

A preliminary version of this paper appears in the proceedings of Eurocrypt 2016. This is the full version.

Hash-Function based PRFs: AMAC and its Multi-User Security

MIHIR BELLARE¹ DANIEL J. BERNSTEIN² STEFANO TESSARO³

2016.02.28

Abstract

AMAC is a simple and fast candidate construction of a PRF from an MD-style hash function which applies the keyed hash function and then a cheap, un-keyed output transform such as truncation. Spurred by its use in the widely-deployed Ed25519 signature scheme, this paper investigates the provable PRF security of AMAC to deliver the following three-fold message: (1) First, we prove PRF security of AMAC. (2) Second, we show that AMAC has a quite unique and attractive feature, namely that its multi-user security is essentially as good as its single-user security and in particular superior in some settings to that of competitors. (3) Third, it is technically interesting, its security and analysis intrinsically linked to security of the compression function in the presence of leakage.

¹ Department of Computer Science & Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, California 92093, USA. URL: <http://cseweb.ucsd.edu/~mihir/>. Supported in part by NSF grants CNS-1526801 and CNS-1228890, ERC Project ERCC FP7/615074 and a gift from Microsoft. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

² University of Illinois at Chicago, USA, and Technische Universiteit Eindhoven, Netherlands. URL: <https://cr.yp.to/djb.html>. Supported in part by NSF grant CNS-1314919 and NWO grant 639.073.005.

³ Department of Computer Science, University of California Santa Barbara, Santa Barbara, California 93106, USA. URL: <http://www.cs.ucsb.edu/~tessarar/>. Supported in part by NSF grant CNS-1423566. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

Contents

1	Introduction	3
2	Related work	7
3	Notation	8
4	Function-family distance framework	9
5	The augmented cascade and its analysis	11
6	Framework for ideal-model cryptography	17
7	Security of the compression function under leakage	19
8	Quantitative bounds for augmented cascades and AMAC	25
9	Comparisons	28
10	Security of the Davies-Meyer construction	30
	References	32
A	Comparing speed of different hash-based MACs	34

1 Introduction

This paper revisits a classical question, namely how can we turn a hash function into a PRF? The canonical answer is HMAC [3], which (1) first applies the keyed hash function to the message and then (2) re-applies, to the result, the hash function keyed with another key. We consider another, even simpler, candidate way, namely to change step (2) to apply a simple *un-keyed* output transform such as truncation. We call this AMAC, for augmented MAC. This paper investigates and establishes provable-security of AMAC, with good bounds, when the hash function is a classical MD-style one like SHA-512.

WHY? We were motivated to determine the security of AMAC by the following. *Usage.* AMAC with SHA-512 is used as a PRF in the Ed25519 signature scheme [7]. (AMAC under a key that is part of the signing key is applied to the hashed message to get coins for a Schnorr-like signature.) Ed25519 is widely deployed, including in SSH, Tor, OpenBSD and dozens of other places [9]. The security of AMAC for this usage was questioned in `cfrg` forum debates on Ed25519 as a proposed standard. Analysis of AMAC is important to assess security of this usage and allow informed choices. *Speed.* AMAC is faster than HMAC, particularly on short messages. See Appendix A. *Context.* Sponge-based PRFs, where truncation is the final step due to its already being so for the hash function, have been proven secure [17, 21, 1, 8, 10]. Our work can be seen as stepping back to ask if truncation works in a similar way for classical MD-style hash functions.

FINDINGS IN A NUTSHELL. Briefly, the message of this paper is the following: (1) First, we are able to prove PRF security of AMAC. (2) Second, AMAC has a quite unique and attractive feature, namely that its multi-user security is essentially as good as its single-user security and in particular superior in some settings to that of competitors. (3) Third, it is technically interesting, its security and analysis intrinsically linked to security of the compression function in the presence of leakage, so that leakage becomes of interest for reasons entirely divorced from side-channel attacks. We now step back to provide some background and discuss our approach and results.

THE BASIC CASCADE. Let $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ represent a compression function taking a c -bit chaining variable and b -bit message block to return a c -bit output. The *basic cascade* of h is the function $h^*: \{0, 1\}^c \times (\{0, 1\}^b)^+ \rightarrow \{0, 1\}^c$ defined by

Basic Cascade $h^*(K, \mathbf{X})$

$Y \leftarrow K$; For $i = 1, \dots, n$ do $Y \leftarrow h(Y, \mathbf{X}[i])$; Return Y

where \mathbf{X} is a vector over $\{0, 1\}^b$ whose length is denoted n and whose i -th component is denoted $\mathbf{X}[i]$. This construct is the heart of MD-style hash functions [12, 22] like MD5, SHA-1, SHA-256 and SHA-512, which are obtained by setting K to a fixed, public value and then applying h^* to the padded message.

Now we want to key h^* to get PRFs. We regard h itself as a PRF on domain $\{0, 1\}^b$, keyed by its c -bit chaining variable. Then h^* is the natural candidate for a PRF on the larger domain $(\{0, 1\}^b)^+$. Problem is, h^* isn't secure as a PRF. This is due to the well-known *extension attack*. If I obtain $Y_1 = h^*(K, X_1)$ for some $X_1 \in \{0, 1\}^b$ of my choice, I can compute $Y_2 = h^*(K, X_1X_2)$ for any $X_2 \in \{0, 1\}^b$ of my choice *without knowing* K , via $Y_2 \leftarrow h(Y_1, X_2)$. This clearly violates PRF security of h^* .

Although h^* is not a PRF, BCK2 [4] show that it is a prefix-free PRF. (A PRF as long as no input on which it is evaluated is a prefix of another. The two inputs X_1, X_1X_2 of the above attack violate this property.) When $b = 1$ and all inputs on which h^* is evaluated are of the same fixed length, the cascade h^* is the GGM construction of a PRF from a PRG [19].

To get a full-fledged PRF, NMAC applies h , under another key, to h^* . The augmented cascade $ACSC = Out \circ h^*$ that we discuss next replaces NMAC’s outer application of a keyed function with a simple un-keyed one.

AUGMENTED CASCADE. The augmented cascade is parameterized by some (keyless) function $Out: \{0, 1\}^c \rightarrow Out.R$ that we call the output transform, and is obtained by simply applying this function to the output of the basic cascade:

Augmented Cascade $(Out \circ h^*)(K, \mathbf{X})$
 $Y \leftarrow h^*(K, \mathbf{X}) ; Z \leftarrow Out(Y) ; \text{Return } Z$

AMAC is obtained from ACSC just as HMAC is obtained from NMAC, namely by putting the key in the input to the hash function rather than directly keying the cascade: $AMAC(K, M) = Out(H(K||M))$. Just as NMAC is the technical core of HMAC, the augmented cascade is the technical core of AMAC, and our analysis will focus in it. We will be able to bridge to AMAC quite simply with the tools we develop.

The ACSC construction was suggested by cryptanalysts with the intuition that “good” choices of Out appear to allow $Out \circ h^*$ to evade the extension attack and thus possibly be a PRF. To understand this, first note that not all choices of Out are good. For example if Out is the identity function then the augmented cascade is the same as the basic one and the attack applies, or if Out is a constant function returning 0^r then $Out \circ h^*$ is obviously not a PRF over range $\{0, 1\}^r$. Cryptanalysts have suggested some specific choices of Out , the most important being (1) truncation, where $Out: \{0, 1\}^c \rightarrow \{0, 1\}^r$ returns, say, the first $r < c$ bits of its input, or (2) the mod function, as in Ed25519, where Out treats its input as an integer and returns the result modulo, say, a public r -bit prime number. Suppose r is sufficiently smaller than c (think $c = 512$ and $r = 256$). An adversary querying X_1 in the PRF game no longer gets back $Y_1 = h^*(K, X_1)$ but rather $Z_1 = Out(Y_1)$, and this does not allow the extension attack to proceed. On this basis, and for the choices of Out just named, the augmented cascade is already seeing extensive usage and is suggested for further usage and standardization.

This raises several questions. First, that $Out \circ h^*$ seems to evade the extension attack does not mean it is a PRF. There may be other attacks. The goal is to get a PRF, not to evade some specific attacks. Moreover we would like a proof that this goal is reached. Second, for which choices of Out does the construction work? We could try to analyze the PRF security of $Out \circ h^*$ in an ad hoc way for the specific choices of Out named above, but it would be more illuminating and useful to be able to establish security in a broad way, for all Out satisfying some conditions. These are the questions our work considers and resolves.

