume 10403 of *LNCS*, pages 3–33. Springer, 2017. Full version at http://eprint.iacr.org/2017/239.

[BDJR97] Mihir Bellare, Anand Desai, E. Jokipii, 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.

[BKR98] Mihir Bellare, Ted Krovetz, and Phillip Rogaway. Luby-Rackoff Backwards: Increasing Security by Making Block Ciphers Non-invertible. In Kaisa Nyberg, editor, *Advances in Cryptology - EUROCRYPT '98*, volume 1403 of *LNCS*, pages 266–280. Springer, 1998.

[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.

[BR06] Mihir Bellare and Phillip Rogaway. The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In Serge Vaudenay, editor, *Advances in Cryptology - EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, 2006. Full version available at http://eprint.iacr.org/2004/331.

[BT16] Mihir Bellare and Björn Tackmann. The Multi-user Security of Authenticated Encryption: AES-GCM in TLS 1.3. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 (Proceedings, Part I)*, volume 9814 of *LNCS*, pages 247–276. Springer, 2016.

[BZD+16] Hanno Böck, Aaron Zauner, Sean Devlin, Juraj Somorovsky, and Philipp Jovanovic. Nonce-Disrespecting Adversaries: Practical Forgery Attacks on GCM in TLS. In *USENIX Workshop on Offensive Technologies, WOOT 2016*. USENIX Association, 2016.

[CAE] CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness. See http://competitions.cr.yp.to/caesar.html.

[CLS17] Benoît Cogliati, Jooyoung Lee, and Yannick Seurin. New Constructions of MACs from (Tweakable) Block Ciphers. *IACR Trans. Symmetric Cryptol.*, 2017(2):27–58, 2017.

[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.

[DHT17] Wei Dai, Viet Tung Hoang, and Stefano Tessaro. Information-theoretic Indistinguishability via the Chi-squared Method. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 (Proceedings, Part III)*, volume 10403 of *LNCS*, pages 497–523. Springer, 2017. Full version at http://eprint.iacr.org/2017/537.

[Dwo07] Morris Dworkin. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D, November 2007. Available at http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-38d.pdf.

[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. Available at http://eprint.iacr.org/2015/102.

[GLL16] Shay Gueron, Adam Langley, and Yehuda Lindell. AES-GCM-SIV: Nonce Misuse-Resistant Authenticated Encryption. CFGR Draft, 2016. Available at https://tools.ietf.org/html/draft-irtf-cfrg-gcmsiv-05.

[GLL17] Shay Gueron, Adam Langley, and Yehuda Lindell. AES-GCM-SIV: Specification and Analysis. IACR Cryptology ePrint Archive, Report 2017/168, 2017. Available at http://eprint.iacr.org/2017/168.

[HT16] Viet Tung Hoang and Stefano Tessaro. Key-Alternating Ciphers and Key-Length Extension: Exact Bounds and Multi-user Security. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO 2016 (Proceedings, Part I)*, volume 9814 of *LNCS*, pages 3–32. Springer, 2016.

[HWKS98] Chris Hall, David Wagner, John Kelsey, and Bruce Schneier. Building PRFs from PRPs. In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98*, volume 1462 of *LNCS*, pages 370–389. Springer, 1998.

[IM16] Tetsu Iwata and Kazuhiko Minematsu. Stronger Security Variants of GCM-SIV. *IACR Trans. Symmetric Cryptol.*, 2016(1):134–157, 2016.

[IMV16] Tetsu Iwata, Bart Mennink, and Damian Vizár. CENC is Optimally Secure. IACR Cryptology ePrint Archive, Report 2016/1087, 2016. Available at http://eprint.iacr.org/2016/1087.

[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.

[Jou06] Antoine Joux. Authentication Failures in NIST Version of GCM. Comments submitted to NIST Modes of Operation Process, 2006. Available at http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/comments/800-38_Series-Drafts/GCM/Joux_comments.pdf.

[Jou09] Antoine Joux. *Algorithmic Cryptanalysis*. CRC Press, 2009.

