

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<https://hdl.handle.net/2066/221723>

Please be advised that this information was generated on 2021-01-26 and may be subject to change.

Security of the Suffix Keyed Sponge

Christoph Dobraunig and Bart Mennink

Digital Security Group, Radboud University, Nijmegen, The Netherlands

cgebraunig@cs.ru.nl, b.mennink@cs.ru.nl

Abstract. We formalize and analyze the general suffix keyed sponge construction, a pseudorandom function built on top of a cryptographic permutation. The construction hashes its data using the (keyless) sponge construction, transforms part of the state using the secret key, and generates the tag from the output of a final permutation call. In its simplest form, if the key and tag size are at most the rate of the sponge, one can see the suffix keyed sponge as a simple sponge function evaluation whose input is the plaintext appended with the key. The suffix keyed sponge is, however, much more general: the key and tag size may exceed the rate without any need to make extra permutation calls. We prove that the suffix keyed sponge construction achieves birthday-bound PRF security in the capacity, even if key and tag size exceed the rate. Furthermore, we prove that if the absorption of the key into the state happens in a leakage resilient manner, the suffix keyed sponge itself is leakage resilient as well. Our findings show that the suffix keyed sponge compares favorably with the hash-then-MAC construction. For instance, to reach a security level of k bits, the side-channel protected component in the suffix keyed sponge just needs to process k bits of input besides the key, whereas schemes following the hash-then-MAC construction need a side-channel protected MAC function that processes $2k$ bits of input besides the key. Moreover, even if we just consider black-box attacks, the MAC function in a hash-then-MAC scheme needs to be cryptographically strong whereas in the suffix keyed sponge the key may be absorbed by a simple XOR. The security proofs are performed using the H-coefficient technique, and make effective use of the multicollision limit function results of Daemen et al. (ASIACRYPT 2017), both for arguing that state manipulation larger than the rate is tolerated after key processing and for upper bounding the amount of leakage an attacker may gain about the secret key.

Keywords: suffix MAC · sponge · SuKS · PRF · leakage resilience · proof

1 Introduction

Whenever a device operates in a hostile environment and side-channel attacks [KJJ99] are a threat, protection against them becomes a necessity. Hence, a lot of effort has been put in the design of countermeasures against side-channel attacks like masking [GP99, CJRR99], threshold implementations [NRR06, NRS11], or special primitives that limit the available data complexity for an attacker [MSJ12, TS14, MSNF16]. However, the application of such countermeasures is not for free and requires additional resources like chip area, energy, power, computation time, randomness, and so on. In the context of message authentication codes (MACs), this has led to an increased popularity of functions that process the bulk of the input in a keyless manner and use the secret key for finalization only. This way, the first and largest part is not prone to side-channel attacks; only the second part must be protected, but as it is reasonably small this is an easier task. This particularly applies to the case where messages may be arbitrarily large (as in our use case).




Figure 1: The suffix keyed sponge. The plaintext P is first injectively padded into r -bit blocks $P_1 \dots P_\ell$.

This approach comes to life in the hash-then-MAC construction, used in [BKP⁺16, BPPS17, BGP⁺19, GPPS19], among others. In the hash-then-MAC construction, one first hashes the plaintext P using a hash function H , and subsequently authenticates the outcome using a MAC function F with key K to obtain a tag T :

$$T = F(K, H(P)).$$

The advantage of hash-then-MAC over traditional MAC functions like PMAC [BR02] and CMAC [Dwo05] in the context of side-channel protection is evident: as such constructions apply a cryptographic primitive with secret key at least once per plaintext block, each of these invocations must have countermeasures against side-channel attacks in place. In contrast, for hash-then-MAC, the amount of exposure of the secret key is independent of the number of processed input blocks.

However, in order for the hash-then-MAC construction to be k -bit secure, where k is the security level, one must necessarily take a hash function with digest size $2k$ [BKP⁺16, BPPS17, BGP⁺19, GPPS19]. This also means that the function $F(K, \cdot)$, which is the critical part in the construction from a side-channel perspective, must be able to process $2k$ -bit inputs. This negatively impacts the resources needed to implement countermeasures; in general, it is cheaper to protect a smaller primitive against side-channel attacks.

An alternative to the hash-then-MAC approach is the suffix keyed sponge approach used in ISAP [DEM⁺17]. A generalized depiction of the suffix keyed sponge is given in Figure 1, and it is described in detail in Section 3. The function operates as a sponge on top of a permutation p with a state of $b = c + r$ bits, split into an outer part of size r (the rate) and an inner part of size c (the capacity). The k -bit key K is absorbed using an s -bit to s -bit keyed function G and one output block T of size t is squeezed. For a specific case where the state size s and the tag size t are at most the rate r and in addition the key is absorbed by a simple XOR, the construction matches the original description of Bertoni et al. [BDPV11a, Section 5.11.2]. However, the suffix keyed sponge construction that we consider is more general.

For the suffix keyed sponge, there is no reason to believe that the function G needs to be $2k$ bits large, like it is required for F in hash-then-MAC. In the case of G , it seems to suffice to use a k -bit keyed function only. Indeed, in the suffix keyed sponge we can resort to the secrecy of the state, and collisions in the input to G are not necessarily harmful. Because for the suffix keyed sponge the function G is the focal point of protection against side-channel attacks, the construction compares favorably over the hash-then-MAC mode.

Unfortunately, no analysis of the suffix keyed sponge has appeared so far, neither in the black-box setting nor in the leakage resilient setting. The best we can do is to fall back to the indistinguishability of the keyless sponge [BDPV08]. This reduction is valid as long as we consider security in a black-box setting, where the key is absorbed using an XOR, key and tag are of size at most the rate, and the capacity is at least twice the security level. In other words, this is how security of the original and more restrictive version of the suffix keyed sponge of Bertoni et al. [BDPV11a, Section 5.11.2] was argued.

1.1 Our Contribution

We present a complete and general analysis of the suffix keyed sponge (SuKS), both in the black-box model and in the leakage resilience model. Noting that G is pivotal in the analysis of leakage resilience (the key comes into play only at the occurrence of G), we demonstrate how the security bound depends on the choice of G .

1.1.1 Suffix Keyed Sponge with Restricted Parameters

As a starter, we consider the restricted version of the suffix keyed sponge, where the state size s of G and the tag size t are at most the rate r of the sponge. This restricted version corresponds to the original suffix keyed sponge description of Bertoni et al. [BDPV11a, Section 5.11.2], though with arbitrary key absorbing function G . In this case, as mentioned above, we can resort to the indistinguishability of the keyless sponge [BDPV08] and argue that the resulting construction behaves as a pseudorandom function. This indistinguishability result states that if there is no inner collision on the c -bit inner part, the sponge function (with random permutation) is indistinguishable from a random oracle, and we can analyze security of an idealized version of the suffix keyed sponge. That construction, in turn, can only be broken by an adversary that ever guesses the output of the function G . The result and analysis are given in Section 4.

1.1.2 Suffix Keyed Sponge with Unrestricted Parameters

We present the main result in Section 5: security of the suffix keyed sponge in an unrestricted setting. In this case, both the state size s of the function G as well as the tag size t may be larger than the rate r . We prove that the construction still achieves $c/2$ -bit security under the assumption that the outputs of G are uniform (hard to guess) and universal (collide with small probability).

Note that one can consider the suffix keyed sponge state right before the evaluation of G as a state with s -bit rate and $(b - s)$ -bit capacity, but $b - s$ may be smaller than c , and hence, it may be that collisions occur on the $(b - s)$ -bit inner part even though the keyless part of the suffix keyed sponge does not have inner collisions. The trick in the proof consists of observing that, even if there are multicollisions on this $(b - s)$ -bit inner part, these multicollisions do not harm security of the construction as long as the input to G is different. A symmetric reasoning is applied on the state from which the tag is derived: it can be split into a t -bit outer part and a $(b - t)$ -bit inner part. Overall, the bound shows that even partial (multi-)collisions on the input to or output of the keyed permutation call in the suffix keyed sponge provide only limited help to the adversary. Hence, the security level usually stays in the area of $2^{c/2}$.

The result holds for an arbitrary G that is uniform and universal, but this does not mean that it needs to be cryptographically strong: a simple XOR suffices. This observation demonstrates another advantage of the suffix keyed sponge over the hash-then-MAC approach, where F must necessarily be cryptographically strong.

1.1.3 Leakage Resilience of Suffix Keyed Sponge

In practice, however, one might want select a stronger function G , for the plain reason that for an XOR even a small leakage usually allows to draw conclusions about the absorbed secret key. We therefore transform the main result to the setting of leakage resilient PRF design in Section 6. We consider non-adaptive leakage, akin to [Pie09, YSPY10, FPS12, SPY⁺10, DP10], and restrict our focus to leakage coming from p . In other words, we assume that G is a protected function (it could be instantiated as a leakage resilient keyed duplex, cf., Dobraunig and Mennink [DM19a]), and consider that the adversary may obtain leakage from the last invocation of p in the suffix keyed sponge. It is easily observed

from Figure 1 that only this invocation of p may leak information; all previous ones process no secret data in the first place. Keeping this in mind, the leakage resilience proof is quite clean. First, we note that for repeated inputs X to G for different $(b - s)$ -bit inner parts, the permutation p may leak different pieces of information about $G(K, X)$. By upper bounding the maximum size of a multicollision in the input to G , we can henceforth upper bound the amount of leakage that an adversary may learn for any particular $G(K, X)$, and then upper bound its success probability of guessing this value.

We stress, again, that G need not be cryptographically strong, unlike F in hash-then-MAC. This provides additional flexibility when considering leakage resilience. For instance, it could be possible to use a masked block cipher as G , but with heavily reduced rounds, since an attacker does not see the full output of the block cipher but only a limited leakage of it.

In Section 7, we map our results to ISAP [DEM⁺17] and demonstrate that the multicollision events only contribute to the bound negligibly, both in the black-box setting and in the leakage resilience setting. In the conclusion in Section 8, we elaborate on what our leakage resilient PRF result means for leakage resilient MAC design.

1.2 Related Work

A significant amount of research has been put in the analysis of the keyed sponge, where the state of the sponge is keyed prior to the absorption of data [BDPV11c, CDH⁺12, ADMV15, NY16, JLM14, GPT15, MRV15, BDPV11b, MRV15, DMV17, Men18]. These results have not only led to better and more fine grained bounds, but in the end also demonstrated that the full bit-width of the state can be used to absorb data in the case of such “prefix” keyed duplex or sponge [DMV17]. The reason for this is that, as the key is used to initialize the state, all subsequent states are secret to an adversary.

By allowing full-state absorption, these constructions are naturally more efficient than the suffix keyed sponge. On the downside, however, secrecy of the states also makes these constructions more expensive from a side-channel protection point of view. Constructions following the prefix keyed duplex or prefix keyed sponge usually allow for the evaluation of the same secret state for varying inputs. This is for instance the case in authenticated encryption schemes, which are initialized using a static secret key and a nonce. Hence, DPA attacks can be mounted if the protection against side-channel attacks is insufficient, as shown by Samwel and Daemen [SD17]. Even stronger: not only the initialization has to be protected, but every stage where known changing inputs are mixed with a static secret. In the case of an attacker that can control all the inputs, e.g. for a MAC based on a prefix keyed sponge during verification, this means that the whole construction has to be protected against side-channel attacks. In this light, it makes sense to limit the exposure of the secret key and rely on a suffix keyed sponge as done in ISAP [DEM⁺17].

It is worth noting that there exist other ways besides hash-then-MAC and the suffix keyed sponge to end up with leakage resilient MAC constructions. One way to do so is to base the scheme on ideas from asymmetric cryptography, an approach endeavored in a proposal of Martin et al. [MOSW15] that is based on the Barreto-Naehrig [BN05] family of pairing-friendly elliptic curves.

