Making GCM Great Again: Toward Full Security and Longer Nonces

Woohyuk Chung[1], Seongha Hwang[1], Seongkwang Kim[2], Byeonghak Lee[2], and Jooyoung Lee[1]*

[1] KAIST, Daejeon, Korea {hephaistus,mathience98,hicalf}@kaist.ac.kr
[2] Samsung SDS, Seoul, Korea, {sk39.kim,byghak.lee}@samsung.com

Abstract. The GCM authenticated encryption (AE) scheme is one of the most widely used AE schemes in the world, while it suffers from risk of nonce misuse, short message length per encryption and an insufficient level of security. The goal of this paper is to design new AE schemes achieving stronger provable security in the standard model and accepting longer nonces (or providing nonce misuse resistance), with the design rationale behind GCM.

As a result, we propose two enhanced variants of GCM and GCM-SIV, dubbed eGCM and eGCM-SIV, respectively. eGCM and eGCM-SIV are built on top of a new CENC-type encryption mode, dubbed eCTR: using $2n$-bit counters, eCTR enjoys beyond-birthday-bound security without significant loss of efficiency. eCTR is combined with an almost uniform and almost universal hash function, yielding a variable input-length variable output-length pseudorandom function, dubbed HteC. GCM and GCM-SIV are constructed using eCTR and HteC as building blocks.

eGCM and eGCM-SIV accept nonces of arbitrary length, and provide almost the full security (namely, n-bit security when they are based on an n-bit block cipher) for a constant maximum input length, under the assumption that the underlying block cipher is a pseudorandom permutation (PRP). Their efficiency is also comparable to GCM in terms of the rate and the overall speed.

Keywords: authenticated encryption, GCM, beyond-birthday-bound security, provable security

1 Introduction

Authenticated Encryption. Authenticated Encryption (AE) aims to achieve the two fundamental security goals of symmetric key cryptography, namely, the confidentiality and the authenticity of data. A significant amount of research has been conducted in this field, resulting in the proposal of numerous AE schemes. Currently, a variety of standard algorithms are in use, including CCM [47],

* Jooyoung Lee was supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (RS-2024-00335568).

GCM [35], AES-GCM-SIV [13], and Chacha20-poly1305 [34]. These AE standards are widely employed to ensure the security of modern communication protocols, such as QUIC [27], TLS [41,42], and SSH [31]. Most of the recent constructions accept associated data (AD), which are authenticated but not encrypted. In this paper, we will consider AE schemes with associated data.

Limitations on GCM. GCM is by far the most widespread AE scheme in the world. It is standardized in NIST Special Publication 800-38D [35] and ISO/IEC 19772:2020 [18], and is the recommended cipher algorithm for numerous communication protocols, including WPA3, IEEE802.11ad, SSH, TLS1.2, and TLS1.3. However, several significant issues have recently been raised regarding the use of GCM (and other standardized AE schemes), drawing attention from both academia and industry.

One of the principal issues on GCM is its nonce misusing risk. Nonces or initial vectors (IVs) are used in most encryption schemes in order to guarantee the variability of the ciphertext and prevent the replay attack. Since formalized by Rogaway [43], nonce-based AE has become a substantial category of AE. It is of significant importance to guarantee the uniqueness of nonces in AEs. When it comes to GCM, it completely loses authenticity as soon as a single nonce is used twice [28]. There have been several attacks that exploit its nonce misuse such as internet-wise nonce-reusing HTTPS server scanning [4] and a key reinstallation attack (KRACK) on WPA2 [45].

However, maintaining the uniqueness of nonces is a challenging task. In addition to external factors such as implementation flaws and low-entropy environments, GCM (and other CTR-based AEs) are susceptible to structural limitations in nonce lengths. When GCM is used with an n-bit block cipher, the combined size of the public nonce and the internal counter is restricted to n bits. In typical applications involving 128-bit block ciphers (such as AES), 96-bit nonces and 32-bit counters are employed. Then, the probability of a nonce collision approaches 2^{-32} for every 2^{32} encryption when using a random nonce. Accordingly, the GCM standard advises that the number of invocations should be restricted to 2^{32} when utilizing random nonces. However, this is a major limitation for large systems, such as high-volume networks. For example, Amazon recently noted the practical challenge of GCM, saying that there is a use-case where 2^{32} invocations can be reached in 2 seconds [29]. Alternatively, a deterministic approach is also recommended, where a nonce is generated through the combination of a device ID and a sequence counter. However, this approach requires the additional expense of managing device IDs and sequence counters, and for large systems where invocation limits are an issue, a 96-bit room might still be insufficient to hold device IDs and sequence counters.

In order to avoid nonce misuse, one might consider using longer nonces. However, reducing the counter size to increase the space for the nonce is not a viable solution, as the counter size, which is already limited to 2^{32}, severely restricts the message length per encryption. If the message length exceeds the counter size, then the same counter should be used twice, leading to a potential vulnerability of the scheme. For this reason, the AES-GCM standard imposes a limit of

$2^{39} - 256$ bits on the message length. This constraint necessitates the splitting of data during encryption of large data sets (e.g., large files, streaming data, databases), while it is incumbent upon developers to consider this when implementing AES-GCM. On the other hand, accepting longer nonces using additional GHash operations is neither efficient nor enhance the overall security [32].

GCM also faces a significant security concern due to its limited level of security. The security of GCM is limited to the birthday bound, which means that with an n-bit block cipher, GCM can be compromised by attacks with data complexity of $O(2^{\frac{n}{2}})$. This limitation not only makes it challenging to use smaller block ciphers, like 64-bit ones, but also puts 128-bit block ciphers, such as AES, at risk. This issue has been highlighted by industry leaders. During a recent NIST Workshop on Block Cipher Modes of Operation[3], companies like Meta and Amazon expressed concerns about the continued use of 128-bit block ciphers with GCM. They argued that this combination potentially poses a significant security risk in cloud-scale systems. Given the exponential growth in data usage, it is expected to become a major threat in the near future. For instance, exabyte-scale ($10^{18} \approx 2^{60}$) data is already in use, with zettabyte-scale ($10^{21} \approx 2^{70}$) data expected soon, further amplifying these security concerns.

Beyond GCM. The issues on GCM listed above are all contingent upon the block size. Consequently, if a larger block cipher were to be employed, these issues would be effectively resolved. However, replacing AES with Rjindael-256 [9] or other wide-block ciphers is challenging due to several factors. AES has widespread hardware support, extensive software optimization, and has undergone thorough security analysis, making it a trusted standard. In contrast, Rjindael-256 lacks hardware acceleration, often resulting in slower performance, and might have unknown vulnerabilities due to less scrutiny. For this reason, it will be suitable to utilize AE schemes that is more reliable and provably more secure than GCM, to address the immediate issues at hand.

One promising approach is to use AE schemes to enjoy beyond-birthday bound (BBB) security. BBB-secure AE schemes have been extensively studied. Iwata proposed the CHM [19] and CIP [21] that combine CENC [20], a BBB-secure nonce-based encryption mode, with a universal hash (UH) function using field multiplications. Bhattacharya and Nandi proposed an almost optimally secure variant of GCM, dubbed mGCM [3], by applying similar methods to GCM. mGCM is proved to be secure up to $O(2^n)$ input blocks only when an adversary makes a single query or it is non-adaptive. On the other hand, there are BBB secure AEs based on tweakable block ciphers (TBCs) or ideal ciphers (ICs) such as ΘCB [30], Romulus [22], and LightOCB [5]. However, using TBCs or ICs is inherently less efficient than BCs, and similar to wide block ciphers, there is a lack of standardized primitives and their comprehensive analysis. Also, BBB-secure AE schemes without large enough nonce spaces still carry nonce-reusing risks.

[3] https://csrc.nist.gov/Events/2023/third-workshop-on-block-cipher-modes-of-operation

An alternative approach is to use misuse resistant AE schemes. Rogaway and Shrimpton [44] formalized the notion of misuse-resistant AE (MRAE) and proposed a method of turning a deterministic AE scheme into a nonce-based MRAE scheme. MRAE schemes include EAX [1], SIV [44], AEZ [16], and GCM-SIV [14]. Later, this notion has been refined by viewing the adversarial distinguishing advantage as a function of the maximum number of multicollisions in nonce values (amongst all encryption queries) [40], or the number of queries with repeated nonces [6]. Although there are MRAE schemes enjoying BBB-security such as SCT [40], ZAE [24], GCM-SIV2 [23] and AES-GCM-SIV [13], none of them is built on top of a standard PRP nor has affordable efficiency loss (for example, in terms of rate).

In conclusion, numerous AE schemes have been proposed so far, yet a superior AE that provides both sufficient security and usability remains elusive. Consequently, this paper aims to develop such an AE, specifically one that meets the following requirements.

1. Beyond-birthday-bound security is provided, and the full security is preferred.
2. Efficiency is comparable to GCM in terms of the rate and the overall speed.
3. Extended nonces are accepted or nonce misuse resistance is guaranteed.
4. The maximum message length should not significantly affect the overall security bound, allowing one to encrypt longer messages.
5. Provable security should be guaranteed under the standard PRP assumption, and hence, standard block ciphers (such as AES) should be supported without frequent rekeying.

1.1 Our Contribution

In this paper, we propose enhanced variants of GCM and GCM-SIV, dubbed eGCM and eGCM-SIV, respectively. eGCM and eGCM-SIV accept nonces of arbitrary length, and provide almost the full security (namely, n-bit security when they are based on an n-bit block cipher) for a constant maximum input length, under the assumption that the underlying block cipher is a pseudorandom permutation (PRP).

As the starting point, we construct an IV-based variable output-length pseudorandom function (VOL-PRF), dubbed eCTR. eCTR follows an CENC-like structure, but it generates output blocks by utilizing $2n$-bit inputs. Precisely, for a fixed positive integer w and n-bit IVs A and B, the $(iw + j)$-th output block of eCTR is defined as

$$E_K(A \oplus 2^{i(w+1)}B) \oplus E_K(A \oplus 2^{i(w+1)+j}B)$$

where E_K denotes an n-bit block cipher with key K (see also Figure 1). If A and B are chosen uniformly at random, then eCTR is secure up to $O(2^n)$ output blocks for a constant maximum output length per query, and secure up to $O(2^{\frac{2n}{3}})$ output blocks with no limit on the maximum output length per query.

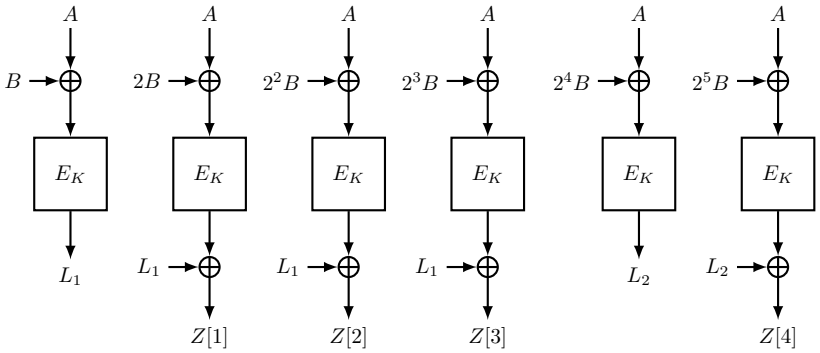

Fig. 1: The first 4 blocks from $\mathsf{eCTR}[E_K, 3](A, B)$ with $w = 3$.

