

Efficient Leakage-Resilient MACs without Idealized Assumptions

Francesco Berti^{1,2}, Chun Guo^{3,4,5}, Thomas Peters¹, François-Xavier Standaert¹

¹ UCLouvain, ICTEAM/ELEN/Crypto Group

² TU Darmstadt, Germany, CAC - Applied Cryptography

³ School of Cyber Science and Technology, Shandong University, Qingdao, Shandong, 266237, China

⁴ Key Laboratory of Cryptologic Technology and Information Security of Ministry of Education, Shandong University, Qingdao, Shandong, 266237, China

⁵ State Key Laboratory of Information Security (Institute of Information Engineering, Chinese Academy of Sciences, Beijing 100093)

Abstract. The security proofs of leakage-resilient MACs based on symmetric building blocks currently rely on idealized assumptions that hardly translate into interpretable guidelines for the cryptographic engineers implementing these schemes. In this paper, we first present a leakage-resilient MAC that is both efficient and secure under standard and easily interpretable black box and physical assumptions. It only requires a collision resistant hash function and a single call per message authentication to a Tweakable Block Cipher (TBC) that is unpredictable with leakage. This construction leverages two design twists: large tweaks for the TBC and a verification process that checks the inverse TBC against a constant. It enjoys beyond birthday security bounds. We then discuss the cost of getting rid of these design twists. We show that security can be proven without them as well. Yet, a construction without large tweaks requires stronger (non idealized) assumptions and may incur performance overheads if specialized TBCs with large tweaks can be exploited, and a construction without twisted verification requires even stronger assumptions (still non idealized) and leads to more involved bounds. The combination of these results makes a case for our first pragmatic construction and suggests the design of TBCs with large tweaks and good properties for side-channel countermeasures as an interesting challenge.

1 Introduction

Ever since its introduction by Dziembowski and Pietrzak [20], leakage-resilient cryptography has been characterized by a quest towards the best tradeoff between weak physical assumptions and efficient cryptographic constructions. Finding good abstractions to limit the informativeness of the leakage function, that can be fulfilled by hardware engineers while also enabling sound security proofs, is a typical example of this challenge. Current assumptions range from various types of “bounded leakage”, as comprehensively discussed in [21], to simpler solutions leveraging the scarce use of “strongly protected components”, modeled

as leak-free in [33]. Unsurprisingly, the most efficient (symmetric) constructions in the literature leverage such strong (idealized) assumptions [6].

While avoiding idealized assumptions is of general interest in cryptography, it is even more desirable in leakage-resilient cryptography, since perfectly ensuring a physical assumption may not be possible, or only at prohibitive cost. For example, instantiating a (128-bit) leak-free component would require masking it at very high security orders, leading to significant performance overheads [22], while security against $< 2^{80}$ measurements may be sufficient for a majority of applications. So despite proofs relying on a leak-free component are a useful guide towards efficient constructions with good leakage properties, interpreting their security bounds in terms of concrete requirements for cryptographic engineers (e.g., in terms of a level of protection to reach in practice) remains difficult.

Given this state-of-the-art, the design of leakage-resilient MACs appears as a first natural target. As put forward by Micali and Reyzin, ensuring unpredictability in the presence of leakage is significantly easier than ensuring indistinguishability in the presence of leakage [32]. This observation led the authors of [10] to propose authenticated encryption schemes for which the integrity holds even if the vast majority of its (ephemeral) secrets are leaked in full to the adversary: a model that we next denote as the “unbounded leakage model”. Yet, these authenticated encryption schemes (and follow ups next listed as related works) still rely on the scarce utilization of a leak-free component.

An intermediate step towards getting rid of the leak-free component has been made by Berti et al. [8]: it shows that it is possible to replace this leak-free component by a (tweakable) block cipher ensuring “strong unpredictability with leakage”. The idea of basing MAC security on unpredictable ciphers is not new. To the best of our knowledge, it dates back to [1] and has been revisited in [18,36,19]. Its leakage generalization is specially appealing since such a game-based definition can then be verified/falsified by evaluation laboratories. Yet, the results in [8] still rely on a random oracle assumption, which seems mostly due to the difficult interaction between an unpredictable (tweakable) block cipher and the hash part of their constructions. As for the leak-free component, such an idealized assumption is in general undesirable, and possibly even more when leakage comes into play. So the main question we tackle in this work is: *can we design leakage-resilient MACs without idealized assumptions (i.e., no leak-free component nor random oracles)?* We answer it positively by exhibiting efficient constructions for which the leakage security holds in the unbounded leakage model, only requiring strong unpredictability with leakage for their tweakable block cipher and (more or less) standard properties for their hash function.

More precisely, our contributions are threefold:

We first propose a pragmatic construction (next denoted as LR-MAC1) that takes advantage of simple design twists. The starting observation for this purpose is that the HTBC construction in [8] performs the verification by checking whether an inverted Tweakable Block Cipher (TBC) matches the output of a hash function. By changing this verification and checking whether the inverted TBC equals a constant (e.g., zero) value, we avoid the difficult interaction be-

tween the hash function and the TBC that has led this previous work to rely on a random oracle assumption. We then show that if the construction is instantiated with a TBC having $2n$ -bit tweaks, it provides tight and beyond-birthday security bounds, under standard assumptions (namely, strong unpredictability with leakage for the TBC and collision-resistance for the hash function). While this construction positively answers our question, it still relies on two design twists, namely a TBC with $2n$ -bit tweaks and the introduction of a constant value in the verification process. So we complement our first pragmatic constructions by two other designs aiming to clarify whether these twists are necessary.

Our second design (next denoted as LR-MAC2) is a variant of LR-MAC1 that only relies on TBCs having n -bit tweaks. We show that it is possible to maintain leakage security without idealized assumptions with this simpler building block. However, designs reaching this goal currently need to rely on two calls to the TBC, which is in general more expensive than a single call to a TBC with $2n$ -bit tweaks if a specialized TBC can be used [27]. Besides, the best construction we reach also requires a less standard (yet non idealized) assumption on its hash function (namely, collision-resistance for one half of its output). The latter may therefore require more rounds in the instances of hash functions it uses.

Eventually, we revisit the leakage security of the HTBC design (which does not use a constant value in its verification), analyzed in [8] under a random oracle assumption. We show that such a construction can be proven secure without this idealized assumption, but at the cost of more involved bounds and stronger (still non idealized) assumptions for the hash function than LR-MAC2.

Overall, the combination of our observations regarding LR-MAC2 and HTBC strengthen the interest of the pragmatic LR-MAC1 solution, and suggest the design of TBCs with $2n$ -bit tweaks and good properties to be protected via masking as natural design targets for leakage-resilient modes of operation.

Related works. Leakage-resilient MACs have first been proposed by Hazay et al. [25] and Martin et al. [31], but imply higher performance overheads than the symmetric constructions we consider in this work. The first leakage-resilient MACs based on symmetric building blocks were proposed in [34,33]. Several (TBC-based or permutation-based) leakage-resilient authenticated encryption schemes have been proposed and embed a leakage-resilient MAC for their integrity guarantees, e.g., [14,16,9,15,24,23,15,30]: they all rely on idealized assumptions to some extent (see [6] for an overview). The specific issue of comparing values (e.g., tags) in a leakage-resilient manner is discussed in [11,17].

2 Background

Notations. With $\{0, 1\}^n$ (resp., $\{0, 1\}^*$), we denote the set of all strings of length n (resp., all finite-length strings). With $|\mathcal{X}|$, we denote the size of the set \mathcal{X} . In some arguments, we denote by n the security parameter, $\text{negl}(n)$ negligible functions and $\text{poly}(n)$ / $\text{superpoly}(n)$ polynomial/super-polynomial functions.

2.1 Primitives: hash functions and TBC

Our schemes use hash functions and TBCs. For hash functions, the minimum property we require is collision-resistance, which we recall next:

Definition 1. Let $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \mathcal{X}$ be a hash function. H is (t, ϵ) -collision resistant (CR) if for every t -bounded adversary A , the probability that $A(s)$ outputs a pair of distinct inputs $(m^0, m^1) \in (\{0, 1\}^*)^2$, such that $H_s(m^0) = H_s(m^1)$ and $m^0 \neq m^1$, is bounded by ϵ , with $s \xleftarrow{\$} \mathcal{HK}$ picked uniformly at random with:

$$\Pr[s \xleftarrow{\$} \mathcal{HK}, A(s) \Rightarrow (m^0, m^1) \in (\{0, 1\}^*)^2 \text{ s.t. } m^0 \neq m^1, H_s(m^0) = H_s(m^1)] \leq \epsilon.$$

We will sometimes require range-oriented preimage resistant hash functions:

Definition 2. Let $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \mathcal{X}$ be a hash function. H is (t, ϵ) -range-oriented preimage resistant (rpre) if, for every t -bounded adversary A :

$$\Pr[s \xleftarrow{\$} \mathcal{HK}, y \xleftarrow{\$} \mathcal{X}, A(s, y) \Rightarrow m \in \{0, 1\}^* \text{ s.t. } H_s(m) = y] \leq \epsilon.$$

Properties of rpre hashing. Following [2], a hash function can be characterized by a table of preimage computing probabilities.

Definition 3. Let A be an adversary against the preimage resistance of a hash function, then the associated success probability matrix $M \in [0, 1]^{|\mathcal{HK}| \times |\mathcal{X}|}$ is:

$$M_{s,y} = \Pr_{A(s,y) \Rightarrow m} [H_s(m) = y],$$

defined entry-wise where $s \in \mathcal{HK}$ and $y \in \mathcal{X}$, and where the probability is taken over the random coins of A .

Consider a hash function $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \mathcal{X}$ and an arbitrary adversary A . The intuition is that the preimage for some “weak” points in \mathcal{X} may be easily computable. But if H is range-oriented preimage resistant, then the total number of such “weak” points should be limited. To formalize this idea, for any $s \in \mathcal{HK}$, we define a set for such “weak” points as:

$$\mathcal{WP}(s, A) = \left\{ y \in \mathcal{X} : \Pr_{A(s,y) \Rightarrow m} [H_s(m) = y] = \text{poly}(n)^{-1} \right\}. \quad (1)$$

With the above definition, we are able to establish the following claim.