Alternatively, one can design a leakage resilient PRF from a (smaller) weak PRF in a sequential fashion, where the key input to the underlying weak PRF depends on the processed data so far. This approach is adopted in [YSPY10, FPS12, SPY⁺10, DP10], among others. Such a PRF can, naturally, also be used as instantiation of G in the suffix keyed sponge.

2 Preliminaries

Let $m, n \in \mathbb{N}$. The set of n -bit strings is denoted $\{0, 1\}^n$, the set of arbitrarily long strings is denoted $\{0, 1\}^*$, the set of n -bit functions is denoted $\text{func}(n)$, and the set of n -bit permutations is denoted $\text{perm}(n)$. We let \mathfrak{s}_n be any injective padding function that transforms an arbitrarily long string $X \in \{0, 1\}^*$ into n -bit blocks X_1, \dots, X_ℓ . For $X \in \{0, 1\}^n$ and if $m \leq n$, we denote by $\text{left}_m(X)$ (resp., $\text{right}_m(X)$) the m leftmost (resp., rightmost) bits of X . If $m \geq n$, we write $(m)_n = m(m-1) \cdots (m-n+1)$ as the falling factorial. For a finite set \mathcal{X} , $X \stackrel{\$}{\leftarrow} \mathcal{X}$ denotes the uniformly random drawing of an element X from \mathcal{X} .

2.1 Distinguishing Advantage

An adversary \mathcal{A} is an algorithm. It is given access to one or more oracles \mathcal{O} , and after interaction with \mathcal{O} it outputs a decision bit b : $b \leftarrow \mathcal{A}^{\mathcal{O}}$. For two oracles \mathcal{O} and \mathcal{P} , the adversarial distinguishing advantage is defined as

$$\Delta_{\mathcal{A}}(\mathcal{O}; \mathcal{P}) = \Pr(1 \leftarrow \mathcal{A}^{\mathcal{O}}) - \Pr(1 \leftarrow \mathcal{A}^{\mathcal{P}}). \quad (1)$$

2.2 PRF Security

Let $b, k, t \in \mathbb{N}$ and $m \in \mathbb{N} \cup \{*\}$. Consider a function $F : \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^t$ that internally uses a permutation $p \in \text{perm}(b)$. We write F instantiated with permutation p and key $K \in \{0, 1\}^k$ as F_K^p . The pseudorandom function (PRF) security of F against an adversary \mathcal{A} is defined as

$$\mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) = \Delta_{\mathcal{A}}(F_K^p, p; R_{m,t}, p), \quad (2)$$

where the randomness is taken over the random drawing of $K \stackrel{\$}{\leftarrow} \{0, 1\}^k$, the random drawing of $p \stackrel{\$}{\leftarrow} \text{perm}(b)$, the definition of a function $R_{m,t} : \{0, 1\}^m \rightarrow \{0, 1\}^t$ that generates a random t -bit response for each new input, and the random coins of \mathcal{A} .

The adversary is typically bounded by three types of complexities: q denotes the construction complexity, the number of queries the adversary may make to F_K^p or $R_{m,t}$; N denotes the total primitive complexity, which accounts for the total number of evaluations of p (either direct forwardly, direct inversely, or via the real-world construction oracle F_K^p); τ denotes the time complexity.

It is a well-known result that a PRF is a secure message authentication code in the black-box model [GGM84, BKR94, BGM04]. More detailed, if $\mathbf{Adv}_F^{\text{mac}}(\mathcal{A})$ denotes the advantage of an adversary \mathcal{A} with query access to (F_K^p, p) to output a non-trivial forgery for F_K^p , then

$$\mathbf{Adv}_F^{\text{mac}}(\mathcal{A}) \leq \mathbf{Adv}_F^{\text{prf}}(\mathcal{A}) + \frac{q}{2^t},$$

where q is the number of forgery attempts that \mathcal{A} makes. We note that this reduction does not necessarily apply in the leakage resilience setting, cf., Section 8.

2.3 Uniform and Universal Functions

Let $k, s \in \mathbb{N}$ and $\delta, \varepsilon \in [0, \infty)$. A function $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$ is $2^{-\delta}$ -uniform if for any $X, Y \in \{0, 1\}^s$,

$$\Pr(G(K, X) = Y) \leq 2^{-\delta},$$

where the randomness is taken over the random drawing of $K \xleftarrow{s} \{0, 1\}^k$. It is $2^{-\varepsilon}$ -universal if for any distinct $X, X' \in \{0, 1\}^s$,

$$\Pr(G(K, X) = G(K, X')) \leq 2^{-\varepsilon},$$

where the randomness is taken over the random drawing of $K \xleftarrow{s} \{0, 1\}^k$. Here, we say that a function is 0-uniform (resp., 0-universal) if above definition applies for $\delta = \infty$ (resp., $\varepsilon = \infty$).

Note that for $k \leq s$, the function $G(K, X) = K \parallel 0^{s-k} \oplus X$ that simply XORs the two input values is 2^{-k} -uniform and 0-universal. Alternatively, we can define G as a random function: $G : |\text{func}(s)| \times \{0, 1\}^s \rightarrow \{0, 1\}^s$ is defined as $G(f, X) = f(X)$. This function G is 2^{-s} -uniform and 2^{-s} -universal.

2.4 Multicollision Limit Function

Daemen et al. [DMV17] introduced the multicollision limit function in the context of keyed sponge proofs. We will use the same notion in our proofs.

Let $q, b, s \in \mathbb{N}$ such that $s \leq b$. Consider the experiment of throwing q balls uniformly at random in 2^{b-s} bins, and denote by μ the maximum number of balls in any single bin. The multicollision limit function $\mu_{b-s,s}^q$ is defined as the smallest natural number x that satisfies

$$\Pr(\mu > x) \leq \frac{x}{2^s}.$$

Daemen et al. also proved that if one does not consider the bins to be uniformly randomly selected, but rather according to a distribution D where the i -th ball ends up in a certain bin with probability

$$\frac{2^s - (i-1)}{2^b - (i-1)} \leq p \leq \frac{2^s}{2^b - (i-1)}, \quad (3)$$

the corresponding multicollision function, defined as $\mu_{b-s,s}^{D,q}$, satisfies $\mu_{b-s,s}^{D,q} \leq \mu_{b-s,s}^{2q}$ [DMV17, Lemma 6].

Furthermore, Daemen et al. [DMV17] gave an in-depth analysis of the term $\mu_{b-s,s}^q$. The analysis is tedious, but the conclusion is that the term behaves as follows:

$$\mu_{b-s,s}^q \lesssim \begin{cases} b / \log_2 \left(\frac{2^{b-s}}{q} \right), & \text{for } q \lesssim 2^{b-s}, \\ b \cdot \frac{q}{2^{b-s}}, & \text{for } q \gtrsim 2^{b-s}. \end{cases}$$

3 Suffix Keyed Sponge

Let $b, c, r, k, s, t \in \mathbb{N}$ such that $c + r = b$ and $k, s, t \leq b$. Let $p \in \text{perm}(b)$ be a permutation and $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$ be a function. The suffix keyed sponge (SuKS) $F : \{0, 1\}^k \times \{0, 1\}^* \rightarrow \{0, 1\}^t$ is defined in Algorithm 1 and depicted in Figure 1. The function first injectively pads its input P , and hashes it the usual sponge way. Only at this point, the key will be used: the outer k bits of the resulting state are transformed through $G(K, \cdot)$. The resulting state is processed by the permutation p once more, and the tag equals the outer t bits of the state.

Algorithm 1 Suffix keyed sponge construction F **Input:** $(K, P) \in \{0, 1\}^k \times \{0, 1\}^*$ **Output:** $T \in \{0, 1\}^t$

- 1: $P_1 \dots P_\ell \leftarrow \mathfrak{s}_{\leftarrow r}(P)$
- 2: $S \leftarrow 0^b$
- 3: **for** $i = 1, \dots, \ell$ **do**
- 4: $S \leftarrow S \oplus P_i \parallel 0^c$
- 5: $S \leftarrow p(S)$
- 6: $S \leftarrow G(K, \text{left}_s(S)) \parallel \text{right}_{b-s}(S)$
- 7: $S \leftarrow p(S)$
- 8: **return** $\text{left}_t(S)$

4 Security of Suffix Keyed Sponge with Restricted Parameters

As a starter, we prove security of the suffix keyed sponge for a restricted case where $s, t \leq r$. This version of the suffix keyed sponge corresponds to the original version suggested by Bertoni et al. [BDPV11a, Section 5.11.2], though with arbitrary key absorbing function G . The result is mainly included to show how, in this case, we can rely on the indistinguishability of the sponge [BDPV08].

Theorem 1. *Let $b, c, r, k, s, t \in \mathbb{N}$ such that $c + r = b$, $k \leq b$, and $s, t \leq r$. Consider the suffix keyed sponge of Section 3 based on random permutation $p \xleftarrow{\$} \text{perm}(b)$ and a function $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$. Assume that G is $2^{-\delta}$ -uniform. For any adversary \mathcal{A} with construction complexity q and primitive complexity $N \leq 2^{b-1}$,*

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) \leq \frac{N^2 + N}{2^c} + \frac{N}{2^\delta}.$$

Proof. Let $K \xleftarrow{\$} \{0, 1\}^k$ and $p \xleftarrow{\$} \text{perm}(b)$. Let $R_{*,t} : \{0, 1\}^* \rightarrow \{0, 1\}^t$ be a function that generates a random t -bit response for each new input. Consider any adversary \mathcal{A} , whose goal is to maximize

$$\Delta_{\mathcal{A}}(F_K^p, p; R_{*,t}, p), \quad (4)$$

As we will argue security based on the uniformity and universality of G (and not on the computational security of it), we assume that \mathcal{A} is computationally unbounded.

Let sponge^p be the sponge construction based on permutation p with rate r , that gets input strings of size a multiple of r (we can assume w.l.o.g. that the message is already padded) and a natural number $z \in \mathbb{N}$, and it outputs a string of z bits. One can equivalently write F_K^p as a function $F_K^{\text{sponge}^p}$ as follows:

Input: $(K, P) \in \{0, 1\}^k \times \{0, 1\}^*$ **Output:** $T \in \{0, 1\}^t$

- 1: $P_1 \dots P_\ell \leftarrow \mathfrak{s}_{\leftarrow r}(P)$
- 2: $Z \leftarrow \text{sponge}^p(P_1 \dots P_\ell, s)$
- 3: $P' \leftarrow (G_K(Z) \oplus Z) \parallel 0^{r-s}$
- 4: $Z \leftarrow \text{sponge}^p(P_1 \dots P_\ell \parallel P', t)$
- 5: **return** Z

We will resort to the indistinguishability of the sponge [BDPV08], which states that there exists a simulator S such that $\Delta_{\mathcal{A}'}(\text{sponge}^p, p; RO, S^{RO}) \leq \binom{N+1}{2}/2^c$ for any adversary \mathcal{A}' with total complexity N , where RO is a random oracle. We write F_K^{RO} as the suffix keyed sponge construction where the two evaluations of sponge^p are replaced by RO , and

write $F_K^{RO,RO'}$ as that construction where the first evaluation of sponge^p is replaced by RO and the second one by RO' . By a game hopping argument,

$$\begin{aligned}
& \Delta_{\mathcal{A}} \left(F_K^{\text{sponge}^p}, p; R_{*,t}, p \right) \\
& \leq \Delta_{\mathcal{A}} \left(F_K^{\text{sponge}^p}, p; F_K^{RO}, S^{RO} \right) \\
& \quad + \Delta_{\mathcal{A}} \left(F_K^{RO}, S^{RO}; F_K^{RO,RO'}, S^{RO} \right) + \Delta_{\mathcal{A}} \left(F_K^{RO,RO'}, S^{RO}; R_{*,t}, p \right) \\
& \leq \Delta_{\mathcal{A}'} \left(\text{sponge}^p, p; RO, S^{RO} \right) \\
& \quad + \Delta_{\mathcal{A}} \left(F_K^{RO}, S^{RO}; F_K^{RO,RO'}, S^{RO} \right) + \Delta_{\mathcal{A}''} \left(S^{RO}; p \right), \tag{5}
\end{aligned}$$

where \mathcal{A}' and \mathcal{A}'' are some adversaries with the same primitive complexity N as \mathcal{A} (noting that any evaluation of $F_K^{\text{sponge}^p}$ indeed consists of two sponge calls, but they jointly consist of $\ell + 1$ unique primitive evaluations). The first and last distance in (5) are at most the indistinguishability of the sponge, $\binom{N+1}{2}/2^c$. For the middle term of (5), we define the following event **bad**:

- **bad**: S^{RO} calls its oracle for input $P_1 \dots P_\ell \parallel (G_K(Z) \oplus Z) \parallel 0^{r-s}$, where $Z = RO(P_1 \dots P_\ell, s)$.