CONNECTION TO LEAKAGE. If we want to prove PRF security of $Out \circ h^*$, a basic question to ask is, under what assumption on the compression function h ? The natural one is that h is itself a PRF, the same assumption as for the proof of NMAC [2, 16]. We observe that this is not enough. Consider an adversary who queries the one-block message X_1 to get back $Z_1 = Out(Y_1)$ and then queries the two-block message X_1X_2 to get back $Z_2 = Out(Y_2)$ where by definition $Y_1 = h^*(K, X_1) = h(K, X_1)$ and $Y_2 = h^*(K, X_1X_2) = h(Y_1, X_2)$. Note that Y_1 is being used as a key in applying h to X_2 . But this key is not entirely unknown to the adversary because the latter knows $Z_1 = Out(Y_1)$. If the application of h with key Y_1 is to provide security, it must be in the face of the fact that some information about this key, namely $Out(Y_1)$, has been “leaked” to the adversary. As a PRF, h must thus be resilient to some leakage on its key, namely that represented by Out viewed as a leakage function.

APPROACH AND QUALITATIVE RESULTS. We first discuss our results at the qualitative level and then later at the (in our view, even more interesting) quantitative level. Theorems 5.2 and 5.3

show that if h is a PRF under Out -leakage then $\text{Out} \circ h^*$ is indistinguishable from the result of applying Out to a random function. (The compression function h being a PRF under Out -leakage means it retains PRF security under key K even if the adversary is given $\text{Out}(K)$. The formal definition is in Section 4.) This result makes no assumptions about Out beyond that implicit in the assumption on h , meaning the result is true for *all* Out , and is in the standard model. As a corollary we establish PRF security of $\text{Out} \circ h^*$ for a large class of output functions Out , namely those that are close to regular. (This means that the distribution of $\text{Out}(Y)$ for random Y is close to the uniform distribution on the range of Out .) In summary we have succeeded in providing conditions on Out, h under which $\text{Out} \circ h^*$ is proven to be PRF. Our conditions are effectively both necessary and sufficient and cover cases proposed for usage and standardization.

The above is a security proof for the augmented cascade $\text{Out} \circ h^*$ under the assumption that the compression function h is resistant to Out leakage. To assess the validity of this assumption, we analyze the security under leakage of an ideal compression function. Theorem 7.2 shows that an ideal compression function is resistant to Out -leakage as long as no range point of Out has too few pre-images. This property is in particular true if Out is close to regular. As a result, in the ideal model, we have a validation of our Out -leakage resilience assumption. Putting this together with the above we have a proof-based validation of the augmented cascade.

MULTI-USER SECURITY. The standard definition of PRF security of a function family F [19] is single user (su), represented by there being a single key K such that the adversary has access to an oracle FN that given x returns either $F(K, x)$ or the result of a random function F on x . But in “real life” there are many users, each with their own key. If we look across the different entities and Internet connections active at any time, the number of users / keys is very large. The more appropriate model is thus a multi-user (mu) one, where, for a parameter u representing the number of users, there are u keys K_1, \dots, K_u . Oracle FN now takes i, x with $1 \leq i \leq u$ and returns either $F(K_i, x)$ or the result of a random function F_i on x . It is in this setting that we should address security.

Multi-user security is typically neglected because it makes no *qualitative* difference: BCK2 [4], who first formalized the notion, also showed by a hybrid argument that the advantage of an adversary relative to u users is not more than u times the advantage of an adversary of comparable resources relative to a single user. Our Lemma 4.1 is a generalization of this result. But this degradation in advantage is quite significant in practice, since u is large, and raises the important question of whether one can do quantitatively better. Clearly one cannot in general, but perhaps one can for specific, special function families F . If so, these function families are preferable in practice. This perspective is reflected in recent work like [23, 26].

These special function families seem quite rare. But we show that the augmented cascade is one of them. In fact we show that mu security gives us a double benefit in this setting, one part coming from the cascade itself and the other from the security of the compression function under leakage, the end result being very good bounds for the mu security of the augmented cascade.

Theorem 5.2 establishes su security of the augmented cascade based not on the su, but on the mu security of the compression function under Out -leakage. The bound is very good, the advantage dropping only by a factor equal to the maximum length of a query. The interesting result is Theorem 5.3, establishing mu security of the augmented cascade under the same assumptions and with essentially the same bounds as Theorem 5.2 establishing its su security. In particular we do not lose a factor of the number of users u in the advantage. This is the first advance.

Now note that the assumption in both of the above-mentioned results is the mu (not su) security of the compression function under Out -leakage. Our final bound will thus depend on this. The second advance is that Theorem 7.2 shows mu security of the compression function under Out -

leakage with bounds almost as good as for su security. This represents an interesting result of independent interest, namely that, under leakage, the mu security of an ideal compression function is almost as good as its su security. This is not true in the absence of leakage. The results are summarized via Fig. 4.

QUANTITATIVE RESULTS. We obtain good quantitative bounds on the mu prf security of the augmented cascade in the ideal compression function model by combining our aforementioned results on the mu prf security under leakage of an ideal compression function with our also aforementioned reduction of the security of the cascade to the security of the compression function under leakage. We illustrate these results for the case where the compression function is of form $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and the output transform Out simply outputs the first r bits of its c -bit input, for $r \leq c$. We consider an attacker making at most q queries to a challenge oracle (that is either the augmented cascade or a random function), each query consisting of at most ℓ b -bit blocks, and q_F queries to the ideal compression function oracle. We show that such an attacker achieves distinguishing advantage at most

$$\frac{\ell^2 q^2 + \ell q q_F}{2^c} + \frac{cr \cdot (\ell^2 q + \ell q_F)}{2^{c-r}}, \quad (1)$$

where we have intentionally omitted constant factors and lower order terms. Note that this bound holds *regardless of the number of users u* . Here c is large, like $c = 512$, so the first term is small. But $c - r$ is smaller, for example $c - r = 256$ with $r = 256$. The crucial merit of the bound of Equation (1) is that the numerator in the second term does not contain quadratic terms like q^2 or $q \cdot q_F$. In practice, q_F and q are the terms we should allow to be large, so this is significant. To illustrate, say for example $\ell = 2^{10}$ (meaning messages are about 128 KBytes if $b = 1024$) and $q_F = 2^{100}$ and $q = 2^{90}$. The bound from Equation (1) is about 2^{-128} , which is very good. But, had the second term been of the form $\ell^2(q_F^2 + q^2)/2^{c-r}$ then the bound would be only 2^{-36} . See Section 8 for more information.

2-TIER CASCADE. We introduce and use an extension of the basic cascade h^* . Our 2-tier cascade is associated to two function families g, h . Under key K , it applies $g(K, \cdot)$ to the first message block to get a sub-key K^* and then applies $h^*(K^*, \cdot)$ to the rest of the message. The corresponding augmented cascade applies Out to the result. Our results about the augmented cascade above are in fact shown for the augmented 2-tier cascade. This generalization has both conceptual and analytical value. We briefly mention two instances. (1) First, we can visualize mu security of $\text{Out} \circ h^*$ as pre-pending the user identity to the message and then applying the 2-tier cascade with first tier a random function. This effectively reduces mu security to su security. With this strategy we prove Theorem 5.3 as a corollary of Theorem 5.2 and avoid a direct analysis of mu security. Beyond providing a modular proof this gives some insight into why the mu security is almost as good as the su security. (2) Second, just as NMAC is the technical core and HMAC the function used (because the latter makes blackbox use of the hash function), in our case the augmented cascade is the technical core but what will be used is AMAC, defined by $\text{AMAC}(K, M) = \text{Out}(H(K, M))$ where H is the hash function derived from compression function $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and K is a k -bit key. For the analysis we note (assuming $k = b$) that this is simply an augmented 2-tier cascade with the first tier being the dual of h , meaning the key and input roles are swapped. We thus directly get an analysis and proof for this case from our above-mentioned results. Obtaining HMAC from NMAC was more work [3, 2] and required assumptions about PRF security of the dual function under related keys.

DAVIES-MEYER. Above we have assessed the PRF security under Out -leakage of the compression function by modeling the latter as ideal (random). But, following CDMP [11], one might say

that the compression functions underlying MD-style hash functions are not un-structured enough to be treated as random because they are built from blockciphers via the Davies-Meyer (DM) construction. To address this we analyze the mu PRF security under Out-leakage of the DM construction in the ideal-cipher model. One’s first thought may be that such an analysis would follow from our analysis for a random compression function and the indistinguishability [20, 11] of DM from a random oracle, but the catch is that DM is *not* indistinguishable from a RO so a direct analysis is needed. The one we give in Section 10 shows mu security with good bounds. Similar analyses can be given for other PGV [25] compression functions.

