# Authenticated Encryption in the Face of Protocol and Side Channel Leakage*

Guy Barwell[1], Daniel P. Martin[2], Elisabeth Oswald[1], and Martijn Stam[1]

[1] Department of Computer Science, University of Bristol,
Merchant Venturers Building, Woodland Road,
Bristol, BS8 1UB, United Kingdom.
rgb.crypto@gmail.com, {elisabeth.oswald, martijn.stam}@bris.ac.uk
[2] School of Mathematics, University of Bristol, Bristol, BS8 1TW, UK,
and the Heilbronn Institute for Mathematical Research, Bristol, UK.
dan.martin@bris.ac.uk

**Abstract.** Authenticated encryption schemes in practice have to be robust against adversaries that have access to various types of leakage, for instance decryption leakage on invalid ciphertexts (protocol leakage), or leakage on the underlying primitives (side channel leakage). This work includes several novel contributions: we augment the notion of nonce-base authenticated encryption with the notion of continuous leakage and we prove composition results in the face of protocol and side channel leakage. Moreover, we show how to achieve authenticated encryption that is simultaneously both misuse resistant and leakage resilient, based on a sufficiently leakage resilient PRF, and finally we propose a concrete, pairing-based instantiation of the latter.

**Keywords:** provable security, authenticated encryption, generic composition, leakage resilience, robustness

## 1 Introduction

Authenticated Encryption (AE) has arisen out of (practical) necessity: historic modes-of-operation for symmetric encryption [33] implicitly target confidentiality against passive adversaries, but most realistic threat models also demand security against active adversaries. Thwarting adversaries trying to modify ciphertexts is best captured by requiring ciphertext integrity; encryption schemes that offer both this and a suitable passive indistinguishability notion are said to provide authenticated encryption. Today, authenticated encryption has become the primitive of choice to enable secure communication. AE schemes can be constructed from components that individually provide either confidentiality or authenticity, both in a traditional probabilistic setting [6] and a more modern nonce-based one [32]. As a result, there exist several black-box constructions of

authenticated encryption schemes based on simpler, keyed primitives such as pseudorandom functions or permutations, including MACs and blockciphers.

Unfortunately, in practice neither the composition nor the underlying components behave as black-boxes: side-channel attacks often leak additional information to an adversary, leading to real-life breaks (e.g. [47]). Invariably, these attacks are possible by exploiting a discrepancy between the capabilities of a theoretical adversary and an actual, real-life one. Thus, these attacks neither violate the security assumptions on the primitive nor do they invalidate the security claims: rather, they render these claims insufficient and the existing security models as inadequate.

In response, a number of works have tried to capture more closely how protocols behave when implemented [10,16,19]. We are particularly interested in *subtle* authenticated encryption [4] which augments the authenticated encryption security game with an implementation-dependent leakage oracle that provides an adversary deterministic decryption leakage on *invalid* ciphertexts only. Subtle authenticated encryption encompasses earlier notions such as multiple decryption errors [9] and the release of unverified plaintexts [2]; it can be regarded as *protocol* leakage.

Orthogonally, *primitives* can leak. Kocher (et al.) [24, 25] showed how both timing and power measurements lead to a side-channel, enabling the extraction of secret data out of cryptographic devices. Primitives believed to be secure, such as AES, were broken without actually violating the assumption that AES is a secure pseudorandom permutation. Such attacks are captured in the framework of leakage resilient cryptography. Here an adversary can adaptively choose a leakage function that is restricted in scope as only computation is assumed to leak information [31], and in size. The latter is captured by leaking only a certain number of bits per call. If the overall leakage remains unbounded the model is referred to as continuous leakage. For a variety of schemes and security notions, resilience against certain classes of leakage can be proven [12, 23, 46], but dealing with adaptivity that allows leakage after an adversary has received a challenge is often problematic.

The current theory of authenticated encryption is not suited to take this additional leakage resource into account. In this work we provide a framework for dealing with AE in the presence of leakage, which then allows us to determine the constraints on primitives and constructions alike to yield AE secure against classes of leakage functions. Moreover, we propose a concrete instantiation of a leakage-resilient pseudorandom function suitable to be used to form the first leakage-resilient, nonce-based authenticated encryption scheme.

### 1.1 Our Contributions

**Augmenting nonce-base authenticated encryption with leakage.** We start by augmenting the nonce-based authenticated encryption security notion (Section 2.1) with leakage (Section 3). This new notion, which we will refer to as LAE, can be regarded as a generalization of the SAE framework by Barwell et al. [4], yet it also captures leakage-resilience as introduced by Dziembowski

and Pietrzak [14]. We provide corresponding leakage notions for the primitives used by the composition results by Namprempre, Rogaway and Shrimpton [32] (henceforth NRS), namely nonce- or iv-based encryption, pseudorandom functions, and message authentication codes.

For the traditional AE notion by Rogaway and Shrimpton [42], an adversary has to distinguish between a world with a real encryption and decryption oracle on the one hand, and a world with a random ciphertext generator and a rejection oracle on the other. In the LAE game the number of oracles available to the adversary increased from two to four: both worlds are augmented with true encryption and decryption oracles and we will allow (only) these additional oracles to leak.

For the leakage mechanism, we adopt the approach originally suggested by Micali and Reyzin [31] and later adapted for leakage resilience [14] where an adversary can provide a leakage function to be evaluated on the internal variables of the oracle, with the leakage output to be returned to the adversary alongside the normal output. The model is very powerful, allowing the adversary to adaptively choose which leakage function they would like evaluated on a query by query basis.

To avoid trivial wins, the leakage functions that are allowed need to be restricted to prevent, for instance, leaking the entire key in one go. We model this by explicitly defining security relative to a class of leakage functions (as is common for instance in the contexts for related-key or key-dependent message attacks). By appropriately setting the class of leakage functions, we show that our notion generalises previous strengthened AE security notions, including SAE, RUP and distinguishable decryption errors [2,4,9], and previous leakage notions, including the simulatable leakage, auxiliary input and probing models [12,20,46].

**Generic composition with leakage.** Our second contribution (Section 5) is an investigation on how to perform generic composition in the presence of leakage by extending the results of NRS [32]. We establish that schemes susceptible to release of unverified plaintext are unsuitable even for much more modest types of leakage and we confirm modern folklore that this affects all schemes that are roughly of the type Encrypt-and-MAC or MAC-then-Encrypt (cf. [2]). Conversely, we show that Encrypt-then-MAC style schemes *are* secure against a large class of leakage functions, where we express this class in terms of the leakage classes against which the underlying primitives are secure. For this composition of leakage from different primitives, we effectively just concatenate the leakage of the constituent parts, which implicitly assumes that only computation leaks (cf. [31]).

In particular, we show security of the N2 and A5 constructions of NRS against nonce-respecting adversaries (Theorem 1 and Corollary 1), and of A6 against adversaries who never repeat a nonce and associated-data pair (Corollary 2).

The above result imply that *none* of the NRS schemes achieve misuse resistant LAE security (mrLAE), hence we propose a novel generic construction that *does* meet this strongest definition of security, albeit at the cost of further

ciphertext expansion (Theorem 3). Our result gives ciphertexts that are two blocks longer than the messages (rather than the single block expansion of an NRS scheme): we leave open whether mrLAE security can be achieved with less ciphertext expansion.

Moreover, we show that instantiating CFB mode with a pseudorandom function yields a secure iv-based encryption scheme even under leakage (Theorem 4). This allows us to apply our generic composition results to construct the first AE scheme secure against continuous leakage based on a pseudorandom function actively secure against continuous leakage and a MAC scheme secure against continuous leakage of both tagging and verification.

**Instantiation using a new leakage resilient PRF.** Our final contribution (in the full version [3]) is the construction of these latter two primitives. To this end, we extend the MAC of Martin et al. [30] in two directions. First, we show how it can be adapted such that it may leak under verification, answering an open question from their work. Then, we show how to implement the tagging function such that it is a PRF in the face of leakage. While the previous implementation of the MAC is a pseudorandom function when no leakage is present, already small amounts of leakage are disastrous for the pseudorandomness property. It turns out that the underlying key update mechanism due to Kiltz and Pietrzak [23] is intrinsically unsuitable to create an actively secure pseudorandom function: the mechanism shares a key out in two which allows a form of leak-in-the-middle attack. The solution we propose is to use three shares instead and we prove that the resulting construction is indeed a pseudorandom function that is leakage-resilient even against adaptive adversaries.

### 1.2 Related Work

*Authenticated encryption.* One of the earliest symmetric works on concrete security of AE was by Bellare and Namprempre [6]. Working within the probabilistic model, they formalised what it meant to be both confidential and authentic, and investigated how one could achieve this through generic composition, combining two schemes (one with each security property) such that their composition achieved both. Yet, modern authenticated encryption is a stateless and deterministic notion, taking in any randomness or state as an extra parameter termed the nonce. It was formalised across a number of papers, culminating in Rogaway and Shrimpton's 2006 work on DAE [42] and only recently a comprehensive study of all the ways one could combine a PRF with an encryption scheme was completed in the nonce-based setting [32].

The CAESAR competition [7] has driven further research into AE, and particularly into the concept of robustness, namely the idea that a scheme should be more resistant to common problems faced in the real-world. One branch of this research has been into designing schemes that are resistant to certain forms of leakage. Prior to the competition, Boldyreva et al. [9] had investigated how to model a scheme from which decryption failures are not identical, such as under

a timing attack. Andreeva et al. [2] (RUP) considered the release of unverified plaintexts, where the decryption oracle releases candidate plaintexts even if they fail verification. The robust authenticated encryption notion of Hoang et al. [19] also implies security against the leakage of these candidate plaintexts, among other goals. Barwell et al. [4] defined the SAE framework as a generalisation of these notions, and used it to compare the three previous works. However, in each of these cases the adversary only receives leakage from decryption, and this leakage is modelled as a fixed, deterministic function, rather than a more general set of functions available to an adaptive side-channel attacker.

*Leakage resilient constructions.* Within the leakage resilient literature, there are several works towards providing leakage resilient encryption, but most of them have been in the bounded leakage model [18,37]. In the bounded retrieval model, Bellare et al. [5] proved the security of a symmetric encryption scheme that provides authenticated encryption in the leak free case, and indistinguishability when leakage is involved. Pereira et al. [34] proposed what is, to our knowledge, the first and only leakage resilient encryption scheme in the simulatable leakage model. However, the construction requires a leak free component and in practice relies on the existence of efficient simulators of the leakage from (e.g.) AES, simulators that Longo et al. [27] demonstrate are unlikely to exist.

Following on from Pereira et al. [34], the recent work by Berti et al. [8] also attempts to construct leakage resilient misuse-resistent authenticated encryption, albeit from a very different direction. In some respects, our work is "top-down", setting a clear objective and evaluating what this demands of the underlying primitives, while theirs is "bottom-up", beginning with well understood primitives and asking what can be constructed. Motivated by this, the two papers adopt very different leakage models: we work in full generality, whereas different sections of Berti follow different leakage models. More generally, their work assumes a single (completely) leak free component, whereas ours allows any of the components to leak as long as the overall leakage is not too great. They hypothesis that (without many leak-free components) leakage resilient misuse resistant authenticated encryption is impossible, while we show that this can be achieved. Furthermore, their work does not consider associated data.

Another manner to ensure that the adversary cannot progressively leak the key material is to update the keys themselves (instead of their representation). Previous leakage resilient works in this direction include the MAC of Schipper [44], or the DH-ratcheting concept [11, 35]. However, these tend to require that all parties to the communication hold modifiable state and remain perfectly in sync, a demand we are able to avoid.

Each of the models above severely restricts the information or computations that an adversary may be able to perform, thereby limiting their utility for modelling active side-channel attacks. The continuous leakage model mitigates these problems, which is why we focus on that when instantiating our AE scheme. To the best of our knowledge, ours is the first leakage resilient encryption scheme in the continuous leakage model.

Our generic composition results allow us to combine leakage resilient components, for which we provide candidates built around a PRF secure against leakage. Currently there are two leakage resilient PRGs, due to Pietrzak (and Dziembowski) [14,36], from which it may be possible to build a leakage resilient stream cipher, although issues arise with restarting using the same key. Works of Dodis and Pietrzak [13], and Faust et al. [15] describe two PRFs secure under non-adaptive leakage: each requires that the leakage (functions) are fixed at the start of the game, while the latter also requires the inputs to be fixed. For a PRF to be used within a composition theorem, adaptive security is required. Finally, Martin et al. [30] provide a MAC which is secure against leakage on the tagging function only. We will use this as the basis of our instantiations, and extend it to achieve security against leakage on verification queries, resolving an open question from their work.