As long as **bad** does not happen, the adversary cannot notice the difference between RO and RO' for the finalization calls in both world. Therefore, (F_K^{RO}, S^{RO}) and $(F_K^{RO,RO'}, S^{RO})$ are indistinguishable, and thus

$$\Delta_{\mathcal{A}} \left(F_K^{RO}, S^{RO}; F_K^{RO,RO'}, S^{RO} \right) \leq \Pr(\text{bad}). \tag{6}$$

The event **bad** requires an adversary to force S into querying the random oracle for value $Y := G_K(Z) \oplus Z$, where Z is fixed by the specific query that \mathcal{A} makes (by design of the actual simulator of [BDPV08]). The adversary makes N attempts, and any of them succeeds with probability at most $2^{-\delta}$. We thus obtain that $\Pr(\text{bad}) \leq N/2^\delta$. \square

5 Security of Suffix Keyed Sponge with Unrestricted Parameters

In this section, we set aside the restriction $s, t \leq r$ of the previous section and present the main result, namely security of the suffix keyed sponge for arbitrary r, s, t .

Theorem 2. *Let $b, c, r, k, s, t \in \mathbb{N}$ such that $c + r = b$ and $k, s, t \leq b$. Consider the suffix keyed sponge of Section 3 based on random permutation $p \xleftarrow{\$} \text{perm}(b)$ and a function $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$. Assume that G is $2^{-\delta}$ -uniform and $2^{-\varepsilon}$ -universal. For any adversary \mathcal{A} with construction complexity $q \geq 2$ and primitive complexity $N \leq 2^{b-1}$,*

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) \leq \frac{2N^2}{2^c} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^{\min\{\delta,\varepsilon\}}} + \frac{\mu_{t,b-t}^q \cdot N}{2^{b-t}}.$$

The proof of Theorem 2 is given in Section 5.1.

Note that G itself may be anything, as long as it is $2^{-\delta}$ -uniform and $2^{-\varepsilon}$ -universal. The simple XOR function, i.e., setting $s = k$ and taking $G(K, X) = K \oplus X$, is 2^{-k} -uniform and 0-universal. A more general instantiation is to have G to be a pseudorandom function independent of the permutation p . In this case, the first step of the proof would be to replace G by a random function at cost $\text{Adv}_G^{\text{prf}}(\mathcal{A}')$, for some adversary \mathcal{A}' . The second step would be to rely on the observation that a random s -bit function is 2^{-s} -uniform and 2^{-s} -universal (see Section 2.3). This leads to the following corollary, for which the formal proof is given in Section 5.2.

Corollary 1. *Let $b, c, r, k, s, t \in \mathbb{N}$ such that $c + r = b$ and $k, s, t \leq b$. Consider the suffix keyed sponge of Section 3 based on random permutation $p \xleftarrow{\$} \text{perm}(b)$ and a function $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$ independent of p . For any adversary \mathcal{A} with construction complexity $q \geq 2$, primitive complexity $N \leq 2^{b-1}$, and time complexity τ ,*

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) \leq \frac{2N^2}{2^c} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^s} + \frac{\mu_{t,b-t}^q \cdot N}{2^{b-t}} + \text{Adv}_G^{\text{prf}}(\mathcal{A}'),$$

for some adversary \mathcal{A}' with construction complexity q , primitive complexity 0, and time complexity $\tau' \approx \tau$.

5.1 Proof of Theorem 2

Let $K \xleftarrow{\$} \{0, 1\}^k$ and $p \xleftarrow{\$} \text{perm}(b)$. Let $R_{*,t} : \{0, 1\}^* \rightarrow \{0, 1\}^t$ be a function that generates a random t -bit response for each new input. Consider any adversary \mathcal{A} , whose goal is to maximize

$$\Delta_{\mathcal{A}}(F_K^p, p; R_{*,t}, p), \quad (7)$$

We will call $\mathcal{O} := (F_K^p, p)$ the real world and $\mathcal{P} := (R_{*,t}, p)$ the ideal world. As we will argue security based on the uniformity and universality of G (and not on the computational security of it), we assume that \mathcal{A} is computationally unbounded. Without loss of generality, it is deterministic: for any probabilistic adversary there is a deterministic one with the same success probability.

5.1.1 Oracle Interaction

The adversary has a total primitive complexity N , meaning that for the *real world* the total number primitive queries (direct or through F_K^p) does not exceed N . Of course, the ideal world $R_{*,t}$ does not query p , but the complexity is counted in terms of the number of evaluations that would be made to p in the real world. Note that in the real world $\mathcal{O} = (F_K^p, p)$, any construction query P of ℓ blocks entails $\ell + 1$ evaluations of p : the first ℓ evaluations are “offline” and the adversary can evaluate them itself using the primitive oracle, and the last evaluation is keyed. We will thus consider \mathcal{A} to be allowed an “offline” complexity $N - q$ and an online complexity q . Without loss of generality, as duplicated queries are not counted doubly, we may assume that before each construction query, the adversary makes *all primitive queries but the last one* itself offline.

The interaction of \mathcal{A} with its oracle (\mathcal{O} or \mathcal{P}) is summarized in a view. All q construction queries (to F_K^p in the real world and to $R_{*,t}$ in the ideal world) are summarized in a view

$$v_c = \{(P_1, T_1), \dots, (P_q, T_q)\}.$$

All primitive evaluations (to p , in both worlds) are summarized in a view

$$v_p = \{(\text{dir}_1, X_1, Y_1), \dots, (\text{dir}_{N-q}, X_{N-q}, Y_{N-q})\},$$

where $\text{dir}_j \in \{+, -\}$ denotes the direction of the j -th query: forward primitive queries are denoted by $\text{dir}_j = +$ and inverse primitive queries by $\text{dir}_j = -$. The variables dir_j are used to distinguish between the direction of the queries in case of inner collisions.

After the conversation of \mathcal{A} with its oracle, but before it outputs its decision bit, we reveal additional information to \mathcal{A} . First of all, we reveal a key K . In the real world, this is the key $K \xleftarrow{\$} \{0, 1\}^k$ that is effectively used by G_K ; in the ideal world, it is a dummy key. In addition, we reveal a tuple $\{Z_1, \dots, Z_q\}$. In the real world, these are the values that are truncated at the end of each construction query to F_K^p (i.e., the values $\text{right}_{b-t}(S)$

with S the state of line 7 of Algorithm 1); in the ideal world, these are random values $Z_i \xleftarrow{\$} \{0, 1\}^{b-t}$.

The values Z_i are appended to v_c :

$$v'_c = \{(P_1, T_1, Z_1), \dots, (P_q, T_q, Z_q)\}.$$

The aggregate view is defined as $v = (v'_c, v_p, K)$. We assume that the adversary never makes any duplicate query, hence $P_i \neq P_{i'}$ for any two distinct queries in v'_c and $(X_j, Y_j) \neq (X_{j'}, Y_{j'})$ for any two distinct queries in v_p .

It is important to note that, as we force \mathcal{A} to make all primitive queries corresponding to the unkeyed part of F_K^p , for each construction query (P_i, T_i) , the state value U_i that is at the end of the unkeyed part of F_K^p (see Figure 1) can be retrieved from v_p . We subsequently define V_i and W_i from U_i , K , and v'_c as follows:

$$V_i = G(K, \text{left}_s(U_i)) \parallel \text{right}_{b-s}(U_i), \quad (8)$$

$$W_i = T_i \parallel Z_i. \quad (9)$$

Note that in the real world (F_K^p, p) , the values V_i and W_i are the actual values indicated in Figure 1. In the ideal world, they are simply defined as in (8) and (9).

5.1.2 H-Coefficient Technique

We denote by $D_{\mathcal{O}}$ the probability distribution of views in interaction with \mathcal{O} , and by $D_{\mathcal{P}}$ the probability distribution of views in interaction with \mathcal{P} . Define \mathcal{V} to be the set of “attainable views”: views v such that $\Pr(D_{\mathcal{P}} = v) > 0$. We will use the H-coefficient technique [Pat08, CS14].

Lemma 1 (H-coefficient technique). *Consider a partition $\mathcal{V} = \mathcal{V}_{\text{good}} \cup \mathcal{V}_{\text{bad}}$ of the set of views into “good” and “bad” views. Let $\epsilon \in [0, 1]$ be such that $\frac{\Pr(D_{\mathcal{O}}=v)}{\Pr(D_{\mathcal{P}}=v)} \geq 1 - \epsilon$ for all $v \in \mathcal{V}_{\text{good}}$. Then, $\Delta_{\mathcal{A}}(\mathcal{O}; \mathcal{P}) \leq \epsilon + \Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}})$.*

5.1.3 Bad Views

Informally, we consider a view *bad* if for any construction query P_i in the real world $\mathcal{O} = (F_K^p, p)$, the evaluation (V_i, W_i) of p is repeated. This is formally covered by the following two events:

- coll_{cc} : there exist distinct $i, i' \in \{1, \dots, q\}$ with

$$V_i = V_{i'} \text{ or } W_i = W_{i'};$$

- coll_{cp} : there exist $i \in \{1, \dots, q\}$ and $j \in \{1, \dots, N - q\}$ with

$$V_i = X_j \text{ or } W_i = Y_j.$$

We indeed need these two bad events, because in the non-occurrence of either of those, $v = (v'_c, v_p, K)$ corresponds to *exactly* N input/output tuples of p (and one key).

However, analyzing bad events coll_{cc} and coll_{cp} is quite involved, and we will define four auxiliary bad events. For this, let $\nu_{\text{right}} = \mu_{b-s, s}^{2(N-q)}$ and $\nu_{\text{tag}} = \mu_{t, b-t}^q$ be two thresholds.

- cap_{fwd} : there exists $j \in \{1, \dots, N - q\}$ with $\text{dir}_j = +$ and

$$\text{right}_c(Y_j) \in \{\text{right}_c(X_1), \dots, \text{right}_c(X_j), \text{right}_c(Y_1), \dots, \text{right}_c(Y_{j-1})\};$$

- cap_{inv} : there exists $j \in \{1, \dots, N - q\}$ with $\text{dir}_j = -$ and

$$\text{right}_c(X_j) \in \{\text{right}_c(Y_1), \dots, \text{right}_c(Y_{j-1}), 0^c\};$$

- mc_{right} : there exist distinct $j_1, \dots, j_{\nu_{\text{right}}+1} \in \{1, \dots, N - q\}$ with $\text{dir}_j = +$ and

$$\text{right}_{b-s}(Y_{j_1}) = \dots = \text{right}_{b-s}(Y_{j_{\nu_{\text{right}}+1}});$$

- mc_{tag} : there exist distinct $i_1, \dots, i_{\nu_{\text{tag}}+1} \in \{1, \dots, q\}$ with

$$\text{left}_t(W_{i_1}) = \dots = \text{left}_t(W_{i_{\nu_{\text{tag}}+1}}).$$

Events cap_{fwd} and cap_{inv} cover the case that there exists a inner collision, and that the adversary has potentially found two different plaintexts leading to the same inner part. At the cost of readability, the event cap_{fwd} can be tightened slightly, noting that it is not a problem if $\text{right}_c(Y_j)$ hits an older inner part if this j -th query is not “rooted”, i.e., if there is no path from 0^b to this query. See also Daemen et al. [DMV18].