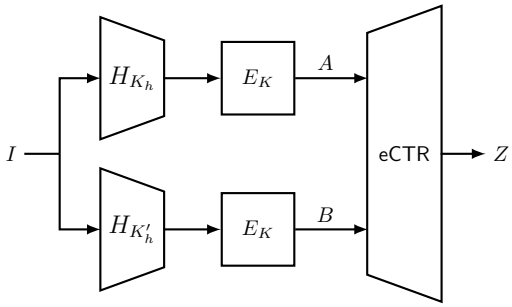

Fig. 2: The HteC VIL-VOL pseudorandom function.

For the next step, we propose a new (almost) fully secure variable input-length variable output-length PRF (VIL-VOL-PRF), dubbed HteC, by combining eCTR with a double-block hash function. In this way, HteC accepts inputs of arbitrary length. Specifically, given a bit string I of arbitrary length, HteC generates intermediate values A and B as follows.

$$A = E_K(H_{K_h}(I)),$$
$$B = E_K(H_{K'_h}(I))$$

where H is an almost universal and almost uniform hash function. Then, the output blocks of HteC are defined as the output blocks from eCTR with A and B being the inputs to eCTR (see also Figure 2). Here, A and B are not perfectly uniform, leading to only negligible loss of security since H is almost universal and almost uniform.

Existing BBB-secure VOL-PRF constructions such as bbb-ddd-AES [10] and a nonce-key derivation function in DNDK-GCM [12] share some similarities with our eCTR/DECK approach as they also follow CENC-like structures. On the other hand, we note that bbb-ddd-AES generates masks by encrypting counter-tweak pairs, requiring twice as many block cipher calls as eCTR's doubling-based

mask updates, and it achieves $2n/3$-bit security compared to eCTR's almost n-bit security. It is also noteworthy that DNDK-GCM's nonce-key derivation requires random or carefully crafted inputs, while our HteC accepts arbitrary inputs using a fewer number of block cipher calls. Due to these advantages, eCTR and HteC outperform the existing constructions.

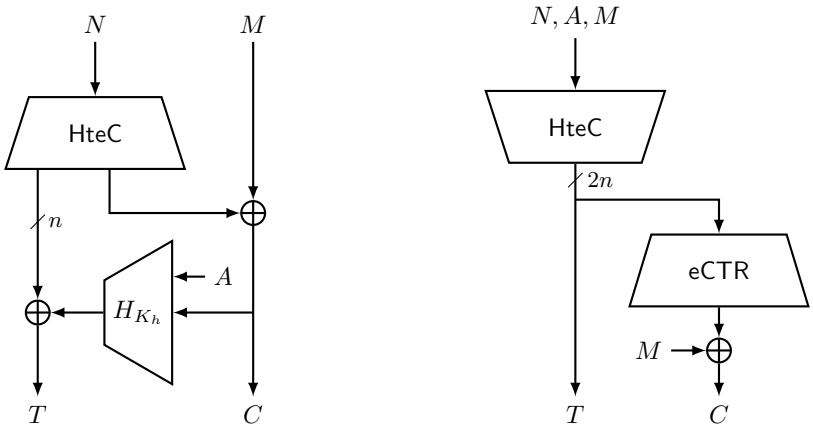

(a) Encryption of eGCM (b) Encryption of eGCM-SIV

Fig. 3: The eGCM and eGCM-SIV AE schemes. A nonce, an associated data, and a message are denoted N, A and M, respectively. Key inputs to HteC and eCTR, and the final tag truncation are omitted.

Since the CTR mode can be viewed as a VOL-PRF, the CTR part in GCM can be replaced by HteC to achieve stronger security and nonce length extension at the same time. The resulting construction is exactly eGCM, and analogously, eGCM-SIV is obtained by replacing the underlying PRF and CTR in GCM-SIV by HteC and eCTR, respectively (see Figure 3).

Figure 4 compares the influence of the maximum message length ℓ to the threshold number of the total length of encryption queries σ for some variants of GCM. Since eGCM and eGCM-SIV do not concatenate counter with nonce for block inputs, they do not have message length limitations by design. Moreover, eGCM and eGCM-SIV have relatively small security impact as message length increases.

Table 1 compares eGCM and eGCM-SIV to well-known AE schemes. eGCM provides stronger security than OCB3, GCM and CWC+. CHM, CIP and mGCM also provides n-bit security as eGCM, while it does not support nonce length extension, which restricts the use of random nonces. On the other hand, eGCM-SIV provides n-bit security in the nonce-misuse setting, while existing MRAE schemes provide only $\frac{n}{2}$-bit security. We emphasize that eGCM and eGCM-SIV are the first

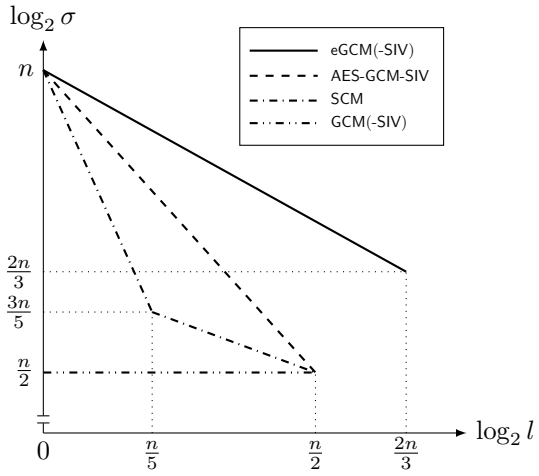

Fig. 4: The threshold number of the total length of the encryption queries σ as a function in ℓ. For MRAE schemes, the number of queries with repeated nonces is fixed as a small constant. For AES-GCM-SIV, the security bound in the ideal cipher model is used.

variants of GCM permitting nonce length extension and providing the full n-bit security at the same time.

Table 2 compares eGCM using POLYVAL[4] [13] as a universal hash function to existing AE schemes in terms of efficiency. In this comparison, we focus on the AE schemes whose reference codes are publicly available. The implementations of ChaCha20-Poly1305 and GCM are taken from BoringSSL[5]. Note that BoringSSL's benchmark includes additional operations for sealing and opening, so the actual performance of ChaCha20-Poly1305 and GCM would be better than given in the table. Our experiments are done in AMD Ryzen 7 2700X Eight-Core Processor (CPU@3.2GHz) which supports PCLMUL, AVX, SSE, and AES instructions, using GCC 11.4.0 with optimization level -O3. We note that eGCM requires an additional doubling operation compared to GCM, while POLYVAL is more implementation-friendly than GHash.

1.2 Tweakable Enciphering Schemes Based on eCTR

A *tweakable enciphering scheme* (TES) is a length-preserving tweakable permutation that accepts a message of variable length and returns the corresponding ciphertext of the same length. The eCTR encryption mode can be employed in a block cipher-based TES with the hash-CTR-hash approach [46], as seen in Figure 5, achieving beyond-birthday-bound security. This construction takes as

[4] POLYVAL is a universal hash function used in AES-GCM-SIV.
[5] https://boringssl.googlesource.com/boringssl

AEAD	Rate	Security		Reference
		NR	NM	
OCB3	1	$n/2$	-	[30]
GCM	1/2	$n/2$	-	[33]
CIP, CHM, mGCM	$\lesssim 1/2^{\ddagger}$	n	-	[21,19,3], Section 4.1
eGCM	$\lesssim 1/2^{\ddagger}$	n	-	Section 4.1
GCM-SIV	1/2	$n/2$	$n/2$	[15]
AES-GCM-SIV	1/2	n	$n/2$	[26]
SCM	1/2	n	$n/2$	[6]
CWC+	$\lesssim 1/2^{\ddagger}$	$3n/4$	$n/2^{\dagger}$	[11]
eGCM-SIV	$\lesssim 1/2^{\ddagger}$	n	n	Section 4.2

† Authenticity only. CWC+ does not provide privacy in the nonce-misuse setting.
‡ Depends on the parameter w, while we write $\lesssim 1/2$ since the rate approaches $1/2$ as w increases and w can be set to a large enough value.

Table 1: Security and efficiency of eGCM, eGCM-SIV and other block cipher based AE schemes. The maximum message length ($= \ell$) is assumed to be a small constant. All the AE schemes are based on a standard block cipher except for AES-GCM-SIV which is based on an ideal cipher. The rate is the number of blocks processed per unit operation, which includes block cipher computation and n-bit field multiplication.

input a message of at least $2n$ bits; let $M = M_1 \| M_2$, where $M_1 \in \{0,1\}^{2n}$ and $M_2 \in \{0,1\}^*$. Then M_1 is encrypted by a $2n$-bit pseudorandom permutation P based on an n-bit block cipher, and the sum of the input and the output of P is used as an IV of eCTR to encrypt M_2. We propose two candidates for P.

- 5-round Feistel cipher based on a block cipher: when each round function is instantiated with a block cipher using a distinct key, the 5-round Feistel cipher becomes a pseudorandom permutation that is secure up to $2^{\frac{2n}{3}}$ queries [2]. In this way, the resulting TES is inverse free, and expected to provide beyond-birthday-bound security.
- CTET$^+$ with $w = 2$ [8]: This construction is a 2-round substitution-permutation cipher using 4 block cipher calls and universal hashing. CTET$^+$ is also secure up to $2^{\frac{2n}{3}}$ queries.

Formal security proof and analysis of the efficiency are left for further research.

2 Preliminary

2.1 Basic Notation

The set $\{0,1\}^n$ is sometimes regarded as a finite field $\mathbf{GF}(2^n)$ with 2^n elements, assuming that 2 cyclically generates all the nonzero elements of $\mathbf{GF}(2^n)$. We

Mode	Message			Reference
	1KB	4KB	64KB	
ChaCha20-Poly1305†	3.92	2.84	2.52	[34]
OCB	0.73	0.61	0.52	[30]
GCM†	2.38	1.30	0.95	[35]
eGCM‡	1.07	1.00	0.98	Section 4.1
AES-GCM-SIV	1.56	1.25	1.15	[13]
SCM	1.42	1.29	1.25	[6]
eGCM-SIV‡	1.55	1.36	1.30	Section 4.2

† Computed by BoringSSL's speed command, so it includes additional operations for sealing.
‡ $w = 24$ is used.

Table 2: Performance of eGCM and other AE schemes. Throughput is measured in cycles per byte.

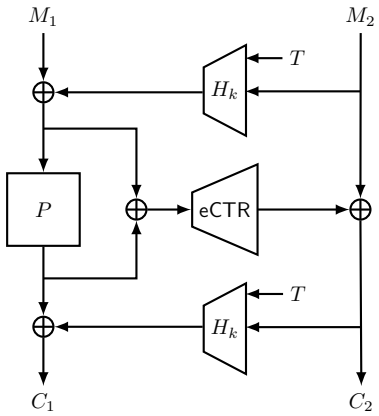

Fig. 5: Hash-CTR-hash type tweakable enciphering scheme based on eCTR.

write $\{0,1\}^*$ to denote the set of all binary strings including the empty string ε. For $A, B \in \{0,1\}^*$, $|A|$ denotes the length of A in bits, and $A \,\|\, B$ denotes the concatenation of A and B.

For a positive integer q, we write $[q] = \{1, \ldots, q\}$. For a non-empty set $\mathcal{X}$, $X \leftarrow_{\$} \mathcal{X}$ denotes that X is drawn uniformly at random from $\mathcal{X}$. The set of all sequences that consist of b pairwise distinct elements of $\mathcal{X}$ is denoted $\mathcal{X}^{*b}$. For positive integers $a \geq b$, let $(a)_b = a(a-1)\ldots(a-b+1)$, and $(a)_0 = 1$ by convention. If $|\mathcal{X}| = a$, then $(a)_b$ becomes the size of $|\mathcal{X}|^{*b}$. $\mathsf{msb}_s(X)$ and $\mathsf{lsb}_s(X)$ denotes the s most significant bits and s least significant bits of X, respectively.

