

Sponges Resist Leakage: The Case of Authenticated Encryption^{*}

Jean Paul Degabriele¹, Christian Janson², and Patrick Struck³

¹ CNS, Technische Universität Darmstadt, Germany
`jeanpaul.degabriele@cysec.de`

² Cryptoplexity, Technische Universität Darmstadt, Germany
`christian.janson@cryptoplexity.de`

³ CDC, Technische Universität Darmstadt, Germany
`pstruck@cdc.tu-darmstadt.de`

Abstract. In this work we advance the study of leakage-resilient Authenticated Encryption with Associated Data (AEAD) and lay the theoretical groundwork for building such schemes from sponges. Building on the work of Barwell et al. (ASIACRYPT 2017), we reduce the problem of constructing leakage-resilient AEAD schemes to that of building fixed-input-length function families that retain pseudorandomness and unpredictability in the presence of leakage. Notably, neither property is implied by the other in the leakage-resilient setting. We then show that such a function family can be combined with standard primitives, namely a pseudorandom generator and a collision-resistant hash, to yield a nonce-based AEAD scheme. In addition, our construction is quite efficient in that it requires only two calls to this leakage-resilient function per encryption or decryption call. This construction can be instantiated entirely from the T-sponge to yield a concrete AEAD scheme which we call SLAE. We prove this sponge-based instantiation secure in the non-adaptive leakage setting. SLAE bears many similarities and is indeed inspired by ISAP, which was proposed by Dobraunig et al. at FSE 2017. However, while retaining most of the practical advantages of ISAP, SLAE additionally benefits from a formal security treatment.

Keywords: AEAD · Leakage Resilience · Side Channels · SLAE · ISAP

1 Introduction

The oldest and most fundamental application of cryptography is concerned with securing the communication between two parties who already share a secret key. The modern cryptographic construct for this application is authenticated encryption with associated data (AEAD), which was the topic of the recent CAESAR competition [6]. Most of the effort in this competition has been directed towards exploring new designs, optimising performance, and offering robust security guarantees. However, there has not been much progress in the development of AEAD constructions that, by design, protect against side-channel attacks. This is a challenging problem that is likely to become a primary focus in the area of AEAD design.

Recently, a handful of AEAD designs with this exact goal have emerged. Each of these is based on a different approach with varying trade-offs between complexity, efficiency, and security guarantees. One notable example is the work of Barwell et al. [4], which proposes AEAD constructions with strong security guarantees but pays a relatively high price in terms of complexity and efficiency. Specifically, their constructions achieve security against adaptive leakage but resort to elliptic-curve pairings and secret sharing in order to realise implementations of a leakage-resilient MAC and a leakage-resilient pseudorandom function (employed in a block-wise fashion for encryption) for instantiating their scheme. A more hands-on approach was adopted by Dobraunig et al. in the design of their proposed AEAD scheme ISAP. It was conceived with the intent to protect against Differential Power Analysis (DPA) [12]. ISAP is entirely sponge-based and follows a fairly conventional design, augmented with a rekeying strategy. Arguably, this simpler approach, employing readily-available symmetric primitives, is more likely to lead to a pragmatic solution. However, ISAP's design rationale is predominantly heuristic, lacking any formal security analysis to justify its claims. As such the efficacy of ISAP's approach in resisting side-channel attacks is unclear, both qualitatively and quantitatively, curtailing any objective comparison with the constructions from [4] and others.

^{*} © IACR 2019. This article is the full version (major revision) of the version published by Springer-Verlag available at DOI: [10.1007/978-3-030-34621-8_8](https://doi.org/10.1007/978-3-030-34621-8_8).

In light of the practical advantages that the sponge-based approach offers, we remedy this state of affairs as follows. We propose SLAE, a derivative of ISAP which retains its main structure and benefits but includes certain modifications to admit a formal security proof. We analyse its security in the framework of leakage-resilient cryptography introduced by Dziembowski and Pietrzak [16], adapted to the random transformation model. Specifically, we prove it secure with respect to the leakage-resilient AEAD definition, put forward in [4] by Barwell et al., in the non-adaptive leakage setting. That is, we assume a leakage function that is fixed a priori and whose output is limited to some number of bits λ .

Admittedly, SLAE achieves qualitatively weaker security than the schemes of Barwell et al., since it only achieves non-adaptive leakage resilience. Nevertheless, we contend that SLAE strikes a more pragmatic balance by improving on efficiency and ease of implementation while still benefiting from a provably-secure design. Indeed, several other works [1, 15, 17, 25, 27] have settled for and argued that non-adaptive leakage security often suffices in practice. Moreover, as discussed in [27], the syntax of primitives like pseudorandom functions makes adaptive-leakage security impossible to achieve. In fact Barwell et al. achieve security against adaptive leakage by resorting to a specialised implementation of a pseudorandom function which requires an additional random input per invocation. In contrast, SLAE adheres to the standard nonce-based AEAD syntax and requires no source of randomness.

When viewed as sponge-based constructions, SLAE and ISAP look very similar and we do not claim any particular novelty in that respect. Nevertheless, the rationale behind their design is rather different. ISAP was conceived as augmenting a standard sponge-based AEAD design with a rekeying strategy, where the rekeying function is in turn also built from sponges, followed by some optimisations. The rekeying is intended to frustrate Differential Power Analysis (which requires several power traces on the same key but distinct inputs) by running the AEAD scheme with a distinct session key each time its inputs change. In turn, the session key is produced by combining a hash of the inputs and the master key through a rekeying function. Ostensibly, the rekeying function is itself strengthened against DPA by reducing its input data complexity through a low sponge absorption rate. In contrast SLAE is understood through a top-down design where we gradually decompose a leakage-resilient AEAD scheme into smaller components which we then instantiate using sponges. In particular there is no mention of rekeying or session keys. Note that there is more to this distinction than mere renaming. For instance, if we compare the MAC components in ISAP and SLAE we notice that the same value that serves as the MAC session key in ISAP is used directly as the MAC tag in SLAE.

At a more general level, the key premise made in [12] is that sponges offer a promising and practical solution to protect against side-channel attacks. Our work serves to provide formal justification to this claim and allows one to calculate concrete parameters for a desired security level.

1.1 Contribution

Below is an outline of our contributions highlighting how we improve on prior works and some of the challenges we face in our analysis.

A Generic Construction (FGHF'). The composition theorem in [4] reconsiders the N2 construction from [22] in the setting of leakage resilience. Specifically they show that given a MAC that is both *leakage-resilient strongly unforgeable* and a *leakage-resilient pseudorandom function*, together with an encryption scheme that is *leakage-resilient against augmented chosen plaintext attacks*, the N2 construction yields a leakage-resilient AEAD scheme. We extend this result, in the non-adaptive setting, by further decomposing the MAC and the encryption scheme into simpler lower-level primitives, ultimately giving rise to the FGHF' construction. In turn this constructs a leakage-resilient AEAD scheme from two *fixed-size* leakage-resilient functions \mathcal{F} and \mathcal{F}' , a standard pseudorandom generator \mathcal{G} , and a collision-resistant vector hash \mathcal{H} . The construction requires that both \mathcal{F} and \mathcal{F}' be leakage-resilient pseudorandom functions and that \mathcal{F}' additionally be a leakage-resilient unpredictable function. The latter is a notion that we introduce.

As pointed out in [4], in the adaptive leakage setting any MAC whose verification algorithm recomputes the tag and checks for equality with the candidate tag, simply cannot be strongly unforgeable. They overcome this issue through an ingenious MAC implementation. However this requires three pairing evaluations per verification and a source of randomness. In the FGHF' construction we show that by settling for non-adaptive leakage security the canonical MAC construction, which recomputes the tag and checks for equality, can be rescued. Specifically, we show that any leakage-resilient unpredictable function gives rise to a canonical MAC which is strongly unforgeable. In contrast to the leakage-free setting, not every pseudorandom function is an unpredictable function. This has to do with the fact that in unpredictability

we give the adversary more freedom in what it can query to its oracles, which is in turn a necessary requirement for composition to hold. In addition, we prove that one can combine a collision-resistant hash function with fixed-input-length leakage-resilient pseudorandom and unpredictable functions to obtain corresponding primitives with extended input domains.

For the encryption part, Barwell et al. use Counter Feedback Mode instantiated with a leakage-resilient pseudorandom function and an additional extra call to generate the initial vector from the nonce. Thus multiple calls to the leakage-resilient pseudorandom function are required for each encryption call. In contrast we show that to meet the required security notion, one can do with just one call to the leakage-resilient pseudorandom function and a pseudorandom generator, thereby resulting in a considerably more efficient scheme. Thus, if one is content with non-adaptive leakage security then the FGHF' construction constitutes a simpler recipe yielding a more efficient AEAD scheme.

All the results needed to prove the security of the FGHF' construction hold in the general adaptive setting. The limitation to the non-adaptive leakage setting comes from the fact that leakage-resilient unpredictable functions are unattainable in the adaptive-leakage setting if no further restriction is imposed on the set of leakage functions.

Non-Adaptively Leakage-Resilient Functions from Sponges. Having reduced the task of constructing a leakage-resilient AEAD scheme to that of constructing suitable leakage-resilient function families, we turn our attention to the latter problem. We instantiate both \mathcal{F} and \mathcal{F}' with the same sponge-based construction, which we refer to as SLFUNC. This construction is essentially the rekeying function employed in ISAP [12] instantiated with a random transformation (T-sponge) instead of a random permutation (P-sponge). In [12] this was proposed without proof, instead its security was argued based on its apparent similarity to the GGM construction [18] and the corresponding results in [17,25] for it yielding a leakage-resilient pseudorandom function family. However, there are clear differences between the sponge construction and the GGM construction and we do not see a way to make a direct connection between the security of the two. In fact our proof follows a fairly different strategy from the ones presented in [17,18,25] – which all rely on a hybrid argument whereas ours does not. Moreover, for the overall security of SLAE we need this function family to additionally be leakage-resilient unpredictable, which, as was discussed above, does not follow from it being leakage-resilient pseudorandom. We show that this construction achieves both security notions when the absorption rate is set to a sufficiently small value.

Another technical challenge that we face here is that we cannot employ the H-coefficient technique which is commonly used to prove the security of various sponge-based constructions. Like most other works on leakage resilience, we resort to arguments based on min-entropy and its chain rule in order to deal with leakage. Unfortunately, such arguments do not combine well with the H-coefficient technique, which precludes us from using it. In turn, this renders the security proof more challenging, as we have to deal with an adversary that may choose its queries (not the leakage function) adaptively. In contrast, the H-coefficient technique would automatically bypass this issue by reducing the security proof to a counting problem.

A Concrete Sponge-Based AEAD Scheme (SLAE). Finally, by instantiating the FGHF' construction with the above sponge-based construction for \mathcal{F} and \mathcal{F}' and matching sponge-based constructions for \mathcal{G} and \mathcal{H} we obtain SLAE. We also present security proofs for the T-sponge instantiations of the pseudorandom generator and the vector hash, which we were unable to readily find in the literature. SLAE is perhaps our most practical contribution – an entirely sponge-based leakage-resilient nonce-based AEAD scheme with provable security guarantees that is simple to implement and reasonably efficient. The efficiency of SLAE could be further optimised using similar techniques to the ones described in [12] for ISAP. Furthermore our security proofs are conducted in the concrete security setting thereby allowing practitioners to easily derive parameter estimates for their desired security level.

1.2 Related Work

To the best of our knowledge, the first authenticated encryption scheme claimed to be leakage-resilient was RCB [3], but it was broken soon after [2].

A series of works [7, 8, 19, 23] have proposed a number of leakage-resilient symmetric encryption schemes, message authentication codes, and authenticated encryption schemes. These constructions assume that a subset of their components (block cipher instances) are leakage-free and that the leakage

in the other components is simulatable, an assumption that is somewhat contentious [21, 26]. Based on these assumptions, they show that the security of their encryption schemes reduces to the security of a single-block variant of the same scheme. However, the security of the corresponding single-block schemes remains an open question that is implicitly assumed to hold.

Abdalla, Belaïd, and Fouque [1] construct a symmetric encryption scheme that is non-adaptively leakage-resilient against chosen-plaintext attacks. Interestingly, their scheme employs a rekeying function that is not a leakage-resilient pseudorandom function. However their encryption scheme is not nonce-based as it necessitates a source of randomness.

In independent and concurrent work [13] Dobraunig and Mennink analyse the leakage resilience of the duplex sponge construction. While their leakage model is closer to ours, they prove something different. Namely they show that the duplex is indistinguishable from an *adjusted ideal extendible input function* (AIXIF) which is an ideal functionality incorporating leakage. In contrast we show that SLFUNC is both a leakage-resilient PRF (LPRF) and a leakage-resilient unpredictable function (LUF), and then leverage this to construct a leakage-resilient AEAD scheme.

Other independent and concurrent work by Guo et al. [20] proposes an AEAD design, TETSPongE, that combines a sponge construction with two tweakable block cipher instances. While their work and ours share the goal of constructing leakage-resilient AEAD schemes, the two works adopt very different approaches. Both the security definitions and the assumptions on which the security of the schemes rely on are significantly different. One notable difference, is that the leakage resilience of TETSPongE relies crucially on the tweakable block cipher instances being leak-free, presumably due to a hardened implementation, whereas our treatment exploits and exposes the inherent leakage resilience of the sponge construction.

1.3 Organization of the Paper

In Section 2 we review the basic concepts and security definitions that we require in the rest of the paper. This is followed by a detailed description of SLAE in Section 3. In Section 4 we cover the security analysis of the generic FGHF' construction. In Section 5 we cover the security of the sponge-based primitives used to instantiate FGHF' and thereby obtain SLAE. We conclude in Section 6 with some remarks on implementing SLAE. The full details of the proofs can be found in the Appendix.

2 Preliminaries

We start by reviewing the basic tools and definitions that we require for our results. We begin by establishing some notation. Some additional standard definitions appear in Appendix A.

2.1 Notation

For any non-negative integer $n \in \mathbb{N}$ we use $[n]$ to denote the set $\{1, \dots, n\}$, where $[n] = \emptyset$ when $n = 0$. For any two strings s_1 and s_2 , $|s_1|$ denotes the size of s_1 and $s_1 \parallel s_2$ denotes their concatenation. For a positive integer $k \leq |s_1|$, we use $[s_1]_k$ to denote the string obtained by truncating s_1 to its leftmost k bits. The empty string is denoted by ε , $\{0, 1\}^n$ denotes the set of bit strings of size n , and $\{0, 1\}^*$ denotes the set of all strings of finite length. We write $x \leftarrow \mathcal{S}$ to denote the process of uniformly sampling a value from the finite set \mathcal{S} and assigning it to x .

We make use of the code-based game-playing framework by Bellare and Rogaway [5], where the interaction between a game and the adversary is implicit. In all games, the adversary is given as its input the output of the initialize procedure, it has oracle access to the other procedures described in the game, and its output is fed into the finalize procedure. The output of the finalize procedure is the output of the game. For a game G and an adversary \mathcal{A} , $G^{\mathcal{A}} \Rightarrow y$ denotes the event that G outputs y when interacting with \mathcal{A} . Similarly, $\mathcal{A}^G \Rightarrow x$ denotes the event that \mathcal{A} outputs x when interacting with G . By convention all boolean variables **Bad** are initialized to **false**, and for any table $p[\cdot]$ its entries are all initialized to \perp . When lazy-sampling a random function with domain \mathcal{X} and co-domain \mathcal{Y} into a table $p[\cdot]$, we use $\text{inset}(p)$ and $\text{outset}(p)$ to denote respectively the sets of input and output values defined up to that point. That is, $\text{inset}(p) = \{X : p[X] \neq \perp \wedge X \in \mathcal{X}\}$ and $\text{outset}(p) = \{p[X] : p[X] \neq \perp \wedge X \in \mathcal{X}\}$. If G_1 and G_2 are games and \mathcal{A} is an adversary we define the corresponding *adversarial advantage* as

$$\text{Adv}(\mathcal{A}^{G_1}, \mathcal{A}^{G_2}) = \Pr[\mathcal{A}^{G_1} \Rightarrow 1] - \Pr[\mathcal{A}^{G_2} \Rightarrow 1],$$

and the corresponding *game advantage* as

$$\mathbf{Adv}(G_1^A, G_2^A) = \Pr[G_1^A \Rightarrow \text{true}] - \Pr[G_2^A \Rightarrow \text{true}].$$

We operate in the random transformation model, where ρ is an idealised random transformation mapping n -bit strings to n -bit strings. For any algorithm \mathcal{F} that uses ρ as a subroutine, we use $Q_{\mathcal{F}}(q, \mu)$ to denote the number of calls to ρ required when evaluating \mathcal{F} q times on a total of μ bits.

2.2 Syntax

Encryption. An *authenticated encryption scheme with associated data* $\text{AEAD} = (\mathcal{E}, \mathcal{D})$ is a pair of efficient algorithms such that:

- The deterministic encryption algorithm $\mathcal{E} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \mathcal{M} \rightarrow \{0, 1\}^*$ takes as input a secret key K , a nonce N , associated data A , and a message M to return a ciphertext C .
- The deterministic decryption algorithm $\mathcal{D} : \mathcal{K} \times \mathcal{N} \times \mathcal{A} \times \{0, 1\}^* \rightarrow \mathcal{M} \cup \{\perp\}$ takes as input a secret key K , a nonce N , associated data A , and a ciphertext C to return either a message in \mathcal{M} or \perp indicating that the ciphertext is invalid.

Sets \mathcal{K} , \mathcal{N} , \mathcal{A} , and \mathcal{M} denote respectively the key space, the nonce space, the associated data space, and the message space associated to the scheme. We assume throughout that \mathcal{E} and \mathcal{D} are never queried on inputs outside of these sets. An authenticated encryption scheme is required to be *correct* and *tidy*. Correctness requires that for all K, N, A, M if $\mathcal{E}(K, N, A, M) = C$ then $\mathcal{D}(K, N, A, C) = M$. Analogously, tidiness requires that for all K, N, A, C if $\mathcal{D}(K, N, A, C) = M \neq \perp$ then $\mathcal{E}(K, N, A, M) = C$. Furthermore we demand that encryption be length regular, i.e for all K, N, A, M it should hold that $|\mathcal{E}(K, N, A, M)|$ is entirely determined by $|N|$, $|A|$, and $|M|$.

We will use the terms *authenticated encryption scheme* and *symmetric encryption scheme* to refer to the analogously defined encryption scheme which does not admit associated data as part of its input. For such schemes, A is implicitly set to the empty string in the security games.

Message Authentication. A *message authentication code* $\text{MAC} = (\mathcal{T}, \mathcal{V})$ is a pair of efficient algorithms with an associated key space \mathcal{K} , domain \mathcal{X} , and tag length t such that:

- The deterministic tagging algorithm $\mathcal{T} : \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^t$ takes as input a key K and a value X to return a tag T of size t .
- The deterministic verification algorithm $\mathcal{V} : \mathcal{K} \times \mathcal{X} \times \{0, 1\}^t \rightarrow \{\top, \perp\}$ takes as input a key K , a value X , and a tag T to return either \top indicating a valid input or \perp otherwise.

We require that for any key $K \in \mathcal{K}$ and any admissible input $X \in \mathcal{X}$, if $T \leftarrow \mathcal{T}(K, X)$, then $\mathcal{V}(K, X, T) = \top$. When $\mathcal{X} = \{0, 1\}^*$ we end up with the usual MAC definition, however we will also consider MACs over tuples of strings, e.g. $\mathcal{X} = \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^*$. Such MACs were considered in [22] and we follow suit in referring to such MACs as vector MACs.

We say that a MAC is *canonical* if it is implicitly defined by \mathcal{T} , where $\mathcal{V}(K, X, T)$ consists of running $T' \leftarrow \mathcal{T}(K, X)$ and returning \top if $T' = T$ and \perp otherwise.

2.3 The Sponge Construction

The sponge construction is a versatile object that can be used to realise various cryptographic primitives. Several variations of the sponge exist, Fig. 1 illustrates the plain version of the sponge as originally introduced by Bertoni et al. [9]. We give here only a brief overview of its operation and the associated nomenclature that we will use throughout this paper.

The sponge operates iteratively on its inputs through a transformation ρ , and generally includes an *absorbing* phase and a *squeezing* phase. The transformation ρ maps strings of size n to strings of size n . Associated to the sponge are two other values called the rate r and the capacity c , where $n = r + c$. At any given iteration we refer to the output of the transformation as the state, which we denote by S . Furthermore, we denote the leftmost r bits of S by \bar{S} and the remaining c bits by \hat{S} . We will at times refer to \bar{S} and \hat{S} as the outer and inner parts of the state, respectively. In the absorbing phase an input M is “absorbed” iteratively r bits at a time. At iteration i input M_i is absorbed by letting $Y_i \leftarrow (M_i \oplus \bar{S}_i) \parallel \hat{S}_i$ and setting $S_{i+1} \leftarrow \rho(Y_i)$. The initial value of S may generally be set to a constant, a concatenation of a secret key and a constant, or by applying the transformation to either of these values.

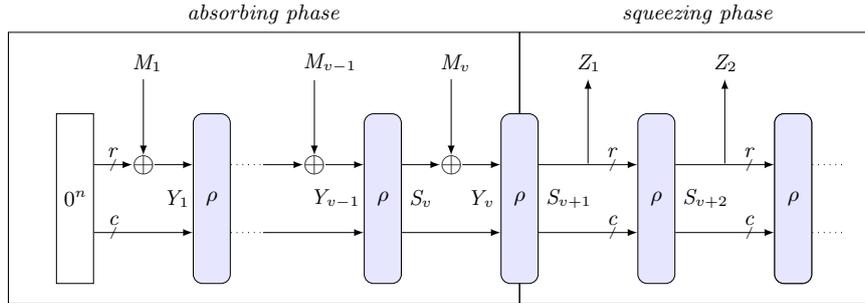


Fig. 1: Illustration of the plain sponge construction.

Output is produced from the sponge during the squeezing phase in one or more iterations, r bits at a time. At iteration i output Z_j is produced by setting $Z_j \leftarrow \bar{S}_i$ and $S_{i+1} \leftarrow \rho(S_i)$. The above variant is normally referred to as the T-sponge, as it employs a fixed-size random transformation. An alternative instantiation, known as the P-sponge, replaces this random transformation with a random permutation.

2.4 The Leakage Model

Our leakage model is based on *leakage resilience* as defined in [16]. This assumes that only computation leaks, and in particular, that only the data that is accessed during computation can leak information. It allows for continuous adaptive leakage, where in each query to a leakage oracle the adversary can specify a leakage function from some predefined set \mathcal{L} that it can choose adaptively based on prior outputs and leakage. Throughout, we restrict ourselves to leakage functions that are deterministic and efficiently computable. While our security definitions are formulated in this general setting, our main results will be in the weaker *granular non-adaptive* leakage setting proposed in [17]. We view the non-adaptive leakage setting as the special case where the leakage set \mathcal{L} is restricted to be a singleton, fixed at the start of the game. In granular leakage, a single time step is with respect to a single computation of some underlying primitive, in our case, the transformation ρ . Correspondingly, in this case the adversary specifies a vector of leakage functions and gets in return the aggregate leakage from the entire evaluation of the higher-level construction. Note that in the granular setting the leakage sets for each iteration can be distinct. Similarly, when studying the leakage resilience of composite constructions we have to consider compositions of leakage functions. For instance, if construction C is composed of primitives A and B with associated leakage sets \mathcal{L}_A and \mathcal{L}_B , then we associate to C the Cartesian product of the two leakage sets, i.e. $\mathcal{L}_C = \mathcal{L}_A \times \mathcal{L}_B$. The actual inputs that get fed to the leakage functions are implicitly defined by the construction and its inputs, whereas the combined output is the aggregate output of all function evaluations.

An analysis of sponge-based constructions compels us to consider leakage resilience in the random transformation model. A similar setting, albeit in the random oracle model, was already considered by Standaert et al. in [25]. A central question that arises in idealised settings like this is whether the leakage function should be given access to the ideal primitive. As in [25], we will not give this access to the leakage function. On the one hand, providing the leakage function with unlimited access to the random oracle gives rise to artificial attacks, such as the “future computation attack” discussed in [25], that would not arise in practice. On the other hand, depriving the leakage function from accessing the ideal primitive, means that the leakage function cannot leak any bits of the ideal primitive’s output, which may seem overly restrictive. However, for the case of sponge-based constructions this is less problematic because from the adversary’s perspective the full output of a transformation call is completely determined by the input to the next transformation call. As such, information about the output of one transformation call can leak as part of the leakage in the next transformation call. Combined with the fact that the only restriction that we will impose on the leakage function is to limit its output length, we think that this leads to a fairly realistic leakage model.