The events mc_{right} and mc_{tag} cover multicollisions on part of the state, and are included to tighten the bounding of the occurrence of $\text{coll}_{\text{cc}} \vee \text{coll}_{\text{cp}}$. We note that the choice of thresholds ν_{right} and ν_{tag} (and our reliance on the multicollision limit function of Section 2.4) is for the sake of generality. It may be that for specific parameter choices b, c, s, t , better thresholds may render a better bound. See also Section 8.

We write $\text{cap} = \text{cap}_{\text{fwd}} \vee \text{cap}_{\text{inv}}$ and

$$\text{bad} = \text{cap} \vee \text{mc}_{\text{right}} \vee \text{mc}_{\text{tag}} \vee \text{coll}_{\text{cc}} \vee \text{coll}_{\text{cp}}. \quad (10)$$

5.1.4 Probability of Bad View

Our goal is to bound the probability that a bad transcript is generated in the ideal world, $\Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}})$. This probability equals the probability that a view v generated by $D_{\mathcal{P}}$ satisfies **bad**. By basic probability theory,

$$\begin{aligned} \Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}}) &= \Pr(\text{bad}) \leq \Pr(\text{cap}) + \Pr(\text{mc}_{\text{right}}) + \Pr(\text{mc}_{\text{tag}}) \\ &\quad + \Pr(\text{coll}_{\text{cc}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}})) \\ &\quad + \Pr(\text{coll}_{\text{cp}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}} \vee \text{mc}_{\text{tag}})). \end{aligned} \quad (11)$$

Here, we recall that in the ideal world, $K \xleftarrow{\$} \{0, 1\}^k$, $W_1, \dots, W_q \xleftarrow{\$} \{0, 1\}^b$, and each tuple $(\text{dir}_j, X_j, Y_j) \in v_p$ has either Y_j random without replacement (if $\text{dir}_j = +$) or X_j random without replacement (if $\text{dir}_j = -$).

cap. Consider the j -th query (dir_j, X_j, Y_j) . If it is a forward query it can only set cap_{fwd} and if it is an inverse query it can only set cap_{inv} . In either case, the response is uniformly randomly generated using a random permutation p from a set of size at least $2^b - (j - 1)$ elements. The query sets cap_{fwd} with probability at most $(2j - 1)2^r / (2^b - (j - 1))$ and cap_{inv} with probability at most $j2^r / (2^b - (j - 1))$. In either case, as any query is either forward or inverse, the success probability is at most $2(2j - 1)/2^c$, using that $j - 1 \leq 2^{b-1}$. Summing over all queries, we obtain

$$\Pr(\text{cap}) = \Pr(\text{cap}_{\text{fwd}} \vee \text{cap}_{\text{inv}}) \leq \sum_{j=1}^{N-q} \frac{2(2j - 1)}{2^c} = \frac{2(N - q)^2}{2^c}.$$

mc_{right}. The values Y_j are randomly generated from $\{0, 1\}^b \setminus \{Y_1, \dots, Y_{j-1}\}$, and at most $N - q$ draws are done. The event mc_{right} is thus a balls-and-bins experiment with $N - q$ balls randomly thrown into 2^{b-s} bins, in such a way that any of the bins contains more than ν_{right} balls. Note that the distribution of balls satisfies the condition of (3). By definition of $\nu_{\text{right}} = \mu_{b-s,s}^{2(N-q)}$, we can resort to the multicollision limit function of Section 2.4 and obtain

$$\Pr(\text{mc}_{\text{right}}) \leq \frac{\mu_{b-s,s}^{2(N-q)}}{2^s}.$$

mc_{tag}. The values W_i are randomly generated from $\{0, 1\}^b$, and q draws are done. The event mc_{tag} is thus a balls-and-bins experiment with q balls uniformly randomly thrown into 2^t bins, in such a way that any of the bins contains more than ν_{tag} balls. By definition of $\nu_{\text{tag}} = \mu_{t,b-t}^q$, we can resort to the multicollision limit function of Section 2.4 and obtain

$$\Pr(\text{mc}_{\text{tag}}) \leq \frac{\mu_{t,b-t}^q}{2^{b-t}}.$$

coll_{cc}. The event $V_i = V_{i'}$ is equivalent to stating that

$$G(K, \text{left}_s(U_i)) \parallel \text{right}_{b-s}(U_i) = G(K, \text{left}_s(U_{i'})) \parallel \text{right}_{b-s}(U_{i'}). \quad (12)$$

Consider any $i, i' \in \{1, \dots, q\}$. If $s \leq r$, then by $\neg\text{cap}$ we necessarily have $\text{right}_{b-s}(U_i) \neq \text{right}_{b-s}(U_{i'})$. Consider the case that $s > r$ and that $\text{right}_{b-s}(U_i) = \text{right}_{b-s}(U_{i'})$ holds. By the non-occurrence of cap , we necessarily have $\text{left}_s(U_i) \neq \text{left}_s(U_{i'})$. By $2^{-\varepsilon}$ -universality of G , the two queries satisfy above equation with probability at most $2^{-\varepsilon}$.

We thus have to count the number of possible choices $i, i' \in \{1, \dots, q\}$ such that $\text{right}_{b-s}(U_i) = \text{right}_{b-s}(U_{i'})$. Consider any i' (q possibilities). Recall that we assumed that the keyless part of F_K^p is evaluated by the adversary. By $\neg\text{cap}$, any construction query must necessarily be formed by making forward evaluations of p starting from 0^b , and none of these collide with an earlier primitive query. This, particularly, means that the last evaluation of p in the keyless part of each construction query is a forward query. By $\neg\text{mc}_{\text{right}}$, there are at most $\nu_{\text{right}} = \mu_{b-s,s}^{2(N-q)}$ possible paths from 0^b .

By eliminating symmetric cases, the number of possible (i, i') 's with $\text{right}_{b-s}(U_i) = \text{right}_{b-s}(U_{i'})$ is at most $\mu_{b-s,s}^{2(N-q)} \cdot q/2$, and we obtain that there exist i, i' such that $V_i = V_{i'}$ with probability at most

$$\frac{\mu_{b-s,s}^{2(N-q)} \cdot q/2}{2^\varepsilon}.$$

For the second event of coll_{cc} , as the values W_i are randomly generated from $\{0, 1\}^b$, there exist i, i' with $W_i = W_{i'}$ with probability at most $\binom{q}{2}/2^b \leq q^2/2^b$.

We thus obtain

$$\Pr(\text{coll}_{\text{cc}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}})) \leq \frac{\mu_{b-s,s}^{2(N-q)} \cdot q/2}{2^\varepsilon} + \frac{q^2}{2^b}.$$

coll_{cp}. Fix any primitive query (dir_j, X_j, Y_j) ($N - q$ possibilities). We consider both equations of the event separately. First consider equation $V_i = X_j$:

- If $\text{dir}_j = +$, the equation $V_i = X_j$ is equivalent to stating that

$$G(K, \text{left}_s(U_i)) \parallel \text{right}_{b-s}(U_i) = \text{left}_s(X_j) \parallel \text{right}_{b-s}(X_j).$$

As for the analysis coll_{cc} , by $\neg\text{cap}$ and $\neg\text{mc}_{\text{right}}$, there are at most $\nu_{\text{right}} = \mu_{b-s,s}^{2(N-q)}$ construction queries with $\text{right}_{b-s}(U_i) = \text{right}_{b-s}(X_j)$. By $2^{-\delta}$ -uniformity of G , any of these satisfies above equation with probability at most $2^{-\delta}$.

- If $\text{dir}_j = -$, the equation happens with probability at most $\mu_{b-s,s}^{2(N-q)}/2^\delta$ as it is symmetric to the case of $V_i = X_j$ in forward queries.

Next consider equation $W_i = Y_j$:

- If $\text{dir}_j = -$, the equation $W_i = Y_j$ is equivalent to stating that

$$T_i \parallel Z_i = \text{left}_s(Y_j) \parallel \text{right}_{b-s}(Y_j).$$

By $\neg\text{mc}_{\text{tag}}$, there are at most $\nu_{\text{tag}} = \mu_{t,b-t}^q$ construction queries with $T_i = \text{left}_s(Y_j)$. As $Z_i \stackrel{\$}{\leftarrow} \{0,1\}^{b-t}$, any of these satisfies above equation with probability at most $1/2^{b-t}$.

- If $\text{dir}_j = +$, the equation $W_i = Y_j$ happens with probability at most $q/(2^b - (j-1)) \leq 2q/2^b$, using that $j-1 \leq 2^{b-1}$.

Aggregating both cases, we obtain

$$\Pr(\text{coll}_{\text{cp}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}} \vee \text{mc}_{\text{tag}})) \leq \frac{\mu_{b-s,s}^{2(N-q)} \cdot (N-q)}{2^\delta} + \frac{\mu_{t,b-t}^q \cdot (N-q)}{2^{b-t}} + \frac{2q(N-q)}{2^b}.$$

Conclusion. Summing the individual terms, we obtain for (11) that

$$\begin{aligned} \Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}}) &\leq \frac{2(N-q)^2}{2^c} + \frac{q^2 + 2q(N-q)}{2^b} \\ &\quad + \frac{\mu_{b-s,s}^{2(N-q)} \cdot (N-q/2+1)}{2^{\min\{\delta,\varepsilon\}}} + \frac{\mu_{t,b-t}^q \cdot (N-q+1)}{2^{b-t}} \\ &\leq \frac{2N^2}{2^c} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^{\min\{\delta,\varepsilon\}}} + \frac{\mu_{t,b-t}^q \cdot N}{2^{b-t}}, \end{aligned} \tag{13}$$

where we used that $\min\{\delta,\varepsilon\} \leq s$ and $q \geq 2$.

5.1.5 Ratio for Good Views

Consider any good view $v \in \mathcal{V}_{\text{good}}$. We will prove that $\Pr(D_{\mathcal{O}} = v) \geq \Pr(D_{\mathcal{P}} = v)$.

Real World. We start with the real world $\mathcal{O} = (F_K^p, p)$. The view v_p corresponds to exactly $N - q$ distinct input/output tuples of p . For any construction query P_i in the extended view v'_c , the unkeyed part of the evaluation of $F_K^p(M)$ is included in v_p . The keyed evaluation of p is the tuple (V_i, W_i) defined in (8)-(9). By the non-occurrence of coll_{cp} , it is different from any tuple in v_p , and by the non-occurrence of coll_{cc} it is different from any earlier tuple $(V_{i'}, W_{i'})$ with $i' < i$. Therefore, the good view $v = (v'_c, v_p, K)$ corresponds to *exactly* N input/output tuples of p and one random key K . Therefore, we obtain:

$$\Pr(D_{\mathcal{O}} = v) = \frac{1}{(2^b)_N} \cdot \frac{1}{2^k}. \tag{14}$$

Ideal World. We next consider the ideal world $\mathcal{P} = (R_{*,t}, p)$. The view v_p corresponds to exactly $N - q$ distinct input/output tuples of p . The extended view v'_c consists of q outputs T_1, \dots, T_q of $R_{*,t}$ and q dummy values $Z_1, \dots, Z_q \xleftarrow{\$} \{0, 1\}^{b-t}$. Finally, the key K is random as before. Therefore, we obtain:

$$\Pr(D_{\mathcal{P}} = v) = \frac{1}{(2^b)_{N-q}} \cdot \frac{1}{(2^t)^q} \cdot \frac{1}{(2^{b-t})^q} \cdot \frac{1}{2^k} = \frac{1}{(2^b)_{N-q}} \cdot \frac{1}{2^{bq}} \cdot \frac{1}{2^k}. \quad (15)$$

Bounding the Ratio. Combining (14) and (15), we obtain the following for any good view $v \in \mathcal{V}_{\text{good}}$:

$$\frac{\Pr(D_{\mathcal{O}} = v)}{\Pr(D_{\mathcal{P}} = v)} = \frac{2^{bq}}{(2^b - (N - q))_q} \geq 1. \quad (16)$$