## 2 Preliminaries

**General notation.** For assignment of a value $U$ to the variable $T$ we will write $T \leftarrow U$, where $U$ may also be the outcome of some computation. If the variable is a set, we use the shorthand $S \xleftarrow{\cup} U$ for $S \leftarrow S \cup \{U\}$. To assign a value drawn uniformly at random from some finite set $B$ to variable $A$, we write $A \xleftarrow{\$} B$. By convention, arrays and lists are initialised empty. We use $=$ for equality testing. We write $\mathbb{A} \rightarrow b$, to denote that adversary $\mathbb{A}$ outputs some value $b$. To define notions etc. we will write $X := Y$ to say that $X$ is defined as some expression $Y$. The distinguished symbol $\lightning$ denotes an invalid query. The symbol $||$ denotes an *unambiguous* encoding, meaning if $Z \leftarrow X||Y$ it must be possible given $Z$ to uniquely recover $X$ and $Y$, notated $X||Y \leftarrow Z$, no matter what types $X, Y$ may take. The length $|A|$ is the length of $A$ when expressed as a string of elements of some underlying alphabet $\Sigma$ (usually $\Sigma = \{0,1\}$).

Whenever a function is described with a subscript, this will define the first parameter, meaning $f_k(\cdot,\cdot) = f(k,\cdot,\cdot)$. For consistency and clarity of notation, we refer to security definitions in capitals (e.g. IND–CPA) and typeset functions in calligraphic ($\mathcal{E}$), spaces in sans serif ($\mathsf{K}$), "secret" elements in lower case ($k$), known elements in upper case ($M$), and adversaries in blackboard bold ($\mathbb{A}$). When we introduce implementations, these will be denoted in bold ($\boldsymbol{\mathcal{E}}$).

**Adversarial advantages.** We will define our security notions through indistinguishability games where an adversary is given access to one of two collections of oracles. The adversary $\mathbb{A}$ may make queries to these oracles, and eventually outputs a bit. Instead of writing the games in code, we adopt shorthand notation [2] so that the *distinguishing advantage* of $\mathbb{A}$ between two collections of $n$ oracles $(\mathcal{O}_1, \ldots, \mathcal{O}_n)$ and $(\mathcal{P}_1, \ldots, \mathcal{P}_n)$ is defined as

$$\underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathcal{O}_1, \ldots, \mathcal{O}_n \\ \mathcal{P}_1, \ldots, \mathcal{P}_n\end{matrix}\right) := \left|\Pr\left[\mathbb{A}^{\mathcal{O}_1,\ldots,\mathcal{O}_n} \rightarrow 1\right] - \Pr\left[\mathbb{A}^{\mathcal{P}_1,\ldots,\mathcal{P}_n} \rightarrow 1\right]\right|,$$

where the probabilities are taken over the randomness of the oracles, and key $k \leftarrow\!\!\$\ \mathsf{K}$ (note that multiple oracles will often use the same key). We may refer to the oracles by their numerical position: the $i^{\text{th}}$ oracle implements either $\mathcal{O}_i$ or $\mathcal{P}_i$ depending which collection the adversary is interacting with.

A scheme is considered secure with respect to a particular security goal if the relevant adversarial advantage is small for all adversaries running within reasonable resources. We do not draw judgement as to what "small" may mean, nor what constitutes "reasonable resources", since these depend heavily on context.

### 2.1 Authenticated Encryption

**Core definitions.** Early works to formalize symmetric encryption (cf. [21]) closely followed the precedent for public key encryption. Over the years understanding of what should be expected of symmetric encryption evolved considerably, both in terms of syntax and security. The basis for our work will be the widely accepted nonce-based model using indistinguishability from random bits for confidentiality [39–41]. After introducing this model, we will briefly refer back to an older, non-authenticated version of encryption as it is one of the building blocks later on.

*Syntax.* An authenticated encryption scheme consists of a pair of deterministic functions $\mathsf{Enc}$ and $\mathsf{Dec}$, called encryption and decryption, respectively. Encryption $\mathsf{Enc}$ takes four inputs, resulting in a single ciphertext $C \in \mathsf{C}$. Besides the key $k \in \mathsf{K}$ and the message $M \in \mathsf{M}$, the inputs are some associated data $A \in \mathsf{A}$ that will be authenticated but not encrypted, and finally a nonce $N \in \mathsf{N}$ used to ensure that repeat encryptions will not result in repeat ciphertexts. Decryption $\mathsf{Dec}$ takes as input again the key, the nonce, and the associated data, in addition to the ciphertext. It outputs a purported message or an error message $\perp \notin \mathsf{M}$.

This syntax can be summarized as

$$\begin{aligned} \mathsf{Enc} &\colon \mathsf{K} \times \mathsf{N} \times \mathsf{A} \times \mathsf{M} \to \mathsf{C} \\ \mathsf{Dec} &\colon \mathsf{K} \times \mathsf{N} \times \mathsf{A} \times \mathsf{C} \to \mathsf{M} \cup \{\perp\} \ . \end{aligned}$$

In practice, the key space $\mathsf{K}$, nonce space $\mathsf{N}$, associated data $\mathsf{A}$, message space $\mathsf{M}$, and ciphertext space $\mathsf{C}$ are generally bitstrings of various lengths. It is common to have $\mathsf{A} = \mathsf{M} = \mathsf{C} = \{0,1\}^*$, and $\mathsf{K} = \mathsf{N} = \{0,1\}^n$ for some security parameter $n$. That said, our implementation in Appendix **??** instantiates the various spaces with more general groups (linked to pairings).

We require that an authenticated encryption scheme is both *correct* and *tidy*. These two properties are satisfied iff, for all $k, N, A, M, C$ in the appropriate spaces:

*Correctness* : $\mathsf{Dec}_k(N, A, \mathsf{Enc}_k, (N, A, M)) = M$

*Tidiness* : if $\mathsf{Dec}_k(N, A, C) \neq \perp$ then $\mathsf{Enc}_k(N, A, \mathsf{Dec}_k(N, A, C)) = C$

Together, tidiness and correctness imply that decryption is wholly specified by the encryption routine.

**function** $\$^F(X)$
  $C_0 \leftarrow F(X)$
  $C_1 \xleftarrow{\$} \Sigma^{|C_0|}$
  **return** $C_1$

**function** $\perp^G(X)$
  **return** $\perp$

Fig. 1: The generic oracles $\$^F$ and $\perp^G$ idealise the output of $F$ as random elements of $\Sigma$, and of $G$ as always rejecting. They are used to define the reference world in our security definitions, for various choices of $(F, G)$, which will be omitted whenever clear. Usually $\Sigma = \{0,1\}$, with $|C_0|$ the length of $C_0$ as a bitstring.

Additionally, we require encryption to be *length regular*, which is satisfied if there exists some stretch function $\tau\colon \mathbb{N} \to \mathbb{N}$ such that for all inputs the ciphertext length $|\mathsf{Enc}_k(N, A, M)| = |M| + \tau(|M|)$.

*Security notions.* Ever since Rogaway and Shrimpton's treatment of deterministic authenticated encryption, it is customary to capture both confidentiality and integrity requirements in a single game. Here the adversary gets oracle access either to the "real" world or to the "ideal" world and needs to distinguish between these two worlds. In the real world, oracle access consists of the encryption and decryption functionalities $\mathsf{Enc}_k$ and $\mathsf{Dec}_k$, using a randomly drawn and secret key $k$. In the ideal world, the encryption oracle is replaced with an oracle \$ that generates randomly drawn ciphertexts and the decryption oracle with an oracle $\perp$ that rejects all ciphertexts. Irrespective of which world the adversary is in, we will refer to the $\mathsf{Enc}_k$ vs. \$ oracle as the challenge encryption oracle or as the first oracle (based on the oracle ordering) and to the $\mathsf{Dec}_k$ vs. $\perp$ oracle as the challenge decryption (or second) oracle.

We will use a slightly different, but equivalent, formulation where an adversary additionally has access to the true encryption and decryption oracles in both worlds. Thus the adversary will have access to *four* oracles in each world: the challenge encryption oracle, the challenge decryption oracle, the true encryption oracle, and finally the true decryption oracle. Having these extra oracles will help us later on to add leakage, which will only ever be on the true oracles and never on one of the challenge oracles. One could even argue that the additional oracles provide a more representative and expressive framework: the honest oracles describe how an adversary may "learn" about a system, while the challenge ones allow them to "prove" they have done so (cf. a similar, more detailed argument for subtle authenticated encryption [4]).

As our reference point we will use the oracles defined in Figure 1, with all probabilities taken over randomness of the key and sampling within the oracle.

*Queries.* Already in the leak-free setting, certain combinations of queries will easily distinguish the two worlds. To avoid these trivial wins, we will therefore prohibit certain queries—or in some cases simply assume adversaries refrain from making prohibited queries. For example, if an adversary can send a challenge encryption to decryption they can trivially win. As a general rule, we prohibit the same query being made to oracles which take the same inputs (such as the honest and challenge encryption oracles), and also prohibit performing the

inverse of previous queries. For example, the ciphertext output from the challenge encryption oracle cannot be passed into the decryption oracle.

If an adversary has made a query $(N, A, M)$ to an encryption oracles (either challenge or true) receiving output $C$, then making the same query again to one of the encryption oracles or making the query $(N, A, C)$ to one of the decryption oracles (either challenge or true) the original and the new queries are deemed *equivalent*. For any query, we refer to the process of later making an equivalent query as *forwarding* the query, i.e. to make a second query whose inputs were inputs or outputs from the first query. A special case of forwarding a query is *repeating* the query, namely making the same query again to the same oracle. Forwarding queries from challenge to true oracles (or vice versa) or from challenge encryption to challenge decryption oracles (or vice versa) will lead to trivial wins unless oracle behaviour is adapted. Without loss of generality, we will restrict the adversary from making problematic queries instead.

*Nonce selection requirements.* Our security games will be agnostic over how the nonce is selected, with this property enforced by restricting the adversary. An adversary against an (authenticated) encryption scheme is called *nonce respecting* if whenever making a new query they do not use a nonce more than once to any oracle matching the syntax of $\mathsf{Enc}_k$ or $\mathcal{E}_k$. They are *random-iv respecting*, or simply *iv respecting*, if for any new query with these oracles their nonce $N$ (which we term an IV and will generally write as $I$ instead) is sampled uniformly from $\mathsf{N}$ immediately prior to querying the oracle (and thus not involved in the logic used to select other elements of the query). These requirements do not apply when interacting with oracles matching the syntax of $\mathsf{Dec}_k$ or $\mathcal{D}_k$. A scheme is called *(nonce) misuse resistant* if the adversary does not have to be nonce respecting, providing that the adversary does not make multiple queries using the same $(N, A, M)$ triple.

**Definition 1.** *Let* $\mathsf{Enc}$ *be an authenticated encryption scheme,* $\mathbb{A}$ *an adversary who does forward queries to or from his first or second oracle (and thus does not repeat first oracle queries). Then, the nAE advantage of an adversary* $\mathbb{A}$ *against* $\mathsf{Enc}$ *is*

$$\mathbf{Adv}^{\mathrm{AE}}_{\mathsf{Enc}}(\mathbb{A}) := \Delta_{\mathbb{A}}\begin{pmatrix}\mathsf{Enc}_k, \mathsf{Dec}_k, \mathsf{Enc}_k, \mathsf{Dec}_k \\ \$\ , \ \perp\ , \mathsf{Enc}_k, \mathsf{Dec}_k\end{pmatrix} .$$

Following our earlier convention, we will refer to a secure *nAE* scheme (or simply nAE) if this nAE advantage is small for all nonce-respecting adversaries running within reasonable resources, and mrAE if it is small for all adversaries running within reasonable resources that might repeat nonces.

**Building blocks: Encryption, MACs and PRFs.** An authenticated encryption scheme is often constructed out of simpler components, with authenticated encryption security derived from that of its constituent parts. The most common of these are "simple" symmetric encryption (ivE), MACs and PRFs. Here we omit the relevant syntax and security notions of these notions, though in the full version [3] we provide a treatment analogous to that for authenticated encryption above.