We conclude our discussion on the leakage model by offering our interpretation of the significance of leakage resilience security with respect to practical side channel attacks. One might object that we model leakage by a deterministic function whose output is of a fixed bit-length whereas in practice the leakage is noisy. However through the leakage function we are really trying to capture the maximum amount of information that an adversary may obtain from evaluating the scheme on a single input. Hence, the underlying assumption is that no matter how many times the scheme is run on the same

input, in order to even out the noise, the information that the adversary can obtain is limited. Put in more practical terms, this roughly translates to assuming that the scheme’s implementation resists Simple Power Analysis (SPA). On the other hand, if the scheme is proven to be leakage-resilient then we are guaranteed that an adversary cannot do much better even if it can observe and accumulate leakage on multiple other (differing) inputs. Thus a proof of leakage resilience can be interpreted as saying that if the scheme’s implementation is secure against SPA then, by the inherent properties of the scheme, it is also secure against Differential Power Analysis (DPA). However, a proof of leakage resilience is of course no guarantee that a scheme’s implementation will be secure against SPA.

2.5 Authenticated Encryption and Leakage Resilience

Recently, Barwell et al. [4] provided a definitional framework augmenting nonce-based authenticated encryption with leakage. Their security notions capture the leakage resilience setting as defined in [16]. Furthermore, they prove composition theorems analogous to [22] that additionally take leakage into account. Below we reproduce their security definitions and composition result which we will employ in this work, with some minor adaptations. We recast their definitions in a style that admits code-based proofs [5]. Unlike [4] we make no distinction between a scheme and its implementation since we are interested in proving security for the actual scheme. When defining these security notions, we only describe the game and the corresponding adversarial advantage. A scheme is understood to be secure if the adversarial advantage is bounded by a sufficiently small value for all reasonably-resourced adversaries. Our security theorems will then establish a bound on the adversarial advantage in terms of the adversary’s resources, without drawing judgement as to what constitutes “small” and “reasonable” since that is a rather subjective matter.

Classifying Adversarial Queries. As usual, the adversary has to be forbidden from making certain queries in order to avoid trivial win conditions. Following the terminology of [4], if an adversary makes a query (N, A, M) to an encryption oracle that returns C , then repeating this query to one of the encryption oracles or querying (N, A, C) to one of the decryption oracles, is considered to be an *equivalent* query. Note that any additional components of a query, such as the leakage function, are ignored for the purpose of determining equivalence between two queries. If an adversary makes equivalent queries across two oracles, it is said to *forward* that query from one oracle to the other. Note that the two oracles do not need to be distinct, and thus forwarded queries include repeated queries to the same oracle.⁴

Let an encryption query refer to any query made to either a challenge encryption oracle or a leakage encryption oracle. Then an adversary against an (authenticated) encryption scheme is said to be *nonce respecting* if it never repeats a nonce in two distinct encryption queries.

Chosen-Plaintext Security with Leakage. Barwell et al. introduce an augmented variant of leakage-resilient chosen-plaintext security called IND-aCPLA, that is required by their composition theorem. Here the adversary is given access to three oracles. A challenge oracle that returns either a valid encryption of a message or a random string of appropriate length. A leakage encryption oracle that, upon being queried on a message and a leakage function, returns the corresponding ciphertext and the evaluated leakage. The adversary is not allowed to forward queries between the two encryption oracles. In addition, it has limited access to a leakage decryption oracle which returns the decryption of the queried ciphertext and the leakage corresponding to the queried leakage function. However, it can only query this oracle on inputs forwarded from the leakage encryption oracle. Thus the adversary can obtain decryption leakage, but only on ciphertexts for which it already knows the corresponding message. Below is the formal definition.

Definition 1 (IND-aCPLA Security). Let $\text{SE} = (\mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme and the INDaCPLA game be as defined in Fig. 2. Then for any nonce-respecting adversary \mathcal{A} that never forwards queries to or from the **Enc** oracle, only makes queries to **LDec** that are forwarded from **LEnc**, and only makes encryption and decryption queries containing leakage functions in the respective sets \mathcal{L}_E and \mathcal{L}_D , its corresponding IND-aCPLA advantage is given by:

$$\text{Adv}_{\text{SE}}^{\text{ind-acpla}}(\mathcal{A}, \mathcal{L}_E, \mathcal{L}_D) = 2 \Pr \left[\text{INDaCPLA}^{\mathcal{A}} \Rightarrow \text{true} \right] - 1.$$

⁴ This is not really required, since contrary to [4] the challenge oracles are not forgetful in our case. Nevertheless we conform to the original definition of forwarded queries.

<p>Game INDaCPLA</p> <hr/> <p>procedure Initialize</p> <p>$b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure Enc(N, M)</p> <p>$C \leftarrow \mathcal{E}(K, N, M)$</p> <p>if $b = 0$</p> <p> if $f[N, M] = \perp$</p> <p> $f[N, M] \leftarrow \{0, 1\}^{ C }$</p> <p> return $f[N, M]$</p> <p>else</p> <p> return C</p>	<hr/> <p>procedure LEnc(N, M, L)</p> <p>$A \leftarrow L(K, N, M)$</p> <p>$C \leftarrow \mathcal{E}(K, N, M)$</p> <p>return (C, A)</p> <hr/> <p>procedure LDec(N, C, L)</p> <p>$A \leftarrow L(K, N, C)$</p> <p>$M \leftarrow \mathcal{D}(K, N, C)$</p> <p>return (M, A)</p> <hr/> <p>procedure Finalize (b')</p> <p>return $(b' = b)$</p>
--	---

Fig. 2: Game used to define IND-aCPLA security.

Leakage-Resilient Function Families. We will distinguish among function families based on their domain \mathcal{X} . We will use the terms *fixed-input-length* function when $\mathcal{X} = \{0, 1\}^l$ for some $l \in \mathbb{N}$, *variable-input-length* function when $\mathcal{X} = \{0, 1\}^*$, and *vector* function when the domain is a cartesian product of string sets, e.g. $\mathcal{X} = \{0, 1\}^* \times \{0, 1\}^*$.

For such function families we will consider two security notions: leakage-resilient pseudorandom functions (LPRF) and leakage-resilient unpredictable functions (LUF). While LPRF security is well-established in the literature, LUF security is new. Below are the formal definitions. Note that no restriction is made on the adversary's queries in the LUF security definition.

<p>Game LPRF</p> <hr/> <p>procedure Initialize</p> <p>$b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure LF(X, L)</p> <p>$y \leftarrow \mathcal{F}(K, X)$</p> <p>$A \leftarrow L(K, X)$</p> <p>return (y, A)</p>	<hr/> <p>procedure F(X)</p> <p>if $b = 0$</p> <p> if $f[X] = \perp$</p> <p> $f[X] \leftarrow \{0, 1\}^t$</p> <p> return $f[X]$</p> <p>else</p> <p> return $\mathcal{F}(K, X)$</p> <hr/> <p>procedure Finalize (b')</p> <p>return $(b' = b)$</p>
---	---

Fig. 3: Game used to define LPRF security.

Definition 2 (LPRF Security). Let $\mathcal{F}: \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^t$ be a function family over the domain \mathcal{X} and indexed by \mathcal{K} , and the LPRF game be as defined in Fig. 3. Then for any adversary \mathcal{A} that never forwards queries to or from the \mathbf{F} oracle and only queries leakage functions in the set \mathcal{L}_F , its corresponding LPRF advantage is given by:

$$\mathbf{Adv}_{\mathcal{F}}^{\text{lprf}}(\mathcal{A}, \mathcal{L}_F) = 2 \Pr \left[\text{LPRF}^{\mathcal{A}} \Rightarrow \text{true} \right] - 1.$$

Definition 3 (LUF Security). Let $\mathcal{F}: \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^t$ be a function family over the domain \mathcal{X} and indexed by \mathcal{K} , and the LUF game be as defined in Fig. 4. Then for any adversary \mathcal{A} its corresponding LUF advantage is given by:

$$\mathbf{Adv}_{\mathcal{F}}^{\text{luf}}(\mathcal{A}, \mathcal{L}_F) = \Pr \left[\text{LUF}^{\mathcal{A}} \Rightarrow \text{true} \right].$$

Game LUF	procedure Lkg(X, L)
procedure Initialize	$\Lambda \leftarrow L(K, X)$ return Λ
$\text{win} \leftarrow \text{false}; K \leftarrow \mathcal{K}$ return	procedure Guess(X, y')
procedure F(X)	$y \leftarrow \mathcal{F}(K, X)$ if $X \notin \mathcal{S} \wedge y = y'$ $\text{win} \leftarrow \text{true}$ return ($y = y'$)
$\mathcal{S} \leftarrow_{\cup} X$ $y \leftarrow \mathcal{F}(K, X)$ return y	procedure Finalize
	return (win)

Fig. 4: Game used to define LUF security.

Unforgeability in the Presence of Leakage. For message authentication we will require the analogue of strong unforgeability in the leakage setting (SUF-CMLA) put forth in [4]. This is essentially strong unforgeability (SUF-CMA) formulated as a distinguishing game, with a challenge verification oracle and additional tagging and verification oracles that leak. Below is the formal definition.

Game SUFCMLA	procedure LTag(X, L)
procedure Initialize	$\Lambda \leftarrow L(K, X)$ $T \leftarrow \mathcal{T}(K, X)$ return (T, Λ)
$b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$ return	procedure LVfy(X, T, L)
procedure Vfy(X, T)	$\Lambda \leftarrow L(K, X, T)$ $v \leftarrow \mathcal{V}(K, X, T)$ return (v, Λ)
if $b = 0$ return \perp else $v \leftarrow \mathcal{V}(K, X, T)$ return v	procedure Finalize (b')
	return ($b' = b$)

Fig. 5: Game used to define SUF-CMLA security.

Definition 4 (SUF-CMLA Security). Let $\text{MAC} = (\mathcal{T}, \mathcal{V})$ be a message authentication code and the SUFCMLA game be as defined in Fig. 5. For any adversary \mathcal{A} that never forwards queries from LTag to Vfy, and only queries leakage functions to its tagging and verification oracles in the respective sets \mathcal{L}_T and \mathcal{L}_V , its corresponding SUF-CMLA advantage is given by:

$$\text{Adv}_{\text{MAC}}^{\text{suf-cmla}}(\mathcal{A}, \mathcal{L}_T, \mathcal{L}_V) = 2 \Pr \left[\text{SUFCMLA}^{\mathcal{A}} \Rightarrow \text{true} \right] - 1.$$

Authenticated Encryption with Leakage. For an authenticated encryption scheme with associated data our target will be LAE security, which is a natural extension of the classical security notion put forth by Rogaway [24] to the leakage setting. This is defined formally below.

Definition 5 (LAE Security). Let $\text{AEAD} = (\mathcal{E}, \mathcal{D})$ be an authenticated encryption scheme with associated data and the LAE game be as defined in Fig. 6. Then for any adversary \mathcal{A} that never forwards queries to or from the Enc and Dec oracles and only makes encryption and decryption queries containing

<p>Game LAE</p> <hr/> <p>procedure Initialize</p> <p>$b \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure Enc(N, A, M)</p> <p>$C \leftarrow \mathcal{E}(K, N, A, M)$</p> <p>if $b = 0$</p> <p> if $f[N, A, M] = \perp$</p> <p> $f[N, A, M] \leftarrow \{0, 1\}^{ C }$</p> <p> return $f[N, A, M]$</p> <p>else</p> <p> return C</p> <hr/> <p>procedure Finalize (b')</p> <p>return ($b' = b$)</p>	<p>procedure Dec(N, A, C)</p> <hr/> <p>$M \leftarrow \mathcal{D}(K, N, A, C)$</p> <p>if $b = 0$</p> <p> return \perp</p> <p>else</p> <p> return M</p> <hr/> <p>procedure LEnc(N, A, M, L)</p> <hr/> <p>$A \leftarrow L(K, N, A, M)$</p> <p>$C \leftarrow \mathcal{E}(K, N, A, M)$</p> <p>return ($C, A$)</p> <hr/> <p>procedure LDec(N, A, C, L)</p> <hr/> <p>$A \leftarrow L(K, N, A, C)$</p> <p>$M \leftarrow \mathcal{D}(K, N, A, C)$</p> <p>return ($M, A$)</p>
--	---

Fig. 6: Game used to define LAE security.

leakage functions in the respective sets \mathcal{L}_{AE} and \mathcal{L}_{VD} , its corresponding LAE advantage is given by:

$$\text{Adv}_{\text{AEAD}}^{\text{lAE}}(\mathcal{A}, \mathcal{L}_{AE}, \mathcal{L}_{VD}) = 2 \Pr[\text{LAE}^{\mathcal{A}} \Rightarrow \text{true}] - 1.$$

Generic Composition in the Leakage Setting. The N2 construction was introduced in [22] and is depicted pictorially in Fig. 7. In [4] Barwell et al. prove a composition theorem for this construction that holds in the leakage setting. We will make use of this theorem and for completeness we reproduce it below, adapted to the random transformation model.

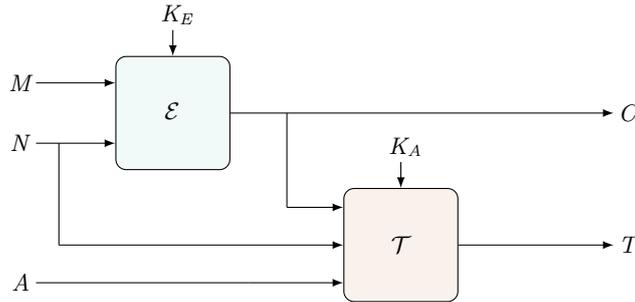


Fig. 7: Graphical representation of the N2 construction.

Theorem 1 (LAE Security of the N2 Construction [4]). Let $\text{SE} = (\mathcal{E}, \mathcal{D})$ be a symmetric encryption scheme with associated leakage sets $(\mathcal{L}_E, \mathcal{L}_D)$ and $\text{MAC} = (\mathcal{T}, \mathcal{V})$ be a MAC with associated leakage sets $(\mathcal{L}_T, \mathcal{L}_V)$. Further let N2 be the composition of SE and MAC described in Fig. 7, with associated leakage sets $(\mathcal{L}_{AE}, \mathcal{L}_{VD})$ where $\mathcal{L}_{AE} = \mathcal{L}_E \times \mathcal{L}_T$ and $\mathcal{L}_{VD} = \mathcal{L}_D \times \mathcal{L}_V$. Then for any LAE adversary \mathcal{A}_{ae} against N2 there exist adversaries \mathcal{A}_{se} , \mathcal{A}_{prf} , and \mathcal{A}_{mac} such that:

$$\begin{aligned} \text{Adv}_{\text{N2}}^{\text{lAE}}(\mathcal{A}_{ae}, \mathcal{L}_{AE}, \mathcal{L}_{VD}) &\leq \text{Adv}_{\text{SE}}^{\text{ind-acpla}}(\mathcal{A}_{se}, \mathcal{L}_E, \mathcal{L}_D) \\ &\quad + \text{Adv}_{\mathcal{T}}^{\text{lprf}}(\mathcal{A}_{prf}, \mathcal{L}_T) + 2\text{Adv}_{\text{MAC}}^{\text{suf-cmla}}(\mathcal{A}_{mac}, \mathcal{L}_T, \mathcal{L}_V). \end{aligned}$$

The resources of the above adversaries are quantified as follows. Let q and μ be such that \mathcal{A}_{ae} makes at most q queries to any of its oracles **Enc**, **Dec**, **LEnc**, and **LDec**, up to a total of μ bits for each oracle. Further let q_ρ denote the number of queries that \mathcal{A}_{ae} makes to ρ . Then \mathcal{A}_{se} queries each of **Enc**, **LEnc**, and **LDec** at most q times up to a total of μ bits, and queries ρ at most $Q_{\mathcal{T}}(2q, 2\mu) + Q_{\mathcal{V}}(q, \mu) + q_\rho$ times. As for \mathcal{A}_{lprf} , it makes at most q and $2q$ queries to **F** and **LF**, totalling μ and 2μ bits respectively, and $Q_{\mathcal{E}}(q, \mu) + Q_{\mathcal{D}}(q, \mu) + q_\rho$ queries to ρ . Finally, \mathcal{A}_{mac} queries each its oracles **Vfy**, **LTag**, and **LVfy** at most $2q$ times, totalling 2μ bits in each case, and queries its ρ oracle $Q_{\mathcal{E}}(2q, 2\mu) + Q_{\mathcal{D}}(q, \mu) + q_\rho$ times.

3 SLAE: A Sponge-Based LAE Construction

SLAE, pronounced “sleigh”, is a **S**ponge-based non-adaptive **L**eakage-resilient **A**EAD scheme. It is based on, and is closely related to, a prior sponge-based AEAD scheme called ISAP [12]. ISAP is a nonce-based AEAD scheme intended to inherently resist side-channel attacks while simultaneously fitting the well-established syntax of AEAD schemes. More specifically, it claims security against Differential Power Analysis (DPA) by employing a rekeying mechanism. An important challenge that ISAP overcomes, is to avoid decrypting distinct ciphertexts under the same key without maintaining a state. Furthermore, as noted by ISAP’s designers, the sponge construction seems markedly well-suited to protect against side-channels. Typically, the sponge employs a large state that is continually evolving, which intuitively endows it with an intrinsic resilience to information leakage. Thus, in contrast to other designs, ISAP potentially offers a fairly efficient LAE solution that can be instantiated with off-the-shelf primitives. However, as we already noted, ISAP’s biggest limitation is that its design is not backed by any formal security analysis, not even in the absence of leakage.

ISAP is composed of a symmetric encryption scheme ISAPENC and a MAC ISAPMAC combined according to the N2 construction. These components were conceived by augmenting established sponge constructs with a rekeying function. In particular the design rationale behind ISAPMAC is to augment a sponge-based suffix MAC with a rekeying function. The rekeying is such that the key fed into the suffix MAC itself depends on the inputs being authenticated and a master authentication key. Similarly ISAPENC is a standard sponge-based encryption scheme whose key is derived from a master encryption key and the nonce. Throughout, the rekeying function is realised from the sponge by setting the absorption rate to be one. Intuitively, ISAP’s resistance to DPA comes from the fact that encryption and authentication never use the same key more than once, and the slow absorption rate employed in the rekeying function. Both of these factors limit the so-called data complexity of computations involving secret values, which in turn encumbers DPA attacks. See [12] for more details on ISAP.

SLAE retains the main structure of ISAP, as well as its benefits, but it includes some changes and restrictions that facilitate its security analysis. While the majority of these differences are conceptual, they are substantial enough, however, to *invalidate* any claim that our security proof applies to ISAP. The design of SLAE can be understood across *three* different levels of abstraction. At the highest level, like ISAP, it is the N2 composition of a symmetric encryption scheme SLENC and a MAC SLMAC. At the second abstraction level, SLMAC and SLENC can be viewed in terms of smaller components. Specifically, we view SLMAC as combining a collision-resistant vector hash function \mathcal{H} and a fixed-input-length function \mathcal{F}' , and we decompose SLENC into a fixed-input-length function \mathcal{F} and a pseudorandom generator with variable output length \mathcal{G} . Indeed this view corresponds to our generic construction of a non-adaptively leakage-resilient AEAD scheme which we refer to as the FGHF’ construction.

Note that there is no explicit idea of rekeying in the FGHF’ construction. The only leakage-resilient primitives are \mathcal{F} and \mathcal{F}' . For security we will require both to be LPRF secure and \mathcal{F}' to additionally be LUF secure. Thus LAE schemes are easy to construct once we have such primitives. Moreover, \mathcal{F} is invoked once for encryption, and likewise \mathcal{F}' is invoked once for authentication, irrespective of the message length. SLAE is obtained by instantiating the four components in the FGHF’ construction with T-sponges. This is the third level view. While the design rationale behind the FGHF’ construction is quite distinct from that of ISAP, once instantiated, SLAE and ISAP suddenly look very similar.

We now describe SLAE in more detail and then elaborate on the differences between SLAE and ISAP in Section 3.4.

3.1 High-Level View of SLAE

As already noted, $\text{SLAE} = (\text{SLAE-}\mathcal{E}, \text{SLAE-}\mathcal{D})$ is a nonce-based AEAD scheme composed from a nonce-based symmetric encryption scheme $\text{SLENC} = (\text{SLENC-}\mathcal{E}, \text{SLENC-}\mathcal{D})$ and a MAC $\text{SLMAC} = (\text{SLMAC-}\mathcal{T}, \text{SLMAC-}\mathcal{V})$. These are combined according to the N2 composition, where the key is split into an encryption

key K_E and an authentication key K_A . During encryption, $\text{SLENC-}\mathcal{E}$ takes the nonce, message, and key K_E to return a ciphertext which is then fed together with the nonce, associated data, and key K_A , to $\text{SLMAC-}\mathcal{T}$ to produce a tag which is then appended onto the ciphertext. Decryption proceeds by reversing these operations in a *verify-then-decrypt* manner, whereby ciphertext decryption using $\text{SLENC-}\mathcal{D}$ proceeds *only if* tag verification under $\text{SLMAC-}\mathcal{V}$ was successful. The pseudocode for this composition is described in Fig. 8.

$\text{SLAE-}\mathcal{E}(K, N, A, M)$	$\text{SLAE-}\mathcal{D}(K, N, A, \bar{C})$
<p>parse K as $K_E \parallel K_A$ $C \leftarrow \text{SLENC-}\mathcal{E}(K_E, N, M)$ $T \leftarrow \text{SLMAC-}\mathcal{T}(K_A, (N, A, C))$ $\bar{C} \leftarrow C \parallel T$ return \bar{C}</p>	<p>parse K as $K_E \parallel K_A$ parse \bar{C} as $C \parallel T$ $v \leftarrow \text{SLMAC-}\mathcal{V}(K_A, (N, A, C), T)$ if $v = \top$ $M \leftarrow \text{SLENC-}\mathcal{D}(K_E, N, C)$ return M else return \perp</p>

Fig. 8: High-level description of SLAE in terms of SLMAC and SLENC.

3.2 The SLMAC Construction

A pseudocode description of $\text{SLMAC} = (\text{SLMAC-}\mathcal{T}, \text{SLMAC-}\mathcal{V})$ can be found in Fig. 9. It is a vector MAC operating on the triple (N, A, C) , where verification works by recomputing the tag for the given triple and checking that it is identical to the given tag. As such, the core functionality of SLMAC is captured in the tagging algorithm $\text{SLMAC-}\mathcal{T}$, which is additionally depicted in Fig. 10. The tagging algorithm can be understood as being composed of a (sponge-based) vector hash function compressing the triple (N, A, C) into a digest of size w bits, which is then fed to the unpredictable function SLFUNC to produce a tag of size t bits. The nonce N is required to be m bits long, whereas A and C can be of arbitrary length. Accordingly, $\text{SLMAC-}\mathcal{T}$ starts by padding both A and C so that their lengths are integer multiples of the sponge rate r . Note that the padding function, lpad , always returns at least a single bit of padding and is always applied, even if the input string is already an integer multiple of r .

To compute the hash digest H , the internal state is initialised to $\rho(N \parallel IV)$, where IV is a constant string of size $n - m$, and the padded associated data \mathbf{A} and the padded ciphertext \mathbf{C} are then absorbed block by block. An input separation mechanism is employed in order to demarcate the boundary between \mathbf{A} and \mathbf{C} . This involves XORing the string $1 \parallel 0^{c-1}$ to the inner part of the state once \mathbf{A} has been absorbed, and ensures that distinct pairs $(\mathbf{A}, \mathbf{C}) \neq (\bar{\mathbf{A}}, \bar{\mathbf{C}})$ for which $\mathbf{A} \parallel \mathbf{C} = \bar{\mathbf{A}} \parallel \bar{\mathbf{C}}$ do not result in the same hash digest.

Once the hash digest is evaluated, it is fed into SLFUNC to compute the final tag. This is also a sponge-based construction for which a graphical representation appears in Fig. 11. Here the state is initialised to $\rho(K_A \parallel IV)$ and the hash digest is then absorbed at a *reduced rate* of rr bits. Once the complete digest has been absorbed the left most t bits of the state are output as the tag.

3.3 The SLENC Construction

This is the sponge-based symmetric encryption scheme $\text{SLENC} = (\text{SLENC-}\mathcal{E}, \text{SLENC-}\mathcal{D})$ described in Fig. 12 and depicted in Fig. 13. It is easy to see that $\text{SLENC-}\mathcal{D}(K_E, N, \cdot) = \text{SLENC-}\mathcal{E}(K_E, N, \cdot)$, and consequently we only describe the operation of $\text{SLENC-}\mathcal{E}$. This algorithm can be viewed as being composed of a pseudorandom function SLFUNC , taking as input the pair (K_E, N) , and whose output is then fed into a pseudorandom generator SPRG . The output of SPRG is then used to encrypt the message.

