

Message Franking via Committing Authenticated Encryption

Paul Grubbs¹(✉), Jiahui Lu², and Thomas Ristenpart¹

¹ Cornell Tech, New York, USA
pag225@cornell.edu

² Shanghai Jiao Tong University, Shanghai, China

Abstract. We initiate the study of message franking, recently introduced in Facebook’s end-to-end encrypted message system. It targets verifiable reporting of abusive messages to Facebook without compromising security guarantees. We capture the goals of message franking via a new cryptographic primitive: compactly committing authenticated encryption with associated data (AEAD). This is an AEAD scheme for which a small part of the ciphertext can be used as a cryptographic commitment to the message contents. Decryption provides, in addition to the message, a value that can be used to open the commitment. Security for franking mandates more than that required of traditional notions associated with commitment. Nevertheless, and despite the fact that AEAD schemes are in general not committing (compactly or otherwise), we prove that many in-use AEAD schemes can be used for message franking by using secret keys as openings. An implication of our results is the first proofs that several in-use symmetric encryption schemes are committing in the traditional sense. We also propose and analyze schemes that retain security even after openings are revealed to an adversary. One is a generalization of the scheme implicitly underlying Facebook’s message franking protocol, and another is a new construction that offers improved performance.

Keywords: Authenticated encryption · Encrypted messaging

1 Introduction

Encrypted messaging systems are now used by more than a billion people, due to the introduction of popular, industry-promoted products including WhatsApp [60], Signal [61], and Facebook Messenger [30]. These use specialized (non-interactive) key exchange protocols, in conjunction with authenticated encryption, to protect messages. Many tools are based off the Signal protocol [44], which itself was inspired by elements of the off-the-record (OTR) messaging protocol [20]. A primary design goal is end-to-end security: intermediaries including the messaging service providers, or those with access to their systems, should not be able to violate confidentiality or integrity of user messages.

Scheme	MO security	Sender binding	Rec. binding	Enc	Dec	Ver
Encode-then-Encipher (Ideal)		✓	✓	–	–	–
Encrypt-then-HMAC (one key)		✓	✓	2+1	2+1	2+1
HMAC-then-CBC		✓	✓	2+1	2+1	2+1
CtE1	✓	✓	✓	3+1	3+1	1+1
CtE2 (Facebook)	✓	✓	✓	3+2	3+2	1+1
CEP	✓	✓	✓	2+1	2+1	1+1

Fig. 1. Summary of schemes investigated in this work. The columns indicate whether the scheme meets multiple-opening (MO) security, sender binding, and receiver binding. The last three columns indicate the number of cryptographic passes over a bit string of length equal to the message plus the number of passes needed to handle the associated data, for each of the three main operations. We omit comparisons with concrete encode-then-encipher constructions, which vary in the number of passes required.

End-to-end security can be at odds with other security goals. A well known example is dealing with filtering and reporting spam in the context of encrypted email [39, 56]. Similar issues arise in modern encrypted messaging systems. For example, in Facebook’s system when one user sends harassing messages, phishing links, malware attachments, etc., the recipient should be able to report the malicious behavior so that Facebook can block or otherwise penalize the sender. But end-to-end confidentiality means that Facebook must rely on users sending examples of malicious messages. How can the provider know that the reported message was the one sent? Reports could, in turn, become a vector for abuse should they allow a malicious reporter to fabricate a message and convince the provider it was the one sent [39].

Facebook messenger recently introduced a seeming solution for verifiable abuse reporting that they refer to as *message franking* [31, 47]. The idea is to include in the report a cryptographic proof that the reported message was the one sent, encrypted, by the particular sender. They offer a protocol (discussed below) and a sensible, but informal and vague, discussion of security goals. At present it is ultimately not clear what message franking provides, whether their approach is secure, and if there exist better constructions. Given the critical role message franking will play for most messaging services moving forward, more study is clearly needed.

We therefore initiate the formal study of message franking. We introduce the notion of compactly committing authenticated encryption with associated data (AEAD) as the cryptographic primitive of merit that serves as the basis for message franking. We provide security definitions, show how several widely used existing AEAD schemes can already serve as compactly committing AEAD, give an analysis of (a generalization of) the scheme underlying Facebook’s protocol, and design a new scheme that has superior performance. A summary of schemes treated in this work, and their efficiency, is shown in Fig. 1.

Facebook’s message franking protocol. Facebook’s protocol works as follows, modulo a few details (see Sect. 3). A sender first generates a fresh key for HMAC [3], and applies HMAC to the message. It then encrypts the HMAC key and message using a conventional AEAD scheme with a symmetric key shared with (just) the recipient, and sends along the resulting ciphertext and the hash value to Facebook’s servers. Facebook signs the hash and forwards on the whole package — signature, HMAC hash, and ciphertext — to the recipient, who decrypts and checks the validity of the HMAC output using the recovered HMAC key. Should the recipient want to report abuse, their software client sends the signature, message, HMAC hash, and HMAC key to Facebook who can now verify the signature and hash.

While descriptions of Facebook’s protocol do not use the term commitment, intuitively that is the role played by HMAC. This may suggest viewing message franking as simply a construction of committing encryption [23]. But committing encryption views the entire ciphertext as the commitment and opens ciphertexts by revealing the secret key. Neither is true of the Facebook scheme.

A new primitive: compactly committing AEAD. We introduce a new cryptographic primitive that captures the properties targeted in verifiable abuse reporting. We refer to it as compactly committing AEAD. This is an AEAD scheme for which a small portion of the ciphertext can be used as a commitment to the message. Decryption reveals an opening for the message, and the scheme comes equipped with an additional verification algorithm that can check the commitment. This formalization has some similarity to one for non-AEAD symmetric encryption due to Gertner and Herzberg [36], but differs in important ways and, in short, their treatment does not suffice for message franking (see Sect. 9 for more detailed comparisons).

Formalizing security for committing AEAD schemes requires care. Informally we want confidentiality, ciphertext integrity, and that the ciphertexts are binding commitments to their underlying plaintexts. While seemingly a straightforward adaptation of real-or-random style confidentiality and ciphertext integrity notions would suffice [52, 54, 55], this turns out to provide only a weaker form of security in which reporting abuse may invalidate security of the encryption moving forward. In short, this is because the opening might reveal cryptographic key material, e.g., if the secret key is itself used as the opening. We refer to this as single-opening (SO) security. We formalize also multiple-opening (MO) security notions which, in addition to the usual challenge oracles, gives the adversary the ability to obtain regular encryptions and decryptions (which, by our syntax, reveals the opening should a ciphertext be valid) under the target key. Analogously to previous AEAD treatments [55], we formalize this both via an all-in-one security game that simultaneously establishes confidentiality and integrity, and as separate notions for confidentiality and integrity. We prove them equivalent.

Standard integrity notions like INT-CTXT do not by themselves imply that the ciphertext is a binding commitment to the underlying message. We introduce a notion called receiver binding, which is similar to the binding notions from the commitment literature, notions from the robust encryption literature [1, 32, 33],

and the prior notion of binding for committing encryption due to Gertner and Herzberg. Importantly, we deal with the fact that only a portion of the ciphertext is committing, and other details such as associated data. Achieving receiver binding means that no computationally limited adversary can find two opening, message pairs that verify for the same committing portion of a ciphertext.

At first glance this seemed like the end of the story with regards to binding security. But in the message franking setting, schemes that are only receiver binding may spectacularly fail to ensure verifiable abuse reporting. In particular, we show how schemes that are receiver binding still allow the following attack: a sender carefully chooses a ciphertext so that an abusive message is correctly decrypted by the receiver, but verification with the resulting opening of that message fails. Such an attack is devastating and arises quickly without careful design. We give an example of a natural performance improvement for the Facebook scheme that probably enjoys confidentiality, ciphertext integrity, and receiver binding, yet subtly falls to this attack. We therefore formalize and target meeting a *sender* binding property that rules out such attacks.

Legacy schemes. With formal notions in place, we start by investigating whether existing, in-use AEAD schemes are compactly committing. For these legacy schemes the opening is taken to be the secret key and per-message randomness used, and in each case we identify a small portion of the ciphertext to take as the committing portion. In this context proving receiver binding also proves the scheme to be committing in the more traditional sense.

As mentioned, AEAD schemes are not in general binding via simple counterexamples. We therefore analyze specific constructions, focusing on three important schemes. The first, Encode-then-Encipher [12], uses a variable-input-length tweakable block cipher to build an authenticated encryption scheme by padding messages with randomness and redundancy information (zero bits). We show that, modeling the underlying tweakable cipher as ideal, one can show that taking a security-parameter number of bits of the ciphertext as the commitment is both receiver and sender binding. Verification re-encrypts the message and checks that the resulting ciphertext properly matches the commitment value.

We next investigate Encrypt-then-MAC constructions [9], which are particularly relevant here given that Signal [44], and in turn Facebook messenger, uses AES-CBC followed by HMAC for authenticated encryption of messages. In practice, one uses a key-derivation function to derive an encryption key and a MAC key. Interestingly, if one uses as opening those two separate keys, then a simple attack shows that this scheme is *not* receiver binding. If, however, one uses the input to the KDF as the opening, we can prove receiver binding assuming the KDF and MAC are collision resistant. Notably this rules out using CMAC [41], PMAC [18], and Carter-Wegman MACs [59], but Encrypt-then-HMAC suffices.

This means that in Facebook messenger the underlying encryption already suffices as a single-opening-secure committing AEAD scheme. Moreover, due to ratcheting [14, 26, 45] Signal never reuses a symmetric key. Thus Facebook could have avoided the dedicated HMAC commitment. Admittedly they may be uncomfortable — for reason of psychological acceptability — with an

architecture that sends decryption keys to Facebook despite the fact that this represents no harm to future or past communications.

We finally investigate MAC-then-Encrypt, the mode of operation underlying TLS 1.2 and before. The binding properties of MAC-then-Encrypt were briefly investigated in a recent paper that used TLS 1.2 records as commitments [58], including a brief proof sketch of receiver binding when taking the entire ciphertext as the commitment. We expand on their proof sketch and provide a full proof for the scheme instantiated with CBC-mode and HMAC (the instantiation used in TLS), taking a small constant number of ciphertext blocks as the committing portion. Interestingly this proof, unlike that of Encrypt-then-MAC, required modeling the block cipher underlying CBC-mode as an ideal cipher and HMAC as a random oracle [28].

Commit-then-Encrypt constructions. We next turn to analyzing generic constructions that combine a commitment with an existing AE scheme. We provide a generalization of the Facebook scheme, and show that it is multiple-opening secure and both sender and receiver binding, assuming only that the underlying AEAD scheme is sound and the commitment is unique. HMAC is a unique commitment, thereby giving us the first formal security analysis of Facebook’s message franking scheme. One can also use a non-malleable commitment [29]. If one instead uses a malleable commitment, then the scheme will not achieve ciphertext integrity.

We also offer an alternative composition that removes the need for non-malleable commitments, and also can improve performance in the case that associated data is relatively long. Briefly, we use a commitment to the associated data and message as the associated data for the underlying AEAD scheme. This indirectly binds the encryption ciphertext to the associated data, without paying the cost of twice processing it.

Both these constructions are multiple-opening secure, since the commitment opening is independent of the underlying AE keys. This is intuitively simple but the proof requires care — commitments play a role in achieving CTXT and so we must take care to show that unopened encryptions, despite using the same keys as opened encryptions, retain ciphertext integrity. See the body for details.

The Committing Encrypt-and-PRF (CEP) scheme. The generic constructions that meet multiple-opening security are slower than existing (single-opening secure) AEAD schemes, since they require an additional cryptographic pass over the message. This represents approximately a 1.5x slowdown both for encryption and decryption. For the expected workload in messaging applications that consists primarily of relatively short plaintexts, this may not matter, but if one wants to use committing AEAD for large plaintexts such as image and video attachments or in streaming settings (e.g., a committing version of TLS) the overhead will add up quickly.

We therefore offer a new AEAD scheme, called Committing Encrypt-and-PRF (CEP) that simultaneously enjoys multiple-opening security while also retaining the two-pass performance of standard AEAD schemes. As an additional bonus we make the scheme nonce-based [54], meaning that it is derandomized

and only needs to be used with non-repeating nonces. (We formalize nonce-based committing AEAD in the body; it is largely similar as the randomized variant.)

The basic idea is to adapt an Encrypt-and-PRF style construction to be compactly committing and multiple-opening. To do so we derive one-time use PRF keys from the nonce, and compute a tag that is two-part. The commitment value for the ciphertext is the output of a keyed hash that is simultaneously a PRF when the key is private and collision resistant when it is adversarially chosen. The latter is critical since receiver binding requires, in this context, a collision-resistance property. If one stopped here, then the scheme would not be secure, since openings reveal the PRF’s key, rendering it only CR, and CR is not enough to prevent future ciphertext forgeries. We therefore additionally run a one-time PRF (with key that is never opened) over this commitment value to generate a tag that is also checked during decryption. Ultimately we prove that the scheme achieves our notions of sender binding, receiver binding, and multiple-opening confidentiality and ciphertext integrity.