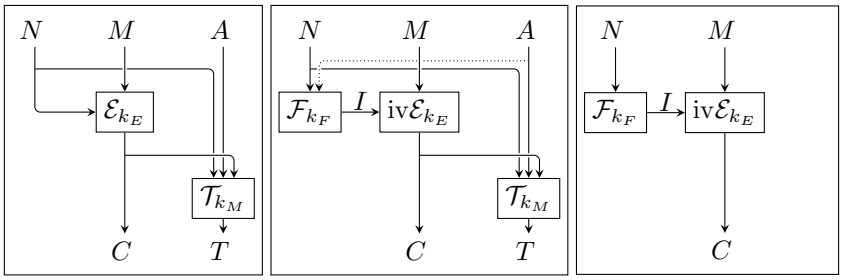


Fig. 2: Graphical representations of the encryption directions of generic composition mechanisms. On the left, N2 converts a nonce-based encryption algorithm $\mathcal{E}$ and MAC scheme $(\mathcal{T}, \mathcal{V})$ into an nAE scheme. On the right, iv2n converts an iv-based encryption scheme iv$\mathcal{E}$ and a PRF into a nonce-based encryption algorithm. Composing these yields A5, shown in the middle ignoring the dotted input, while A6 includes the dotted input. Overall decryption of A5, A6, and N2 will recompute and verify the tag first, only proceeding with further decryption of $C$ if this verification is successful.

**Generic composition for nAE.** NRS [32] investigated how to construct an nAE scheme by composing two PRFs with an ivE scheme. The IV of the ivE scheme is derived from the nAE's inputs using the first PRF call; the optional second PRF call may be used to create an authentication tag. Different schemes emerge by changing which variables are provided to each of the components. NRS identify eight schemes, dubbed A1–A8, with strong security bounds. For a further four schemes (A9–A12) neither strong security bounds nor insecurity was established. Additionally, NRS investigated mechanisms for combining a PRF with an nE scheme. Three schemes (N1–N3) were found secure, with that of a fourth (N4) remaining unresolved.

Figure 2's middle panel shows the schemes A5 and A6. For these two schemes, as well as for N2 (on the left), the ciphertext is input to the second PRF, which means they classify as Encrypt-then-MAC (EtM). The schemes A4, A7–A12, as well as N3 and N4 only use a single PRF and release the IV as tag; for that reason we refer to them as MAC-then-Encrypt (MtE). Finally, the schemes A1–A3 and N1 use two PRFs that can be called in parallel, leading to their classification as Encrypt-and-MAC (E&M). We refer to NRS for full descriptions and graphical illustrations of all schemes mentioned above.

## 3 Security Notions Involving Leakage

Authenticated encryption, as defined above, is deterministic. In a leakage-free setting, this provides a stronger notion than the older probabilistic notion of encryption (as implicitly still used for ivE). When introducing leakage, deterministic schemes are problematic both from a practical and a theoretical perspective.

On the one hand, a practical side-channel attack such as differential power analysis can effectively recover keys from unprotected blockciphers and their AE modes with near certainty. Randomized masking based on secret sharing is one of the main countermeasures against these attacks.

On the other hand, theoretical leakage is often modelled as a function on the inputs of the computation, which will include the key. If with each invocation of

the scheme an adversary can let the scheme leak a different key bit of its choice, the full key is easily recovered. To prevent such devastating yet simple leakage, a typical design strategy is to split the key in two shares and update the shares on-the-fly using fresh randomness, mimicking the practical approach.

### 3.1 Implementations Versus Functions

In both the practical and the theoretical approaches mentioned above, a deterministic scheme is implemented in a randomized fashion in order to provide resistance against leakage. Therefore, when arguing about leakage, we will need to make a distinction between the *scheme* (a collection of deterministic functions) and its probabilistic *implementation*.

For our definition of the implementations of a function we take our cue from the secret-sharing approach, where a redundant representation of the key is used and this representation is rerandomized as part of the implementation. To enable this rerandomization, we provide the implementation of a function with explicit randomness in Definition 2 below, where we use a bold font to denote either the implementation of a function or the representation of a key used by the implementation.

**Definition 2.** *An implementation of a function* $f : \mathsf{K}\times\mathsf{X} \to \mathsf{Y}$ *is a deterministic function* $\boldsymbol{f}\colon \boldsymbol{\mathsf{K}} \times \mathsf{X} \times \mathsf{R} \to \boldsymbol{\mathsf{K}} \times \mathsf{Y}$ *along with a probabilistic key initialisation function* $\iota\colon \mathsf{K} \to \boldsymbol{\mathsf{K}}$ *such that* $\iota(k) = \iota(l) \Rightarrow k = l$*. We define the inverse of* $\iota$ *as the function* $\iota^{-1}\colon \boldsymbol{\mathsf{K}} \to \mathsf{K} \cup \{\perp\}$ *such that* $\iota^{-1}(\boldsymbol{k}) = k$ *if* $\iota(k)$ *could have resulted in* $\boldsymbol{k}$*, and* $\perp$ *if no such* $k$ *exists.*

*The implementation is* correct *iff for all* $k \in \mathsf{K}, X \in \mathsf{X}$*, and* $r \in \mathsf{R}$*, setting* $\boldsymbol{k} \leftarrow \iota(k)$ *and* $(\boldsymbol{k}', Y) \leftarrow \boldsymbol{f}(\boldsymbol{k}, X; r)$ *guarantees both* $Y = f(k, X)$ *and* $\iota^{-1}(\boldsymbol{k}') = k$*.*

The initial representation of the key is generated using the function $\iota$, which maps a key $k \in \mathsf{K}$ to a suitable representation $\boldsymbol{k} \in \boldsymbol{\mathsf{K}}$ for the implementation. We assume that $\iota$ is performed only once, and in a leak-free manner, during setup (straight after key generation). Moreover, its inverse $\iota^{-1}$ induces an equivalence relation on the space $\boldsymbol{\mathsf{K}}$; in other words, the implementation keys $\boldsymbol{k}$ can be thought of as alternative representations of the key. During evaluation of $\boldsymbol{f}$ the auxiliary input $r \in \mathsf{R}$ is used to refresh the representation; typically this requires a good randomness source to draw $r$ from.

**Discussion.** Correctness implies that an implementation is identical to the original function when restricted to the second output and that the new key representation $\boldsymbol{k}'$ is equivalent to the initial one $\boldsymbol{k}$. We make no demands of $\boldsymbol{k}$ or $\boldsymbol{k}'$ beyond these, so it is permissible to set $\boldsymbol{k} = \boldsymbol{k}' = k$ and thus recover the traditional syntax. Our security definitions will be such that for correct schemes and assuming "trivial" leakage, the corresponding leak-free security notions from the preceding section will emerge.

Definition 2 can be linked to practice in a straightforward manner. Recall that practical implementations of blockciphers often use masking based on secret

sharing schemes. In this case, the implementation of the blockcipher describes how to evaluate the blockcipher based on the shares of the key as well as how the sharing is refreshed using external randomness $r$ (which need not be leak-free). Furthermore, $\iota$ is exactly the function that creates the initial secret sharing of the key.

Syntactically the implementation $\boldsymbol{f}$ may appear stateful: after all they take in some $\boldsymbol{k}$ and output an updated $\boldsymbol{k}'$ for the next invocation. However, since the implementation is of a stateless function $f$, there is no need to synchronize state between communication parties. Instead, each party can use its own, independent representation of the key.

**Implementation of an nAE Scheme.** For concreteness, we now explicitly define the implementation of an nAE scheme. We assume that **Enc** and **Dec** syntactically use the same representations $\mathbf{K}$ (and key initialisation function $\iota$), which we later use for expressing our security notions.

By correctness of the implementation, one can see that the ciphertext output by **Enc** (resp. message by **Dec**) will always be independent of the randomness $r$, since they are equal to the corresponding output of $\mathsf{Enc}$ (resp. $\mathsf{Dec}$). Definitions for the implementations of other security primitives are written accordingly.

**Definition 3.** *Let* $(\mathsf{Enc}, \mathsf{Dec})$ *be an authenticated encryption scheme. An* AE implementation *is a pair of deterministic functions*

$$\begin{aligned}\mathbf{Enc}&\colon \mathbf{K} \times \mathsf{N} \times \mathsf{A} \times \mathsf{M} \times \mathsf{R} \to \mathbf{K} \times \mathsf{C}\\ \mathbf{Dec}&\colon \mathbf{K} \times \mathsf{N} \times \mathsf{A} \times \mathsf{C} \times \mathsf{R} \to \mathbf{K} \times (\mathsf{M} \cup \{\perp\})\end{aligned}$$

*along with* $\iota\colon \mathsf{K} \to \mathbf{K}$ *satisfying* $\iota(k) = \iota(l) \Rightarrow k = l$ *and* $\iota^{-1}\colon \mathbf{K} \to \mathsf{K} \cup \{\perp\}$ *such that* $\iota^{-1}(\boldsymbol{k}) = k$ *if* $\iota(k)$ *could have resulted in* $\boldsymbol{k}$*, and* $\perp$ *if no such* $k$ *exists. The implementation is* correct *iff for any* $k, N, A, M, C, r$ *from the appropriate spaces and* $\boldsymbol{k} \leftarrow\!\!\$\ \iota(k)$*, setting*

$$(\boldsymbol{k}', C') \leftarrow \mathbf{Enc}(\boldsymbol{k}, N, A, M; r) \text{ and } (\boldsymbol{k}'', M') \leftarrow \mathbf{Dec}(\boldsymbol{k}, N, A, C; r),$$

$(\boldsymbol{k}', C') \leftarrow \mathbf{Enc}(\boldsymbol{k}, N, A, M; r)$ *and* $(\boldsymbol{k}'', M') \leftarrow \mathbf{Dec}(\boldsymbol{k}, N, A, C; r)$*, the following properties hold:*

$$\begin{aligned}&k = \iota^{-1}(\boldsymbol{k}) = \iota^{-1}(\boldsymbol{k}') = \iota^{-1}(\boldsymbol{k}'')\\ &C' = \mathsf{Enc}_k(N, A, M) \text{ and } M' = \mathsf{Dec}_k(N, A, C)\ .\end{aligned}$$

### 3.2 What Constitutes Leakage

Following Micali and Reyzin's approach, we will model leakage by allowing an adversary to specify a leakage function in conjunction with an oracle query. The input signature of the leakage function matches that of the implementation $\boldsymbol{f}$ it relates to, allowing it to wholly simulate the implementation. A leakage set is a collection of leakage functions for an implementation.

**Definition 4.** *A* leakage function *of an implementation* $\boldsymbol{f}\colon \mathsf{K} \times \mathsf{X} \times \mathsf{R} \to \mathsf{K} \times \mathsf{Y}$ *is a function* $L\colon \mathsf{K} \times \mathsf{X} \times \mathsf{R} \to \mathsf{L}$ *for some output leakage space* $\mathsf{L}$. *A* leakage set *of an implementation* $\boldsymbol{f}$ *is a set of leakage functions.*

The choice of leakage set should contain all plausible (functions of) inputs to the implementation that an adversary can compute, and may be probabilistic. This might include functions of any intermediate variables, since these are computable from the inputs simply by simulating the construction. Broadly speaking, the larger the leakage set the more powerful the adversary is likely to be. The leakage set $\emptyset$ allows us to model the leak-free case. Technically we define it to be the set containing just the null function, meaning the adversary can always select a leakage function, thus maintaining the correct syntax for our security games.

### 3.3 Security Notions Incorporating Leakage

We are now in a position to define the security of an implementation in the presence of leakage. We do so by reframing the classical notions given to work on the implementation of a function, and by extending the notions such that the honest oracles are allowed to leak. The adversary wins the game if they can distinguish whether their leak-free challenge oracles implement the scheme honestly or are idealised. We differentiate our notions from the classic variant by prefixing an "L", for *leakage*.

In the classical setting, each oracle simply evaluates the appropriate function with the game's secret key. For an implementation, a similar, but slightly more complicated, approach is required. The oracle must draw randomness, and provide this to the implementation to update the key representation. This same randomness, along with all other inputs, must be provided to the leakage function. The new representation must then be stored, and the two outputs returned to the adversary. For any implementation $\boldsymbol{f}$, the corresponding leakage oracle is denoted $\ell[\boldsymbol{f}]_k$, when initialised with representation $\boldsymbol{k} = \iota(k)$. Code-based descriptions for certain leaky implementations related to authenticated encryption are given in Figure 3. If an adversary has access to multiple oracles based on the same key, say $\mathsf{Enc}_k$ and $\mathsf{Dec}_k$, then we will assume that their respective implementation oracles (so $\ell[\mathbf{Enc}]_k$ and $\ell[\mathbf{Dec}]_k$) will operate on the same representation $\boldsymbol{k}$, which hence will be initialized only once. Such a shared representation corresponds to a setting where both **Enc** and **Dec** are implemented on the same device. Needless to say, our security definitions below can be strengthened by allowing an adversary to interact with multiple implementations each using their own representation of the same key.