The nonce N is required to be m bits long and we do not require any additional padding for the message. The evaluation of SLFUNC proceeds by initialising the internal state to $\rho(K_E \parallel IV)$, with a constant IV of size $n - k$, and then absorbing the nonce at a reduced rate of rr bits. Once the nonce is absorbed, the output state S_{l+1} serves as the seed to the pseudorandom generator SPRG . A separate

<p>SLMAC-$\mathcal{T}(K_A, (N, A, C))$</p> <hr/> <p>$A \leftarrow A \parallel \text{lpad}(A, r)$ parse A as $A_1 \parallel \dots \parallel A_u$ st $\forall i A_i = r$ $C \leftarrow C \parallel \text{lpad}(C, r)$ parse C as $C_1 \parallel \dots \parallel C_v$ st $\forall i C_i = r$</p> <p>$Y_0 \leftarrow N \parallel IV$ $S_1 \leftarrow \rho(Y_0)$ // Absorb Associated Data for i in $\{1, \dots, u\}$ $Y_i \leftarrow (\bar{S}_i \oplus A_i) \parallel \hat{S}_i$ $S_{i+1} \leftarrow \rho(Y_i)$ // Separate Inputs $S_{u+1} \leftarrow \bar{S}_{u+1} \parallel (\hat{S}_{u+1} \oplus (1 \parallel 0^{c-1}))$ // Absorb Ciphertext for i in $\{u+1, \dots, u+v\}$ $Y_i \leftarrow (\bar{S}_i \oplus C_{i-u}) \parallel \hat{S}_i$ $S_{i+1} \leftarrow \rho(Y_i)$ // Generate Tag $H \leftarrow \lfloor S_{u+v+1} \rfloor_w$ $T \leftarrow \lfloor \text{SLFUNC}(K_A, H) \rfloor_t$ return T</p>	<p>SLMAC-$\mathcal{V}(K_A, (N, A, C), T)$</p> <hr/> <p>$T' \leftarrow \text{SLMAC-}\mathcal{T}(K_A, (N, A, C))$ if $T = T'$ return \top return \perp</p> <p>SLFUNC(K_A, H)</p> <hr/> <p>parse H as $H_1 \parallel \dots \parallel H_l$ st $\forall i H_i = rr$</p> <p>$Y_0 \leftarrow K_A \parallel IV$ $S_1 \leftarrow \rho(Y_0)$ for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus H_i) \parallel \hat{S}_i$ $S_{i+1} \leftarrow \rho(Y_i)$ return $\lfloor S_{l+1} \rfloor_t$</p> <p>lpad($A, r$)</p> <hr/> <p>$x \leftarrow A \bmod r$ return $1 \parallel 0^{r-x-1}$</p>
---	--

Fig. 9: Pseudocode description of SLMAC and SLFUNC.

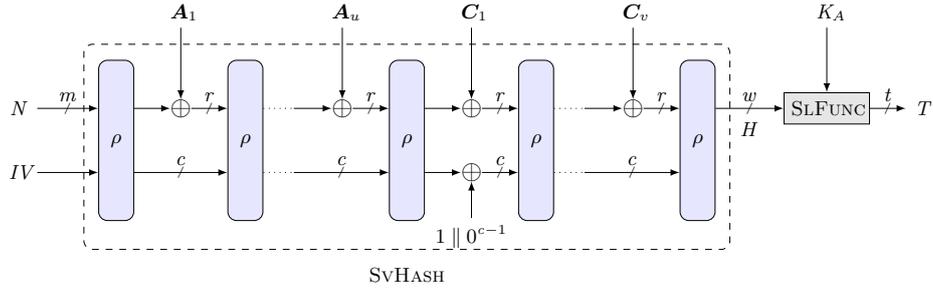


Fig. 10: Graphical illustration of SLMAC- \mathcal{T} .

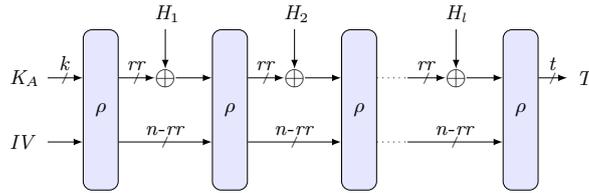


Fig. 11: Graphical illustration of SLFUNC.

pseudocode description of SPRG can be found in Fig. 14. The first ciphertext block is generated by XORing the outer part of this state with the first message block. Afterwards the initial state is given as input to the random transformation outputting a new state which is then used to derive the next ciphertext block by simply XORing again the outer state with the next message block. This process is repeated until the whole message has been processed. If the last message block is smaller than r bits, we simply truncate the outer state to the required size and XOR both parts to obtain the last ciphertext block.

$\text{SLENC-}\mathcal{E}(K_E, N, M)$	$\text{SLENC-}\mathcal{D}(K_E, N, C)$
<p>parse N as $N_1 \parallel \dots \parallel N_l$ st $\forall i \ N_i = rr$ parse M as $M_1 \parallel \dots \parallel M_v$ st $\forall i < v \ M_i = r$ and $M_v \leq r$</p> <p>// First Sponge Iteration $Y_0 \leftarrow K_E \parallel IV$ $S_1 \leftarrow \rho(Y_0)$</p> <p>// Absorb Nonce for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus N_i) \parallel \hat{S}_i$ $S_{i+1} \leftarrow \rho(Y_i)$</p> <p>// Squeeze and Encrypt for i in $\{l+1, \dots, l+v-1\}$ $C_{i-l} \leftarrow \bar{S}_i \oplus M_{i-l}$ $S_{i+1} \leftarrow \rho(S_i)$ $C_v \leftarrow \lfloor \bar{S}_{l+v} \rfloor_{ M_v } \oplus M_v$ return $C_1 \parallel \dots \parallel C_v$</p>	<p>parse N as $N_1 \parallel \dots \parallel N_l$ st $\forall i \ N_i = rr$ parse C as $C_1 \parallel \dots \parallel C_v$ st $\forall i < v \ C_i = r$ and $C_v \leq r$</p> <p>// First Sponge Iteration $Y_0 \leftarrow K_E \parallel IV$ $S_1 \leftarrow \rho(Y_0)$</p> <p>// Absorb Nonce for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus N_i) \parallel \hat{S}_i$ $S_{i+1} \leftarrow \rho(Y_i)$</p> <p>// Squeeze and Decrypt for i in $\{l+1, \dots, l+v-1\}$ $M_{i-l} \leftarrow \bar{S}_i \oplus C_{i-l}$ $S_{i+1} \leftarrow \rho(S_i)$ $M_v \leftarrow \lfloor \bar{S}_{l+v} \rfloor_{ C_v } \oplus C_v$ return $M_1 \parallel \dots \parallel M_v$</p>

Fig. 12: Pseudocode description of the SLENC encryption scheme.

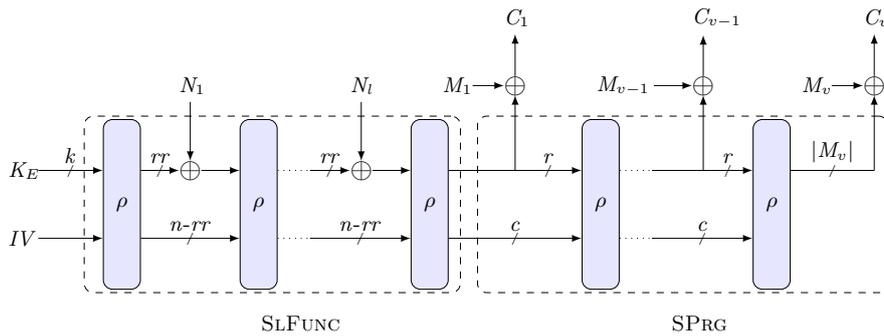


Fig. 13: Graphical illustration of SLENC- \mathcal{E} .

3.4 Differences Between SLAE and ISAP

We have already described in passing some of the differences between SLAE and ISAP, but for clarity, we summarise these distinctions below and discuss them in more detail.

The most prominent difference is that SLAE is based on the T-sponge whereas ISAP employs the P-sponge. In particular the security proofs of SLAE rely on treating ρ as a non-invertible transformation.

<pre style="margin: 0;"> SPRG(<i>seed</i>, <i>L</i>) ----- <i>v</i> ← ⌈$\frac{L}{r}$⌉ <i>S</i>₁ ← <i>seed</i> for <i>i</i> in {1, ..., <i>v</i> - 1} <i>S</i>_{<i>i</i>+1} ← ρ(<i>S</i>_{<i>i</i>}) <i>R</i> ← $\bar{S}_1 \parallel \dots \parallel \bar{S}_v$ return $\lfloor R \rfloor_L$ </pre>

Fig. 14: Pseudocode description of SPRG.

Treating ρ as an invertible random permutation would add another layer of complexity to the security analysis and we chose not to pursue this route at this point.

The description of ISAP actually specifies three distinct permutations, each obtained from the same round function but with a varying number of rounds. These are used in the different components of ISAP as a means of optimisation. Such heuristic optimisations could be employed in SLAE as well, but in our security analysis we instantiate SLAE with the same random transformation throughout. Indeed this is the more conservative assumption, since otherwise we would be treating these variants as being sampled independently at random when in fact they are intimately related.

Another difference between SLAE and ISAP can be seen in their MAC components SLMAC and ISAPMAC. The design of ISAPMAC is based on combining a rekeying function ISAPRK with a sponge-based suffix MAC. In turn, ISAPRK takes as input a hash of the MAC inputs. As a design optimisation, it is then noted that this hash is already being computed as part of the suffix MAC, at which point it is extracted, fed into ISAPRK, and its output (the session key) is fed back into the last permutation of the suffix MAC to yield the MAC tag. In contrast, in SLMAC, the value corresponding to the session key in ISAPMAC is output directly as the MAC tag.

Finally there are some differences in the way we set parameters in SLAE as opposed to ISAP. For instance, ISAP sets the size of the key and the nonce to be equal. On the other hand, our analysis indicates that the limiting factor in the security of SLAE is the key size. As such it makes sense to set the key size k equal to the width of the sponge n while setting the nonce to be much smaller, say between 64 and 128 bits.

4 The Security of FGHF'

In this section we establish the security of the FGHF' construction which is depicted in Fig. 15. This is an abstraction of SLAE, and proving its security brings us halfway towards proving the security of SLAE. At the same time, we believe the FGHF' construction to be of independent interest as it serves as a generic blueprint for constructing efficient AEAD schemes that are non-adaptively leakage-resilient.

The FGHF' construction is a refinement of the N2 construction [22] which builds a nonce-based AEAD scheme from a nonce-based symmetric encryption scheme and a vector MAC. Barwell et al. [4] showed that the security of this construction extends to the setting of leakage resilience. Specifically they showed that if the encryption component is IND-aCPLA secure and the vector MAC is both LPRF and SUF-CMLA secure, then the composition is LAE secure. In turn the FGHF' construction further breaks down the encryption component, denoted by $\text{SE}[\mathcal{F}, \mathcal{G}]$, and the vector MAC component, denoted by $\text{MAC}[\mathcal{H}, \mathcal{F}']$, of N2 into smaller parts. Namely encryption is realised from a fixed-input-length leakage-resilient PRF \mathcal{F} and a standard PRG \mathcal{G} , whereas the vector MAC is built from a vector hash function \mathcal{H} , and a fixed-input-length function \mathcal{F}' that is both leakage-resilient pseudorandom and leakage-resilient unpredictable.

Since FGHF' is an instance of N2 we can apply the composition theorem of Barwell et al. [4], which we reproduced and adapted to the random transformation model in Section 2.5. Moreover, since we can view non-adaptive leakage as a special case of adaptive leakage where the leakage set is a singleton, the theorem carries over to that setting which is what we are interested in here. Thus to prove that the FGHF' construction is LAE secure we only need to show that the encryption and MAC components meet the requirements of Theorem 1.

As it turns out, we can realise an IND-aCPLA secure encryption directly from an LPRF and a variable-output-length PRG. Here the PRG serves only to extend the range of the LPRF in order for

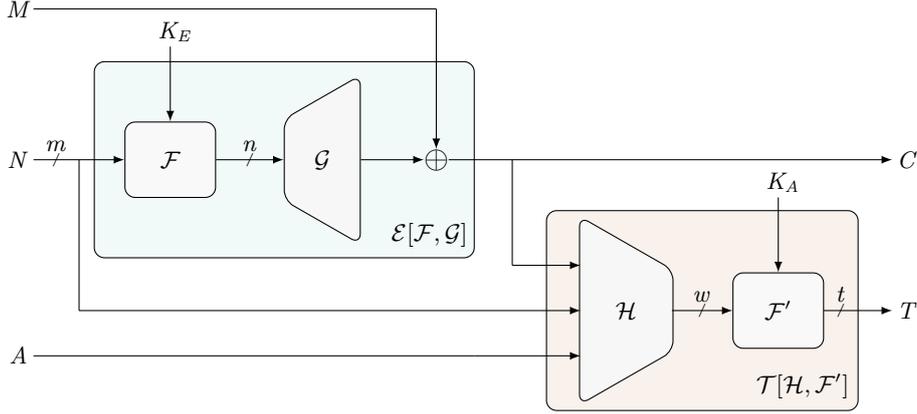


Fig. 15: Graphical representation of the FGHF' construction. It corresponds to the N2 composition of $\text{SE}[\mathcal{F}, \mathcal{G}] = (\mathcal{E}, \mathcal{D})$ and $\text{MAC}[\mathcal{H}, \mathcal{F}'] = (\mathcal{T}, \mathcal{V})$ which are in turn composed of a fixed-input-length LPRF \mathcal{F} , a PRG \mathcal{G} , a vector hash \mathcal{H} , and a fixed-input-length function \mathcal{F}' that is both a LUF and an LPRF.

the encryption scheme to accommodate variable-length messages. Surprisingly, a standard PRG without any leakage resilience suffices. As for the vector MAC component it needs to be an LPRF over a vector of strings and simultaneously satisfy SUF-CMLA security. Contrary to the leakage-free setting, the latter property is not automatically implied by the former when a MAC is constructed from an LPRF through the canonical construction. This is because the SUFCMLA game is more permissive than the LPRF game with respect to the adversary's queries. Namely, the adversary can forward queries from the LVfy to Vfy, whereas in the LPRF game the adversary is not allowed to forward queries from LF to F. This precludes reducing SUF-CMLA security to LPRF security due to our inability of simulating the verification oracles via the respective LPRF oracles. Note that SUF-CMLA needs to be defined this way for Theorem 1 to hold whereas lifting the restriction in the LPRF game would make it unsatisfiable. We overcome this problem by noting that, in the non-adaptive leakage setting, unpredictability suffices to achieve SUF-CMLA security, and at the same time we can allow the adversary to forward queries between its leakage and challenge oracles while maintaining satisfiability. This leads to our notion of a LUF which we prove to be sufficient to yield SUF-CMLA security. As we will see in the next section we can construct fixed-input-length function families satisfying both notions rather easily from sponges. Given such a function family \mathcal{F}' , we can turn it into the required vector MAC by composing it with a collision-resistant vector hash function. Specifically we show that we can extend the domain of LPRFs and LUFs, rather efficiently, by composing them with standard collision-resistant hash functions over appropriate domains.

Combining the results in this section, leads to the LAE security of the FGHF' construction against non-adaptive leakage. We like this construction as it strikes a practical balance between security and efficiency. By settling for non-adaptive leakage, which seems to suffice for many practical applications, it only requires one call to each of the leakage-resilient primitives, \mathcal{F} and \mathcal{F}' , per encryption query. In this work we focused on SLAE which is a specific sponge-based instantiation of FGHF' , but other instantiations, possibly based on different techniques, are of course possible. Thus this construction essentially reduces the problem of designing non-adaptively leakage-resilient AEAD schemes to that of designing function families over small domains that are good LPRFs and LUFs, which conceptually is a much simpler target.

4.1 $\text{SE}[\mathcal{F}, \mathcal{G}]$ is IND-aCPLA Secure

We begin by proving the security of the encryption component of FGHF' . Note that for this part security holds in more general setting of adaptive leakage. Below is the formal theorem statement and its proof is presented in Appendix B.1.

Theorem 2. *Let $\text{SE}[\mathcal{F}, \mathcal{G}]$ be the encryption scheme depicted in Fig. 15, composed of the function family $\mathcal{F}: \mathcal{K} \times \{0, 1\}^m \rightarrow \{0, 1\}^n$ and the PRG $\mathcal{G}: \{0, 1\}^n \rightarrow \{0, 1\}^*$ with respective associated leakage sets $\mathcal{L}_{\mathcal{F}}$ and $\mathcal{L}_{\mathcal{G}}$. Then for any IND-aCPLA adversary \mathcal{A}_{se} against $\text{SE}[\mathcal{F}, \mathcal{G}]$ and associated leakage sets $\mathcal{L}_{\mathcal{E}} = \mathcal{L}_{\mathcal{D}} = \mathcal{L}_{\mathcal{F}} \times \mathcal{L}_{\mathcal{G}}$, there exist an LPRF adversary \mathcal{A}_{lprf} against \mathcal{F} and a PRG adversary \mathcal{A}_{prg}*

against \mathcal{G} such that:

$$\mathbf{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{ind-acpla}}(\mathcal{A}_{se}, \mathcal{L}_E, \mathcal{L}_D) \leq 2 \mathbf{Adv}_{\mathcal{F}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_F) + 2 \mathbf{Adv}_{\mathcal{G}}^{\text{prg}}(\mathcal{A}_{prg}) .$$

Let q and μ be such that \mathcal{A}_{se} makes at most q queries totalling μ bits to each of its oracles **Enc**, **LEnc**, and **LDec**, and let q_ρ denote the number of queries it makes to ρ . Then \mathcal{A}_{lprf} makes at most q and $2q$ queries to its oracles **F** and **LF**, totalling qm and $2qm$, respectively, and at most $Q_{\mathcal{G}}(2q, 2\mu) + q_\rho$ to ρ . As for \mathcal{A}_{prg} , it makes at most q queries to its oracle **G** totalling μ bits and $Q_{\mathcal{F}}(2q, 2qm) + Q_{\mathcal{G}}(2q, 2\mu) + q_\rho$ queries to ρ .

4.2 MAC $[\mathcal{H}, \mathcal{F}']$ is SUF-CMLA Secure

Next we reduce the SUF-CMLA security of $\text{MAC}[\mathcal{H}, \mathcal{F}']$ to the LUF security of \mathcal{F}' and the collision resistance of \mathcal{H} . Towards this end, we first show that any LUF $\hat{\mathcal{F}}$ over domain \mathcal{X} yields a SUF-CMLA secure MAC with message space \mathcal{X} via the canonical construction. Then we show that such a function $\hat{\mathcal{F}}$ can be constructed from a fixed-input-length LUF \mathcal{F}' and a collision-resistant hash function with domain \mathcal{X} . The formal theorem statements now follow. Their proofs can be found in Appendices B.2 and B.3 respectively.

Theorem 3. *Let $\hat{\mathcal{F}}: \mathcal{K} \times \mathcal{X} \rightarrow \{0, 1\}^t$ be a function family with associated leakage set $\mathcal{L}_{\hat{\mathcal{F}}}$, and let $\text{MAC}[\hat{\mathcal{F}}]$ be the corresponding canonical MAC with associated leakage sets $\mathcal{L}_T, \mathcal{L}_V$ where $\mathcal{L}_{\hat{\mathcal{F}}} = \mathcal{L}_T = \mathcal{L}_V$. Then for any SUF-CMLA adversary \mathcal{A}_{mac} against $\text{MAC}[\hat{\mathcal{F}}]$, there exists an adversary \mathcal{A}_{luf} against $\hat{\mathcal{F}}$ such that:*

$$\mathbf{Adv}_{\text{MAC}[\hat{\mathcal{F}}]}^{\text{suf-cmla}}(\mathcal{A}_{mac}, \mathcal{L}_T, \mathcal{L}_V) \leq \mathbf{Adv}_{\hat{\mathcal{F}}}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\hat{\mathcal{F}}}) .$$

Let q and μ be such that \mathcal{A}_{mac} makes at most q queries totalling μ bits to each of its oracles **Vfy**, **LTag**, and **LVfy**. Then \mathcal{A}_{luf} makes at most q , $2q$, and $2q$ queries to **F**, **Lkg**, and **Guess**, totalling μ , 2μ , and 2μ bits, respectively.

Theorem 4. *Let $\mathcal{F}': \mathcal{K} \times \{0, 1\}^w \rightarrow \{0, 1\}^t$ be a function family with associated leakage set $\mathcal{L}_{\mathcal{F}'}$, and let $\mathcal{H}: \mathcal{X} \rightarrow \{0, 1\}^w$ be a hash function over any domain \mathcal{X} . Further let their composition $\hat{\mathcal{F}}$ be defined as*

$$\hat{\mathcal{F}}(K, X) = \mathcal{F}'(K, \mathcal{H}(X))$$

where $X \in \mathcal{X}$, $K \in \mathcal{K}$, and $\mathcal{L}_{\hat{\mathcal{F}}} = \mathcal{L}_{\mathcal{F}'} \times \mathcal{L}_{\mathcal{H}}$ for any set of efficiently computable functions $\mathcal{L}_{\mathcal{H}}$. Then for any LUF adversary \mathcal{A}_{luf} against $\hat{\mathcal{F}}$, there exists a corresponding LUF adversary \mathcal{A}'_{luf} against \mathcal{F}' and an adversary \mathcal{A}_{hash} against \mathcal{H} such that:

$$\mathbf{Adv}_{\hat{\mathcal{F}}}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\hat{\mathcal{F}}}) \leq 2 \mathbf{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}) + \mathbf{Adv}_{\mathcal{F}'}^{\text{luf}}(\mathcal{A}'_{luf}, \mathcal{L}_{\mathcal{F}'}) .$$

Let q and μ be such that \mathcal{A}_{luf} makes at most q queries totalling μ bits to each of its oracles **F**, **Lkg**, and **Guess**, and let q_ρ denote the number of queries it makes to ρ . Then \mathcal{A}'_{luf} makes at most q queries totalling qw bits to each of its oracles **F**, **Lkg**, and **Guess**, and at most $Q_{\mathcal{H}}(3q, 3\mu) + q_\rho$ queries to ρ . As for \mathcal{A}_{hash} , it requires at most $Q_{\mathcal{F}'}(3q, 3qw) + Q_{\mathcal{H}}(3q, 3\mu)$ queries to ρ in order to simulate \mathcal{F}' and \mathcal{H} .

Combining both theorems, we obtain the following simple corollary reducing the SUF-CMLA security of $\text{MAC}[\mathcal{H}, \mathcal{F}']$ to that of its building blocks \mathcal{H} and \mathcal{F}' .

Corollary 1. *Let $\text{MAC}[\mathcal{H}, \mathcal{F}']$ be the MAC component depicted in Fig. 15, composed of the hash function \mathcal{H} and the function family \mathcal{F}' with respective leakage sets $\mathcal{L}_{\mathcal{H}}$ and $\mathcal{L}_{\mathcal{F}'}$. Then for any SUF-CMLA adversary \mathcal{A}_{mac} against $\text{MAC}[\mathcal{H}, \mathcal{F}']$ with associated leakage sets $\mathcal{L}_T = \mathcal{L}_V = \mathcal{L}_{\mathcal{F}'} \times \mathcal{L}_{\mathcal{H}}$, there exists a LUF adversary \mathcal{A}_{luf} against \mathcal{F}' and an adversary \mathcal{A}_{hash} against \mathcal{H} such that:*

$$\mathbf{Adv}_{\text{MAC}[\mathcal{H}, \mathcal{F}']}^{\text{suf-cmla}}(\mathcal{A}_{mac}, \mathcal{L}_T, \mathcal{L}_V) \leq 2 \mathbf{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}) + \mathbf{Adv}_{\mathcal{F}'}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\mathcal{F}'}) .$$

Suppose \mathcal{A}_{mac} makes at most q queries totalling at most μ bits to each of its oracles **Vfy**, **LTag**, and **LVfy**, and q_ρ to ρ . Then \mathcal{A}_{luf} makes at most $2q$ queries totalling at most $2qw$ bits to each of the oracles in the LUF game, and $Q_{\mathcal{H}}(6q, 6\mu) + q_\rho$ queries to ρ . As for \mathcal{A}_{hash} it needs at most $Q_{\mathcal{F}'}(6q, 6qw) + Q_{\mathcal{H}}(6q, 6\mu)$ queries to ρ to simulate \mathcal{F}' and \mathcal{H} .

4.3 MAC $[\mathcal{H}, \mathcal{F}']$ is LPRF Secure

The final piece needed to apply Theorem 1 is to show that MAC $[\mathcal{H}, \mathcal{F}']$, or rather its tagging algorithm $\mathcal{T}[\mathcal{H}, \mathcal{F}']$, is a leakage-resilient PRF. Since by assumption \mathcal{F}' is already an LPRF, this result is analogous to Theorem 4 in that it provides us with simple technique for extending the domain of an LPRF. The proof can be found in Appendix B.4.