[LMP17] Atul Luykx, Bart Mennink, and Kenneth G. Paterson. Analyzing Multi-key Security Degradation. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 (Proceedings, Part II)*, volume 10625 of *LNCS*, pages 575–605. Springer, 2017. Available at http://eprint.iacr.org/2017/435.

[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.

[McG12] David McGrew. Impossible plaintext cryptanalysis and probable-plaintext collision attacks of 64-bit block cipher modes. IACR Cryptology ePrint Archive, Report 2012/623, 2012. Available at http://eprint.iacr.org/2012/623.

[ML15] Nicky Mouha and Atul Luykx. Multi-key Security: The Even-Mansour Construction Revisited. In Rosario Gennaro and Matthew Robshaw, editors, *Advances in Cryptology - CRYPTO 2015 (Proceedings, Part I)*, volume 9215 of *LNCS*, pages 209–223. Springer, 2015.

[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.

[NSA17] NSA Information Insurance. Key Recovery Attacks on AES-GCM-SIV. CFRG mailing list, 2017. Available at https://www.ietf.org/mail-archive/web/cfrg/current/pdfTJvccyJc1O.pdf.

[Pat05] Jacques Patarin. On Linear Systems of Equations with Distinct Variables and Small Block Size. In Dongho Won and Seungjoo Kim, editors, *Information Security and Cryptology - ICISC 2005*, volume 3935 of *LNCS*, pages 299–321. Springer, 2005.

[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.

[Pat10] Jacques Patarin. Introduction to Mirror Theory: Analysis of Systems of Linear Equalities and Linear Non Equalities for Cryptography. IACR Cryptology ePrint Archive, Report 2010/287, 2010. Available at http://eprint.iacr.org/2010/287.

[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.

[QUI] QUIC, a multiplexed stream transport over UDP. See https://www.chromium.org/quic.

[Rog04] 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.

[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.

A Proof of Theorem 3

We will make use of the H-coefficients technique [Pat08b], that we recall very briefly here. See e.g. [CS14] for more details. Consider an adversary A interacting with one out of two possible systems (i.e., tuples of oracles) S_{re} and S_{id}, called by convention respectively the *real world* and the *ideal world*, and outputting a single bit. The interaction of A with either system defines a *transcript* τ which lists all queries made by A together with their answers. A transcript is said attainable if it can be obtained with non-zero probability when A interacts with the ideal world. We let X_{re}, resp. X_{id} denote the random variable for the transcript in the real, resp. ideal world. The fundamental lemma of the H-coefficients technique is the following one.

Lemma 1. *Fix an adversary A trying to distinguish two systems S_{re} and S_{id}. Let $\mathcal{T}_{\mathrm{bad}}$ be a subset of the set $\mathcal{T}$ of attainable transcripts and $\mathcal{T}_{\mathrm{good}} = \mathcal{T} \setminus \mathcal{T}_{\mathrm{bad}}$ be its complement. Assume that there exists ν such that for any $\tau \in \mathcal{T}_{\mathrm{good}}$,*

$$\frac{\Pr\left[X_{\mathrm{re}} = \tau\right]}{\Pr\left[X_{\mathrm{id}} = \tau\right]} \geq 1 - \nu.$$

Then

$$\left|\Pr\left[\mathsf{A}^{S_{\mathrm{re}}} = 1\right] - \Pr\left[\mathsf{A}^{S_{\mathrm{id}}} = 1\right]\right| \leq \Pr\left[X_{\mathrm{id}} \in \mathcal{T}_{\mathrm{bad}}\right] + \nu.$$

We will need the following "constrained" multi-user PRP-PRF switching lemma for a multi-user adversary making queries according to a specific pattern. Note that we cannot use the nice result that PRP-PRF switching does not suffer multi-user degradation [LMP17] (according to which we could use the single-user PRP-PRF switching bound with the total number Q_{tot} of adversarial queries) since the only upper bound we have for Q_{tot} is $QR \cdot 2^k + q_D + \sigma_D$, and this would yield a bound which is quadratic in Q.