As in the leakage free definitions, security is taken over the randomness of the initial keys, and of the oracles. Notice that this choice includes the sampling from $\mathsf{R}$. We assume the adversary only ever makes queries for which his inputs are selected from the appropriate spaces. For leakage, this means some leakage set that will be specified in the security notion.

**function** $\ell[\boldsymbol{\mathcal{E}}]_k(M; L)$
  $r \leftarrow\!\!\$\ \mathsf{R}$
  $\Lambda \leftarrow L(\boldsymbol{k}, M; r)$
  $C, \boldsymbol{k} \leftarrow \boldsymbol{\mathcal{E}}(\boldsymbol{k}, M; r)$
  **return** $(C, \Lambda)$

**function** $\ell[\boldsymbol{\mathcal{D}}]_k(C; L)$
  $r \leftarrow\!\!\$\ \mathsf{R}$
  $\Lambda \leftarrow L(\boldsymbol{k}, C; r)$
  $M, \boldsymbol{k} \leftarrow \boldsymbol{\mathcal{D}}(\boldsymbol{k}, C; r)$
  **return** $(\perp, \Lambda)$

**function** $\ell[\mathbf{Enc}]_k(N, A, M; L)$
  $r \leftarrow\!\!\$\ \mathsf{R}$
  $\Lambda \leftarrow L(\boldsymbol{k}, N, A, M; r)$
  $C, \boldsymbol{k} \leftarrow \mathbf{Enc}(\boldsymbol{k}, N, A, M; r)$
  **return** $(C, \Lambda)$

**function** $\ell[\mathbf{Dec}]_k(N, A, C; L)$
  $r \leftarrow\!\!\$\ \mathsf{R}$
  $\Lambda \leftarrow L(\boldsymbol{k}, N, A, C; r)$
  $M, \boldsymbol{k} \leftarrow \mathbf{Dec}(\boldsymbol{k}, N, A, C; r)$
  **return** $(M, \Lambda)$

Fig. 3: Honest leakage oracles an adversary may use to help them distinguish. All inputs are taken from the appropriate spaces, with leakage functions chosen from the relevant leakage set. For $\mathcal{L}_{\mathsf{E}}$-IND–CPLA, the adversary has access to $\ell[\boldsymbol{\mathcal{E}}]_k$, and for the augmented notion $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-IND-aCPLA they are also given very limited access to $\ell[\boldsymbol{\mathcal{D}}]_k$. LAE security, $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$-LAE provides access to $(\ell[\mathbf{Enc}]_k, \ell[\mathbf{Dec}]_k)$.

For the purposes of defining forwarding of queries, we will ignore the additional input associated to the leakage. For instance, after a query $(N, A, M)$ to the challenge encryption oracle, the query $(N, A, M, L)$ to the true encryption oracle will be considered equivalent—and would constitute forwarding—irrespective of $L$.

**Definition 5.** *Let* $(\mathbf{Enc}, \mathbf{Dec})$ *be an implementation of an authenticated encryption scheme* $\mathsf{Enc}, \mathsf{Dec}$*, and* $\mathbb{A}$ *an adversary who does not forward queries to or from his first or second oracles (and thus does not repeat such queries). Then, the* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$–LAE advantage *of an adversary* $\mathbb{A}$ *against* $(\mathbf{Enc}, \mathbf{Dec})$ *under leakage* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$ *is*

$$\mathbf{Adv}^{\mathrm{LAE}}_{\mathbf{Enc},\mathbf{Dec};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) := \underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathsf{Enc}_k, \mathsf{Dec}_k, \ell[\mathbf{Enc}]_k, \ell[\mathbf{Dec}]_k \\ \$\ ,\ \perp\ , \ell[\mathbf{Enc}]_k, \ell[\mathbf{Dec}]_k\end{matrix}\right).$$

**Definition 6.** *Let* $\boldsymbol{\mathcal{E}}$ *be an implementation of an encryption scheme* $\mathcal{E}$*, and* $\mathbb{A}$ *an adversary who never forwards queries to or from his first oracle (and thus does not repeat first oracle queries). The* $\mathcal{L}_{\mathsf{E}}$*-IND–CPLA advantage (named for chosen-plaintext-with-leakage-attack) of* $\mathbb{A}$ *against* $\boldsymbol{\mathcal{E}}$ *is*

$$\mathbf{Adv}^{\mathrm{IND-CPLA}}_{\boldsymbol{\mathcal{E}};\mathcal{L}_{\mathsf{E}}}(\mathbb{A}) := \underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathcal{E}_k, \ell[\boldsymbol{\mathcal{E}}]_k \\ \$\ , \ell[\boldsymbol{\mathcal{E}}]_k\end{matrix}\right).$$

We next provide an additional encryption notion, IND–aCPLA, that will be required for our composition results later. It describes a modified version of the IND–CPLA game in which the adversary is also allowed leakage from the decryption implementation $\ell[\boldsymbol{\mathcal{D}}]_k$ (see Figure 3), but *only* on ciphertexts they have previously received from $\ell[\boldsymbol{\mathcal{E}}]_k$. At first glance, this appears to be more similar to an IND–CCA style notion, but we emphasise this is not the case since the possible decryption queries are heavily restricted. Thus it should be thought of as IND–CPA under the most general form of leakage. Indeed, when the leakage sets are empty, the resulting security notion is equivalent to IND–CPA.

**Definition 7.** *Let $(\boldsymbol{\mathcal{E}}, \boldsymbol{\mathcal{D}})$ be an implementation of an encryption scheme, $\mathbb{A}$ an adversary who does not forward queries to or from his first oracle, and only makes queries to their third oracle that were forwarded from the second. Then the $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-IND–aCPLA advantage of $\mathbb{A}$ against $\boldsymbol{\mathcal{E}}$ is*

$$\mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\boldsymbol{\mathcal{E}},\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{E}},\mathcal{L}_{\mathsf{D}}}(\mathbb{A}) := \underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathcal{E}_k, \ell[\boldsymbol{\mathcal{E}}]_k, \ell[\boldsymbol{\mathcal{D}}]_k \\ \$\ , \ell[\boldsymbol{\mathcal{E}}]_k, \ell[\boldsymbol{\mathcal{D}}]_k\end{matrix}\right) .$$

The IND–aCPLA notion is required for the general composition, where the goal is to construct an LAE scheme from an ivLE scheme (and other components). However, for decryption of the LAE scheme to leak (as we want the leakage to be as powerful as possible), the decryption of ivLE scheme would have to leak. The IND–CPLA security notion does not capture this. Consider an IND–CPA scheme where encryption does not leak, but the leakage from decrypting the zero string returns the key. Clearly the scheme is also IND–CPLA but will trivially break when the adversary is given decryption leakage. The IND–aCPLA notion is trying to capture that decryption "does not leak too much information", so that limited decryption queries made by the LAE scheme will be able to leak.

Against many natural choices of leakage sets, $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-IND–aCPLA and $\mathcal{L}_{\mathsf{E}}$-IND–CPLA are equivalent, since the encryption oracle often suffices to simulate any leakage from decryption. In the nonce-abusing setting (where the adversary is free to select nonces however they wish) there is an obvious mechanism for proving the equivalence, using repeat encryption queries to simulate leaking decryption queries, but even this requires rather strong assumptions on the leakage sets.

In the nonce respecting or iv respecting scenarios such a general reduction is not possible, because there is no way to allow the adversary to use the same nonce multiple times, something a decryption oracle would allow. If the leakage is independent of the nonce (for example) similar results can be recovered, but these are much more restrictive scenarios. It is an interesting open problem to describe sets $\mathcal{L}_{\mathsf{ED}}$ that are in some sense "minimal" for various pairs of leakage sets $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$ taken from some general function classes.

**LMAC and LPRF.** Here we give the PRF and MAC notions a similar treatment to the encryption definitions by enhancing the standard definitions to incorporate leakage.

The LPRF definition below strengthens earlier definitions by Dodis and Pietrzak [13], and by Faust et al. [15]: in our definition both the leakage functions and the inputs can be chosen adaptively based on outputs already seen by the adversary.

**Definition 8.** *Let $\boldsymbol{\mathcal{F}}$ be an implementation of a function $\mathcal{F}$, and $\mathbb{A}$ an adversary who never forwards or repeats queries. Then the $\mathcal{L}_{\mathsf{F}}$-PRLF advantage of $\mathbb{A}$ against $\boldsymbol{\mathcal{F}}$ under leakage $\mathcal{L}_{\mathsf{F}}$ is*

$$\mathbf{Adv}^{\mathrm{PRLF}}_{\boldsymbol{\mathcal{F}};\mathcal{L}_{\mathsf{F}}}(\mathbb{A}) := \underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathcal{F}_k, \ell[\boldsymbol{\mathcal{F}}]_k \\ \$\ , \ell[\boldsymbol{\mathcal{F}}]_k\end{matrix}\right).$$

**function** $\ell[\boldsymbol{\mathcal{T}}]_k(M; L)$
$r \leftarrow\!\!\$\ \mathsf{R}$
$\Lambda \leftarrow L(\boldsymbol{k}, M; r)$
$T, \boldsymbol{k} \leftarrow \boldsymbol{\mathcal{T}}(\boldsymbol{k}, M; r)$
**return** $(T, \Lambda)$

**function** $\ell[\boldsymbol{\mathcal{V}}]_k(M, T; L)$
$r \leftarrow\!\!\$\ \mathsf{R}$
$\Lambda \leftarrow L(\boldsymbol{k}, M, T; r)$
$V, \boldsymbol{k} \leftarrow \boldsymbol{\mathcal{V}}(\boldsymbol{k}, M, T; r)$
**return** $(V, \Lambda)$

Fig. 4: Honest leakage oracles an adversary may use to help them distinguish. All inputs are taken from the appropriate spaces, with leakage functions chosen from the relevant leakage set. $(\mathcal{L}_\mathsf{T}, \mathcal{L}_\mathsf{V})$-LMAC security gives access to $(\ell[\boldsymbol{\mathcal{T}}]_k, \ell[\boldsymbol{\mathcal{V}}]_k)$. Since PRFs and the tagging function of a MAC have the same syntax, the LPRF game provides access to $\ell[\boldsymbol{\mathcal{F}}]_k$, which is identical to $\ell[\boldsymbol{\mathcal{T}}]_k$.

Our notion of strong existential unforgeability under chosen message with leakage (below) strengthens both the classical definition, and the leakage definition of Martin et al. [30] (they only allow tagging to leak; setting $\mathcal{L}_\mathsf{V} = \emptyset$ recovers their definition). Allowing meaningful leakage on $\boldsymbol{\mathcal{T}}$ hampers direct use of a secure LPRF as a MAC as typically during verification the "correct" tag would be recomputed as output of the PRF and could consequently be leaked upon (effectively yielding a surreptitious tagging algorithm).

**Definition 9.** *Let $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ be an implementation of a MAC $(\mathcal{T}, \mathcal{V})$, and $\mathbb{A}$ an adversary who does not forward queries from his second oracle to the first. Then the $(\mathcal{L}_\mathsf{T}, \mathcal{L}_\mathsf{V})$-sEUF-CMLA advantage of $\mathbb{A}$ against $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ under leakage $(\mathcal{L}_\mathsf{T}, \mathcal{L}_\mathsf{V})$ is*

$$\mathbf{Adv}^{\mathrm{sEUF-CMLA}}_{\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{V}};\mathcal{L}_\mathsf{T},\mathcal{L}_\mathsf{V}}(\mathbb{A}) := \underset{\mathbb{A}}{\Delta}\left(\begin{matrix}\mathcal{V}_k, \ell[\boldsymbol{\mathcal{T}}]_k, \ell[\boldsymbol{\mathcal{V}}]_k \\ \perp, \ell[\boldsymbol{\mathcal{T}}]_k, \ell[\boldsymbol{\mathcal{V}}]_k\end{matrix}\right).$$

Note that we cast unforgeability as a distinguishing game, rather than as a more usual computational game ("adversary must forge a tag"), but it is straightforward to show equivalence (even in the face of leakage).