For a real number t, $\lceil t \rceil$ is the smallest integer that is the same as or bigger than t.

Let $\mathsf{Perm}(n)$ be the set of all permutations from $\{0,1\}^n$ to $\{0,1\}^n$, and let $\mathsf{Func}(n,m)$ be the set of all functions from $\{0,1\}^n$ to $\{0,1\}^m$. Given two strings A, B and injective encoding encode, we write $\mathsf{encode}(A,B)$ for encoded string. Throughout this paper, we fix encode function, and for $F : \{0,1\}^* \to \mathcal{Y}$, we simply write $F(A,B) = F(\mathsf{encode}(A,B))$.

2.2 Security Notions

Almost Xor Universal and Almost Uniform Hash Functions. Let $\delta > 0$, and let $H : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ be a keyed function for three non-empty sets $\mathcal{K}$, $\mathcal{X}$, and $\mathcal{Y}$. H is said to be *δ-almost XOR universal* (AXU) if for any distinct $X, X' \in \mathcal{X}$ and $Y \in \mathcal{Y}$,

$$\Pr\left[K \leftarrow_{\$} \mathcal{K} : H_K(X) \oplus H_K(X') = Y\right] \leq \delta.$$

Moreover, H is said to be *δ'-almost uniform* (AU) if for any $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$,

$$\Pr\left[K \leftarrow_{\$} \mathcal{K} : H_K(X) = Y\right] \leq \delta'.$$

PRPs and PRFs. Let $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ be a keyed permutation with key space $\mathcal{K}$, where $E(K,\cdot)$ is a permutation for each $K \in \mathcal{K}$. We will write $E_K(X)$ to denote $E(K,X)$. A (q,t)-distinguisher against E is an algorithm $\mathcal{D}$ with oracle access to an n-bit permutation, making at most q oracle queries, running in time at most t, and outputting a single bit. The advantage of $\mathcal{D}$ in breaking the PRP-security of E, i.e., in distinguishing E from a uniform random permutation $\pi \leftarrow_{\$} \mathsf{Perm}(n)$, is defined as

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

$\mathsf{Adv}^{\mathrm{prp}}_E(q,t)$ is defined as the maximum of $\mathsf{Adv}^{\mathrm{prp}}_E(\mathcal{D})$ over all (q,t)-distinguishers $\mathcal{D}$.

Let

$$F : \mathcal{K} \times \mathcal{X} \times \mathbb{N} \to \mathcal{Y}$$

be a keyed function with key space $\mathcal{K}$, input space $\mathcal{X}$, length space $\mathbb{N}$, and output space $\mathcal{Y}$. Then $F(K,\cdot,m)$ is a function from $\mathcal{X}$ to $\{0,1\}^m$ for $K \in \mathcal{K}$ and $m \in \mathbb{N}$. We will write $F_K(\cdot,\cdot)$ to denote $F(K,\cdot,\cdot)$.

A distinguisher $\mathcal{D}$ against the PRF-security of F is an algorithm that is allowed to make an oracle query with $(X,s) \in \mathcal{X} \times \mathbb{N}$, where $\mathcal{D}$ is supposed to choose a distinct X for every query; in the real world, $F_K(X,s)$ is returned for a secret key $K \in \mathcal{K}$, and in the ideal world, an independent random string $Z \in \{0,1\}^s$ is returned. Let \$ denote such an oracle in the ideal world. Then the advantage of $\mathcal{D}$ in breaking the PRF-security of F is defined as

$$\mathsf{Adv}^{\mathsf{prf}}_F(\mathcal{D}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathcal{D}^{F_K(\cdot,\cdot)} = 1\right] - \Pr\left[\mathcal{D}^{\$} = 1\right]\right|.$$

Throughout this paper, we will assume that F is based on an n-bit block cipher (or simply an n-bit permutation). Then a distinguisher $\mathcal{D}$ against the PRF-security of F is called a (q, σ, ℓ, t)-distinguisher if $\mathcal{D}$ runs in time at most t, making at most q oracle queries where the total output length (over all the queries) is at most σ blocks of n bits, and the output length of each query is at most ℓ blocks of n bits. $\mathsf{Adv}^{\mathsf{prf}}_F(q, \sigma, \ell, t)$ is defined to be the maximum of $\mathsf{Adv}^{\mathsf{prf}}_F(\mathcal{D})$ over all (q, σ, ℓ, t)-distinguishers $\mathcal{D}$. When the running time is unlimited, we will simply omit the parameter t, writing "(q, σ, ℓ)-distinguisher", $\mathsf{Adv}^{\mathsf{prf}}_F(q, \sigma\ell)$.

The goal of an adversary $\mathcal{D}'$ against IV-PRF security of F is defined similarly:

$$\mathsf{Adv}^{\mathsf{iv\text{-}prf}}_F(\mathcal{D}') = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathcal{D}'^{F_K(\cdot,\cdot)}\right] - \Pr\left[\mathcal{D}'^{\$'} = 1\right]\right|.$$

The difference is that $\mathcal{D}'$ chooses $X \in \mathcal{X}$ uniformly and independently at random. $\mathsf{Adv}^{\mathsf{iv\text{-}prf}}_F(q, \sigma, \ell, t)$ and $\mathsf{Adv}^{\mathsf{iv\text{-}prf}}_F(q, \sigma, \ell)$ are also defined similarly.

Nonce based and Misuse resistant AEs. Given five non-empty sets $\mathcal{K}$, $\mathcal{N}$, $\mathcal{A}$, $\mathcal{M}$ and $\mathcal{T}$, a nonce-based authenticated encryption (AE) scheme is a tuple

$$\Pi = (\mathcal{K}, \mathcal{N}, \mathcal{A}, \mathcal{M}, \mathcal{T}, \mathsf{Enc}, \mathsf{Dec}),$$

where Enc and Dec are called encryption and decryption algorithms, respectively. The encryption algorithm Enc takes as input a key $K \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, an associated data $A \in \mathcal{A}$, and a message $M \in \mathcal{M}$, and outputs a ciphertext $C \in \mathcal{M}$ and a tag $T \in \mathcal{T}$. The decryption algorithm Dec takes as input a tuple $(K, N, A, C, T) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \times \mathcal{T}$, and outputs either a message $M \in \mathcal{M}$ or a special symbol $\perp$. We require that

$$\mathsf{Enc}(K, N, A, M) = (C, T) \Rightarrow \mathsf{Dec}(K, N, A, C, T) = M$$

for any tuple $(K, N, A, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M}$. We will write $\mathsf{Enc}_K(N, A, M)$ and $\mathsf{Dec}_K(N, A, C, T)$ to denote $\mathsf{Enc}(K, N, A, M)$ and $\mathsf{Dec}(K, N, A, C, T)$, respectively.

The goal of an adversary $\mathcal{D}$ against the nAE security of Π is to distinguish the real world $(\mathsf{Enc}_K, \mathsf{Dec}_K)$ (using a random key K, unknown to $\mathcal{D}$) and the ideal world. The ideal world oracles are $(\mathsf{Rand}, \mathsf{Rej})$, where Rand returns an independent random string of length $|\mathsf{Enc}_K(N, A, M)|$ and Rej always returns $\perp$ for every decryption query. We assume that $\mathcal{D}$ does not make a decryption query by reusing any previous encryption query. The advantage of $\mathcal{D}$ breaking the nAE-security of Π is defined as

$$\mathsf{Adv}^{\mathsf{nAE}}_{\Pi}(\mathcal{D}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathcal{D}^{\mathsf{Enc}_K, \mathsf{Dec}_K} = 1\right] - \Pr\left[\mathcal{D}^{\mathsf{Rand}, \mathsf{Rej}} = 1\right]\right|.$$

A $(q_e, q_d, \sigma, \ell, t)$-adversary against the nonce-based AE security of Π is an algorithm that makes at most q_e encryption queries to its first oracle with at most q_d decryption queries to its second oracle, and running in time at most t, where the length of each encryption/decryption query is at most l blocks of n bits, and the total length of the encryption queries (nonce excluded) is at most σ

blocks of n bits. However, the adversary is allowed to repeat nonces in its Dec oracle. We define $\mathsf{Adv}^{\mathsf{nAE}}_{\Pi}(q_e, q_d, \sigma, \ell, t)$ as the maximum of $\mathsf{Adv}^{\mathsf{nAE}}_{\Pi}(\mathcal{D})$ over all $(q_e, q_d, \sigma, \ell, t)$-adversaries $\mathcal{D}$ against Π. When we consider information theoretic security, we will drop the parameter t.

The advantage of $\mathcal{D}$ breaking the mrAE security of Π is defined as

$$\mathsf{Adv}^{\mathsf{mrAE}}_{\Pi}(\mathcal{D}) = \left|\Pr\left[K \leftarrow_{\$} \mathcal{K} : \mathcal{D}^{\mathsf{Enc}_K, \mathsf{Dec}_K} = 1\right] - \Pr\left[\mathcal{D}^{\$^{\mathsf{Rand}, \mathsf{Rej}}} = 1\right]\right|,$$

where an adversary $\mathcal{D}$ is allowed to repeat nonces even in encryption queries, and we can define $\mathsf{Adv}^{\mathsf{mrAE}}_{\Pi}(q_e, q_d, \sigma, \ell)$ similarly to $\mathsf{Adv}^{\mathsf{nAE}}_{\Pi}(q_e, q_d, \sigma, \ell, t)$.

2.3 Coefficient-H Technique

We will use Patarin's coefficient-H technique [39]. The goal of this technique is to upper bound the adversarial distinguishing advantage between a real construction and its ideal counterpart. In the real and the ideal worlds, an information-theoretic adversary $\mathcal{D}$ is allowed to make queries to certain oracles (with the same oracle interfaces), denoted $\mathcal{O}_{\mathsf{real}}$ and $\mathcal{O}_{\mathsf{ideal}}$, respectively. The interaction between the adversary $\mathcal{D}$ and the oracle determines a "transcript"; it contains all the information obtained by $\mathcal{D}$ during the interaction. We call a transcript τ *attainable* if the probability of obtaining τ in the ideal world is non-zero. We also denote T_{id} (resp. T_{re}) the probability distribution of the transcript τ induced by the ideal world (resp. the real world). By extension, we use the same notation to denote a random variable distributed according to each distribution.

We partition the set of attainable transcripts Γ into a set of "good" transcripts Γ_{good} such that the probabilities to obtain some transcript $\tau \in \Gamma_{\mathsf{good}}$ are close in the real world and the ideal world, and a set Γ_{bad} of "bad" transcripts such that the probability to obtain any $\tau \in \Gamma_{\mathsf{bad}}$ is small in the ideal world. With this partition, the coefficient-H technique is summarized by the following lemma.

Lemma 1. *Let $\Gamma = \Gamma_{\mathsf{good}} \sqcup \Gamma_{\mathsf{bad}}$ be a partition of the set of attainable transcripts, where there exists a non-negative real number ε_1 such that for any $\tau \in \Gamma_{\mathsf{good}}$,*

$$\frac{\Pr\left[\mathsf{T}_{\mathsf{re}} = \tau\right]}{\Pr\left[\mathsf{T}_{\mathsf{id}} = \tau\right]} \geq 1 - \varepsilon_1,$$

and there exists ε_2 such that $\Pr[\mathsf{T}_{\mathsf{id}} \in \Gamma_{\mathsf{bad}}] \leq \varepsilon_2$. Then for any distinguisher $\mathcal{D}$, one has

$$\left|\Pr\left[\mathcal{D}^{\mathcal{O}_{\mathsf{real}}} = 1\right] - \Pr\left[\mathcal{D}^{\mathcal{O}_{\mathsf{ideal}}} = 1\right]\right| \leq \varepsilon_1 + \varepsilon_2.$$

We refer to [17] for the proof of Lemma 1.