2 Related work

SPONGES. SHA-3 already internally incorporates a truncation output transform. The construction itself is a sponge. The suggested way to obtain a PRF is to simply key the hash function via its IV, so that the PRF is a keyed, truncated sponge. The security of this construct has been intensively analyzed [17, 21, 1, 8, 10] with Gaži, Pietrzak and Tessaro (GPT) [17] establishing PRF security with tight bounds. Our work can be seen as stepping back to ask whether the same truncation method would work for MD-style hash functions like SHA-512. Right now these older hash functions are much more widely deployed than SHA-3, and current standardization and deployment efforts continue to use them, making the analysis of constructions based on them important with regard to security in practice. The underlying construction in this case is the cascade, which is quite different from the sponge. The results and techniques of GPT [17] do not directly apply but were an important inspiration for our work.

We note that keyed sponges with truncation to an r -bit output from a c -bit state can easily be distinguished from a random function with advantage roughly $q^2/2^{c-r}$ or $qq_F/2^{c-r}$, as shown for example in [17]. The bound of Equation (1) is better, meaning the augmented cascade offers greater security. See Section 9 for more information.

CASCADE. BCK2 [4] show su security of the basic cascade (for prefix-free queries) in two steps. First, they show su security of the basic cascade (for prefix-free queries) assuming not su, but mu security of the compression function. Second, they apply the trivial bound mentioned above to conclude su security of the basic cascade for prefix-free queries assuming su security of the compression function. We follow their approach to establish su security of the augmented cascade, but there are differences as well: They have no output transform while we do, they assume prefix-free queries and we do not, we have leakage and they do not. They neither target nor show mu security of the basic cascade in any form, mu security arising in their work only as an intermediate technical step and only for the compression function, not for the cascade.

CHOP-MD. The chop-MD construction of CDMP [11] is the case of the augmented cascade in which the output transform is truncation. They claim this is indistinguishable from a RO when the compression function is ideal. This implies PRF security but their bound is $O(\ell^2(q + q_F)^2/2^{c-r})$ which as we have seen is significantly weaker than our bound of Equation (1). Also, they have no standard-model proofs or analysis for this construction. In contrast our results in Section 5 establish standard-model security.

NMAC AND HMAC. NMAC takes keys $K_{\text{in}}, K_{\text{out}}$ and input \mathbf{X} to return $\text{h}(K_{\text{out}}, \text{h}^*(K_{\text{in}}, \mathbf{X}) \parallel \text{pad})$ where pad is some $(b - c)$ -bit constant and $b \geq c$. Through a series of intensive analyses, the PRF security of NMAC has been established based only on the assumed PRF security of the compression function h , and with tight bounds [3, 2, 16]. Note that NMAC is not a special case of the augmented cascade because Out is not keyed but the outer application of h in NMAC is keyed.

In the model where the compression function is ideal, one can show bounds for NMAC that are somewhat better than for the augmented cascade. This is not surprising. Indeed, when attacking the augmented cascade, the adversary can learn far more information about the internal states of the hash computation. What is surprising (at least to us) is that the gap is actually quite small. See Section 9 for more information. We stress also that this is in the ideal model. In the standard model, there is no proof that NMAC has the type of good mu prf security we establish for the augmented cascade in Section 5.

AES AND OTHER MACS. Why consider new MACs? Why not just use an AES-based MAC like CMAC? The 128 bit key and block size limits security compared to $c = 512$ for SHA-512. A Schnorr signature takes the result of the PRF modulo a prime; the PRF output must have at least as many bits as the prime, and even more bits for most primes, to avoid the Bleichenbacher attack discussed in [24]. Also in that context a hash function is already being used to hash the message before signing so it is convenient to implement the PRF also with the same hash function. HMAC-SHA-512 will provide the desired security but AMAC has speed advantages, particularly on short messages, as discussed in Appendix A, and is simpler. Finally, the question is in some sense moot since AMAC is already deployed and in widespread use via Ed25519 and we need to understand its security.

LEAKAGE. Leakage-resilience of a PRF studies the PRF security of a function h when the attacker can obtain the result of an *arbitrary* function, called the leakage function, applied to the key [15, 13]. This is motivated by side-channel attacks. We are considering a much more restricted form of leakage where there is just one, very specific leakage function, namely `Out`. This arises naturally, as we have seen, in the PRF security of the augmented cascade. We are not considering side-channel attacks.

3 Notation

If \mathbf{x} is a vector then $|\mathbf{x}|$ denotes its length and $\mathbf{x}[i]$ denotes its i -th coordinate. (For example if $\mathbf{x} = (10, 00, 1)$ then $|\mathbf{x}| = 3$ and $\mathbf{x}[2] = 00$.) We let ε denote the empty vector, which has length 0. If $0 \leq i \leq |\mathbf{x}|$ then we let $\mathbf{x}[1..i] = (\mathbf{x}[1], \dots, \mathbf{x}[i])$, this being ε when $i = 0$. We let S^n denote the set of all length n vectors over the set S . We let S^+ denote the set of all vectors of positive length over the set S and $S^* = S^+ \cup \{\varepsilon\}$ the set of all finite-length vectors over the set S . As special cases, $\{0, 1\}^n$ and $\{0, 1\}^*$ denote vectors whose entries are bits, so that we are identifying strings with binary vectors and the empty string with the empty vector.

For sets A_1, A_2 we let $\llbracket A_1, A_2 \rrbracket$ denote the set of all vectors \mathbf{X} of length $|\mathbf{X}| \geq 1$ such that $\mathbf{X}[1] \in A_1$ and $\mathbf{X}[i] \in A_2$ for $2 \leq i \leq |\mathbf{X}|$.

We let $x \leftarrow_s X$ denote picking an element uniformly at random from a set X and assigning it to x . For infinite sets, it is assumed that a proper measure can be defined on X to make this meaningful. Algorithms may be randomized unless otherwise indicated. Running time is worst case. If A is an algorithm, we let $y \leftarrow A(x_1, \dots; r)$ denote running A with random coins r on inputs x_1, \dots and assigning the output to y . We let $y \leftarrow_s A(x_1, \dots)$ be the result of picking r at random and letting $y \leftarrow A(x_1, \dots; r)$. We let $[A(x_1, \dots)]$ denote the set of all possible outputs of A when invoked with inputs x_1, \dots .

We use the code based game playing framework of [5]. (See Fig. 1 for an example.) By $\Pr[G]$ we denote the probability that game G returns `true`.

For an integer n we let $[1..n] = \{1, \dots, n\}$.

<u>Game $\text{DIST}_{F_0, F_1}(\mathcal{A})$</u> $v \leftarrow 0$ $c \leftarrow_s \{0, 1\}; c' \leftarrow_s \mathcal{A}^{\text{NEW}, \text{FN}}$ Return ($c = c'$) <u>NEW()</u> $v \leftarrow v + 1; F_v \leftarrow_s F_c$ <u>FN(i, x)</u> Return $F_i(x)$	<u>Game $\text{DIST}_{F_0, F_1, \text{Out}}(\mathcal{A})$</u> $v \leftarrow 0$ $c \leftarrow_s \{0, 1\}; c' \leftarrow_s \mathcal{A}^{\text{NEW}, \text{FN}}$ Return ($c = c'$) <u>NEW()</u> $v \leftarrow v + 1; K_v \leftarrow_s F_1.K$ If ($c = 1$) then $F_v \leftarrow F_1(K_v, \cdot)$ else $F_v \leftarrow_s F_0$ Return $\text{Out}(K_v)$ <u>FN(i, x)</u> Return $F_i(x)$
--	--

Figure 1: **Games defining distance metric between function families F_0, F_1 .** In the basic (left) case there is no leakage, while in the extended (right) case there is leakage represented by Out .

4 Function-family distance framework

We will be considering various generalizations and extensions of standard prf security. This includes measuring proximity not just to random functions but to some other family, multi-user security and leakage on the key. We also want to allow an easy later extension to a setting with ideal primitives. To enable all this in a unified way we introduce a general distance metric on function families and then derive notions of interest as special cases.

FUNCTION FAMILIES. A *function family* is a two-argument function $F: F.K \times F.D \rightarrow F.R$ that takes a key K in the key space $F.K$ and an input x in the domain $F.D$ to return an output $y \leftarrow F(K, x)$ in the range $F.R$. We let $f \leftarrow_s F$ be shorthand for $K \leftarrow_s F.K; f \leftarrow F(K, \cdot)$, the operation of picking a function at random from family F .