5.1.6 Conclusion

The H-coefficient technique of Lemma 1 gives, using (13) and (16):

$$\Delta_{\mathcal{A}}(F_K^p, p; R_{*,t}, p) \leq \frac{2N^2}{2^c} + \frac{\mu_{b-s}^{2(N-q)} \cdot N}{2^{\min\{\delta, \varepsilon\}}} + \frac{\mu_{t,b-t}^q \cdot N}{2^{b-t}}.$$

5.2 Proof of Corollary 1

Let $K \xleftarrow{\$} \{0, 1\}^k$ and $p \xleftarrow{\$} \text{perm}(b)$. Let $R_{*,t} : \{0, 1\}^* \rightarrow \{0, 1\}^t$ be a function that generates a random t -bit response for each new input. Consider any adversary \mathcal{A} , whose goal is to maximize

$$\Delta_{\mathcal{A}}(F_K^p, p; R_{*,t}, p), \quad (17)$$

As a first step, we replace G_K by a function $R_s : \{0, 1\}^s \rightarrow \{0, 1\}^s$ that generates a random s -bit response for each new input. Denote the resulting suffix keyed sponge construction by F^{p, R_s} . By a straightforward reduction,

$$\begin{aligned} \Delta_{\mathcal{A}}(F_K^p, p; R_{*,t}, p) &\leq \Delta_{\mathcal{A}}(F^{p, R_s}, p; R_{*,t}, p) + \Delta_{\mathcal{A}'}(G_K; R_s) \\ &= \Delta_{\mathcal{A}}(F^{p, R_s}, p; R_{*,t}, p) + \text{Adv}_G^{\text{prf}}(\mathcal{A}'), \end{aligned} \quad (18)$$

where \mathcal{A}' is some adversary with construction complexity q , primitive complexity 0 (as G does not depend on p), and time complexity $\tau' \approx \tau$.

The remaining distance $\Delta_{\mathcal{A}}(F^{p, R_s}, p; R_{*,t}, p)$ of (18) is bounded using Theorem 2, noting that R_s is an instance of a random function family, which is 2^{-s} -uniform and 2^{-s} -universal (see Section 2.3).

6 Leakage Resilience of Suffix Keyed Sponge

We will demonstrate how Theorem 2 carries over to leakage resilience. We first describe the model of leakage resilient PRFs in Section 6.1. The leakage resilience of the suffix keyed sponge is stated in Section 6.2.

6.1 Leakage Resilient PRF Security

We transform the definition of PRF security of Section 2.2 to security in case of leakage resilience. We will restrict our focus to non-adaptive \mathcal{L} -resilience of Dodis and Pietrzak [DP10], where the adversary receives leakage under any leakage $L \in \mathcal{L}$ of the

scheme under investigation. We adopt the well-established approach [Pie09, YSPY10, FPS12, SPY+10, DP10, BMOS17] where the adversary has access to a leak-free version of the construction which it has to distinguish from random, and a leaky version, which it may use to obtain leakage.

Let $b, k, t, \lambda, \lambda' \in \mathbb{N}$ and $m \in \mathbb{N} \cup \{*\}$. Consider the suffix keyed sponge function $F : \{0, 1\}^k \times \{0, 1\}^m \rightarrow \{0, 1\}^t$ that internally uses a permutation $p \in \text{perm}(b)$ and a function $G : \{0, 1\}^k \times \{0, 1\}^s \rightarrow \{0, 1\}^s$. Let $\mathcal{L}_p = \{L_p : \{0, 1\}^b \times \{0, 1\}^b \rightarrow \{0, 1\}^{\lambda'}\}$ be a class of leakage functions independent of p , i.e., that do not internally evaluate p or p^{-1} , and let $\mathcal{L}_G = \{L_G : \{0, 1\}^k \times \{0, 1\}^s \times \{0, 1\}^s \rightarrow \{0, 1\}^{\lambda'}\}$ be a fixed, predetermined set that captures all allowed leakage functions on the instantiation of G . Write $\mathcal{L} = \mathcal{L}_p \times \mathcal{L}_G$. For any leakage function $L = (L_p, L_G) \in \mathcal{L}$, define by $[F_K^p]_L$ an evaluation of F that leaks $L_G(K, \text{left}_s(U), \text{left}_s(V))$ and $L_p(V, W)$, where U is the state before evaluation of G and where V and W are the state before and after the last evaluation of p in F_K^p (see Figure 1). Note that L_G must also operate on $\text{left}_s(U)$ and L_p must also operate on $\text{right}_{b-s}(V)$ and $\text{left}_t(W)$, despite that the adversary knows these values: the leakage functions may operate differently depending on these values. The leakage functions give the same leakage whenever the inputs are the same, but possibly different leakages for different inputs. The non-adaptive leakage resilient pseudorandom function (NALR-PRF) security of F against an adversary \mathcal{A} is defined as

$$\text{Adv}_F^{\text{nalr-prf}}(\mathcal{A}) = \max_{L \in \mathcal{L}} \Delta_{\mathcal{A}}([F_K^p]_L, F_K^p, p; [F_K^p]_L, R_{m,t}, p), \quad (19)$$

where the randomness is taken over the random drawing of $K \xleftarrow{\$} \{0, 1\}^k$, the random drawing of $p \xleftarrow{\$} \text{perm}(b)$, the definition of a function $R_{m,t} : \{0, 1\}^m \rightarrow \{0, 1\}^t$ that generates a random t -bit response for each new input, and the random coins of \mathcal{A} .

The adversary never repeats construction queries (both to the leaky and leak-free oracle). The adversary is typically bounded by three types of complexities: q denotes the construction complexity, the number of queries it may make to $([F_K^p]_L, F_K^p)$ or $([F_K^p]_L, R_{m,t})$; N denotes the total primitive complexity, which accounts for the total number of evaluations of p (either direct or via the real-world construction oracle); τ denotes the time complexity.

We briefly elaborate on the scope of our model. Any leakage function L that gives a bounded amount of information about its inputs is fine. For instance, a probe placed on a wire within a circuit falls into this category. The only restriction we have is that L is chosen at the beginning of the experiment, and hence stays the same for every evaluation. E.g., a function that gives for a permutation call $p(K||A_1)$ the first λ bits of K and for $p(K||A_2)$ the second λ bits of K is fine. However, the function is not allowed to be change, e.g., getting for $p(K||A_1)$ the first λ bits of K and for a later same call $p(K||A_1)$ the second λ bits of K is not allowed. The parameter λ is basically used to upper bound the *total amount* an adversary might learn from this single call $p(K||A_1)$, over the entire duration of the experiment.

6.2 Security of Suffix Keyed Sponge under Leakage

We will prove that the suffix keyed sponge is a leakage resilient PRF under the assumption that p is a random permutation and G is a $2^{-\delta}$ -uniform and $2^{-\varepsilon}$ -universal function that is strongly protected against side-channel attacks. Stated differently, we inherit the assumptions of Theorem 2, and *in addition* require that G is strongly protected.

In a bit more detail, we say that G is strongly protected if it is $2^{-\delta}$ -uniform and $2^{-\varepsilon}$ -universal *even under internal leakage*, i.e., if for any leakage function $L_G \in \mathcal{L}_G$, any auxiliary leakage function $L_{\text{aux}} : \{0, 1\}^s \rightarrow \{0, 1\}^{\nu}$ (for some ν), and any $X, X', Y \in \{0, 1\}^s$

with $X \neq X'$,

$$\Pr(G(K, X) = Y \mid \{L_G(K, Z, G(K, Z))\}_{Z \in \{0,1\}^s} \wedge L_{\text{aux}}(G(K, X))) \leq 2^{-(\delta-\nu)}, \quad (20)$$

$$\Pr(G(K, X) = G(K, X') \mid \{L_G(K, Z, G(K, Z))\}_{Z \in \{0,1\}^s}) \leq 2^{-\varepsilon}, \quad (21)$$

where the randomness is taken over the random drawing of $K \xleftarrow{\$} \{0,1\}^k$. The condition that earlier leakage must be considered over any possible Z is generous: it can be refined to only consider values Z for which G actually gets evaluated in the suffix keyed sponge construction. However, in many applications, including the one that we will discuss in the next paragraph, the generous conditioning is tolerable. The auxiliary leakage function L_{aux} is required to capture leakage on $G(K, X)$ that might be gathered from outside G , i.e., from leaky evaluations of p . Indeed, in our construction, the leaky evaluation of p gets V as input, where $\text{left}_s(V) = G(K, \text{left}_s(U))$, and we must take this leakage into account when investigating leakage coming from G .

The assumption that G is strongly protected is well-established, see [BKP⁺18, BPPS17, GPPS19] among others, with the difference that we apply it to the definition of universal hashing. By adopting the generous definition of strongly protected universal G , we can separate leakage of G from leakage of p , and thus discard of leakage incurred by evaluations of G in the proof: they are implicit in the uniformity terms δ and ε .

One way to construct a suitable G is by instantiating it with an ‘‘adjusted ideal extendable input function (AIXIF)’’ in light of the formalism of Dobraunig and Mennink [DM19a]. Simply said, an AIXIF is the ideal equivalent of a leakage resilient keyed duplex construction, and Dobraunig and Mennink proved that the keyed duplex is indistinguishable from an AIXIF (under certain conditions) even if the construction leaks upon each duplexing call. Such a function would, indeed, do the job in the instantiation of G , as an AIXIF leaks no useful information, except, possibly, a constant amount of information about the key that is the same for each evaluation of the AIXIF. Clearly, upon instantiation of the AIXIF as a leakage resilient keyed duplex [DM19a], one has to take into account the distinguishing advantage with respect to the AIXIF. A concrete example of how a leakage resilient keyed duplex can be used to instantiate G in the suffix keyed sponge is given by Dobraunig and Mennink [DM19b, Proposition 6]. In this example, G is instantiated with a k -bit keyed AIXIF that is strongly protected $2^{-(k-\lambda)}$ -uniform and $2^{-(k-\lambda)}$ -universal, where λ denotes the maximum amount of leakage incurred by evaluations of the cryptographic permutation within the mode.

We are ready to state the theorem on the leakage resilience of the suffix keyed sponge.

Theorem 3. *Let $b, c, r, k, s, t, \lambda \in \mathbb{N}$ such that $c + r = b$ and $k, s, t \leq b$. Consider the suffix keyed sponge of Section 3 based on random permutation $p \xleftarrow{\$} \text{perm}(b)$ and a function $G : \{0,1\}^k \times \{0,1\}^s \rightarrow \{0,1\}^s$. Assume that G is strongly protected $2^{-\delta}$ -uniform and $2^{-\varepsilon}$ -universal. For any adversary \mathcal{A} with construction complexity $q \geq 2$ and primitive complexity $N \leq 2^{b-1}$,*

$$\text{Adv}_F^{\text{nalr-prf}}(\mathcal{A}) \leq \frac{2N^2}{2^c} + \frac{\mu_{s,b-s}^{2(N-q)}}{2^{b-s}} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^{\min\{\delta,\varepsilon\} - \mu_{s,b-s}^{2(N-q)} \lambda}} + \frac{\mu_{t,b-t}^{2q} \cdot N}{2^{b-t-\lambda}},$$

where we recall that λ denotes the maximum amount of information leaked in one leaky evaluation of p .

The proof is given in Section 6.3.

6.3 Proof of Theorem 3

The proof is a mere extension of that of Section 5.1, but as the adversary is given access to an additional oracle, and gets leakage for that oracle, care must be taken in the formal application of the H-coefficient technique. In current proof, we highlight the changes.