2.4 Mirror Theory

Mirror theory was first proposed by Patarin [37,38] as a useful tool to lower bound the number of solutions to a multi-variable system of equations, and then

to prove the security of Feistel ciphers and the sum of random permutations. Recently, Cogliati et al. [7] proved Mirror theory for a wide range of $\xi_{\max}$ (which will be defined later) with a new formal approach using the link-deletion equations.

We will represent a system of equations by a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex corresponds to an n-bit *distinct* unknown, and each edge is labeled by an element in $\{0,1\}^n$. In particular, an edge with label $\lambda \in \{0,1\}^n$ connecting two vertices P and Q, denoted $P \overset{\lambda}{-} Q$, represents an equation $P \oplus Q = \lambda$. An assignment of *distinct* values to the vertices of $\mathcal{V}$ satisfying all the equations in $\mathcal{E}$ is called a *solution* to $\mathcal{G}$.

Suppose that $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ satisfies the following properties.

1. $\mathcal{G}$ contains no cycle.
2. For $s \geq 1$ and any trail[6] $P_0 \overset{\lambda_1}{-} P_1 \overset{\lambda_2}{-} \cdots \overset{\lambda_s}{-} P_s$ in $\mathcal{G}$,

$$\lambda_1 \oplus \lambda_2 \oplus \cdots \oplus \lambda_s \neq 0^n.$$

Lemma 2. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph such that $\mathcal{G}$ satisfies the above conditions, $|\mathcal{V}| = p$, and $|\mathcal{E}| = m$. Let $h(\mathcal{G})$ be the number of solutions to $\mathcal{G}$ and $\xi_{\max}$ be the size of the largest connected component of $\mathcal{G}$. Then*

$$h(\mathcal{G}) \geq \frac{(2^n)_p}{(2^n)^m},$$

where either $p \leq 2^{\frac{n}{2}}$ or $n\xi_{\max}^2 + \xi_{\max} \leq 2^{\frac{n}{2}}$, and $1 \leq p \leq \frac{2^n}{12\xi_{\max}^2}$.

We refer to [7] for the proof of Lemma 2.

3 Almost Optimally Secure VIL-VOL PRF

3.1 eCTR: CTR-type Mode of Operation with Full Security

In this section, we propose a new block cipher-based encryption mode eCTR, and prove its security. It uses a pair of n-bit blocks (A, B) as an initial vector, and it is chosen uniformly at random from $\{0,1\}^n \times \{0,1\}^n$. The eCTR encryption mode is formally described in Algorithm 1 (see also Figure 1).

Up to the PRP-security of E, the keyed block cipher E_K can be replaced by an n-bit random permutation π in the eCTR encryption mode. Then the security of eCTR based on a random permutation π is given as follows.

Theorem 1. *Let $\mathcal{D}$ be a (q, σ, ℓ)-adversary against the* iv-prf *security of* $\mathsf{eCTR}[\pi, w]$, $\bar{\sigma} = \left\lceil \frac{(w+1)\sigma}{w} \right\rceil$ *and* $\bar{\ell} = \left\lceil \frac{(w+1)\ell}{w} \right\rceil$. *If $n(nw+1)^2 + (nw+1) \leq 2^{\frac{n}{2}}$ and $12(nw+1)^2\bar{\sigma} \leq 2^n$, we have*

$$\mathsf{Adv}^{\mathsf{iv\text{-}prf}}_{\mathsf{eCTR}[\pi,w]}(\mathcal{D}) \leq \frac{(2w+3)\bar{\sigma} + q}{2^n} + \frac{3w\bar{\ell}\bar{\sigma}^2}{2^{2n}}.$$

[6] A trail is a walk in which all edges are distinct.

Algorithm 1: $\mathsf{eCTR}[E, w]$

Input: $K \in \mathcal{K}_b$, $A \in \{0,1\}^n$, $B \in \{0,1\}^n$, s: output length in bits
Output: $Z \in \{0,1\}^*$

$Z \leftarrow \epsilon$
for $j = 1, \ldots, \lceil \frac{s}{nw} \rceil$ **do**
  $L \leftarrow E_K(A \oplus 2^{(w+1)(j-1)} B)$
  for $\alpha = 1, \ldots, w$ **do**
    $Z \leftarrow Z \,\|\, \left(L \oplus E_K(A \oplus 2^{(w+1)(j-1)+\alpha} B)\right)$
return $\mathsf{msb}_s(Z)$

3.2 Proof of Theorem 1

Transcript. Let $\mathcal{D}$ be a (q, σ, ℓ)-adversary against the iv-prf security of $\mathsf{eCTR}[\pi, w]$. Then for $i \in [q]$, $\mathcal{D}$ chooses (A_i, B_i) uniformly and independently at random from $\{0,1\}^n \times \{0,1\}^n$, makes the i-th query with (A_i, B_i, s_i), and receives Z_i of s_i bits as the response. For simplicity of proof, we will assume that s_i are multiple of n for $i \in [q]$, and let $\ell_i = s_i/n$. Then the transcript that $\mathcal{D}$ obtains at the end of the interaction is defined as

$$\tau = (A_i, B_i, Z_i[1], \ldots, Z_i[\ell_i])_{i \in [q]}$$

where $Z_i = Z_i[1] \,\|\, \cdots \,\|\, Z_i[\ell_i]$. Let $\bar{\ell}_i = \lceil (w+1)\ell_i / w \rceil$ for $i \in [q]$ and note that $\sum_{i \in [q]} \bar{\ell}_i \le \bar{\sigma}$ and $e(w+1)\bar{\sigma} \le 2^{n-1}$.

Bad Transcripts. A transcript τ satisfies bad_1 if $B_i = 0^n$ for some $i \in [q]$. Then we have

$$\Pr[\mathsf{bad}_1] \le \frac{q}{2^n}. \tag{1}$$

Assuming that $B_i \neq 0^n$ for every $i = 1, \ldots, q$ (without bad_1), we can represent a transcript as a graph with labeled edges; let

$$X_{i,j}[\alpha] = A_i \oplus 2^{(w+1)(j-1)+\alpha} B_i$$

for $i \in [q]$, $j \in [\ell_i / w]$ and $\alpha \in [w]$. Then we can define a graph $\mathcal{G}_{i,j} = (\mathcal{V}_{i,j}, \mathcal{E}_{i,j})$, where

$$\mathcal{V}_{i,j} = \{X_{i,j}[0], X_{i,j}[1], \ldots, X_{i,j}[w]\},$$
$$\mathcal{E}_{i,j} = \left\{\overline{X_{i,j}[0] X_{i,j}[\alpha]} : \alpha = 1, \ldots, w\right\},$$

and each edge $\overline{X_{i,j}[0] X_{i,j}[\alpha]}$ is labeled by $Z_i[w(j-1)+\alpha]$ for $\alpha \in [w]$. In this way, each graph $\mathcal{G}_{i,j}$ becomes a *star* of $w+1$ (distinct) vertices centered at $X_{i,j}[0]$. We will say that a star $\mathcal{G}_{i,j}$ is of *query index* i. Then all the vertices of the stars of the same query index are distinct. Finally, we define $\mathcal{G} = (\mathcal{V}, \mathcal{E})$,

where $\mathcal{V} = \bigcup \mathcal{V}_{i,j}$ and $\mathcal{E} = \bigcup \mathcal{E}_{i,j}$. Sometimes we will write $\mathcal{G}[\tau]$ to denote that it has been determined by a transcript τ.

Consider a transcript whose graph does not satisfy bad_1. With this assumption, we say that τ satisfies bad_2 if any pair of stars share two vertices in common. For two stars $\mathcal{G}_{i,j}$ and $\mathcal{G}_{i',j'}$; (with different query indices), bad_2 happens if and only if

$$A_i \oplus 2^{(w+1)(j-1)+\alpha_1} B_i = A_{i'} \oplus 2^{(w+1)(j'-1)+\alpha'_1} B_{i'},$$
$$A_i \oplus 2^{(w+1)(j-1)+\alpha_2} B_i = A_{i'} \oplus 2^{(w+1)(j'-1)+\alpha'_2} B_{i'}$$

for some α_1, α_2, α'_1, α'_2 such that $\alpha_1 \neq \alpha_2$ and $\alpha'_1 \neq \alpha'_2$. When we arbitrarily fix $A_{i'}$ and $B_{i'}$, the above system of equations holds with probability $\frac{1}{2^{2n}}$ over the random choice of A_i and B_i since the coefficients of B_i are distinct. Therefore we have

$$\Pr\left[\mathsf{bad}_2 \wedge \neg \mathsf{bad}_1\right] \leq \sum_{(i,i') \in [q]^{*2}} \left(\frac{\bar{\ell}_i w}{2}\right) \left(\frac{\bar{\ell}_{i'} w}{2}\right) \frac{1}{2^{2n}} \leq \frac{w^2 \bar{\sigma}^2}{2^{2n+2}} \tag{2}$$

If a transcript τ does not satisfy any of bad_1 and bad_2, then we can define a graph on stars $\mathcal{G}_{i,j}$, denoted $\mathcal{S}$, as following. Let $\mathcal{G}_{i_1,j_1}, \ldots, \mathcal{G}_{i_h,j_h}$ be the sequence of stars where $(i_1, j_1) \leq (i_a, j_a)$ with lexicographical order for $a \in [h]$. Then, if all $\mathcal{G}_{i_a,j_a}$ shares a (unique) vertex in common, which is called a *connecting vertex*, $\mathcal{G}_{i_1,j_1}$ and $\mathcal{G}_{i_a,j_a}$ are connected by an edge for $a = \{2, \ldots, h\}$.

Now we say that τ satisfies bad_3 if there is a trail

$$\mathcal{G}_{i_1,j_1} - \mathcal{G}_{i_2,j_2} - \cdots - \mathcal{G}_{i_{h-1},j_{h-1}} - \mathcal{G}_{i_h,j_h}$$

in $\mathcal{S}$, where $h \geq 3$, $i_1, i_2, \ldots, i_{h-1}$ are all distinct, $i_1 = i_h$, and connecting vertices are all distinct. If we fix $B_{i_2}, \ldots, B_{i_{h-1}}$ and A_{i_2} are fixed, $\mathcal{G}_{i_2,j_2}, \ldots, \mathcal{G}_{i_{h-1},j_{h-1}}$ are connected with probability at most

$$\frac{|\mathcal{G}_{i_2,j_2}| \cdot |\mathcal{G}_{i_{h-1},j_{h-1}}|}{2^{(h-3)n}} \prod_{a=3}^{h-2} |\mathcal{G}_{i_a,j_a}| \cdot (|\mathcal{G}_{i_a,j_a}| - 1)$$

Again, once $A_{i_2}, \ldots, A_{i_{h-1}}$ and $B_{i_2}, \ldots, B_{i_{h-1}}$ are fixed, in order for the first and the last connections to be made, A_{i_1} and B_{i_1} should satisfy the following system of equations.

$$A_{i_1} \oplus 2^{(w+1)(j_1-1)+\alpha_1} B_{i_1} = A_{i_2} \oplus 2^{(w+1)(j_2-1)+\alpha'} B_{i_2},$$
$$A_{i_1} \oplus 2^{(w+1)(j_h-1)+\alpha_2} B_{i_1} = A_{i_{h-1}} \oplus 2^{(w+1)(j_{h-1}-1)+\alpha''} B_{i_{h-1}}$$

for some α_1, α_2, α' and α''. If α_1, α_2, α' and α'' are fixed, the above equations hold with probability $\frac{1}{2^{2n}}$ over the random choice of A_{i_1} and B_{i_1}. So for a fixed sequence of h stars, they are connected as a chain with probability at most

$$\frac{|\mathcal{G}_{i_1,j_1}| \cdot |\mathcal{G}_{i_h,j_h}|}{2^{(h-1)n}} \prod_{a=2}^{h-1} |\mathcal{G}_{i_a,j_a}| \cdot (|\mathcal{G}_{i_a,j_a}| - 1) \leq \frac{w^{h-2} \bar{\ell}}{2^{(h-1)n}} \prod_{a=1}^{h-1} |\mathcal{G}_{i_a,j_a}|$$

Therefore we have

$$\Pr\left[\mathsf{bad}_3 \wedge \neg(\mathsf{bad}_1 \vee \mathsf{bad}_2)\right] \le \sum_{h=3}^{\infty} \frac{w^{h-2}\bar{\ell}}{2^{(h-1)n}} \sum_{(i_1,\ldots,i_{h-1})\in[q]^{(h-1)*}} \bar{\ell}_{i_1}\ldots\bar{\ell}_{i_{h-1}}$$
$$\le \frac{\bar{\ell}}{w}\cdot\sum_{h=3}^{\infty}\left(\frac{w\bar{\sigma}}{2^n}\right)^{h-1} \le \frac{2w\bar{\ell}\bar{\sigma}^2}{2^{2n}} \tag{3}$$

since $w\bar{\sigma} \le 2^{n-1}$.