An example of a function family that is important for us is the compression function underlying a hash function, in which case $F.K = F.R = \{0, 1\}^c$ and $F.D = \{0, 1\}^b$ for integers c, b called the length of the chaining variable and the block length, respectively. Another example is a block cipher. However, families of functions do not have to be efficiently computable or have short keys. For sets D, R the *family* $A: A.K \times D \rightarrow R$ of all functions from D to R is defined simply as follows: let $A.K$ be the set of all functions mapping D to R and let $A(f, x) = f(x)$. (We can fix some representation of f as a key, for example the vector whose i -th component is the value f takes on the i -th input under some ordering of D . But this is not really necessary.) In this case $f \leftarrow_s A$ denotes picking at random a function mapping D to R .

Let $F: F.K \times F.D \rightarrow F.R$ be a function family and let $\text{Out}: F.R \rightarrow \text{Out}.R$ be a function with domain the range of F and range $\text{Out}.R$. Then the composition $\text{Out} \circ F: F.K \times F.D \rightarrow \text{Out}.R$ is the function family defined by $(\text{Out} \circ F)(K, x) = \text{Out}(F(K, x))$. We will use composition in some of our constructions.

BASIC DISTANCE METRIC. We define a general metric of distance between function families that will allow us to obtain other metrics of interest as special cases. Let F_0, F_1 be families of functions such that $F_0.D = F_1.D$. Consider game DIST on the left of Fig. 1 associated to F_0, F_1 and an adversary \mathcal{A} . Via oracle NEW , the adversary can create a new instance F_v drawn from F_c where c is the challenge bit. It can call this oracle multiple times, reflecting a multi-user setting. It can

obtain $F_i(x)$ for any i, x of its choice with the restriction that $1 \leq i \leq v$ (instance i has been initialized) and $x \in F_1.D$. It wins if it guesses the challenge bit c . The advantage of adversary \mathcal{A} is

$$\text{Adv}_{F_0, F_1}^{\text{dist}}(\mathcal{A}) = 2 \Pr[\text{DIST}_{F_0, F_1}(\mathcal{A})] - 1 \quad (2)$$

$$= \Pr[\text{DIST}_{F_0, F_1}(\mathcal{A}) \mid c = 1] - (1 - \Pr[\text{DIST}_{F_0, F_1}(\mathcal{A}) \mid c = 0]) . \quad (3)$$

Equation (2) is the definition, while Equation (3) is a standard alternative formulation that can be shown equal via a conditioning argument. We often use the second in proofs.

Let F be a function family and let A be the family of all functions from $F.D$ to $F.R$. Let $\text{Adv}_F^{\text{prf}}(\mathcal{A}) = \text{Adv}_{F, A}^{\text{dist}}(\mathcal{A})$. This gives a metric of multi-user prf security. The standard (single user) prf metric is obtained by restricting attention to adversaries that make exactly one NEW query.

DISTANCE UNDER LEAKAGE. We extend the framework to allow leakage on the key. Let $\text{Out}: F_1.K \rightarrow \text{Out}.R$ be a function with domain $F_1.K$ and range a set we denote $\text{Out}.R$. Consider game DIST on the right of Fig. 1, now associated not only to F_0, F_1 and an adversary \mathcal{A} but also to Out . Oracle NEW picks a key K_v for F_1 and will return as leakage the result of Out on this key. The instance F_v is either $F_1(K_v, \cdot)$ or a random function from F_0 . Note that the leakage is on a key for a function from F_1 regardless of the challenge bit, meaning even if $c = 0$, we leak on the key K_v drawn from $F_1.K$. The second oracle is as before. The advantage of adversary \mathcal{A} is

$$\text{Adv}_{F_0, F_1, \text{Out}}^{\text{dist}}(\mathcal{A}) = 2 \Pr[\text{DIST}_{F_0, F_1, \text{Out}}(\mathcal{A})] - 1 \quad (4)$$

$$= \Pr[\text{DIST}_{F_0, F_1, \text{Out}}(\mathcal{A}) \mid c = 1] - (1 - \Pr[\text{DIST}_{F_0, F_1, \text{Out}}(\mathcal{A}) \mid c = 0]) . \quad (5)$$

This generalizes the basic metric because $\text{Adv}_{F_0, F_1}^{\text{dist}}(\mathcal{A}) = \text{Adv}_{F_0, F_1, \text{Out}}^{\text{dist}}(\mathcal{A})$ where Out is the function that returns ε on all inputs.

As a special case we get a metric of multi-user prf security under leakage. Let F be a function family and let A be the family of all functions from $F.D$ to $F.R$. Let $\text{Out}: F.K \rightarrow \text{Out}.R$. Let $\text{Adv}_{F, \text{Out}}^{\text{prf}}(\mathcal{A}) = \text{Adv}_{F, A, \text{Out}}^{\text{dist}}(\mathcal{A})$.

NAIVE MU TO SU REDUCTION. Multi-user security for PRFs was first explicitly considered in [4]. They used a hybrid argument to show that the prf advantage of an adversary \mathcal{A} against u users is at most u times the prf advantage of an adversary of comparable resources against a single user. The argument extends to the case where instead of prf advantage we consider distance and where leakage is present. This is summarized in Lemma 4.1 below.

We state this lemma to emphasize that mu security is not qualitatively different from su security, at least in this setting. The question is what is the quantitative difference. The lemma represents the naive bound, which always holds. The interesting element is that for the 2-tier augmented cascade, Theorem 5.3 shows that one can do better: the mu advantage is not a factor u less than the single-user advantage, but about the same. In the proof of the lemma below we specify the adversary for the sake of making the reduction concrete but we omit the standard hybrid argument that establishes that this works.

Lemma 4.1 *Let F_0, F_1 be function families with $F_0.D = F_1.D$ and let $\text{Out}: F_1.K \rightarrow \text{Out}.R$ be an output transform. Let \mathcal{A} be an adversary making at most u queries to its NEW oracle and at most q queries to its FN oracle. The proof specifies an adversary \mathcal{A}_1 making one query to its NEW oracle and at most q queries to its FN oracle such that*

$$\text{Adv}_{F_0, F_1, \text{Out}}^{\text{dist}}(\mathcal{A}) \leq u \cdot \text{Adv}_{F_0, F_1, \text{Out}}^{\text{dist}}(\mathcal{A}_1) . \quad (6)$$

The running time of \mathcal{A}_1 is that of \mathcal{A} plus the time for u computations of F_0 or F_1 .

Proof of Lemma 4.1: Adversary \mathcal{A}_1 runs \mathcal{A} , simulating the latter's NEW, FN oracles via sub-

routines $\text{NEW}^*, \text{FN}^*$, as follows:

<u>Adversary $\mathcal{A}_1^{\text{NEW}, \text{FN}}$</u> $v \leftarrow 0$ $g \leftarrow_{\$} [1..u]$ $c' \leftarrow_{\$} \mathcal{A}^{\text{NEW}^*, \text{FN}^*}$ Return c'	<u>$\text{NEW}^*(\cdot)$</u> $v \leftarrow v + 1 ; K_v \leftarrow_{\$} \mathbf{F}_1.K$ $L \leftarrow \text{Out}(K_v)$ If $(v < g)$ then $F_v \leftarrow \mathbf{F}_1(K_v, \cdot)$ If $(v = g)$ then $L \leftarrow \text{NEW}()$ If $(v > g)$ then $F_v \leftarrow_{\$} \mathbf{F}_0$ Return L	<u>$\text{FN}^*(i, x)$</u> If $(v = g)$ then $y \leftarrow \text{FN}(1, x)$ Else $y \leftarrow F_i(x)$ Return y
--	--	--

We omit the standard analysis establishing Equation (6). \blacksquare

5 The augmented cascade and its analysis

We first present a generalization of the basic cascade construction that we call the 2-tier cascade. We then present the augmented (2-tier) cascade construction and analyze its security.