Theorem 5. *Let $\mathcal{F}' : \mathcal{K} \times \{0, 1\}^w \rightarrow \{0, 1\}^t$ be a function family with associated leakage set $\mathcal{L}_{\mathcal{F}'}$, and let $\mathcal{H} : \mathcal{X} \rightarrow \{0, 1\}^w$ be a hash function over the domain \mathcal{X} . Further let their composition $\hat{\mathcal{F}}$ be defined as*

$$\hat{\mathcal{F}}(K, X) = \mathcal{F}'(K, \mathcal{H}(X))$$

where $X \in \mathcal{X}$, $K \in \mathcal{K}$, and $\mathcal{L}_{\hat{\mathcal{F}}} = \mathcal{L}_{\mathcal{F}'} \times \mathcal{L}_{\mathcal{H}}$ for any set of efficiently computable functions $\mathcal{L}_{\mathcal{H}}$. Then for any LPRF adversary \mathcal{A}_{lprf} against $\hat{\mathcal{F}}$, there exists a corresponding LPRF adversary \mathcal{A}'_{lprf} against \mathcal{F}' and an adversary \mathcal{A}_{hash} against \mathcal{H} such that:

$$\mathbf{Adv}_{\hat{\mathcal{F}}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_{\hat{\mathcal{F}}}) \leq 2 \mathbf{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}) + \mathbf{Adv}_{\mathcal{F}'}^{\text{lprf}}(\mathcal{A}'_{lprf}, \mathcal{L}_{\mathcal{F}'}) .$$

Let q and μ be such that \mathcal{A}_{lprf} makes at most q queries totalling μ bits to each of its oracles \mathbf{F} and \mathbf{LF} , and let q_ρ denote the number of queries it makes to ρ . Then \mathcal{A}'_{lprf} makes at most q queries totalling qw bits to each of its oracles \mathbf{F} and \mathbf{LF} , and at most $Q_{\mathcal{H}}(2q, 2\mu) + q_\rho$ queries to ρ . As for \mathcal{A}_{hash} , it requires at most $Q_{\mathcal{F}'}(2q, 2qw) + Q_{\mathcal{H}}(2q, 2\mu)$ queries to ρ in order to simulate \mathcal{F}' and \mathcal{H} .

4.4 The FGHF' Composition Theorem

Collecting the results from this section and combining it with the N2 composition theorem we get the following composition theorem for the FGHF' construction.

Theorem 6 (LAE Security of the FGHF' Construction). *Let \mathcal{F} be a fixed-input-length LPRF, \mathcal{G} a PRG, \mathcal{H} a vector hash function, and \mathcal{F}' be a fixed-input-length function that is both an LUF and an LPRF with associated leakage sets $\mathcal{L}_{\mathcal{F}}$, $\mathcal{L}_{\mathcal{G}}$, $\mathcal{L}_{\mathcal{H}}$, and $\mathcal{L}_{\mathcal{F}'}$, respectively. Let FGHF' be the composition of \mathcal{F} , \mathcal{G} , \mathcal{H} , and \mathcal{F}' with associated leakage sets $\mathcal{L}_{AE} = \mathcal{L}_{VD} = \mathcal{L}_{\mathcal{F}} \times \mathcal{L}_{\mathcal{G}} \times \mathcal{L}_{\mathcal{H}} \times \mathcal{L}_{\mathcal{F}'}$. Then for any LAE adversary \mathcal{A}_{ae} against FGHF' there exist adversaries \mathcal{A}_{lprf} , \mathcal{A}'_{lprf} , \mathcal{A}_{prg} , \mathcal{A}_{hash} , and \mathcal{A}_{luf} such that:*

$$\begin{aligned} \mathbf{Adv}_{\text{FGHF}'}^{\text{lae}}(\mathcal{A}_{ae}, \mathcal{L}_{AE}, \mathcal{L}_{VD}) &\leq 2 \mathbf{Adv}_{\mathcal{F}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_{\mathcal{F}}) + 2 \mathbf{Adv}_{\mathcal{F}'}^{\text{lprf}}(\mathcal{A}'_{lprf}, \mathcal{L}_{\mathcal{F}'}) \\ &\quad + 2 \mathbf{Adv}_{\mathcal{G}}^{\text{prg}}(\mathcal{A}_{prg}) + 6 \mathbf{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}) \\ &\quad + 2 \mathbf{Adv}_{\mathcal{F}'}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\mathcal{F}'}) . \end{aligned}$$

Now suppose \mathcal{A}_{ae} makes at most q queries totalling at most μ bits to each of its \mathbf{Enc} , \mathbf{LEnc} , \mathbf{Dec} , and \mathbf{LDec} oracles, and let q_ρ denote its number of queries to ρ . Then, \mathcal{A}_{lprf} makes at most $2q$ queries totalling $2qm$ bits to each of its oracles \mathbf{F} and \mathbf{LF} , and at most $2Q_{\mathcal{H}}(2q, 2\mu) + 2Q_{\mathcal{F}'}(2q, 2qw) + Q_{\mathcal{G}}(2q, 2\mu)$ queries to ρ . Similarly, \mathcal{A}'_{lprf} makes at most $2q$ queries totalling $2qw$ bits to each of its oracles \mathbf{F} and \mathbf{LF} , and at most $2Q_{\mathcal{F}}(q, qm) + 2Q_{\mathcal{G}}(q, \mu) + Q_{\mathcal{H}}(4q, 4\mu)$ queries to ρ . \mathcal{A}_{luf} makes at most $4q$ queries, totalling $4qw$ bits to each of its oracles \mathbf{F} and \mathbf{Lkg} , and at most $2Q_{\mathcal{F}}(2q, 2qm) + 2Q_{\mathcal{G}}(2q, 2\mu) + Q_{\mathcal{H}}(12q, 12\mu)$ to ρ . As for \mathcal{A}_{prg} , it makes at most q queries, totalling μ bits, to its oracle \mathbf{G} and at most $2Q_{\mathcal{H}}(2q, 2\mu) + 2Q_{\mathcal{F}'}(2q, 2qw) + Q_{\mathcal{F}}(2q, 2qm) + Q_{\mathcal{G}}(2q, 2\mu)$ queries to ρ . Finally, \mathcal{A}_{hash} requires at most $2Q_{\mathcal{F}}(2q, 2qm) + 2Q_{\mathcal{G}}(2q, 2\mu) + Q_{\mathcal{F}'}(12q, 12qw) + Q_{\mathcal{H}}(12q, 12\mu)$ queries to ρ .

5 Security of Sponge-Based Primitives

We now turn our attention to instantiating the constituent blocks of the FGHF' construction using sponge-based primitives. Specifically we prove the security of the vector hash function SVHASH, the pseudorandom generator SPRG, and the leakage-resilient function family SLFUNC for instantiating both \mathcal{F} and \mathcal{F}' . All primitives are based on the T-sponge and this particular instantiation of the FGHF' construction gives rise to SLAE. The most interesting results are Theorems 7 and 8 which substantiate our claim that sponges offer an inherent resistance to non-adaptive leakage. Informally these two theorems state that by using a reduced absorption rate (e.g. $rr = 1$), the impact of leakage on security can be limited to λ^{2rr} . In particular, note that the aggregate security loss is independent of the number of

queries, meaning that the effect of leakage cannot be amplified by repeated use. While sponge-based hash functions and pseudorandom generators have been studied quite extensively, SVHASH and SPRG are non-standard constructions. Firstly, they are based on a transformation rather than a permutation which is not common in the literature. Secondly, unlike other constructions SPRG treats the whole initial state as the seed, and SVHASH takes a triple of strings as its input. Thus while not particularly novel, we include their security proofs for completeness.

5.1 Revised Security Bounds

The security bounds presented in the proceedings version of this paper [11] were overly optimistic, and in fact, incorrect. This was pointed out to us by Bart Mennink and Krzysztof Pietrzak. The error was due to a wrong enumeration of the support when applying Lemma 1 in the proofs.

5.2 A Sponge-Based Leakage-Resilient Function Family

Although LPRF and LUF security are incomparable notions, it is still possible to meet both notions simultaneously through a single primitive. Indeed the FGHF' construction requires that such a primitive exist since \mathcal{F}' is required to satisfy both security notions. We now show that the SLFUNC construction is well-suited for this role, and in fact that it can be used to instantiate both the \mathcal{F} and \mathcal{F}' components – as is the case in SLAE. Moreover, the most extensively studied leakage-resilient object is that of a pseudorandom function due to its versatility in several potential applications. SLFUNC yields a practical construction of this primitive against non-adaptive leakage and as such we think it may be of independent interest. The security of SLFUNC is stated formally in the following two theorems. Their proofs can be found in Sections C.1 and C.2 of the Appendix.

Theorem 7. *Let SLFUNC be the function family described in Fig. 9 taking as input strings of size $(l \cdot rr)$ bits and returning t -bit strings. Then for any LPRF adversary \mathcal{A} against SLFUNC and any vector of leakage functions $[L_0, \dots, L_l]$ where each component maps n bits to λ bits such that $\mathcal{L}_\lambda = \{[L_0, \dots, L_l]\}$, it holds that:*

$$\text{Adv}_{\text{SLFUNC}}^{\text{lprf}}(\mathcal{A}, \mathcal{L}_\lambda) \leq \frac{q_T(q_T + 2) + (q_F + q_{\text{LF}})q_\rho}{2^{n-rr-1}} + \frac{2q_\rho}{2^{k-\lambda 2^{rr}}} + \frac{2lq_Fq_\rho}{2^{n-\lambda 2^{rr}}}.$$

In the above q_ρ , q_F , and q_{LF} denote respectively the number of queries \mathcal{A} makes to its oracles ρ , F , and LF and $q_T = (l + 1)(q_{\text{LF}} + q_F) + q_\rho$. Moreover it is required that $q_\rho + l(q_F + q_{\text{LF}}) \leq 2^{k-1}$ and $(2^{rr})q_\rho + l(q_F + q_{\text{LF}}) \leq 2^{n-1}$.

The next Theorem shows that SLFUNC is a good LUF. Its proof bears some similarity to that of Theorem 7 as it uses similar ideas. However one important difference lies in the leakage model that is used in this theorem. Since the Lkg oracle returns only the leakage and no output, we add here an extra leakage function that returns the leakage on the output of SLFUNC. In the LPRF case this was not required since in that game the leakage oracle returns the full output anyway.

Theorem 8. *Let SLFUNC be the function family described in Fig. 9 taking as input strings of size $(l \cdot rr)$ bits and returning t -bit long strings. Then for any LUF adversary \mathcal{A} against SLFUNC, and any vector of leakage functions $[L_0, \dots, L_{l+1}]$ where each component maps n bits to λ bits and letting $\mathcal{L}_\lambda = \{[L_0, \dots, L_{l+1}]\}$, it holds that:*

$$\text{Adv}_{\text{SLFUNC}}^{\text{luf}}(\mathcal{A}, \mathcal{L}_\lambda) \leq \frac{q_T(q_T + 2)}{2^{n-rr}} + \frac{2q_\rho}{2^{k-\lambda 2^{rr}}} + \frac{2lq_{\text{Lkg}}q_\rho}{2^{n-\lambda 2^{rr}}} + \frac{q_{\text{Guess}}}{2^{t-\lambda-1}}.$$

In the above q_ρ , q_F , q_{Lkg} and q_{Guess} denote respectively the number of queries \mathcal{A} makes to its oracles ρ , F , Lkg , and Guess and $q_T = (l + 1)(q_F + q_{\text{Lkg}} + q_{\text{Guess}}) + q_\rho$. Moreover it is required that the following conditions be satisfied $q_\rho + (l + 1)(q_F + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{k-1}$, $(2^{rr})q_\rho + (l + 1)(q_F + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{n-1}$, and $q_{\text{Guess}} + (l + 1)(q_F + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{n-1}$.

5.3 The Security of SPRG

As explained in Section 3.3, SLENC can be decomposed into the cascade of SLFUNC and SPRG, matching the encryption component of the FGHF' construction. A pseudocode description of the variable-output-length pseudorandom generator SPRG is given in Fig. 14. Decomposing SLENC this way requires us

to treat all of SPRG’s initial state as the seed, which deviates from the more conventional ways of constructing sponge-based pseudorandom generators. Moreover we consider a security definition which allows the adversary to make multiple queries to the PRG, each with differing output lengths.

The security of SPRG is stated formally in Theorem 9. Its proof follows from a standard hybrid argument and can be found in Appendix C.3.

Theorem 9. *Let SPRG be the pseudorandom generator described in Fig. 14. Then for any PRG adversary \mathcal{A} , it holds that:*

$$\text{Adv}_{\text{SPRG}}^{\text{prg}}(\mathcal{A}) \leq \frac{v q_{\mathbf{G}} q_{\rho}}{2^c - q_{\rho}} + \frac{v q_{\rho} q_{\mathbf{G}} + v^2 q_{\mathbf{G}}^2}{2^n} + \frac{q_{\mathbf{G}}^2}{2^n}.$$

In the above, \mathcal{A} makes at most q_{ρ} queries to the random transformation ρ and $q_{\mathbf{G}}$ queries to the challenge oracle \mathbf{G} . Moreover, $v = \lceil \frac{L_{\max}}{r} \rceil$, where L_{\max} is an upper bound on the adversary’s inputs to \mathbf{G} .

5.4 A Sponge-Based Vector Hash Function

The final building block is the sponge-based vector hash function SVHASH which is graphically represented in Fig. 10. It takes as input a triple of strings, namely a nonce, associated data and a ciphertext to return a string digest. A salient feature of this construction is the xoring of $1 \parallel 0^{c-1}$ into the inner state in order to separate the (padded) associated data from the (padded) ciphertext. We prove the security of SVHASH in a modular fashion, by first reducing its security to that of a plain hash function taking a single input and then prove the collision-resistance of this latter construction in the random transformation model. The collision-resistance of SVHASH is stated formally in the following theorem, and the full proof details can be found in Appendix C.4.

Theorem 10. *Let SVHASH be the vector hash function described in Fig. 10. Then for any adversary \mathcal{A} making q queries to ρ , it holds that:*

$$\text{Adv}_{\text{SVHASH}}^{\text{cr}}(\mathcal{A}) \leq \frac{q(q-1)}{2^{w+1}} + \frac{q(q+2)}{2^{c-1}}.$$

5.5 Concrete Security of SLAE

A bound for the security of SLAE is obtained by directly combining the FGHF’ composition theorem (Theorem 6) with Theorems 7–10. It then only remains to derive concrete bounds for the expressions $Q_{\mathcal{F}}$, $Q_{\mathcal{G}}$, $Q_{\mathcal{H}}$, $Q_{\mathcal{F}'}$ for the specific case of SLAE. Assuming a nonce size of m bits and that the output of \mathcal{H} is w bits long, the following expressions are easily derived from the algorithm definitions. Namely, we have that:

$$\begin{aligned} Q_{\mathcal{F}}(q, qm) &= q \left\lceil \frac{m+1}{rr} \right\rceil & Q_{\mathcal{F}'}(q, qw) &= q \left\lceil \frac{w+1}{rr} \right\rceil \\ Q_{\mathcal{G}}(q, \mu) &= \left\lceil \frac{\mu}{r} \right\rceil & Q_{\mathcal{H}}(q, \mu) &= \left\lceil \frac{\mu}{r} \right\rceil + 3q. \end{aligned}$$

6 Concluding Remarks and Implementation Aspects

In this work we proposed the FGHF’ construction as a template for constructing non-adaptively leakage-resilient AEAD schemes from relatively simpler primitives – requiring only two calls to the leakage-resilient functions per encryption or decryption call. We then presented SLAE as a sponge-based instantiation of this construction, offering good performance and simplicity. Our security analysis shows that if the absorption rate is set sufficiently low, the transformation sponge yields a leakage-resilient function with the desired properties. However some care is needed in interpreting these results. Like most treatments of leakage resilience we assume that the leakage per evaluation is limited and does not drain the entropy in the secret state. Thus it is implicitly assumed that an implementation is good enough to withstand basic side-channel attacks like Simple Power Analysis (SPA) attacks. The benefit of our leakage-resilience security proof is that resistance to basic attacks automatically translates to resistance against more sophisticated attacks like Differential Power Analysis (DPA).

In the FGHF' construction and SLAE, authenticity is verified by recomputing the MAC tag and testing for equality between the recomputed tag and the one included in the ciphertext. While our leakage model accounts for the leakage that may take place during the tag recomputation, equality testing is assumed to be leak-free. Thus any implementation of SLAE (or any other realisation of the FGHF' construction) needs to ensure that equality testing does not leak, or take additional measures, such as masking, to protect against leakage from this component.

Finally the security of SLAE relies on it being instantiated with a non-invertible transformation rather than a permutation. On the other hand, most practical schemes employ permutations, such as KECCAK- p and XODOO- p . While in this work we did not specify any concrete transformation, a natural candidate is to use $\rho(x) = p(x) \oplus x$ for $p \in \{\text{KECCAK-}p, \text{XODOO-}p\}$. Although this construction is known to be differentiable from a random transformation when given access to p , this should not preclude it from being a suitable candidate for instantiating constructions in the random transformation model. Indeed, KECCAK- p and XODOO- p are also differentiable from a random permutation when given access to their underlying building blocks.

Acknowledgements

We are grateful to Bart Mennink and Krzysztof Pietrzak for pointing out an error in a previous version of this paper. We thank Daniel Baur and Christian Schuller for initial discussions during the early stages of this project, and our anonymous reviewers for their constructive comments. Degabriele was supported by the German Federal Ministry of Education and Research (BMBF) as well as by the Hessian State Ministry for Higher Education, Research and Arts (HMWK) within CRISP. Janson and Struck have been (partially) funded by the Deutsche Forschungsgemeinschaft (DFG) – SFB 1119 – 236615297.

References

1. M. Abdalla, S. Belaïd, and P.-A. Fouque. Leakage-resilient symmetric encryption via re-keying. In G. Bertoni and J.-S. Coron, editors, *CHES 2013*, volume 8086 of *LNCS*, pages 471–488. Springer, Heidelberg, Aug. 2013.
2. F. abed, F. Berti, and S. Lucks. Insecurity of RCB: Leakage-resilient authenticated encryption. Cryptology ePrint Archive, Report 2016/1121, 2016. <http://eprint.iacr.org/2016/1121>.
3. M. Agrawal, T. K. Bansal, D. Chang, A. K. Chauhan, S. Hong, J. Kang, and S. K. Sanadhya. Rcb: leakage-resilient authenticated encryption via re-keying. *The Journal of Supercomputing*, 74(9):4173–4198, Sep 2018.
4. G. Barwell, D. P. Martin, E. Oswald, and M. Stam. Authenticated encryption in the face of protocol and side channel leakage. In T. Takagi and T. Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 693–723. Springer, Heidelberg, Dec. 2017.
5. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
6. D. J. Bernstein. CAESAR: Competition for Authenticated Encryption: Security, Applicability, and Robustness, 2014.
7. F. Berti, F. Koeune, O. Pereira, T. Peters, and F.-X. Standaert. Leakage-resilient and misuse-resistant authenticated encryption. Cryptology ePrint Archive, Report 2016/996, 2016. <http://eprint.iacr.org/2016/996>.
8. F. Berti, O. Pereira, T. Peters, and F.-X. Standaert. On leakage-resilient authenticated encryption with decryption leakages. *IACR Trans. Symm. Cryptol.*, 2017(3):271–293, 2017.
9. G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche. Sponge functions. In *ECRYPT Hash Workshop, 2007*. <https://keccak.team/files/SpongeFunctions.pdf>.
10. D. Boneh and V. Shoup. *A Graduate Course in Applied Cryptography*. draft 0.4.
11. J. P. Degabriele, C. Janson, and P. Struck. Sponges resist leakage: The case of authenticated encryption. In S. D. Galbraith and S. Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 209–240. Springer, Heidelberg, Dec. 2019.
12. C. Dobraunig, M. Eichlseder, S. Mangard, F. Mendel, and T. Unterluggauer. ISAP – towards side-channel secure authenticated encryption. *IACR Trans. Symm. Cryptol.*, 2017(1):80–105, 2017.
13. C. Dobraunig and B. Mennink. Leakage resilience of the duplex construction. In S. D. Galbraith and S. Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 225–255. Springer, Heidelberg, Dec. 2019.
14. Y. Dodis, R. Ostrovsky, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. Cryptology ePrint Archive, Report 2003/235, 2003. <http://eprint.iacr.org/2003/235>.
15. Y. Dodis and K. Pietrzak. Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 21–40. Springer, Heidelberg, Aug. 2010.

16. S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pages 293–302. IEEE Computer Society Press, Oct. 2008.
17. S. Faust, K. Pietrzak, and J. Schipper. Practical leakage-resilient symmetric cryptography. In E. Prouff and P. Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 213–232. Springer, Heidelberg, Sept. 2012.
18. O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, Oct. 1986.
19. C. Guo, O. Pereira, T. Peters, and F.-X. Standaert. Leakage-resilient authenticated encryption with misuse in the leveled leakage setting: Definitions, separation results, and constructions. Cryptology ePrint Archive, Report 2018/484, 2018. <https://eprint.iacr.org/2018/484>.
20. C. Guo, O. Pereira, T. Peters, and F.-X. Standaert. Towards lightweight side-channel security and the leakage-resilience of the duplex sponge. Cryptology ePrint Archive, Report 2019/193, 2019. <https://eprint.iacr.org/2019/193>.
21. J. Longo, D. P. Martin, E. Oswald, D. Page, M. Stam, and M. Tunstall. Simulatable leakage: Analysis, pitfalls, and new constructions. In P. Sarkar and T. Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 223–242. Springer, Heidelberg, Dec. 2014.
22. C. Namprempre, P. Rogaway, and T. Shrimpton. Reconsidering generic composition. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 257–274. Springer, Heidelberg, May 2014.
23. O. Pereira, F.-X. Standaert, and S. Vivek. Leakage-resilient authentication and encryption from symmetric cryptographic primitives. In I. Ray, N. Li, and C. Kruegel, editors, *ACM CCS 2015*, pages 96–108. ACM Press, Oct. 2015.
24. P. Rogaway. Authenticated-encryption with associated-data. In V. Atluri, editor, *ACM CCS 2002*, pages 98–107. ACM Press, Nov. 2002.
25. F. Standaert, O. Pereira, Y. Yu, J. Quisquater, M. Yung, and E. Oswald. Leakage resilient cryptography in practice. In A. Sadeghi and D. Naccache, editors, *Towards Hardware-Intrinsic Security - Foundations and Practice*, Information Security and Cryptography, pages 99–134. Springer, 2010.
26. F.-X. Standaert, O. Pereira, and Y. Yu. Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 335–352. Springer, Heidelberg, Aug. 2013.
27. Y. Yu, F.-X. Standaert, O. Pereira, and M. Yung. Practical leakage-resilient pseudorandom generators. In E. Al-Shaer, A. D. Keromytis, and V. Shmatikov, editors, *ACM CCS 2010*, pages 141–151. ACM Press, Oct. 2010.

A Additional Preliminary Material

A.1 Standard Cryptographic Primitives

We make use of the following definition of a pseudorandom generator. Note that our syntax defines a pseudorandom generator with variable output length, where the output length (in bits) is specified as part of the input. In addition, in our security definition we allow the adversary to make multiple queries to the challenge oracle \mathcal{G} .

Definition 6 (Pseudorandom Generators). *Let $\mathcal{G}: \mathcal{S} \times \mathbb{N} \rightarrow \{0, 1\}^*$ be a pseudorandom generator with an associated seed space \mathcal{S} , and let the PRG game be as defined in Fig. 16. Then for any adversary \mathcal{A} , its corresponding PRG advantage is given by:*

$$\text{Adv}_{\mathcal{G}}^{\text{prg}}(\mathcal{A}) = 2 \Pr \left[\text{PRG}^{\mathcal{A}} \Rightarrow \text{true} \right] - 1.$$

We define collision-resistant hash functions over a generic domain \mathcal{X} . Letting $\mathcal{X} = \{0, 1\}^*$ results in the usual syntax but we can also, for instance, model a vector hash function over a triple of strings by setting $\mathcal{X} = \{0, 1\}^* \times \{0, 1\}^* \times \{0, 1\}^*$. For simplicity we only consider the random transformation model.

Definition 7 (Collision-Resistant Hash Functions). *Let $\mathcal{H}: \mathcal{X} \rightarrow \{0, 1\}^w$ be a hash function constructed from a random transformation ρ , with domain \mathcal{X} and output length w . Then for any adversary \mathcal{A} with oracle access to ρ , its corresponding advantage is given by:*

$$\text{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}) = 2 \Pr[(X_0, X_1) \leftarrow \mathcal{A}^\rho : \mathcal{H}(X_0) = \mathcal{H}(X_1) \wedge X_0 \neq X_1 \wedge X_0, X_1 \in \mathcal{X}].$$

Game PRG	procedure $\mathsf{G}(L)$
procedure Initialize	if $b = 0$
$b \leftarrow \{0, 1\}$	$R \leftarrow \{0, 1\}^L$
return	return R
	else
	$seed \leftarrow \mathcal{S}$
	$R \leftarrow \mathcal{G}(seed, L)$
	return R
	procedure Finalize (b')
	return ($b' = b$)

Fig. 16: Game used to define PRG security.