Lemma 1. If H is range-oriented preimage resistant, then with high probability, the size of the set $\mathcal{WP}(s, A)$ of “weak” images is polynomial with:

$$\Pr[s \xleftarrow{\$} \mathcal{HK} : |\mathcal{WP}(s, A)| = \text{superpoly}(n)] = \text{negl}(n).$$

Proof. To see this, we expand the expression as follows:

$$\begin{aligned}
& \Pr[s \xleftarrow{\$} \mathcal{HK}, y \xleftarrow{\$} \mathcal{X}, \mathbf{A}(s, y) \Rightarrow m \in \{0, 1\}^* \text{ s.t. } \mathbf{H}_s(m) = y], \\
&= \sum_{s \in \mathcal{HK}} \frac{1}{|\mathcal{HK}|} \sum_{y \in \mathcal{X}} \frac{1}{|\mathcal{X}|} \Pr[\mathbf{A}(s, y) \Rightarrow m \in \{0, 1\}^* \text{ s.t. } \mathbf{H}_s(m) = y], \\
&\geq \sum_{s \in \mathcal{HK}} \frac{1}{|\mathcal{HK}|} \sum_{y \in \mathcal{WP}(s, \mathbf{A})} \frac{1}{|\mathcal{X}|} \cdot \text{superpoly}(n), \\
&\geq \sum_{s \in \mathcal{HK}, |\mathcal{WP}(s, \mathbf{A})| = \text{superpoly}(n)} \frac{1}{|\mathcal{HK}|} \frac{|\mathcal{WP}(s, \mathbf{A})|}{|\mathcal{X}|} \cdot \text{superpoly}(n), \\
&\geq \Pr[s \xleftarrow{\$} \mathcal{HK} : |\mathcal{WP}(s, \mathbf{A})| = \text{superpoly}(n)] \cdot \frac{\text{superpoly}(n)}{|\mathcal{X}|}. \tag{2}
\end{aligned}$$

By this, if $\Pr[s \xleftarrow{\$} \mathcal{HK} : |\mathcal{WP}(s, \mathbf{A})| = \text{superpoly}(n)]$ is not $\text{negl}(n)$, then Eq. (2) cannot be negligible, contradicting the range-oriented preimage resistance assumption. The claim thus follows. \square

Thanks to Lemma 1, it *may* be feasible to decide the set $\mathcal{WP}(s, \mathbf{A})$ for any adversary \mathbf{A} and any seed s . In Section 5, Theorem 3, this result will allow us to assume that the set $\mathcal{WP}(s, \mathbf{A})$ for the hash function used in our design is computable in Probabilistic Polynomial-Time (PPT).

For TBC, the minimum property we require is strong pseudorandomness:

Definition 4 (stPRP). A tweakable block cipher $\mathbf{F} : \mathcal{K} \times \mathcal{TW} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a (q, t, ϵ) -strong tweakable pseudorandom permutation (stPRP) if $\forall (k, tw) \in \mathcal{K} \times \mathcal{TW}, \mathbf{F}_k^{tw} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a permutation and if for every (q, t) -adversary \mathbf{A} , the advantage :

$$\text{Adv}_{\mathbf{F}}^{\text{stPRP}}(\mathbf{A}) := \left| \Pr[\mathbf{A}^{\mathbf{F}_k(\cdot, \cdot), \mathbf{F}_k^{-1}(\cdot, \cdot)} \Rightarrow 1] - \Pr[\mathbf{A}^{\mathbf{f}(\cdot, \cdot), \mathbf{f}^{-1}(\cdot, \cdot)} \Rightarrow 1] \right|,$$

is upper bounded by ϵ , where k and \mathbf{f} are chosen uniformly at random from their domains, namely \mathcal{K} and the space $\mathcal{TPERM}(\mathcal{TW}, \{0, 1\}^n)$ of tweakable permutations (i.e., the space of functions $\mathbf{f} : \mathcal{TW} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ s.t. $\forall tw \in \mathcal{TW}, \mathbf{f}(tw, \cdot) : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a permutation). The adversary can do at most q' queries to the first oracle and $q - q'$ queries to the second one for any $q' \leq q$.

For simplicity, an (n, n, n) -TBC is a TBC with $\mathcal{K} = \mathcal{TW} = \{0, 1\}^n$.

2.2 MAC (Message Authentication Code)

We recall the definition of Message Authentication Code (MAC) as follows:

Definition 5. A MAC is a triple $\Pi = (\text{Gen}, \text{Mac}, \text{Vrfy})$ where:

Gen. The key-generation algorithm Gen picks a key in the keyspace \mathcal{K} .

Mac. The tag-generation algorithm Mac takes as input a couple $(k, m) \in \mathcal{K} \times \{0, 1\}^*$ and outputs a tag $\tau \leftarrow \text{Mac}_k(m)$ from the tag space \mathcal{TAG} .

Vrfy. The verification algorithm Vrfy takes as input a triple (k, m, τ) in $\mathcal{K} \times \{0, 1\}^* \times \mathcal{TAG}$ and outputs either “ \top ” (“accept”) or “ \perp ” (“reject”).

We require correctness: $\forall (k, m) \in \mathcal{K} \times \{0, 1\}^*, \text{Vrfy}(k, m, \text{Mac}(k, m)) = \top$.

Since we will only consider the security for MACs in the presence of leakage, we omit the security definition in the black-box model (that is, when there is no leakage). This definition can be found in many works, for example [29].

2.3 Leakage models and security definitions with leakage

Notations. When an adversary has access not only to the outputs of an oracle but also to its leakage, we denote it with \mathbf{A}^{OL} . In this case, queries to the oracle OL on input x are answered with $y = \text{O}(x)$ and the leakage $\text{l}_O := \text{L}_O(x)$. If the oracle is keyed with the key k , we write the leakage function as $\text{L}_O(x; k)$. Finding good abstractions to restrict L_O is in general a hard problem [21].⁶

Strong unforgeability with leakage (sUF-L2). We start by introducing the security for MACs in the presence of leakage. We want that it is hard to forge valid tags even having access to the leakage of the tag-generation and the verification algorithms (that is, finding a fresh and valid couple message tag (m, τ) such that $\text{Vrfy}_k(m, \tau) = \top$ should be hard). We use the sUF-L2 definition of Berti et al. [8] for this purpose, which we recall next:

Definition 6 (sUF-L2). A MAC = (Gen, Mac, Vrfy) with tag-generation leakage function L_M and verification leakage function L_V is $(q_L, q_M, q_V, t, \epsilon)$ -strongly existentially unforgeable against chosen message and verification attacks with leakage in the tag-generation and the verification (sUF-L2) if for all (q_L, q_M, q_V, t) -adversaries \mathbf{A}^{L} , we have:

$$\Pr \left[\text{FORGEL2}_{\text{MAC}, \text{L}_M, \text{L}_V, \mathbf{A}}^{\text{suf-vcma-L2}} \Rightarrow 1 \right] \leq \epsilon,$$

where the $\text{FORGEL2}_{\text{MAC}, \text{L}_M, \text{L}_V, \mathbf{A}}^{\text{suf-vcma-L2}}$ experiment is defined in Table 1.

For simplicity, we consider the verification query induced by the final output of the adversary as the $(q_V + 1)$ th verification query.

⁶ Adversaries are sometimes allowed to “model” the leakage. For this purpose, we grant them access to the oracle L . This oracle is peculiar since it allows the adversary to make queries not only on inputs x but also of keys k' of its choice.

The FORGEL2 ^{suf-vcma-L2} _{MAC,L_M,L_V,A^L} experiment	
Initialization: $k \leftarrow \text{Gen}$ $\mathcal{S} \leftarrow \emptyset$	Oracle MacL_k(m): $\tau = \text{Mac}_k(m)$ $\mathcal{S} \leftarrow \mathcal{S} \cup \{(m, \tau)\}$ Return $(\tau, L_M(m; k))$
Finalization: $(m, \tau) \leftarrow A^{\text{L}, \text{MacL}_k(\cdot), \text{VrfyL}_k(\cdot, \cdot)}$ If $(m, \tau) \in \mathcal{S}$ or $\perp = \text{Vrfy}_k(m, \tau)$ Return 0 Return 1	Oracle VrfyL_k(m, τ): Return $(\text{Vrfy}_k(m, \tau), L_V(m, \tau; k))$

Table 1: The FORGEL2^{suf-vcma-L2} experiment.

The unbounded leakage model. We next need to specify the functions we will use for L_M and L_V . Based on the aforementioned difficulty to restrict the leakage in a meaningful manner, we first follow the observation made in [33] that it is possible to implement leakage-resilient cryptographic functionalities so that the execution of most underlying building blocks can leak in an unrestricted manner. The resulting “leveled implementations” only require a few calls to a strongly protected component (frequently modeled as leak-free) to ensure the desired security property (here, sUF-L2). So we will next consider the unbounded leakage model, where the leakage function yields all the internal states produced during each execution of the scheme under investigation, at the exclusion of the strongly protected components that are used to manipulate long-term secrets. More precisely, in the unbounded leakage model:

- Unprotected building blocks leak their inputs, outputs and keys in full;
- Building blocks with strongly protected implementation leak their inputs and outputs in full and their key only leaks in a restricted manner.

In practice, the only strongly protected component we will use in our construction is a TBC, and we next specify how its leakage will be restricted by asking implementers to ensure its strong unpredictability with leakage.

Strong unpredictability with leakage (sUP-L2). Unpredictability is among the simplest requirements for (tweakable) block ciphers. As mentioned in introduction, its application in leakage-resilient cryptography is appealing since it corresponds to a game-based definition that can directly be tested by an evaluation laboratory. We next give its definition which formalizes that it should be hard for an adversary to find a triple (tw, x, y) which is fresh and valid (i.e., $y = F_k^{tw}(x)$) even if the leakage function $L = (L_{\text{Eval}}, L_{\text{Inv}})$ associated to an implementation of the TBC can be queried, with $L_{\text{Eval}}(tw, x; k)$ (resp., $L_{\text{Inv}}(tw, z; k)$) the leakage resulting from the computation of $F_k(tw, x)$ (resp., $F_k^{-1}(tw, z)$).

Definition 7 (sUP-L2). A tweakable block cipher $F : \mathcal{K} \times \mathcal{TW} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ with leakage function pair $L = (L_{\text{Eval}}, L_{\text{Inv}})$ is $(q_L, q_E, q_I, t, \epsilon)$ -strongly un-

predictable with leakage in evaluation and inversion (sUP-L2), or $(q_L, q_E, q_I, t, \epsilon)$ -sUP-L2, if for any (q_L, q_E, q_I, t) -adversary A , we have:

$$\Pr[\text{sUP-L2}_{A,F,L} \Rightarrow 1] \leq \epsilon,$$

where the sUP-L2 experiment is defined in Table 2, and A^L makes at most q_L (offline) queries to L (used to model the leakage, cf. Footnote 6).