2-TIER CASCADE CONSTRUCTION. Let \mathcal{K} be a set. Let \mathbf{g}, \mathbf{h} be function families such that $\mathbf{g}: \mathbf{g}.K \times \mathbf{g}.D \rightarrow \mathcal{K}$ and $\mathbf{h}: \mathcal{K} \times \mathbf{h}.D \rightarrow \mathcal{K}$. Thus, outputs of both \mathbf{g} and \mathbf{h} can be used as keys for \mathbf{h} . This is the basis of our 2-tier version of the cascade. This is a function family $\mathbf{CSC}[\mathbf{g}, \mathbf{h}]: \mathbf{g}.K \times \llbracket \mathbf{g}.D, \mathbf{h}.D \rrbracket \rightarrow \mathcal{K}$. That is, a key is one for \mathbf{g} . An input —as per the notation $\llbracket \cdot, \cdot \rrbracket$ defined in Section 3— is a vector \mathbf{X} of length at least one whose first component is in $\mathbf{g}.D$ and the rest of whose components are in $\mathbf{h}.D$. Outputs are in \mathcal{K} . The function itself is defined as follows:

Function $\mathbf{CSC}[\mathbf{g}, \mathbf{h}](K, \mathbf{X})$
 $n \leftarrow |\mathbf{X}| ; Y \leftarrow \mathbf{g}(K, \mathbf{X}[1])$
 For $j = 2, \dots, n$ do $Y \leftarrow \mathbf{h}(Y, \mathbf{X}[j])$
 Return Y

We say that a function family \mathbf{G} is a 2-tier cascade if $\mathbf{G} = \mathbf{CSC}[\mathbf{g}, \mathbf{h}]$ for some \mathbf{g}, \mathbf{h} . If $\mathbf{f}: \mathcal{K} \times \mathbf{f}.D \rightarrow \mathcal{K}$ then its basic cascade is recovered as $\mathbf{CSC}[\mathbf{f}, \mathbf{f}]: \mathcal{K} \times \mathbf{f}.D^+ \rightarrow \mathcal{K}$. We will also denote this function family by \mathbf{f}^* .

Recall that even if $\mathbf{f}: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ is a PRF, \mathbf{f}^* is not a PRF due to the extension attack. It is shown by BCK2 [4] to be a PRF when the adversary is restricted to prefix-free queries. When $b = 1$ and the adversary is restricted to queries of some fixed length ℓ , the cascade \mathbf{f}^* is the GGM construction of a PRF from a PRG [19]. Bernstein [6] considers a generalization of the basic cascade in which the function applied depends on the block index and proves PRF security for any fixed number ℓ of blocks.

Our generalization to the 2-tier cascade has two motivations and corresponding payoffs. First, it will allow us to reduce mu security to su security in a simple, modular and tight way, the idea being that mu security of the basic cascade is su security of the 2-tier one for a certain choice of the 1st tier family. Second, it will allow us to analyze the blackbox AMAC construction in which the cascade is not keyed directly but rather the key is put in the input to the hash function.

THE AUGMENTED CASCADE. With $\mathcal{K}, \mathbf{g}, \mathbf{h}$ as above let $\text{Out}: \mathcal{K} \rightarrow \text{Out}.R$ be a function we call the output transform. The augmented (2-tier) cascade $\mathbf{ACSC}[\mathbf{g}, \mathbf{h}, \text{Out}]: \mathbf{g}.K \times \llbracket \mathbf{g}.D, \mathbf{h}.D \rrbracket \rightarrow \text{Out}.R$ is the composition of Out with $\mathbf{CSC}[\mathbf{g}, \mathbf{h}]$, namely $\mathbf{ACSC}[\mathbf{g}, \mathbf{h}, \text{Out}] = \text{Out} \circ \mathbf{CSC}[\mathbf{g}, \mathbf{h}]$, where composition was defined above. In code:

Function $\mathbf{ACSC}[g, h, \text{Out}](K, \mathbf{X})$

$Y \leftarrow \mathbf{CSC}[g, h](K, \mathbf{X}) ; Z \leftarrow \text{Out}(Y)$

Return Z

We say that a function family G^+ is an augmented (2-tier) cascade if $G^+ = \mathbf{ACSC}[g, h, \text{Out}]$ for some g, h, Out .

The natural goal is that an augmented cascade G^+ be a PRF. This however is clearly not true for all Out . For example Out may be a constant function, or a highly irregular one. Rather than restrict Out at this point we target a general result that would hold for any Out . Namely we aim to show that $\mathbf{ACSC}[g, h, \text{Out}]$ is close under our distance metric to the result of applying Out to a random function. Next we formalize and prove this.

SINGLE-USER SECURITY OF 2-TIER AUGMENTED CASCADE. Given g, h, Out defining the 2-tier augmented cascade $\text{Out} \circ \mathbf{CSC}[g, h]$, we want to upper bound $\text{Adv}_{\text{Out} \circ A, \text{Out} \circ \mathbf{CSC}[g, h]}^{\text{dist}}(\mathcal{A})$ for an adversary \mathcal{A} making one NEW query, where A is the family of all functions with the same domain as $\mathbf{CSC}[g, h]$. We will do this in two steps. First, in Lemma 5.1, we will consider the case that the first tier is a random function, meaning $g = r$ is the family of all functions with the same domain and range as g . Then, in Theorem 5.2, we will use Lemma 5.1 to analyze the general case where g is a PRF. Most interestingly we will later use these single-user results to easily obtain, in Theorem 5.3, bounds for multi-user security that are essentially as good as for single-user security. This showcases a feature of the 2-tier cascade that is rare amongst PRFs. We now proceed to the above-mentioned lemma.

Lemma 5.1 *Let \mathcal{K}, \mathcal{D} be non-empty sets. Let $h: \mathcal{K} \times h.D \rightarrow \mathcal{K}$ be a function family. Let r be the family of all functions with domain \mathcal{D} and range \mathcal{K} . Let $\text{Out}: \mathcal{K} \rightarrow \text{Out.R}$ be an output transform. Let A be the family of all functions with domain $[[\mathcal{D}, h.D]]$ and range \mathcal{K} . Let \mathcal{A} be an adversary making exactly one query to its NEW oracle followed by at most q queries to its FN oracle, the second argument of each of the queries in the latter case being a vector $\mathbf{X} \in [[\mathcal{D}, h.D]]$ with $2 \leq |\mathbf{X}| \leq \ell + 1$. Let adversary \mathcal{A}_h be as in Fig. 2. Then*

$$\text{Adv}_{\text{Out} \circ A, \text{Out} \circ \mathbf{CSC}[r, h]}^{\text{dist}}(\mathcal{A}) \leq \ell \cdot \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_h). \quad (7)$$

Adversary \mathcal{A}_h makes at most q queries to its NEW oracle and at most q queries to its FN oracle. Its running time is that of \mathcal{A} plus the time for $q\ell$ computations of h .

With the first tier being a random function, Lemma 5.1 is bounding the single-user (\mathcal{A} makes one NEW query) distance of the augmented 2-tier cascade to the result of applying Out to a random function under our distance metric. The bound of Equation (7) is in terms of the multi-user security of h as a PRF and grows linearly with one less than the maximum number of blocks in a query.

We note that we could apply Lemma 4.1 to obtain a bound in terms of the single-user PRF security of h , but this is not productive. Instead we will go the other way, later bounding the multi-user security of the 2-tier augmented cascade in terms of the multi-user PRF security of its component functions.

The proof below follows the basic paradigm of the proof of BCK2 [4], which is itself an extension of the classic proof of GGM [19]. However there are several differences: (1) The cascade in BCK2 is single-tier and non-augmented, meaning both the r component and Out are missing (2) BCK2 assume the adversary queries are prefix-free, meaning no query is a prefix of another, an assumption we do not make (3) BCK2 bounds prf security, while we bound the distance.