A.2 Min-Entropy

Below are some definitions and results on min-entropy that we will need in some of the proofs.

Definition 8. Let X , Y , and Z be random variables. The average-case guessing probability of X conditioned on Y and $Z = z$ is given by

$$\text{GP}(X | Y, Z = z) := \mathbb{E}_{y \leftarrow Y} \left[\max_x \Pr[X = x | Y = y \wedge Z = z] \right].$$

The average-case min-entropy of X conditioned on Y and $Z = z$ is defined to be:

$$\text{H}_\infty(X | Y, Z = z) := -\log(\text{GP}(X | Y, Z = z)).$$

Lemma 1 ([14]). Let X , Y , and Z be random variables such that the support of Y is bounded above by 2^λ . Then

$$\text{H}_\infty(X | Y, Z = z) \geq \text{H}_\infty(X, Y | Z = z) - \lambda \geq \text{H}_\infty(X | Z = z) - \lambda.$$

B Security Proofs for Section 4

B.1 Proof of Theorem 2

We prove the theorem through the sequence of four games which are described in Fig. 17. The initialize and finalize procedure as well as the leakage oracles are identical for all games, hence we only list them once. Only the Enc oracle is changed in each game. Now, game G_0 corresponds to the INDaCPLA game instantiated with $\text{SE}[\mathcal{F}, \mathcal{G}]$, whereas game G_3 is defined so that for any adversary its probability of winning is exactly $\frac{1}{2}$.

It then follows that

$$\begin{aligned} \text{Adv}_{\text{SE}[\mathcal{F}, \mathcal{G}]}^{\text{ind-acpla}}(\mathcal{A}_{se}, \mathcal{L}_E, \mathcal{L}_D) &= 2 \text{Adv}(\mathsf{G}_0^{\mathcal{A}_{se}}, \mathsf{G}_3^{\mathcal{A}_{se}}) + 2 \Pr[\mathsf{G}_3^{\mathcal{A}_{se}} \Rightarrow \text{true}] - 1 \\ &= 2 \sum_{i=0}^2 \text{Adv}(\mathsf{G}_i^{\mathcal{A}_{se}}, \mathsf{G}_{i+1}^{\mathcal{A}_{se}}). \end{aligned}$$

We will bound the IND-aCPLA advantage of \mathcal{A}_{se} by providing bounds on the distinguishing advantages between subsequent games. We reduce the ability to distinguish between games G_0 and G_1 to the LPRF security of \mathcal{F} , and that between games G_1 and G_2 to the PRG security of \mathcal{G} . As for G_2 and G_3 , we will show that no adversary can distinguish between them, and therefore the corresponding advantage term is also zero.

Let us begin by bounding the distinguishing advantage between game G_0 and game G_1 . Recall that G_0 is the INDaCPLA game instantiated with $\text{SE}[\mathcal{F}, \mathcal{G}]$. Game G_1 is the same, except that the challenge oracle samples $seed$ at random instead of using the pseudorandom function \mathcal{F} . We now construct an

<pre> procedure Initialize <hr/> $d \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$ return procedure LEnc($N, M, (L_F, L_G)$) <hr/> $seed \leftarrow \mathcal{F}(K, N)$ $R \leftarrow \mathcal{G}(seed, M)$ $C \leftarrow R \oplus M$ $A_F \leftarrow L_F(K, N)$ $A_G \leftarrow L_G(seed, M)$ return ($C, (A_F, A_G)$) procedure LDec($N, C, (L_F, L_G)$) <hr/> $seed \leftarrow \mathcal{F}(K, N)$ $R \leftarrow \mathcal{G}(seed, C)$ $M \leftarrow R \oplus C$ $A_F \leftarrow L_F(K, N)$ $A_G \leftarrow L_G(seed, C)$ return ($M, (A_F, A_G)$) </pre>	<pre> procedure Enc(N, M) in G_0 and $\boxed{\mathsf{G}_1}$ <hr/> if $d = 0$ if $f[N, M] = \perp$ $f[N, M] \leftarrow \{0, 1\}^{ M }$ return $f[N, M]$ else $seed \leftarrow \mathcal{F}(K, N)$ $seed \leftarrow \{0, 1\}^n$ $R \leftarrow \mathcal{G}(seed, M)$ $C \leftarrow R \oplus M$ return C procedure Finalize (d') <hr/> return ($d' = d$) </pre>
<pre> procedure Enc(N, M) in G_2 <hr/> if $d = 0$ if $f[N, M] = \perp$ $f[N, M] \leftarrow \{0, 1\}^{ M }$ return $f[N, M]$ else $R \leftarrow \{0, 1\}^{ M }$ $C \leftarrow R \oplus M$ return C </pre>	<pre> procedure Enc(N, M) in G_3 <hr/> if $d = 0$ $C \leftarrow \{0, 1\}^{ M }$ return C else $C \leftarrow \{0, 1\}^{ M }$ return C </pre>

Fig. 17: Games G_0 , G_1 , G_2 , and G_3 used to prove Theorem 2. Game G_1 contains the boxed code, G_0 does not. Games G_2 and G_3 are identical to Game G_1 except that oracle Enc is replaced with the one described.

adversary \mathcal{A}_{lprf} against the LPRF security of \mathcal{F} from any adversary \mathcal{A}_{se} trying to distinguish between games G_0 and G_1 . Note that there will be two secret bits at play, b and d . The former is the one used in the LPRF game (cf. Fig. 3) which \mathcal{A}_{lprf} is trying to guess. Depending on this bit, \mathcal{A}_{lprf} will simulate either G_0 or G_1 to \mathcal{A}_{se} . On the other hand, bit d is the secret bit sampled by \mathcal{A}_{lprf} when simulating game G_0 or G_1 which \mathcal{A}_{se} is trying to guess.

The adversary \mathcal{A}_{lprf} starts by sampling d at random and then proceeds as follows. For any leakage query $(N, M, (L_F, L_G))$ that \mathcal{A}_{se} makes to LEnc, \mathcal{A}_{lprf} queries (N, L_F) to LF to obtain $(seed, A_F)$ and computes $A_G = L_G(seed, M)$. It then computes the ciphertext $C = \mathcal{G}(seed, |M|) \oplus M$ and forwards $(C, (A_F, A_G))$ to \mathcal{A}_{se} . Queries of the form $(N, C, (L_F, L_G))$ which \mathcal{A}_{se} makes to LDec are answered just as queries to LEnc, except that the message M is obtained by XORing the ciphertext C to the output of the PRG \mathcal{G} and that L_G takes C instead of M as input. For any challenge query (N, M) by \mathcal{A}_{se} , the behaviour of \mathcal{A}_{lprf} depends on the bit d . If it is equal to 0, it returns a random bit string of length equal to the length of the message and records this bit string using the table f . If the bit d is equal to 1, \mathcal{A}_{lprf} queries N to its own challenge oracle F obtaining the seed $seed$ for the pseudorandom generator \mathcal{G} . Subsequently, it computes $C = \mathcal{G}(seed, |M|) \oplus M$ and sends C to \mathcal{A}_{se} . Eventually, \mathcal{A}_{se} outputs a bit d' to indicate whether it was in the ideal or in the real world. If \mathcal{A}_{se} guessed correctly, i.e. $d' = d$, \mathcal{A}_{lprf} outputs $b' \leftarrow 0$. Otherwise, \mathcal{A}_{lprf} outputs $b' \leftarrow 1$.

From its leakage oracle, \mathcal{A}_{lprf} obtains the leakage $\Lambda_F = L_F(K, N)$ and the seed for the pseudorandom generator. It then uses the seed to compute the ciphertext C or the message M together with the leakage Λ_G . When $d = 0$ (ideal world), \mathcal{A}_{lprf} can easily simulate the challenge oracle by sampling random ciphertexts. However, when $d = 1$ (real world), \mathcal{A}_{lprf} uses its own challenge oracle to obtain the seed for \mathcal{G} . Then depending on the bit b , the seed will be the output of a pseudorandom function ($b = 1$) or sampled randomly ($b = 0$). Since the games G_0 and G_1 only differ in how the seed for the PRG is generated, depending on the value of b , \mathcal{A}_{lprf} provides \mathcal{A}_{se} with a perfect simulation of either G_0 or G_1 .

It remains to verify that \mathcal{A}_{lprf} is a valid LPRF adversary. By construction, as long as \mathcal{A}_{se} queries leakage functions from the permitted leakage set then so will \mathcal{A}_{lprf} . For any leakage query $(N, \cdot, L = (L_F, \cdot))$ to either LEnc or LDec by \mathcal{A}_{se} , \mathcal{A}_{lprf} queries LF on (N, L_F) . It remains to argue that \mathcal{A}_{lprf} does not forward queries to or from its F oracle. Note that \mathcal{A}_{lprf} only makes queries to F to answer \mathcal{A}_{se} 's queries to Enc . Then since \mathcal{A}_{se} must be nonce-respecting, it cannot reuse the same nonce across Enc and LEnc . Moreover, all of its queries to LDec must be forwarded from LEnc . Thus the nonces between Enc and the leakage oracles have to be disjoint and therefore \mathcal{A}_{lprf} 's queries to F and LF will also be disjoint. We are now ready to derive our first bound. From the foregoing discussion we have that

$$\begin{aligned} \mathbf{Adv} \left(\mathsf{G}_0^{\mathcal{A}_{se}}, \mathsf{G}_1^{\mathcal{A}_{se}} \right) &= \Pr[\mathsf{G}_0^{\mathcal{A}_{se}} \Rightarrow \text{true}] - \Pr[\mathsf{G}_1^{\mathcal{A}_{se}} \Rightarrow \text{true}] \\ &\leq \Pr[\mathcal{A}_{lprf}^{\text{LPRF}} \Rightarrow 0 \mid b = 1] - \Pr[\mathcal{A}_{lprf}^{\text{LPRF}} \Rightarrow 0 \mid b = 0] \\ &\leq \mathbf{Adv}_{\mathcal{F}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_F). \end{aligned}$$

Next, we bound the advantage in distinguishing G_1 from G_2 . The only difference between these two games is how R is generated, in G_1 this is generated as $\mathcal{G}(\text{seed})$ whereas in G_2 it is sampled uniformly at random. For any adversary \mathcal{A}_{se} trying to distinguish between G_1 and G_2 we construct a PRG adversary \mathcal{A}_{prg} against \mathcal{G} as follows. It starts by sampling a random key K for the PRF \mathcal{F} and a random bit d . For any query $(N, M, (L_F, L_G))$ to LEnc that \mathcal{A}_{se} makes, \mathcal{A}_{prg} computes locally $\text{seed} \leftarrow \mathcal{F}(K, N)$, $R \leftarrow \mathcal{G}(\text{seed}, |M|)$, and $C \leftarrow R \oplus M$. It then computes $\Lambda_F \leftarrow L_F(K, N)$ as well as $\Lambda_G \leftarrow L_G(\text{seed}, M)$ and returns $(C, (\Lambda_F, \Lambda_G))$ to \mathcal{A}_{se} . Leakage queries to LDec are handled similarly, only XORing the ciphertext to the output of the PRG \mathcal{G} and feeding the ciphertext as input to L_G . For any challenge query (N, M) by \mathcal{A}_{se} , when $d = 0$ \mathcal{A}_{prg} returns a random bit string of length equal to $|M|$. Otherwise, if $d = 1$ \mathcal{A}_{prg} queries its own challenge oracle G on $|M|$ to obtain R and returns $C \leftarrow R \oplus M$ to \mathcal{A}_{se} . Then, when \mathcal{A}_{se} outputs a guess d' , \mathcal{A}_{prg} outputs 0 if $d' = d$ and outputs 1 otherwise.

Now note that if $b = 0$ then \mathcal{A}_{prg} provides \mathcal{A}_{se} with simulation of G_2 and if $b = 1$ it simulates G_1 . However when $b = 1$, \mathcal{A}_{prg} evaluates \mathcal{G} both locally and through its oracle G on independently sampled seeds. Since \mathcal{A}_{se} cannot reuse the same nonce across queries to Enc and its leakage oracles, these two sets of queries are always answered using independently-sampled seeds and its simulation of G_2 is perfect. We then have that

$$\begin{aligned} \mathbf{Adv} \left(\mathsf{G}_1^{\mathcal{A}_{se}}, \mathsf{G}_2^{\mathcal{A}_{se}} \right) &= \Pr[\mathsf{G}_1^{\mathcal{A}_{se}} \Rightarrow \text{true}] - \Pr[\mathsf{G}_2^{\mathcal{A}_{se}} \Rightarrow \text{true}] \\ &\leq \Pr[\mathcal{A}_{prg}^{\text{PRG}} \Rightarrow 0 \mid b = 1] - \Pr[\mathcal{A}_{prg}^{\text{PRG}} \Rightarrow 0 \mid b = 0] \\ &\leq \mathbf{Adv}_{\mathcal{G}}^{\text{prg}}(\mathcal{A}_{prg}). \end{aligned}$$

We now conclude the proof by bounding the distinguishing advantage between the games G_2 and G_3 . These two games differ in two aspects. In the ideal world ($d = 0$), the former outputs a random bit string using a table f while the latter forgetfully outputs a random bit string of the same size. However as long as \mathcal{A}_{se} is nonce-respecting it will not be able to distinguish these two cases since every output will be freshly sampled. The other difference is how C is generated in the real world ($d = 1$). However both methods yield identically distributed ciphertexts. Thus the two games are identical and therefore

$$\mathbf{Adv} \left(\mathsf{G}_2^{\mathcal{A}_{se}}, \mathsf{G}_3^{\mathcal{A}_{se}} \right) = 0.$$

Collecting all of the above bounds together yields

$$\mathbf{Adv}_{\mathsf{SE}[\mathcal{F}, \mathcal{G}]}^{\text{ind-acpla}}(\mathcal{A}_{se}, \mathcal{L}_E, \mathcal{L}_D) \leq 2 \mathbf{Adv}_{\mathcal{F}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_F) + 2 \mathbf{Adv}_{\mathcal{G}}^{\text{prg}}(\mathcal{A}_{prg}),$$

which proves the theorem. \square

B.2 Proof of Theorem 3

For any SUF-CMLA adversary \mathcal{A}_{mac} against $\text{MAC}[\hat{\mathcal{F}}]$ we construct a LUF adversary \mathcal{A}_{luf} against $\hat{\mathcal{F}}$. We proceed in two steps, we first bound the SUF-CMLA advantage of \mathcal{A}_{mac} by the probability of a particular event occurring and then we bound this probability by the LUF advantage of \mathcal{A}_{luf} . Let E be the event that the \mathcal{A}_{mac} makes a query to Vfy which returns \top . Then

$$\begin{aligned} \text{Adv}_{\text{MAC}[\hat{\mathcal{F}}]}^{\text{suf-cmla}}(\mathcal{A}_{mac}, \mathcal{L}_T, \mathcal{L}_V) &= 2 \Pr \left[\text{SUF-CMLA}^{\mathcal{A}_{mac}} \Rightarrow \text{true} \right] - 1 \\ &= 2 \Pr \left[\text{SUF-CMLA}^{\mathcal{A}_{mac}} \Rightarrow \text{true} \cap E \right] \\ &\quad + 2 \Pr \left[\text{SUF-CMLA}^{\mathcal{A}_{mac}} \Rightarrow \text{true} \cap \bar{E} \right] - 1 \\ &\leq 2 \Pr[E] + 2 \Pr \left[\text{SUF-CMLA}^{\mathcal{A}_{mac}} \Rightarrow \text{true} \mid \bar{E} \right] - 1. \end{aligned}$$

Now, conditioned on \bar{E} the Vfy oracle will always returns \perp irrespective of the value of b and hence the probability of \mathcal{A}_{mac} winning is exactly one half. Thus the above can be re-written as

$$\leq 2 \Pr[E \cap b = 0] + 2 \Pr[E \cap b = 1].$$

When $b = 0$ the Vfy oracle will always returns \perp and thus E simply cannot occur, i.e. the first term in the above expression is zero. Thus

$$\begin{aligned} &= 2 \Pr[E \mid b = 1] \Pr[b = 1] \\ &= \Pr[E \mid b = 1]. \end{aligned}$$

We now bound this last probability by the LUF advantage of \mathcal{A}_{luf} . By construction we have that $\mathcal{L}_T = \mathcal{L}_V = \mathcal{L}_{\hat{\mathcal{F}}}$. Then \mathcal{A}_{luf} runs \mathcal{A}_{mac} and provides it with a simulation of the SUF-CMLA game with the bit b fixed to 1. Whenever \mathcal{A}_{mac} makes a query (X, L) to LTag , \mathcal{A}_{luf} queries X to F and (X, L) to Lkg to obtain respectively $y = \hat{\mathcal{F}}(K, X)$ and $\Lambda = L(K, X)$, and returns (y, Λ) back to \mathcal{A}_{mac} . Similarly if \mathcal{A}_{mac} queries (X, T, L) to its LVfy oracle, \mathcal{A}_{luf} queries (X, L) to Lkg and (X, T) to Guess to obtain Λ and v respectively, and returns (v, Λ) to \mathcal{A}_{mac} . Every query (X, T) that \mathcal{A}_{mac} makes to its Vfy oracle, is forwarded by \mathcal{A}_{luf} to its own Guess oracle and returns \top to \mathcal{A}_{mac} if Guess returns true and \perp otherwise.

Recall that \mathcal{A}_{luf} 's queries to F are recorded in the set \mathcal{S} , and it only wins the LUF game if it makes a query to Guess which returns true where the corresponding X value is not contained in \mathcal{S} . However since \mathcal{A}_{mac} does not forward queries from LTag to Vfy , it is guaranteed that \mathcal{A}_{luf} 's queries to Guess are not contained in \mathcal{S} . It then follows that \mathcal{A}_{luf} wins the LUF game whenever E occurs, and hence

$$\Pr[E \mid b = 1] \leq \text{Adv}_{\hat{\mathcal{F}}}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\hat{\mathcal{F}}})$$

which proves the theorem. \square

B.3 Proof of Theorem 4

We prove the Theorem through the two games, G and $\boxed{\text{G}}$, as described in Fig. 18. Namely:

- **Game G** is the LUF game instantiated with $\hat{\mathcal{F}}$ and its associated leakage set $\mathcal{L}_{\hat{\mathcal{F}}} = \mathcal{L}_{F'} \times \mathcal{L}_H$, and some additional bookkeeping to maintain the set Ω and test for the event **Bad**.
- **Game $\boxed{\text{G}}$** is identical to game G except that the winning condition is restricted to only hold if **Bad** does not occur.

Moreover G and $\boxed{\text{G}}$ are identical until **Bad**, and therefore by the game playing lemma we have that

$$\begin{aligned} \text{Adv} \left(\text{G}^{\mathcal{A}_{luf}}, \boxed{\text{G}}^{\mathcal{A}_{luf}} \right) &= \Pr[\text{G}^{\mathcal{A}_{luf}} \Rightarrow \text{true}] - \Pr[\boxed{\text{G}}^{\mathcal{A}_{luf}} \Rightarrow \text{true}] \\ &\leq \Pr[\text{Bad}]. \end{aligned}$$

Then, through some simple algebraic manipulation and substituting for the LUF advantage of $\hat{\mathcal{F}}$ we get

$$\text{Adv}_{\hat{\mathcal{F}}}^{\text{luf}}(\mathcal{A}_{luf}, \mathcal{L}_{\hat{\mathcal{F}}}) \leq 2 \Pr[\text{Bad}] + \left(2 \Pr[\boxed{\text{G}}^{\mathcal{A}_{luf}} \Rightarrow \text{true}] - 1 \right). \quad (1)$$

<p>Game G $\boxed{\mathsf{G}}$</p> <hr/> <p>procedure Initialize</p> <p>win \leftarrow false; $K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure F (X)</p> <p>$Z \leftarrow \mathcal{H}(X)$</p> <p>$\Omega \leftarrow_{\cup} (X, Z)$</p> <p>if $\exists (\bar{X}, \bar{Z}) \in \Omega$ st $X \neq \bar{X} \wedge Z = \bar{Z}$</p> <p style="padding-left: 20px;">Bad \leftarrow true</p> <p>$\mathcal{S} \leftarrow_{\cup} X$</p> <p>$y \leftarrow \mathcal{F}'(K, Z)$</p> <p>return y</p> <hr/> <p>procedure Finalize</p> <p>return win $\wedge \neg \text{Bad}$</p>	<p>procedure Lkg($X, (L_{F'}, L_H)$)</p> <hr/> <p>$\Lambda_{F'} \leftarrow L_{F'}(K, Z)$</p> <p>$\Lambda_H \leftarrow L_H(X)$</p> <p>return $\Lambda = (\Lambda_{F'}, \Lambda_H)$</p> <hr/> <p>procedure Guess(X, y')</p> <hr/> <p>$Z \leftarrow \mathcal{H}(X)$</p> <p>$\Omega \leftarrow_{\cup} (X, Z)$</p> <p>if $\exists (\bar{X}, \bar{Z}) \in \Omega$ st $X \neq \bar{X} \wedge Z = \bar{Z}$</p> <p style="padding-left: 20px;">Bad \leftarrow true</p> <p>$y \leftarrow \mathcal{F}'(K, Z)$</p> <p>if $X \notin \mathcal{S} \wedge y = y'$</p> <p style="padding-left: 20px;">win \leftarrow true</p> <p>return $(y = y')$</p>
--	---

Fig. 18: Games G and $\boxed{\mathsf{G}}$ used to prove Theorem 4. Game $\boxed{\mathsf{G}}$ includes the boxed code whereas G does not.

We now bound the two terms on the right-hand side of the above inequality. We start by bounding the probability that \mathcal{A}_{luf} sets the flag **Bad** to **true** in game G . Towards this end we construct from \mathcal{A}_{luf} another adversary \mathcal{A}_{hash} for breaking the collision resistance of \mathcal{H} . The adversary \mathcal{A}_{hash} samples a key for \mathcal{F}' , runs \mathcal{A}_{luf} and provides it with a simulation of G . Throughout the game it maintains the set Ω and constantly checks for colliding pairs in this set. As soon as it finds a collision, i.e. **Bad** is set to **true**, it halts and outputs this colliding pair. Clearly

$$\Pr[\text{Bad}] \leq \mathbf{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}) . \quad (2)$$

Now, to bound the second term of inequality (1), we construct from \mathcal{A}_{luf} another adversary \mathcal{A}'_{luf} against the LUF security of \mathcal{F}' . Specifically, \mathcal{A}'_{luf} runs \mathcal{A}_{luf} and provides it with a simulation of the $\boxed{\mathsf{G}}$ game. For any query that \mathcal{A}_{luf} makes to any of its oracles which includes the value X , \mathcal{A}'_{luf} simply replaces this value with $\mathcal{H}(X)$ and forwards this query to its corresponding oracle and then passes the reply back to \mathcal{A}_{luf} . In the case of a leakage query, \mathcal{A}'_{luf} additionally evaluates $L_H(X)$ locally and combines it with the response it obtains from its oracle before forwarding it back to \mathcal{A}_{luf} . By construction, the leakage queries of \mathcal{A}'_{luf} will be in the permitted leakage set as long as \mathcal{A}_{luf} 's queries are. Moreover, whenever \mathcal{A}_{luf} wins the $\boxed{\mathsf{G}}$ game then \mathcal{A}'_{luf} also wins the LUF game with respect to \mathcal{F}' . This is because if \mathcal{A}_{luf} wins $\boxed{\mathsf{G}}$ then **Bad** must not have occurred, thus if the winning query (X, y') is such that $X \notin \mathcal{S}$ then $\mathcal{H}(X)$ is also not contained in the corresponding set in the LUF game of \mathcal{A}'_{luf} . It then follows that

$$\left(2 \Pr[\boxed{\mathsf{G}}^{\mathcal{A}_{luf}} \Rightarrow \text{true}] - 1 \right) \leq \mathbf{Adv}_{\mathcal{F}'}^{\text{luf}}(\mathcal{A}'_{luf}, \mathcal{L}_{F'}) . \quad (3)$$

Combining inequalities (1), (2), and (3) yields the desired result. \square

B.4 Proof of Theorem 5

We prove the Theorem through the two games, G and $\boxed{\mathsf{G}}$, described in Fig. 19. Namely:

- **Game G** is the LPRF game instantiated with $\hat{\mathcal{F}}$ and its associated leakage set $\mathcal{L}_{\hat{F}} = \mathcal{L}_{F'} \times \mathcal{L}_H$, and some additional bookkeeping to maintain the sets Ω_1 and Ω_2 and test for events **Bad**¹ and **Bad**².
- **Game $\boxed{\mathsf{G}}$** is identical to game G except that the **F** oracle returns the same output for any two inputs which hash to the same value, even when $d = 0$, and the winning condition is restricted to only hold if **Bad**¹ does not occur.