When the adversary has access only to the l_e oracle, the TBC is *unpredictable with leakage* (as previously defined by Dodis and Steinberger [18]).⁷

The sUP-L2 _{A,F,L} experiment.	
Initialization: $k \xleftarrow{\$} \mathcal{K}$ $\mathcal{L} \leftarrow \emptyset$ Finalization: $(x, tw, z) \leftarrow A^{L, \text{LEval}(\cdot, \cdot), \text{LInv}(\cdot, \cdot)}$ If $(x, tw, z) \in \mathcal{L}$ Return 0 If $z = F_k(tw, x)$ Return 1 Return 0	Oracle LEval(tw, x): $z = F_k(tw, x)$ $l_e = L_{\text{Eval}}(tw, x; k)$ $\mathcal{L} \leftarrow \mathcal{L} \cup \{(x, tw, z)\}$ Return (z, l_e) Oracle LInv(tw, z): $x = F_k^{-1}(tw, z)$ $l_i = L_{\text{Inv}}(tw, z; k)$ $\mathcal{L} \leftarrow \mathcal{L} \cup \{(x, tw, z)\}$ Return (x, l_i)

Table 2: Strong unpredictability with leakage in evaluation and inversion.

Concretely, this assumption is significantly closer to practice than idealized ones. In particular, breaking it requires one to fully compute the value of the block cipher on a new point with non-trivial probability [18]. Therefore, we expect that satisfying unpredictability with leakage mainly requires protecting the long-term key, and that testing block cipher implementations against side-channel key recovery attacks, which is the current focus of evaluation laboratories [4], should give a good indication of their unpredictability with leakage.

3 Design and analysis of LR-MAC1

In this section, we present LR-MAC1, which is sUF-L2 in the unbounded leakage model, assuming a collision-resistant hash function and a sUP-L2 TBC.

⁷ Degabriele et al. give an alternative definition in [14]. The main difference is that the set of inputs (\mathcal{S} in their definition, \mathcal{L} in ours) is not increased for inputs X for which only the leakage is observed, while we always give access to both the primitive’s output and their leakage. Hence, this definition cannot be satisfied in the unbounded leakage model as we aim, since the adversary can then get a valid tag in full. (Their motivation was also different from ours and specially tailored for analyzing constructions leveraging leakage-resilient PRFs).

The main design idea used by LR-MAC1 is to avoid that a value obtained by leakage in verification may be used for a future forgery. This happened in the HTBC leakage-resilient MAC of [8] because the output of an (inverse) TBC is compared with a hash value $h = H(m)$. As illustrated in Figure 1, LR-MAC1 prevents this by doing the verification check with a fixed value that does not depend on the message m . In this figure, strongly protected components are in gray, unprotected components are in white, inputs and outputs of the scheme (together with values that can be computed publicly from them) are in green, intermediate values that leak (in an unbounded manner in our case) are in orange and long-term secrets manipulated by strongly protected components are in red.

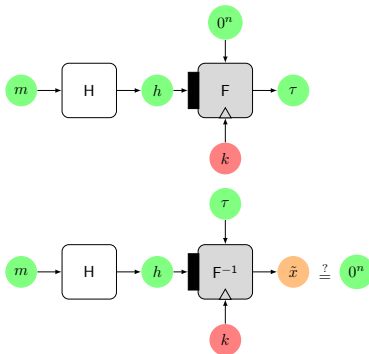


Fig. 1: LR-MAC1.

Concretely, the hash of the message is used only as a tweak of the TBC, while the input of the TBC is always a fixed value (i.e., 0^n). In decryption, we check if the inverse of the tag $\tilde{x} := F_k^{-1}(\tau)$ is 0^n . This way, even an unbounded leakage during an invalid verification cannot lead to efficient forgeries since \tilde{x} cannot be reused (contrary to what happens for HTBC). Moreover, since the hash of the message h is used only as a tweak, there is no target to find a preimage (without producing a collision). As a result, a verification query $\text{Vrfy}_k(m, \tau) \rightarrow \top$ immediately results in a forgery $(H_s(m), 0^n, \tau)$, which does not rely on the distribution of $H_s(m)$. This eliminates the necessity of the random oracle in our proofs. The detailed specifications of LR-MAC1 can be found in Algorithm 1.

3.1 sUF-L2-security of LR-MAC1

We now prove that LR-MAC1 is sUF-L2 in the unbounded leakage model.

Unbounded leakage specification. Before assessing the sUF-L2 of LR-MAC1, we have to define the leakage that an adversary can collect:

Algorithm 1 The LR-MAC1 algorithm.

It uses a strongly protected TBC $F : \mathcal{K} \times \mathcal{TW} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a hash function $H : \mathcal{HK} \times \{0, 1\}^n \rightarrow \mathcal{TW}$.

- Gen
 - $k \xleftarrow{s} \mathcal{K}$
 - $s \xleftarrow{s} \mathcal{HK}$
 - $\text{Mac}_k(m)$:
 - $h = H_s(m)$
 - $\tau = F_k^h(0^n)$
 - Return τ
 - $\text{Vrfy}_k(m, \tau)$:
 - $h = H_s(m)$
 - $\tilde{x} = F_k^{h,-1}(\tau)$
 - If $\tilde{x} == 0^n$ Return \top
 - Else Return \perp
-

- $L_M(m; k)$ returns $\text{LEval}(0^n, h; k)$ with $h = H_s(m)$.⁸
- $L_V(m, \tau; k)$ returns \tilde{x} and $\text{LInv}(\tau, h; k)$ with $h = H_s(m)$.

Our main result for LR-MAC1 is then formalized by the following theorem:

Theorem 1. *Let H be a $(t_1, \epsilon_{\text{CR}})$ -collision resistant hash function. Let F be a $(q_L, q_M, q_V, t_2, \epsilon_{\text{sUP-L2}})$ -sUP-L2 TBC. Then, LR-MAC1 is (q_M, q_V, t, ϵ) -sUF-L2-secure in the unbounded leakage model with:*

$$\epsilon \leq \epsilon_{\text{CR}} + (q_V + 1) \epsilon_{\text{sUP-L2}},$$

with $t_1 = t + (q_M + q_V + 1)t_H + (q_M + q_V)(t_F + t_{L(F)})$ and $t_2 = t + (q_M + q_V + 1)t_H$, where t_H is the time needed to execute once the hash function H , t_F is the time needed to execute once the TBC F and $t_{L(F)}$ is the time needed to collect the leakage of one execution of the TBC F .

Proof. We use a sequence of games. To make the proof simpler, we give only a sketch of the adversaries used. Details are in the ePrint version of the paper.

Game 0. Let Game 0 be the sUF-L2 game where the (q_M, q_V, t) -adversary A^L tries to produce a forgery when she plays against LR-MAC1. Let E_0 be the event that the adversary wins the game (i.e., the output of the game is 1).

Game 1. Game 1 is Game 0 except that we abort if there is a collision in the hash function. Let E_1 be the event that the adversary wins the game.

⁸ h can be computed by the adversary since the key s of the hash function is public.

Transition between Game 1 and Game 2. Clearly, Game 0 and Game 1 are identical if the following event HC (*Hash Collision*) does not happen:

$$HC := \{\exists i, j \in \{1, \dots, q_M\} \cup \{1, \dots, q_V + 1\} \text{ with } i \stackrel{\%}{=} j \text{ s.t. } h^i = h^j \text{ and } m^i \neq m^j\}.$$
⁹

To compute this event, we build a t_1 -CR-adversary B .

The t_1 -CR-adversary B is an adversary playing against the hash function, based on A , which does the hash queries induced by the queries of A and then correctly computes the answer for A . At the end of the game, B sees if her hash queries have produced a collision. If it is the case, she outputs it. This adversary needs time $t + (q_M + q_V + 1)t_H + (q_M + q_V)(t_F + t_{L(F)})$.

Bounding $|\Pr[E_0] - \Pr[E_1]|$. Observe that if event HC happens, B wins. Note that B simulates correctly Game 0 for A^L . Since B is a t_1 -adversary and H is a (t_1, ϵ_{CR}) -collision resistant hash function, we can bound:

$$|\Pr[E_0] - \Pr[E_1]| \leq \epsilon_{CR}.$$

Game 2. Game 2 is Game 1 except that we abort if one verification query (or the final one) made by A is fresh and valid. Let E_2 be the event that the adversary wins the game (i.e., the output of the game is 1).

Transition between Game 1 and Game 2. We build a sequence of $q_V + 2$ games: Game $1^0, \dots, \text{Game } 1^{q_V + 1}$ as follows.

Game 1^j . Game 1^j is Game 1 where we abort if one of the first j verification queries is fresh and valid. Thus, Game 1^0 is Game 1 while Game $1^{q_V + 1}$ is Game 2. Let E_1^j be the event that the adversary wins the game.

Transition between Game 1^j and Game 1^{j+1} . Clearly, Game 1^j and Game 1^{j+1} are identical if the $j+1$ -th verification query is either invalid or not fresh. If this query is fresh and valid, we say that event GT^j (Good Tag) happens. In order to bound the probability that event GT^j happens we build a $(q_L, q_{Eval}, q_{Inv}, t_2)$ -sUP-L2-adversary C^j against F as follows.

The $(q_L, q_{Eval}, q_{Inv}, t_2)$ -sUP-L2-adversary C^j has to find a valid triple (x, tw, y) for F_k which is fresh and valid. She simulates Game 1^j for A until the $j+1$ th verification query. Then, when A asks his $j+1$ -th verification query on input $(m^{q_V + 1}, \tau^{q_V + 1})$, C^j computes (1) $h^{j+1} = H_s(m^{j+1})$, and (2) she outputs $(0^n, h^{j+1}, \tau^{j+1})$. This takes time t_H . In total, C^j does at most q_L query to L , q_M to $LEval$ and $j \leq q_V$ to $LInv$. She needs time at most $t + (q_M + j + 1)t_H \leq t_2$.

⁹ $i \stackrel{\%}{=} j$ means that if i comes from a tag-generation query and j from a verification query, or vice-versa, then they are considered differently.