Lemma 2. *Let E be a block cipher with n-bit blocks. Let A be an adversary against the mu-PRF-security of E making at most q queries for at most Q distinct users and distributing at most q' additional queries as it wishes among users. Then*

$$\mathbf{Adv}_E^{\text{mu-prf}}(\mathsf{A}) \leq \mathbf{Adv}_E^{\text{mu-prp}}(\mathsf{A}) + \frac{Q \cdot q^2}{2^n} + \frac{q'(q+q')}{2^n}.$$

Proof. First, we replace E by a family of independent and uniformly random permutations, at the cost of the advantage of A against the mu-PRP-security of E. We must now upper bound A's advantage in distinguishing a family of independent and uniformly random permutations from a family of independent and uniformly random functions. For this, we use the H-coefficients technique. We assume *wlog* that the adversary makes the maximal number of allowed queries and never repeats queries. We let the real world be the family of functions and the ideal world be the family of permutations. There is no bad transcript. Let τ be an attainable transcript. This implies that for each identifier $i \in \mathcal{I}$, all queries with identifier i are distinct (by the convention that the adversary never repeats queries) and all corresponding answers are distinct (since in the ideal world the adversary interacts with permutations). Then, in the real (function) world,

$$\Pr\left[X_{\mathrm{re}} = \tau\right] = \frac{1}{(2^n)^{Q \cdot q + q'}}.$$

On the other hand, in the ideal (permutation) world,

$$\Pr[X_{\mathrm{id}} = \tau] = \prod_{i \in \mathcal{I}} \frac{1}{(2^n)_{q_i}},$$

where q_i is the number of queries for identifier i in the transcript and $(2^n)_{q_i} = 2^n(2^n - 1)\cdots(2^n - q_i + 1)$ with the convention that $(2^n)_0 = 1$. Let $\mathcal{J}$ be the set of Q identifiers for which the adversary makes q queries, and for $i \in \mathcal{I}$ let q'_i be the number of additional queries made for identifier i. Then

$$\begin{aligned}\Pr[X_{\mathrm{id}} = \tau] &\leq \prod_{i \in \mathcal{J}} \frac{1}{(2^n)_q} \prod_{i \in \mathcal{I}} \frac{1}{(2^n - q)_{q'_i}} \\ &\leq \left(\frac{1}{(2^n)_q}\right)^Q \cdot \frac{1}{(2^n - q)_{q'}},\end{aligned}$$

where we used that $\sum_{i \in \mathcal{I}} q'_i = q'$. Then we obtain for the ratio

$$\begin{aligned}\frac{\Pr[X_{\mathrm{re}} = \tau]}{\Pr[X_{\mathrm{id}} = \tau]} &\geq \left(\frac{(2^n)_q}{(2^n)^q}\right)^Q \cdot \frac{(2^n - q)_{q'}}{(2^n)^{q'}} \\ &\geq 1 - \frac{Q \cdot q^2}{2^n} - \frac{q'(q + q')}{2^n}.\end{aligned}$$

Combined with Lemma 1, this concludes the proof. □

We are now ready to prove our improved security bound for AES-GCM-SIV.

Proof of Theorem 3. We use a game-based approach [Sho04, BR06], i.e., we gradually modify the behavior of the two oracles (that we call "worlds" rather than games) to which the adversary has access. World W_1 corresponds to the real encryption and decryption oracles. The changes in each of the successive worlds, which are formally specified in Figure 2 and Figure 3, are as follows:

- in world W_2, we replace the key derivation function $\mathsf{KeyDer}[E](K, \cdot)$ by a uniformly random function $\rho_{\mathrm{kd}} : \{0,1\}^{\mathtt{nl}} \to \mathcal{K}_1 \times \mathcal{K}_2$;
- in world W_3, we replace the block cipher E with a uniformly random function $F^* : \mathcal{K}_2 \times \{0,1\}^n \to \{0,1\}^n$;
- in world W_4, we use two independent random functions $F^*, G^* : \{0,1\}^{\mathtt{nl}} \times \{0,1\}^n \to \{0,1\}^n$ keyed by the nonce for resp. the tag generation part and the encryption part;[10]
- in world W_5, we replace the tag generation function by a random function ρ_{tag} from $\{0,1\}^{\mathtt{nl}} \times \{0,1\}^* \times \{0,1\}^*$ to $\{0,1\}^n$ and let the decryption oracle always reject;
- the final world W_6 is simply the ideal world $(\$, \perp)$.