We strove to make the scheme simple and fast. Instantiated with a stream cipher such as AES-CTR-mode or ChaCha20, we require just a single secret key and use the stream cipher to generate not only the one-time keys for the PRFs but also a pad for encrypting the message. Because we need a collision-resistant PRF, our suggested instantiation is HMAC, though other multi-property hash functions [10] would work as well.

Future directions. Our work has focused on the symmetric encryption portion of messaging protocols, but one can also ask how the landscape changes if one holistically investigates the public-key protocols or key exchange in particular. Another important direction is to understand the potential tension between committing AEAD and security in the face of selective opening attacks (SOA) [7,8]. Our current definitions do not model SOAs. (An SOA would allow, for example, a compromise of the full cryptographic key, not just the ability to get openings.) While it may seem that committing encryption and SOA security are at odds, we actually conjecture that this is not fundamental (particularly in the random oracle model), and future work will be able to show SOA-secure compactly committing AEAD.

2 Preliminaries

We fix some alphabet Σ , e.g., $\Sigma = \{0,1\}$. For any $x \in \Sigma^*$ let $|x|$ denote its length. We write $x \leftarrow_s X$ to denote uniformly sampling from a set X . We write $X \parallel Y$ to denote concatenation of two strings. For a string X of n bits, we will write $X[i, \dots, j]$ for $i < j \leq n$ to mean the substring of X beginning at index i and ending at index j . For notational simplicity, we assume that one can unambiguously parse $Z = X \parallel Y$ into its two parts, even for strings of varying length. For strings $X, Y \in \{0,1\}^*$ we write $X \oplus Y$ to denote taking the XOR of $X[1, \dots, \min\{|X|, |Y|\}] \oplus Y[1, \dots, \min\{|X|, |Y|\}]$.

We use code-based games [13] to formalize security notions. A game G is a sequence of pseudocode statements, with variables whose type will be clear from

context. Variables are implicitly initialized to appropriate defaults for their type (zero for integers, empty set for sets, etc.). Each variable is a random variable in the probability distribution defined by the random coins used to execute the game. We write $\Pr[G \Rightarrow y]$ to denote the event that the game outputs a value y . Associated to this pseudocode is some fixed RAM model of computation where most operations are unit cost. We will use “big-O” notation $\mathcal{O}(\cdot)$ to hide only small constants that do not materially impact the interpretation of our results.

We will work in the random oracle model (ROM) [11] and the ideal cipher model (ICM). In the ROM, algorithms and adversaries are equipped with an oracle that associates to each input a random output of some length that will vary by, and be clear from, context. In the ICM, algorithms and adversaries are equipped with a pair of oracles. The first takes input a key, a tweak, and a message, all bit strings of some lengths k , t , and n , respectively. Each key, tweak pair selects a random permutation on $\{0, 1\}^n$. The second oracle takes as input a key, a tweak, and an n -bit value, and returns the inverse of the permutation selected by the key and tweak applied to the value.

Below, we will only discuss the time complexity of a reduction if bounding it is non-trivial. Otherwise we will omit discussions of time complexity.

Symmetric encryption. A nonce-based authenticated encryption (AE) scheme $\text{SE} = (\text{Kg}, \text{enc}, \text{dec})$ consists of a triple of algorithms. Associated to it are a key space $\mathcal{K} \subseteq \Sigma^*$, nonce space $\mathcal{N} \subseteq \Sigma^*$, header space $\mathcal{H} \subseteq \Sigma^*$, message space $\mathcal{M} \subseteq \Sigma^*$, and ciphertext space $\mathcal{C} \subseteq \Sigma^*$. The randomized key generation algorithm Kg outputs a secret key $K \in \mathcal{K}$. Canonically Kg selects $K \leftarrow_s \mathcal{K}$ and outputs K . Encryption enc is deterministic and takes as input a four-tuple $(K, N, H, M) \in (\Sigma^*)^4$ and outputs a ciphertext C or a distinguished error symbol \perp . We require that $\text{enc}(K, N, H, M) \neq \perp$ if $(K, N, H, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{M}$. Decryption dec is deterministic and takes as input a tuple $(K, N, H, C) \in (\Sigma^*)^4$ and outputs a message M or \perp .

An SE scheme is correct if for any $(K, N, H, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{M}$ it holds that $\text{dec}(K, N, H, \text{enc}(K, N, H, M)) = M$.

Some schemes that we will analyze predate the viewpoint of nonce-based encryption, including generic compositions that utilize CTR or CBC mode. A randomized SE scheme $\text{SE} = (\text{Kg}, \text{enc}, \text{dec})$ is the same as a nonce-based SE scheme except that we omit nonces everywhere, and have enc take an additional input, the coins, that are assumed to be drawn from some coin space $\mathcal{R} \subseteq \sigma^*$. Correctness now is met if for any $(K, H, M, R) \in \mathcal{K} \times \mathcal{H} \times \mathcal{M} \times \mathcal{R}$ it holds that $\text{dec}(K, H, \text{enc}(K, H, M; R)) = M$. We will focus on schemes that are public-coin, meaning the ciphertext includes R explicitly. This is true, for example, of CTR or CBC mode encryption. For notational simplicity, we will assume for such schemes that enc outputs R concatenated with the remainder of the ciphertext.

Pseudorandom functions. For a function $F : \mathcal{K} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ and adversary \mathcal{A} we define the *pseudorandom function* (PRF) advantage of \mathcal{A} to be

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) = \left| \Pr \left[K \leftarrow_s \mathcal{K} : \mathcal{A}^{F(K, \cdot)} = 1 \right] - \Pr \left[R \leftarrow_s \text{Func} : \mathcal{A}^{R(\cdot)} = 1 \right] \right| .$$

We define **Func** to be the space of all functions that output n bits.¹ Informally, we say the function F is a PRF if $\mathbf{Adv}_F^{\text{prf}}()$ is small for all efficient adversaries. Below we will sometimes refer to the left-hand experiment as the “real world” and the other as the “ideal world”.

In proofs it will be convenient to use multi-user PRF security [4]. We define the MU-PRF advantage of an adversary \mathcal{A} to be

$$\mathbf{Adv}_F^{\text{mu-prf}}(\mathcal{A}) = \left| \Pr \left[\mathcal{A}^{\overline{F}(\cdot, \cdot)} \Rightarrow 1 \right] - \Pr \left[\mathcal{A}^{\overline{R}(\cdot, \cdot)} \Rightarrow 1 \right] \right| .$$

where \overline{F} on input a key identifier $S \in \{0, 1\}^*$ and a message M , checks if there is a key associated to S , and if not chooses a fresh one $K[S] \leftarrow_s \{0, 1\}^k$. It then returns $F(K[S], M)$. The oracle \overline{R} on input a key identifier $S \in \{0, 1\}^*$ and a message M , checks if there is a random function associated to S , and if not chooses a fresh one $R[S] \leftarrow_s \text{Func}$. It returns $R[S](M)$. Note that MU-PRF security is implied by PRF security via a standard argument.

Collision-resistance. For a function $F : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ and adversary \mathcal{A} , define the *collision-resistance* (CR) advantage as

$$\mathbf{Adv}_F^{\text{cr}}(\mathcal{A}) = \Pr \left[((x_1, x_2), (x'_1, x'_2)) \leftarrow_s \mathcal{A} : \begin{array}{l} F(x_1, x_2) = F(x'_1, x'_2), \\ (x_1, x_2) \neq (x'_1, x'_2) \end{array} \right] .$$

Informally, we say F is collision-resistant if $\mathbf{Adv}_F^{\text{cr}}()$ is small for all efficient adversaries.

Commitment schemes with verification. A commitment scheme with verification $\text{CS} = (\text{Com}, \text{VerC})$ consists of two algorithms. Associated to any commitment scheme is an opening space $\mathcal{K}_f \subseteq \Sigma^*$, a message space $\mathcal{M} \subseteq \Sigma^*$, and a commitment space $\mathcal{C} \subseteq \Sigma^*$. The algorithm **Com** is randomized and takes as input a $M \in \Sigma^*$ and outputs a pair $(K, C) \in \mathcal{K}_f \times \mathcal{C}$ or an error symbol \perp . We assume that **Com** returns \perp with probability one if $M \notin \mathcal{M}$. The algorithm **VerC** is deterministic. It takes input a tuple $(K, C, M) \in \Sigma^*$ and outputs a bit. We assume that **VerC** returns 0 if its input $(K, C, M) \notin \mathcal{K}_f \times \mathcal{C} \times \mathcal{M}$. We assume that the commitment values C are of some fixed length (typically denoted by t).

A commitment scheme (with verification) is correct if for all $M \in \mathcal{M}$ it holds that $\Pr[\text{VerC}(\text{Com}(M), M) = 1] = 1$ where the probability is over the coins used by **Com**. We will not use the alternate definition of commitments with opening [21]. We can formalize the binding security notion of our commitment scheme as a game. Formally, the game $\text{vBIND}_{\text{CS}, \mathcal{A}}$ first runs the adversary \mathcal{A} who outputs a tuple (K_c, M, K'_c, M', C) . The game then runs $b \leftarrow \text{VerC}(K_c, C, M)$ and $b' \leftarrow \text{VerC}(K'_c, C, M')$. The game outputs true if $M \neq M'$ and $b = b' = 1$ and false otherwise. To a commitment scheme CS and adversary \mathcal{A} we associate the vBIND advantage

$$\mathbf{Adv}_{\text{CS}}^{\text{v-bind}}(\mathcal{A}) = \Pr[\text{vBIND}_{\text{CS}, \mathcal{A}} \Rightarrow \text{true}] .$$

¹ We are abusing the formalism here by sampling R from an infinite set; we do so for notational consistency and simplicity.

The probability is over the coins used by the game.

Commitment schemes should enjoy a hiding property as well. Traditionally this is formalized as a left-or-right indistinguishability notion (q.v., [6]). For our purposes we will target a stronger notion, analogous to real-or-random (ROR) security for symmetric encryption. It asks that a commitment be indistinguishable from a random bit string while the opening remaining secret. Game $\text{ROR1}_{\text{CS},\mathcal{A}}$ runs an adversary \mathcal{A} and gives it access to an oracle **Com** to which it can query messages. The oracle computes $(K_c, C) \leftarrow^* \text{Com}(M)$ and returns C . The adversary outputs a bit, and the game outputs **true** if the bit is one. Game $\text{ROR0}_{\text{CS},\mathcal{A}}$ is similar except that the oracle returns a string of random bits of length $|C|$ and the game outputs **true** if the adversary outputs zero. We define the advantage by $\text{Adv}_{\text{CS}}^{\text{cs-ror}}(\mathcal{A}) = |\Pr[\text{ROR1}_{\text{CS},\mathcal{A}} \Rightarrow \text{true}] - \Pr[\text{ROR0}_{\text{CS},\mathcal{A}} \Rightarrow \text{true}]|$.

HMAC is a good commitment. Any PRF that is also collision-resistant meets our security goals for commitments. In particular, one can build a commitment scheme $\text{CS}[F] = (\text{Com}, \text{VerC})$ works from any function $F: \mathcal{K} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ as follows. Commitment $\text{Com}(M)$ chooses a fresh value $K \leftarrow^* \mathcal{K}$, computes $C \leftarrow F(K, M)$ and outputs (K, C) . Verification $\text{VerC}(K, C, M)$ outputs one if $F(K, M) = C$ and zero otherwise. Then the following theorem captures the security of this commitment scheme, which rests on the collision resistance and PRF security of F . A proof of this theorem appears in the full version of the paper.

Theorem 1. *Let F be a function and $\text{CS}[F]$ be the commitment scheme built from it as described above. Then for any efficient adversaries \mathcal{A} making at most q queries in game ROR and \mathcal{A}' in game vBIND respectively, there exists a pair of adversaries $\mathcal{B}, \mathcal{B}'$ so that*

$$\text{Adv}_{\text{CS}[F]}^{\text{cs-ror}}(\mathcal{A}) \leq q \cdot \text{Adv}_F^{\text{prf}}(\mathcal{B}) \quad \text{and} \quad \text{Adv}_{\text{CS}[F]}^{\text{v-bind}}(\mathcal{A}') \leq \text{Adv}_F^{\text{cr}}(\mathcal{B}').$$

The adversary \mathcal{B} runs in time that of \mathcal{A} and makes the same number of oracle queries as \mathcal{A} . Adversary \mathcal{B}' runs in time that of \mathcal{A}' .