Proof of Lemma 5.1: Consider the hybrid games and adversaries in Fig. 2. The following chain

<p><u>Game H_s ($0 \leq s \leq \ell$)</u> $b' \leftarrow_s \mathcal{A}^{\text{NEW}^*, \text{FN}^*}$ Return ($b' = 1$)</p> <p><u>NEW[*]()</u> $f \leftarrow_s \mathbf{A}$</p> <p><u>FN[*]($i, \mathbf{X}$)</u> $n \leftarrow \mathbf{X}$ If ($n \leq s$) then $Y \leftarrow f(\mathbf{X})$ Else $Y \leftarrow f(\mathbf{X}[1..s+1])$ For $j = s+2, \dots, n$ do $Y \leftarrow h(Y, \mathbf{X}[j])$ $T_1[\mathbf{X}] \leftarrow Y$; $T_2[\mathbf{X}] \leftarrow \text{Out}(T_1[\mathbf{X}])$ Return $T_2[\mathbf{X}]$</p> <hr/> <p><u>Adversary $\mathcal{A}_h^{\text{NEW}, \text{FN}}$</u> $g \leftarrow_s \{1, \dots, \ell\}$; $b' \leftarrow_s \mathcal{A}_g^{\text{NEW}, \text{FN}}$ Return b'</p>	<p><u>Adversary $\mathcal{A}_g^{\text{NEW}, \text{FN}}$ ($1 \leq g \leq \ell$)</u> $v \leftarrow 0$; $b' \leftarrow_s \mathcal{A}^{\text{NEW}^*, \text{FN}^*}$; Return b'</p> <p><u>NEW[*]()</u></p> <p><u>FN[*](i, \mathbf{X})</u> $n \leftarrow \mathbf{X}$ If ($n \leq g-1$) then If (not $T_1[\mathbf{X}]$) then $T_1[\mathbf{X}] \leftarrow_s \mathcal{K}$; $T_2[\mathbf{X}] \leftarrow \text{Out}(T_1[\mathbf{X}])$ If ($n \geq g$) then If (not $U[\mathbf{X}[1..g]]$) then $v \leftarrow v+1$; $U[\mathbf{X}[1..g]] \leftarrow v$ $T_2[\mathbf{X}[1..g]] \leftarrow \text{NEW}()$ If ($n \geq g+1$) then $T_1[\mathbf{X}[1..g+1]] \leftarrow \text{FN}(U[\mathbf{X}[1..g]], \mathbf{X}[g+1])$ For $j = g+2, \dots, n$ do $T_1[\mathbf{X}[1..j]] \leftarrow h(T_1[\mathbf{X}[1..j-1]], \mathbf{X}[j])$ $T_2[\mathbf{X}] \leftarrow \text{Out}(T_1[\mathbf{X}])$ Return $T_2[\mathbf{X}]$</p>
---	--

Figure 2: **Games and adversaries for proof of Theorem 5.1.**

of equalities establishes Equation (7) and will be justified below:

$$\ell \cdot \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_h) = \sum_{g=1}^{\ell} \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_g) \quad (8)$$

$$= \sum_{g=1}^{\ell} \Pr[\mathbf{H}_{g-1}] - \Pr[\mathbf{H}_g] \quad (9)$$

$$= \Pr[\mathbf{H}_0] - \Pr[\mathbf{H}_\ell] \quad (10)$$

$$= \text{Adv}_{\text{Out} \circ \mathbf{A}, \text{Out} \circ \text{CSC}_{[r, h]}}^{\text{dist}}(\mathcal{A}) \quad (11)$$

Adversary \mathcal{A}_h (bottom left of Fig. 2) picks g at random in the range $1, \dots, \ell$ and runs adversary \mathcal{A}_g (right of Fig. 2) so $\text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_h) = (1/\ell) \cdot \sum_{g=1}^{\ell} \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_g)$, which explains Equation (8). For the rest we begin by trying to picture what is going on.

We imagine a tree of depth $\ell+1$, meaning it has $\ell+2$ levels. The levels are numbered $0, 1, \dots, \ell+1$, with 0 being the root. The root has $|\mathcal{D}|$ children while nodes at levels $1, \dots, \ell$ have $|\mathbf{h}, \mathcal{D}|$ children each. A query \mathbf{X} of \mathcal{A} in game $\text{DIST}_{\text{Out} \circ \mathbf{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A})$ specifies a path in this tree starting at the root and terminating at a node at level $n = |\mathbf{X}|$. Both the path and the final node are viewed as named by \mathbf{X} . To a queried node \mathbf{X} we associate two labels, an internal label $T_1[\mathbf{X}] \in \mathcal{K}$ and an external label $T_2[\mathbf{X}] = \text{Out}(T_1[\mathbf{X}]) \in \text{Out.R}$. The external label is the response to query \mathbf{X} . Since the first component of our 2-tier cascade is the family r of all functions from \mathcal{D} to \mathcal{K} , we can view $\text{DIST}_{\text{Out} \circ \mathbf{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A})$ as picking $T_1[\mathbf{X}[1]]$ at random from \mathcal{K} and then setting $T_1[\mathbf{X}] = h^*(T_1[\mathbf{X}[1]], \mathbf{X}[2..n])$ for all queries \mathbf{X} of \mathcal{A} .

Now we consider the hybrid games H_0, \dots, H_ℓ of Fig. 2. They simulate \mathcal{A} 's NEW, FN oracles via procedures NEW^{*}, FN^{*}, respectively. By assumption \mathcal{A} makes exactly one NEW^{*} query, and this will

have to be its first. In response H_s picks at random a function $f: \llbracket \mathcal{D}, \mathcal{K} \rrbracket \rightarrow \mathcal{K}$. A query FN^* has the form (i, \mathbf{X}) but here i can only equal 1 and is ignored in responding. By assumption $2 \leq |\mathbf{X}| \leq \ell$. The game populates nodes at levels $2, \dots, s$ of the tree with $T_1[\cdot]$ values that are obtained via f and thus are random elements of \mathcal{K} . For a node \mathbf{X} at level $n \geq s + 1$, the $T_1[\mathbf{X}[1..s + 1]]$ value is obtained at random and then further values (if needed, meaning if $n \geq s + 2$) are computed by applying the cascade h^* with key $T_1[\mathbf{X}[1..s + 1]]$ to input $\mathbf{X}[s + 2..n]$.

Consider game H_0 , where $s = 0$. By assumption $n \geq 2$ so we will always be in the case $n \geq s + 1$. In the Else statement, $Y \leftarrow f(\mathbf{X}[1])$ is initialized as a random element of \mathcal{K} . With this Y as the key, h^* is then applied to $\mathbf{X}[2..n]$ to get $T_1[\mathbf{X}]$. This means H_0 exactly mimics the $c = 1$ case of game $\text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A})$, so that

$$\Pr[H_0] = \Pr[\text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A}) \mid c = 1] . \quad (12)$$

At the other extreme, consider game H_ℓ , where $s = \ell$. By assumption $n \leq \ell + 1$, yielding two cases. If $n \leq \ell$ we are in the $n \leq s$ case and the game, via f , the assigns $T_1[\mathbf{X}]$ a random value. If $n = \ell + 1$ we are in the $n \geq s + 1$ case, but the For loop does nothing so $T_1[\mathbf{X}]$ is again random. This means H_ℓ mimics the $c = 0$ case of game $\text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A})$, except returning `true` exactly when the latter returns `false`. Thus

$$\Pr[H_\ell] = 1 - \Pr[\text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[r, h], \text{Out}}}(\mathcal{A}) \mid c = 0] . \quad (13)$$

We will justify Equation (9) in a bit but we can now dispense with the rest of the chain. Equation (10) is obvious because the sum “telescopes.” Equation (11) follows from Equations (12) and (13) and the formulation of dist advantage of Equation (5).

It remains to justify Equation (9), for which we consider the adversaries $\mathcal{A}_1, \dots, \mathcal{A}_\ell$ on the right side of Fig. 2. Adversary \mathcal{A}_g is playing the PRF, formally game $\text{DIST}_{\mathcal{B}, h}$ on the left of Fig. 1 in our notation, with \mathcal{B} the family of all functions from $h.D$ to \mathcal{K} . It thus has oracles `NEW`, `FN`. It will make crucial use of the assumed multi-user security of h , meaning its ability to query `NEW` many times, keeping track in variable u of the number of instances it creates. It simulates the oracles of \mathcal{A} of the same names via procedures `NEW*`, `FN*`, sampling functions lazily rather than directly as in the games. Arrays T_1, T_2, U are assumed initially to be everywhere \perp and get populated as the adversary assigns values to entries. A test of the form “If (not $T_1[\mathbf{X}]$) ... ” returns `true` if $T_1[\mathbf{X}] = \perp$, meaning has not yet been initialized. In response to the (single) `NEW*` query of \mathcal{A} , adversary \mathcal{A}_g does nothing. Following that, its strategy is to have the $T_1[\cdot]$ values of level g nodes populated, not explicitly, but implicitly by the keys in game $\text{DIST}_{\mathcal{B}, h}$ created by the adversary’s own `NEW` queries, using array U to keep track of the user index associated to a node. $T_1[\cdot]$ values for nodes at levels $1, \dots, g - 1$ are random. At level $g + 1$, the $T_1[\cdot]$ values are obtained via the adversary’s `FN` oracle, and from then on via direct application of the cascade h^* . One crucial point is that, if \mathcal{A}_g does not know the $T_1[\cdot]$ values at level g , how does it respond to a length g query \mathbf{X} with the right $T_2[\cdot]$ value? This is where the leakage enters, the response being the leakage provided by the `NEW` oracle. The result is that for every $g \in \{1, \dots, \ell\}$ we have

$$\Pr[\text{DIST}_{\mathcal{B}, h}(\mathcal{A}_g) \mid c = 1] = \Pr[H_{g-1}] \quad (14)$$

$$1 - \Pr[\text{DIST}_{\mathcal{B}, h}(\mathcal{A}_g) \mid c = 0] = \Pr[H_g] , \quad (15)$$

where c is the challenge bit in game $\text{DIST}_{\mathcal{B}, h}$. Thus

$$\begin{aligned} \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_g) &= \Pr[\text{DIST}_{\mathcal{B}, h}(\mathcal{A}_g) \mid c = 1] - (1 - \Pr[\text{DIST}_{\mathcal{B}, h}(\mathcal{A}_g) \mid c = 0]) \\ &= \Pr[H_{g-1}] - \Pr[H_g] . \end{aligned} \quad (16)$$

This justifies Equation (9). \blacksquare

We now extend the above to the case where the first tier g of the 2-tier cascade is a PRF rather than a random function. We will exploit PRF security of g to reduce this to the prior case.