Let $K \xleftarrow{\$} \{0,1\}^k$ and $p \xleftarrow{\$} \text{perm}(b)$. Let $R_{*,t} : \{0,1\}^* \rightarrow \{0,1\}^t$ be a function that generates a random t -bit response for each new input. Let $L = (L_p, L_G) \in \mathcal{L}$ be any leakage function. Consider any adversary \mathcal{A} , whose goal is to maximize

$$\Delta_{\mathcal{A}}([F_K^p]_L, F_K^p, p; [F_K^p]_L, R_{*,t}, p), \quad (22)$$

We will call $\mathcal{O} := ([F_K^p]_L, F_K^p, p)$ the real world and $\mathcal{P} := ([F_K^p]_L, R_{*,t}, p)$ the ideal world. As before, we assume that \mathcal{A} is deterministic.

6.3.1 Oracle Interaction

As before, the adversary has a total primitive complexity N , meaning that for the *real world* the total number primitive queries (direct or through $[F_K^p]_L$ or F_K^p) does not exceed N . We again require \mathcal{A} to make all unkeyed primitive queries itself offline, and all primitive evaluations (to p , in both worlds) are summarized in a view

$$v_p = \{(\text{dir}_1, X_1, Y_1), \dots, (\text{dir}_{N-q}, X_{N-q}, Y_{N-q})\}.$$

Any construction query to $[F_K^p]_L$ leaks information, namely through evaluations of G_K and through evaluations of the last permutation call. As discussed above, we can restrict our focus to leakage from p : leakage from G_K is captured implicitly in the universality terms δ and ϵ , as we will also detail later on. Thus, in our setting, any construction query to $[F_K^p]_L$ leaks the value $L_p(V, W) \in \{0,1\}^\lambda$. Noting that, in fact, the adversary also learns the values $\text{right}_{b-s}(V)$ and $\text{left}_t(W)$, we will model the leakage more generously: we assume that for each query the adversary learns λ bits of $\text{left}_s(V)$, called $\ell_V \in \{0,1\}^\lambda$, and λ bits of $\text{right}_{b-t}(W)$, called $\ell_W \in \{0,1\}^\lambda$. For consistency, we simply define $\ell_V = \ell_W = \perp$ if the corresponding query is made to the challenge oracle (F_K^p in the real world or $R_{*,t}$ in the ideal world). All q construction queries (to $([F_K^p]_L, F_K^p)$ in the real world and to $([F_K^p]_L, R_{*,t})$ in the ideal world) are summarized in a view

$$v_c = \{(P_1, T_1, \ell_{V_1}, \ell_{W_1}), \dots, (P_q, T_q, \ell_{V_q}, \ell_{W_q})\}.$$

(One can deduce from (ℓ_{V_i}, ℓ_{W_i}) whether the i -th query was made to the leaky or leak-free oracle.) After the conversation of \mathcal{A} with its oracle, but before it outputs its decision bit, we reveal the key K and the values Z_i as before, but now the definition of these values Z_i is slightly more delicate. For queries to $[F_K^p]_L$ (in either world) or to F_K^p (in the real world only), the values Z_i are the values that are truncated at the end of the construction query; for $R_{*,t}$, these are random values $Z_i \xleftarrow{\$} \{0,1\}^{b-t}$. The values Z_i are appended to v_c :

$$v'_c = \{(P_1, T_1, Z_1, \ell_{V_1}, \ell_{W_1}), \dots, (P_q, T_q, Z_q, \ell_{V_q}, \ell_{W_q})\}.$$

The aggregate view is defined as $v = (v'_c, v_p, K)$.

Note that the definitions of U_i , V_i , and W_i (see also (8)-(9)) carry over. In particular, if the i -th query was made to the leaky oracle $[F_K^p]_L$, the leakages ℓ_{V_i} and ℓ_{W_i} are consistent with the values V_i and W_i .

6.3.2 Bad Views

We inherit the bad events of Section 5.1, but now with thresholds $\nu_{\text{right}} = \mu_{b-s,s}^{2(N-q)}$ and $\nu_{\text{tag}} = \mu_{t,b-t}^{2q}$. We furthermore need to split event coll_{cp} into input and output collisions:

- $\text{coll}_{\text{cp-in}}$: there exist $i \in \{1, \dots, q\}$ and $j \in \{1, \dots, N - q\}$ with

$$V_i = X_j;$$

- $\text{coll}_{\text{cp-out}}$: there exist $i \in \{1, \dots, q\}$ and $j \in \{1, \dots, N - q\}$ with

$$W_i = Y_j.$$

The reason for this separation is that in the proof of Section 5.1, $\text{coll}_{\text{cp-out}}$ is proven under the non-occurrence of mc_{tag} , but as the view v'_c also includes queries to $[F_K^p]_L$, one can only upper bound the probability that mc_{tag} occurs under the condition that $\text{coll}_{\text{cp-in}}$ does not occur.

Finally, we define a helping event used to bound the amount of leakage. Let $\nu_{\text{left}} = \mu_{s, b-s}^{2(N-q)}$ be a threshold.

- mc_{left} : there exist distinct $j_1, \dots, j_{\nu_{\text{left}}+1} \in \{1, \dots, N - q\}$ with $\text{dir}_j = +$ and

$$\text{left}_s(Y_{j_1}) = \dots = \text{left}_s(Y_{j_{\nu_{\text{left}}+1}}).$$

The event mc_{left} is motivated by the observation that any time G is evaluated for the same input $\text{left}_s(U_i)$ but alongside a different inner part $\text{right}_{b-s}(U_i)$ (see (8)), the adversary may learn different information about $\text{left}_s(V_i) = G(K, \text{left}_s(U_i))$. The bad event mc_{left} upper bounds the maximum number of times G is invoked on the same input $\text{left}_s(U_i)$.

We write $\text{cap} = \text{cap}_{\text{fwd}} \vee \text{cap}_{\text{inv}}$ and

$$\text{bad} = \text{cap} \vee \text{mc}_{\text{left}} \vee \text{mc}_{\text{right}} \vee \text{mc}_{\text{tag}} \vee \text{coll}_{\text{cc}} \vee \text{coll}_{\text{cp-in}} \vee \text{coll}_{\text{cp-out}}. \quad (23)$$

6.3.3 Probability of Bad View

Our goal is to bound the probability that a bad view is generated in the ideal world, $\Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}})$. This probability equals the probability that a view v generated by $D_{\mathcal{P}}$ sets bad . By basic probability theory,

$$\begin{aligned} \Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}}) &= \Pr(\text{bad}) \\ &\leq \Pr(\text{cap}) + \Pr(\text{mc}_{\text{left}}) + \Pr(\text{mc}_{\text{right}}) + \Pr(\text{mc}_{\text{tag}} \mid \neg \text{coll}_{\text{cp-in}}) \\ &\quad + \Pr(\text{coll}_{\text{cc}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}})) \\ &\quad + \Pr(\text{coll}_{\text{cp-in}} \mid \neg(\text{cap} \vee \text{mc}_{\text{left}} \vee \text{mc}_{\text{right}})) \\ &\quad + \Pr(\text{coll}_{\text{cp-out}} \mid \neg(\text{cap} \vee \text{mc}_{\text{tag}} \vee \text{coll}_{\text{cp-in}})). \end{aligned} \quad (24)$$

Much of the analysis of the proof of Section 5.1 carries over verbatim, but there are some differences, most notably in the analysis of mc_{left} (the new bad event) and coll_{cp} (leakages must be taken into account).

cap. The analysis is identical to the one in Section 5.1, and we obtain

$$\Pr(\text{cap}) = \Pr(\text{cap}_{\text{fwd}} \vee \text{cap}_{\text{inv}}) \leq \sum_{j=1}^{N-q} \frac{2(2j-1)}{2^c} = \frac{2(N-q)^2}{2^c}.$$

mc_{left}. The analysis is symmetric to the one of mc_{right} in Section 5.1, and we obtain

$$\Pr(\text{mc}_{\text{left}}) \leq \frac{\mu_{s, b-s}^{2(N-q)}}{2^{b-s}}.$$

mc_{right}. The analysis is identical to the one in Section 5.1, and we obtain

$$\Pr(\text{mc}_{\text{right}}) \leq \frac{\mu_{b-s, s}^{2(N-q)}}{2^s}.$$

mc_{tag}. Note that the tuples in v'_c either come from $R_{*,t}$ or $[F_K^p]_L$. In the former case, the values W_i are randomly generated from $\{0,1\}^b$. In the latter case, by $\neg\text{coll}_{\text{cp-in}}$, they are generated uniformly randomly without replacement. The event mc_{tag} is thus a balls-and-bins experiment with q balls randomly (with some distribution) thrown into 2^t bins, in such a way that any of the bins contains more than ν_{tag} balls. Note that the distribution of balls satisfies the condition of (3). By definition of $\nu_{\text{tag}} = \mu_{t,b-t}^{2q}$, we can resort to the multicollision limit function of Section 2.4 and obtain

$$\Pr(\text{mc}_{\text{tag}} \mid \neg\text{coll}_{\text{cp-in}}) \leq \frac{\mu_{t,b-t}^{2q}}{2^{b-t}}.$$

coll_{cc}. The analysis for the case of $V_i = V_{i'}$ is identical to the one in Section 5.1. Here, we note that, particularly, the probability analysis does not change as it is independent of the leakage that happens outside G . That leakage, in turn, is captured within the $2^{-\varepsilon}$ -universality of G .

For the case of $W_i = W_{i'}$, we have to be a bit more careful as v'_c contains queries to $[F_K^p]_L$, where the W_i 's are the outputs of a permutation, as well as $R_{*,t}$, where the W_i 's are random b -bit values. If both queries i and i' are made to $[F_K^p]_L$, we have $W_i = W_{i'}$ if and only if $V_i = V_{i'}$, but that case was already covered above. If at least one of the two queries is made to $R_{*,t}$, equation $W_i = W_{i'}$ holds with probability at most $1/(2^b - (j-1)) \leq 2/2^b$, using that $j-1 \leq 2^{b-1}$. There are at most $\binom{q}{2} \leq q^2/2$ choices $i, i' \in \{1, \dots, q\}$ such that at least one of the two queries is made to $R_{*,t}$.

We obtain

$$\Pr(\text{coll}_{\text{cc}} \mid \neg(\text{cap} \vee \text{mc}_{\text{right}})) \leq \frac{\mu_{b-s,s}^{2(N-q)} \cdot q/2}{2^\varepsilon} + \frac{q^2}{2^b}.$$

coll_{cp-in}. Fix any primitive query (dir_j, X_j, Y_j) . We make a distinction depending on the direction of this query.

If $\text{dir}_j = +$, the equation $V_i = X_j$ is equivalent to stating that

$$G(K, \text{left}_s(U_i)) \parallel \text{right}_{b-s}(U_i) = \text{left}_s(X_j) \parallel \text{right}_{b-s}(X_j).$$

As for the analysis coll_{cc} , by $\neg\text{cap}$ and $\neg\text{mc}_{\text{right}}$, there are at most $\nu_{\text{right}} = \mu_{b-s,s}^{2(N-q)}$ construction queries with $\text{right}_{b-s}(U_i) = \text{right}_{b-s}(X_j)$. By $\neg\text{mc}_{\text{left}}$, the adversary might have seen at most $\nu_{\text{left}} = \mu_{s,b-s}^{2(N-q)}$ leakages for $G(K, \text{left}_s(U_i))$. Each of these leakages reduces the entropy of the state by at most λ bits. As G itself is $2^{-\delta}$ -uniform (and this covers leakage within G itself), any of the queries satisfies above equation with probability at most $2^{-(\delta - \mu_{s,b-s}^{2(N-q)} \lambda)}$. Here, we explicitly note that (20) is defined conditioned on auxiliary leakage, which in current analysis is of size $\nu := \mu_{s,b-s}^{2(N-q)} \lambda$ bits.

If $\text{dir}_j = -$, the case of equation $V_i = X_j$ is symmetric, as in the proof of Section 5.1, but now with probability $2^{-\delta}$ as no leakage can be exploited.