As the underlying function needs to be both CR and a good PRF, a suitable candidate would be HMAC [5], i.e., $F(K, M) = \text{HMAC}(K, M)$. Other multi-property hash functions [10] could be used as well. The Facebook franking scheme (discussed in Sect. 3) uses a non-standard HMAC-based commitment based on $F(K, M) = \text{HMAC}(K, M \parallel K)$. We will assume HMAC remains a PRF when used in this non-standard way. One can substantiate this assumption directly in the random oracle model, or using techniques from the key-dependent message literature [19, 37].

3 Message Franking and End-to-End Encryption

In end-to-end encrypted messaging services there exists a tension between message privacy and reporting abusive message contents to service providers.

The latter is important to flag abusive accounts, but reports need to be verifiable, meaning that the provider can check the contents of the allegedly abusive message *and* be certain that it was the message sent. Otherwise abuse-reporting mechanisms could themselves be abused to make false accusations.

A recipient can send the allegedly abusive plaintext to the service provider, but message privacy guarantees that the provider does not know whether the alleged message was in fact the one sent.² A seeming solution would be for the service to log ciphertexts, and have the recipient disclose the secret key to allow the provider to decrypt the ciphertext. Not only is this impractical due to the storage requirements, but it also does *not* guarantee that the decrypted message is correct. It could be that the recipient chose a key that somehow decrypts the (legitimate) ciphertext to a fake message. Ultimately what is required for this to work is for the encryption to be committing: no computationally efficient adversary can find a secret key that decrypts the ciphertext to anything but the originally encrypted message.

Facebook’s approach. Facebook recently detailed a new cryptographic mechanism [31, 47] targeting verifiable abuse reporting on Facebook messenger, which uses end-to-end encryption based on Signal [61]. The basic idea is to force the sender to provide a commitment, sent in the clear, to the plaintext message. A diagram of Facebook’s protocol, that they call “message franking” (as in “speaking frankly”), is shown in Fig. 2. The sender first applies HMAC with a fresh key K_f to the concatenation of the message and K_f to produce a value C_2 , and then encrypts using an AEAD scheme the message and K_f to produce a ciphertext C_1 using a key K_r shared with the recipient. Then (C_1, C_2) is sent to Facebook. Facebook applies HMAC with its own secret key K_{FB} to C_2 to get a tag a , and sends to the recipient (C_1, C_2, a) . The recipient decrypts C_1 , recovers the message M and key K_f and checks the value $C_2 = \text{HMAC}(K_f, M \parallel K_f)$. To report abuse, the recipient sends M , K_f , and a to Facebook. Facebook recomputes $\text{HMAC}(K_f, M \parallel K_f)$ and checks the tag a .

It is clear that the sender is using HMAC as a cryptographic commitment to the message. (This terminology is not used in their technical specifications.) The use of HMAC by Facebook to generate the tag a is simply to forego having to store commitments, instead signing them so that they can be outsourced to recipients for storage and verified should an abuse report come in.

There are interesting security issues that could arise with Facebook’s scheme, and cryptographic abuse reporting in general, that are orthogonal to the ones discussed here. In particular, binding Facebook’s tag to the communicating parties seems crucial: otherwise a malicious party could create a sock-puppet (i.e. fake) account, send itself an abusive message, then accuse a victim of having sent it.

While the design looks reasonable, and the Facebook white paper provides some informal discussion about security, there has been no formal analysis to

² Of course, if the recipient is running a trusted client, then this assertion could be trusted. We are concerned with the case that the client is subverted.




Fig. 2. Facebook’s message franking protocol [47]. The key K_r is a one-time-use symmetric key derived as part of the record layer protocol. The top portion is the sending of an encrypted message to the recipient. The bottom portion is the abuse reporting protocol.

date. It is also not clear what security properties the main cryptographic construction — combining a commitment with AEAD — should satisfy. We rectify this by introducing, in the following section, the notion of committing AEAD. This will allow us not only to analyze Facebook’s franking scheme, but to suggest alternative designs, including ones that are legacy-compatible with existing deployed AEAD schemes and do not, in particular, require adding an additional dedicated commitment.

4 Committing AEAD

Formally, a committing AEAD scheme $\text{CE} = (\text{Kg}, \text{Enc}, \text{Dec}, \text{Ver})$ is a four-tuple of algorithms. Associated to a scheme is a key space $\mathcal{K} \subseteq \Sigma^*$, header space $\mathcal{H} \subseteq \Sigma^*$, message space $\mathcal{M} \subseteq \Sigma^*$, ciphertext space $\mathcal{C} \subseteq \Sigma^*$, opening space $\mathcal{K}_f \subseteq \Sigma^*$, and franking tag space $\mathcal{T} \subseteq \Sigma^*$.

- **Key generation:** The randomized key generation algorithm Kg outputs a secret key $K \in \mathcal{K}$. We write $K \leftarrow \text{Kg}$ to denote executing key generation.
- **Encryption:** Encryption Enc is randomized. The input to encryption is a triple $(K, H, M) \in (\Sigma^*)^3$ and the output is a pair $(C_1, C_2) \in \mathcal{C} \times \mathcal{T}$ or a distinguished error symbol \perp . Unlike with regular symmetric encryption, the output includes two components: a ciphertext C_1 and a franking tag C_2 . We also refer to C_2 as the commitment. We require that $\text{Enc}(K, H, M) \neq \perp$ if $(K, H, M) \in \mathcal{K} \times \mathcal{H} \times \mathcal{M}$. We write $(C_1, C_2) \leftarrow \text{Enc}(K, H, M)$ to denote executing encryption.
- **Decryption:** Decryption, which is deterministic, takes as input a tuple $(K, H, C_1, C_2) \in (\Sigma^*)^4$ and outputs a message, opening value pair $(M, K_f) \in \mathcal{M} \times \mathcal{K}_f$ or \perp . We write $(M, K_f) \leftarrow \text{Dec}(K, H, C_1, C_2)$ to denote executing decryption.
- **Verification:** Verification, which is deterministic, takes as input a tuple $(H, M, K_f, C_2) \in (\Sigma^*)^4$ and outputs a bit. For $(H, M, K_f, C_2) \notin \mathcal{H} \times \mathcal{M} \times \mathcal{K}_f \times \mathcal{T}$, we assume that Ver outputs 0. We write $b \leftarrow \text{Ver}(H, M, K_f, C_2)$ to denote executing verification.

We will often place K in the subscript of relevant algorithms. For example, $\text{Enc}_K(H, M) = \text{Enc}(K, H, M)$ and $\text{Dec}_K(H, C_1, C_2) = \text{Dec}(K, H, C_1, C_2)$.

We require that CE schemes output ciphertexts whose lengths are determined solely by the length of the header and message. Formally this means that there exists a function $\text{clen} : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}$ such that for all $(K, H, M) \in \mathcal{K} \times \mathcal{H} \times \mathcal{M}$ it holds that $\Pr[(|C_1|, |C_2|) = \text{clen}(|H|, |M|)] = 1$ where $(C_1, C_2) \leftarrow^s \text{Enc}_K(H, M)$ and the probability is over the coins used by encryption.

We say a CE scheme has *decryption correctness* if for all $(K, H, M) \in \mathcal{K} \times \mathcal{H} \times \mathcal{M}$ it holds that $\Pr[\text{Dec}(K, H, C_1, C_2) = M] = 1$ where the probability is taken over the coins used to compute $(C_1, C_2) \leftarrow^s \text{Enc}(K, H, M)$.

We say that a scheme has *commitment correctness* if for all $(K, H, M) \in \mathcal{K} \times \mathcal{H} \times \mathcal{M}$ it holds that $\Pr[\text{Ver}(H, M, K_f, C_2) = 1] = 1$ where the probability is taken over the random variables used in the experiment

$$(C_1, C_2) \leftarrow^s \text{Enc}_K(H, M) ; (M, K_f) \leftarrow \text{Dec}_K(H, C_1, C_2) ; \text{Return } (K_f, C_2)$$

Our formulation of CE schemes is a generalization of that for conventional (randomized) AE schemes in the following sense. One can consider an AE scheme as a CE scheme that has encryption output the entire ciphertext as C_2 , decryption output an empty string for the opening value, and has verify always return one.

Compactly committing AEAD. In our formalism, a ciphertext has two components. A scheme may output $C_1 = \varepsilon$ and a C_2 value that therefore consists of the entire ciphertext. This embodies the traditional viewpoint on committing AEAD, in which the entire ciphertext is viewed as the commitment. But we are more general, and in particular our formalism allows schemes with *compact* commitments, by which we mean schemes for which $|C_2|$ is small. In particular we will want $|C_2|$ to be linear in the security-parameter, rather than linear in the message length. One can make any CE scheme compact by hashing the ciphertext with a collision-resistant (CR) hash function, as we show formally in a moment. But we will also show compact schemes that have better performance.

Single versus multiple openings. In some protocols, we may wish to use a CE scheme so that multiple different ciphertexts, encrypted under the same secret key, can be opened without endangering the privacy or integrity of other unopened ciphertexts. In other contexts, the CE scheme’s opening need only be “single-use” — the secret key will not continue to be used after an opening. An example of the latter is Signal, which due to ratcheting effectively has a fresh secret key per message. As we will now discuss, whether one wants single-opening or multiple-opening CE must be reflected in the security definitions.

Confidentiality. We want our CE schemes to provide message confidentiality. We will in fact adapt the stronger real-or-random notion from the AE literature (q.v., [55]) to CE. At a high level we ask that no adversary can distinguish between legitimate CE encryptions and (pairs of) random bit strings. A complexity arises in the multi-opening case, where we want confidentiality to hold even after openings occur. We handle this by giving the attacker an additional pair of oracles, one for encryption and decryption. We must take care to avoid trivial

<p><u>MO-REAL_{CE,A}:</u> $K \leftarrow \\$ \text{Kg}$ $b' \leftarrow \\$ \mathcal{A}^{\text{Enc,Dec,ChalEnc}}$ Return b'</p> <p><u>Enc(H, M)</u> $(C_1, C_2) \leftarrow \\$ \text{Enc}_K(H, M)$ $\mathcal{Y}_1 \leftarrow \mathcal{Y}_1 \cup \{(H, C_1, C_2)\}$ Return (C_1, C_2)</p> <p><u>Dec(H, C_1, C_2)</u> If $(H, C_1, C_2) \notin \mathcal{Y}_1$ then Return \perp $(M, K_f) \leftarrow \text{Dec}_K(H, C_1, C_2)$ Return (M, K_f)</p> <p><u>ChalEnc(H, M)</u> $(C_1, C_2) \leftarrow \\$ \text{Enc}_K(H, M)$ Return (C_1, C_2)</p>	<p><u>MO-RAND_{CE,A}:</u> $K \leftarrow \\$ \text{Kg}$ $b' \leftarrow \\$ \mathcal{A}^{\text{Enc,Dec,ChalEnc}}$ Return b'</p> <p><u>Enc(H, M)</u> $(C_1, C_2) \leftarrow \\$ \text{Enc}_K(H, M)$ $\mathcal{Y}_1 \leftarrow \mathcal{Y}_1 \cup \{(H, C_1, C_2)\}$ Return (C_1, C_2)</p> <p><u>Dec(H, C_1, C_2)</u> If $(H, C_1, C_2) \notin \mathcal{Y}_1$ then Return \perp $(M, K_f) \leftarrow \text{Dec}_K(H, C_1, C_2)$ Return (M, K_f)</p> <p><u>ChalEnc(H, M)</u> $(\ell_1, \ell_2) \leftarrow \text{clen}(H , M)$ $(C_1, C_2) \leftarrow \\$ \{0, 1\}^{\ell_1} \times \{0, 1\}^{\ell_2}$ Return (C_1, C_2)</p>	<p><u>MO-CTXT_{CE,A}:</u> $K \leftarrow \\$ \text{Kg}$; win \leftarrow false $\mathcal{A}^{\text{Enc,Dec,ChalDec}}$ Return win</p> <p><u>Enc(H, M)</u> $(C_1, C_2) \leftarrow \\$ \text{Enc}_K(H, M)$ $\mathcal{Y} \leftarrow \mathcal{Y} \cup \{(H, C_1, C_2)\}$ Return (C_1, C_2)</p> <p><u>Dec(H, C_1, C_2)</u> Return $\text{Dec}_K(H, C_1, C_2)$</p> <p><u>ChalDec($H, C_1, C_2$)</u> If $(H, C_1, C_2) \in \mathcal{Y}$ then Return \perp $(M, K_f) \leftarrow \text{Dec}_K(H, C_1, C_2)$ If $M \neq \perp$ then win \leftarrow true Return (M, K_f)</p>
---	---	--

Fig. 3. Confidentiality (left two games) and ciphertext integrity (rightmost) games for committing AEAD.