Next, we define bad_4: a transcript τ satisfies bad_4 if there is a tree in $\mathcal{S}$ with $n+1$ stars. If there is such a tree, then we can fix a sequence of $n+1$ stars

$$\mathcal{G}_{i_1,j_1}, \mathcal{G}_{i_2,j_2}, \cdots, \mathcal{G}_{i_n,j_n}, \mathcal{G}_{i_{n+1},j_{n+1}}$$

where for each $h = 2,\ldots,n+1$, $\mathcal{G}_{i_h,j_h}$ is connected with exactly one star $\mathcal{G}_{i_{h'},j_{h'}}$ such that $h' < h$. If τ does not satisfy bad_3, then all the stars of the tree should have distinct query indices. In this case, the probability of connection is upper bounded by

$$\left(\frac{w+1}{2^n}\right)^n \cdot \prod_{a=1}^{n+1} |\mathcal{G}_{i_a,j_a}|$$

Since there are $(n+1)^{n-1}$ trees on $n+1$ stars by Cayley's formula, we have

$$\Pr\left[\mathsf{bad}_4 \wedge \neg(\mathsf{bad}_1 \vee \mathsf{bad}_2 \vee \mathsf{bad}_3)\right] \le \binom{\bar{\sigma}}{n+1}\frac{(n+1)^{n-1}(w+1)^n}{2^{n\cdot n}}$$
$$\le \left(\frac{e\bar{\sigma}}{n+1}\right)^{n+1}\frac{(n+1)^{n-1}(w+1)^n}{2^{n\cdot n}}$$
$$\le \frac{e\bar{\sigma}}{(n+1)^2}\left(\frac{e(w+1)\bar{\sigma}}{2^n}\right)^n \le \frac{\bar{\sigma}}{2^n} \tag{4}$$

where the last inequality comes from $e(w+1)\bar{\sigma} \le 2^{n-1}$.

Finally, we define bad_5: a transcript τ satisfies bad_5 if there is a trail in $\mathcal{G}$ whose "label sum" is zero. More precisely, let $\mathcal{P} = (Y_1,\ldots,Y_h)$ denote such a trail of distinct vertices, where $h \ge 2$, Y_i and Y_{i+1} are connected by an edge labeled W_i for $i = 1,\ldots,h-1$, and $W_1 + W_2 + \cdots + W_{h-1} = 0^n$. Since each edge of $\mathcal{P}$ is contained in a unique star in $\mathcal{S}$, $\mathcal{P}$ induces a trail in $\mathcal{S}$, and conversely, from a trail of $\mathcal{S}$, we can construct a trail in $\mathcal{G}$ by appropriately choosing the connecting vertices.

If τ does not satisfy any of bad_3 and bad_4, then there cannot be a trail of stars of length $n+1$ in $\mathcal{S}$. So we consider a sequence $\mathcal{G}_{i_1,j_1},\ldots,\mathcal{G}_{i_h,j_h}$ of h stars with distinct query indices for each h such that $1 \le h \le n$. For a fixed sequence, we have a trail with distinct connecting vertices

$$\mathcal{G}_{i_1,j_1} - \mathcal{G}_{i_2,j_2} - \cdots - \mathcal{G}_{i_{h-1},j_{h-1}} - \mathcal{G}_{i_h,j_h}$$

with probability at most

$$\prod_{a=1}^{h-1}\frac{|\mathcal{G}_{i_a,j_a}|\cdot|\mathcal{G}_{i_{a+1},j_{a+1}}|}{2^n}$$

When the stars are connected with this probability, we can choose additional two vertices, each of which is contained in $\mathcal{G}_{i_1,j_1}$ and $\mathcal{G}_{i_h,j_h}$, respectively, yielding a trail in $\mathcal{G}$. For each trail, its label sum is zero with probability $\frac{1}{2^n}$ since the labels are chosen uniformly and independently at random from $\{0,1\}^n$ (in the ideal world). Therefore, we have

$$\Pr\left[\mathsf{bad}_5 \wedge \neg \bigvee_{i=1}^{4} \mathsf{bad}_i\right] \leq \sum_{h=1}^{\infty} \left(\frac{(w+1)\bar{\sigma}}{2^n}\right)^h \leq \frac{2(w+1)\bar{\sigma}}{2^n} \tag{5}$$

since $(w+1)\sigma \leq 2^{n-1}$.

A transcript τ is defined to be *bad* if it satisfies one of bad_i, $i = 1, \ldots, 5$. Then by (1), (2), (3), (4) and (5), the probability of obtaining a bad transcript in the ideal world is upper bounded as follows.

$$\Pr\left[\mathsf{T}_{\mathsf{id}} \in \Gamma_{\mathsf{bad}}\right] \leq \frac{(2w+3)\bar{\sigma}+q}{2^n} + \frac{3(w+1)\bar{\ell}\bar{\sigma}^2}{2^{2n}}. \tag{6}$$

Analyzing Good Transcripts. If a transcript is not bad, then such a transcript is called *good*. For a good transcript τ, $\mathcal{G}[\tau]$ satisfies the following properties.

- $\mathcal{G}[\tau]$ contains no cycle since otherwise $\mathcal{G}[\tau]$ satisfies either bad_1 or bad_2 or bad_3.
- The number of vertices in the largest component of $\mathcal{G}[\tau]$ is at most $nw+1$ since otherwise $\mathcal{G}[\tau]$ satisfies bad_4.
- For any trail of $\mathcal{G}[\tau]$, its label sum is nonzero since otherwise $\mathcal{G}[\tau]$ satisfies bad_5.

Then, in the real world, $\pi(X_{i,j}[\alpha])$ should be a solution to the system of equations defined by graph $\mathcal{G}[\tau]$ and they should be all distinct. The number of such solutions is at least

$$\frac{(2^n)_p}{2^{n\sigma}}$$

by Lemma 2, if $n(nw+1)^2 + (nw+1) \leq 2^{\frac{n}{2}}$ and $p \leq \frac{2^n}{12(nw+1)^2}$, where p denotes the number of vertices in $\mathcal{G}[\tau]$ and σ is the number of edges in $\mathcal{G}[\tau]$.

The probability that π realizes each solution in the real world is $\frac{1}{(2^n)_p}$, and the probability of obtaining A_i and B_i (in the transcript), $i = 1, \ldots, q$, is $\left(\frac{1}{2^{2n}}\right)^q$. Therefore, we have

$$\Pr\left[\mathsf{T}_{\mathsf{re}} = \tau\right] \geq \frac{(2^n)_p}{2^{n\sigma}} \cdot \frac{1}{(2^n)_p} \cdot \left(\frac{1}{2^{2n}}\right)^q = \frac{1}{2^{n\sigma+2nq}}.$$

Since $\Pr\left[\mathsf{T}_{\mathsf{id}} = \tau\right] = \frac{1}{2^{n\sigma+2nq}}$, we have

$$\frac{\Pr\left[\mathsf{T}_{\mathsf{re}} = \tau\right]}{\Pr\left[\mathsf{T}_{\mathsf{id}} = \tau\right]} \geq 1. \tag{7}$$

We have Theorem 1 by Lemma 1, (6), and (7) , and since $p \leq \overline{\sigma}$.

3.3 HteC: Almost Optimally Secure VIL-VOL PRF

By hashing an arbitrary length message, encrypting the hash value with a block cipher, and using the output as an initial vector of eCTR, we can obtain a PRF. This construction is dubbed HteC. The HteC PRF is formally described in Algorithm 2 (see also Figure 2) .

Algorithm 2: $\mathsf{HteC}[H, E, w]$

Input: $(K_1, K_2, K, K') \in \mathcal{K}_h^2 \times \mathcal{K}_b^2$, $I \in \mathcal{N}$, s: output length in bits
Output: $Z \in \{0,1\}^*$

$A \leftarrow E_K(H_{K_1}(I))$
$B \leftarrow E_K(H_{K_2}(I))$
$Z \leftarrow \mathsf{eCTR}[E, w](K', A, B, s)$
return Z

Up to the PRP-security of E, the keyed block ciphers E_K and $E_{K'}$ can be replaced by two independent random permutations π and π', respectively, in the construction of HteC. When HteC is based on random permutations, denoted $\mathsf{HteC}[H, \pi, \pi', w]$, its security is given as follows.

Theorem 2. *Let H be a δ-almost universal and δ'-almost uniform hash function, and let $\mathcal{D}$ be a (q, σ, ℓ)-adversary against the* prf *security of* $\mathsf{HteC}[H, \pi, \pi', w]$. *Let* $\bar{\sigma} = \left\lceil \frac{(w+1)\sigma}{w} \right\rceil$ *and* $\bar{\ell} = \left\lceil \frac{(w+1)\ell}{w} \right\rceil$. *If* $n(nw+1)^2 + (nw+1) \le 2^{\frac{n}{2}}$ *and* $12(nw+1)^2\bar{\sigma} \le 2^n$, *we have*

$$\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{HteC}[H,\pi,\pi',w]}(\mathcal{D}) \le \delta' q + \delta^2 q^2 + \frac{2(\delta+\delta')\bar{\sigma}^2 + (2w+3)\bar{\sigma} + q}{2^n - 2q} + \frac{3w\bar{\ell}\bar{\sigma}^2}{(2^n-2q)^2}.$$

3.4 Proof of Theorem 2

Transcript. Suppose that $\mathcal{D}$ is a (q, σ, ℓ)-adversary against the prf security of $\mathsf{HteC}[H, \pi, \pi', w]$. Then for $i = 1, \ldots, q$, $\mathcal{D}$ makes the i-th query with (I_i, s_i), and obtains the output Z_i of s_i bits, where I_i is distinct for every $i = 1, \ldots, q$. For simplicity of proof, we will assume that the hash keys K_1 and K_2, and the inputs (A_i, B_i) to eCTR are given to $\mathcal{D}$ for free at the end of the interaction. Let $\ell_i = s_i/n$. In the ideal world, dummy keys K_1 and K_2 are chosen uniformly and independently at random from $\mathcal{K}_h$. Then $U_i = H_{K_1}(I_i)$ and $V_i = H_{K_2}(I_i)$ are determined for $i = 1, \ldots, q$, and (A_i, B_i) are sampled by faithfully simulating truly random permutations π, which are returned to $\mathcal{D}$. In this way, $\mathcal{D}$ will not be able to distinguish the real and ideal worlds using the additional information.