Bounding $|\Pr[E_1^j] - \Pr[E_1^{j+1}]|$. Note that C^i simulates correctly Game 1^j for A^L . Since C^j is a (q_L, q_M, q_V, t_2) -adversary and F is $(q_L, q_M, q_V, t_2, \epsilon_{\text{sUP-L2}})$ -sUP-L2, we can bound:

$$|\Pr[E_1^j] - \Pr[E_1^{j+1}]| \leq \Pr[GT^{j+1}] \leq \epsilon_{\text{sUP-L2}}.$$

Bounding $|\Pr[E_1] - \Pr[E_2]|$. Iterating we can bound:

$$|\Pr[E_1] - \Pr[E_2]| \leq \sum_{j=0}^{q_V} |\Pr[E_1^j] - \Pr[E_1^{j+1}]| \leq (q_V + 1)\epsilon_{\text{sUP-L2}}.$$

Concluding the proof. Since the probability of E_2 is 0, we can conclude the proof putting everything together:

$$\Pr[E_0] = \sum_{i=0}^1 |\Pr[E_i] - \Pr[E_{i+1}]| + \Pr[E_2] = \epsilon_{\text{CR}} + (q_V + 1)\epsilon_{\text{sUP-L2}} + 0.$$

3.2 Instantiation and concrete security

Built upon a “good enough” hash function, the term ϵ_{CR} is expected to be in $O(t_1^2/2^{|\mathcal{TW}|})$. On the other hand, the term $\epsilon_{\text{sUP-L2}}$ depends on the concrete side-channel strength of the TBC implementation. Since we have a birthday bound only on the size of the tweak space \mathcal{TW} and not on the block space $\mathcal{B} = \{0, 1\}^n$, a natural idea is to use a TBC with the tweak space bigger than the block space. For example, a good instantiation would be to use SHA256 as hash function and Deoxys-384 (which has a 128-bit block and 256-bit tweak [28]) as TBC. In this case, finding a collision requires the computation of around 2^{128} hash functions, and therefore, the $(q_V + 1)\epsilon_{\text{sUP-L2}}$ term dominates the bound.¹⁰

4 Design and analysis of LR-MAC2

In this section, we question the possibility to design a leakage-resilient MAC with a (more standard) TBC design having only n -bit tweaks rather than $2n$ for LR-MAC1. We show that this is achievable at the cost of a second call per message authentication to the TBC and by requiring an additional property of the hash function that we develop after the description of the scheme. Our security analysis also holds without idealized assumption.

4.1 Description of LR-MAC2

To get high security we start again from $H_s(m) = u||v$, where $u, v \in \{0, 1\}^n$, as illustrated in Figure 2. As in HTBC, we first compute $F_k^v(u)$ with $k = k_1$ to get a

¹⁰ In the black box setting (i.e., without leakage), the security of this construction is beyond birthday since $(q_V + 1)\epsilon_{\text{sUP-L2}} \leq \epsilon_{\text{sPRP}} + \frac{q_V + 1}{2^{128} - q_M - q_V}$, which is optimal.

value dependent on both n -bit strings u and v . To verify the validity of a tag τ for message m , we know that we should rely on a TBC inversion call to avoid leaking useful information about the good tag if (m, τ) is actually an invalid message-tag pair. However, we also have to circumvent the difficulty encountered when dealing with the verification of HTBC tags, as discussed in the introduction. That means that we should use our second TBC call in such a way that the equality check must not directly involve either u or v (so, a half hash value) and an output of the TBC (thus requiring to rely on a backwards computation). As a result, we need a serial evaluation of the TBC calls to compute a tag, e.g., as $\tau = (F_{k_2}^v \circ F_{k_1}^v)(u)$, with an equality check performed in the middle to compare $F_{k_1}^v(u)$ and $F_{k_2}^{v,-1}(\tau')$ in the verification for a candidate tag τ' .

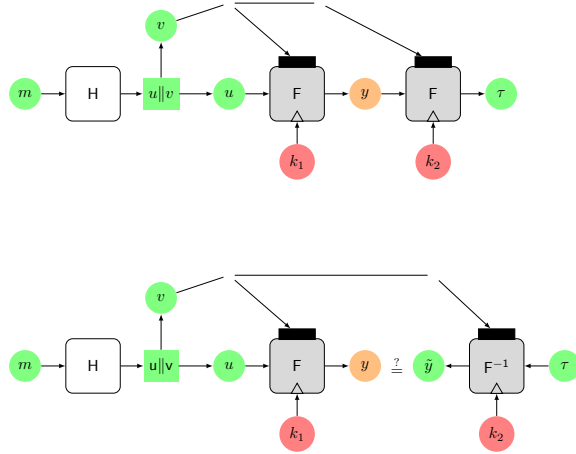


Fig. 2: LR-MAC2.

Up to switching the role between u and v , we now explain why $F_{k_2}^v \circ F_{k_1}^v$ is the good choice beyond the fact that using only a secret key primitive for the equality check allows us to learn what is computed and when, contrarily to the hash function. Since F_{k_2} and F_{k_1} are independent instances of the TBC, using the same tweak still enable us to use the output of one call against each other. The advantage of re-using the same tweak is twofold. First, we can easily partition all the queries with respect to v , the “lower half” of $H_s(m) = u||v$. This avoids considering many intermediate cases in the proof since with another computation of the form $F_{k_2}^u$, with tweak u , we would have to deal with many “crossed” half hashes like $u_1||v_1, u_1||v_2, u_2||v_1\dots$ Second, given v , $F_{k_2}^v \circ F_{k_1}^v$ is a permutation between u and τ . We will see that this also simplifies the distribution of the potential forgeries, which in turn improves the security bound.

The complete specification of LR-MAC2 is given in Algorithm 2.

Algorithm 2 The LR-MAC2 algorithm.

It uses a strongly protected TBC $F : \mathcal{K} \times \mathcal{TW} \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a hash function $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$.

- Gen
 - $k = (k_1, k_2) \xleftarrow{\$} \mathcal{K}^2$
 - $s \xleftarrow{\$} \mathcal{HK}$
 - $\text{Mac}_k(m)$:
 - $u||v \xleftarrow{\$} H_s(m)$
 - $y = F_{k_1}^v(u)$
 - $\tau = F_{k_2}^v(y)$
 - Return τ
 - $\text{Vrfy}_k(m, \tau)$:
 - $u||v \xleftarrow{\$} H_s(m)$
 - $y = F_{k_1}^v(u)$
 - $\tilde{y} = F_{k_2}^{v,-1}(\tau)$
 - If $y == \tilde{y}$ Return \top
 - Else Return \perp
-

Single key version. For the sake of simplicity, we described LR-MAC2 with two uniform and independent keys. However, with some extra work we can rely on a single secret key k . The careful reader will have noted that we cannot derive k_1 and k_2 from F_k as the TBC is only unpredictable with leakage, and not pseudorandom with leakage which is a controversy assumption in the standard model. So the keys k_1 and k_2 will not have the right distribution to assume that F_{k_1} and F_{k_2} are themselves unpredictable with leakage. Nevertheless, the other known alternatives work. We can either assume that the TBC supports one more bit-tweak, i.e., tweaks of size $n + 1$, or we can drop one bit of v in order to force in both cases the use of a domain separation bit in the tweaks between the two calls of the TBC. For completeness, we depict the construction where we compute $\tau = F_k^{v||1} \circ F_k^{v||0}(u)$ in the ePrint version of this work.

4.2 Half multi-collision resistance

It is natural to combine a double output length hash function $H_s(m) = h = u||v \in \{0, 1\}^{2n}$ with an (n, n, n) -TBC F to target high MAC unforgeability. But since the hash value h is too long it cannot be used directly, and we split it into two n -bit pieces u and v to input $F_k^v(u)$, for instance, where k is the n -bit key. The values u and v thus have distinct purposes in the design and it is equally natural to investigate the properties of both halves separately. At high level, this therefore requires that $H_s = H_s^1 || H_s^2$ comes with additional guarantees for H_s^1 and H_s^2 , which can be seen as a strengthened security requirement for H_s .

General purpose standard hash functions should be as close as possible to public random functions. Since our goal is to avoid relying on idealized assumption, we will use the intuition that H_s^1 and H_s^2 should be asymptotically as good

as H_s . That is, even if in practice hash functions are built for a precise output-length (e.g., 256 bits), it just corresponds to asking that H_s^1 and H_s^2 are 128-bit output-length secure (enough) hash functions.

Obviously, we cannot simply concatenate any two 128-bit output collision resistant hash functions to build H_s and hope boosting the collision resistance of the resulting construction, say, up to 2^{96} queries or more. Still, $H_s^1(m) = u$ and $H_s^2(m) = v$ have to behave as securely as possible and we can expect them to be range-oriented/second/...preimage resistant and/or collision resistant up to $2^{64\theta}$ queries, for some $\theta \in (0, 1)$ as close as possible to 1.

Building some good hash $H_s = H_s^1 || H_s^2$ enjoying additional “half-security” is out of the scope of this paper, and this additional requirement will be part of the discussion confirming the interest of the pragmatic LR-MAC1 construction at the end of the section. Yet, we see the independent security of H_s^1 and H_s^2 as a valuable asset and a natural target as many constructions might benefit from using the two halves u and v of h at different places in a cryptographic scheme, leading to a better overall security. We assume that it could be ensured by standard constructions, at the cost of slightly increased parameters.

Among the usual security guarantees of hash functions we can ask for the halves H_s^i , for $i = 1, 2$, we require the hardness of computing many collisions on H_s^2 . Bounding the probability of multi-collisions is reminiscent of deriving beyond-birthday security from symmetric (T)BC-based constructions of MAC or AEAD in the ideal-cipher model. In our construction, $H_s^2(m) = v$ is used as a tweak and we want to have a probability that decreases fast in the number of additional colliding v -values. By controlling the probability of multi-collisions on H_s^2 , we will have only few messages m_i in the adversarial view such that $H_s(m_i) = u_i || v$ for distinct u_i (otherwise we have a collision on the full H_s). That is, few possible preimages of F_k^v for any chosen v to compute a forgery.

We follow [7] to define a natural notion of multi-collision resistance before adapting it to the “half output” case as we require for our proofs.

Definition 8. *Let $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a hash function, and $\mu \geq 2$ be an integer. We say that H is μ -multi collision resistant if, for any efficient adversary A , there is a negligible function negl such that:*

$$\Pr_{\substack{s \leftarrow \mathcal{HK} \\ m_1, \dots, m_\mu \leftarrow A(s)}} [\forall i, j \in [\mu] : i \neq j \rightarrow m_i \neq m_j \wedge H_s(m_i) = H_s(m_j)] \leq \text{negl}(n).$$

We further say that H is multi collision resistant if there is a value $\mu \geq 2$ such that H is μ -multi collision resistant.

If H is modeled as a random oracle, H is clearly μ -multi collision resistant for any $\mu \geq 2$. In the standard model, H might only be multi collision resistant from some lower bound, say $\mu \geq 10$ or $\mu \geq n/4$. Still, if H is μ -multi collision resistant it is also μ' -multi collision resistant, for any $\mu' \geq \mu$. We elaborate more on the possible behavior of μ in relation with a decreasing family of negligible functions to derive more concrete bounds in Section 4.4, after the security analysis.

We now simply define the *half multi-collision resistance* of $H_s: \{0, 1\}^* \mapsto \{0, 1\}^{2n}$ as the multi-collision resistance of H_s^1 or H_s^2 , where $H_s = H_s^1 \parallel H_s^2$.