wins, of course, separating use of the real oracles from the challenge ones. We also additionally require that the adversary can only query its decryption oracle on valid ciphertexts returned from the encryption oracle. This all is formalized in the games MO-REAL_{CE,A} and MO-RAND_{CE,A} shown in Fig. 3. We measure the multiple-openings real-or-random (MO-ROR) advantage of an adversary \mathcal{A} against a scheme CE by

$$\text{Adv}_{\text{CE}}^{\text{mo-ror}}(\mathcal{A}) = |\Pr[\text{MO-REAL}_{\text{CE,A}} \Rightarrow 1] - \Pr[\text{MO-RAND}_{\text{CE,A}} \Rightarrow 1]| .$$

The single-opening ROR (SO-ROR) games REAL_{CE,A} and RAND_{CE,A} are identical to MO-REAL_{CE,A} and MO-RAND_{CE,A} in Fig. 3 except that we omit the **Enc** and **Dec** oracles. We measure the single-openings real-or-random (ROR) advantage of an adversary \mathcal{A} against a scheme CE by

$$\text{Adv}_{\text{CE}}^{\text{ror}}(\mathcal{A}) = |\Pr[\text{REAL}_{\text{CE,A}} \Rightarrow 1] - \Pr[\text{RAND}_{\text{CE,A}} \Rightarrow 1]| .$$

Ciphertext integrity. We also want our CE schemes to enjoy ciphertext integrity. As with confidentiality, we will lift the standard (randomized) AEAD security notions to the multiple-opening and single-opening CE settings. The game MO-CTXT_{CE,A} is shown in Fig. 3. The adversary can obtain encryptions and decryptions under the secret key, and its goal is to query a valid ciphertext to a challenge decryption oracle. That ciphertext must not have been returned by the encryption oracle. We measure the multiple-openings ciphertext integrity (MO-CTXT) advantage of an adversary \mathcal{A} against a scheme CE by

$$\text{Adv}_{\text{CE}}^{\text{mo-ctxt}}(\mathcal{A}) = \Pr[\text{MO-CTXT}_{\text{CE,A}} \Rightarrow \text{true}] .$$

<p>s-BIND_{CE}^A:</p> <p>$(K, H, C_1, C_2) \leftarrow \mathcal{A}$</p> <p>$(M', K_f) \leftarrow \text{Dec}(K, H, C_1, C_2)$</p> <p>If $M' = \perp$ then Return false</p> <p>$b \leftarrow \text{Ver}(H, M', K_f, C_2)$</p> <p>If $b = 0$ then</p> <p style="padding-left: 20px;">Return true</p> <p>Return false</p>	<p>r-BIND_{CE}^A:</p> <p>$((H, M, K_f), (H', M', K'_f), C_2) \leftarrow \mathcal{A}$</p> <p>$b \leftarrow \text{Ver}(H, M, K_f, C_2)$</p> <p>$b' \leftarrow \text{Ver}(H', M', K'_f, C_2)$</p> <p>If $(H, M) = (H', M')$ then</p> <p style="padding-left: 20px;">Return false</p> <p>Return $(b = b' = 1)$</p>
---	---

Fig. 4. Binding security games for committing AEAD. Sender binding (left game) models a setting where a malicious sender wants to send a message, but prevent commitment opening from succeeding. Receiver binding (right game) models a setting where a sender and recipient collude to open a ciphertext to different messages.

As with confidentiality, we can also specify a single-opening version of security by removing the decryption oracle **Dec** from game MO-CTXT_{CE,A}. Let the resulting game be CTXT_{CE,A}. We measure the single-openings ciphertext integrity (CTXT) advantage of an adversary \mathcal{A} against a scheme CE by

$$\text{Adv}_{\text{CE}}^{\text{ctxt}}(\mathcal{A}) = \Pr[\text{CTXT}_{\text{CE},\mathcal{A}} \Rightarrow \text{true}].$$

All-in-one notions. We have given separate confidentiality and ciphertext integrity notions. As with traditional AEAD security, however, we can alternatively give an all-in-one notion that simultaneously captures confidentiality and integrity goals. We defer the details to the full version of this work.

Security for AEAD. Given the fact that CE schemes encompass (randomized) AEAD schemes as well (see our comments above), we note that the RORand CTXTnotions apply to standard (randomized) AE schemes. As a slight abuse of notation, we will therefore use RORand CTXTand their associated games and advantage measures for the security of traditional AE schemes.

Binding security notions. We introduce two security notions for binding: *sender binding* and *receiver binding*. Sender binding ensures the sender of a message is bound to the message it actually sent. In abuse-reporting scenarios, this prevents the sender of an abusive message from generating a bogus commitment that does not give the receiver the ability to report the message. The pseudocode game s-BINDon the left-hand-side of Fig. 4 formalizes this requirement. To an adversary \mathcal{A} and CE scheme CE we associate the “sender binding” advantage

$$\text{Adv}_{\text{CE}}^{\text{s-bind}}(\mathcal{A}) = \Pr[\text{s-BIND}_{\text{CE},\mathcal{A}} \Rightarrow \text{true}].$$

The probability is over the coins used in the game.

A CE scheme can generically meet sender binding by running **Ver** during **Dec** and having **Dec** return \perp if **Ver** returns 0. We omit the proof of this, which follows by inspection. But legacy AEAD schemes do not do this, and one needs to check sender binding. For new schemes we will see more efficient ways to achieve sender binding.

The second security notion, receiver binding, is a lifting of the more traditional binding notion from commitment schemes (see Sect. 2). This definition is important in abuse reporting, where it formalizes the intuition that a malicious receiver should not be able to accuse a non-abusive sender of having said something abusive. A malicious receiver could do this by opening one of the sender’s ciphertexts to an abusive message instead of the one the sender intended.

The pseudocode game r-BIND is shown on the right in Fig. 4. It has an adversary output a pair of triples containing associated data, a message, and an opening. The adversary outputs a franking tag C_2 as well. The adversary wins if verification succeeds on both triples with C_2 and the header/message pairs differ. To a CE scheme CE and adversary \mathcal{A} we associate the “receiver binding” advantage

$$\text{Adv}_{\text{CE}}^{\text{r-bind}}(\mathcal{A}) = \Pr[\text{r-BIND}_{\text{CE},\mathcal{A}} \Rightarrow \text{true}] .$$

The probability is over the coins used in the game.

It is important to note that r-BIND security does not imply s-BIND security. These notions are, in fact, orthogonal. Moreover, our MO-RORand MO-CTXT notions do not generically imply either of the binding notions.

Discussion. Our definitions also allow associated data, sometimes referred to as headers. This puts committing AEAD on equal footing with modern authenticated encryption with associated data (AEAD) schemes [52], which require it. That said, modern AEAD schemes are most often formalized as nonce-based, meaning that instead of allowing internal randomness, a non-repeating value (the nonce) is an explicit input and encryption is deterministic. Existing systems relevant to abuse complaints use randomized AEAD (e.g., Signal [44]) that do not meet nonce-based AEAD security. That said, we will explore nonce-based committing AEAD in Sect. 7.

5 Are Existing AEAD Schemes Committing?

In this section we will study whether existing AEAD schemes meet our security goals for CE. We believe it is important to study legacy schemes for several reasons. If existing AEAD schemes are also committing, it will have important positive and negative implications for deployed protocols (such as OTR or Facebook’s franking scheme) that implicitly rely on binding (or non-binding) properties of symmetric encryption. It is also helpful for protocol designers who may want to build a protocol on top of existing legacy encryption. If well-tested, mature implementations of AEAD can be used as CE schemes without code changes, the attack surface of new protocol implementations is minimized.

In this section we only examine the binding properties of schemes, since past work has shown they meet standard definitions for confidentiality and integrity. We will prove that encode-then-encipher and encrypt-then-MAC (EtM) satisfy our binding notions in the ideal cipher model, with the additional requirement that the MAC used in EtM is a collision-resistant PRF. We will prove MAC-then-encrypt meets our binding notions in the random oracle and ideal cipher

$\text{Enc}(K, H, M):$ $R \leftarrow \{0, 1\}^r$ $C \leftarrow \tilde{E}_K^H(M \parallel R \parallel 0^s)$ $\ell \leftarrow l + r + s - t$ $C_1 \leftarrow C[1, \dots, \ell]$ $C_2 \leftarrow C[\ell + 1, \dots, l + r + s]$ $\text{Return } (C_1, C_2)$	$\text{Dec}(K, H, C_1, C_2):$ $M' \parallel R' \parallel Z \leftarrow \tilde{D}_K^H(C_1 \parallel C_2)$ $\text{If } Z \neq 0^s \text{ then}$ $\quad \text{Return } \perp$ $\text{Return } (M', (R', K))$	$\text{Ver}(H, M, K_f, C_2):$ $R \parallel K \leftarrow K_f$ $\ell \leftarrow l + r + s - t$ $C \leftarrow \tilde{E}_K^H(M \parallel R \parallel 0^s)$ $\text{Return } C[\ell + 1, \dots, l + r + s] = C_2$
--	--	---

Fig. 5. Encode-then-encipher as a committing AEAD scheme where the commitment is the final t bits of the ciphertext. \tilde{E}^H and \tilde{D}^H refer to encryption and decryption for a tweakable blockcipher where the header H is the tweak.

model. We will also show simple attacks that break binding for real-world modes using Carter-Wegman MACs (GCM and ChaCha20/Poly1305).

5.1 Committing Encode-then-Encipher

The Encode-then-Encipher (EtE) construction of Bellare and Rogaway shows how to achieve AE security for messages given only a variable-input-length PRP [12]. Their construction is quite simple: given a key $K \in \mathcal{K}$, encrypt a message $M \in \mathcal{M}$ ($|\mathcal{M}| = 2^l$) with header $H \in \mathcal{H}$ by first drawing a random string $R \leftarrow \{0, 1\}^r$ and computing $c = \tilde{E}_K^H(M \parallel R \parallel 0^s)$ where \tilde{E}^H is a tweakable, variable-input length cipher with the header as the tweak. Decrypting a ciphertext M works by first running $M' = \tilde{D}_K^H(C)$ and checking whether the last s bits of M' are all zero. If they are, we call the message “valid” and output M , else we output \perp . For compactness, we commit to only the last t bits of the ciphertext. We must include the randomness used to encrypt in the opening of the commitment. Detailed pseudocode is given in Fig. 5. We will assume that E is an ideal tweakable cipher in our proof of r-BIND security.

Theorem 2. *Let $\text{EtE}[E]$ be the scheme defined above using an ideal tweakable cipher E and parameters $s, t > 0$. Let \mathcal{A} be any adversary making at most q queries to its ideal cipher oracles. Then $\mathbf{Adv}_{\text{EtE}}^{\text{r-bind}}(\mathcal{A}) \leq \frac{q+1}{2^s} + \frac{q^2}{2^r}$.*

The proof will appear in the full version of this work. The scheme achieves perfect s-BIND security: the advantage of any adversary for is zero because the output of decryption is simply re-computed in Ver .

5.2 Encrypt-then-MAC

The classic Encrypt-then-MAC (EtM) construction composes a symmetric encryption scheme and a message authentication code (MAC), by first encrypting the message, then computing the MAC over the ciphertext and any associated data.

Committing EtM. We analyze EtM as a committing AEAD scheme in the case that the encryption and authentication keys are derived via a key derivation function (KDF) that is a collision-resistant pseudorandom function. The scheme $\text{EtM}[\text{KDF}, F, \text{SE}]$ is detailed in Fig. 6. Beyond the functions F and KDF, the scheme also makes use of a public-coin randomized symmetric encryption algorithm $\text{SE} = (\text{Kg}, \text{enc}, \text{dec})$ that does not use associated data and whose key generation is a random selection of some fixed-length bit string. It is important that the scheme is public coin, as we require the randomness to be recoverable during decryption to be included in the opening.

$\text{Enc}(K, H, M):$ $K^e \leftarrow \text{KDF}_K(0)$ $K^m \leftarrow \text{KDF}_K(1)$ $R \leftarrow \mathcal{R}$ $R \parallel C \leftarrow \text{enc}_{K^e}(M; R)$ $T \leftarrow F_{K^m}(H \parallel R \parallel C)$ $\text{Return } (R \parallel C, T)$	$\text{Dec}(K, H, C_1, C_2):$ $R \parallel C \leftarrow C_1$ $K^e \leftarrow \text{KDF}_K(0)$ $K^m \leftarrow \text{KDF}_K(1)$ $T' \leftarrow F_{K^m}(H \parallel R \parallel C_1)$ $\text{If } T' \neq C_2 \text{ then Return } \perp$ $M \leftarrow \text{dec}_{K^e}(C_1)$ $\text{If } M = \perp \text{ then Return } \perp$ $\text{Return } (M, (R, K))$	$\text{Ver}(H, M, (R, K), C_2):$ $K^e \leftarrow \text{KDF}_K(0)$ $K^m \leftarrow \text{KDF}_K(1)$ $C \leftarrow \text{enc}_{K^e}(M; R)$ $T \leftarrow F_{K^m}(H \parallel R \parallel C)$ $\text{Return } T = C_2$
---	---	---

Fig. 6. Committing AEAD scheme $\text{EtM}[\text{KDF}, F, \text{SE}]$ that composes an encryption scheme $\text{SE} = (\text{Kg}, \text{enc}, \text{dec})$ using random coins from \mathcal{R} , a MAC F , and that derives keys via a function KDF.