<p>Game \mathbb{G}</p> <hr/> <p>procedure Initialize</p> <p>$d \leftarrow \{0, 1\}; K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure LF($X, (L_{F'}, L_H)$)</p> <hr/> <p>$Z \leftarrow \mathcal{H}(X)$</p> <p>$\Omega_1 \leftarrow_{\cup} (X, Z)$</p> <p>if $\exists(\bar{X}, \bar{Z}) \in \Omega_2$ st $X \neq \bar{X} \wedge Z = \bar{Z}$</p> <p style="padding-left: 20px;">Bad¹ \leftarrow true</p> <p>$\Lambda_{F'} \leftarrow L_{F'}(K, Z)$</p> <p>$\Lambda_H \leftarrow L_H(X)$</p> <p>$y \leftarrow \mathcal{F}'(K, Z)$</p> <p>return $(y, (\Lambda_{F'}, \Lambda_H))$</p> <hr/> <p>procedure Finalize</p> <hr/> <p>return $(d' = d) \boxed{\wedge \neg \text{Bad}^1}$</p>	<p>procedure F(X)</p> <hr/> <p>$Z \leftarrow \mathcal{H}(X)$</p> <p>$\Omega_2 \leftarrow_{\cup} (X, Z)$</p> <p>if $\exists(\bar{X}, \bar{Z}) \in \Omega_1$ st $X \neq \bar{X} \wedge Z = \bar{Z}$</p> <p style="padding-left: 20px;">Bad¹ \leftarrow true</p> <p>if $\exists(\bar{X}, \bar{Z}) \in \Omega_2$ st $X \neq \bar{X} \wedge Z = \bar{Z}$</p> <p style="padding-left: 20px;">Bad² \leftarrow true</p> <p style="padding-left: 20px;">$f[X] \leftarrow f[\bar{X}]$</p> <p>if $d = 0$</p> <p style="padding-left: 20px;">if $f[X] = \perp$</p> <p style="padding-left: 40px;">$f[X] \leftarrow \{0, 1\}^t$</p> <p style="padding-left: 20px;">return $f[X]$</p> <p>else</p> <p style="padding-left: 20px;">return $\mathcal{F}'(K, Z)$</p>
--	---

Fig. 19: Games \mathbb{G} and $\boxed{\mathbb{G}}$ used to prove Theorem 5. Game $\boxed{\mathbb{G}}$ includes the boxed code whereas \mathbb{G} does not.

Moreover \mathbb{G} and $\boxed{\mathbb{G}}$ are identical until $\text{Bad} = \text{Bad}^1 \cup \text{Bad}^2$, and therefore by the game playing lemma we have that

$$\begin{aligned} \text{Adv} \left(\mathbb{G}^{\mathcal{A}_{lprf}}, \boxed{\mathbb{G}}^{\mathcal{A}_{lprf}} \right) &= \Pr[\mathbb{G}^{\mathcal{A}_{lprf}} \Rightarrow \text{true}] - \Pr[\boxed{\mathbb{G}}^{\mathcal{A}_{lprf}} \Rightarrow \text{true}] \\ &\leq \Pr[\text{Bad}]. \end{aligned}$$

Then, through some simple algebraic manipulation and substituting for the LPRF advantage of $\hat{\mathcal{F}}$ we get

$$\text{Adv}_{\hat{\mathcal{F}}}^{\text{lprf}}(\mathcal{A}_{lprf}, \mathcal{L}_{\hat{\mathcal{F}}}) \leq 2 \Pr[\text{Bad}^1 \cup \text{Bad}^2] + \left(2 \Pr[\boxed{\mathbb{G}}^{\mathcal{A}_{lprf}} \Rightarrow \text{true}] - 1 \right). \quad (4)$$

We now bound the two terms on the right-hand side of the above inequality. We start by bounding the probability that Bad occurs, i.e. \mathcal{A}_{lprf} sets Bad^1 or Bad^2 to **true** in game \mathbb{G} . Towards this end we construct from \mathcal{A}_{lprf} another adversary \mathcal{A}_{hash} for breaking the collision resistance of \mathcal{H} . The adversary \mathcal{A}_{hash} samples a key for \mathcal{F}' , runs \mathcal{A}_{lprf} and provides it with a simulation of \mathbb{G} . Throughout the game it maintains the sets Ω_1 and Ω_2 and constantly checks for the conditions that set Bad^1 or Bad^2 to **true**. When either of these conditions occur a collision has been found, at which point \mathcal{A}_{hash} halts and outputs the colliding pair. It then follows that

$$\Pr[\text{Bad}^1 \cup \text{Bad}^2] \leq \text{Adv}_{\mathcal{H}}^{\text{cr}}(\mathcal{A}_{hash}). \quad (5)$$

Now, to bound the second term of inequality (4), we construct from \mathcal{A}_{lprf} another adversary \mathcal{A}'_{lprf} against the LPRF security of \mathcal{F}' . Specifically, \mathcal{A}'_{lprf} runs \mathcal{A}_{lprf} and provides it with a simulation of the $\boxed{\mathbb{G}}$ game. For any query that \mathcal{A}_{lprf} makes to either of its oracles with the value X , \mathcal{A}'_{lprf} simply replaces this value with $\mathcal{H}(X)$ and forwards this query to its corresponding oracle and then passes the reply back to \mathcal{A}_{lprf} . In the case of a leakage query, \mathcal{A}'_{lprf} additionally evaluates $L_H(X)$ locally and combines it with the response it obtains from its oracle before forwarding it back to \mathcal{A}_{lprf} . Now, since in $\boxed{\mathbb{G}}$ any two queries to F that hash to the same value will return the same output (irrespective of the value d) \mathcal{A}'_{lprf} 's simulation is perfect. The leakage queries of \mathcal{A}'_{lprf} will be in the permitted leakage set as long as \mathcal{A}_{lprf} 's queries are. Moreover, as long as Bad^1 does not occur the queries of \mathcal{A}'_{lprf} to its F and LF oracles will be disjoint as long as \mathcal{A}_{lprf} 's queries are disjoint. Thus whenever \mathcal{A}_{lprf} wins game $\boxed{\mathbb{G}}$, \mathcal{A}'_{lprf} will also win

the LPRF game with respect to \mathcal{F}' , because if \mathcal{A}_{lprf} wins $\boxed{\mathbb{G}}$ then Bad^1 must not have occurred. It then follows that

$$\left(2 \Pr[\boxed{\mathbb{G}}^{\mathcal{A}_{lprf}} \Rightarrow \text{true}] - 1\right) \leq \text{Adv}_{\mathcal{F}'}^{\text{LPRF}}(\mathcal{A}'_{lprf}, \mathcal{L}_{F'}) . \quad (6)$$

Combining inequalities (4), (5), and (6) yields the desired result. \square

C Security Proofs for Section 5

C.1 Proof of Theorem 7

<p>Game \mathbb{G}_1</p> <hr/> <p>Initialize</p> <p>$K \leftarrow \{0, 1\}^k$</p> <hr/> <p>procedure $\rho(Z)$</p> <hr/> <p>if $p[Z] = \perp$ $p[Z] \leftarrow \{0, 1\}^n$ return $p[Z]$</p> <hr/> <p>procedure $\mathbb{F}(X)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$ if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$ $S_1 \leftarrow p[Y_0]$ for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$ return S_{l+1}</p>	<p>procedure $\text{LF}(X, L_0, \dots, L_l)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$ $A \leftarrow L_0(Y_0)$ if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$ $S_1 \leftarrow p[Y_0]$ for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ $A \leftarrow A \parallel L_i(Y_i)$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$ return (S_{l+1}, A)</p>
--	--

Fig. 20: Game \mathbb{G}_1 used in the proof of Theorem 7.

We prove the theorem for the case where SLFUNC takes inputs of fixed length ($l \cdot rr$) and returns outputs of length $t = n$. The general case then follows by means of a simple reduction which truncates the output of SLFUNC to t bits. The proof proceeds by considering the following sequence of games.

\mathbb{G}_1 This game is described in Fig. 20. It is the LPRF game instantiated with SLFUNC and the bit b set to 1, i.e. \mathbb{F} always returns real evaluations of SLFUNC. Furthermore the adversary is also given oracle access to the random transformation ρ which SLFUNC depends on. This random transformation is sampled lazily across all oracles and the corresponding values are stored in a global array $p[\cdot]$. All entries of $p[\cdot]$ are initialised to \perp . Thus,

$$\Pr[\mathcal{A}^{\text{LPRF}} \Rightarrow 1 \mid b = 1] = \Pr[\mathcal{A}^{\mathbb{G}_1} \Rightarrow 1]. \quad (7)$$

$\boxed{\mathbb{G}_2}$ In this game, described in Fig. 21, we introduce three flags Bad_1^1 , Bad_2^2 , and Bad_3^3 and evaluate the last round of the \mathbb{F} oracle in a different but equivalent way. Here the transformation used in the last round of \mathbb{F} is considered to be special and its values are stored in $p^*[\cdot]$ instead. However, as we lazy-sample $p[\cdot]$ and $p^*[\cdot]$ we ensure that they are aligned on common input values. When one array is to be sampled on an input for which the other array has already been defined, then one of the three

Bad_2 flags will be set and the output for that array is copied from the one that has already been defined. Note that in comparison to G_1 the sampling of the transformation's outputs is unaltered, we are merely storing them across two arrays instead of one. Thus, the functionality of the game and the distribution of its random variables is unchanged, and hence

$$\Pr[\mathcal{A}^{\boxed{\mathsf{G}_2}} \Rightarrow 1] = \Pr[\mathcal{A}^{\mathsf{G}_1} \Rightarrow 1]. \quad (8)$$

G_2 This game is the game described in Fig. 21 without the boxed code. By removing the boxed code we no longer maintain consistency between $p[\cdot]$ and $p^*[\cdot]$ on common inputs. In particular, the last round of F is now evaluated using an independent random transformation. Clearly, G_2 and $\boxed{\mathsf{G}_2}$ are identical until the event $\text{Bad}_2^1 \cup \text{Bad}_2^2 \cup \text{Bad}_2^3$ occurs. Then, by the fundamental lemma of game-playing it follows that:

$$\Pr[\mathcal{A}^{\mathsf{G}_2} \Rightarrow 1] - \Pr[\mathcal{A}^{\boxed{\mathsf{G}_2}} \Rightarrow 1] \leq \Pr[\text{Bad}_2^1 \cup \text{Bad}_2^2 \cup \text{Bad}_2^3].$$

Moreover, for events C and E such that $\text{Bad}_2^3 \subseteq C$ and $\text{Bad}_2^1 \subseteq E$, we have that

$$\begin{aligned} &\leq \Pr[E \cup \text{Bad}_2^2 \cup C], \\ &\leq \Pr[C] + \Pr[E \mid \overline{C}] + \Pr[\text{Bad}_2^2 \mid \overline{C}, \overline{E}]. \end{aligned} \quad (9)$$

We now define event C appropriately and bound the above three probabilities. Let q_ρ , q_{LF} , and q_{F} denote the number of queries \mathcal{A} makes to its oracles ρ , LF , and F , respectively. We can further assume, without loss of generality, that \mathcal{A} never repeats a query to any of its oracles.

If Bad_2^3 is set \mathcal{A} must have made two distinct queries, X to F and X' to LF , such that $\text{SLFUNC}(K, X) = \text{SLFUNC}(K, X')$. Then, note that $\text{SLFUNC}(K, X)$ can be viewed as the sponge-based hash function from Section 5.4 evaluated on input $0 \parallel X$ with an initial state of $K \parallel IV$ (instead of 0^n) and capacity $n - rr$. We will use this observation to bound the probability of event C (and Bad_2^3) occurring, but let us first define event C . Let $\mathcal{H}^{K \parallel IV}$ be the variable-input-length hash function just described, then C is the event that one of the following conditions is satisfied.

- (i) \mathcal{A} makes queries X and X' across F and LF such that $\mathcal{H}^{K \parallel IV}(0 \parallel X) = \mathcal{H}^{K \parallel IV}(0 \parallel X')$ and $X \neq X'$.
- (ii) \mathcal{A} makes queries X and X' across F and LF such that $\mathcal{H}^{K \parallel IV}(0 \parallel X) = \mathcal{H}^{K \parallel IV}(0 \parallel X'')$ for some X'' that is a prefix of X' and $|X''| < |X'|$.
- (iii) \mathcal{A} makes a query X to F or LF such that $\mathcal{H}^{K \parallel IV}(0 \parallel X') = K \parallel IV$ for some X' that is a prefix of X .
- (iv) \mathcal{A} makes queries X and X' across F and LF such that $\mathcal{H}^{K \parallel IV}(0 \parallel X'') = \mathcal{H}^{K \parallel IV}(0 \parallel X''')$, where X'' is a prefix of X , X''' is a prefix of X' , $X'' \neq \varepsilon \neq X'''$, and $X'' \neq X'''$.

Note that $\text{Bad}_2^3 \subseteq C$ is implied by the first condition. Now for any adversary \mathcal{A} that causes C , we can construct an adversary \mathcal{B} that finds a collision in $\mathcal{H}^{K \parallel IV}$. Specifically, \mathcal{B} knows K and uses this together with its access to ρ to simulate game G_2 to \mathcal{A} while constantly checking if event C has occurred. Note that any of the sub-cases in C will directly enable \mathcal{B} to produce a collision on $\mathcal{H}^{K \parallel IV}$. We can therefore use the security bound derived in Section 5.4 for \mathcal{H} since it applies to any initial state that \mathcal{H} is assigned, as long as it is fixed at the start of the game. Thus, applying inequality (49) from Appendix C.4 while setting $w = n$ and $c' = n - rr$ yields

$$\Pr[\text{Bad}_2^3] \leq \Pr[C] \leq \frac{q_T(q_T + 2)}{2^{n-rr}}, \quad (10)$$

where $q_T = (l + 1)(q_{\text{LF}} + q_{\text{F}}) + q_\rho$.

We now define event E and bound $\Pr[E \mid \overline{C}]$. At any point in time, let \mathcal{V}_{F} and \mathcal{V}_{LF} denote the set unions of intermediate values $\{Y_1, \dots, Y_l\}$ evaluated in F and LF respectively over all queries made up to that point. Further let E_j be the event that the j^{th} query Z that \mathcal{A} makes to ρ is such that $Z = Y_0$ or $Z \oplus A \parallel 0^{n-rr} \in \mathcal{V}_{\text{F}}$ for some $A \in \{0, 1\}^{rr}$. Then E represents the union of events E_j for $j \in \{1, \dots, q_\rho\}$. Note that $\text{Bad}_2^1 \subseteq E$ as required, and we also have that

$$\begin{aligned} \Pr[E \mid \overline{C}] &\leq \Pr[E_1 \cup E_2 \cup \dots \cup E_{q_\rho} \mid \overline{C}] \\ &= \Pr[E_1 \mid \overline{C}] + \Pr[E_2 \mid \overline{E}_1, \overline{C}] + \dots \\ &\quad + \Pr[E_{q_\rho} \mid \overline{E}_1, \dots, \overline{E}_{q_\rho-1}, \overline{C}]. \end{aligned} \quad (11)$$

<p>Games $\boxed{\mathbb{G}_2}, \mathbb{G}_2$</p> <hr/> <p>Initialize</p> <hr/> <p>$K \leftarrow \{0, 1\}^k$</p> <hr/> <p>procedure $\rho(Z)$</p> <hr/> <p>if $p[Z] = \perp$ $p[Z] \leftarrow \{0, 1\}^n$</p> <p>if $Z \in \text{inset}(p^*)$ $\text{Bad}_2^1 \leftarrow \text{true}$ $p[Z] \leftarrow p^*[Z]$</p> <p>return $p[Z]$</p> <hr/> <p>procedure $\mathbb{F}(X)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$</p> <p>if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$ $S_1 \leftarrow p[Y_0]$</p> <p>for i in $\{1, \dots, l-1\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$</p> <p>$Y_l \leftarrow (\bar{S}_l \oplus X_l) \parallel \hat{S}_l$ if $p^*[Y_l] = \perp$ $p^*[Y_l] \leftarrow \{0, 1\}^n$</p> <p>if $Y_l \in \text{inset}(p)$ $\text{Bad}_2^2 \leftarrow \text{true}$ $p^*[Y_l] \leftarrow p[Y_l]$</p> <p>$S_{l+1} \leftarrow p^*[Y_l]$ return S_{l+1}</p>	<hr/> <p>procedure $\text{LF}(X, L_0, \dots, L_l)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$ $A \leftarrow L_0(Y_0)$</p> <p>if $S_1 = \perp$ $S_1 \leftarrow \{0, 1\}^n$</p> <p>for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ $A \leftarrow A \parallel L_i(Y_i)$</p> <p>if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$</p> <p>if $i = l \wedge Y_i \in \text{inset}(p^*)$ $\text{Bad}_2^3 \leftarrow \text{true}$ $p[Y_i] \leftarrow p^*[Y_i]$</p> <p>$S_{i+1} \leftarrow p[Y_i]$ return (S_{l+1}, A)</p>
---	--

Fig. 21: Games $\boxed{\mathbb{G}_2}$ and \mathbb{G}_2 used in the proof of Theorem 7. Game $\boxed{\mathbb{G}_2}$ includes the boxed code whereas \mathbb{G}_2 does not.

Now as long as C does not occur no two queries will result in values Y_i and $Y_{i'}$ such that $Y_i = Y_{i'}$ and $i \neq i'$. Thus no state variable $p[Y_i]$ occurs in more than one position (when considering both \mathbf{F} and \mathbf{LF}), meaning that no state variable is subject to more than one leakage function and that no internal state variable is ever exposed as the output of \mathbf{F} or \mathbf{LF} . Moreover, as long as E does not occur, the outputs of ρ will never expose the internal state variables of \mathbf{F} and \mathbf{LF} . Then, for $i \geq 1$, guessing a value $Y_i \in \mathcal{V}_{\mathbf{F}}$ is tantamount to guessing the corresponding S_i , since $Y_i = S_i \oplus (X_i \parallel 0^{n-rr})$ and X_i is known to the adversary. Then by the union bound and the above observations it follows that

$$\begin{aligned} \Pr[E_j | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}] &\leq \sum_{Y_i \in \mathcal{V}_{\mathbf{F}}} 2^{-\mathsf{H}_{\infty}(Y_i | A, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C})} \\ &+ 2^{-\mathsf{H}_{\infty}(K | A, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C})} \\ &= \sum_{Y_i \in \mathcal{V}_{\mathbf{F}}} 2^{-\mathsf{H}_{\infty}(S_i | A, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C})} \\ &+ 2^{-\mathsf{H}_{\infty}(K | A, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C})}. \end{aligned} \quad (12)$$

Although \mathbf{F} does not leak we have to take into account that the same state variable may result internally in a distinct query to \mathbf{LF} and thereby be subject to leakage. With respect to the key the adversary only learns $L_0(K \parallel IV)$ from the leakage. However, for a given S_i the adversary may obtain (at most) the leakage $L_i(S_i \oplus (X_i \parallel 0^{n-rr}))$ for each possible value of X_i . If we consider the aggregate leakage for a given S_i over all possible values of X_i , i.e. $L_i(S_i \oplus 0^{rr} \parallel 0^{n-rr}) \parallel \dots \parallel L_i(S_i \oplus 1^{rr} \parallel 0^{n-rr})$, the support of this aggregate random variable is bounded by $2^{\lambda 2^{rr}}$. Then by applying Lemma 1 we have that

$$\mathsf{H}_{\infty}(K | A_0, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq \mathsf{H}_{\infty}(K | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) - \lambda 2^{rr}$$

and

$$\mathsf{H}_{\infty}(S_i | A_i, \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq \mathsf{H}_{\infty}(S_i | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) - \lambda 2^{rr}. \quad (13)$$

It now remains to bound the min-entropies of K and S_i conditioned on E_1, \dots, E_{j-1} and C not occurring. For each event E_j , conditioning on it not occurring excludes one value for the variable K and 2^{rr} possible values for the variable S_i . On the other hand, conditioning on \bar{C} excludes (at most) a further $(l+1)(q_{\mathbf{F}} + q_{\mathbf{LF}})$ from the possible values that K and S_i may take. Then by constraining the adversary's queries such that $q_{\rho} + (l+1)(q_{\mathbf{F}} + q_{\mathbf{LF}}) \leq 2^{k-1}$ and $(2^{rr})q_{\rho} + (l+1)(q_{\mathbf{F}} + q_{\mathbf{LF}}) \leq 2^{n-1}$, we obtain, for all possible j , the following bounds

$$\mathsf{H}_{\infty}(K | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq k - 1$$

and

$$\mathsf{H}_{\infty}(S_i | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq n - 1. \quad (14)$$

Combining (12), (13), and (14) and applying the bound $|\mathcal{V}_{\mathbf{F}}| \leq l q_{\mathbf{F}}$ yields

$$\Pr[E_j | \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}] \leq \frac{1}{2^{k-\lambda 2^{rr}-1}} + \frac{l q_{\mathbf{F}}}{2^{n-\lambda 2^{rr}-1}}.$$

We then substitute the above in (11) to obtain

$$\Pr[\mathbf{Bad}_2^1 | \bar{C}] \leq \frac{q_{\rho}}{2^{k-\lambda 2^{rr}-1}} + \frac{l q_{\mathbf{F}} q_{\rho}}{2^{n-\lambda 2^{rr}-1}}. \quad (15)$$

Finally, we bound the last term of inequality (9). The flag \mathbf{Bad}_2^2 is set if \mathcal{A} makes some query to \mathbf{F} such that the corresponding Y_l value was already contained in $\mathsf{inset}(p)$. Now if it were the case that Y_l was already in $\mathcal{V}_{\mathbf{F}} \cup \mathcal{V}_{\mathbf{LF}}$ when \mathbf{Bad}_2^2 occurs it would mean that event C has also occurred. Thus, conditioned on C not occurring, the only way for \mathbf{Bad}_2^2 to be set is if Y_l was already queried to ρ by \mathcal{A} . Let X^* represent \mathcal{A} 's query to \mathbf{F} that sets \mathbf{Bad}_2^2 , resulting in the corresponding values $Y_i^* = S_i^* \oplus (X_i^* \parallel 0^{n-rr})$ for $1 \leq i \leq l$, such that Y_l^* was already queried to ρ . Then, conditioned on E not occurring, the variable S_l^* cannot be sampled during a query to ρ . That is, $p[Y_{l-1}^*]$ must be set during some query to \mathbf{F} or \mathbf{LF} after that Y_l^* was queried to ρ . Now let F be the event that \mathcal{A}

makes a query to **F** or **LF** resulting in S_i^* such that for some $A \in \{0, 1\}^{rr}$ the value $S_i^* \oplus (A \parallel 0^{n-rr})$ was already queried to ρ . It then follows that

$$\Pr[\text{Bad}_2^2 \mid \overline{C}, \overline{E}] \leq \Pr[F \mid \overline{C}, \overline{E}]. \quad (16)$$

We now bound $\Pr[F \mid \overline{C}, \overline{E}]$. Let F_j denote the event that F occurs at the j^{th} query that \mathcal{A} makes to **F** or **LF**. Then, by the union bound we have that

$$\Pr[F \mid \overline{C}] \leq \Pr[F_1 \mid \overline{C}, \overline{E}] + \Pr[F_2 \mid \overline{C}, \overline{E}] + \dots + \Pr[F_{q_{\mathbf{F}}+q_{\mathbf{LF}}} \mid \overline{C}, \overline{E}].$$

Then for each query there are at most $(2^{rr})q_\rho$ values that S_i^* can take to set F out of $2^n - (l+1)(q_{\mathbf{F}} + q_{\mathbf{LF}}) - (2^{rr})q_\rho$ (due to conditioning on \overline{C} and \overline{E}). Thus

$$\begin{aligned} \Pr[F \mid \overline{C}] &\leq \sum_{j=1}^{q_{\mathbf{F}}+q_{\mathbf{LF}}} \frac{(2^{rr})q_\rho}{2^n - (l+1)(q_{\mathbf{F}} + q_{\mathbf{LF}}) - (2^{rr})q_\rho} \\ &\leq \frac{2^{rr}(q_{\mathbf{F}} + q_{\mathbf{LF}})q_\rho}{2^{n-1}} \\ &= \frac{(q_{\mathbf{F}} + q_{\mathbf{LF}})q_\rho}{2^{n-rr-1}}. \end{aligned} \quad (17)$$

Then, combining inequalities (9), (10), (15), and (17) we obtain

$$\Pr[\mathcal{A}^{\mathbf{G}_2} \Rightarrow 1] - \Pr[\mathcal{A}^{\boxed{\mathbf{G}_2}} \Rightarrow 1] \leq \frac{q_T(q_T + 2)}{2^{n-rr}} + \frac{q_\rho}{2^{k-\lambda 2^{rr}-1}} + \frac{l q_{\mathbf{F}} q_\rho}{2^{n-\lambda 2^{rr}-1}} + \frac{(q_{\mathbf{F}} + q_{\mathbf{LF}})q_\rho}{2^{n-rr-1}}, \quad (18)$$

where $q_T = (l+1)(q_{\mathbf{LF}} + q_{\mathbf{F}}) + q_\rho$.