Theorem 5.2 *Let \mathcal{K} be a non-empty set. Let $g: g.K \times g.D \rightarrow \mathcal{K}$ and $h: \mathcal{K} \times h.D \rightarrow \mathcal{K}$ be function families. Let $\text{Out}: \mathcal{K} \rightarrow \text{Out.R}$ be an output transform. Let \mathcal{A} be the family of all functions with domain $[[g.D, h.D]]$ and range \mathcal{K} . Let \mathcal{A} be an adversary making exactly one query to its NEW oracle followed by at most q queries to its FN oracle, the second argument of each of the queries in the latter case being a vector $\mathbf{X} \in [[g.D, h.D]]$ with $2 \leq |\mathbf{X}| \leq \ell + 1$. The proof shows how to construct adversaries $\mathcal{A}_h, \mathcal{A}_g$ such that*

$$\text{Adv}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[g,h]}}^{\text{dist}}(\mathcal{A}) \leq \ell \cdot \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_h) + 2 \text{Adv}_g^{\text{prf}}(\mathcal{A}_g). \quad (17)$$

Adversary \mathcal{A}_h makes at most q queries to its NEW oracle and at most q queries to its FN oracle. Adversary \mathcal{A}_g makes one query to its NEW oracle and at most q queries to its FN oracle. The running time of both constructed adversaries is about that of \mathcal{A} plus the time for $q\ell$ computations of h .

Proof of Theorem 5.2: Let $\mathcal{D} = g.D$. Let r be the family of all functions with domain \mathcal{D} and range \mathcal{K} . As shorthand for the relevant games, let

$$G_0 = \text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[g,h]}}(\mathcal{A}) \quad \text{and} \quad G_1 = \text{DIST}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[r,h]}}(\mathcal{A}).$$

Then

$$\begin{aligned} \text{Adv}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[g,h]}}^{\text{dist}}(\mathcal{A}) &= 2 \Pr[G_0] - 1 \\ &= 2 (\Pr[G_0] - \Pr[G_1] + \Pr[G_1]) - 1 \\ &= 2 (\Pr[G_0] - \Pr[G_1]) + 2 \Pr[G_1] - 1. \end{aligned} \quad (18)$$

Lemma 5.1 provides adversary \mathcal{A}_h such that

$$2 \Pr[G_1] - 1 \leq \ell \cdot \text{Adv}_{h, \text{Out}}^{\text{prf}}(\mathcal{A}_h). \quad (19)$$

Let adversary \mathcal{A}_g be as follows:

<p style="margin: 0;">Adversary $\mathcal{A}_g^{\text{NEW, FN}}$</p> <p style="margin: 0;">$c \leftarrow_{\\$} \{0, 1\}; c' \leftarrow_{\\$} \mathcal{A}^{\text{NEW, FN}*}$</p> <p style="margin: 0;">If $(c = c')$ then return 1 else return 0</p>	<p style="margin: 0;">$\text{FN}^*(i, \mathbf{X})$</p> <p style="margin: 0;">$Y \leftarrow \text{FN}(1, \mathbf{X}[1])$</p> <p style="margin: 0;">For $j = 2, \dots, \mathbf{X}$ do $Y \leftarrow h(Y, \mathbf{X}[j])$</p> <p style="margin: 0;">$Z \leftarrow \text{Out}(Y)$; Return Z</p>
--	---

Adversary \mathcal{A}_g responds to the single NEW query that \mathcal{A} makes directly via its own NEW oracle. It responds to FN queries of \mathcal{A} via procedure FN^* . The latter applies \mathcal{A}_g 's own FN oracle to the first component of \mathbf{X} to get a key Y , and then applies h^* with key Y to the rest of \mathbf{X} , finally applying the output function to get the value returned. Then we have

$$\begin{aligned} \text{Adv}_g^{\text{prf}}(\mathcal{A}_g) &= \Pr[\text{DIST}_{r,g}(\mathcal{A}_g) | c = 1] - (1 - \Pr[\text{DIST}_{r,g}(\mathcal{A}_g) | c = 0]) \\ &= \Pr[G_0] - \Pr[G_1] \end{aligned} \quad (20)$$

Putting together Equations (18), (19) and (20) yields Equation (17). \blacksquare

MULTI-USER SECURITY OF 2-TIER AUGMENTED CASCADE. We now want to assess the multi-user security of a 2-tier augmented cascade. This means we want to bound $\text{Adv}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[g,h]}}^{\text{dist}}(\mathcal{A})$

with everything as in Theorem 5.2 above except that \mathcal{A} can now make any number u of NEW queries rather than just one. We could do this easily by applying Lemma 4.1 to Theorem 5.2, resulting in a bound that is u times the bound of Equation (17). We consider Theorem 5.3 below the most interesting result of this section. It says one can do much better, and in fact the bound for the multi-user case is not much different from that for the single-user case.

Theorem 5.3 *Let \mathcal{K} be a non-empty set. Let $\mathbf{g}: \mathbf{g.K} \times \mathbf{g.D} \rightarrow \mathcal{K}$ and $\mathbf{h}: \mathcal{K} \times \mathbf{h.D} \rightarrow \mathcal{K}$ be function families. Let $\text{Out}: \mathcal{K} \rightarrow \text{Out.R}$ be an output transform. Let \mathcal{A} be the family of all functions with domain $\llbracket \mathbf{g.D}, \mathbf{h.D} \rrbracket$ and range \mathcal{K} . Let \mathcal{A} be an adversary making at most u queries to its NEW oracle and at most q queries to its FN oracle, the second argument of each of the queries in the latter case being a vector $\mathbf{X} \in \llbracket \mathbf{g.D}, \mathbf{h.D} \rrbracket$ with $2 \leq |\mathbf{X}| \leq \ell + 1$. The proof shows how to construct adversaries $\mathcal{A}_h, \mathcal{A}_g$ such that*

$$\text{Adv}_{\text{Out} \circ \mathcal{A}, \text{Out} \circ \text{CSC}_{[\mathbf{g}, \mathbf{h}]}}^{\text{dist}}(\mathcal{A}) \leq \ell \cdot \text{Adv}_{\mathbf{h}, \text{Out}}^{\text{prf}}(\mathcal{A}_h) + 2 \text{Adv}_{\mathbf{g}}^{\text{prf}}(\mathcal{A}_g). \quad (21)$$

Adversary \mathcal{A}_h makes at most q queries to its NEW oracle and at most q queries to its FN oracle. Adversary \mathcal{A}_g makes u queries to its NEW oracle and at most q queries to its FN oracle. The running time of both constructed adversaries is about that of \mathcal{A} plus the time for $q\ell$ computations of \mathbf{h} .

A comparison of Theorems 5.2 and 5.3 shows that the bound of Equation (21) is the same as that of Equation (17). So where are we paying for u now not being one? It is reflected only in the resources of adversary \mathcal{A}_g , the latter in Theorem 5.3 making u queries to its NEW oracle rather than just one in Theorem 5.2.

The proof below showcases one of the advantages of the 2-tier cascade over the basic single-tier one. Namely, by appropriate choice of instantiation of the first tier, we can reduce multi-user security to single-user security in a modular way. In this way we avoid re-entering the proofs above. Indeed, the ability to do this is one of the main reasons we introduced the 2-tier cascade.