In all the following, we let

$$\Delta_{i,j} = \left|\Pr\left[\mathsf{A}^{W_i} = 1\right] - \Pr\left[\mathsf{A}^{W_j} = 1\right]\right|.$$

We are interested in upper bounding $\Delta_{1,6}$ and will consider each transition in turn.

[10] Since F^* and G^* are secret random functions, using the nonce directly as the key is equivalent (but syntactically simpler) to drawing *distinct* keys $K_{2,N}$ for keying F^* in the tag generation part and distinct keys $K'_{2,N}$ for keying G^* in the encryption part.

TRANSITION $\mathbf{W_1}$-$\mathbf{W_2}$. It is easy to see that $\Delta_{1,2}$ can be upper bounded by the PRF-advantage against KeyDer of an adversary making at most $Q+q_D$ oracle queries. By [GLL17, Lemma 5], there exists an adversary A'' against the PRP-security of E making at most $6(Q+q_D)$ queries such that

$$\Delta_{1,2} \leq \mathbf{Adv}_E^{\mathrm{prp}}(\mathsf{A}'') + \min\left\{\frac{36(Q+q_D)^2}{2^{n+1}}, \frac{6(Q+q_D)}{2^{3n/4}}\right\}.$$

TRANSITION $\mathbf{W_2}$-$\mathbf{W_3}$. We construct an adversary A' against the mu-PRF-security of E. Let O be the oracle to which A' has access. It runs A and answers its queries according to the pseudocode of worlds $\mathsf{W}_{2/3}$ by drawing all necessary keys K_1 itself and replacing all calls to E/F^* by calls to its oracle, using the nonce as the "user identifier". Then A' makes at most $R \cdot 2^k$ queries for at most Q distinct users and distributes at most $q_D + \sigma_D$ additional queries as it wishes among users and it perfectly simulates W_2, resp. W_3, when its oracle is E, resp. F^*, so that

$$\Delta_{2,3} \leq \mathbf{Adv}_E^{\text{mu-prf}}(\mathsf{A}').$$

Combined with Lemma 2 with $q = R \cdot 2^k$ and $q' = q_D + \sigma_D$, we obtain

$$\Delta_{2,3} \leq \mathbf{Adv}_E^{\text{mu-prp}}(\mathsf{A}') + \frac{QR^2}{2^{n-2k}} + \frac{(q_D+\sigma_D)(R\cdot 2^k + q_D + \sigma_D)}{2^n}.$$

TRANSITION $\mathbf{W_3}$-$\mathbf{W_4}$. Note that in W_3, all calls to F^* in the tag generation part, resp. encryption part, have their most significant bit set to 0, resp. 1, which is equivalent to having two independent families of random functions. Hence, as long as all keys K_2 generated in W_3 are distinct, worlds W_3 and W_4 are perfectly equivalent (if two keys collide in W_3, the same function is used for two distinct nonces, whereas in W_4 functions associated with distinct nonces are independent by construction). By the fundamental lemma of game playing, the indistinguishability advantage is upper bounded by the probability that two keys K_2 collide in W_3, so that

$$\Delta_{3,4} \leq \frac{(Q+q_D)^2}{2^{\mathtt{k}\mathtt{l}+1}}.$$

TRANSITION $\mathbf{W_4}$-$\mathbf{W_5}$. Let TagGen and Ver be the oracles defined in Figure 3. By simulating the encryption part, we construct an adversary B having access to a pair of oracles $(O_1, O_2) \in \{(\mathsf{TagGen}, \mathsf{Ver}), (\rho_{\mathrm{tag}}, \bot)\}$ which runs A and answers its queries as follows, lazily sampling random functions G_N^* when needed: on an encryption query (N, A, M), it returns $C\|T$ where $T = O_1(N, A, M)$ and