Without loss of generality, assume that s_i are multiple of n for $i \in [q]$ and let $\ell_i = s_i/n$ and $\bar{\ell}_i = \lceil (w+1)\ell/w \rceil$. Then the transcript that $\mathcal{D}$ obtains at the

end of the interaction is defined as

$$\tau = (K_1, K_2, I_i, A_i, B_i, Z_i[1], \ldots, Z_i[\ell_i])_{i\in[q]}$$

where $Z_i = Z_i[1] \,\|\, \cdots \,\|\, Z_i[\ell_i]$. Suppose that

$$\tau' = (A_i, B_i, Z_i[1], \ldots, Z_i[\ell_i])_{i\in[q]}$$

is a good (partial) transcript as defined in the proof of eCTR. Then we can also define stars $\mathcal{G}_{i,j}$, and $\mathcal{G}$ as their union, using the same notations and definitions as in the proof of eCTR.

Bad Transcripts. Besides bad_1, bad_2, bad_3, bad_4 and bad_5, we need additional conditions for a bad transcript as follows.

- $\mathsf{bad}_6 \Leftrightarrow U_i = V_i$ for some $i \in [q]$.
- $\mathsf{bad}_7 \Leftrightarrow (U_i, V_i) = (U_j, V_j)$ for some $i, j \in [q]$ such that $i < j$.
- $\mathsf{bad}_8 \Leftrightarrow$ there is a trail of stars with distinct query indices

  $$\mathcal{G}_{i_1,j_1} - \mathcal{G}_{i_2,j_2} - \cdots - \mathcal{G}_{i_{h-1},j_{h-1}} - \mathcal{G}_{i_h,j_h}$$

  for any integer $h \geq 2$, where $U_{i_1}, \ldots, U_{i_{h-1}}, V_{i_1}, \ldots, V_{i_{h-1}}$ are all distinct, and $\{U_{i_1}, V_{i_1}\} \cap \{U_{i_h}, V_{i_h}\} \neq \emptyset$.

Since H is δ-almost universal and δ'-almost uniform, we have

$$\Pr[\mathsf{bad}_6] \leq \delta' q,$$
$$\Pr[\mathsf{bad}_7] \leq \delta^2 q^2.$$

For each sequence $(i_a, j_a)_{a\in[h]}$, the probability that $\{U_{i_1}, V_{i_1}\} \cap \{U_{i_h}, V_{i_h}\} \neq \emptyset$ is at most $2(\delta + \delta')$ since H is δ-almost universal and δ'-almost uniform. Assuming that $U_{i_1}, \ldots, U_{i_{h-1}}, V_{i_1}, \ldots, V_{i_{h-1}}$ are all distinct and bad_1, bad_6 and bad_7 does not happen, $\mathcal{G}_{i_1,j_1}, \ldots, \mathcal{G}_{i_h,j_h}$ are connected with probability at most

$$\prod_{a=1}^{h-1} \frac{|\mathcal{G}_{i_a,j_a}| \cdot |\mathcal{G}_{i_{a+1},j_{a+1}}|}{2^n - 2q} \leq \left(\frac{w+1}{2^n - 2q}\right)^{h-1} \prod_{a=1}^{h} |\mathcal{G}_{i_a,j_a}|$$

since A_i's and B_i's are sampled by simulating a random permutation π and for each sampling, there are at least $2^n - 2q$ possible choices. Overall, we have

$$\begin{aligned}\Pr[\mathsf{bad}_8 \wedge \neg(\mathsf{bad}_1 \vee \mathsf{bad}_6 \vee \mathsf{bad}_7)] &\leq 2(\delta + \delta') \sum_{h=2}^{\infty} \left(\frac{w+1}{2^n - 2q}\right)^{h-1} \cdot \frac{\bar{\sigma}^h}{2} \\ &\leq \frac{2(\delta + \delta')(w+1)\bar{\sigma}^2}{2^n - 2q} \qquad (8)\end{aligned}$$

since $2(w+1)\bar{\sigma} \leq 2^n/2 \leq 2^n - 2q$.

A transcript τ is defined to be *bad* if it satisfies one of bad_i, $i = 1, \ldots, 8$. The analysis of bad_i, $i = 1, \ldots, 5$, is similar to the proof of eCTR; in particular, by

avoiding bad_6, bad_7 and bad_8, we can assume that A_i's and B_i's are all sampled *independently* by the simulation of π. Overall, the probability of obtaining a bad transcript in the ideal world is upper-bounded as

$$\Pr\left[\mathsf{T}_{\mathsf{id}} \in \Gamma_{\mathsf{bad}}\right] \leq \delta' q + \delta^2 q^2 + \frac{2(\delta+\delta')\bar{\sigma}^2 + (2w+3)\bar{\sigma} + q}{2^n - 2q} + \frac{3w\bar{\ell}\bar{\sigma}^2}{(2^n-2q)^2} \quad (9)$$

since $2e(w+1)\bar{\sigma} \leq 2^{n-1} \leq 2^n - 2q$

Analyzing Good Transcripts. For a good transcript τ, we have

$$\frac{\Pr\left[\mathsf{T}_{\mathsf{re}} = \tau\right]}{\Pr\left[\mathsf{T}_{\mathsf{id}} = \tau\right]} \geq 1 \quad (10)$$

using the same argument as in the proof of eCTR. Then we have Theorem 2 by Lemma 1, (9), and (10).

4 Highly Secure Variants of GCM and GCM-SIV

In this section, we propose new authenticated encryption schemes, dubbed eGCM and eGCM-SIV, which follow the structure of GCM and GCM-SIV, and achieve stronger security by replacing the CTR mode by eCTR.

4.1 eGCM AE Scheme

The eGCM AE scheme is based on a keyed hash function $H : \mathcal{K}_h \times \{0,1\}^* \to \{0,1\}^n$ and a block cipher $E : \mathcal{K}_b \times \{0,1\}^n \to \{0,1\}^n$. Then the key space of eGCM is defined as $\mathcal{K}_f = \mathcal{K}_b^2 \times \mathcal{K}_h^3$.

Given a key $(K_b, K_b', K_h, K_h', K_h'') \in \mathcal{K}_f$, let $\mathbf{K}_f = (K_b, K_b', K_h, K_h')$. Then eGCM encrypts a triple of a nonce, an associated data and a message $(N, A, M) \in (\{0,1\}^*)^3$ by computing $(C, T) \in \{0,1\}^* \times \{0,1\}^\tau$ such that $|M| = |C|$ as follows (see also Figure 6).

1. Generate keystreams $Z_L \in \{0,1\}^n$ and $Z_R \in \{0,1\}^*$:

$$Z_L \,\|\, Z_R = \mathsf{HteC}[H, E, w]_{\mathbf{K}_f}(N, |M| + n).$$

2. Compute ciphertext C and tag T:

$$C = M \oplus Z_R,$$
$$T = \mathsf{lsb}_\tau(Z_L \oplus H_{K_h''}(A, C)).$$

The nAE security of eGCM is summarized by the following theorem.

Theorem 3. *For $\delta > 0$, let $H : \mathcal{K}_h \times \{0,1\}^* \to \{0,1\}^n$ be δ-almost XOR universal and δ-almost uniform. Fix non-negative integers w, q_e, q_d, σ and ℓ, and let*

Algorithm 3: Encryption of $\mathsf{eGCM}[H, E, w]$

Input: $(K_b, K'_b, K_h, K'_h, K''_h) \in \mathcal{K}_b^2 \times \mathcal{K}_h^3$, $N, A, M \in \{0,1\}^*$
Output: $T \in \{0,1\}^\tau, C \in \{0,1\}^*$

// Generate keystreams
$Z \leftarrow \mathsf{HteC}[H, E, w]_{K_b, K'_b, K_h, K'_h}(N, |M| + n)$
$Z_L \leftarrow \mathsf{lsb}_n(Z)$
$Z_R \leftarrow \mathsf{msb}_{|M|}(Z)$
// Compute tag and ciphertext
$C = M \oplus Z_R$
$T = \mathsf{lsb}_\tau(Z_L \oplus H_{K''_h}(A, C))$
return C, T

Algorithm 4: Decryption of $\mathsf{eGCM}[H, E, w]$

Input: $(K_b, K'_b, K_h, K'_h, K''_h) \in \mathcal{K}_b^2 \times \mathcal{K}_h^3$, $N, A, C \in \{0,1\}^*$, $T \in \{0,1\}^\tau$
Output: $M \in \{0,1\}^*$ or $\perp$

// Generate keystreams
$Z \leftarrow \mathsf{HteC}[H, E, w]_{K_b, K'_b, K_h, K'_h}(N, |M| + n)$
$Z_L \leftarrow \mathsf{lsb}_n(Z)$
$Z_R \leftarrow \mathsf{msb}_{|M|}(Z)$
// Compute tag and plaintext
$M = C \oplus Z_R$
$T' = \mathsf{lsb}_\tau(Z_L \oplus H_{K''_h}(A, C))$
if $T = T'$ **then**
  return M
else
  return $\perp$

Fig. 6: Encryption and decryption of eGCM.

$q = q_e + q_d$ *and* $\bar{x} \lceil (w+1)x/w \rceil$ *for* $x \in \{q, \ell, \sigma\}$. *Then, we have*

$$\begin{aligned}\mathsf{Adv}^{\mathsf{nAE}}_{\mathsf{eGCM}[H,E,w]}(q_e, q_d, \sigma, \ell) \le\ & \frac{q_d}{2^\tau} \cdot (2 + \delta \cdot 2^n) + \delta q + \delta^2 q^2 \\ & + \frac{4\delta(\bar{\sigma} + \bar{q})^2 + (2w+3)(\bar{\sigma} + \bar{q}) + q}{2^n - 2q} \\ & + \frac{3w(\bar{\ell} + 1)(\bar{\sigma} + \bar{q})^2}{(2^n - 2q)^2} \\ & + 2\mathsf{Adv}^{\mathsf{prp}}_E(3q_e + 3q_d + \sigma),\end{aligned}$$

provided that

$$n(nw+1)^2 + (nw+1) \le 2^{\frac{n}{2}},$$
$$12(nw+1)^2(\bar{\sigma}+\bar{q}) \le 2^n.$$

As eGCM and GCM share the same structure, one can easily adapt the security proof of GCM [25,33,36] to eGCM, while for completeness, we give a full security proof here.