Definition 9. Let $H: \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$ be a hash function, and define $H^i: \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$, for $i = 1, 2$, such that for any $s \leftarrow \mathcal{HK}$ and any message $m \in \{0, 1\}^*$, we have $H_s(m) = H_s^1(m) \parallel H_s^2(m)$. We say that H is half multi-collision resistant for $i = 1, 2$ if H^i is multi-collision resistant. To avoid indexes, if $i = 1$, we also say that H is upper-half multi-collision resistant, and, if $i = 2$, that H is lower-half multi-collision resistant.

Assuming that H satisfies this notion is a falsifiable assumption.

4.3 Security analysis of LR-MAC2

Unbounded leakage specification. Before assessing the sUF-L2 security of LR-MAC2, we have to define the leakage that an adversary can collect:

- $L_M(m; k)$ returns $(y, \text{LEval}(v, u; k_1), \text{LEval}(v, y; k_2))$, where $u \parallel v \stackrel{r}{\leftarrow} H_s(m)$ and the intermediate value $y = F_{k_1}^v(u)$ is given.
- $L_V(m, \tau; k)$ returns $(y, \tilde{y}, \text{LEval}(v, u; k_1), \text{LInv}(v, \tau; k_2))$, where $u \parallel v \stackrel{r}{\leftarrow} H_s(m)$, and the intermediate values $y = F_{k_1}^v(u)$ and $\tilde{y} = F_{k_2}^{-1, v}(\tau)$ are given.

Our main result for LR-MAC2 is then formalized by the following theorem:

Theorem 2. Let $H: \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$ be a $(t_1, \epsilon_{\text{CR}(2n)})$ -collision resistant hash function and $(t_1, \epsilon_{\mu\text{-CR}(n)})$ - μ -lower-half multi-collision resistant, for some $\mu \geq 2$. Let F be a $(q_L, q_{\text{Eval}}, q_{\text{Inv}}, t_2, \epsilon_{\text{sUP-L2}})$ -strongly unpredictable with leakage TBC, then, LR-MAC2 is $(q_L, q_M, q_V, t, \epsilon)$ -sUF-L2 with:

$$\epsilon \leq \epsilon_{\text{CR}(2n)} + \epsilon_{\mu\text{-CR}(n)} + 2\mu q_V \epsilon_{\text{sUP-L2}},$$

where $q_{\text{LF}} \leq q_L$, $q_{\text{Eval}} \leq q_M + q_V + 1$, $q_{\text{Inv}} \leq q_V$, $t_1 \leq t + (q_M + q_V + 1)t_H + 2(q_M + q_V)(t_F + t_{\text{L}(F)})$, and $t_2 \leq t + (q_M + j + 1)t_H + (q_M + q_V + 1)(t_F + t_{\text{L}(F)})$.

Below, we give a detailed sketch of the proof. We defer the full proof to the ePrint version of the paper and discuss the security bound in Section 4.4.

Proof (Sketch). To simplify our exposition, we assume the adversary gets the unbounded leakage without being explicit. Let (m, τ) be the first valid message-tag pair involved in a verification query which can be used by the adversary as a forgery. That is, the pair (m, τ) is *fresh* at the time of that verification query in the sense that it was never involved in an earlier verification query and (m, τ) is not the result of any previous tag-generation query for m . Given such a fresh and valid pair, let $u \parallel v \stackrel{r}{\leftarrow} H_s(m)$, $y = F_{k_1}^v(u)$ and $\tilde{y} = F_{k_2}^{-1, v}(\tau)$. Below, we consider different cases depending on whether the TBC triples (v, u, y) for k_1 and (v, \tilde{y}, τ) for k_2 have already been defined before the verification query of (m, τ) or not. We note that if the computation of $\tau = F_{k_2}^v(\tilde{y})$ occurred before (necessarily in a tag-generation query) our second triple is already defined.

- F_1 The triples (v, u, y) for k_1 and (v, \tilde{y}, τ) for k_2 have not been defined yet.
- F_2 The triple (v, u, y) for k_1 has already been defined, while (v, \tilde{y}, τ) for k_2 has not been defined yet.
- F_3 Conversely, the triple (v, u, y) for k_1 has not been defined yet, while (v, \tilde{y}, τ) for k_2 has already been defined.
- F_4 Both triples (v, u, y) for k_1 and (v, \tilde{y}, τ) for k_2 have already been defined. We split that case into three sub-cases:
 - F_{41} Both triples (v, u, y) for k_1 and (v, \tilde{y}, τ) for k_2 have been defined in the same tag-generation or verification query.
 - F_{42} The triple (v, u, y) for k_1 was defined in some query (of any type) precedent to the one in which (v, \tilde{y}, τ) was defined for k_2 .
 - F_{43} The triple (v, \tilde{y}, τ) for k_2 was defined in some query (of any type) precedent to the one in which (v, u, y) was defined for k_1 .

We now give intuition about how we bound the probability of these events. The main difficulty arises in F_4 , where we would like to rely on the sUP-L2 security of the TBC. In F_{42} (resp., F_{43}), we would like to use the second-call triple (v, \tilde{y}, τ) (resp., the first-call triple (v, u, y)) as our winning prediction against the TBC. The trouble comes while we should avoid querying an evaluation/inversion that settles all the triples associated to (m, τ) as defined, in which case the attack will not be successful (to win the triple must not be defined). Fortunately, half-multi-collisions on v allow detecting the potential triples for which we have to be careful. Since their amounts are bounded by μ , unless \mathbf{H} is not lower-half multi-collision resistant, we only get an additional factor μ in $\epsilon_{\text{sUP-L2}}$ for each case.

- F_1 We build a sUP-L2 adversary against the TBC, conceptually against the “second call” \mathbf{F}_{k_2} . It picks $k_1 \xleftarrow{\$} \mathcal{K}$ for itself and simulates the computation of \mathbf{F}_{k_2} by querying its own TBC oracle. In the verification query (m, τ) , it computes y itself and outputs (v, y, τ) as its prediction. Given that (m, τ) , the probability that this event occurs is bounded by $\epsilon_{\text{sUP-L2}}$.
- F_2 We build a sUP-L2 adversary against the TBC as in the case of F_1 .
- F_3 We build a sUP-L2 adversary against the TBC, conceptually against the “first call” \mathbf{F}_{k_1} . It picks $k_2 \xleftarrow{\$} \mathcal{K}$ for itself and simulates the computation of \mathbf{F}_{k_1} by querying its own TBC oracle. In the verification query (m, τ) , it computes \tilde{y} itself and outputs (v, u, \tilde{y}) as its prediction. Given that (m, τ) , the probability that this event occurs is bounded by $\epsilon_{\text{sUP-L2}}$.
- F_4 This case is less straightforward. We deal with its sub-cases independently.
 - F_{41} Let (m', τ') be the message-tag pair involved in the earlier query that defines (v, u, y) and (v, \tilde{y}, τ) . Then, $\tau' = \tau$ (regardless of the query type).
 - F_{411} If $m^i = m$, it contradicts the freshness of (m, τ) . This case is void.
 - F_{412} If $m^i \neq m$, since the defined triples implies $u \parallel v \xleftarrow{\tau} \mathbf{H}(m')$, we found a collision. The probability that this event occurs is bounded by ϵ_{CR} .
 - F_{42} Let (m^i, τ^i) be the message-tag pair involved in the earlier query that defines (v, u, y) for k_1 and (m^j, τ^j) be the message-tag pair involved in a subsequent query that defines (v, \tilde{y}, τ) for k_2 , and both (regardless of their respective type of query) occurring before our (m, τ) . By the

meaning of the triples, we have $u \| v \stackrel{n}{\leftarrow} \mathbf{H}(m^i)$ and $\tau^j = \tau$, and we can assume that $m^i = m$ as we already dealt with collisions. However, we note that $\tau^i \neq \tau$ and $m^j \neq m$ as otherwise it would contradict the freshness of (m, τ) . So (m^i, τ^i) must be an invalid pair in a verification query as well as (m^j, τ^j) (if it is valid it can only be a tag-generation query by definition of (m, τ) , but then $u^j = \mathbf{F}_{k_1}^{v, -1} \circ \mathbf{F}_{k_2}^{v, -1}(\tau^j) = \mathbf{F}_{k_1}^{v, -1} \circ \mathbf{F}_{k_2}^{v, -1}(\tau) = u$).¹¹ This implies that we already have a lower-half 2-multi-collision on v . Indeed, we see that $\mathbf{H}_s^2(m^i) = \mathbf{H}_s^2(m^j)$ with $m^i \neq m^j$.

Then, to avoid defining the second-triple (v, \tilde{y}, τ) we must use it as a winning prediction against the TBC at the time the query (m^j, τ^j) is made. Moreover, the situation will be the same with any subsequent query (m^j, τ^j) before (m, τ) where the second-triple is given by (v, \tilde{y}, τ) , i.e., $v^l = v$, $\tau^l = \tau$ and distinct m^l (otherwise it is a repeating query).

We can bound the probability of all these events given (m^i, τ^i) . Each time a new verification query is made for some (m^l, τ^l) such that $m^i \neq m^l$ and $v^i = v^l$ (we include $l = j$ here), for the at most $\mu - 1$ possible cases, we make an hybrid on the ordered such l , and a reduction to the sUP-L2 experiment against the TBC in the “second call” by picking k_1 and outputting the prediction (v^i, y^i, τ^l) (while the TBC oracle is called for all the previous verification queries to emulate the second MAC call).

F_{43} We are in the same case as case F_{42} with the difference that the query (m^j, τ^j) defining (v, \tilde{y}, τ) for k_2 is made before the query (m^i, τ^i) defining (v, u, y) for k_1 . Since our argument showing that these queries in F_{42} are of verification type and the input pairs are invalid does not depend on which query happens before, we are exactly in the dual situation where we somehow switched the first-call and the second-call.

We can bound the probability of all these events given (m^j, τ^j) . Each time a new verification query is made for some (m^l, τ^l) such that $m^l \neq m^j$ and $v^l = v^j$ (we include $l = i$ here), for the at most $\mu - 1$ possible cases, we make an hybrid on the ordered such l , and a reduction to the sUP-L2 experiment against the TBC in the “first call” by picking k_2 and outputting the prediction (v^j, u^l, \tilde{y}^j) (while the TBC oracle is called for all the previous verification queries to emulate the second MAC call).

In the full proof, we first deal with the collision resistance of \mathbf{H}_s and the μ -multi-collision resistance of \mathbf{H}_s^2 . For the events F_{42} and F_{43} we of course do not know when we will be in face of the queries (m^i, τ^i) and (m^j, τ^j) which could define the triples of the the first potential forgery (m, τ) . However, by considering all the verification queries as potentially being (m^i, τ^i) or (m^j, τ^j) that defines the first triple of (m, τ) , we cover all the cases with a term like $2\mu q_V \epsilon_{\text{sUP-L2}}$. That way, we also cover the case that (m, τ) might be any of the verification queries. Indeed, once we analyze a potential (m^i, τ^i) or (m^j, τ^j) , we will consider all the

¹¹ The choice to reuse v as the only tweak of the TBC in our design is crucial. With distinct tweaks the security bound will be much more loose. Among other advantages, here, we only have to deal with the verification queries.

next verification queries sharing v^i or v^j among which all the possible (m, τ) lie when F_{42} or F_{43} occur. In the remaining case, it is straightforward to deal with F_1 , F_2 and F_3 . Overall, we find $[2(\mu - 1)q_V + q_V + 1]\epsilon_{\text{sUP-L2}}$ for the TBC term.