Aggregating both cases, we obtain

$$\Pr(\text{coll}_{\text{cp-in}} \mid \neg(\text{cap} \vee \text{mc}_{\text{left}} \vee \text{mc}_{\text{right}})) \leq \frac{\mu_{b-s,s}^{2(N-q)} \cdot (N-q)}{2^{\delta - \mu_{s,b-s}^{2(N-q)} \lambda}}.$$

coll_{cp-out}. Fix any primitive query (dir_j, X_j, Y_j) . We make a distinction depending on the direction of this query.

If $\text{dir}_j = -$, the equation $W_i = Y_j$ is equivalent to stating that

$$T_i \parallel Z_i = \text{left}_s(Y_j) \parallel \text{right}_{b-s}(Y_j).$$

By $\neg \text{mc}_{\text{tag}}$, there are at most $\nu_{\text{tag}} = \mu_{t,b-t}^{2q}$ construction queries with $T_i = \text{left}_s(Y_j)$. We have $Z_i \stackrel{\$}{\leftarrow} \{0,1\}^{b-t}$, but the adversary may have learned information about this value through leakage. There could be at most one leakage for this value, and this leakage reduces the entropy by at most λ bits. Therefore any of these construction queries satisfies above equation with probability at most $1/2^{b-t-\lambda}$.

If $\text{dir}_j = +$ and the i -th query is to $[F_K^p]_L$, the equation cannot be set by condition $\neg \text{coll}_{\text{cp-in}}$. If the i -th query is to $R_{*,t}$, the equation $W_i = Y_j$ happens with probability at most $q/(2^b - (j-1)) \leq 2q/2^b$, using that $j-1 \leq 2^{b-1}$.

Aggregating both cases, we obtain

$$\Pr(\text{coll}_{\text{cp-out}} \mid \neg(\text{cap} \vee \text{mc}_{\text{tag}} \vee \text{coll}_{\text{cp-in}})) \leq \frac{\mu_{t,b-t}^{2q} \cdot (N-q)}{2^{b-t-\lambda}} + \frac{2q(N-q)}{2^b}.$$

Conclusion. Summing the individual terms, we obtain for (24) that

$$\Pr(D_{\mathcal{P}} \in \mathcal{V}_{\text{bad}}) \leq \frac{2N^2}{2^c} + \frac{\mu_{s,b-s}^{2(N-q)}}{2^{b-s}} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^{\min\{\delta,\varepsilon\} - \mu_{s,b-s}^{2(N-q)} \lambda}} + \frac{\mu_{t,b-t}^{2q} \cdot N}{2^{b-t-\lambda}}, \quad (25)$$

where we used that $\min\{\delta,\varepsilon\} \leq s$ and $q \geq 2$.

6.3.4 Ratio for Good Views

Consider any good view $v \in \mathcal{V}_{\text{good}}$. We will prove that $\Pr(D_{\mathcal{O}} = v) \geq \Pr(D_{\mathcal{P}} = v)$.

Real World. For the real world, the analysis is exactly as that of Section 5.1, using that the leakage values ℓ_{V_i}, ℓ_{W_i} in v'_c – the only difference compared with the analysis in Section 5.1 – contain no additional information after the values Z_i and K have been revealed. Therefore, we obtain:

$$\Pr(D_{\mathcal{O}} = v) = \frac{1}{(2^b)_N} \cdot \frac{1}{2^k}. \quad (26)$$

Ideal World. For the ideal world, the analysis is more involved, as construction queries may be to the real but leaky oracle $[F_K^p]_L$ or to the ideal oracle $R_{*,t}$. Assume that v contains $q' \leq q$ queries to the ideal oracle. By independence of $R_{*,t}$ from $([F_K^p]_L, p)$, we can resort to the analysis for the real world for all elements in v except for the q' queries to $R_{*,t}$. For these q' queries, we can observe that the outputs $T_i \stackrel{\$}{\leftarrow} \{0,1\}^t$ and dummy values $Z_i \stackrel{\$}{\leftarrow} \{0,1\}^{b-t}$ are randomly generated, as before. Therefore, we obtain:

$$\Pr(D_{\mathcal{P}} = v) = \frac{1}{(2^b)_{N-q'}} \cdot \frac{1}{(2^t)^{q'}} \cdot \frac{1}{(2^{b-t})^{q'}} \cdot \frac{1}{2^k} = \frac{1}{(2^b)_{N-q'}} \cdot \frac{1}{2^{bq'}} \cdot \frac{1}{2^k}. \quad (27)$$

Bounding the Ratio. Combining (26) and (27), we obtain the following for any good view $v \in \mathcal{V}_{\text{good}}$:

$$\frac{\Pr(D_{\mathcal{O}} = v)}{\Pr(D_{\mathcal{P}} = v)} = \frac{2^{bq'}}{(2^b - (N - q'))_{q'}} \geq 1. \quad (28)$$

6.3.5 Conclusion

The H-coefficient technique of Lemma 1 gives, using (25) and (28):

$$\begin{aligned} \Delta_{\mathcal{A}}([F_K^p]_L, F_K^p, p; [F_K^p]_L, R_{m,t}, p) \\ \leq \frac{2N^2}{2^c} + \frac{\mu_{s,b-s}^{2(N-q)}}{2^{b-s}} + \frac{\mu_{b-s,s}^{2(N-q)} \cdot N}{2^{\min\{\delta,\varepsilon\} - \mu_{s,b-s}^{2(N-q)} \lambda}} + \frac{\mu_{t,b-t}^{2q} \cdot N}{2^{b-t-\lambda}}. \end{aligned}$$

As this holds for any leakage function $L \in \mathcal{L}$, this completes the proof.

7 Application

One clear application of our bound is the message authentication part of ISAP [DEM⁺17]. The MAC of ISAP is based on a suffix keyed sponge construction that uses a 400-bit permutation with rate $r = 144$, capacity $c = 256$, and $s = t = 128$. According to Daemen et al. [DMV17, Equation (32)], we compute $\mu_{t,b-t}^q = \mu_{128,272}^{2^{128}} \leq 80$ and $\mu_{b-s,s}^{2(N-q)} = \mu_{272,128}^{2^{129}} \leq 3$, where, w.l.o.g., we can assume that $q, N \leq 2^{c/2} = 2^{128}$. If we assume that G is 2^{-128} -uniform, and 2^{-128} -universal, then an attacker has advantage

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) \leq \frac{2N^2}{2^{256}} + \frac{3N}{2^{128}} + \frac{80N}{2^{272}}.$$

If we consider a smaller permutation of, e.g., 320 bits like the one used for Ascon [DEMS16, DEMS19] and aim to achieve a similar security level, we have to change the rate to $r = 64$, but can leave $c = 256$ and $s = t = 128$. Then, we get $\mu_{t,b-t}^q = \mu_{128,192}^{2^{128}} \leq 67$ and $\mu_{b-s,s}^{2(N-q)} = \mu_{192,128}^{2^{129}} \leq 5$. In this case, we get a bound

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) \leq \frac{2N^2}{2^{256}} + \frac{5N}{2^{128}} + \frac{67N}{2^{272}}.$$

Comparing both construction, we see that they achieve a similar security level, namely 128 bits up to small constant factors. The biggest difference in both bounds is the factor $\mu_{b-s,s}^{2(N-q)}$, which changes from 3 to 5. This is a marginal change considering that the inner part of the state for the 320-bit construction is just 192 bits, which is significantly smaller than the 256-bit capacity one would need in order to rely on the indistinguishability result of the sponge [BDPV08].

In previous examples, we upper bounded the values of $\mu_{t,b-t}^q$ and $\mu_{b-s,s}^{2(N-q)}$ by upper bounding the query complexities q and N to 2^{128} . However, if we consider the leakage resilience bound of Theorem 3, doing the same thing would distort the advantage, since the advantage an attacker gains by exploiting the leakage is tightly coupled with $\mu_{s,b-s}^{2(N-q)}$. Hence, we have to consider $\mu_{s,b-s}^{2(N-q)}$ as a function, which grows for a fixed b and s if $2(N-q)$ increases. To get some insight what this means for practical cases, let us consider an attacker with complexities $N = 2^{65}$ and $q = 2^{64}$. Furthermore, assume that $r = 64$, $c = 256$, and $s = t = 128$ as in the example above. In this case, we get $\mu_{s,b-s}^{2(N-q)} = \mu_{t,b-t}^{2q} = \mu_{128,192}^{2^{65}} \leq 5$ and $\mu_{b-s,s}^{2(N-q)} = \mu_{192,128}^{2^{65}} \leq 3$. If we assume that G is strongly protected 2^{-128} -uniform and 2^{-128} -universal, we get

$$\text{Adv}_F^{\text{nalr-prf}}(\mathcal{A}) \leq \frac{1}{2^{61}} + \frac{5}{2^{192}} + \frac{3 \cdot 2^{65}}{2^{128-5\lambda}} + \frac{5 \cdot 2^{65}}{2^{192-\lambda}},$$

meaning that in this case λ can be up to 12 bits while still having an advantage smaller than one.

8 Conclusion and Discussion

Towards Leakage Resilient Message Authentication. In this paper, we provided the – to the best of our knowledge – first dedicated analysis of the suffix keyed sponge and bounded the advantage of an attacker in distinguishing this construction from a PRF in the black-box security model and in the leakage resilience model. However, we stress that such a leakage resilient PRF not always results in a leakage resilient MAC.

This remark seems counter-intuitive, but recently, Berti et al. [BGP⁺19] and Guo et al. [GPPS19] detailed a problem with MACs during the verification. In most cases, this is done by computing a tag T' from the transmitted data M by using the PRF and compare it with the transmitted tag T . If this comparison is not done in a leakage resilient manner, information about T' can leak. Thus, an attacker could learn information about T' by repeated queries of the same M with different T . If all the information about T' is learned, a forgery $M \parallel T'$ can be created.

In their schemes, Berti et al. [BGP⁺19] and Guo et al. [GPPS19] overcome this issue in their hash-then-MAC construction by computing the inverse of a perfectly protected tweakable block cipher on the transmitted T to compare directly the outcome of the hash function, so no information about T' can leak. The approach is also used in Spook [BBB⁺19]. Clearly, this can also be done for a MAC based on a suffix keyed sponge if we do not use T directly, but also put T in a perfectly protected block cipher and transmit $E_K(T)$, and also doing the comparison via the inverse. Another way to protect the comparison is, as suggested by Dobraunig et al. [DEM⁺19], via one additional permutation call during the verification. In this case, T and T' are computed and transmitted as normal, but instead of a direct comparison, e.g., $\text{left}_t(p(T' \parallel 0^{b-t}))$ is compared with $\text{left}_t(p(T \parallel 0^{b-t}))$ first. In such a comparison, quite in a similar manner to [GPPS19, BGP⁺19], the comparison only reveals mostly useless information, namely only at most t out of $b - t$ bits of $p(T' \parallel 0^{b-t})$, which gives no advantage in retrieving T' .

Estimation of Multicollision Probabilities. The leakage resilience security bound of Theorem 3 has a term that bounds the probability of a large multicollision in the outer part of the state:

$$\frac{\mu_{s,b-s}^{2(N-q)}}{2^{b-s}}.$$

The bound is based on the multicollision limit function bounding from Section 2.4. This bound, however, shows an unexpected behavior if $s = b$, i.e., if one opts to absorb the key over the full state: the bound of Theorem 3 becomes meaningless. However, this is purely due to generality. For specific cases, in particular when $b - s \ll c$, other choices for the threshold value ν_{left} would give a better bound. For example, in case $b = s$ one would simply not even have collisions in the first place (by assumption that $\neg\text{cap}$ holds).

Suffix Keyed Sponge for Short Messages. The description of the suffix keyed sponge of Section 3 is given in a general manner in such a way that it still resembles as much as possible from the sponge. One can improve the efficiency of the suffix keyed sponge slightly by absorbing more data bits in the first round. Let $a \in \mathbb{N}$ satisfy $a \leq c/2$, and consider an adjustment of the suffix keyed sponge where the first block is of size $r + a$ bits, but all remaining blocks are of size r bits. In the security proofs of both the black-box and the leakage resilience case, this *only affects the definition and analysis of cap_{inv}* , which is set if an inverse query satisfies $\text{right}_{c-a}(X_j) = 0^{c-a}$. This, in turn, happens with probability at most $2(N - q)/2^{c-a}$. Summarizing, one can increase the rate for the first block by $a \leq c/2$ bits, and this only adds $2(N - q)/2^{c-a}$ to the bound of Theorem 2, Corollary 1, or Theorem 3. This term would not dominate the bound as long as $a \leq c/2$.