Proof. Let $\mathcal{D}$ be a (q_e, q_d, σ, ℓ)-adversary against the nAE security of eGCM. We assume that $\mathcal{D}$ does not make any redundant query and makes exactly q_e encryption queries and q_d decryption queries without loss of generality. Up to the prf-security of HteC, VIL-VOL keyed function $\mathsf{HteC}[E,H]_{\mathbf{K}_f}$ can be replaced by truly random function F. By Theorem 2, the cost of this replacement is upper bounded by

$$\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{HteC}[H,\pi,\pi',w]}(q_e+q_d, \sigma+q_e+q_d, \ell+1) \le \delta q + \delta^2 q^2 + \frac{4\delta(\bar{\sigma}+\bar{q})^2 + (2w+3)(\bar{\sigma}+\bar{q}) + q}{2^n - 2q} + \frac{3(w+1)(\bar{\ell}+1)(\bar{\sigma}+\bar{q})^2}{(2^n-2q)^2}.$$

Let

$$\tau_e = (N_i, A_i, M_i, C_i, T_i)_{i\in[q_e]},$$
$$\tau_d = (N'_j, A'_j, C'_j, T'_j, b'_j)_{j\in[q_d]}$$

denote the list of encryption queries and decryption queries, respectively. Note that $\mathcal{D}$ always has $b'_j = \perp$ for $j \in [q_d]$ if $\mathcal{D}$ interacts with the ideal oracle. At the end of the interaction, we give $\{K_h, F_n(N_1), \ldots, F_n(N_{q_e}), F_n(N'_1), \ldots, F_n(N'_{q_d})\}$ to $\mathcal{D}$ for free, where F_n takes only the first n bits from the output of F. In the ideal world,

- a dummy key K_h is selected uniformly at random from $\mathcal{K}_h$,
- for each encryption query (N, A, M, C, T), the ideal oracle sets $F_n(N) = H_{K_h}(A, C) \oplus (T \,\|\, s)$ where $s \leftarrow_{\$} \{0,1\}^{n-\tau}$,
- for each decryption query $(N', A', C', T', \perp)$, the ideal oracle sets $F_n(N') = F_n(N_i)$ if $N' = N_i$ for some $i \in [q_e]$ and $F_n(N') \leftarrow_{\$} \{0,1\}^n$ otherwise, and gives it to $\mathcal{D}$.

A transcript

$$\tau = \big(\tau_e, \tau_d, K_h, F_n(N_1), \ldots, F_n(N_{q_e}), F_n(N'_1), \ldots, F_n(N'_{q_d})\big)$$

is defined as *bad* if one of the following conditions holds.

- $\mathsf{bad}_1 \Leftrightarrow$ there exist $i \in [q_e]$ and $j \in [q_d]$ such that

$$(N_i, A_i, C_i, T_i) = (N'_j, A'_j, C'_j, T'_j).$$

– $\mathsf{bad}_2 \Leftrightarrow$ there exist $i \in [q_e]$, $j \in [q_d]$ and $s \in \{0,1\}^{n-\tau}$ such that $N_i = N'_j$, $(A_i, C_i) \neq (A'_j, C'_j)$, and

$$H_{K_h}(A_i, C_i) \oplus H_{K_h}(A'_j, C'_j) = (T_i \oplus T'_j) \parallel s$$

– $\mathsf{bad}_3 \Leftrightarrow$ there exist $j \in [q_d]$ and $s \in \{0,1\}^{n-\tau}$ such that $N'_j \notin \{N_i\}_{i\in[q_e]}$ and

$$F_n(N'_j) \oplus H_{K_h}(A'_j, C'_j) = T'_j \parallel s.$$

If a transcript τ is not bad, then it will be called a *good* transcript. The probability of obtaining a bad transcript in the ideal world is upper bounded as follows.

1. Suppose that there exist $i \in [q_e]$ and $j \in [q_d]$ such that

   $$(N_i, A_i, C_i) = (N'_j, A'_j, C'_j).$$

   Note that there are at most q_d pairs of such i and j. Since $\mathcal{D}$ does not make any redundant query, the i-th encryption query is made later than the j-th decryption query. Then, since T_i is chosen uniformly at random from $\{0,1\}^\tau$, we have

   $$\Pr[\mathsf{bad}_1] \leq \frac{q_d}{2^\tau}.$$

2. Suppose that there exist $i \in [q_e]$ and $j \in [q_d]$ such that $N_i = N'_j$ and $(A_i, C_i) \neq (A'_j, C'_j)$. Since H is δ-AXU,

   $$\Pr\left[H_{K_h}(A_i, C_i) \oplus H_{K_h}(A'_j, C'_j) = (T_i \oplus T'_j) \parallel s\right] \leq \delta$$

   for any $s \in \{0,1\}^{n-\tau}$ and we have

   $$\Pr[\mathsf{bad}_2] \leq q_d \cdot \delta \cdot 2^{n-\tau}.$$

3. Suppose that there exists $j \in [q_d]$ such that $N'_j \notin \{N_i\}_{i\in[q_e]}$. Since $F_n(N'_j)$ is chosen uniformly at random from $\{0,1\}^n$,

   $$\Pr\left[F_n(N'_j) \oplus H_{K_h}(A'_j, C'_j) = T'_j \parallel s\right] = \frac{1}{2^n}$$

   for any $s \in \{0,1\}^{n-\tau}$. Thus, we have

   $$\Pr[\mathsf{bad}_3] \leq \frac{q_d}{2^\tau}.$$

All in all, we have

$$\Pr[\mathsf{T}_{\mathsf{id}} \in \Gamma_{\mathsf{bad}}] \leq \frac{q_d}{2^\tau} \cdot (2 + \delta \cdot 2^n). \tag{11}$$

Fix a good transcript

$$\tau = \left(\tau_e, \tau_d, K_h, F_n(N_1), \ldots, F_n(N_{q_e}), F_n(N'_1), \ldots, F_n(N'_{q_d})\right),$$

and let $\mathcal{N}_e = \{N_i\}_{i\in[q_e]}$ and $\mathcal{N}_d = \{N'_j\}_{j\in[q_d]}$. Then, one can easily see that

$$\Pr[\tau = \mathsf{T}_{\mathsf{id}}] = \left(\prod_{i\in[q_e]} \frac{1}{2^{|M_i|+\tau}}\right) \cdot \frac{1}{|\mathcal{K}_h|} \cdot \left(\frac{1}{2^{n-\tau}}\right)^{q_e} \cdot \left(\frac{1}{2^n}\right)^{|\mathcal{N}_d \setminus \mathcal{N}_e|}.$$

Since we have $\neg(\mathsf{bad}_1 \vee \mathsf{bad}_2 \vee \mathsf{bad}_3)$, for all $j \in [q_d]$,

$$\mathsf{lsb}_\tau(F_n(N'_j) \oplus H_{K_h}(A'_j, C'_j)) \neq T'_j,$$

and hence, $b'_j = \perp$ is always compatible with F. Thus, we have

$$\Pr[\tau = \mathsf{T}_{\mathsf{re}}] = \left(\prod_{i\in[q_e]} \frac{1}{2^{|M_i|+n}}\right) \cdot \frac{1}{|\mathcal{K}_h|} \cdot \left(\frac{1}{2^n}\right)^{|\mathcal{N}_d \setminus \mathcal{N}_e|}$$

and hence,

$$\frac{\Pr[\mathsf{T}_{\mathsf{re}} = \tau]}{\Pr[\mathsf{T}_{\mathsf{id}} = \tau]} = 1. \tag{12}$$

The proof is complete by (11), (12), and Lemma 1. □

Remark 1. We can prove stronger security of CHM, CIP and mGCM in a similar way as they all can be seen as a combination of CENC and a universal hash function. Given an n-bit block cipher E, a key K, and a nonce $N \in \{0,1\}^{n-s}$, an $(iw+j)$-th output block of $\mathsf{CENC}[E,w]$ is defined as

$$E_K(N \,\|\, \langle i(w+1)\rangle_s) \oplus E_K(N \,\|\, \langle i(w+1)+j\rangle_s)$$

where $\langle x\rangle_s$ denotes the s-bit representation of x (when $x < 2^s$). By Lemma 2 the output blocks cannot be distinguished from random strings except when 0^n is returned, provided that $n(w+1)^2+(w+1) \le 2^{n/2}$ and $\frac{(w+1)^3}{w}\sigma \le \frac{2^n}{12}$. Therefore, for a small enough w, one can easily prove that

$$\mathsf{Adv}^{\mathsf{prf}}_{\mathsf{CENC}[E,w]}(q,\sigma,\ell) \le O\left(\frac{\sigma}{2^n}\right)$$

and

$$\mathsf{Adv}^{\mathsf{nAE}}_{\Pi}(q_e,q_d,\sigma,\ell) \le O\left(\frac{\sigma}{2^n} + \frac{q_d\ell}{2^\tau}\right)$$

for $\Pi \in \{\mathsf{CHM}, \mathsf{CIP}, \mathsf{mGCM}\}$, where the underlying hash functions in these schemes are assumed to have $O(\ell/2^n)$-almost XOR universality.

4.2 eGCM-SIV AE Scheme

The eGCM-SIV AE scheme is built on top of a keyed hash function $H : \mathcal{K}_h \times \{0,1\}^* \to \{0,1\}^n$ and a block cipher $E : \mathcal{K}_b \times \{0,1\}^n \to \{0,1\}^n$. Then the key space of eGCM is defined as $\mathcal{K}_f = \mathcal{K}_b^3 \times \mathcal{K}_h^2$.

Given a key $(K_b, K'_b, K''_b, K_h, K'_h) \in \mathcal{K}_f$, let $\mathbf{K}_f = (K_b, K'_b, K_h, K'_h)$. Then eGCM-SIV encrypts a triple of a nonce, an associated data and a message $(N, A, M) \in (\{0,1\}^*)^3$ by computing $(C,T) \in \{0,1\}^* \times \{0,1\}^\tau$ such that $|M| = |C|$ as follows (see also Figure 7).

Algorithm 5: Encryption of eGCM-SIV$[H, E, w]$

Input: $(K_b, K'_b, K''_b, K_h, K'_h) \in \mathcal{K}_b^3 \times \mathcal{K}_h^2$, $N, A, M \in \{0,1\}^*$
Output: $T \in \{0,1\}^\tau, C \in \{0,1\}^*$

// Generate the tag
$T \leftarrow \mathsf{HteC}[H, E, 2]_{K_b, K'_b, K_h, K'_h}((N, A, M), 2n)$

// Compute the keystream and the ciphertext
$Z = \mathsf{eCTR}[E, w]_{K''_b}(T, |M|)$
$C = M \oplus Z$
return C, T

Algorithm 6: Decryption of eGCM-SIV$[H, E]$

Input: $(K_b, K'_b, K''_b, K_h, K'_h) \in \mathcal{K}_b^3 \times \mathcal{K}_h^2$, $N, A, C \in \{0,1\}^*$, $T \in \{0,1\}^\tau$
Output: $M \in \{0,1\}^*$ or $\perp$

// Compute the keystream and plaintext
$Z = \mathsf{eCTR}[E, w]_{K''_b}(T, |C|)$
$M = C \oplus Z$

// Compute the tag
$T \leftarrow \mathsf{HteC}[H, E, w]_{K_b, K'_b, K_h, K'_h}((N, A, M), 2n)$

if $T = T'$ **then**
  return M
else
  return $\perp$

Fig. 7: Encryption and decryption of eGCM-SIV.

1. Generate tag $T \in \{0,1\}^{2n}$:
$$T = \mathsf{HteC}[H, E]_{\mathbf{K}_f}(N, A, M; 2n).$$
2. Compute keystream Z and ciphertext C:
$$Z = \mathsf{eCTR}[E]_{K''_b}(T; |M|),$$
$$C = M \oplus Z.$$

The mrAE security of eGCM-SIV is summarized by the following theorem.