## 4 Applying LAE to Attacks in Theory and Practice

A security framework is not much use if it does not highlight the difference between schemes for which strong security results are known, and those against which efficient attacks exist. In this section we discuss the types of leakage normally considered within the literature. We show how previous leakage models can be captured by our leakage set style notion. In the literature there is focus on two types of leakage; protocol leakage (by the AE literature) and side channel leakage (by the leakage resilient literature). We believe that these two notions are highly related and thus we discuss how to capture both. For example, termination of an algorithm at different points (distinguishable decryption failures) is normally detected by a side-channel; timing can be used to capture this if the failures terminate the algorithm at different points in time and power can be used to detect if conditional branches were taken.

Below we recast existing leakage resilience work within our general framework. For completeness, in the full version [3] we describe an existing attack (against GCM) within our setting.

### 4.1 Theoretical Leakage Models

We observe that our model is in many ways the most general possible, and that many previous leakage notions can be captured as version of the $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-LAE security game for suitable choice of leakage sets $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$. Reassuringly, by setting $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}}) = (\emptyset, \emptyset)$ we recover the traditional leakage-free security notions, with $(\emptyset, \emptyset)$-nLAE equivalent to nAE, and both $\emptyset$-IND–CPLA and $(\emptyset, \emptyset)$-IND–aCPLA equivalent to IND–CPA, meaning a secure nE scheme is $\emptyset$-nLE secure.

The deterministic decryption leakage notions from the AE literature can be recovered by choosing the appropriate leakage set. The SAE framework generalises both the RUP model, and (nonce-based analogues of) the Distinguishable Decryption Failure notions of Boldyreva et al. [2, 4, 9]. The security notions are parametrised by a deterministic decryption leakage function $\Lambda$, corresponding to security under the leakage sets $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}}) = (\emptyset, \{\Lambda\})$. Thus the strongest notions available in these settings are equivalent to $(\emptyset, \{\Lambda\})$–LAE. Several of their weaker notions translate to the corresponding weakening of this, including authenticity under deterministic leakage, (known variously as CTI–sCPA, INT–RUP or an extended form of INT–CTXT), which translates to a variant of $(\emptyset, \{\Lambda\})$–LAE in which the adversary cannot query the encryption challenge oracle (and thus does not interact with either $\mathcal{E}_k$ or \$).

In the simulatable leakage model (e.g. [46]), the adversary receives leakage in addition to their query, but is restricted to leakage functions that can be simulated without the key. The simulatable model considered by Standaert et al. (for example) can be captured by our model by having set of leakage functions contain the single function which provides the power trace to the adversary. The auxiliary input model [12] gives the adversary the output of a hard to invert function applied to the key, alongside the normal security notion interactions. The only computation leaks model [31] (discussed in more detail in Section 5.1) restricts the adversary to leakage functions that can be locally computed: any step of the algorithm can only leak on variables being used at that point. In the following sections we show how this leakage set can be defined for our given constructions.

In the probing model [20] the adversary can gain access to the values of $t$ of the internal wires from the computation. A scheme is secure if an adversary with the knowledge of $t$ internal wires can do no better than if they had access to the function in a black box manner. If there are $n$ internal wires, this leakage can be captured by our set notation by constructing a set with $n$ choose $t$ leakage functions, each giving the complete value of the relevant wires.

Our leakage sets incorporate the bounded leakage model (e.g. [18, 22, 26]) by restricting the set of allowable adversaries to those who only make sufficiently few queries to the leakage oracles.

One mechanism that need not rely on randomness is to instead use a leak-free component [48]. Although instantiating such components in practice is between hard and impossible [29], our framework nonetheless supports it (by suitable choice of leakage set).

Another idea to provide security is frequent rekeying. However, such a solution relies on synchronized states between encryption and decryption which can be difficult to maintain, thereby restricting applicability of this approach. However, in specific contexts such as secure channels, synchronization might not be too onerous.

## 5 Generic Composition for LAE

### 5.1 Modelling Composed Leakage

Our challenge is to establish to what extent the NRS schemes remain secure when taking leakage into account. Ideally, we would like to claim that if both the ivE and the PRFs are secure in the presence of leakage, then so will the composed nAE be. To make such a statement precise, the leakage classes involved need to be specified. We opt for an approach where the leakage classes for the components are given (and can be arbitrary) and then derive a leakage class for the resulting nAE for which we can prove security.

*Encryption leakage.* In a nutshell, we define the leakage of the composition as the composition of the leakage. As an example, consider an implementation of $\mathsf{A5}$ (Figure 2). When encrypting, the leakage may come from any of the components: the PRF $\boldsymbol{\mathcal{F}}$ may leak some information $L_F(\boldsymbol{k}_F, N; r_F)$; the IV-encryption routine $\mathbf{iv}\boldsymbol{\mathcal{E}}$ might leak some information $L_E(\boldsymbol{k}_E, I, M; r_E)$; the Tag function $\boldsymbol{\mathcal{T}}$ may leak some information $L_T(\boldsymbol{k}_M, N, A, C_e; r_M)$. To ease notation, we will use the shorthand $L_F(\star), L_E(\star)$, and $L_T(\star)$ respectively for these leakages. In that case, we say that the leakage on the authenticated encryption operation as a whole consists of the triple $(L_F(\star), L_E(\star), L_T(\star))$. Under the hood, this implies some parsing and forwarding of the AE's key $(\boldsymbol{k}_F, \boldsymbol{k}_E, \boldsymbol{k}_M)$, randomness $(r_F, r_E, r_M)$ and inputs $N, A, M$, including the calculated values $I$ and $C_e$, to the component leakage functions $L_F, L_E$, and $L_T$.

Expanding the above to classes of functions is as follows. Let $\mathcal{L}_{\mathsf{F}}, \mathcal{L}_{\mathsf{E}}$, and $\mathcal{L}_{\mathsf{T}}$ be the respective leakage classes for $\boldsymbol{\mathcal{F}}, \mathbf{iv}\boldsymbol{\mathcal{E}}$, and $\boldsymbol{\mathcal{T}}$. Then the leakage class $\mathcal{L}_{\mathsf{Enc}}$ for the resulting authenticated encryption scheme is defined as

$$\{(L_F, L_E, L_T) | L_F \in \mathcal{L}_{\mathsf{F}}, L_E \in \mathcal{L}_{\mathsf{E}}, L_T \in \mathcal{L}_{\mathsf{T}}\}.$$

Since an adversary has to select a leakage function in $\mathcal{L}_{\mathsf{Enc}}$ the moment it queries the encryption oracle, it will not be possible to adaptively select for instance the leakage function $L_T$ based on the leakage received from $L_E$ *of that encryption query.*

*Decryption leakage.* In order to describe leakage from decryption, we take a closer look at the role of the two PRFs in the generic constructions. The first one, $\boldsymbol{\mathcal{F}}$, computes the initial vector which is needed both for encryption and decryption. This makes it inevitable that during decryption $\mathcal{F}$ is again computed as a PRF, presumably using the same implementation $\boldsymbol{\mathcal{F}}$. On the other hand, the

| Structure | Leakage | Inverse | Inverse Leakage |
|---|---|---|---|
| MtE | $L_T(\star)$, $L_F(\star)$, $L_E(\star)$ | DtV | $L_F(\star)$, $L_D(\star)$, $L_V(\star)$ |
| M&E | $L_T(\star)$, $L_F(\star)$, $L_E(\star)$ | | |
| EtM | $L_F(\star)$, $L_E(\star)$, $L_T(\star)$ | D&V | $L_F(\star)$, $L_D(\star)$, $L_V(\star)$ |
| | | VtD | $\begin{cases} L_V(\star) & \text{if } \mathcal{V}(\star) = \bot \\ L_V(\star), L_F(\star), L_D(\star) & \text{if } \mathcal{V}(\star) = \top \end{cases}$ |

Fig. 5: The structure of a leakage function from a composition scheme based on the order of its primitives. The exact input parameters to the leakage function vary per scheme, so have been replaced with $\star$: the different $\star$ variables are not the same. On the left are the encryption structures MtE, M&E and EtM, along with the associated leakage function. The right gives the associated inverse: DtV (Decrypt then Verify) is the only way of inverting MtE or M&E schemes. EtM schemes can be inverted in any order, as DtV, D&V (Decrypt and Verify) or VtD (Verify then Decrypt). All constructions have the same encryption leakage, and most have the same decryption leakage. The only one that is different is an EtM–VtD scheme, where the decryption leakage format depends on the validity of the ciphertext.

second PRF, $\boldsymbol{\mathcal{T}}$, is used to create a tag $T$ during encryption. Normally during decryption one would recompute the tag (again using $\boldsymbol{\mathcal{T}}$) and check whether the recomputed tag $T'$ equals the received tag $T$. Yet, in the leakage setting this approach is problematic: $T'$ is the correct tag and its recomputation might well leak it, even when used (repeatedly) to check an incorrect and completely unrelated $T$. Hence, during decryption we will not use a recompute-and-check model, but rather refer directly to a tag-verification implementation $\boldsymbol{\mathcal{V}}$ (that hopefully leaks less).

When considering the decryption leakage of A5, we will assume that, on invalid ciphertexts, the computation terminates as soon as the verification algorithm returns $\bot$. This implies that for invalid ciphertexts only leakage on $\boldsymbol{\mathcal{V}}$ will be available, whereas for valid ciphertexts all three components ($\boldsymbol{\mathcal{V}}, \boldsymbol{\mathcal{F}}$, and $\mathbf{iv}\boldsymbol{\mathcal{E}}$) might leak.

**Overview and interpretation.** Recall that we divided the NRS schemes in three categories: MtE, M&E, and EtM. Figure 5 shows how the composed leakage will leak for each of these schemes. For completeness, we also listed the leakage for the EtM scheme (such as A5) in case full decryption will always take place, even for invalid ciphertexts (where one could have aborted early).

Our choice to model the leakage from the authenticated encryption scheme as completely separate components from the three underlying primitives is rooted in the assumption that only computation leaks. This assumption was first formalized by Micali and Reyzin [31] and, although there are counterexamples to the assumption at for instance the gate level [38], we believe that implementations of the three primitives result in large enough physical components, which can be suitably segregated to avoid cross-leakage.

Leakage on the wire (for instance of the initial vector $I$) can be captured as leakage of the PRF computing the $I$ or alternatively as that of the ivE. In particular, by letting the decryption of the $\mathbf{iv}\boldsymbol{\mathcal{E}}$ component leak its full output (while not allowing any further leakage), we capture the release of unverified plaintext. Furthermore, distinguishable decryption failures on MtE and M&E

schemes invariably arise from verification, which might incorporate a padding check as well. This is modelled by allowing $\mathcal{V}$ to leak, but not any of the other components.

### 5.2 MAC-and/then-Encrypt are Brittle under Leakage

For schemes where the plaintext is input to the MAC (i.e. MtE and M&E schemes), decryption is inevitably of the form DtV. Consequently, during decryption a purported message $M$ is computed before the tag can be verified. Leaking this message $M$ corresponds to the release of unverified plaintext [2], but even more modest leakage, such as the first bit of the candidate message, can be insecure as we show by the following example.

Let us assume for a moment that the encryption routine $\mathbf{iv}\mathcal{E}$ is online, so that reencrypting a slightly modified plaintext using the same $I$ will only affect a change in the ciphertext after the modification in the plaintext. CBC and CFB modes are well-known examples of online $\mathbf{iv}\mathcal{E}$ schemes. Additionally, assume that $\mathbf{iv}\mathcal{E}$'s decryption routine indeed leaks the first bit of the message. Then the authenticated encryption scheme is not secure in the presence of leakage (for the leakage class derived according to the principles outlined previously), which an attack demonstrates.

The adversary first submits a message $M$ to its challenge encryption oracle, receiving a ciphertext $C$ which either is an encryption $\mathcal{E}_k(M||T)$ or, in the ideal world, a uniformly random string. The adversary subsequently queries its decryption-with-leakage oracle on $C$ with its final bit flipped. In the real world, where $C = \mathcal{E}_k(M||T)$, the leakage will then equal the first bit of $M$ with probability 1. Yet in the ideal world, $C$ is independent of $M$, so the leakage will equal the first bit of $M$ with probability half. Thus, testing whether the decryption leakage equals the first bit of $M$ leads to a distinguisher with a significant advantage. However, this does not invalidate IND–aCPLA security of $\mathbf{iv}\mathcal{E}$ as in that game decryption only leaks on valid ciphertexts with known plaintexts.