4.4 Deriving concrete bounds

Deriving concrete bounds from asymptotic security is not an easy task in the standard model. However, we can argue why it is realistic to assume the existence of building blocks which will confer high security to LR-MAC2. Since the sUP-L2 of the TBC can be argued in the same lines as in the rest of the paper, we next focus on the stronger properties that we require for the hash function.

Multi-collision resistant hashing. The security bound shows that it is reasonable to take $\mu \in O(n)$ without blowing up the term related to the unpredictability of the TBC. Moreover, increasing μ can only decrease the probability of finding more and more multi-collisions. Indeed, μ -multi collision resistance implies $(\mu + 1)$ -multi collision resistance. And, even if the definition does not imply that the negligible probability decreases when μ increases, it is reasonable to assume so for “good” hash functions, as for SHA2 or SHA3.

In the ROM, it is a well-known fact that n -multi collision resistance allows up to $q \approx 2^n/n$ hash queries. Of course, we will not assume in the standard model that H remains secure up to that number of evaluations. However, it might be the case that H is $3n$ -multi collision resistant up to $q \approx 2^{96}$, which is much less demanding and fully compatible with our other terms in our security bound.

As another justification, Berti et al. [9, Lemma 5] proved that the Merkle-Damgård iteration of the Hirose’s DBL compression function [26] gives rise to a half-multi-collision resistant hash function in the ideal cipher model.

Multi-collision independence. By more closely mimicking the behavior of random functions, we derive another relaxed notion that could be realized without idealized assumption on half outputs of H . This is a stronger definition than the (half-) multi-collision resistance notion, but it offers a way to more accurately model a decreasing negligible functions’ family depending on the multi-collision parameter μ . This allows us to soundly analyze how the negligible functions related to a (half-) μ -multi-collision resistant hash function H might become better and better as μ grows, as expected from any “good” hash function, and without relying on any ideal argument. Of course, the decreasing rate in μ can be much slower than a truly random function but our security bounds of LR-MAC2 shows how comfortable we are to increase μ even in a leakage setting.

First, we recall a result of [35] showing that for a random function with 2^n possible outputs, the multi-collision event $\text{MultiColl}(2^n, q) \geq \mu$ that at least $\mu \geq 2$ inputs give a same output among a total of q function evaluations satisfies:

$$\Pr[\text{MultiColl}(2^n, q) \geq \mu] \leq \frac{1}{2^{n(\mu-1)}} \binom{q}{\mu}.$$

Since the right-hand side is also bounded by $q^\mu/2^{n(\mu-1)}$, we can simplify this result by saying that any additional collision comes at a marginal probability of $\approx q/2^n$ starting from the standard collision bound $q^2/2^n$, where $\mu = 2$.

Definition 10. Let $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$ be a hash function. We say that H is θ -powerwise collision independent if, for $s \leftarrow \mathcal{HK}$, the probability that an efficient adversary $A(s)$ making at most q hash evaluations computes at least μ distinct $m_1, \dots, m_\mu \in \{0, 1\}^*$ such that $H_s(m_1) = \dots = H_s(m_\mu)$ is bounded by:

$$2^n \left(\frac{q^{1+\theta}}{2^n} \right)^\mu,$$

with $\mu \geq 2$, $0 \leq \theta \leq 1$. The factor $q^{1+\theta}/2^n$ is the marginal collision probability and θ is called the non-idealization power. We say that H is multi-collision independent if it is θ -powerwise collision independent for some $\theta \in (0, 1)$.

If the non-idealization power is null, the hash function is a random function and the marginal collision probability matches the one discussed above.

We stress that finding a μ -multi-collision for any $\mu \geq 2$ gives the same chance of finding “one more” collision in the sense of a $(\mu + 1)$ -multi-collision. Said otherwise, the marginal collision probability is independent of the number of multi-collisions already found, hence the name. For instance, finding more multi-collisions does not help in finding even more multi-collisions. This might be seen as a strong collision resistant flavor as the hash function should resist further collision attacks even if some previous attacks already succeeded. However, standard hash functions should fulfill such a collision resilient condition. At least, they have to offer a graceful degradation and, in particular, given a collision it should not become easy to find some others. Still, we require here that the degradation factor be constant, given the security parameter n and the non-idealization power θ . But we might only have a middle non-idealization power $\theta = 1/2$. In that case, the collision probability (i.e., the 2-multi-collision) is only upper-bounded by $q^3/2^n$ ensuring collision security up to $q \approx 2^{42}$, for $n = 128$. Nevertheless, we can reach the 20-multi-collision security up to $q \approx 2^{81}$. By letting $\theta = 1/\rho$, we can reach the n -multi-collision security up to $q \approx 2^{n\rho/(1+\rho)}$.

We now simply define *half-multi-collision independence* of $H_s : \{0, 1\}^* \mapsto \{0, 1\}^{2n}$ as the multi-collision independence of H_s^1 or H_s^2 , where $H_s = H_s^1 \parallel H_s^2$.

Definition 11. Let $H : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$ be a hash function, and define $H^i : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^n$, for $i = 1, 2$, such that for any $s \leftarrow \mathcal{HK}$ and any message $m \in \{0, 1\}^*$, we have $H_s(m) = H_s^1(m) \parallel H_s^2(m)$. We say that H is half-multi-collision independent for $i = 1, 2$ if H^i is multi-collision independent. To avoid indexes, if $i = 1$, we also say that H is upper-half-multi-collision independent, and, if $i = 2$, that H is lower-half-multi-collision independent.

Armed with this stronger theoretical background, we can re-evaluate the security of LR-MAC2 in the light of the non-idealization power θ , assuming that H is lower-half-multi-collision independent (where H_s^2 outputs tweaks). Let us

first recall the security bound $\epsilon_{\text{CR}(2n)} + \epsilon_{\mu\text{-CR}(n)} + 2\mu q_V \epsilon_{\text{sUP-L2}}$ of Theorem 2, and assume that $\epsilon_{\text{sUP-L2}}$ supports 2^{85} (online) queries, for $n = 128 = 2^7$. Setting $q_V = 2^{75}$ then allows us to take $\mu = 2^8$. That means that if the number of evaluation of H remains below 2^{90} , the non-idealization power could be $\theta = 2/5$ for H_s^2 since $\epsilon_\mu \leq 2^{128}(2^{90 \cdot 7/5}/2^{128})^{256} < 2^{-128}$, even if H_s^2 would only be (2-multi-) collision resistant up to $2^{63/(1+\theta)} = 2^{45}$ hash evaluations.

So overall, we conclude this section by observing that LR-MAC2 could provide strong leakage resilience guarantees. Yet, on the one hand, it requires stronger assumptions which, despite not idealized, are less standard than the collision resistance that LR-MAC1 requires. On the other hand, it also requires two calls to a TBC with n -bit tweaks which, as mentioned in introduction, should in general be more expensive than one call to a TBC with $2n$ -bit tweaks if specialized constructions can be used. We note there are constructions that achieve the same leakage-resilience guarantees without half collision resistance, but the best solution we are aware of (given in the ePrint version of the paper) requires three TBC calls for this purpose. So despite theoretically interesting, we believe these results also amplify the pragmatic interest of using large tweaks.

5 Analysis of HTBC

In this section, we finally show that the popular Hash-then-PRF MAC construction is provable without idealized assumptions in a leakage setting. For consistency with the LR-MAC1 and LR-MAC2 constructions, we focus on the Hash-then-TBC scheme described in Figure 3 and specified in Algorithm 3.

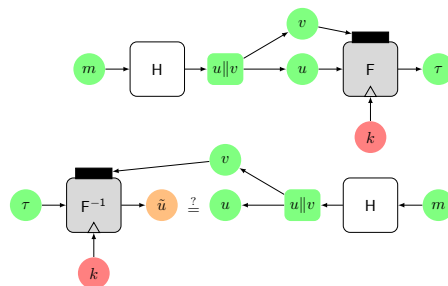


Fig. 3: Hash-then-TBC.

5.1 A first analysis in the standard model

The security proof requires the hash H to be collision resistant and range-oriented preimage resistant. A stronger (and less standard) assumption is *weak image set computable*, that is, there exists a PPT algorithm M that outputs the set

Algorithm 3 The HTBC algorithm.

It uses a strongly protected TBC $F : \mathcal{K} \times \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a hash function $H : \mathcal{HK} \times \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$.

- Gen
 - $k \xleftarrow{\$} \mathcal{K}$
 - $s \xleftarrow{\$} \mathcal{HK}$
 - $\text{Mac}_k(m)$:
 - $u \| v = H_s(m)$
 - $\tau = F_k^v(u)$
 - Return τ
 - $\text{Vrfy}_k(m, \tau)$:
 - $u \| v = H_s(m)$
 - $\tilde{u} = F_k^{v,-1}(\tau)$
 - If $\tilde{u} == u$ Return \top
 - Else Return \perp
-

$\mathcal{WP}(s, A)$ defined in Eq. (1) for any s and A . By Lemma 1, for range-oriented preimage resistant hash functions it is feasible to output the set $\mathcal{WP}(s, A)$ in PPT. We will serve more intuitions on this assumption at the end of this section.

The assumption of *weak image set computability* admittedly renders the concrete security analysis hard to interpret. As a result, below we eschew the concrete security approach (which is followed by Theorems 1 and 2) in favor of the *asymptotic approach*. We justify this choice by the fact that the goal of this section is to show a theoretical possibility, potentially opening a path towards more advanced and concrete analyzes in the future.

Theorem 3. *Let the hash function H be collision resistant, range-oriented preimage resistant, and weak image set computable. Let F be a TBC that is sUF-L2. Then, HTBC is sUF-L2-secure in the unbounded leakage model.*

Proof. In the unbounded leakage model, all the intermediate computations leak (except the long-term key). Wlog, we assume that the forgery adversary A has been “normalized”. That is, before making every oracle query, A outputs the list of all images $\mathcal{S}(s) \subseteq \{0, 1\}^n$ for which she knows the preimage. Clearly, this cannot decrease the success probability of A . Moreover, this only induces a polynomial blow-up in A ’s running time: $\mathcal{S}(s)$ must be polynomial, as otherwise A gets a superpolynomial number of input/output relations $H_s(m) = y$ and would be able to break the collision resistance of H .