$$C = \mathsf{CTR}[G_N^*](\mathsf{Trunc}_{n-1}(T), M);$$

on a decryption query $(N, A, C\|T)$, it computes

$$M = \mathsf{CTR}[G_N^*](\mathsf{Trunc}_{n-1}(T), C),$$

queries $O_2(N, A, M, T)$, and returns M if O_2 returns $\top$ and $\bot$ if O_2 returns $\bot$. Then one can check that B perfectly simulates W_4 when $(O_1, O_2) = (\mathsf{TagGen}, \mathsf{Ver})$ and W_5 when $(O_1, O_2) = (\rho_{\mathrm{tag}}, \bot)$, so that

$$\Delta_{4,5} = \left|\Pr\left[\mathsf{B}^{\mathsf{TagGen},\mathsf{Ver}} = 1\right] - \Pr\left[\mathsf{B}^{\rho_{\mathrm{tag}},\bot} = 1\right]\right|. \tag{7}$$

Moreover, note that B's queries to its left, resp. right oracle have the same characteristics as A's queries to its encryption, resp. decryption oracle. In particular, B never submits

a right query (N, A, M, T) if a previous left query (N, A, M) returned T (this can only happen if A makes a decryption query $(N, A, C\|T)$ such that a previous encryption query (N, A, M) returned $C\|T$, which is forbidden by definition of MRAE-security). We will refer to B's queries to the left, resp. right oracle as *tag queries*, resp. *verification queries*.

We must now upper bound B's distinguishing advantage. For this, we use the H-coefficients technique. More specifically, our approach is very similar to [CLS17]. We refer to $(\mathsf{TagGen}, \mathsf{Ver})$, resp. $(\rho_{\mathrm{tag}}, \perp)$ as the *real*, resp. *ideal* world. From the interaction of B with its oracles, we build the *queries transcript* which consists of all tag queries together with their answer, which we denote generically $(N, A, M) \to T$, and all verification queries (N', A', M', T'). Note that for an attainable transcript (whose probability in the ideal world is non-zero), all answers to verification queries are $\perp$, and hence we omit these answers from the queries transcript.

In order to define good and bad transcripts easily, we reveal to B, after it has made all its queries, the hashing keys $K_{1,N}$ for nonces appearing in the queries transcript (in the ideal world, we simply reveal "dummy" keys that are uniformly random and independent from the queries transcript). By appending these keys to the queries transcript, we obtain what we simply call the *attack transcript*. We let X_{re}, resp. X_{id} denote the random variable for the attack transcript in the real, resp. ideal world.

We say that a transcript is bad if one of the two following conditions is fulfilled (otherwise we say that it is good):

(C-1) there exists two distinct tag queries $(N, A_1, M_1) \to T_1$ and $(N, A_2, M_2) \to T_2$ with the same nonce such that

$$H_{K_{1,N}}(A_1, M_1) = H_{K_{1,N}}(A_2, M_2);$$

(C-2) there exists a tag query $(N, A, M) \to T$ and a verification query (N, A', M', T') with the same nonce such that

$$\begin{cases} H_{K_{1,N}}(A, M) = H_{K_{1,N}}(A', M') \\ T = T'. \end{cases}$$

Note that the second condition cannot happen in the real world since it would imply that the verification query is valid (i.e., should have returned $\top$).

Probability of Bad Transcripts. We consider each condition in turn, using the ε-AU property of H (recall that in the ideal world, hashing keys are random and independent from the queries transcript). Consider the first condition. For each of the Q possible values of the nonce, and for each of the $R(R-1)/2$ possible pairs of tag queries for this nonce, the probability of a hash output collision is at most ε. By a union bound and Equation (6) with $\ell_m = 2^k - 1$, we obtain that the probability that the first condition is met is at most

$$\frac{QR^2(2^k + \ell_a)}{2^n}.$$

Consider now the second condition. Fix any verification query (N, A', M', T'). There are at most R tag queries $(N, A, M) \to T$ with the same nonce. Let us fix one of them. We distinguish two cases. If the verification query comes after the tag query, then either $T \neq T'$ or $(A, M) \neq (A', M')$ since otherwise this would mean that B submitted a verification query (N, A, M, T) after having received tag T to tag query (N, A, M), which is forbidden. In the first case, the condition cannot be fulfilled, while in the later case the probability that the hash outputs collide is at most ε. If the tag query comes after the verification query, then T is uniformly random and independent from T', and hence the condition is fulfilled with probability 2^{-n}. Since $\varepsilon \geq 2^{-n}$, in all cases, the condition is fulfilled with