The above observation implies that for schemes where decryption follows a DtV or D&V structure proving generic composition secure in the presence of leakage is impossible. This affects the NRS compositions A1–A4, A7–A12, N1, N3 and N4; none of which can be regarded as generically secure under leakage and *all* are insecure when using online iv$\mathcal{E}$ and releasing unverified plaintext.

Less general composition results might still be possible, for instance by restricting the leakage classes of the primitives. After all, in the trivial case that the leakage classes are all $\emptyset$, the original NRS results hold directly. We leave open whether significantly larger realistic leakage classes exist leading to secure MtE constructions.

Alternatively, stronger assumptions on $\mathcal{E}$ could help. For instance, if $\mathcal{E}$'s security matches that of a tweakable (variable input length) cipher, the MAC-then-Encrypt constructions become a sort of encode-then-encipher. The latter is secure against release of unverified plaintext [19]. We leave open the identification of sufficient conditions on $\mathcal{E}$ for a generic composition result in the

presence of leakage to pull through for EtM or E&M; relatedly, we leave open the extension of our work to the encode-then-encipher setting.

### 5.3 Encrypt-then-MAC is Secure under Leakage

The iv-based schemes A5 and A6, as well as the nonce-based N2, all fall under the EtM design. The inverse of an EtM scheme can be D&V or VtD, but as just discussed for the D&V variant no meaningful generic security is possible; henceforth we restrict attention to the VtD variant only. These schemes, along with the iv2n mechanism for building a nonce-based encryption scheme out of an iv-based one, are all represented in Figure 2. Before proving their security, we begin with some observations about EtM–VtD designs in the face of leakage.

**Initial observations.** Since the final ciphertext will be formed from an encryption ciphertext and a tag, if the overall output is to be indistinguishable from random bits, then so must the tag. Thus we require both that $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ is a secure $(\mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$-LMAC, and that $\boldsymbol{\mathcal{T}}$ is a secure $\mathcal{L}_{\mathsf{T}}$-LPRF. Shrimpton and Terashima [45] defined a (weaker) authenticated encryption notion where the "recovery information" does not need to be random—only the ciphertext—in which case one may drop the second requirement.

In the traditional case, it is possible to build secure EtM schemes from an encryption scheme that is IND–CPA secure. After all, by assumption on the security of the MAC, the only output the adversary can ever receive from the internal decryption function $\mathcal{D}$ is a plaintext corresponding to a previous $\mathcal{E}$ query. However, when leakage is involved, this previously harmless decryption query suddenly allows the adversary to evaluate a leakage function $L \in \mathcal{L}_{\mathsf{D}}$, albeit on a $(N, C)$ pair for which they already know the corresponding plaintext. If $\mathcal{L}_{\mathsf{D}}$ contained functions revealing sufficient information about the key, this would render the composed scheme completely broken, notwithstanding any IND–CPLA security. Luckily, the augmented IND–aCPLA game in which the adversary is allowed to leak on select decryption queries, is sufficiently nuanced to capture relevant weaknesses in the decryption's implementation.

**Security of EtM composition schemes.** We now describe the security of the composition schemes A5, A6 and N2, and the iv2n construction. Working under the assumption of OCLI-style leakage, as described in Section 5.1, we will reduce the security of the composition to the security of its components. Technically the bound includes a term quantifying any additional weaknesses due to the composition scheme, but in all cases this term is zero. The proofs can be found in the full version [3]. We begin with N2, and show it is essentially as secure as the weakest of its components, by constructing explicit adversaries against each.

**Theorem 1.** *Let* $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}}, \mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$ *be leakage sets for the appropriate primitives, and define* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$ *as in Section 5.1. Let* $\mathbb{A}$ *be an adversary against the* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$*-nLAE security of* $\mathsf{N2}[\boldsymbol{\mathcal{E}}, \boldsymbol{\mathcal{D}}; \boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}}]$*. Then, there exist adversaries*

*$\mathbb{A}_{\mathrm{CPA}}$, $\mathbb{A}_{\mathrm{PRF}}$ and $\mathbb{A}_{\mathrm{MAC}}$ against the $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-nLE security of $(\boldsymbol{\mathcal{E}}, \boldsymbol{\mathcal{D}})$, the $\mathcal{L}_{\mathsf{T}}$-LPRF security of $\boldsymbol{\mathcal{T}}$ and the $(\mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$-LMAC security of $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ such that:*

$$\begin{aligned}&\mathbf{Adv}^{\mathrm{nLAE}}_{\mathsf{N2};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) \leq \\ &\quad \mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\boldsymbol{\mathcal{E}},\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{E}},\mathcal{L}_{\mathsf{D}}}(\mathbb{A}_{\mathrm{CPA}}) + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{T}};\mathcal{L}_{\mathsf{T}}}(\mathbb{A}_{\mathrm{PRF}}) + 2 \cdot \mathbf{Adv}^{\mathrm{sEUF-CMLA}}_{\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{V}};\mathcal{L}_{\mathsf{T}},\mathcal{L}_{\mathsf{V}}}(\mathbb{A}_{\mathrm{MAC}}).\end{aligned}$$

As the following result shows, the intuitive mechanism for building a nLE scheme from a secure ivLE scheme and a secure LPRF is itself secure. While unsurprising, this will allow us to instantiate the N2 construction with the more common object of an ivLE scheme.

**Theorem 2.** *Let $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}}, \mathcal{L}_{\mathsf{F}})$ be leakage sets for the appropriate primitives, and define $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$ as in Section 5.1. Let $\mathbb{A}$ be an adversary against the $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-nLE security of $\mathsf{iv2n}[\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}}; \boldsymbol{\mathcal{F}}]$. Then, there exist $\mathbb{A}_{\mathrm{CPA}}$, $\mathbb{A}_{\mathrm{PRF}}$ against the $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$-ivLE security of $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$, and the $\mathcal{L}_{\mathsf{F}}$-LPRF security of $\boldsymbol{\mathcal{F}}$ respectively, such that:*

$$\mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\mathsf{iv2n};\mathcal{L}_{\mathsf{E}},\mathcal{L}_{\mathsf{D}}}(\mathbb{A}) \leq \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{F}};\mathcal{L}_{\mathsf{F}}}(\mathbb{A}_{\mathrm{PRF}}) + \mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\mathbf{iv}\boldsymbol{\mathcal{E}},\mathbf{iv}\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{ivE}},\mathcal{L}_{\mathsf{ivD}}}(\mathbb{A}_{\mathrm{CPA}}) \ .$$

Pulling these two results together and taking the maximum over the similar adversaries, we are able to prove the security of the A5 construction. The security of A6 against adversaries who never repeat the pair $(N, A)$ can be easily recovered from this by considering it as an equivalent representation of the A5 scheme acting on nonce space $\mathsf{N}' = \mathsf{N} \times \mathsf{A}$ but with no associated data.

**Corollary 1 (nLAE from ivLE and LPRF via A5 composition).** *Let $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}}, \mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}}, \mathcal{L}_{\mathsf{F}})$ be leakage sets for the appropriate primitives, and define $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$ as in Section 5.1. Let $\mathbb{A}$ be an adversary against the $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$-nLAE security of $\mathsf{A5}[\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}}; \boldsymbol{\mathcal{F}}; \boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}}]$. Then, there exist adversaries $\mathbb{A}_{\mathrm{CPA}}$, $\mathbb{A}_{\mathrm{PRF}}$, $\mathbb{A}'_{\mathrm{PRF}}$, and $\mathbb{A}_{\mathrm{MAC}}$ against the $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$-ivLE security of $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$, the $\mathcal{L}_{\mathsf{F}}$-LPRF security of $\boldsymbol{\mathcal{F}}$, the $\mathcal{L}_{\mathsf{T}}$-LPRF security of $\boldsymbol{\mathcal{T}}$ and the $(\mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$-LMAC security of $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ such that*

$$\begin{aligned}\mathbf{Adv}^{\mathrm{nLAE}}_{\mathsf{A5};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) \leq \ & \mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\boldsymbol{\mathcal{E}},\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{ivE}},\mathcal{L}_{\mathsf{ivD}}}(\mathbb{A}_{\mathrm{CPA}}) + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{F}};\mathcal{L}_{\mathsf{F}}}(\mathbb{A}_{\mathrm{PRF}}) \\ & + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{T}};\mathcal{L}_{\mathsf{T}}}(\mathbb{A}'_{\mathrm{PRF}}) + 2 \cdot \mathbf{Adv}^{\mathrm{EUF-CMLA}}_{\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{V}};\mathcal{L}_{\mathsf{T}},\mathcal{L}_{\mathsf{V}}}(\mathbb{A}_{\mathrm{MAC}}) \ .\end{aligned}$$

**Corollary 2 (nLAE from ivLE and LPRF via A6 composition).** *Let $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}}, \mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}}, \mathcal{L}_{\mathsf{F}})$ be leakage sets for the appropriate primitives, and define $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$ as in Section 5.1. Let $\mathbb{A}$ be an adversary against the $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$-LAE security of $\mathsf{A6}[\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}}; \boldsymbol{\mathcal{F}}; \boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}}]$ who does not make two encryption queries with the same $(N, A)$ pair. Then, there exist explicit adversaries $\mathbb{A}_{\mathrm{CPA}}$, $\mathbb{A}_{\mathrm{PRF}}$, $\mathbb{A}'_{\mathrm{PRF}}$, and $\mathbb{A}_{\mathrm{MAC}}$ against the $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$-ivLE security of $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$, the $\mathcal{L}_{\mathsf{F}}$-LPRF security of $\boldsymbol{\mathcal{F}}$, the $\mathcal{L}_{\mathsf{T}}$-LPRF security of $\boldsymbol{\mathcal{T}}$ and the $(\mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$-LMAC security of $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ such that*

$$\begin{aligned}\mathbf{Adv}^{\mathrm{nLAE}}_{\mathsf{A6};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) \leq \ & \mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\boldsymbol{\mathcal{E}},\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{ivE}},\mathcal{L}_{\mathsf{ivD}}}(\mathbb{A}_{\mathrm{CPA}}) + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{F}};\mathcal{L}_{\mathsf{F}}}(\mathbb{A}_{\mathrm{PRF}}) \\ & + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{T}};\mathcal{L}_{\mathsf{T}}}(\mathbb{A}'_{\mathrm{PRF}}) + 2 \cdot \mathbf{Adv}^{\mathrm{sEUF-CMLA}}_{\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{V}};\mathcal{L}_{\mathsf{T}},\mathcal{L}_{\mathsf{V}}}(\mathbb{A}_{\mathrm{MAC}}) \ .\end{aligned}$$

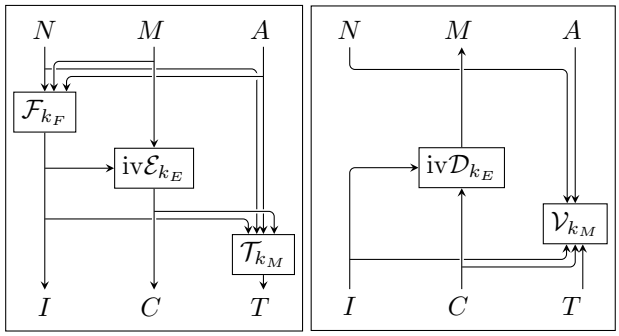


Fig. 6: The Synthetic-IV-and-Tag (SIVAT) scheme. On the left, the encryption routine runs from top to bottom, outputting a ciphertext $I||C||T$. Decryption (on the right) runs from bottom to top. If during decryption verification fails, and $\mathcal{V}_{\boldsymbol{k}_m}$ returns $\perp$, no further computations are performed. In the decryption direction, the PRF $\mathcal{F}$ is not required.

### 5.4 Achieving Misuse Resistant LAE Security

In Section 5.2 we discussed why no composition scheme can be (generically) secure against leakage if its decryption begins by calculating a candidate plaintext. This meant ruling out every NRS construction secure in the nonce misuse model, an important feature for a modern robust AE schemes [7,19,42]. Roughly speaking, for MRAE security a scheme must be MtE (to ensure maximum diffusion) yet for leakage resilience it must be EtM (to ensure minimal leakage).