$\boxed{\mathbf{G}_3}$ This game is described in Fig. 22. Here we remove the three flags Bad_2^1 , Bad_2^2 , and Bad_2^3 , we introduce flag Bad_3 , and change how the last round of **F** is evaluated. Notably, in tandem with $p^*[\cdot]$ we maintain an additional array $f[\cdot]$, where the latter is indexed by strings of size l instead of n . Then in every evaluation of **F** we sample a random string, store it in $f[X]$, and then copy this value to $p^*[Y_l]$. However if $p^*[Y_l]$ was already defined, then Bad_3 is set and the random string in $f[X]$ is replaced with $p^*[Y_l]$. This ensures that no entry in $p^*[\cdot]$ is ever overwritten. Thus the array $p^*[\cdot]$ is sampled in an equivalent manner as in \mathbf{G}_2 . Moreover, the functionality of \mathbf{G}_2 is not dependent on the flags Bad_2^1 , Bad_2^2 , and Bad_2^3 and their removal is merely a cosmetic alteration. Hence $\boxed{\mathbf{G}_3}$ is identical to \mathbf{G}_2 and it follows that

$$\Pr[\mathcal{A}^{\boxed{\mathbf{G}_3}} \Rightarrow 1] = \Pr[\mathcal{A}^{\mathbf{G}_2} \Rightarrow 1]. \quad (19)$$

\mathbf{G}_3 This game is the game described in Fig. 22 without the boxed code. By removing the boxed code we now have that $\mathbf{F}(X) = f[X]$, and therefore $\mathbf{F}(X)$ behaves as a random function. Therefore

$$\Pr[\mathcal{A}^{\text{LPRF}} \Rightarrow 1 \mid b = 0] = \Pr[\mathcal{A}^{\mathbf{G}_3} \Rightarrow 1]. \quad (20)$$

Furthermore, \mathbf{G}_3 and $\boxed{\mathbf{G}_3}$ are identical until Bad_3 occurs. Now, if Bad_3 is set then \mathcal{A} must have queried X and X' (where $X \neq X'$ since \mathcal{A} is assumed to not repeat queries to **F**) such that they result in the same value for Y_l . As before we can view this as a collision on $\mathcal{H}^{K \parallel IV}$ with l rounds. Thus, applying the fundamental lemma of game-playing and using once again inequality (49) from Appendix C.4, we have that

$$\begin{aligned} \Pr[\mathcal{A}^{\mathbf{G}_3} \Rightarrow 1] - \Pr[\mathcal{A}^{\boxed{\mathbf{G}_3}} \Rightarrow 1] &\leq \Pr[\text{Bad}_3] \\ &\leq \frac{q'_T(q'_T + 2)}{2^{n-rr}}. \end{aligned} \quad (21)$$

where $q'_T = l(q_{\mathbf{LF}} + q_{\mathbf{F}}) + q_\rho$.

Combining inequalities (7), (8), (18), (19), (21), and (20) yields the desired result. \square

<p>Games $\boxed{\mathbb{G}_3}, \mathbb{G}_3$</p> <hr/> <p>Initialize</p> <hr/> <p>$K \leftarrow \{0, 1\}^k$</p> <hr/> <p>procedure $\rho(Z)$</p> <hr/> <p>if $p[Z] = \perp$ $p[Z] \leftarrow \{0, 1\}^n$ return $p[Z]$</p> <hr/> <p>procedure $\mathbb{F}(X)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$ if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$ $S_1 \leftarrow p[Y_0]$ for i in $\{1, \dots, l-1\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$ $Y_l \leftarrow (\bar{S}_l \oplus X_l) \parallel \hat{S}_l$ $f[X] \leftarrow \{0, 1\}^n$ if $p^*[Y_l] \neq \perp$ $\mathbf{Bad}_3 \leftarrow \mathbf{true}$ <div style="border: 1px solid black; padding: 2px; display: inline-block;">$f[X] \leftarrow p^*[Y_l]$</div> $p^*[Y_l] \leftarrow f[X]$ $S_{l+1} \leftarrow p^*[Y_l]$ return S_{l+1}</p>	<p>procedure $\mathbb{LF}(X, L_0, \dots, L_l)$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$ $Y_0 \leftarrow K \parallel IV$ $A \leftarrow L_0(Y_0)$ if $S_1 = \perp$ $S_1 \leftarrow \{0, 1\}^n$ for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ $A \leftarrow A \parallel L_i(Y_i)$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$ return (S_{l+1}, A)</p>
--	--

Fig. 22: Games $\boxed{\mathbb{G}_3}$ and \mathbb{G}_3 used in the proof of Theorem 7. Game $\boxed{\mathbb{G}_3}$ includes the boxed code whereas \mathbb{G}_3 does not.

C.2 Proof of Theorem 8

We assume that SLFUNC takes inputs of fixed length ($l \cdot rr$) and returns outputs of length t . Now consider the game described in Fig. 23. It is the LUF game instantiated with SLFUNC where the adversary is also given oracle access to the random transformation ρ which SLFUNC depends on. This random transformation is sampled lazily across all oracles and the corresponding values are stored in a global array $p[\cdot]$. All entries of $p[\cdot]$ are initialised to \perp . Thus,

$$\Pr[\text{LUF}^{\mathcal{A}} \Rightarrow \text{true}] = \Pr[\text{G}^{\mathcal{A}} \Rightarrow \text{true}]. \quad (22)$$

At any point in time, let \mathcal{V}_{Lkg} be the set union of intermediate values $\{Y_1, \dots, Y_l\}$ evaluated in Lkg over all queries made by \mathcal{A} up to that point. Similarly, let \mathcal{U}_{Lkg} be the set of values S_{l+1} evaluated in Lkg over all queries made by \mathcal{A} . Further let W be the event that \mathcal{A} wins and B be the event that \mathcal{A} makes a query to ρ that is contained in $\mathcal{V}_{\text{Lkg}} \cup \{Y_0\}$. Then for any event C it follows that

$$\begin{aligned} \Pr[\text{G}^{\mathcal{A}} \Rightarrow \text{true}] &\leq \Pr[W \cup B \cup C] \\ &\leq \Pr[C] + \Pr[B \mid \overline{C}] + \Pr[W \mid \overline{C}, \overline{B}]. \end{aligned} \quad (23)$$

Recall that $\text{SLFUNC}(K, X)$ can be viewed as the sponge-based hash function from Section 5.4 evaluated on input $0 \parallel X$ with an initial state of $K \parallel IV$ (instead of 0^n) and capacity $n - rr$. Let $\mathcal{H}^{K \parallel IV}$ denote the variable-input-length hash function just described and C be the event that one of the following conditions is satisfied.

- (i) \mathcal{A} makes queries X and X' across F, Lkg, and Guess such that $\mathcal{H}^{K \parallel IV}(0 \parallel X) = \mathcal{H}^{K \parallel IV}(0 \parallel X')$ and $X \neq X'$.
- (ii) \mathcal{A} makes queries X and X' across F, Lkg, and Guess such that $\mathcal{H}^{K \parallel IV}(0 \parallel X) = \mathcal{H}^{K \parallel IV}(0 \parallel X'')$ for some X'' that is a prefix of X' and $|X''| < |X'|$.
- (iii) \mathcal{A} makes a query X to F, Lkg, or Guess such that $\mathcal{H}^{K \parallel IV}(0 \parallel X') = K \parallel IV$ for some X' that is a prefix of X .
- (iv) \mathcal{A} makes queries X and X' across F, Lkg, and Guess such that $\mathcal{H}^{K \parallel IV}(0 \parallel X'') = \mathcal{H}^{K \parallel IV}(0 \parallel X''')$, where X'' is a prefix of X , X''' is a prefix of X' , $X'' \neq \varepsilon \neq X'''$, and $X'' \neq X'''$.

Now for any adversary \mathcal{A} that causes C , we can construct an adversary \mathcal{B} that finds a collision in $\mathcal{H}^{K \parallel IV}$. Specifically, \mathcal{B} knows K and uses this together with its access to ρ to simulate game G_1 to \mathcal{A} while constantly checking if event C has occurred. Note that any of the sub-cases in C will directly enable \mathcal{B} to produce a collision on $\mathcal{H}^{K \parallel IV}$. We can therefore use the security bound derived in Section 5.4 for \mathcal{H} since it applies to any initial state that \mathcal{H} is assigned, as long as it is fixed at the start of the game. Thus, applying inequality (49) from Appendix C.4 and setting $w = n$, $c' = n - rr$, and $q_T = (l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}}) + q_{\rho}$, yields

$$\Pr[C] \leq \frac{q_T(q_T + 2)}{2^{n-rr}}. \quad (24)$$

We now bound $\Pr[B \mid \overline{C}]$. Let B_j be the event that B occurs on the j^{th} query that \mathcal{A} makes to ρ . Then we have that

$$\begin{aligned} \Pr[B \mid \overline{C}] &\leq \Pr[B_1 \cup B_2 \cup \dots \cup B_{q_{\rho}} \mid \overline{C}] \\ &= \Pr[B_1 \mid \overline{C}] + \Pr[B_2 \mid \overline{B}_1, \overline{C}] + \dots \\ &\quad + \Pr[B_{q_{\rho}} \mid \overline{B}_1, \dots, \overline{B}_{q_{\rho}-1}, \overline{C}]. \end{aligned} \quad (25)$$

Now as long as C does not occur no two queries will result in values Y_i and $Y_{i'}$ such that $Y_i = Y_{i'}$ and $i \neq i'$. Thus no state variable $p[Y_i]$ occurs in more than one position (across F, Lkg, and Guess), meaning that no state variable is subject to more than one leakage function and that no internal state variable is ever exposed as the output of F. Moreover, as long as B does not occur, the outputs of ρ will never expose any of the state variables contained in $\mathcal{V}_{\text{Lkg}} \cup \mathcal{U}_{\text{Lkg}} \cup \{Y_0\}$. Then, for $i \geq 1$, guessing a value $Y_i \in \mathcal{V}_{\text{Lkg}}$ is tantamount to guessing the corresponding S_i , since $Y_i = S_i + (X_i \parallel 0^{n-rr})$ and X_i is known to the adversary. Then, by the union bound and the above observations, for all $j \in \{1, \dots, q_{\rho}\}$, we have that

<p>Game G</p> <hr/> <p>Initialize</p> <p>win \leftarrow false; $K \leftarrow \mathcal{K}$</p> <p>return</p> <hr/> <p>procedure $\rho(Z)$</p> <hr/> <p>if $p[Z] = \perp$ $p[Z] \leftarrow \{0, 1\}^n$</p> <p>return $p[Z]$</p> <hr/> <p>procedure $\text{Guess}(X, y')$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$</p> <p>$Y_0 \leftarrow K \parallel IV$</p> <p>if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$</p> <p>$S_1 \leftarrow p[Y_0]$</p> <p>for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$</p> <p>$y \leftarrow \lfloor S_{l+1} \rfloor_t$</p> <p>if $X \notin \mathcal{S} \wedge y = y'$ win \leftarrow true</p> <p>return $(y = y')$</p> <hr/> <p>procedure Finalize</p> <hr/> <p>return win</p>	<p>procedure $F(X)$</p> <hr/> <p>$\mathcal{S} \leftarrow_{\cup} X$</p> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$</p> <p>$Y_0 \leftarrow K \parallel IV$</p> <p>if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$</p> <p>$S_1 \leftarrow p[Y_0]$</p> <p>for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$</p> <p>return $\lfloor S_{l+1} \rfloor_t$</p> <hr/> <p>procedure $\text{Lkg}(X, L_0, \dots, L_{l+1})$</p> <hr/> <p>parse X as $X_1 \parallel \dots \parallel X_l \forall i X_i = rr$</p> <p>$Y_0 \leftarrow K \parallel IV$</p> <p>$A \leftarrow L_0(Y_0)$</p> <p>if $p[Y_0] = \perp$ $p[Y_0] \leftarrow \{0, 1\}^n$</p> <p>$S_1 \leftarrow p[Y_0]$</p> <p>for i in $\{1, \dots, l\}$ $Y_i \leftarrow (\bar{S}_i \oplus X_i) \parallel \hat{S}_i$ $A \leftarrow A \parallel L_i(Y_i)$ if $p[Y_i] = \perp$ $p[Y_i] \leftarrow \{0, 1\}^n$ $S_{i+1} \leftarrow p[Y_i]$</p> <p>$A \leftarrow A \parallel L_{l+1}(S_{l+1})$</p> <p>return (A)</p>
---	---

Fig. 23: Game G used in the proof of Theorem 8.

$$\begin{aligned}
\Pr[B_j \mid \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}] &\leq \sum_{Y_i \in \mathcal{V}_{\text{Lkg}}} 2^{-H_{\infty}(Y_i \mid \Lambda, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C})} \\
&+ 2^{-H_{\infty}(K \mid \Lambda, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C})} \\
&= \sum_{Y_i \in \mathcal{V}_{\text{Lkg}}} 2^{-H_{\infty}(S_i \mid \Lambda, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C})} \\
&+ 2^{-H_{\infty}(K \mid \Lambda, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C})}. \tag{26}
\end{aligned}$$

Now with respect to the key the adversary only learns $L_0(K \parallel IV)$ from the leakage. However, for a given S_i the adversary may obtain (at most) the leakage $L_i(S_i \oplus X_i \parallel 0^{n-rr})$ for each possible value of X_i . If we consider the aggregate leakage for a given S_i over all possible values of X_i , i.e. $L_i(S_i \oplus 0^{rr} \parallel 0^{n-rr}) \parallel \dots \parallel L_i(S_i \oplus 1^{rr} \parallel 0^{n-rr})$, the support of this aggregate random variable is bounded by $2^{\lambda 2^{rr}}$. Then by applying Lemma 1 we have that

$$H_{\infty}(K \mid \Lambda_0, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}) \geq H_{\infty}(K \mid \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}) - \lambda 2^{rr}$$

and

$$H_{\infty}(S_i \mid \Lambda_i, \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}) \geq H_{\infty}(S_i \mid \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}) - \lambda 2^{rr}. \tag{27}$$

It now remains to bound the min-entropies of K and S_i conditioned on B_1, \dots, B_{j-1} and C not occurring. For each event B_j , conditioning on it not occurring excludes one value for the variable K and 2^{rr} possible values for the variable S_i . This is because for any query Z that the adversary makes to ρ , the values $(Z \oplus 0^{rr} \parallel 0^{n-rr}), \dots, (Z \oplus 1^{rr} \parallel 0^{n-rr})$ may all be contained \mathcal{V}_{Lkg} . On the other hand, conditioning on \bar{C} excludes (at most) a further $(l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}})$ from the possible values that K and S_i may take. Thus if we constrain the adversary's queries such that $q_{\rho} + (l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{k-1}$ and $(2^{rr})q_{\rho} + (l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{n-1}$, we obtain, for all possible j , the following bounds

$$H_{\infty}(K \mid \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq k - 1$$

and

$$H_{\infty}(S_i \mid \bar{E}_1, \dots, \bar{E}_{j-1}, \bar{C}) \geq n - 1. \quad (28)$$

Combining (26), (27), and (28) and applying the bound $|\mathcal{V}_{\text{Lkg}}| \leq lq_{\text{Lkg}}$ yields

$$\Pr[B_j \mid \bar{B}_1, \dots, \bar{B}_{j-1}, \bar{C}] \leq \frac{1}{2^{k-\lambda 2^{rr}-1}} + \frac{lq_{\text{Lkg}}}{2^{n-\lambda 2^{rr}-1}}.$$

We then substitute the above in (25) to obtain

$$\Pr[B \mid \bar{C}] \leq \frac{q_{\rho}}{2^{k-\lambda 2^{rr}-1}} + \frac{lq_{\text{Lkg}}q_{\rho}}{2^{n-\lambda 2^{rr}-1}}. \quad (29)$$

We now conclude the proof by bounding $\Pr[W \mid \bar{C}, \bar{B}]$. Let W_j be the event that W occurs on the j^{th} query that \mathcal{A} makes to **Guess**. Then we have that

$$\begin{aligned} \Pr[W \mid \bar{C}, \bar{B}] &\leq \Pr[W_1 \cup W_2 \cup \dots \cup W_{q_{\rho}} \mid \bar{C}, \bar{B}] \\ &= \Pr[W_1 \mid \bar{C}, \bar{B}] + \Pr[W_2 \mid \bar{W}_1, \bar{C}, \bar{B}] + \dots \\ &\quad + \Pr[W_{q_{\text{Guess}}} \mid \bar{W}_1, \dots, \bar{W}_{q_{\text{Guess}}-1}, \bar{C}, \bar{B}]. \end{aligned} \quad (30)$$

Now conditioning on neither of events B and C occurring, it must be that if W occurs the value of S_{l+1} corresponding to that query must be either new (freshly sampled) or contained in \mathcal{U}_{Lkg} . In either case, for all $j \in \{1, \dots, q_{\text{Guess}}\}$, we have that

$$\Pr[W_j \mid \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B}] \leq 2^{-H_{\infty}(\lfloor S_{l+1} \rfloor_t \mid A, \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B})}, \quad (31)$$

where the case where S_{l+1} is freshly sampled corresponds to $A = \varepsilon$. Thus since A is at most $L_{l+1}(S_{l+1})$, by Lemma 1 we have that

$$H_{\infty}(S_{l+1} \mid A, \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B}) \geq H_{\infty}(S_{l+1} \mid \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B}) - \lambda. \quad (32)$$

Once again conditioning on \bar{C} excludes (at most) $(l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}})$ from the possible values that S_{l+1} may take, whereas conditioning on \bar{B} does not affect S_{l+1} . Moreover, for each event W_j , conditioning on it not occurring excludes one value from the pool of values that S_{l+1} may take. To cater for this we require that the adversary's queries satisfy $q_{\text{Guess}} + (l+1)(q_{\text{F}} + q_{\text{Lkg}} + q_{\text{Guess}}) \leq 2^{n-1}$. Then from the above and (32) we have that

$$H_{\infty}(S_{l+1} \mid A, \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B}) \geq n - 1 - \lambda. \quad (33)$$

To cater for the fact that S_{l+1} is truncated to the leftmost t bits, we apply Lemma 1 once more. The probability of guessing $\lfloor S_{l+1} \rfloor_t$ can only increase when given the rightmost $n - t$ bits of S_{l+1} . Thus

$$H_{\infty}(\lfloor S_{l+1} \rfloor_t \mid A, \bar{W}_1, \dots, \bar{W}_{j-1}, \bar{C}, \bar{B}) \geq t - \lambda - 1. \quad (34)$$

Combining inequalities (30), (31), and (34) yields

$$\Pr[W \mid \bar{C}, \bar{B}] \leq \frac{q_{\text{Guess}}}{2^{t-\lambda-1}}. \quad (35)$$

The theorem then follows by combining (23), (24), (29), and (35). \square

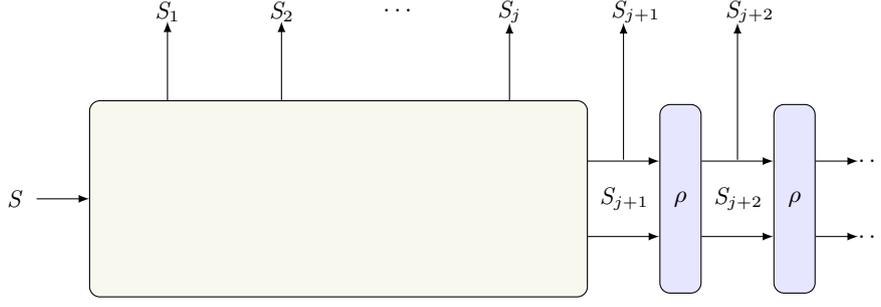


Fig. 24: A graphical illustration of the hybrid game H_j used to prove Theorem 9. The large box represents a random function with input S and output $(\bar{S}_2, \dots, \bar{S}_j, S_{j+1})$. The additional output \bar{S}_1 is simply copied from the input, i.e. $\bar{S}_1 = \bar{S}$.

C.3 Proof of Theorem 9

The proof follows through a standard hybrid argument, where we gradually replace the SPRG construction with a random function that takes as its input the seed S . In each game hop we remove one call to the random transformation and extend the output of the random function. This is illustrated in Fig. 24.

Let L_{max} be an upper bound on the output bit-length specified by the adversary in its queries to G . Then each evaluation will require at most $v - 1$ call to the random transformation, where $v = \lceil \frac{L_{max}}{r} \rceil$. Accordingly we specify the hybrid games H_0, \dots, H_{v-1} , as shown in Fig. 25. Following the hybrid argument, a further game hop to game G_0 , displayed in Fig. 26, is required to complete the proof.

In game H_j (cf. Fig. 25), the G oracle samples the states S_1, \dots, S_{j+1} ideally and independent of the transformation ρ , while the remaining states S_{j+2}, \dots, S_v are sampled by querying ρ . The output of the G oracle is the concatenation of the outer states truncated to L bits. The boxed version, that is, $\boxed{H_j}$, ensures that the state S_{j+1} is also evaluated through ρ , and thus $H_j = \boxed{H_{j+1}}$. Note that H_0 is equivalent to the PRG game instantiated with SPRG with the challenge bit b fixed to 1. Thus,

$$\Pr[\mathcal{A}^{\text{PRG}} \Rightarrow 1 \mid b = 1] = \Pr[\mathcal{A}^{H_0} \Rightarrow 1]. \quad (36)$$

Moreover, we have that

$$\Pr[\mathcal{A}^{H_0} \Rightarrow 1] - \Pr[\mathcal{A}^{H_{v-1}} \Rightarrow 1] = \sum_{j=1}^{v-1} \Pr[\mathcal{A}^{H_{j-1}} \Rightarrow 1] - \Pr[\mathcal{A}^{H_j} \Rightarrow 1],$$

and since $H_{j-1} = \boxed{H_j}$, this reduces to

$$= \sum_{j=1}^{v-1} \Pr[\mathcal{A}^{\boxed{H_j}} \Rightarrow 1] - \Pr[\mathcal{A}^{H_j} \Rightarrow 1].$$

Let Bad_j^1 , Bad_j^2 , and Bad_j^3 each denote the event that the corresponding flag in game $\boxed{H_j}$ is set. Further note that games $\boxed{H_j}$ and H_j are identical until one of these bad events occurs. Then, applying the fundamental lemma of game playing and the union bound yields

$$\leq \sum_{j=1}^{v-1} \Pr[\text{Bad}_j^1] + \Pr[\text{Bad}_j^2] + \Pr[\text{Bad}_j^3]. \quad (37)$$

The flag Bad_j^1 is set if the adversary makes a query S to ρ , such that $S = S_j[S^*]$ for some $S^* \in \mathcal{D}$. Now, for any random variable $S_j[S^*]$, only the outer part $\bar{S}_j[S^*]$ is known to the adversary (through the G oracle) and the inner c bits remain hidden. Moreover, at any point in time, there are only $|\mathcal{D}|$ such variables and $|\mathcal{D}| \leq q_G$. Hence, by the union bound

$$\Pr[\text{Bad}_j^1] \leq \sum_{i=0}^{q_\rho-1} \frac{|\mathcal{D}|}{2^{c-i}} \leq \sum_{i=0}^{q_\rho-1} \frac{q_G}{2^{c-i}} \leq \frac{q_G q_\rho}{2^c - q_\rho}. \quad (38)$$

<p>Games $\boxed{H_j}, H_j$</p> <hr/> <p>procedure $\rho(S)$</p> <hr/> <p>if $p[S] = \perp$ $p[S] \leftarrow \{0, 1\}^n$ <i>// keep random function and ρ consistent</i> if $\exists S^* \in \mathcal{D}$ s.t. $S = S_j[S^*]$ $\text{Bad}_j^1 \leftarrow \text{true}$ <div style="border: 1px solid black; padding: 2px; display: inline-block;">$p[S] \leftarrow S_{j+1}[S^*]$</div> return $p[S]$</p>	<p>procedure $G(L)$</p> <hr/> <p>$S \leftarrow \{0, 1\}^n$ $\mathcal{D} \leftarrow \cup S$ $S_1[S] \leftarrow S$ <i>// evaluate random function</i> for i in $\{2, \dots, j+1\}$ if $S_i[S] = \perp$ $S_i[S] \leftarrow \{0, 1\}^n$ <i>// keep random function and ρ consistent</i> if $S_j[S] \in \text{inset}(p)$ $\text{Bad}_j^2 \leftarrow \text{true}$ <div style="border: 1px solid black; padding: 2px; display: inline-block;">$S_{j+1}[S] \leftarrow p[S_j[S]]$</div> <i>// evaluate remaining part of SPRG</i> for i in $\{j+1, \dots, v-1\}$ if $p[S_i[S]] = \perp$ $p[S_i[S]] \leftarrow \{0, 1\}^n$ <i>// keep random function and ρ consistent</i> if $\exists S^* \in \mathcal{D}$ s.t. $S_i[S] = S_j[S^*]$ $\text{Bad}_j^3 \leftarrow \text{true}$ <div style="border: 1px solid black; padding: 2px; display: inline-block;">$p[S_i[S]] \leftarrow S_{j+1}[S^*]$</div> $S_{i+1}[S] \leftarrow p[S_i[S]]$ $R \leftarrow \bar{S}_1[S] \parallel \bar{S}_2[S] \parallel \dots \parallel \bar{S}_v[S]$ return $\lfloor R \rfloor_L$</p>
--	---

Fig. 25: The hybrid games used in the proof of Theorem 9. Game $\boxed{H_j}$ contains the boxed code whereas H_j does not.