Theorem 4. *For $\delta > 0$, let $H : \mathcal{K}_h \times \{0,1\}^* \to \{0,1\}^n$ be δ-almost XOR universal and δ-almost uniform. Fix non-negative integers w, q_e, q_d, σ and ℓ, and let $q = q_e + q_d$ and $\bar{x} \lceil (w+1)x/w \rceil$ for $x \in \{\ell, \sigma\}$. Then we have*

$$\mathsf{Adv}^{\mathsf{mrAE}}_{\mathsf{eGCM\text{-}SIV}[H,E,w]}(q_e, q_d, \sigma, \ell) \le \delta q + \delta^2 q^2 + \frac{144\delta q^2 + 23q}{2^n - 4q} + \frac{648q^2}{(2^n - 4q)^2} + \frac{(2w+3)\bar{\sigma} + q_e}{2^n} + \frac{3(w+1)\bar{\ell}\bar{\sigma}^2}{2^{2n}} + \frac{q_d}{2^{2n}}$$

provided that

$$n(nw+1)^2 + (nw+1) \le 2^{\frac{n}{2}},$$
$$12(nw+1)\bar{\sigma} \le 2^n,$$
$$72(2n+1)q \le 2^n.$$

As eGCM-SIV follows the generic structure of SIV, one can easily modify the security proof of SIV [15], obtaining

$$\mathsf{Adv}^{\mathsf{mrAE}}_{\mathsf{eGCM\text{-}SIV}[H,E]}(q_e, q_d, \sigma, \ell) \le \mathsf{Adv}^{\mathsf{prf}}_{\mathsf{HteC}[H,E,2]}(2(q_e+q_d), 4(q_e+q_d), 2) + \mathsf{Adv}^{\mathsf{iv\text{-}prf}}_{\mathsf{eCTR}[E,w]}(q_e, \sigma, \ell) + \frac{q_d}{2^{2n}}$$

and combine it with Theorem 1 and Theorem 2.

References

1. Bellare, M., Rogaway, P., Wagner, D.: The EAX Mode of Operation. In: Roy, B., Meier, W. (eds.) Fast Software Encryption - FSE 2004. LNCS, vol. 3017, pp. 389–407. Springer (2004). https://doi.org/10.1007/978-3-540-25937-4_25, https://iacr.org/archive/fse2004/30170391/30170391.pdf
2. Bhattacharjee, A., Bhaumik, R., Dutta, A., Nandi, M., Raychaudhuri, A.: Bbb security for 5-round even-mansour-based key-alternating feistel ciphers. Designs, Codes and Cryptography **92**(1), 13–49 (2024)
3. Bhattacharya, S., Nandi, M.: Revisiting variable output length xor pseudorandom function. IACR Transactions on Symmetric Cryptology pp. 314–335 (2018)
4. Böck, H., Zauner, A., Devlin, S., Somorovsky, J., Jovanovic, P.: Nonce-Disrespecting adversaries: Practical forgery attacks on GCM in TLS. In: 10th USENIX Workshop on Offensive Technologies (WOOT 16). USENIX Association, Austin, TX (Aug 2016), https://www.usenix.org/conference/woot16/workshop-program/presentation/bock
5. Chakraborti, A., Datta, N., Jha, A., Mancillas-López, C., Nandi, M.: Light-ocb: parallel lightweight authenticated cipher with full security. In: International Conference on Security, Privacy, and Applied Cryptography Engineering. pp. 22–41. Springer (2021)
6. Choi, W., Lee, B., Lee, J., Lee, Y.: Toward a fully secure authenticated encryption scheme from a pseudorandom permutation. In: Advances in Cryptology – ASIACRYPT 2021. pp. 407–434. Springer International Publishing, Cham (2021)
7. Cogliati, B., Dutta, A., Nandi, M., Patarin, J., Saha, A.: Proof of mirror theory for a wide range of ξ max. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 470–501. Springer (2023)
8. Cogliati, B., Ethan, J., Lallemand, V., Lee, B., Lee, J., Minier, M.: Ctet+: A beyond-birthday-bound secure tweakable enciphering scheme using a single pseudorandom permutation. IACR Transactions on Symmetric Cryptology pp. 1–35 (2021)
9. Daemen, J.: Aes proposal: Rijndael (1999)
10. Dobraunig, C., Matusiewicz, K., Mennink, B., Tereschenko, A.: Efficient instances of docked double decker with AES, and application to authenticated encryption. Cryptology ePrint Archive, Paper 2024/084 (2024), https://eprint.iacr.org/2024/084

11. Dutta, A., Nandi, M., Talnikar, S.: Beyond birthday bound secure mac in faulty nonce model. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 437–466. Springer (2019)
12. Gueron, S.: Double Nonce Derive Key AES-GCM (DNDK-GCM). Internet-Draft draft-gueron-cfrg-dndkgcm-01, Internet Engineering Task Force (Oct 2024), https://datatracker.ietf.org/doc/draft-gueron-cfrg-dndkgcm/01/, work in Progress
13. Gueron, S., Langley, A., Lindell, Y.: AES-GCM-SIV: Nonce Misuse-Resistant Authenticated Encryption. RFC 8452 (Apr 2019). https://doi.org/10.17487/RFC8452, https://rfc-editor.org/rfc/rfc8452.txt
14. Gueron, S., Lindell, Y.: GCM-SIV: Full Nonce Misuse-Resistant Authenticated Encryption at Under One Cycle per Byte. In: Ray, I. (ed.) ACM SIGSAC Conference on Computer and Communications Security - CCS 2015. pp. 109–119. Association for Computing Machinery (2015)
15. Gueron, S., Lindell, Y.: Gcm-siv: Full nonce misuse-resistant authenticated encryption at under one cycle per byte. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. p. 109–119. CCS '15, Association for Computing Machinery, New York, NY, USA (2015). https://doi.org/10.1145/2810103.2813613
16. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust Authenticated-Encryption AEZ and the Problem That It Solves. In: Oswald, E., Fischlin, M. (eds.) Advances in Cryptology - EUROCRYPT 2015 (Proceedings, Part I). LNCS, vol. 9056, pp. 15–44. Springer (2015). https://doi.org/10.1007/978-3-662-46800-5_2
17. Hoang, V.T., Tessaro, S.: Key-Alternating Ciphers and Key-Length Extension: Exact Bounds and Multi-user Security. In: Advances in Cryptology - CRYPTO 2016 (Proceedings, Part I). LNCS, vol. 9814, pp. 3–32. Springer (2016). https://doi.org/10.1007/978-3-662-53018-4_1
18. Iso/iec 19772:2020 information security — authenticated encryption (2020)
19. Iwata, T.: New blockcipher modes of operation with beyond the birthday bound security. In: Robshaw, M. (ed.) Fast Software Encryption. pp. 310–327. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
20. Iwata, T.: New blockcipher modes of operation with beyond the birthday bound security. In: Fast Software Encryption: 13th International Workshop, FSE 2006, Graz, Austria, March 15-17, 2006, Revised Selected Papers 13. pp. 310–327. Springer (2006)
21. Iwata, T.: Authenticated Encryption Mode for Beyond the Birthday Bound Security. In: Vaudenay, S. (ed.) Progress in Cryptology - AFRICACRYPT 2008. LNCS, vol. 5023, pp. 125–142. Springer (2008)
22. Iwata, T., Khairallah, M., Minematsu, K., Peyrin, T.: Duel of the titans: the romulus and remus families of lightweight aead algorithms. IACR Transactions on Symmetric Cryptology pp. 43–120 (2020)
23. Iwata, T., Minematsu, K.: Stronger security variants of gcm-siv. IACR Transactions on Symmetric Cryptology pp. 134–157 (2016)
24. Iwata, T., Minematsu, K., Peyrin, T., Seurin, Y.: Zmac: a fast tweakable block cipher mode for highly secure message authentication. In: Annual international cryptology conference. pp. 34–65. Springer (2017)
25. Iwata, T., Ohashi, K., Minematsu, K.: Breaking and repairing gcm security proofs. In: Safavi-Naini, R., Canetti, R. (eds.) Advances in Cryptology – CRYPTO 2012. pp. 31–49. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
26. Iwata, T., Seurin, Y.: Reconsidering the security bound of aes-gcm-siv. Cryptology ePrint Archive (2017)

27. Iyengar, J., Thomson, M.: QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000 (May 2021). https://doi.org/10.17487/RFC9000, https://www.rfc-editor.org/info/rfc9000
28. Joux, A.: 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
29. Kampanakis, P., Campagna, M., Crocket, E., Petcher, A., Gueron, S.: Practical challenges with aes-gcm and the need for a new cipher. In: Third NIST Workshop on Block Cipher Modes of Operation (2023)
30. Krovetz, T., Rogaway, P.: The software performance of authenticated-encryption modes. In: Fast Software Encryption: 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers 18. pp. 306–327. Springer (2011)
31. Lonvick, C.M., Ylonen, T.: The Secure Shell (SSH) Transport Layer Protocol. RFC 4253 (Jan 2006). https://doi.org/10.17487/RFC4253, https://www.rfc-editor.org/info/rfc4253
32. Mattsson, J.P.: Collision attacks on galois/counter mode (GCM). Cryptology ePrint Archive, Paper 2024/1111 (2024), https://eprint.iacr.org/2024/1111
33. McGrew, D.A., Viega, J.: The security and performance of the galois/counter mode (gcm) of operation. In: Progress in Cryptology - INDOCRYPT 2004. pp. 343–355. Springer Berlin Heidelberg, Berlin, Heidelberg (2005)
34. Nir, Y., Langley, A.: ChaCha20 and Poly1305 for IETF protocols. RFC 7539 (2015)
35. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800-38D (2007), National Institute of Standards and Technology.
36. Niwa, Y., Ohashi, K., Minematsu, K., Iwata, T.: Gcm security bounds reconsidered. In: FSE. pp. 385–407. Springer (2015). https://doi.org/10.1007/978-3-662-48116-5_19, https://www.iacr.org/archive/fse2015/85400168/85400168.pdf
37. Patarin, J.: Luby-rackoff: 7 rounds are enough for 2 n (1- ε) security. In: Advances in Cryptology-CRYPTO 2003: 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003. Proceedings 23. pp. 513–529. Springer (2003)
38. Patarin, J.: On linear systems of equations with distinct variables and small block size. In: Information Security and Cryptology-ICISC 2005: 8th International Conference, Seoul, Korea, December 1-2, 2005, Revised Selected Papers 8. pp. 299–321. Springer (2006)
39. Patarin, J.: The coefficients h technique. In: International Workshop on Selected Areas in Cryptography. pp. 328–345. Springer (2008)
40. Peyrin, T., Seurin, Y.: Counter-in-Tweak: Authenticated Encryption Modes for Tweakable Block Ciphers. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology - CRYPTO 2016 (Proceedings, Part I). LNCS, vol. 9814, pp. 33–63. Springer (2016). https://doi.org/10.1007/978-3-662-53018-4_2
41. Rescorla, E.: The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446 (Aug 2018). https://doi.org/10.17487/RFC8446, https://www.rfc-editor.org/info/rfc8446
42. Rescorla, E., Dierks, T.: The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (Aug 2008). https://doi.org/10.17487/RFC5246, https://www.rfc-editor.org/info/rfc5246

43. Rogaway, P.: Authenticated-encryption with associated-data. In: Proceedings of the 9th ACM Conference on Computer and Communications Security. p. 98–107. CCS '02, Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/586110.586125, https://doi.org/10.1145/586110.586125
44. Rogaway, P., Shrimpton, T.: A Provable-Security Treatment of the Key-Wrap Problem. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer (2006). https://doi.org/10.1007/11761679_23, https://iacr.org/archive/eurocrypt2006/40040377/40040377.pdf
45. Vanhoef, M., Piessens, F.: Key reinstallation attacks: Forcing nonce reuse in WPA2. In: Proceedings of the 24th ACM Conference on Computer and Communications Security (CCS). ACM (2017)
46. Wang, P., Feng, D., Wu, W.: Hctr: A variable-input-length enciphering mode. In: International Conference on Information Security and Cryptology. pp. 175–188. Springer (2005)
47. Whiting, D., Housley, R., Ferguson, N.: Counter with CBC-MAC (CCM). Submission to NIST (2002), available at https://csrc.nist.gov/groups/ST/toolkit/BCM/documents/proposedmodes/ccm/ccm.pdf