This scheme arises in practice. The Signal protocol [44], for example, uses HKDF to derive keys for use with CTR mode encryption combined with HMAC. The following theorem proves the committing EtM construction in Fig. 6 meets r-BIND if the MAC and key derivation function are both collision-resistant PRFs.

Theorem 3. *Let $\text{EtM} = \text{EtM}[\text{KDF}, F, \text{SE}]$ be the EtM construction using functions F and KDF as well as encryption scheme SE. Let \mathcal{A} be any $\text{r-BIND}_{\text{EtM}}$ adversary. Then there exist adversaries \mathcal{B} and \mathcal{C} , each that run in time that of \mathcal{A} , such that $\text{Adv}_{\text{EtM}}^{\text{r-bind}}(\mathcal{A}) < \text{Adv}_F^{\text{cr}}(\mathcal{B}) + \text{Adv}_{\text{KDF}}^{\text{cr}}(\mathcal{C})$.*

The proof of this theorem will appear in the full version of this work. The s-BIND security of $\text{EtM}[\text{KDF}, F, \text{SE}]$ is perfect because verification re-encrypts the plaintext to check the tag.

Two-key EtM is not binding. The use of a KDF to derive the encryption and MAC keys above is requisite to achieve receiver binding security. Consider omitting the KDF steps, and instead letting keys be a pair (K^e, K^m) where each component is chosen randomly. The opening output by encryption and used by verification is instead $(R, (K^e, K^m))$. The rest of the scheme remains the same as that in Fig. 6. But it is easy to break the receiver binding for this two-key variant: simply have an adversary \mathcal{A} that chooses an arbitrary header H , message M , keys (K^e, K^m) , and randomness R , and computes $R \parallel C \leftarrow \text{enc}_{K^e}(M; R)$ and then

$\text{Enc}(K, H, M):$ $K^e, K^m \leftarrow K$ $IV \leftarrow \{0, 1\}^n$ $T \leftarrow \text{RO}_{K^m}(H \parallel M)$ $C \leftarrow \text{CBC}_{K^e}(\text{Pad}_n(M \parallel T); IV)$ $\ell \leftarrow \text{Pad}_n(M \parallel T)/n$ $C' \parallel C_{\ell-2} \parallel C_{\ell-1} \parallel C_\ell \leftarrow C$ $\text{Return}(C', IV \parallel C_{\ell-2} \parallel C_{\ell-1} \parallel C_\ell)$	$\text{Dec}(K, H, C_1, C_2):$ $K^e, K^m \leftarrow K$ $IV \parallel C_{\ell-2} \parallel C_{\ell-1} \parallel C_\ell \leftarrow C_2$ $C_f \leftarrow C_{\ell-2} \parallel C_{\ell-1} \parallel C_\ell$ $M \parallel T \leftarrow \text{CBC}_{K^e}^{-1}(C_1 \parallel C_f; IV)$ $T' \leftarrow \text{RO}_{K^m}(H \parallel M)$ $\text{If } T \neq T' \text{ then Return } \perp$ $\text{Return}(M, (K^e, K^m))$	$\text{Ver}(H, M, K_f, C_2):$ $K^e, K^m \leftarrow K_f$ $IV \parallel C'_{\ell-2} \parallel C'_{\ell-1} \parallel C'_\ell \leftarrow C_2$ $T \leftarrow \text{RO}_{K^m}(H \parallel M)$ $P \leftarrow \text{Pad}_n(M \parallel T)$ $C \leftarrow \text{CBC}_{K^e}(P; IV)$ $C' \parallel C'_{\ell-2} \parallel C'_{\ell-1} \parallel C'_\ell \leftarrow C$ $\text{Return} \bigwedge_{i=\ell-2}^{\ell} (C'_i = C'_i)$
--	--	---

Fig. 7. Committing authenticated encryption based on MtE composition of CBC mode and a MAC modeled as a random oracle. The length ℓ is defined to be $\text{Pad}_n(M \parallel T)/n$. The function Pad is the standard PKCS#7 padding used in TLS. The notation $\text{CBC}_K(\cdot; IV)$ and $\text{CBC}_K^{-1}(\cdot; IV)$ means CBC mode encryption and decryption with key K and initialization vector IV .

$T \leftarrow F_{K^m}(H \parallel R \parallel C)$. It then chooses another key $\tilde{K}^e \neq K^e$, and computes $\tilde{M} \leftarrow \text{dec}_{\tilde{K}^e}(R \parallel C)$. Finally, it outputs $(H, (R, (K^e, K^m))), (H, (R, (\tilde{K}^e, K^m))), (T)$. It is easy to check that this adversary will win the r-BIND game with probability close to one, assuming SE is such that decrypting the same ciphertext under different keys yields distinct plaintexts with overwhelming probability.

5.3 MAC-then-Encrypt

The MAC-then-encrypt mode generically composes a MAC and an encryption scheme by first computing the MAC of the header and message, then appending the MAC to the message and encrypting them both. The pseudocode in Fig. 7 uses for concreteness CBC mode encryption and we refer to this committing AEAD scheme as MtE. We will also assume the MAC is suitable to be modeled as a keyed random oracle; HMAC-SHA256 is one such [28]. CBC with HMAC in an MtE mode is a common cipher suite for modern TLS connections, which motivated these choices. Prior work has investigated the security of MtE in the sense of CTXT [42, 51] and its ROR security is inherited directly from the encryption mode. Below we will assume that the block size of n bits for the cipher underlying CBC mode, and that our MACs have output length $2n$ bits.

Unlike with Encrypt-then-MAC, we are able to prove the two-key version of MtE secure in the sense of receiver binding. The binding security of MtE in the case where keys are derived via a KDF follows as a corollary, though we believe better bounds can be achieved in this case.

A sketch of an argument that MtE is binding (in the traditional sense where the entire ciphertext is the commitment) appeared in [58]. Their approach, which only relied on modeling the MAC as a RO and made no assumptions about CBC mode, led to a rather loose bound. We instead additionally model the cipher underlying CBC as ideal. This results in a simpler and tighter proof. Our proof, given in the full version of the paper, can also be readily adapted to when CTR mode is used instead of CBC.

Theorem 4. *Let MtE be the scheme defined above using a random oracle and an ideal cipher within CBC mode. For any r -BIND_{MtE} adversary \mathcal{A} making at most q_i queries to its ideal cipher and q_r queries to its random oracle, it holds that $\text{Adv}_{\text{MtE}}^{r\text{-bind}}(\mathcal{A}) < q_i q_r / 2^{2n}$.*

The s-BIND advantage against compactly-committing MtE is zero, since the commitment along with the output of a successful call to Dec uniquely defines the inputs to Ver. Thus, no other ciphertext can be computed in Ver other than the one previously decrypted in Dec, because the inputs to Ver are fixed by Dec.

5.4 Some Non-binding AEAD Schemes

In this section we will briefly detail attacks which break the receiver binding security of some deployed AEAD schemes. In particular, typical schemes that use MACs which are not collision resistant, such as Carter-Wegman MACs, do not suffice. For completeness we spell out an example of breaking the receiver binding of GCM [46], an encrypt-then-MAC style construction that uses a Carter-Wegman MAC.

A slight simplification of the GCM MAC is the function F shown in Fig. 8 applied to a ciphertext. (We ignore associated data for simplicity.) It uses a key K for a block cipher E with block size n , as well as a nonce N . An initial point $P_0 \leftarrow E_K(0^n)$ and a pad $R \leftarrow E_K(N)$ are computed. GCM uses an ϵ -almost XOR universal (ϵ -AXU) [57] hash function computed by considering a ciphertext of m encrypted message blocks an m -degree polynomial defined over a finite field \mathbb{F} . The field is a particular representation of $\text{GF}(2^{128})$. This polynomial is evaluated at the encryption point P_0 and the result is XOR'd with the pad R . The GCM AEAD scheme encrypts the message using CTR mode encryption using E_K and a random 96-bit IV concatenated with a 32-bit counter initially set at one, and then MACs the resulting ciphertext $C = C_1, \dots, C_m$ to generate a tag $T = F(K, IV \parallel 0^{32}, C_1, \dots, C_m)$.

A straightforward way to consider GCM as a compactly committing AEAD is to have encryption output as the commitment portion C_2 the tag T , and the rest of the ciphertext as the first portion C_1 . Decryption works as usual for GCM, but additionally outputs (IV, K) as the opening. Verification works by recomputing encryption and checking that the resulting tag matches the commitment value C_2 . We denote this scheme simply by $\text{GCM} = (\text{Kg}, \text{Enc}, \text{Dec}, \text{Ver})$ below.

We now give an r -BIND_{GCM} adversary \mathcal{A} . We ignore associated data for simplicity. To win, \mathcal{A} must output $((M, (IV, K), (M', (IV', K'), T))$ so that $\text{Ver}(M, (IV, K), T) = \text{Ver}(M', (IV', K'), T) = 1$. We will build an \mathcal{A} that chooses messages such that $|M| = |M'|$. The adversary \mathcal{A} will start by choosing a ciphertext C_1, \dots, C_m such that

$$F(K, IV, C_1, \dots, C_m) = F(K', IV', C_1, \dots, C_m) \quad (1)$$

$F(K, N, (C_1, \dots, C_m)):$ $P_0 \leftarrow E_K(0^n)$ $R \leftarrow E_K(N)$ $S \leftarrow \sum_{i=1}^m C_i P_0^{m-i}$ $T \leftarrow R \oplus S$ $\text{Return } T$
--

Fig. 8. A simplified description of the CW MAC used in GCM.

$\text{CtE1-Enc}(K, H, M)$ $(K_f, C_2) \leftarrow \text{Com}(H \parallel M)$ $C_1 \leftarrow \text{enc}_K(C_2, M \parallel K_f)$ Return (C_1, C_2)	$\text{CtE1-Dec}(K, H, C_1, C_2)$ $(M \parallel K_f) \leftarrow \text{dec}_K(C_2, C_1)$ If $M = \perp$ then Return \perp $b \leftarrow \text{VerC}(K_f, C_2, H \parallel M)$ If $b = 0$ then Return \perp Return (M, K_f)	$\text{CtE2-Dec}(K, H, C_1, C_2)$ $(M \parallel K_f) \leftarrow \text{dec}_K(H, C_1)$ If $M = \perp$ then Return \perp $b \leftarrow \text{VerC}(K_f, C_2, H \parallel M)$ If $b = 0$ then Return \perp Return (M, K_f)
$\text{CtE2-Enc}(K, H, M)$ $(K_f, C_2) \leftarrow \text{Com}(H \parallel M)$ $C_1 \leftarrow \text{enc}_K(H, M \parallel K_f)$ Return (C_1, C_2)		

Fig. 9. Algorithms for two Commit-then-Encrypt variants. Facebook’s scheme uses CtE2 with an HMAC-based commitment. CtE1-Ver and CtE2-Ver both just output $\text{VerC}(H, M, K_f, C_2)$.

and letting M (resp. M') be the CTR-mode decryption of C_1, \dots, C_m under IV, K (resp. IV', K'). Choosing the ciphertext such that condition 1 holds is straightforward, as plugging in for the definition of F and rearranging, the adversary must solve the equation

$$\left[\sum_{i=1}^m C_i (P^{m-i} + (P')^{m-i}) \right] + (E_K(N) + E_{K'}(N')) = 0$$

where $P \leftarrow E_K(0^n)$ and $P' \leftarrow E_{K'}(0^n)$. For example, pick arbitrary C_1, \dots, C_{m-1} and solve for the C_m that satisfies the equation.

This attack works even if associated data is used, or if the whole ciphertext is used as the commitment. A very similar attack works on ChaCha20/Poly1305 [15]; a small tweak is required to handle the fact that not every member of $\mathbb{F}_{2^{130-5}}$ is a valid ciphertext block.

6 Composing Commitment and AEAD

In the last section we saw that existing AEAD schemes already realize (compactly) committing AEAD in some cases. These schemes, however, only realize single-opening security as the opening includes the secret key. We now turn to schemes that achieve multi-opening committing AEAD, and focus specifically on schemes that generically compose AEAD with a commitment scheme.