To prove the theorem, we follow a sequence of games.

Game 0. Let Game 0 be the sUF-L2 game where the adversary A^L tries to produce a forgery when she plays against HTBC. Let E_0 be the event that the adversary wins the game; that is, the output of the game is 1.

Game 1. Game 1 is Game 0, except that we abort if there is a full collision in the hash function. Clearly, Game 0 and Game 1 are identical if the following event

HC (*Hash collision*) does not happen: $HC := \{\exists i, j \in \{1, \dots, q_M\} \cup \{1, \dots, q_V + 1\} \text{ with } i \neq j \text{ s.t. } h^i = h^j \text{ and } m^i \neq m^j\}$ (see Footnote 9). Since H is collision resistant, we have $\Pr[\text{Game 0} \Rightarrow 1] \approx \Pr[\text{Game 1} \Rightarrow 1]$.

Linking to unpredictability. To complete the argument, we exhibit a predictor B against F_k using an adversary A that forges in Game 1. In detail, B picks a hash seed s uniformly at random, passes s to A and simulates Game 1 against A . B also picks an index $\ell \xleftarrow{\$} \{1, \dots, q_v + 1\}$ at random. Assume that A outputs the forgery (m^*, τ^*) at the end of the game, where $H_s(m^*) = u^* \| v^*$.

Now, if Game 1 outputs 1, i.e., $\text{VrfyL}_k(m^*, \tau^*)$ returns \top at the end, then the probability that the ℓ -th verification query $\text{VrfyL}_k(m^{(\ell)}, \tau^{(\ell)})$ constitutes the first time $(F_k^{v^*})^{-1}(\tau^*)$ is queried is at least $1/q_v$. It can be seen that either of the following bad events necessarily occurred during the game:

- Case 1: intuitively, at some time, the oracle $(F_k^{v^*})^{-1}(\tau^*)$ returns a target u^* and A later solves its preimage $H_s(m^*) = u^* \| v^*$; or,
- Case 2: intuitively, at some time, A has been aware of the relation $H_s(m^*) = u^* \| v^*$, and the oracle $(F_k^{v^*})^{-1}(\tau^*)$ later returns a target $u' = u^*$.

But B cannot predict which case will be encountered. Fortunately, in either case the number of candidates for $(F_k^{v^*})^{-1}(\tau^*)$ is polynomial, and thus B can simply take a union. In more detail:

If the first case is encountered, then the concatenation $u^* \| v^*$ of the returned target $u^* = (F_k^{v^*})^{-1}(\tau^*)$ and the tweak v^* necessarily falls in the set $\mathcal{WP}(s, A)$ (otherwise A won't be able to solve the preimage). Then, since we assumed that the set $\mathcal{WP}(s, A)$ is computable given s and A , the predictor B could make a guess u' among $\mathcal{WP}(s, A)$ and take (v^*, u', τ^*) as a forgery for F_k .

If the second case is encountered, then as we assumed that A always outputs the list $\mathcal{S}(s)$ of all images that she knows the preimage, the predictor B could make a guess u' among $\mathcal{S}(s)$ and takes (v^*, u', τ^*) as a forgery for F_k .

By the above, the predictor B first runs the algorithm M to obtain $\mathcal{WP}(s, A)$, and then reads the set $\mathcal{S}(s)$ from A 's outputs. B then picks $u' \xleftarrow{\$} \mathcal{WP}(s, A) \cup \mathcal{S}(s)$ uniformly and outputs (v^*, u', τ^*) as a forgery for F_k .

Summing over the two cases, the success probability of B is at least:

$$\begin{aligned} & \left(\frac{\Pr[\text{Game 1} \Rightarrow 1]}{q_v + 1} \right) \cdot \left(\frac{\Pr[\text{Case 1}]}{|\mathcal{WP}(s, A)| + |\mathcal{S}(s)|} + \frac{\Pr[\text{Case 2}]}{|\mathcal{WP}(s, A)| + |\mathcal{S}(s)|} \right), \\ & \geq \left(\frac{\Pr[\text{Game 1} \Rightarrow 1]}{q_v + 1} \right) \cdot \left(\Pr[\text{Case 1}] + \Pr[\text{Case 2}] \right) \cdot \frac{1}{|\mathcal{WP}(s, A)| + |\mathcal{S}(s)|}, \\ & = \frac{\Pr[\text{Game 1} \Rightarrow 1]}{(q_v + 1)(|\mathcal{WP}(s, A)| + |\mathcal{S}(s)|)}. \end{aligned} \tag{3}$$

By the unpredictability assumption on F , the final term $\Pr[\text{Game 1} \Rightarrow 1]/(q_v + 1)(|\mathcal{WP}(s, A)| + |\mathcal{S}(s)|)$ should be negligible, which means $\Pr[\text{Game 1} \Rightarrow 1] = \text{negl}(n)$. This complete the proof. \square

Discussion. To facilitate understanding, we provide further intuition on the proof of Theorem 3. A first impression may suggest the infeasibility to decide and output the weak set $\mathcal{WP}(s, A)$ since the seed space is exponential, which (seems) to gain support from the formalism of collision resistant hashing. However, this intuition is incorrect: the seed space being exponential does not *necessarily* mean the computation is infeasible. In particular, it may be easy to derive $\mathcal{WP}(s, A)$ from the seed, or even seed-independent (i.e., the algorithm does not necessarily need to keep a table for all the seeds) as we now detail.

First, assume using a keyed random oracle $\text{RO} : \mathcal{HK} \times \{0, 1\}^* \rightarrow \{0, 1\}^{2n}$ as the hash function. In this case, $\Pr_{A(s, y) \Rightarrow m} [\text{RO}_s(m) = y]$ is negligible for all $y \in \{0, 1\}^{2n}$, meaning that $\mathcal{WP}(s, A) = \emptyset$. By this, the probability that Case 1 in the proof occurs is negligible (which is indeed the case, as the probability is in $O(q/2^n)$, with q the number of adversarial RO queries). On the other hand, the “normalized” adversary A simply outputs its RO query/response list before the ℓ -th verification query, and this could be used to predict $(F_k^v)^{-1}(\tau)$.¹²

Let us now assume using a keyed SHA3 variant (denoted sha) producing 256-bit hash digests. According to known cryptanalytic results, for any A with running time t , $\Pr_{A(s, y) \Rightarrow m} [\text{sha}_s(m) = y] = O(t/2^{256})$ for all $y \in \{0, 1\}^{256}$, meaning that $\mathcal{WP}(s, A)$ remains \emptyset in some sense.

On the one hand, it seems that the number of SHA3 input/output pairs that could be derived in time t is only $O(t)$, and the “normalized” adversary A simply outputs a list of $O(t)$ input/output pairs. On the other hand, the double-block-length hash function Tandem-DM has a “weak” image 0 that is easily invertible [3]. If Tandem-DM is used, then $\mathcal{WP}(s, A) = \{0\}$, the size of which is just 1. Further consider the keyed hash $\text{KeyedTDM}_s(M) = \text{Tandem-DM}(M) \oplus s$. Despite being contrived, for this instructive example the weak set $\mathcal{WP}(s, A) = \{s\}$ which is seed-dependent but very easy to derive.

Note that we do not need the hash function to be non-malleable. The proof follows even if the hash admits the weakness that $H_s(m \oplus \delta_1) = H_s(m \oplus \delta_1) \oplus \delta_2$ for a *specific pair* of differences (δ_1, δ_2) and *any* m : in this case, the number of “known” input/output pairs remains in $O(t)$. We additionally remark that:

- the weak image set computability assumption doesn’t contradict the range-oriented preimage resistant assumption, because the latter concerns with random images that are unlikely to fall in $\mathcal{WP}(s, A)$;
- this assumption doesn’t contradict the collision resistant assumption either, because weak images don’t necessarily have multiple preimages.

Finally, we note that our proof approach easily extends to the constructions of Hash-then-block cipher of [11] and Hash-then-PRF (enhanced with the leakage-resilient value comparison tricks [17]). It is also intriguing to ask whether more “common” assumptions on hashing could suffice, e.g., even the UCE hashing [5] or some weak form of correlation intractable hashing [13, 12]. We leave the investigation of these alternative approaches for future investigations.

¹² The latter trick was also used in [8].

5.2 Towards another analysis

As an opening, we eventually discuss whether adaptating the security proofs of [8] could lead to another analysis. The simplest solution for this purpose is to find the appropriate assumption on H that could hold in the standard model to replace the random oracle. This is also what is done in the previous analysis but we next suggest another (stronger) possibility that is simple to express.

Essentially, what we need from the hash function H is that any *targeted* image y that has never been defined yet by computing $H(x)$ for some x is still unattainable. Of course, by computing $H(x)$ for a fresh input x we always attain a new image with overwhelming probability, unless we find a collision. Obviously, all the undefined images cannot be our targets. So, the main difficulty is to define a natural meaning of *targets* that can be chosen *adaptively* by the adversary. This is where the new definition will anyway be less standard than usual pre-image resistance. But this is precisely what we need in the security proof of HTBC to adapt [8]. Indeed, let $\tilde{u} = F_k^{v,-1}(\tau)$, and let us consider again the cases where either $\tilde{u}\|v$ was already computed from $H(m)$ for some m or it was not. In the first case, if we know that $H(m) = \tilde{u}\|v$ has already been computed in the unforgeability game, we can easily build an adversary against TBC and compute a prediction. In the second case, if we know that $\tilde{u}\|v$ was never defined as an output of H , then we want to define $\tilde{u}\|v$ as a target. That means that if we know which hashing has already been computed, we can adaptively define a set of targets of polynomial size. Since the number of TBC calls is bounded by the number of queries against the unforgeability of HTBC, knowing which value is already a defined image of H is precisely what we need, and what is done in the random oracle model. Still, we do not need to have uniformly random output.

As a consequence, as long as H is half-multi-collision independent or even half-multi-collision resistant as defined in Section 4, since the proof in [8] already needs such kind of assumption even if it comes for free from the random oracle, the proof can go through if we require H to satisfy the notion of *essentially pre-image resistance* which precisely models the hardness of computing a pre-image of a single target among any polynomial set of targets that can be adaptively updated by the adversary before trying to find it.

Essentially-pre-image resistance would simply model which hashing is already defined or not and maintain the list of targets that the adversary announced she will try to attack. To do so, even if the hash function is a public-key primitive, we should demand that computing hash values requires “querying” H even if the hash values are computed *faithfully* (without any other restriction about the distribution of the outputs). Hashing becomes interactive in that model but *only* for the sake of defining fresh targets and nothing else. The interaction does not occur in practice: it is just a way to soundly define what are the fresh targets and to avoid trivial attacks in the standard model.