The Synthetic IV and Tag (SIVAT) scheme, defined in Figure 6, addresses the combined mrLAE goal, by essentially using an MtEtM approach. It can be seen as composing the SIV construction [42] (referred to as A4 in NRS) with a secure MAC, or alternatively as the natural strengthening of A6 towards nonce misuse security, by adding the message to the IV calculation and making the appropriate modifications to enable decryption.

Our additional feature does come at a cost. While schemes in the traditional setting achieve misuse resistance for the same ciphertext expansion as non-resistant schemes, the SIVAT scheme requires essentially twice the expansion. It also has a large number of internal wires, with each function taking in a large number of inputs, although removing any one leads to incorrectness or insecurity. For encryption calls, all inputs must go into the LPRF (for misuse resistance) and for decryption they must go into verification (to prevent RUP attacks).

The proof (in the full version)is very similar to that for A5 or A6 (Corollaries 1 and 2), since the additional element of a SIVAT ciphertext ($I$) is present in those settings, and might already be available to the adversary through leakage.

**Theorem 3.** *Let* $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}}, \mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}}, \mathcal{L}_{\mathsf{F}})$ *be leakage sets for the appropriate primitives, and define* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$ *as in Section 5.1. Let* $\mathbb{A}$ *be an adversary against the* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$*-mrLAE security of* $\mathsf{SIVAT}[\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}}; \boldsymbol{\mathcal{F}}; \boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}}]$*. Then, there exist explicit adversaries* $\mathbb{A}_{\mathrm{CPA}}$*,* $\mathbb{A}_{\mathrm{PRF}}$*,* $\mathbb{A}'_{\mathrm{PRF}}$*, and* $\mathbb{A}_{\mathrm{MAC}}$ *against the* $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$*-ivLE security of* $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$*, the* $\mathcal{L}_{\mathsf{F}}$*-LPRF security of* $\boldsymbol{\mathcal{F}}$*, the* $\mathcal{L}_{\mathsf{T}}$*-LPRF security of* $\boldsymbol{\mathcal{T}}$ *and the* $(\mathcal{L}_{\mathsf{T}}, \mathcal{L}_{\mathsf{V}})$*-LMAC security of* $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$ *such that*

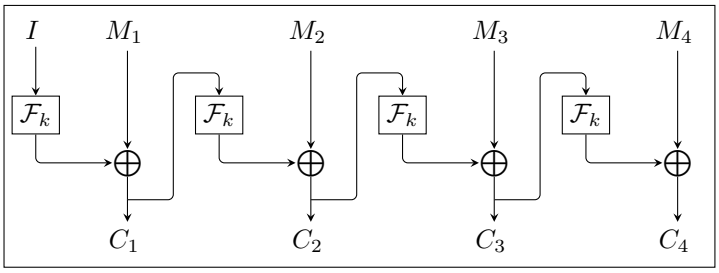


Fig. 7: CFB Mode of Operation based on $\mathcal{F} : \mathsf{K} \times \mathsf{X} \to \mathsf{X}$. The message $M$ is parsed into blocks or elements $M_1 || \ldots || M_m$, and fed through to output ciphertext $C_1 || \ldots || C_m$. The operation $\oplus$ can be any group operation on $\mathsf{X}$.

$$\begin{aligned}\mathbf{Adv}^{\mathrm{nLAE}}_{\mathsf{SIVAT};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) \leq & \mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\boldsymbol{\mathcal{E}},\boldsymbol{\mathcal{D}};\mathcal{L}_{\mathsf{ivE}},\mathcal{L}_{\mathsf{ivD}}}(\mathbb{A}_{\mathrm{CPA}}) + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{F}};\mathcal{L}_{\mathsf{F}}}(\mathbb{A}_{\mathrm{PRF}}) \\ & + \mathbf{Adv}^{\mathrm{LPRF}}_{\boldsymbol{\mathcal{T}};\mathcal{L}_{\mathsf{T}}}(\mathbb{A}'_{\mathrm{PRF}}) + 2 \cdot \mathbf{Adv}^{\mathrm{sEUF-CMLA}}_{\boldsymbol{\mathcal{T}},\boldsymbol{\mathcal{V}};\mathcal{L}_{\mathsf{T}},\mathcal{L}_{\mathsf{V}}}(\mathbb{A}_{\mathrm{MAC}}) \ .\end{aligned}$$

### 5.5 A Leakage Resilient IV-based Encryption Scheme

A crucial component required for our composition is an encryption scheme iv$\mathcal{E}$, whose implementation $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$ is IND-aCPLA secure against a rich class $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$ of leakage functions. As generic composition relies on a secure PRLF implementation $\boldsymbol{\mathcal{F}}$ anyway, we will investigate to what extent this PRLF can be used to bootstrap some iv$\mathcal{E}$ implementation as well. Here we turn to the classical mode of operation CFB (Fig. 7), which has the advantage that only the forward direction of the underlying primitive $\mathcal{F}$ is required, even for decryption (relevant if one would instantiate with a blockcipher). When we move from the standard CFB[$\mathcal{F}$] to its implementation CFB[$\boldsymbol{\mathcal{F}}$] (by replacing $\mathcal{F}$ with its implementation $\boldsymbol{\mathcal{F}}$), processing multi-block plaintexts (or ciphertexts) will result in multiple refreshes of the key's representation $\boldsymbol{k}$ (one for each call to $\boldsymbol{\mathcal{F}}$). We will show that CFB is secure against leakage when instantiated with a PRLF, using an adaptation of the classical proof for CFB security [1].

Our first task is to express the leakage sets $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivD}})$ for scheme $\mathbf{iv}\boldsymbol{\mathcal{E}}$ in terms of that of the PRF $\boldsymbol{\mathcal{F}}$, namely $\mathcal{L}_{\mathsf{F}}$. When tracing through the operation of CFB-encryption, we will make two assumptions. Firstly, that leakage for each of the $\boldsymbol{\mathcal{F}}$ calls is local (cf. OCLI), which in particular means leakage will be restricted to the representation of $\boldsymbol{k}$ specific for the $\boldsymbol{\mathcal{F}}$ call at hand (and $\boldsymbol{k}$ is expected to be refreshed during a single $\mathbf{iv}\boldsymbol{\mathcal{E}}$ call). Secondly, that all visible wires in Fig. 7, corresponding to the $\mathbf{iv}\boldsymbol{\mathcal{E}}$'s public inputs and outputs, will leak. Note that longer messages will lead to more leakage for an adversary, which matches practice (where the size of the power trace might be linear in the size of the message.

Decryption closely matches encryption and, under the same assumptions as above, leakage on decryption of a ciphertext can be expressed instead as leakage on the encryption of the corresponding plaintext. Hence we refer to decryption

leakage as $\mathcal{L}_{\mathsf{ivE}}{}'$ (where the prime connotes the syntactical malarkey to deal with the different input spaces for encryption and decryption).

Concluding, we define the leakage set $\mathcal{L}_{\mathsf{ivE}}$ to be the collection of all functions $L_{\mathrm{CFB}} : \mathsf{K} \times \mathsf{N} \times \mathsf{M} \times \mathsf{R} \to \{0,1\}^*$ that are of the form

$$L_{\mathrm{CFB}}(\boldsymbol{k}, I, M; r) = (M, C, L_i(\boldsymbol{k}_i, C_i; r_i)_{i \in \{0..n-1\}})$$

with $L_i \in \mathcal{L}_{\mathsf{T}}$ (for $i \in \{0..n-1\}$) and where $M$ is an $n$-block message, $C = \mathrm{iv}\mathcal{E}_k(I, M)$ is an $(n+1)$-block ciphertext constituted of blocks $C_i$ ($i \in \{0, ..n\}$), $r$ is the concatenation of the random values $r_i$ passed to the $i^{\mathrm{th}}$ $\boldsymbol{\mathcal{F}}$-call ($i \in \{1..n\}$), and $\boldsymbol{k}_{i-1}$ is the key representation for the $i^{\mathrm{th}}$ $\boldsymbol{\mathcal{F}}$-call ($i \in \{1..n\}$).

**Theorem 4.** *Let $\boldsymbol{\mathcal{F}} : \mathsf{T}^* \to \mathsf{T}$ be a PRF with leakage class $\mathcal{L}_{\mathsf{F}}$ and let $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$ be the symmetric encryption scheme* CFB$[\boldsymbol{\mathcal{F}}]$ *with derived leakage $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivE}}{}')$. Let $\mathbb{A}$ be an iv-respecting adversary against the $(\mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivE}}{}')$-IND–aCPLA security of $(\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}})$. Then there exists an adversary $\mathbb{A}_{\mathrm{PRF}}$ of similar complexity to $\mathbb{A}$ against the $\mathcal{L}_{\mathsf{F}}$-PRLF security of $\boldsymbol{\mathcal{F}}$ such that*

$$\mathbf{Adv}^{\mathrm{IND-aCPLA}}_{\mathbf{iv}\boldsymbol{\mathcal{E}}, \mathbf{iv}\boldsymbol{\mathcal{D}}; \mathcal{L}_{\mathsf{ivE}}, \mathcal{L}_{\mathsf{ivE}}{}'}(\mathbb{A}) \leq 2 \cdot \mathbf{Adv}^{\mathrm{PRLF}}_{\boldsymbol{\mathcal{F}}; \mathcal{L}_{\mathsf{F}}}(\mathbb{A}_{\mathrm{PRF}}) + \frac{3}{4} \cdot \frac{\sigma^2}{|\mathsf{T}|} \ ,$$

*where $\sigma$ is the total number of blocks encrypted, and the blocksize is $|\mathsf{T}|$.*

The proof can be found in the full version [3].

## 6 mrLAE Security by Instantiating the PRF and MAC

The A5 and SIVAT composition mechanisms can be instantiated with any suitably secure primitives to yield secure nLAE or mrLAE schemes. Together with using CFB$[\boldsymbol{\mathcal{F}}]$ as underlying $\mathbf{iv}\boldsymbol{\mathcal{E}}$, these allow us to construct a secure mrLAE scheme through any PRF $\mathcal{F}$ with a secure implementation $\boldsymbol{\mathcal{F}}$ and a secure MAC implementation $(\boldsymbol{\mathcal{T}}, \boldsymbol{\mathcal{V}})$. The remaining questions therefore are what can be said about securely implementing these primitives and what conclusions for the overall scheme can subsequently be drawn. We will answer these questions from two perspectives: a practical side-channel one (for those favouring masked AES) and a more theoretical, yet eminently implementable one in the continuous leakage model.

**A side-channel perspective.** Our result provides a roadmap for obtaining a side-channel misuse-resistant AE scheme by selecting reasonable practical primitives (and implementations) for the PRF and the MAC (say a suitably masked AES, respectively KMAC) and subsequently gauging to what extent actual leakage on the *primitive implementations* can be used to break the relevant PRLF or EUF–CLMA notions as well as whether leakage on the *full* implementation is cleanly segregated or whether undesired correlation indicates bleeding of leakage from the values or variables from one component into say part of the power trace associated with another component.

The result above no longer explicitly takes into account leakage classes; these have effectively become implicit artefacts of the attack. We assume that a successful attack on the full scheme will be recognized as such: our result essentially says that if such an attack is found then either the leakage is not cleanly separated or one of the primitive implementations is already insecure (or both).

**A leakage resilience perspective.** A complementary approach to the practical one above is to design the primitives and their implementations with a provable level of resistance against leakage functions from a specific class. As already explained in the introduction, a multitude of models exist depending on the class of functions under consideration. One of the stronger models is that of continuous leakage: here the leakage functions can be arbitrary, subject to the constraint that their range is bounded. A usual refinement is to use a split-state model, where the key's representation $\boldsymbol{k}$ is operated upon in two (or more) tranches and each tranch can only leak on that part of $\boldsymbol{k}$ in scope for the operation at hand (assuming only computation leaks, as usual).

While there are PRFs that have been proven secure in the continuous leakage model, as far as we can tell this has always come at the price of adaptivity. In order for our constructions to be implemented a new PRF is called for, with an implementation secure in the stronger, adaptive continuous leakage model. In the full version [3] we provide such a function and implementation, and prove the latter secure in the generic group model (against adaptive continous leakage in a split-state setting). Additionally, we show how to create a related MAC such that leaking on the verification's implementation is ok.

Our construction is an evolution of the MAC of Martin et al. [30], itself inspired by a scheme by Kiltz and Pietrzak [23]. The key enabling novelty is the use of *three* shares instead of the customary two. A thorough discussion of the design choices, specifications, and security justification can be found in the full version [3] but for completeness we provide the final theorem statement below.