The flag Bad_j^2 is set whenever the G oracle samples a value S_j contained in $\text{inset}(p)$. Since the G oracle samples the values S_j uniformly at random from $\{0, 1\}^n$, it follows that

$$\Pr[\text{Bad}_j^2] \leq \sum_{i=1}^{q_G} \frac{|\text{inset}(p)|}{2^n}.$$

At any point in time, $|\text{inset}(p)| \leq q_\rho + q_G(v-j-1)$, since each query to ρ and G adds at most 1 and $(v-j-1)$, new values to $\text{inset}(p)$ respectively. This leads to

$$\leq \sum_{i=1}^{q_G} \frac{q_\rho + q_G(v-j-1)}{2^n} \leq \frac{q_\rho q_G + q_G^2(v-j-1)}{2^n}. \quad (39)$$

The flag Bad_j^3 is set if, for some $i \in \{j+1, \dots, v-1\}$ and some $S^* \in \mathcal{D}$, the oracle G samples a value $S_i[S]$ equal to $S_i[S^*]$. The probability of this bad event depends solely on the number of queries that \mathcal{A} makes to its oracle G . In each query at most, $(v-j-1)$ states are sampled which can set this flag, hence

$$\Pr[\text{Bad}_j^3] \leq \sum_{i=1}^{q_G} \frac{(v-j-1)|\mathcal{D}|}{2^n}.$$

Since at any point in time $|\mathcal{D}| \leq q_G$, this reduces to

$$\leq \sum_{i=1}^{q_G} (v-j-1) \frac{q_G}{2^n} \leq \frac{(v-j-1)q_G^2}{2^n}. \quad (40)$$

<p>Game G_0</p> <p>procedure $\rho(S)$</p> <hr/> <p>if $p[S] = \perp$ $p[S] \leftarrow \{0, 1\}^n$ return $p[S]$</p>	<p>procedure $G(L)$</p> <hr/> <p>$S \leftarrow \{0, 1\}^n$ $R \leftarrow \{0, 1\}^{vr}$ if $S \in \mathcal{D}$ $\mathbf{Bad} \leftarrow \mathbf{true}$</p> <div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 5px auto;"> $R \leftarrow \bar{S}_1[S] \parallel \dots \parallel \bar{S}_v[S]$ </div> <p>$\mathcal{D} \leftarrow \cup S$ $(\bar{S}_1[S], \dots, \bar{S}_v[S]) \leftarrow R$ return $\lfloor R \rfloor_L$</p>
--	---

Fig. 26: Games G_0 and $\boxed{G_0}$ used in the proof of Theorem 9. Game $\boxed{G_0}$ contains the boxed code, G_0 does not.

Combining (37), (38), (39), and (40) yields

$$\begin{aligned}
\Pr[\mathcal{A}^{H_0} \Rightarrow 1] - \Pr[\mathcal{A}^{H_{v-1}} \Rightarrow 1] &\leq \sum_{j=1}^{v-1} \Pr[\mathbf{Bad}_j^1] + \Pr[\mathbf{Bad}_j^2] + \Pr[\mathbf{Bad}_j^3] \\
&\leq \sum_{j=1}^{v-1} \frac{q_G q_\rho}{2^c - q_\rho} + \frac{q_\rho q_G + q_G^2 (v-j-1)}{2^n} + \frac{(v-j-1) q_G^2}{2^n} \\
&\leq \frac{v q_G q_\rho}{2^c - q_\rho} + \frac{v q_\rho q_G + v^2 q_G^2}{2^n}
\end{aligned} \tag{41}$$

To conclude the proof we require one more game hop. Consider the game $\boxed{G_0}$ displayed in Fig. 26. In this game, the G oracle first samples a value S , followed by sampling a random output R , independent of ρ . The boxed code ensures that the output R is identical if a value for S is sampled twice. This game is equivalent to H_{v-1} which yields

$$\Pr[\mathcal{A}^{H_{v-1}} \Rightarrow 1] = \Pr[\mathcal{A}^{\boxed{G_0}} \Rightarrow 1]. \tag{42}$$

Now, games $\boxed{G_0}$ and G_0 (also displayed in Fig. 26) are identical until \mathbf{Bad} occurs. The flag is set to true if a value S is sampled that is already in the set \mathcal{D} . At any point in time $|\mathcal{D}| \leq q_G$, and S is sampled by the G independently of \mathcal{A} , thus

$$\Pr[\mathcal{A}^{G_0} \Rightarrow 1] - \Pr[\mathcal{A}^{\boxed{G_0}} \Rightarrow 1] \leq \Pr[\mathbf{Bad}] \leq \sum_{i=1}^{q_G} \frac{|\mathcal{D}|}{2^n} \leq \sum_{i=1}^{q_G} \frac{q_G}{2^n} \leq \frac{q_G^2}{2^n}. \tag{43}$$

Moreover G_0 and the PRG game with its challenge bit fixed to 0 are functionally equivalent, and therefore

$$\Pr[\mathcal{A}^{G_0} \Rightarrow 1] = \Pr[\mathcal{A}^{\text{PRG}} \Rightarrow 1 \mid b = 0]. \tag{44}$$

By combining (36), (41), (42), (43), and (44) we obtain the desired result. \square

C.4 Proof of Theorem 10

We prove the theorem in two stages. First, we reduce the collision resistance of the vector hash function to the collision resistance of the standard sponge-based hash function. Towards this end we introduce a padding scheme which we call *input separation padding* (isPad). This will allow us to view the evaluation of the vector hash over a triple of inputs as the evaluation of the standard sponge-based hash over an encoding of this triple of inputs into a single string – see Fig. 27. We show that isPad is injective, and consequently that any collision in the vector hash yields a collision in the standard sponge-based hash function. In the second part, we prove that the standard sponge-based hash function, instantiated with a random transformation, is collision-resistant. This part is based on the proof described by Boneh and Shoup in [10]. However, since their proof is in the random permutation model, it requires minor changes to adapt it to the random transformation model. We state it here for completeness.

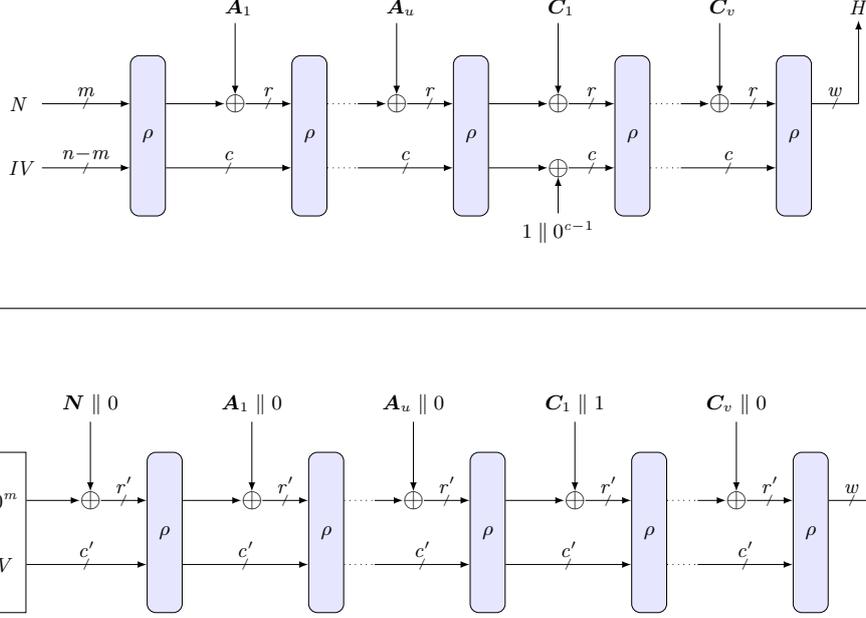


Fig. 27: **Above:** the sponge-based vector hash function SVHASH (used in SLAE). **Below:** SVHASH viewed as the standard sponge-based hash function \mathcal{H}' with its inputs encoded via isPad.

The padding isPad is described in Fig. 28, which in turn makes use of lpad*. The latter takes the triple (N, A, C) and an integer r (the rate) as input and outputs $(\mathbf{N}, \mathbf{A}, \mathbf{C})$ such that the size of each is a multiple of r . The first input (N) is simply padded with 0's to a length r , whereas A and C are padded with lpad(A, r) and lpad(C, r). Here lpad simply applies the 10* padding described in Fig. 9 and is already used in SVHASH— see Fig. 10. Then $(\mathbf{N}, \mathbf{A}, \mathbf{C})$ is mapped to a string $Z = Z_1 \parallel Z_2 \parallel Z_3$ whose length is a multiple of $r + 1$, where Z_1 is \mathbf{N} appended with a 0, Z_2 is \mathbf{A} with a 0 inserted after each r bits, and Z_3 is \mathbf{C} where a 1 is inserted after the first r bits and a 0 after all other r bits. Note that nonces are assumed to be of fixed size.

First note that lpad is itself an injective encoding. Assume now, towards a contradiction, that the isPad encodings of two distinct triples are equal, i.e. $(N, A, C) \neq (\bar{N}, \bar{A}, \bar{C})$ and $Z = \bar{Z}$. It then immediately follows that $N = \bar{N}$, since the nonces are of fixed size, and that $|Z| = |\bar{Z}|$. There are then only two possible cases: (1) $|Z_2| \neq |\bar{Z}_2|$ or (2) $|Z_2| = |\bar{Z}_2|$. For the first case, assume, without loss of generality, that $|Z_2| < |\bar{Z}_2|$, i.e. $u < \bar{u}$, where u and \bar{u} are the number of r -bit blocks of \mathbf{A} and $\bar{\mathbf{A}}$, respectively. Then when writing the strings Z and \bar{Z} directly below each other, the block \mathbf{C}_1 is above $\bar{\mathbf{A}}_{u+1}$. However, isPad requires that \mathbf{C}_1 is followed by a 1 whereas $\bar{\mathbf{A}}_{u+1}$ is followed by a 0 which contradicts the assumption that $Z = \bar{Z}$. As for the second case, it follows that $Z_2 = \bar{Z}_2$ and $Z_3 = \bar{Z}_3$. In turn this means that $\mathbf{A} = \bar{\mathbf{A}}$ and $\mathbf{C} = \bar{\mathbf{C}}$, and by the injectivity of lpad it also follows that $A = \bar{A}$ and $C = \bar{C}$. Since we already established that $N = \bar{N}$, we have that $(N, A, C) = (\bar{N}, \bar{A}, \bar{C})$, which contradicts our assumption that the inputs are distinct.

isPad($(N, A, C), r$)	lpad* ($(N, A, C), r$)
$(\mathbf{N}, \mathbf{A}, \mathbf{C}) \leftarrow \text{lpad}^*((N, A, C), r)$	$\mathbf{N} \leftarrow N \parallel 0^{r- N }$
$Z_1 \leftarrow \mathbf{N} \parallel 0$	$\mathbf{A} \leftarrow A \parallel \text{lpad}(A, r)$
$Z_2 \leftarrow \mathbf{A}_1 \parallel 0 \parallel \mathbf{A}_2 \parallel 0 \parallel \dots \parallel \mathbf{A}_u \parallel 0$ st $\forall i \ A_i = r$	$\mathbf{C} \leftarrow C \parallel \text{lpad}(C, r)$
$Z_3 \leftarrow \mathbf{C}_1 \parallel 1 \parallel \mathbf{C}_2 \parallel 0 \parallel \dots \parallel \mathbf{C}_v \parallel 0$ st $\forall i \ C_i = r$	return $(\mathbf{N}, \mathbf{A}, \mathbf{C})$
return $Z \leftarrow Z_1 \parallel Z_2 \parallel Z_3$	

Fig. 28: The padding schemes isPad and lpad*.

Now we can use `isPad` to reduce the collision resistance of `SVHASH` over triples to the collision resistance of the standard sponge-based hash \mathcal{H}' over strings. Both hash functions are depicted in Fig. 27. Furthermore, note that `SVHASH` can be expressed as $\text{SVHASH}(N, A, C) = \mathcal{H}'(\text{isPad}(N, A, C))$. This mapping is such that \mathcal{H}' requires an extra call to the random transformation and has capacity $c' = c - 1$, where c is the capacity of \mathcal{H}' . Now, since `isPad` is injective, any collision that is found on `SVHASH` must correspond to a collision on \mathcal{H}' . Moreover since `isPad` is efficiently computable, any collision on `SVHASH` can easily be translated into a collision on \mathcal{H}' . This leads to a straightforward reduction which for any adversary \mathcal{A} against `SVHASH` yields an adversary with similar resources $\bar{\mathcal{A}}$ against \mathcal{H}' such that

$$\mathbf{Adv}_{\text{SVHASH}}^{cr}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{H}'}^{cr}(\bar{\mathcal{A}}). \quad (45)$$

This concludes the first part of the proof and we now go on to prove the collision-resistance of the standard sponge-based hash function \mathcal{H}' with capacity c' .

We assume, without loss of generality, that the adversary makes queries to the random transformation which correspond to its final output. That is, it queries the random transformation on all intermediate states that occur while computing the output. Furthermore, we assume that no redundant queries are made by the adversary, i.e. whenever the adversary makes a ρ query on Y yielding $S = \rho(Y)$, the adversary will never make another ρ query on Y . This assumption is justified by the fact that an adversary gets no additional information from redundant queries and any adversary that makes redundant queries can easily be turned into an adversary which makes no redundant queries.

Similar to the proof in [10], we use a directed graph G to visualize the attack. At the start, the graph contains all bit strings of length n as nodes and no edges. Let Y and S be two nodes. An edge from Y to S is added to the graph, if the adversary makes a query to ρ on Y resulting in S . Since there are no redundant queries, an edge is added to the graph exactly once. In contrast to the proof in [10], we have to deal with the fact that the transformation ρ is not injective and that a node may have more than one incoming edge. For $k \geq 1$, a path of length k can be described by a sequence of $2k$ nodes

$$Y_1, S_2, Y_2, S_3, Y_3, \dots, S_k, Y_k, S_{k+1}.$$

Moreover it holds that $\hat{Y}_1 = IV$, $\hat{S}_i = \hat{Y}_i$ for $i = 2, \dots, k$, and the graph G contains edges $Y_i \rightarrow S_{i+1}$ for $i = 1, \dots, k$. The message of this path is defined as a tuple (M_1, \dots, M_k) of bit strings of length r , where $M_1 = \bar{Y}_1$ and $M_i = \bar{S}_i \oplus \bar{Y}_i$ for $i = 2, \dots, k$. The result of this path is $M_{k+1} = \bar{S}_{k+1}$. Such a path corresponds to the computation of $\mathcal{H}'(M_1 \parallel \dots \parallel M_k)$ resulting in output M_{k+1} . We write such paths as

$$M_1|Y_1 \rightarrow S_2|M_2|Y_2 \rightarrow \dots \rightarrow S_k|M_k|Y_k \rightarrow S_{k+1}|M_{k+1}.$$

We can then define a collision in the hash function in terms of the graph G . A collision corresponds to finding a pair of *colliding paths*, which are paths on different messages $(M_1, \dots, M_k) \neq (M'_1, \dots, M'_l)$ resulting in messages M_{k+1} and M'_{l+1} which agree in their first w bits, i.e. $\lfloor M_{k+1} \rfloor_w = \lfloor M'_{l+1} \rfloor_w$.

As in [10], we use the notion of *problematic paths*. Consider the following two paths on messages $(M_1, \dots, M_k) \neq (M'_1, \dots, M'_l)$ of length k and l , respectively.

$$\begin{aligned} M_1|Y_1 \rightarrow S_2|M_2|Y_2 \rightarrow \dots \rightarrow S_{k-1}|M_{k-1}|Y_{k-1} \rightarrow S_k|M_k|Y_k \rightarrow S_{k+1}|M_{k+1} \\ M'_1|Y'_1 \rightarrow S'_2|M'_2|Y'_2 \rightarrow \dots \rightarrow S'_{l-1}|M'_{l-1}|Y'_{l-1} \rightarrow S'_l|M'_l|Y'_l \rightarrow S'_{l+1}|M'_{l+1} \end{aligned}$$

Intuitively, we call a pair of paths *problematic* if their states are equal before the last random transformation is applied. More formally, using the above representation of a path, if $Y_k = Y'_l$ then the paths are called problematic.

Let us denote by CP the event that the adversary finds a pair of colliding paths and by PP the event that the adversary finds a pair of problematic paths. Then the probability of finding a pair of colliding paths, i.e. a collision in the hash function, is bounded as follows

$$\Pr[CP] \leq \Pr[CP \wedge \neg PP] + \Pr[PP]. \quad (46)$$

We start by proving an upper bound for $\Pr[CP \wedge \neg PP]$. This argument closely follows the argument given in [10], however, we obtain a different bound due to the fact that the proof is in the random transformation model. That a pair of paths are both colliding and not problematic means that the final edges correspond to ρ queries on different inputs resulting in outputs where the first w bits are equal. That is, the adversary makes queries to ρ on $Y \neq Y'$ which result in S and S' , respectively, such that $\lfloor \bar{S} \rfloor_w = H = \lfloor \bar{S}' \rfloor_w$. For $i \leq j$, let X_{ij} denote the event that the i -th query of the adversary is some value

Y resulting in $S = \rho(Y)$, while the j -th query is on some value $Y' \neq Y$, yielding output $S' = \rho(Y')$ with $[\bar{S}]_w = [\bar{S}']_w$. Fix i, j , the random coins of the adversary, and the outputs of all queries made before the j -th query. Then Y, S , and Y' are fixed while S' is distributed uniformly at random over a set of size 2^n . Since there are 2^{n-w} nodes W for which it holds that $[W]_w = [\bar{S}]_w$, S' must be equal to one of these, thus

$$\Pr[X_{ij}] \leq \frac{2^{n-w}}{2^n} \leq \frac{1}{2^w},$$

which leads to

$$\Pr[CP \wedge \neg PP] \leq \sum_{j=1}^q \sum_{i=1}^{j-1} \Pr[X_{ij}] \leq \sum_{j=1}^q \sum_{i=1}^{j-1} \frac{1}{2^w} \leq \frac{q(q-1)}{2^{w+1}}.$$

Next, we prove an upper bound for $\Pr[PP]$. The probability of finding a pair of problematic paths is closely related to two basic attacks against random sponges introduced by Bertoni et al. [9], which allows us to upper bound the probability. The first one, called *path to an inner state*, asks, for a given bit string $x \in \{0, 1\}^c$, to find a path to a node S with inner part equal to x , i.e. $\hat{S} = x$. To match this to our case, we define $E1$ to be the event that an adversary finds a path to the inner state IV . The second one, called *inner collision*, asks to find two different paths to nodes S and S' with equal inner states, i.e. $\hat{S} = \hat{S}'$, hence we define $E2$ to be the event that an adversary finds an inner collision. In the following, we show that

$$\Pr[PP] \leq \Pr[E1] + \Pr[E2], \quad (47)$$

i.e. every adversary that finds a pair of problematic paths either finds a path to the inner state IV ($E1$) or an inner collision ($E2$). Consider an adversary that finds a pair of problematic paths which are of length k and l , respectively. Write these paths as

$$\begin{aligned} M_1|Y_1 \rightarrow S_2|M_2|Y_2 \rightarrow \cdots \rightarrow S_{k-1}|M_{k-1}|Y_{k-1} \rightarrow S_k|M_k|Y_k \rightarrow S_{k+1}|M_{k+1} \\ M'_1|Y'_1 \rightarrow S'_2|M'_2|Y'_2 \rightarrow \cdots \rightarrow S'_{l-1}|M'_{l-1}|Y'_{l-1} \rightarrow S'_l|M'_l|Y'_l \rightarrow S'_{l+1}|M'_{l+1} \end{aligned}$$

Due to the fact that the paths are problematic, it holds that $Y_k = Y'_l$ and $(M_1, \dots, M_k) \neq (M'_1, \dots, M'_l)$. We further assume that the pair is the shortest among all pairs of problematic paths. This means that $k+l$ is minimal, and without loss of generality, that $k \leq l$. There are three cases to consider, depending on the length of the paths:

1. ($k = 1$ and $l = 1$) In this case, the paths are simply $M_1|Y_1 \rightarrow S_2|M_2$ and $M'_1|Y'_1 \rightarrow S'_2|M'_2$. The messages of these paths are M_1 and M'_1 , respectively. However, it is impossible that both $Y_1 = Y'_1$ and $M_1 \neq M'_1$. By construction it holds that

$$\bar{Y}_1 = M_1 \neq M'_1 = \bar{Y}'_1$$

which leads to a contradiction.

2. ($k = 1$ and $l \geq 2$) In this case, we have that $Y_1 = Y'_l$. The adversary has made a ρ query on Y'_{l-1} which resulted in $S'_l = \rho(Y'_{l-1})$. Therefore, it holds that $\hat{S}'_l = \hat{Y}'_l = \hat{Y}_1 = IV$. The first equality follows from the construction, where the inner state between two evaluations of ρ does not change. The second one trivially follows from $Y_1 = Y'_l$, while the third follows again from the construction. By taking the path up to node S'_l , the adversary has found a path to the inner state IV , i.e. event $E1$ occurs.

3. ($k \geq 2$ and $l \geq 2$) Here it holds that $Y_k = Y'_l$. Consider now the last but one edges which are

$$Y_{k-1} \rightarrow S_k|M_k|Y_k$$

$$Y'_{l-1} \rightarrow S'_l|M'_l|Y'_l$$

There are two cases depending on the absorbed messages M_k and M'_l :

- (a) The absorbed messages are equal, i.e. $M_k = M'_l$. Hence, it holds that $S_k = S'_l$. Due to the fact that ρ is a random transformation we have to further distinguish between the following sub-cases:

- i. ($Y_{k-1} = Y'_{l-1}$) This contradicts the minimality of $k + l$, as we can find a shorter pair of problematic paths by throwing away the last edge in each path, while the truncated messages (M_1, \dots, M_{k-1}) and (M'_1, \dots, M'_{l-1}) still differ.
 - ii. ($Y_{k-1} \neq Y'_{l-1}$) In this case, both Y_{k-1} and Y'_{l-1} map to $S_k = S'_l$, therefore the adversary has found an inner collision on the nodes S_k and S'_l , i.e. event $E2$ occurs.
- (b) The absorbed messages are different, i.e. $M_k \neq M'_l$. This yields $\bar{S}_k \neq \bar{S}'_l$ which further implies $S_k \neq S'_l$ and $Y_{k-1} \neq Y'_{l-1}$. On the other hand, the absorbed messages do not affect the inner state of S_k and S'_l which ensures $\hat{S}_k = \hat{S}'_l$. Combining these, it holds that the adversary has found an inner collision on the nodes S_k and S'_l by taking the paths up to S_k and S'_l , i.e. event $E2$ occurs. We emphasise that $Y_{k-1} \neq Y'_{l-1}$ implies that the paths are indeed distinct.

This proves inequality (47). Thus, we can upper bound $\Pr[PP]$ using the following bounds on $\Pr[E1]$ and $\Pr[E2]$ by Bertoni et al. [9], where q denotes the number of queries the adversary makes to ρ .

$$\Pr[E1] \leq \frac{q}{2^{c'}} \quad \Pr[E2] \leq \frac{q(q+1)}{2^{c'+1}}$$

Substituting everything into (46) yields

$$\Pr[CP] \leq \frac{q(q-1)}{2^{w+1}} + \frac{q}{2^{c'}} + \frac{q(q+1)}{2^{c'+1}} \leq \frac{q(q-1)}{2^{w+1}} + \frac{q(q+2)}{2^{c'}}. \quad (48)$$

By combining (45) and (48) we get

$$\mathbf{Adv}_{\text{SVHASH}}^{cr}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{H}'}^{cr}(\bar{\mathcal{A}}) \leq \Pr[CP] \leq \frac{q(q-1)}{2^{w+1}} + \frac{q(q+2)}{2^{c'}}, \quad (49)$$

and then replacing $c' = c - 1$ we obtain the desired result:

$$\mathbf{Adv}_{\text{SVHASH}}^{cr}(\mathcal{A}) \leq \frac{q(q-1)}{2^{w+1}} + \frac{q(q+2)}{2^{c-1}}.$$

□