Acknowledgments

We thank Joan Daemen, Maria Eichlseder, and Florian Mendel for the fruitful discussions. Christoph Dobraunig is supported by the Austrian Science Fund (FWF): J 4277-N38. Bart Mennink is supported by a postdoctoral fellowship from the Netherlands Organisation for

Scientific Research (NWO) under Veni grant 016.Veni.173.017. We thank the reviewers and shepherds of ToSC 2019 for their valuable comments.

References

- [ADMV15] Elena Andreeva, Joan Daemen, Bart Mennink, and Gilles Van Assche. Security of Keyed Sponge Constructions Using a Modular Proof Approach. In Gregor Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 364–384. Springer, 2015.
- [BBB⁺19] Davide Bellizia, Francesco Berti, Olivier Bronchain, Gaétan Cassiers, Sébastien Duval, Chun Guo, Gregor Leander, Gaétan Leurent, Itamar Levi, Charles Momin, Olivier Pereira, Thomas Peters, François-Xavier Standaert, and Friedrich Wiemer. Spook: Sponge-Based Leakage-Resilient Authenticated Encryption with a Masked Tweakable Block Cipher. Submission to NIST Lightweight Cryptography, 2019.
- [BDPV08] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the Indifferentiability of the Sponge Construction. In Nigel P. Smart, editor, *EUROCRYPT 2008*, volume 4965 of *LNCS*, pages 181–197. Springer, 2008.
- [BDPV11a] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Cryptographic sponge functions, January 2011.
- [BDPV11b] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Duplexing the Sponge: Single-Pass Authenticated Encryption and Other Applications. In Ali Miri and Serge Vaudenay, editors, *SAC 2011*, volume 7118 of *LNCS*, pages 320–337. Springer, 2011.
- [BDPV11c] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. On the Security of the Keyed Sponge Construction. Symmetric Key Encryption Workshop, February 2011.
- [BGM04] Mihir Bellare, Oded Goldreich, and Anton Mityagin. The Power of Verification Queries in Message Authentication and Authenticated Encryption. Cryptology ePrint Archive, Report 2004/309, 2004.
- [BGP⁺19] Francesco Berti, Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. TEDT, a Leakage-Resilient AEAD mode for High (Physical) Security Applications. Cryptology ePrint Archive, Report 2019/137, 2019.
- [BKP⁺16] Francesco Berti, François Koeune, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Leakage-Resilient and Misuse-Resistant Authenticated Encryption. Cryptology ePrint Archive, Report 2016/996, 2016.
- [BKP⁺18] Francesco Berti, François Koeune, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Ciphertext Integrity with Misuse and Leakage: Definition and Efficient Constructions with Symmetric Primitives. In Jong Kim, Gail-Joon Ahn, Seungjoo Kim, Yongdae Kim, Javier López, and Taesoo Kim, editors, *AsiaCCS 2018*, pages 37–50. ACM, 2018.
- [BKR94] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of cipher block chaining. In Yvo Desmedt, editor, *CRYPTO '94*, volume 839 of *LNCS*, pages 341–358. Springer, 1994.

- [BMOS17] Guy Barwell, Daniel P. Martin, Elisabeth Oswald, and Martijn Stam. Authenticated Encryption in the Face of Protocol and Side Channel Leakage. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017*, volume 10624 of *LNCS*, pages 693–723. Springer, 2017.
- [BN05] Paulo S. L. M. Barreto and Michael Naehrig. Pairing-Friendly Elliptic Curves of Prime Order. In Bart Preneel and Stafford E. Tavares, editors, *SAC 2005*, volume 3897 of *LNCS*, pages 319–331. Springer, 2005.
- [BPPS17] Francesco Berti, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. On Leakage-Resilient Authenticated Encryption with Decryption Leaks. *IACR Trans. Symmetric Cryptol.*, 2017(3):271–293, 2017.
- [BR02] John Black and Phillip Rogaway. A Block-Cipher Mode of Operation for Parallelizable Message Authentication. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 384–397. Springer, 2002.
- [CDH⁺12] Donghoon Chang, Morris Dworkin, Seokhie Hong, John Kelsey, and Mridul Nandi. A Keyed Sponge Construction with Pseudorandomness in the Standard Model. NIST SHA-3 Workshop, March 2012.
- [CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards Sound Approaches to Counteract Power-Analysis Attacks. In Wiener [Wie99], pages 398–412.
- [CS14] Shan Chen and John P. Steinberger. Tight Security Bounds for Key-Alternating Ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, 2014.
- [DEM⁺17] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, and Thomas Unterluggauer. ISAP - Towards Side-Channel Secure Authenticated Encryption. *IACR Trans. Symmetric Cryptol.*, 2017(1):80–105, 2017.
- [DEM⁺19] Christoph Dobraunig, Maria Eichlseder, Stefan Mangard, Florian Mendel, Bart Mennink, Robert Primas, and Thomas Unterluggauer. ISAP v2. Submission to NIST Lightweight Cryptography, 2019.
- [DEMS16] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer. Ascon v1.2. Submission to Round 3 of the CAESAR competition, 2016.
- [DEMS19] Christoph Dobraunig, Maria Eichlseder, Florian Mendel, and Martin Schl affer. Ascon v1.2. Submission to NIST Lightweight Cryptography, 2019.
- [DM19a] Christoph Dobraunig and Bart Mennink. Leakage Resilience of the Duplex Construction. In Steven Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019*, LNCS, 2019. to appear.
- [DM19b] Christoph Dobraunig and Bart Mennink. Leakage Resilience of the ISAP Mode: a Vulgarized Summary, 2019. NIST Lightweight Cryptography Workshop 2019.
- [DMV17] Joan Daemen, Bart Mennink, and Gilles Van Assche. Full-State Keyed Duplex with Built-In Multi-user Support. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017*, volume 10625 of *LNCS*, pages 606–637. Springer, 2017.

- [DMV18] Joan Daemen, Bart Mennink, and Gilles Van Assche. Sound Hashing Modes of Arbitrary Functions, Permutations, and Block Ciphers. *IACR Trans. Symmetric Cryptol.*, 2018(4):197–228, 2018.
- [DP10] Yevgeniy Dodis and Krzysztof Pietrzak. Leakage-Resilient Pseudorandom Functions and Side-Channel Attacks on Feistel Networks. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 21–40. Springer, 2010.
- [Dwo05] Morris Dworkin. NIST SP 800-38B: Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication , 2005.
- [FPS12] Sebastian Faust, Krzysztof Pietrzak, and Joachim Schipper. Practical Leakage-Resilient Symmetric Cryptography. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 213–232. Springer, 2012.
- [GGM84] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the Cryptographic Applications of Random Functions. In G. R. Blakley and David Chaum, editors, *CRYPTO '84*, volume 196 of *LNCS*, pages 276–288. Springer, 1984.
- [GP99] Louis Goubin and Jacques Patarin. DES and Differential Power Analysis (The “Duplication” Method). In Çetin Kaya Koç and Christof Paar, editors, *CHES'99*, volume 1717 of *LNCS*, pages 158–172. Springer, 1999.
- [GPPS19] Chun Guo, Olivier Pereira, Thomas Peters, and François-Xavier Standaert. Towards Lightweight Side-Channel Security and the Leakage-Resilience of the Duplex Sponge. Cryptology ePrint Archive, Report 2019/193, 2019.
- [GPT15] Peter Gazi, Krzysztof Pietrzak, and Stefano Tessaro. The Exact PRF Security of Truncation: Tight Bounds for Keyed Sponges and Truncated CBC. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO 2015*, volume 9215 of *LNCS*, pages 368–387. Springer, 2015.
- [JLM14] Philipp Jovanovic, Atul Luykx, and Bart Mennink. Beyond $2^{c/2}$ Security in Sponge-Based Authenticated Encryption Modes. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014*, volume 8873 of *LNCS*, pages 85–104. Springer, 2014.
- [KJJ99] Paul C. Kocher, Joshua Jaffe, and Benjamin Jun. Differential Power Analysis. In Wiener [Wie99], pages 388–397.
- [Men18] Bart Mennink. Key Prediction Security of Keyed Sponges. *IACR Trans. Symmetric Cryptol.*, 2018(4):128–149, 2018.
- [MOSW15] Daniel P. Martin, Elisabeth Oswald, Martijn Stam, and Marcin Wójcik. A leakage resilient MAC. In Jens Groth, editor, *IMACC 2015*, volume 9496 of *LNCS*, pages 295–310. Springer, 2015.
- [MRV15] Bart Mennink, Reza Reyhanitabar, and Damian Vizár. Security of Full-State Keyed Sponge and Duplex: Applications to Authenticated Encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015*, volume 9453 of *LNCS*, pages 465–489. Springer, 2015.
- [MSJ12] Marcel Medwed, François-Xavier Standaert, and Antoine Joux. Towards Super-Exponential Side-Channel Security with Efficient Leakage-Resilient PRFs. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 193–212. Springer, 2012.

- [MSNF16] Marcel Medwed, François-Xavier Standaert, Ventsislav Nikov, and Martin Feldhofer. Unknown-Input Attacks in the Parallel Setting: Improving the Security of the CHES 2012 Leakage-Resilient PRF. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016*, volume 10031 of *LNCS*, pages 602–623, 2016.
- [NRR06] Svetla Nikova, Christian Rechberger, and Vincent Rijmen. Threshold Implementations Against Side-Channel Attacks and Glitches. In Peng Ning, Sihan Qing, and Ninghui Li, editors, *ICICS 2006*, volume 4307 of *LNCS*, pages 529–545. Springer, 2006.
- [NRS11] Svetla Nikova, Vincent Rijmen, and Martin Schl affer. Secure Hardware Implementation of Nonlinear Functions in the Presence of Glitches. *J. Cryptology*, 24(2):292–321, 2011.
- [NY16] Yusuke Naito and Kan Yasuda. New Bounds for Keyed Sponges with Extendable Output: Independence Between Capacity and Message Length. In Thomas Peyrin, editor, *FSE 2016*, volume 9783 of *LNCS*, pages 3–22. Springer, 2016.
- [Pat08] Jacques Patarin. The “Coefficients H” Technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, 2008.
- [Pie09] Krzysztof Pietrzak. A Leakage-Resilient Mode of Operation. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 462–482. Springer, 2009.
- [SD17] Niels Samwel and Joan Daemen. DPA on hardware implementations of Ascon and Keyak. In *CF’17*, pages 415–424. ACM, 2017.
- [SPY+10] François-Xavier Standaert, Olivier Pereira, Yu Yu, Jean-Jacques Quisquater, Moti Yung, and Elisabeth Oswald. Leakage Resilient Cryptography in Practice. In Ahmad-Reza Sadeghi and David Naccache, editors, *Towards Hardware-Intrinsic Security - Foundations and Practice*, Information Security and Cryptography, pages 99–134. Springer, 2010.
- [TS14] Mostafa M. I. Taha and Patrick Schaumont. Side-channel countermeasure for SHA-3 at almost-zero area overhead. In *HOST 2014*, pages 93–96. IEEE Computer Society, 2014.
- [Wie99] Michael J. Wiener, editor. *CRYPTO ’99*, volume 1666 of *LNCS*. Springer, 1999.
- [YSPY10] Yu Yu, François-Xavier Standaert, Olivier Pereira, and Moti Yung. Practical leakage-resilient pseudorandom generators. In Ehab Al-Shaer, Angelos D. Keromytis, and Vitaly Shmatikov, editors, *CCS 2010*, pages 141–151. ACM, 2010.