**Theorem 5.** *Let* $\mathsf{SIVAT}$ *be the SIVAT mechanism instantiated with the implementations described in the full version [3] over a generic group of* $p$ *elements, and assume that each share of the internal PRF leaks at most* $\lambda$ *per call following the associated leakage functions, as described by leakage sets* $(\mathcal{L}_{\mathsf{Enc}}, \mathcal{L}_{\mathsf{Dec}})$*. Let* $\mathbb{A}$ *be an adversary making at most* $g$ *direct queries to the generic group oracle (including the complexity of all chosen leakage queries) and making* $q$ *construction queries totalling* $\sigma$ *blocks. Then,*

$$\mathbf{Adv}^{\mathrm{LAE}}_{\mathsf{SIVAT};\mathcal{L}_{\mathsf{Enc}},\mathcal{L}_{\mathsf{Dec}}}(\mathbb{A}) \leq \frac{7}{p}\left(2^{4\lambda}\cdot\sigma^2\cdot(g+9q+5\sigma)^2 + 8(g+9q+5\sigma)^2\right).$$

To get a feel for the practical security level, let's look at parameters if the schemes are instantiated over a 512 bit elliptic curves, and we want the keep the attack success probability below $2^{-60}$ (a common limit in the real world, e.g. [28]). Let's assume that each internal leakage function leaks at most $\lambda = 85$ bits, which is approximately a sixth of a group element. Then the scheme would

remain secure until the adversary has encrypted or decrypted around $2^{25}$ blocks, and made a similar number of queries to the generic group.

This result comes with a few caveats, covered in more detail by the full version [3]. For instance, to ensure security against the leakage of arbitrary functions of the key, to process $q$ queries of total $\sigma$ blocks the construction must sample $4q + \sigma$ random group elements in a leakage-resilient manner, which can be complicated [30]. Nonetheless, our construction is proof positive of the existence of leakage resilient authenticated encryption in a very strong sense.

## 7 Conclusions and Open Problems

We introduced notions for strengthened AE when considering leakage, discussed generic composition under leakage, and showed the EtM type constructions can be proven secure in this context. We give a new scheme, SIVAT, that achieves misuse resistance and leakage resilience simultaneously, and show how this can be bootstrapped from a PRF secure against leakage. Finally, we give a concrete instantiation for the SIVAT mechanism. Our research unveils several interesting open problems, which we summarise subsequently.

*IND–aCPLA.* If one allows nonce-reuse, then for any leakage set $\mathcal{L}_{\mathsf{E}}$ security against $\mathcal{L}_{\mathsf{E}}$-IND–CPLA adversary implies $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{E}}{}')$-IND–aCPLA security, where $\mathcal{L}_{\mathsf{E}}{}'$ is the essentially the same set as $\mathcal{L}_{\mathsf{E}}$ with some minor bookkeeping to ensure correct syntax. The implication is trivial as the leakage on any valid $\mathcal{D}$-query can be perfectly simulated by repeating the corresponding $\mathcal{E}$-query instead. In the the nonce or iv respecting cases the implication remains open (as repeating encryption queries including nonce is no longer allowed). Nonetheless, we conjecture that even in these two settings for many reasonable leakage sets $\mathcal{L}_{\mathsf{E}}$, $\mathcal{L}_{\mathsf{E}}$-IND–CPLA does imply $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{E}}{}')$-IND–aCPLA. We leave it as an interesting question to formalise this or find a counter-example. More generally, is there some way of defining $\mathcal{L}_{\mathsf{ED}}$ as a function of some general sets $\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}}$ such that $\mathcal{L}_{\mathsf{ED}}$-IND–CPLA $\implies$ $(\mathcal{L}_{\mathsf{E}}, \mathcal{L}_{\mathsf{D}})$-IND–aCPLA?

*MtE with restricted leakage sets.* The insecurity of the majority of the MtE schemes when considering leakage comes from a generic attack against any schemes whose inverse follows the decrypt-then-verify or decrypt-and-verify structure. We leave it as an interesting open question to investigate the leakage security under other more restricted leakage sets.

*Misuse resistance with minimal message expansion.* We demonstrate that misuse resistance can be achieved through generic composition, at the cost of additional message expansion, using a MAC-then-Encrypt-then-MAC structure (leading to SIVAT). We believe that dedicated constructions are likely to exist that can achieve mrLAE security with minimal expansion, or more generally LAE without requiring independent keys.

**Acknowledgements.** Initial work was conducted while Dan Martin was employed by the Department of Computer Science, University of Bristol. Guy Barwell was supported by an EPSRC grant; Elisabeth Oswald and Dan Martin were in part supported by EPSRC via grants EP/I005226/1 (SILENT) and EP/N011635/1 (LADA).

## References

1. Alkassar, A., Geraldy, A., Pfitzmann, B., Sadeghi, A.R.: Optimized self-synchronizing mode of operation. In: Matsui, M. (ed.) FSE 2001. LNCS, vol. 2355, pp. 78–91. Springer, Heidelberg (Apr 2002)
2. Andreeva, E., Bogdanov, A., Luykx, A., Mennink, B., Mouha, N., Yasuda, K.: How to securely release unverified plaintext in authenticated encryption. In: Sarkar and Iwata [43], pp. 105–125
3. Barwell, G., Martin, D.P., Oswald, E., Stam, M.: Authenticated encryption in the face of protocol and side channel leakage. Cryptology ePrint Archive, Report 2017/068 (2017), http://eprint.iacr.org/2017/068
4. Barwell, G., Page, D., Stam, M.: Rogue decryption failures: Reconciling AE robustness notions. In: Groth [17], pp. 94–111
5. Bellare, M., Kane, D., Rogaway, P.: Big-key symmetric encryption: Resisting key exfiltration. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 373–402. Springer, Heidelberg (Aug 2016)
6. 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 (Dec 2000)
7. Bernstein, D.J.: CAESAR competition call (2013), http://competitions.cr.yp.to/caesar-call-3.html
8. Berti, F., Koeune, F., Pereira, O., Peters, T., Standaert, F.X.: Leakage-resilient and misuse-resistant authenticated encryption. Cryptology ePrint Archive, Report 2016/996 (2016), http://eprint.iacr.org/2016/996
9. Boldyreva, A., Degabriele, J.P., Paterson, K.G., Stam, M.: On symmetric encryption with distinguishable decryption failures. In: Moriai, S. (ed.) FSE 2013. LNCS, vol. 8424, pp. 367–390. Springer, Heidelberg (Mar 2014)
10. Boldyreva, A., Degabriele, J.P., Paterson, K.G., Stam, M.: Security of symmetric encryption in the presence of ciphertext fragmentation. Cryptology ePrint Archive, Report 2015/059 (2015), http://eprint.iacr.org/2015/059
11. Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. Cryptology ePrint Archive, Report 2016/1013 (2016), http://eprint.iacr.org/2016/1013
12. Dodis, Y., Kalai, Y.T., Lovett, S.: On cryptography with auxiliary input. In: Mitzenmacher, M. (ed.) 41st ACM STOC. pp. 621–630. ACM Press (May / Jun 2009)
13. Dodis, Y., Pietrzak, K.: Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 21–40. Springer, Heidelberg (Aug 2010)
14. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS. pp. 293–302. IEEE Computer Society Press (Oct 2008)
15. Faust, S., Pietrzak, K., Schipper, J.: Practical leakage-resilient symmetric cryptography. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 213–232. Springer, Heidelberg (Sep 2012)

16. Fischlin, M., Günther, F., Marson, G.A., Paterson, K.G.: Data is a stream: Security of stream-based channels. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 545–564. Springer, Heidelberg (Aug 2015)
17. Groth, J. (ed.): 15th IMA International Conference on Cryptography and Coding, LNCS, vol. 9496. Springer, Heidelberg (Dec 2015)
18. Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 160–176. Springer, Heidelberg (May 2013)
19. Hoang, V.T., Krovetz, T., Rogaway, P.: Robust authenticated-encryption AEZ and the problem that it solves. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 15–44. Springer, Heidelberg (Apr 2015)
20. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (Aug 2003)
21. Katz, J., Lindell, Y.: Introduction to Modern Cryptography. Chapman & Hall/CRC (2008)
22. Katz, J., Vaikuntanathan, V.: Signature schemes with bounded leakage resilience. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 703–720. Springer, Heidelberg (Dec 2009)
23. Kiltz, E., Pietrzak, K.: Leakage resilient ElGamal encryption. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 595–612. Springer, Heidelberg (Dec 2010)
24. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO'96. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (Aug 1996)
25. Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (Aug 1999)
26. Kurosawa, K., Phong, L.T.: Leakage resilient IBE and IPE under the DLIN assumption. In: Jacobson Jr., M.J., Locasto, M.E., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 13. LNCS, vol. 7954, pp. 487–501. Springer, Heidelberg (Jun 2013)
27. Longo, J., Martin, D.P., Oswald, E., Page, D., Stam, M., Tunstall, M.: Simulatable leakage: Analysis, pitfalls, and new constructions. In: Sarkar and Iwata [43], pp. 223–242
28. Luykx, A., Paterson, K.: Limits on authenticated encryption use in tls (2016), http://www.isg.rhul.ac.uk/ kp/TLS-AEbounds.pdf
29. Mangard, S., Oswald, E., Popp, T.: Power Analysis Attacks: Revealing the Secrets of Smart Cards. Springer (2008)
30. Martin, D.P., Oswald, E., Stam, M., Wójcik, M.: A leakage resilient MAC. In: Groth [17], pp. 295–310
31. Micali, S., Reyzin, L.: Physically observable cryptography (extended abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (Feb 2004)
32. Namprempre, C., Rogaway, P., Shrimpton, T.: Reconsidering generic composition. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 257–274. Springer, Heidelberg (May 2014)
33. NIST: FIPS 81: DES Modes of Operation. Issued December 2, 63 (1980)
34. Pereira, O., Standaert, F.X., Vivek, S.: Leakage-resilient authentication and encryption from symmetric cryptographic primitives. In: Ray, I., Li, N., Kruegel:, C. (eds.) ACM CCS 15. pp. 96–108. ACM Press (Oct 2015)
35. Perrin, T.: Double ratchet algorithm (2014), https://github.com/trevp/double_ratchet/wiki, Retrieved 2016-09-01

36. Pietrzak, K.: A leakage-resilient mode of operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (Apr 2009)
37. Qin, B., Liu, S.: Leakage-flexible CCA-secure public-key encryption: Simple construction and free of pairing. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 19–36. Springer, Heidelberg (Mar 2014)
38. Renauld, M., Standaert, F.X., Veyrat-Charvillon, N., Kamel, D., Flandre, D.: A formal study of power variability issues and side-channel attacks for nanoscale devices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 109–128. Springer, Heidelberg (May 2011)
39. Rogaway, P.: Authenticated-encryption with associated-data. In: Atluri, V. (ed.) ACM CCS 02. pp. 98–107. ACM Press (Nov 2002)
40. Rogaway, P.: Nonce-based symmetric encryption. In: Roy, B.K., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 348–359. Springer, Heidelberg (Feb 2004)
41. Rogaway, P., Bellare, M., Black, J., Krovetz, T.: OCB: A block-cipher mode of operation for efficient authenticated encryption. In: ACM CCS 01. pp. 196–205. ACM Press (Nov 2001)
42. 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 (May / Jun 2006)
43. Sarkar, P., Iwata, T. (eds.): ASIACRYPT 2014, Part I, LNCS, vol. 8873. Springer, Heidelberg (Dec 2014)
44. Schipper, J.: Leakage-Resilient Authentication. Ph.D. thesis, Utrecht University (2010)
45. Shrimpton, T., Terashima, R.S.: A modular framework for building variable-input-length tweakable ciphers. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part I. LNCS, vol. 8269, pp. 405–423. Springer, Heidelberg (Dec 2013)
46. Standaert, F.X., Pereira, O., Yu, Y.: Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 335–352. Springer, Heidelberg (Aug 2013)
47. Yau, A.K.L., Paterson, K.G., Mitchell, C.J.: Padding oracle attacks on CBC-mode encryption with secret and random IVs. In: Gilbert, H., Handschuh, H. (eds.) FSE 2005. LNCS, vol. 3557, pp. 299–319. Springer, Heidelberg (Feb 2005)
48. Yu, Y., Standaert, F.X., Pereira, O., Yung, M.: Practical leakage-resilient pseudorandom generators. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM CCS 10. pp. 141–151. ACM Press (Oct 2010)