It is easy to see that for any Turing Machine A without interaction with H , we can build a Turing Machine A' that explicitly returns all the inputs it uses for H , and directly after the output of that computation. Based on this discussion, we can see that the security of HTBC can be essentially good.

6 Conclusions

This paper describes and analyzes three leakage-resilient MACs that can be proven based on minimum physical assumptions (namely the unpredictability of a TBC with leakage) and more or less standard (anyway non idealized) black box assumptions for their hash part. We believe these results make an important step in building cryptographic primitives enabling strong security against side-channel attacks at limited cost, by leveraging the leveled implementation concept. The constructions we analyze range from pragmatic to theoretically more involved. The difficulty to maintain tight bounds with minimum assumptions when getting rid of the design twists used in our pragmatic solution (i.e., large tweaks for the TBC and a verification process that checks the inverse TBC against a constant) confirms that selecting good assumptions is critical for such constructions, as generally observed in cryptography but possibly even more with leakage. Avoiding idealized assumptions is important in this respect, since they make it impossible to translate security proofs and bounds into concrete requirements for the engineers implementing the corresponding schemes.

These results suggest different tracks for further investigations. In view of the limited overheads that its design twists imply, further instantiating the LR-MAC1 construction and designing a TBC with large tweaks that is suited to side-channel countermeasures appears as a practically relevant goal. In parallel, investigating whether the performances' and assumptions' gaps between LR-MAC1 and the LR-MAC2 or HTBC constructions can be tightened is an interesting question as well (for example using some of the directions outlined at the end of Section 5). Eventually, extending our results to sponge-based constructions (e.g., ISAP-MAC [15]) is another challenge that would deserve attention: we expect that such an analysis will require some adaptation and different assumptions. For example, secure authentication in this case cannot directly work in the unbounded leakage model, as per the recent work of Dobraunig and Menning [17]. The same holds for extending the quest for leakage-resilient constructions without idealized assumptions from MACs to (authenticated) encryption schemes.

Acknowledgments

We thank the ASIACRYPT 2021 reviewers for their insightful feedback. Francesco Berti was funded by the Emmy Noether Program FA 1320/1-1 of the German Research Foundation (DFG). Chun Guo was supported in parts by the Program of Taishan Young Scholars of the Shandong Province, the Program of Qilu Young Scholars (Grant No. 6158008996 3177) of Shandong University, the National Natural Science Foundation of China (Grant No. 62002202), and the Shandong Nature Science Foundation of China (Grant No. ZR2020MF053). Thomas Peters and François-Xavier Standaert are research associate and senior research associate of the Belgian Fund for Scientific Research (F.R.S.-FNRS). This work has been funded in parts by the EU through the ERC project SWORD (724725).

References

1. J. H. An and M. Bellare. Constructing vil-macs from fil-macs: Message authentication under weakened assumptions. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 252–269. Springer, 1999.
2. E. Andreeva and M. Stam. The symbiosis between collision and preimage resistance. In L. Chen, editor, *Cryptography and Coding - 13th IMA International Conference, IMACC 2011, Oxford, UK, December 12-15, 2011. Proceedings*, volume 7089 of *Lecture Notes in Computer Science*, pages 152–171. Springer, 2011.
3. F. Armknecht, E. Fleischmann, M. Krause, J. Lee, M. Stam, and J. P. Steinberger. The preimage security of double-block-length compression functions. In D. H. Lee and X. Wang, editors, *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, volume 7073 of *Lecture Notes in Computer Science*, pages 233–251. Springer, 2011.
4. M. Azouaoui, D. Bellizia, I. Buhan, N. Debande, S. Duval, C. Giraud, É. Jaulmes, F. Koeune, E. Oswald, F. Standaert, and C. Whittall. A systematic appraisal of side channel evaluation strategies. In *SSR*, volume 12529 of *Lecture Notes in Computer Science*, pages 46–66. Springer, 2020.
5. M. Bellare, V. T. Hoang, and S. Keelveedhi. Instantiating random oracles via uces. In R. Canetti and J. A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, volume 8043 of *Lecture Notes in Computer Science*, pages 398–415. Springer, 2013.
6. D. Bellizia, O. Bronchain, G. Cassiers, V. Grosso, C. Guo, C. Momin, O. Pereira, T. Peters, and F. Standaert. Mode-level vs. implementation-level physical security in symmetric cryptography - A practical guide through the leakage-resistance jungle. In *CRYPTO (1)*, volume 12170 of *Lecture Notes in Computer Science*, pages 369–400. Springer, 2020.
7. I. Berman, A. Degwekar, R. D. Rothblum, and P. N. Vasudevan. Multi-collision resistant hash functions and their applications. In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, pages 133–161. Springer, 2018.
8. F. Berti, C. Guo, O. Pereira, T. Peters, and F. Standaert. Strong authenticity with leakage under weak and falsifiable physical assumptions. In *Inscrypt*, volume 12020 of *Lecture Notes in Computer Science*, pages 517–532. Springer, 2019.
9. F. Berti, C. Guo, O. Pereira, T. Peters, and F. Standaert. Tedt, a leakage-resist AEAD mode for high physical security applications. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2020(1):256–320, 2020.
10. F. Berti, F. Koeune, O. Pereira, T. Peters, and F. Standaert. Ciphertext integrity with misuse and leakage: Definition and efficient constructions with symmetric primitives. In *AsiaCCS*, pages 37–50. ACM, 2018.
11. F. Berti, O. Pereira, T. Peters, and F. Standaert. On leakage-resilient authenticated encryption with decryption leakages. *IACR Trans. Symmetric Cryptol.*, 2017(3):271–293, 2017.
12. R. Canetti, Y. Chen, L. Reyzin, and R. D. Rothblum. Fiat-shamir and correlation intractability from strong kdm-secure encryption. In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International*

- Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 91–122. Springer, 2018.
13. R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. *J. ACM*, 51(4):557–594, 2004.
 14. J. P. Degabriele, C. Janson, and P. Struck. Sponges resist leakage: The case of authenticated encryption. In *ASIACRYPT (2)*, volume 11922 of *Lecture Notes in Computer Science*, pages 209–240. Springer, 2019.
 15. C. Dobraunig, M. Eichlseder, S. Mangard, F. Mendel, B. Mennink, R. Primas, and T. Unterluggauer. Isap v2.0. *IACR Trans. Symmetric Cryptol.*, 2020(S1):390–416, 2020.
 16. C. Dobraunig and B. Mennink. Leakage resilience of the duplex construction. In *ASIACRYPT (3)*, volume 11923 of *Lecture Notes in Computer Science*, pages 225–255. Springer, 2019.
 17. C. Dobraunig and B. Mennink. Leakage resilient value comparison with application to message authentication. In *EUROCRYPT (2)*, volume 12697 of *Lecture Notes in Computer Science*, pages 377–407. Springer, 2021.
 18. Y. Dodis and J. P. Steinberger. Message authentication codes from unpredictable block ciphers. In *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 267–285. Springer, 2009.
 19. Y. Dodis and J. P. Steinberger. Domain extension for macs beyond the birthday barrier. In *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 323–342. Springer, 2011.
 20. S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *FOCS*, pages 293–302. IEEE Computer Society, 2008.
 21. B. Fuller and A. Hamlin. Unifying leakage classes: Simulatable leakage and pseudoentropy. In *ICITS*, volume 9063 of *Lecture Notes in Computer Science*, pages 69–86. Springer, 2015.
 22. D. Goudarzi and M. Rivain. How fast can higher-order masking be in software? In *EUROCRYPT (1)*, volume 10210 of *Lecture Notes in Computer Science*, pages 567–597, 2017.
 23. C. Guo, O. Pereira, T. Peters, and F. Standaert. Towards low-energy leakage-resistant authenticated encryption from the duplex sponge construction. *IACR Trans. Symmetric Cryptol.*, 2020(1):6–42, 2020.
 24. C. Guo, F. Standaert, W. Wang, and Y. Yu. Efficient side-channel secure message authentication with better bounds. *IACR Trans. Symmetric Cryptol.*, 2019(4):23–53, 2019.
 25. C. Hazay, A. López-Alt, H. Wee, and D. Wichs. Leakage-resilient cryptography from minimal assumptions. In *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 160–176. Springer, 2013.
 26. S. Hirose. Some plausible constructions of double-block-length hash functions. In M. J. B. Robshaw, editor, *Fast Software Encryption, 13th International Workshop, FSE 2006, Graz, Austria, March 15-17, 2006, Revised Selected Papers*, volume 4047 of *Lecture Notes in Computer Science*, pages 210–225. Springer, 2006.
 27. J. Jean, I. Nikolic, and T. Peyrin. Tweaks and keys for block ciphers: The TWEAKE framework. In *ASIACRYPT (2)*, volume 8874 of *Lecture Notes in Computer Science*, pages 274–288. Springer, 2014.
 28. J. Jean, I. Nikolic, T. Peyrin, and Y. Seurin. Deoxys v1. 41. *CAESAR Competition, Final Portfolio*, 2016.
 29. J. Katz and Y. Lindell. *Introduction to Modern Cryptography, Second Edition*. CRC Press, 2014.

30. J. Krämer and P. Struck. Leakage-resilient authenticated encryption from leakage-resilient pseudorandom functions. In *COSADE*, volume 12244 of *Lecture Notes in Computer Science*, pages 315–337. Springer, 2020.
31. D. P. Martin, E. Oswald, M. Stam, and M. Wójcik. A leakage resilient MAC. In *IMACC*, volume 9496 of *Lecture Notes in Computer Science*, pages 295–310. Springer, 2015.
32. S. Micali and L. Reyzin. Physically observable cryptography (extended abstract). In *TCC*, volume 2951 of *Lecture Notes in Computer Science*, pages 278–296. Springer, 2004.
33. O. Pereira, F. Standaert, and S. Vivek. Leakage-resilient authentication and encryption from symmetric cryptographic primitives. In *CCS*, pages 96–108. ACM, 2015.
34. J. Schipper. Leakage-resilient authentication. Master’s thesis, 2011.
35. K. Suzuki, D. Tonien, K. Kurosawa, and K. Toyota. Birthday paradox for multi-collisions. In M. S. Rhee and B. Lee, editors, *Information Security and Cryptology - ICISC 2006, 9th International Conference, Busan, Korea, November 30 - December 1, 2006, Proceedings*, volume 4296 of *Lecture Notes in Computer Science*, pages 29–40. Springer, 2006.
36. L. Zhang, W. Wu, P. Wang, L. Zhang, S. Wu, and B. Liang. Constructing rate-1 macs from related-key unpredictable block ciphers: PGV model revisited. In *FSE*, volume 6147 of *Lecture Notes in Computer Science*, pages 250–269. Springer, 2010.