probability at most ε. By summing over the at most Rq_D possible pairs of verification and tag queries and using Equation (6) with $\ell_m = 2^k - 1$, we obtain that the probability that the ideal transcript satisfies condition (C-2) is at most

$$\frac{Rq_D(2^k + \ell_a)}{2^{n-1}}.$$

All in all,

$$\Pr\left[X_{\text{id}} \text{ is bad}\right] \leq \frac{QR^2(2^k + \ell_a)}{2^n} + \frac{Rq_D(2^k + \ell_a)}{2^{n-1}}. \tag{8}$$

Probability of Good Transcripts. Fix any good transcript τ. It remains to lower bound the ratio $\Pr\left[X_{\text{re}} = \tau\right] / \Pr\left[X_{\text{id}} = \tau\right]$. We omit the probability that hashing keys take some particular value since it is the same in both worlds and cancels in the ratio. Then one simply has $\Pr\left[X_{\text{id}} = \tau\right] = 1/(2^n)^{QR}$. The probability to obtain τ in the real world is exactly the probability (over functions F_N^*) that for each tag query $(N, A, M) \to T$,

$$F_N^*(H_{K_{1,N}}(A, M) \oplus N) = T$$

and for each verification query (N', A', M', T')

$$F_{N'}^*(H_{K_{1,N'}}(A', M') \oplus N') \neq T'.$$

Note that for each nonce N used in encryption queries, values $H_{K_{1,N}}(A, M) \oplus N$ for tag queries (N, A, M) made with this nonce are distinct, as otherwise condition (C-1) would be fulfilled, and that no decryption query is "incompatible" with the tag queries transcript, as otherwise condition (C-2) would be fulfilled. This implies that

$$\Pr\left[X_{\text{re}} = \tau\right] \geq \left(\frac{1}{2^n}\right)^{QR} \left(1 - \frac{q_D}{2^n}\right).$$

Thus,

$$\frac{\Pr\left[X_{\text{re}} = \tau\right]}{\Pr\left[X_{\text{id}} = \tau\right]} \geq 1 - \frac{q_D}{2^n}. \tag{9}$$

Finally, combining Equation (7), Lemma 1, Equation (8), and Equation (9), we obtain

$$\Delta_{4,5} \leq \frac{QR^2(2^k + \ell_a)}{2^n} + \frac{Rq_D(2^k + \ell_a)}{2^{n-1}} + \frac{q_D}{2^n}.$$

Transition $\mathbf{W_5}$-$\mathbf{W_6}$. Clearly, the distinguishing advantage from W_5 to W_6 is upper bounded by the mu-ivE-advantage against the counter encryption mode of an adversary making at most R queries of maximal length 2^k to at most Q users. By a straightforward hybrid argument and Equation (5), one has

$$\Delta_{5,6} \leq \frac{QR^2}{2^{n-k-1}}.$$

Combining all bounds on $\Delta_{i,i+1}$ for $i = 1, \ldots, 5$ yields the result. □

algorithm KeyGen — World W_1
  $K \leftarrow_{\$} \mathcal{K}_2$
algorithm $\mathsf{Enc}(N, A, M)$
  $(K_1, K_2) := \mathsf{KeyDer}(K, N)$
  $T := E_{K_2}(H_{K_1}(A, M) \oplus N)$
  $C := \mathsf{CTR}[E_{K_2}](\mathsf{Trunc}_{n-1}(T), M)$
  return $C \| T$
algorithm $\mathsf{Dec}(N, A, C \| T)$
  $(K_1, K_2) := \mathsf{KeyDer}(K, N)$
  $M := \mathsf{CTR}[E_{K_2}](\mathsf{Trunc}_{n-1}(T), C)$
  $T' := E_{K_2}(H_{K_1}(A, M) \oplus N)$
  if $T = T'$ **then return** M **else return** $\perp$