Commit-then-Encrypt. We start with a simple general construction, what we call the Commit-then-Encrypt scheme.³ It combines a commitment scheme $\text{CS} = (\text{Com}, \text{VerC})$ with an AEAD scheme $\text{SE} = (\text{Kg}, \text{enc}, \text{dec})$. Formally the scheme $\text{CtE1}[\text{CS}, \text{SE}] = (\text{Kg}, \text{CtE1-Enc}, \text{CtE1-Dec}, \text{CtE1-Ver})$ works as shown in Fig. 9.

The CtE1 scheme produces a commitment value to the message and associated data H , and then encrypts the message along with the opening of the commitment. It uses as associated data during encryption the commitment value,

³ This name was also used in [36], but the scheme is distinct. See Sect. 9.

but not H . This nevertheless binds the underlying AEAD ciphertext to H as well as C_2 — as we will show tampering with either will be detected and rejected during decryption. One could additionally include H in the associated data for enc, but this would be less efficient. Should a protocol want H to not be in the commitment scope, one can instead include H only as associated data within enc and omit it from the commitment.

The proofs of the next two theorems will appear in the full version.

Theorem 5 (CtE1 confidentiality). *Let $\text{CtE1} = \text{CtE1}[\text{CS}, \text{SE}]$. Let \mathcal{A} be an $\text{MO-ROR}_{\text{CtE1}}$ adversary making at most q queries to its oracles. Then we give adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{C}$ such that*

$$\mathbf{Adv}_{\text{CtE1}}^{\text{mo-ror}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{SE}}^{\text{ror}}(\mathcal{B}_1) + \mathbf{Adv}_{\text{SE}}^{\text{ror}}(\mathcal{B}_2) + \mathbf{Adv}_{\text{CS}}^{\text{cs-ror}}(\mathcal{C}).$$

The adversaries $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{C} all make the same number of queries as \mathcal{A} and all run in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead.

Theorem 6 (CtE1 ciphertext integrity). *Let $\text{CtE1} = \text{CtE1}[\text{CS}, \text{SE}]$. Let \mathcal{A} be an $\text{MO-CTXT}_{\text{CtE1}}$ adversary making at most q queries to its oracles. Then we give adversaries \mathcal{B}, \mathcal{C} such that*

$$\mathbf{Adv}_{\text{CtE1}}^{\text{mo-ctxt}}(\mathcal{A}) \leq \mathbf{Adv}_{\text{SE}}^{\text{ctxt}}(\mathcal{B}) + \mathbf{Adv}_{\text{CS}}^{\text{v-bind}}(\mathcal{C}).$$

Adversary \mathcal{B} makes the same number of queries as \mathcal{A} and runs in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead. Adversary \mathcal{C} runs in time that of \mathcal{A} .

The receiver binding security of CtE1 is trivially implied by the security of the underlying commitment scheme, as captured by the next theorem.

Theorem 7 (CtE1 receiver binding). *Let $\text{CtE1} = \text{CtE1}[\text{CS}, \text{SE}]$. Let \mathcal{A} be an $\text{r-BIND}_{\text{CtE1}}$ adversary. Then $\mathbf{Adv}_{\text{CtE1}}^{\text{r-bind}}(\mathcal{A}) = \mathbf{Adv}_{\text{CS}}^{\text{v-bind}}(\mathcal{A})$.*

We conclude the section by noting CtE1 meets s-BIND security, since it runs Ver during decryption.

Facebook’s scheme. The Facebook franking scheme (Sect. 3) is almost, but not quite, an instantiation of CtE1 using HMAC as the commitment scheme CS. One difference is that their franking scheme does not bind C_2 to C_1 by including C_2 in the associated data during encryption. The other difference is that the Facebook scheme builds a commitment from HMAC by first generating a random secret key, then using it to evaluate HMAC on the concatenation of the message and the key itself (see Fig. 2 for a diagram). Assuming HMAC remains a collision-resistant PRF when evaluated on its own key, we can prove Facebook’s non-standard construction is a secure commitment (see Theorem 1).

To analyze Facebook’s scheme, then, we introduce the scheme $\text{CtE2}[\text{SE}, \text{CS}] = (\text{Kg}, \text{CtE2-Enc}, \text{CtE2-Dec}, \text{CtE2-Ver})$ that works as shown in Fig. 9. Note that Facebook does not discuss how to handle associated data, and so their scheme is CtE2 using CS instantiated with HMAC and requiring $H = \varepsilon$.

There are two benefits to the approach of CtE1: (1) proving ciphertext integrity does not require any special properties of the commitment scheme, and (2) it is more efficient because associated data is cryptographically processed once, rather than twice. We therefore advocate CtE1, but analyze CtE2 here since it is already in use.

CtE2 is *not* secure assuming just that CS is hiding and binding. The reason is that such commitments can be malleable and this allows easy violation of ciphertext integrity. Specifically, consider a commitment scheme CSBad = (ComBad, VerBad) built using a standard commitment scheme CS = (Com, VerC). Algorithm ComBad(M) runs $(K_c, C) \leftarrow_s \text{Com}(M)$ and then outputs $(K_c, C \parallel 1)$. Algorithm VerBad($M, K_f, C \parallel b$) runs $\text{VerC}(M, K_f, C)$ and outputs the result. An easy reduction shows that CSBad is both hiding and binding, assuming CS is too. But it's clear that CtE2[SE, CSBad] does *not* enjoy ciphertext integrity. The adversary simply obtains one ciphertext, flips the last bit, and submits to the challenge decryption oracle to win.

This shows that standard commitments with hiding and binding properties are insufficient to instantiate CtE2. But if a scheme CS has unique commitments, then we can in fact show security of CtE2. A scheme has *unique commitments* if for any pair $(K_c, M) \in \mathcal{K}_c \times \mathcal{M}$ it holds that there is a single commitment value $C \in \mathcal{C}$ for which $\text{Ver}(K_c, C, M) = 1$. All hash-based CS schemes, including the HMAC one used by Facebook's franking scheme, have unique commitments. If one wanted to use a scheme that does not have unique commitments, then one would need the commitment to satisfy a form of non-malleability [29].

The following sequence of theorems captures the security of CtE2 assuming a unique commitment scheme. Proofs appear in the full version.

Theorem 8 (CtE2 confidentiality). *Let $\text{CtE2} = \text{CtE2}[\text{CS}, \text{SE}]$. Let \mathcal{A} be an MO-ROR_{CtE2} adversary making at most q oracle queries. Then we give adversaries $\mathcal{B}_1, \mathcal{B}_2, \mathcal{C}$ such that*

$$\text{Adv}_{\text{CtE2}[\text{CS}, \text{SE}]}^{\text{mo-ror}}(\mathcal{A}) \leq \text{Adv}_{\text{SE}}^{\text{ror}}(\mathcal{B}_1) + \text{Adv}_{\text{SE}}^{\text{ror}}(\mathcal{B}_2) + \text{Adv}_{\text{CS}}^{\text{cs-ror}}(\mathcal{C})$$

Adversaries $\mathcal{B}_1, \mathcal{B}_2$, and \mathcal{C} all run in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead and make at most q queries.

Theorem 9 (CtE2 ciphertext integrity). *Let $\text{CtE2} = \text{CtE2}[\text{CS}, \text{SE}]$ and assume CS has unique commitments. Let \mathcal{A} be an MO-CTXT_{CtE2} adversary making at most q queries. Then we give adversaries \mathcal{B}, \mathcal{C} such that*

$$\text{Adv}_{\text{CtE2}[\text{CS}, \text{SE}]}^{\text{mo-ctxt}}(\mathcal{A}) \leq \text{Adv}_{\text{SE}}^{\text{ctxt}}(\mathcal{B}) + \text{Adv}_{\text{CS}}^{\text{v-bind}}(\mathcal{C}).$$

Adversaries \mathcal{B} and \mathcal{C} both run in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead. Adversary \mathcal{B} makes at most q queries to its oracles.

Theorem 10 (CtE2 receiver binding). *Let $\text{CtE2} = \text{CtE2}[\text{CS}, \text{SE}]$. Let \mathcal{A} be an r-BIND_{CtE2} adversary. Then we give an adversary \mathcal{B} such that*

$$\text{Adv}_{\text{CtE1}[\text{CS}, \text{SE}]}^{\text{r-bind}}(\mathcal{A}) \leq \text{Adv}_{\text{CS}}^{\text{v-bind}}(\mathcal{B}).$$

Adversary \mathcal{B} runs time that of \mathcal{A} .

Finally, note that CtE2 achieves s-BIND security because it verifies the commitment during decryption.

7 Nonce-Based Committing AEAD and the CEP Construction

The committing AEAD schemes thus far have all been randomized. Cryptographers have advocated that modern AEAD schemes, however, be designed to be nonce-based. Here one replaces internal randomness during encryption with an input, called the nonce. Security should hold as long as the nonce never repeats throughout the course of encrypting messages with a particular key.

We formalize nonce-based committing AEAD and provide a construction of it that additionally achieves a number of valuable properties. It will achieve a multiple-opening security notion suitably modified to the nonce-based setting. It is faster than the other multiple-opening schemes, requiring only two cryptographic passes during encryption and decryption, and a single one during verification. It also reduces ciphertext stretch compared to the schemes of Sect. 6, since the opening will be recomputed in the course of decryption and so does not need to be sent in the encryption.

Nonce-based committing AEAD. A nonce-based CE scheme is a tuple of algorithms $\text{nCE} = (\text{Kg}, \text{Enc}, \text{Dec}, \text{Ver})$. We define it exactly like CE schemes (Sect. 4) except for the following differences. In addition to the other sets, we associate to any nCE scheme a nonce space $\mathcal{N} \subseteq \Sigma^*$. Encryption and decryption are now defined as follows:

- *Encryption:* Encryption Enc is deterministic and takes as input a tuple $(K, N, H, M) \in (\Sigma^*)^4$ and outputs a pair $(C_1, C_2) \in \mathcal{C} \times \mathcal{T}$ or a distinguished error symbol \perp . We require that for any $(K, N, H, M) \in \mathcal{K} \times \mathcal{N} \times \mathcal{H} \times \mathcal{M}$ it is the case that $\text{Enc}(K, N, H, M) \neq \perp$.
- *Decryption:* Decryption Dec is deterministic. It takes as input a quintuple $(K, N, H, C_1, C_2) \in (\Sigma^*)^5$ and outputs a message, opening value pair $(M, K_f) \in \mathcal{M} \times \mathcal{K}_f$ or \perp .

Key generation and verification are unchanged relative to randomized CE schemes. As for randomized schemes, we assume that the length of ciphertexts are dictated only by the lengths of the header and message. We will often write $\text{Enc}_K^N(H, M)$ for $\text{Enc}(K, N, H, M)$ and $\text{Dec}_K^N(H, C_1, C_2)$ for $\text{Dec}(K, N, H, C_1, C_2)$.

Nonce-based security. We adapt the confidentiality and integrity security notions from Sect. 4 to the nonce-based setting. Let game $\text{MO-nREAL}_{\text{nCE}}^A$ be the same as the game $\text{MO-RAND}_{\text{nCE}}^A$ (Fig. 3), except that all oracles take an additional input N , Enc and Dec executions use that value N as the nonce, the sets $\mathcal{Y}_1, \mathcal{Y}_2$ are instead updated with (N, H, C_1, C_2) , and the decryption oracle checks if $(N, H, C_1, C_2) \in \mathcal{Y}_1$. Similarly let game $\text{MO-nRAND}_{\text{nCE}}^A$ be the same as $\text{MO-RAND}_{\text{nCE}}^A$ (Fig. 3), except that all oracles take an additional

input N , and Enc and Dec use that value N as the nonce, and \mathcal{Y}_2 is updated with (N, H, C_1, C_2) . For a scheme nCE, we measure the nonce-based multiple-openings real-or-random $\text{MO-nROR}_{\text{nCE}}$ advantage of an adversary \mathcal{A} by

$$\text{Adv}_{\text{nCE}}^{\text{mo-nror}}(\mathcal{A}) = \left| \Pr \left[\text{MO-nREAL}_{\text{nCE}}^{\mathcal{A}} \Rightarrow 1 \right] - \Pr \left[\text{MO-nRAND}_{\text{nCE}}^{\mathcal{A}} \Rightarrow 1 \right] \right|.$$

An adversary is *nonce-respecting* if its queries never repeat the same N across a pair of encryption queries (two queries to **Enc**, two to **ChalEnc**, or one to each). We will assume nonce-respecting $\text{MO-nRAND}_{\text{nCE}}$ adversaries.

Let $\text{MO-nCTXT}_{\text{nCE}}^{\mathcal{A}}$ be the same as the game $\text{MO-CTXT}_{\text{nCE}}^{\mathcal{A}}$ (Fig. 3), except that all oracles take an additional input N , **Enc** and **Dec** executions use that value N as the nonce, and the set \mathcal{Y} is instead updated with (N, H, C_1, C_2) . For a scheme nCE, we measure the nonce-based multiple-openings real-or-random $\text{MO-nCTXT}_{\text{nCE}}$ advantage of an adversary \mathcal{A} by

$$\text{Adv}_{\text{nCE}}^{\text{mo-nctxt}}(\mathcal{A}) = \Pr \left[\text{MO-nCTXT}_{\text{nCE}}^{\mathcal{A}} \Rightarrow 1 \right].$$

As with randomized committing AEAD, we can provide single-opening versions of the above definitions, and can give an all-in-one version of nonce-based MO and SO security. We omit the details for the sake of brevity.

The sender binding notion s-BIND for nonce-based schemes is the same as for randomized schemes except that the adversary also outputs a nonce N , which is used with Dec. Because verification is unchanged, receiver binding security is formalized exactly the same for randomized and nonce-based committing AEAD.

The Committing Encrypt-and-PRF scheme. One can analyze some traditional nonce-based AEAD schemes to show they are compactly committing. As one example, it is easy to see that the EtE construction (Sect. 5.1) works just as well with non-repeating nonces, but with only single-opening security. The other schemes in Sect. 5 do not, but can be easily modified to by replacing IV with $E_K(N)$. Here we focus on a new scheme that will have better overall performance and security than previous ones. Unlike the legacy schemes studied in Sect. 5 it will be provably secure for multiple openings. At the same time, it will be more efficient than the schemes in Sect. 6.

The scheme $\text{CEP}[\text{G}[\text{K}], \text{F}, \text{F}^{\text{cr}}] = (\text{Kg}, \text{CEP-Enc}, \text{CEP-Dec}, \text{CEP-Ver})$ is in the style of an Encrypt-and-PRF construction. It uses an underlying stream cipher $G[E]$ built from a block cipher $E: \{0, 1\}^k \times \{0, 1\}^n \times \{0, 1\}^n$ and functions $F, F^{\text{cr}}: \{0, 1\}^n \times \{0, 1\}^t \rightarrow \{0, 1\}^t$. The key space is $\mathcal{K} = \{0, 1\}^k$ and key generation simply outputs a random draw from it. Encryption starts by using the nonce with the key K to derive one-time keys for the keyed cryptographic hash F^{cr} and a PRF F , as well as to generate an encryption pad to XOR with the message. We use a block cipher E in CTR mode to generate these values. Finally it computes a binding value for H, M and applies F to that commitment value to generate a tag. Detailed pseudocode is given in Fig. 10.

We will need F^{cr} to both be CR (for binding) as well as secure as a one-time PRF (for confidentiality). This rules out some otherwise desirable choices such as CMAC [41], PMAC [53] and Carter-Wegman-style PRFs such as Poly1305 [16]

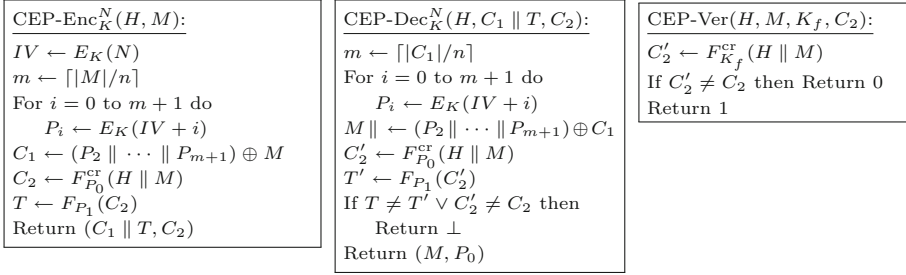


Fig. 10. A nonce-based committing AEAD.

and UMAC [17]. These PRFs are some of the fastest available, but would make CEP vulnerable to binding attacks. (See also the discussion in Sect. 5.4.)

The most obvious choice is HMAC, for which formal analyses support it being a secure PRF for a key secret [2, 3] and CR for adversarially chosen keys of the same length (assuming the underlying hash function is CR). Other multi-property hash functions [10] would also suffice.

The reason we use E_K both for CTR mode and for key derivation is speed. This ensures that we need ever only use a single key with E ; in some environments rekeying can be almost as expensive as another invocation of E . In fact we are simply using E_K to build a stream cipher, and any nonce-based secure stream cipher would do, e.g., ChaCha-20 [15].

One might wonder why have a tag T as well as the commitment value C_2 . The reason is that to achieve multi-opening security, we must disclose the key used with F^{cr} , rendering the unforgeability of C_2 values moot. If one instead omitted T and only checked $C_2' = C_2$ to attempt to achieve unforgeability, then there exists a straightforward MO-nCTXT attack that obtains a ciphertext for a nonce N , queries it to **Dec** to get the key for F^{cr} , and then uses that to forge a new ciphertext to be submitted to **ChalDec**. The application of F under a distinct key provides ciphertext integrity even after an adversary obtains openings (keys for F^{cr}). Similarly, dropping the check during decryption that $C_2' = C_2$ also leads to an attack, but this time on sender binding.

Comparisons. Before getting into the formal security analysis in the next section, we first compare CEP to the generic composition constructions that also achieve multiple-opening security. The first benefit over other schemes is that it is nonce-based, making it suitable for stateful as well as randomized settings (see also Rogaway’s discussion of the benefits of nonce-based encryption [54]).

The second is that ciphertext expansion is reduced by a security parameter number of bits compared to the generic composition constructions, because in CEP we do not need to transport an explicit opening — the recipient recomputes it pseudorandomly from the secret key. Consequently, CEP ciphertexts are shorter than Facebook’s by 256 bits.

The third is that encryption and decryption both save an entire cryptographic pass over the associated data and message. For Facebook’s chosen

algorithms (HMAC for the commitment, plus Encrypt-then-MAC using AES-CBC and HMAC), this means that CEP offers more than a 50% speed-up for both algorithms.⁴ While in some messaging settings encryption and decryption may not be particularly performance-sensitive operations, any cost savings is desirable. In other contexts, such as if one starts using committing encryption on larger files (images, videos) sent over messaging applications or if one wants abuse reporting for streaming communications, performance will be very important.

CEP achieves the stronger multiple-opening security goal, setting it apart from the legacy committing AEAD schemes from Sect. 5. At the same time, CEP has equivalent or better performance than those schemes. With respect to EtM and MtE, verification is reduced from two cryptographic passes to one.

8 Analysis of CEP

Useful abstractions. We will introduce some intermediate abstractions of the underlying primitives. First, a nonce-based stream cipher G takes as input a key K , a nonce N , and an output length ℓ . It outputs a string of length ℓ bits. The second abstraction is of the implicit MAC used within CEP. It is the composition $F \circ F^{\text{cr}}(P_1, P_0, H \| M) = F_{P_1}(F_{P_0}^{\text{cr}}(H \| M))$ for random keys P_0, P_1 and any strings H, M . The output is a t -bit string. We defer a discussion of the security properties required from these abstractions to the full version of this work. There, we define a nonce-based pseudorandom generator (PRG) security notion that mandates attackers cannot distinguish between G 's output and random bit strings, as well as a multi-user unforgeability notion $\text{MU-UF-CMA}_{F \circ F^{\text{cr}}}$ that captures the unforgeability of $F \circ F^{\text{cr}}$ when adversaries can attack it under multiple keys.

Security of CEP. We are now in position to formally analyze the confidentiality, ciphertext integrity, and binding of CEP. We give theorems for each in turn, with proofs deferred to the full version of the paper.

Theorem 11 (CEP confidentiality). *Let $\text{CEP} = \text{CEP}[G, F]$. Let \mathcal{A} be an $\text{MO-nROR}_{\text{CEP}}$ adversary making at most q queries and whose queried messages total at most σ bits. Then we give adversaries $\mathcal{B}, \mathcal{C}, \mathcal{D}$ such that*

$$\text{Adv}_{\text{CEP}}^{\text{mo-nror}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_G^{\text{prg}}(\mathcal{B}) + 2 \cdot \text{Adv}_F^{\text{prf}}(\mathcal{C}) + \cdot \text{Adv}_{F^{\text{cr}}}^{\text{prf}}(\mathcal{D})$$

Adversary \mathcal{B} makes at most q queries to its oracle, the sum of its total outputs requested is σ bits. Adversary \mathcal{C} makes at most q queries to its oracle, and never repeats a key identifier. Adversary \mathcal{D} make at most q queries to its oracle and never repeats a key identifier. All adversaries run in time at most that of \mathcal{A} plus an overhead of at most $\mathcal{O}(q)$.

⁴ HMAC is slower than AES. If AES-NI is available, then the speed-up will be even larger, since the HMAC passes will be the performance bottleneck.

Theorem 12 (CEP ciphertext integrity). *Let $\text{CEP} = \text{CEP}[G, F]$. Let \mathcal{A} be an $\text{MO-nCTXT}_{\text{CEP}}$ adversary making at most q queries with query inputs totalling at most σn bits. Let F^2 be the tagging scheme described earlier. Then we give adversaries \mathcal{B}, \mathcal{C} such that*

$$\text{Adv}_{\text{CEP}}^{\text{mo-nctxt}}(\mathcal{A}) \leq \text{Adv}_G^{\text{prg}}(\mathcal{B}) + \text{Adv}_{F^2}^{\text{mu-uf-cma}}(\mathcal{C}) .$$

Adversary \mathcal{B} runs in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead and makes q queries totaling at most σn bits. Adversary \mathcal{C} makes at most q queries and runs in time that of \mathcal{A} plus at most $\mathcal{O}(q)$ overhead.

Finally we turn to binding. Recall that any scheme that effectively runs commitment verification during decryption achieves sender binding. The check that $C'_2 = C_2$ during decryption accomplishes this, and so the scheme is perfectly sender binding. For receiver binding, a simple reduction gives the following theorem showing that the CR of F^{cr} implies binding of CEP.

Theorem 13 (CEP receiver binding). *Let $\text{CEP} = \text{CEP}[G, F]$. Let \mathcal{A} be any $\text{r-BIND}_{\text{CEP}}$ adversary. Then we give an adversary \mathcal{B} such that $\text{Adv}_{\text{CEP}}^{\text{r-bind}}(\mathcal{A}) \leq \text{Adv}_{F^{\text{cr}}}^{\text{cr}}(\mathcal{B})$ and \mathcal{B} runs in time that of \mathcal{A} .*

9 Related Work

The primary viewpoint in the literature has been that committing encryption is undesirable either because one wants deniability [20, 22, 50] or due to the theoretical challenges associated with proving encryption confidentiality in the face of adaptive compromises [23]. Thus while *non*-committing encryption has received significant attention (q.v., [22–25, 27, 34, 35, 40, 43, 49, 50, 62–65]), there is a dearth of literature on building purposefully committing encryption.

We are aware of only one previous work on building committing encryption schemes, due to Gertner and Herzberg [36]. They give definitions that are insufficient for the message franking setting (in particular they do not capture server binding or multiple opening security). They do not analyze AE schemes, and focus only on building asymmetric primitives.

Our receiver binding security property is related to the concept of robust encryption, introduced by Abdalla et al. [1]. They give two security notions for public-key encryption (PKE). The stronger, called strong robustness, asks that an adversarially-chosen ciphertext should only correctly decrypt under at most one legitimate secret key. Mohassel [48] showed efficient ways of adapting existing PKE schemes to be robust. Farshim et al. [32] subsequently pointed out that some applications require robustness to adversarially generated secret keys, and introduced a notion called complete robustness. In a later work, Farshim, Orlandi, and Rosie [33] adapt these robustness definitions to the setting of authenticated encryption, message authentication codes (MACs), and pseudorandom functions (PRFs). They show that in this context, the simpler full robustness notion of [32] is the strongest of those considered.

These prior notions, in particular the full robustness for AE notion from [33], do not suffice for formalizing binding for AEAD. First, it does not capture sender binding. Second, for receiver binding, it turns out that the most straightforward adaptation of full robustness to handle associated data fails to imply receiver binding. We defer a more detailed explanation to the full version of this work.

Abdalla et al. [1] propose a generic composition of a commitment scheme and PKE scheme to achieve robustness and Farshim et al. [33] show a variant of this for the symmetric encryption setting. The latter construction commits to the key, not the message, and could not be used to achieve the multiple opening security targeted by our generic composition constructions.

Selective-opening security allows an adversary to adaptively choose to corrupt some senders that sent (correlated) encrypted messages [8] or to compromise the keys of a subset of receivers [38]. Bellare et al. [8] gave the first constructions of schemes secure against selective-opening attacks for sender corruptions. Non-committing encryption can be used to realize security for receiver corruptions. Our definitions do not model selective-opening attacks, and as mentioned in the introduction, assessing the viability of committing AEAD in selective-opening settings is an interesting open problem.

Acknowledgments. The authors would like to thank the anonymous reviewers of Crypto 2017 for their thoughtful comments, as well as Mihir Bellare for discussions about robust encryption and its relation to binding. This work was funded in part by NSF grant CNS-1330308.