algorithm KeyGen — World W_2
  $\rho_{\mathrm{kd}} \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}}, \mathcal{K}_1 \times \mathcal{K}_2)$
algorithm $\mathsf{Enc}(N, A, M)$
  $(K_1, K_2) := \rho_{\mathrm{kd}}(N)$
  $T := E_{K_2}(H_{K_1}(A, M) \oplus N)$
  $C := \mathsf{CTR}[E_{K_2}](\mathsf{Trunc}_{n-1}(T), M)$
  return $C \| T$
algorithm $\mathsf{Dec}(N, A, C \| T)$
  $(K_1, K_2) := \rho_{\mathrm{kd}}(N)$
  $M := \mathsf{CTR}[E_{K_2}](\mathsf{Trunc}_{n-1}(T), C)$
  $T' := E_{K_2}(H_{K_1}(A, M) \oplus N)$
  if $T = T'$ **then return** M **else return** $\perp$

algorithm KeyGen — World W_3
  $\rho_{\mathrm{kd}} \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}}, \mathcal{K}_1 \times \mathcal{K}_2)$
  $F^* \leftarrow_{\$} \mathsf{Func}(\mathcal{K}_2 \times \{0,1\}^n, \{0,1\}^n)$
algorithm $\mathsf{Enc}(N, A, M)$
  $(K_1, K_2) := \rho_{\mathrm{kd}}(N)$
  $T := F^*_{K_2}(H_{K_1}(A, M) \oplus N)$
  $C := \mathsf{CTR}[F^*_{K_2}](\mathsf{Trunc}_{n-1}(T), M)$
  return $C \| T$
algorithm $\mathsf{Dec}(N, A, C \| T)$
  $(K_1, K_2) := \rho_{\mathrm{kd}}(N)$
  $M := \mathsf{CTR}[F^*_{K_2}](\mathsf{Trunc}_{n-1}(T), C)$
  $T' := F^*_{K_2}(H_{K_1}(A, M) \oplus N)$
  if $T = T'$ **then return** M **else return** $\perp$

Figure 2: Worlds W_1-W_3 used in the proof of Theorem 3. The keyed hash function H is defined as in Equation (4).

algorithm KeyGen — World W_4

$\rho_{\mathrm{kd}} \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}}, \mathcal{K}_1)$
$F^* \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}} \times \{0,1\}^n, \{0,1\}^n)$
$G^* \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}} \times \{0,1\}^n, \{0,1\}^n)$
algorithm $\mathsf{Enc}(N, A, M)$
$T := \mathsf{TagGen}(N, A, M)$
$C := \mathsf{CTR}[G^*_N](\mathsf{Trunc}_{n-1}(T), M)$
return $C \| T$
algorithm $\mathsf{Dec}(N, A, C\|T)$
$M := \mathsf{CTR}[G^*_N](\mathsf{Trunc}_{n-1}(T), C)$
if $\mathsf{Ver}(N, A, M, T) = \top$ **then return** M **else return** $\bot$
algorithm $\mathsf{TagGen}(N, A, M)$
$K_{1,N} := \rho_{\mathrm{kd}}(N)$
$T := F^*_N(H_{K_{1,N}}(A, M) \oplus N)$
return T
algorithm $\mathsf{Ver}(N, A, M, T)$
$K_{1,N} := \rho_{\mathrm{kd}}(N)$
$T' := F^*_N(H_{K_{1,N}}(A, M) \oplus N)$
if $T = T'$ **then return** $\top$ **else return** $\bot$

algorithm KeyGen — World W_5

$\rho_{\mathrm{tag}} \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}} \times \{0,1\}^* \times \{0,1\}^*, \{0,1\}^n)$
$G^* \leftarrow_{\$} \mathsf{Func}(\{0,1\}^{\mathtt{nl}} \times \{0,1\}^n, \{0,1\}^n)$
algorithm $\mathsf{Enc}(N, A, M)$
$T := \rho_{\mathrm{tag}}(N, A, M)$
$C := \mathsf{CTR}[G^*_N](\mathsf{Trunc}_{n-1}(T), M)$
return $C \| T$
algorithm $\mathsf{Dec}(N, A, C\|T)$
return $\bot$

Figure 3: Worlds W_4-W_5 used in the proof of Theorem 3. The keyed hash function H is defined as in Equation (4).