References

1. Abdalla, M., Bellare, M., Neven, G.: Robust encryption. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 480–497. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-11799-2_28](https://doi.org/10.1007/978-3-642-11799-2_28)
2. Bellare, M.: New proofs for NMAC and HMAC: security without collision-resistance. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 602–619. Springer, Heidelberg (2006). doi:[10.1007/11818175_36](https://doi.org/10.1007/11818175_36)
3. Bellare, M., Canetti, R., Krawczyk, H.: Keying hash functions for message authentication. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 1–15. Springer, Heidelberg (1996). doi:[10.1007/3-540-68697-5_1](https://doi.org/10.1007/3-540-68697-5_1)
4. Bellare, M., Canetti, R., Krawczyk, H.: Pseudorandom functions revisited: The cascade construction and its concrete security. In: Proceedings of the 37th Annual Symposium on Foundations of Computer Science, 1996, pp. 514–523. IEEE (1996)
5. Bellare, M., Canetti, R., Krawczyk, H.: HMAC: Keyed-hashing for message authentication. Internet Request for Comment RFC, 2104 (1997)
6. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: Foundations of Computer Science (FOCS) (1997)
7. Bellare, M., Dowsley, R., Waters, B., Yilek, S.: Standard security does not imply security against selective-opening. In: Pointcheval, D., Johansson, T. (eds.) EURO-CRYPT 2012. LNCS, vol. 7237, pp. 645–662. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-29011-4_38](https://doi.org/10.1007/978-3-642-29011-4_38)

8. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-01001-9_1](https://doi.org/10.1007/978-3-642-01001-9_1)
9. Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000). doi:[10.1007/3-540-44448-3_41](https://doi.org/10.1007/3-540-44448-3_41)
10. Bellare, M., Ristenpart, T.: Multi-property-preserving hash domain extension and the EMD transform. In: Lai, X., Chen, K. (eds.) ASIACRYPT 2006. LNCS, vol. 4284, pp. 299–314. Springer, Heidelberg (2006). doi:[10.1007/11935230_20](https://doi.org/10.1007/11935230_20)
11. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: ACM CCS (1993)
12. Bellare, M., Rogaway, P.: Encode-then-encipher encryption: how to exploit nonces or redundancy in plaintexts for efficient cryptography. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 317–330. Springer, Heidelberg (2000). doi:[10.1007/3-540-44448-3_24](https://doi.org/10.1007/3-540-44448-3_24)
13. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: EUROCRYPT (2006)
14. Bellare, M., Singh, A.C., Jaeger, J., Nyayapati, M., Stepanovs, I.: Ratcheted encryption and key exchange: The security of messaging. Cryptology ePrint Archive, Report 2016/1028 (2016). <http://eprint.iacr.org/2016/1028>
15. Bernstein, D.J.: ChaCha, a variant of Salsa20. <https://cr.yp.to/chacha/chacha-20080128.pdf>
16. Bernstein, D.J.: The Poly1305-AES message-authentication code. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 32–49. Springer, Heidelberg (2005). doi:[10.1007/11502760_3](https://doi.org/10.1007/11502760_3)
17. Black, J., Halevi, S., Krawczyk, H., Krovetz, T., Rogaway, P.: UMAC: fast and secure message authentication. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 216–233. Springer, Heidelberg (1999). doi:[10.1007/3-540-48405-1_14](https://doi.org/10.1007/3-540-48405-1_14)
18. Black, J., Rogaway, P.: A block-cipher mode of operation for parallelizable message authentication. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 384–397. Springer, Heidelberg (2002). doi:[10.1007/3-540-46035-7_25](https://doi.org/10.1007/3-540-46035-7_25)
19. Black, J., Rogaway, P., Shrimpton, T.: Encryption-scheme security in the presence of key-dependent messages. In: Nyberg, K., Heys, H. (eds.) SAC 2002. LNCS, vol. 2595, pp. 62–75. Springer, Heidelberg (2003). doi:[10.1007/3-540-36492-7_6](https://doi.org/10.1007/3-540-36492-7_6)
20. Borisov, N., Goldberg, I., Brewer, E.: Off-the-record communication, or, why not to use PGP. In: ACM Workshop on Privacy in the Electronic Society (2004)
21. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. *J. Comput. Syst. Sci.* **37**, 156–189 (1988)
22. Canetti, R., Dwork, C., Naor, M., Ostrovsky, R.: Deniable encryption. In: Kaliski, B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 90–104. Springer, Heidelberg (1997). doi:[10.1007/BFb0052229](https://doi.org/10.1007/BFb0052229)
23. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: STOC (1996)
24. Canetti, R., Poburinnaya, O., Raykova, M.: Optimal-rate non-committing encryption in a CRS model. IACR Cryptology ePrint Archive (2016)
25. Choi, S.G., Dachman-Soled, D., Malkin, T., Wee, H.: Improved non-committing encryption with applications to adaptively secure protocols. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 287–302. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10366-7_17](https://doi.org/10.1007/978-3-642-10366-7_17)

26. Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. IACR ePrint Archive (2016)
27. Damgård, I., Nielsen, J.B.: Improved non-committing encryption schemes based on a general complexity assumption. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 432–450. Springer, Heidelberg (2000). doi:[10.1007/3-540-44598-6_27](https://doi.org/10.1007/3-540-44598-6_27)
28. Dodis, Y., Ristenpart, T., Steinberger, J., Tessaro, S.: To hash or not to hash again? (in)differentiability results for H^2 and HMAC. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 348–366. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-32009-5_21](https://doi.org/10.1007/978-3-642-32009-5_21)
29. Dolev, D., Dwork, C., Naor, M.: Nonmalleable cryptography. SIAM Review (2003)
30. Facebook. Facebook Messenger app (2016). <https://www.messenger.com/>
31. Facebook. Messenger Secret Conversations technical whitepaper (2016). <https://fbnewsroomus.files.wordpress.com/2016/07/secret-conversations-whitepaper-1.pdf>
32. Farshim, P., Libert, B., Paterson, K.G., Quaglia, E.A.: Robust encryption, revisited. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 352–368. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-36362-7_22](https://doi.org/10.1007/978-3-642-36362-7_22)
33. Farshim, P., Orlandi, C., Rosie, R.: Security of symmetric primitives under incorrect usage of keys. IACR Trans. Symmetric Cryptology **2017**(1), 449–473 (2017)
34. Garay, J.A., Wichs, D., Zhou, H.-S.: Somewhat non-committing encryption and efficient adaptively secure oblivious transfer. IACR Cryptology ePrint Archive (2008)
35. Garay, J.A., Wichs, D., Zhou, H.-S.: Somewhat non-committing encryption and efficient adaptively secure oblivious transfer. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 505–523. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-03356-8_30](https://doi.org/10.1007/978-3-642-03356-8_30)
36. Gertner, Y., Herzberg, A.: Committing encryption and publicly-verifiable signcryption. IACR Cryptology ePrint Archive (2003)
37. Halevi, S., Krawczyk, H.: Security under key-dependent inputs. In: ACM CCS (2007)
38. Hazay, C., Patra, A., Warinschi, B.: Selective opening security for receivers. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015. LNCS, vol. 9452, pp. 443–469. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-48797-6_19](https://doi.org/10.1007/978-3-662-48797-6_19)
39. Hearn, M.: Modern anti-spam and E2E crypto. <https://moderncrypto.org/mail-archive/messaging/2014/000780.html>
40. Hemenway, B., Ostrovsky, R., Rosen, A.: Non-committing encryption from Φ -hiding. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015. LNCS, vol. 9014, pp. 591–608. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-46494-6_24](https://doi.org/10.1007/978-3-662-46494-6_24)
41. Iwata, T., Kurosawa, K.: OMAC: one-key CBC MAC. In: Johansson, T. (ed.) FSE 2003. LNCS, vol. 2887, pp. 129–153. Springer, Heidelberg (2003). doi:[10.1007/978-3-540-39887-5_11](https://doi.org/10.1007/978-3-540-39887-5_11)
42. Krawczyk, H.: The order of encryption and authentication for protecting communications (or: How Secure Is SSL?). In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 310–331. Springer, Heidelberg (2001). doi:[10.1007/3-540-44647-8_19](https://doi.org/10.1007/3-540-44647-8_19)
43. Lei, F., Chen, W., Chen, K.: A non-committing encryption scheme based on quadratic residue. In: Levi, A., Savaş, E., Yenigün, H., Balcişoy, S., Saygın, Y. (eds.) ISCIS 2006. LNCS, vol. 4263, pp. 972–980. Springer, Heidelberg (2006). doi:[10.1007/11902140_101](https://doi.org/10.1007/11902140_101)
44. Marlinspike, M.: libsignal protocol (Java) (2016). <https://github.com/WhisperSystems/libsignal-protocol-java>

45. Marlinspike, M., Perrin, T.: The Double Ratchet algorithm. <https://whisper-systems.org/docs/specifications/doubleratchet/doubleratchet.pdf>
46. McGrew, D., Viega, J.: The galois/counter mode of operation (gcm). Submission to NIST Modes of Operation Process 20 (2004)
47. Millican, J.: Challenges of E2E Encryption in Facebook Messenger. Real World Cryptography Conference (2017). <https://www.realworldcrypto.com/rwc2017/program>
48. Mohassel, P.: A closer look at anonymity and robustness in encryption schemes. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 501–518. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-17373-8_29](https://doi.org/10.1007/978-3-642-17373-8_29)
49. Nielsen, J.B.: Separating random oracle proofs from complexity theoretic proofs: the non-committing encryption case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (2002). doi:[10.1007/3-540-45708-9_8](https://doi.org/10.1007/3-540-45708-9_8)
50. O’Neill, A., Peikert, C., Waters, B.: Bi-deniable public-key encryption. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 525–542. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-22792-9_30](https://doi.org/10.1007/978-3-642-22792-9_30)
51. Paterson, K.G., Ristenpart, T., Shrimpton, T.: Tag Size *Does* matter: attacks and proofs for the TLS record protocol. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 372–389. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-25385-0_20](https://doi.org/10.1007/978-3-642-25385-0_20)
52. Rogaway, P.: Authenticated-encryption with associated-data. In: ACM CCS (2002)
53. Rogaway, P.: Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In: Lee, P.J. (ed.) ASIACRYPT 2004. LNCS, vol. 3329, pp. 16–31. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-30539-2_2](https://doi.org/10.1007/978-3-540-30539-2_2)
54. Rogaway, P.: Nonce-based symmetric encryption. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 348–358. Springer, Heidelberg (2004). doi:[10.1007/978-3-540-25937-4_22](https://doi.org/10.1007/978-3-540-25937-4_22)
55. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006). doi:[10.1007/11761679_23](https://doi.org/10.1007/11761679_23)
56. Ryan, M.D.: Enhanced certificate transparency and end-to-end encrypted mail. In: NDSS. The Internet Society (2014)
57. Stinson, D.R.: Universal hashing and authentication codes. Designs, Codes and Cryptography (1994)
58. Wang, L., Pass, R., Shelat, A., Ristenpart, T.: Secure channel injection and anonymous proofs of account ownership. Cryptology ePrint Archive, Report 2016/925 (2016). <http://eprint.iacr.org/2016/925>
59. Wegman, M.N., Lawrence Carter, J.: New hash functions and their use in authentication and set equality. Journal of computer and system sciences (1981)
60. Whatsapp. Whatsapp (2016). <https://www.whatsapp.com/>
61. Wikipedia. Signal (software) (2016). [https://en.wikipedia.org/wiki/Signal_\(software\)](https://en.wikipedia.org/wiki/Signal_(software))
62. Zhu, H., Araragi, T., Nishide, T., Sakurai, K.: Adaptive and composable non-committing encryptions. In: Steinfeld, R., Hawkes, P. (eds.) ACISP 2010. LNCS, vol. 6168, pp. 135–144. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-14081-5_9](https://doi.org/10.1007/978-3-642-14081-5_9)
63. Zhu, H., Araragi, T., Nishide, T., Sakurai, K.: Universally composable non-committing encryptions in the presence of adaptive adversaries. In: SECRYPT (2010)

64. Zhu, H., Bao, F.: Non-committing encryptions based on oblivious naor-pinkas cryptosystems. In: Roy, B., Sendrier, N. (eds.) INDOCRYPT 2009. LNCS, vol. 5922, pp. 418–429. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-10628-6_27](https://doi.org/10.1007/978-3-642-10628-6_27)
65. Zhu, H., Bao, F.: Error-free, multi-bit non-committing encryption with constant round complexity. In: Lai, X., Yung, M., Lin, D. (eds.) Inscrypt 2010. LNCS, vol. 6584, pp. 52–61. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-21518-6_4](https://doi.org/10.1007/978-3-642-21518-6_4)