Proof of Theorem 5.3: Let $\mathcal{D} = [1..u]$. Let $\bar{\mathbf{r}}$ be the family of all functions with domain \mathcal{D} and range $\mathbf{g.K}$. Let function family $\bar{\mathbf{g}}: \bar{\mathbf{r}}.\mathbf{K} \times (\mathcal{D} \times \mathbf{g.D}) \rightarrow \mathcal{K}$ be defined by $\bar{\mathbf{g}}(f, (i, x)) = \mathbf{g}(f(i), x)$. Let \mathcal{B} be the family of all functions with domain $\llbracket \mathcal{D} \times \mathbf{g.D}, \mathbf{h.D} \rrbracket$ and range \mathcal{K} . The main observation is as follows. Suppose $i \in \mathcal{D}$ and $\mathbf{X} \in \llbracket \mathbf{g.D}, \mathbf{h.D} \rrbracket$. Let $\mathbf{Y} \in \llbracket \mathcal{D} \times \mathbf{g.D}, \mathbf{h.D} \rrbracket$ be defined by $\mathbf{Y}[1] = (i, \mathbf{X}[1])$ and $\mathbf{Y}[j] = \mathbf{X}[j]$ for $2 \leq j \leq |\mathbf{X}|$. Let $f: \mathcal{D} \rightarrow \mathbf{g.K}$ be a key for $\bar{\mathbf{g}}$. Then $f(i) \in \mathbf{g.K}$ is a key for \mathbf{g} , and

$$\text{CSC}_{[\bar{\mathbf{g}}, \mathbf{h}]}(f, \mathbf{Y}) = \text{CSC}_{[\mathbf{g}, \mathbf{h}]}(f(i), \mathbf{X}). \quad (22)$$

Think of $f(i)$ as the key for instance i . Then Equation (22) allows us to obtain values of $\text{CSC}_{[\mathbf{g}, \mathbf{h}]}$ for different instances $i \in \mathcal{D}$ via values of $\text{CSC}_{[\bar{\mathbf{g}}, \mathbf{h}]}$ on a single instance with key f . This will allow us to reduce the multi-user security of $\text{CSC}_{[\mathbf{g}, \mathbf{h}]}$ to the single-user security of $\text{CSC}_{[\bar{\mathbf{g}}, \mathbf{h}]}$. Theorem 5.2 will allow us to measure the latter in terms of the prf security of \mathbf{h} under leakage and the (plain) prf security of $\bar{\mathbf{g}}$. The final step will be to measure the prf security of $\bar{\mathbf{g}}$ in terms of that of \mathbf{g} .

Proceeding to the details, let adversary \mathcal{B} be as follows:

$$\begin{array}{l|l} \text{Adversary } \mathcal{B}^{\text{NEW}, \text{FN}} & \\ \hline \text{NEW}() & \text{FN}^*(i, \mathbf{X}) \\ b' \leftarrow_s \mathcal{A}^{\text{NEW}^*, \text{FN}^*}; \text{ Return } b' & \mathbf{Y}[1] \leftarrow (i, \mathbf{X}[1]) \\ & \text{For } j = 2, \dots, |\mathbf{X}| \text{ do } \mathbf{Y}[j] \leftarrow \mathbf{X}[j] \\ \text{NEW}^*() & Z \leftarrow \text{FN}(1, \mathbf{Y}); \text{ Return } Z \end{array}$$

Then we have

$$\text{Adv}_{\text{Out} \circ \mathbf{A}, \text{Out} \circ \mathbf{CSC}[\mathbf{g}, \mathbf{h}]}^{\text{dist}}(\mathcal{A}) = \text{Adv}_{\text{Out} \circ \mathbf{B}, \text{Out} \circ \mathbf{CSC}[\bar{\mathbf{g}}, \mathbf{h}]}^{\text{dist}}(\mathcal{B}) \quad (23)$$

$$\leq \ell \cdot \text{Adv}_{\mathbf{h}, \text{Out}}^{\text{prf}}(\mathcal{A}_{\mathbf{h}}) + 2 \text{Adv}_{\bar{\mathbf{g}}}^{\text{prf}}(\mathcal{A}_{\bar{\mathbf{g}}}) \quad (24)$$

Adversary \mathcal{B} is allowed only one NEW query, and begins by making it so as to initialize instance 1 in its game. It answers queries of \mathcal{A} to its NEW oracle via procedure NEW*. Adversary \mathcal{A} can make up to u queries to NEW*, but, as the absence of code for NEW* indicates, this procedure does nothing, meaning no action is taken when \mathcal{A} makes a NEW* query. When \mathcal{A} queries its FN oracle, \mathcal{B} answers via procedure FN*. The query consists of an instance index i with $1 \leq i \leq u$ and a vector \mathbf{X} . Adversary \mathcal{B} creates \mathbf{Y} from \mathbf{X} as described above. Namely it modifies the first component of \mathbf{X} to pre-pend i , so that $\mathbf{Y}[1] \in \mathcal{D} \times \mathbf{g.D}$ is in the domain of $\bar{\mathbf{g}}$. It leaves the rest of the components unchanged, and then calls its own FN oracle on vector $\mathbf{Y} \in [\mathcal{D} \times \mathbf{g.D}, \mathbf{h.D}]$. The instance used is 1, regardless of i , since \mathcal{B} has only one instance active. The result Z of FN is returned to \mathcal{A} as the answer to its query. Equation (23) is now justified by Equation (22), thinking of $f(i)$ as the key K_i chosen in game $\text{DIST}_{\text{Out} \circ \mathbf{A}, \text{Out} \circ \mathbf{CSC}[\mathbf{g}, \mathbf{h}]}(\mathcal{A})$ where f is the (single) key chosen in game $\text{DIST}_{\text{Out} \circ \mathbf{B}, \text{Out} \circ \mathbf{CSC}[\bar{\mathbf{g}}, \mathbf{h}]}(\mathcal{B})$. Theorem 5.2 applied to $\bar{\mathbf{g}}, \mathbf{h}$ and adversary \mathcal{B} provides the adversaries $\mathcal{A}_{\mathbf{h}}, \mathcal{A}_{\bar{\mathbf{g}}}$ of Equation (24).

Now consider adversary $\mathcal{A}_{\bar{\mathbf{g}}}$ defined as follows:

<u>Adversary $\mathcal{A}_{\bar{\mathbf{g}}}^{\text{NEW}, \text{FN}}$</u> For $i = 1, \dots, u$ do NEW() $b' \leftarrow_{\$} \mathcal{A}_{\bar{\mathbf{g}}}^{\text{NEW}^*, \text{FN}^*}$; Return b' <u>NEW*()</u>	<u>FN*(j, X)</u> $(i, x) \leftarrow X$; $Y \leftarrow \text{FN}(i, x)$ Return Y
--	---

Adversary $\mathcal{A}_{\bar{\mathbf{g}}}$ begins by calling its NEW oracle u times to initialize u instances. It then runs $\mathcal{A}_{\bar{\mathbf{g}}}$, answering the latter's oracle queries via procedures NEW*, FN*. By Theorem 5.2 we know that $\mathcal{A}_{\bar{\mathbf{g}}}$ makes only one NEW* query. In response the procedure NEW* above does nothing. When $\mathcal{A}_{\bar{\mathbf{g}}}$ makes query j, X to FN* we know that $j = 1$ and $X \in \mathcal{D} \times \mathbf{g.D}$. Procedure FN* parses X as (i, x) . It then invokes its own FN oracle with instance i and input x and returns the result Y to $\mathcal{A}_{\bar{\mathbf{g}}}$. We have

$$\text{Adv}_{\bar{\mathbf{g}}}^{\text{prf}}(\mathcal{A}_{\bar{\mathbf{g}}}) = \text{Adv}_{\bar{\mathbf{g}}}^{\text{prf}}(\mathcal{A}_{\bar{\mathbf{g}}}) . \quad (25)$$

Equations (24) and (25) imply Equation (21). \blacksquare

One might ask why prove Theorem 5.3 for a 2-tier augmented cascade $\text{Out} \circ \mathbf{CSC}[\mathbf{g}, \mathbf{h}]$ instead of a single tier one $\text{Out} \circ \mathbf{CSC}[\mathbf{h}, \mathbf{h}]$. Isn't the latter the one of ultimate interest in usage? We establish a more general result in Theorem 5.3 because it allows us to analyze AMAC itself by setting \mathbf{g} to the dual of \mathbf{h} [2], and also for consistency with Theorem 5.2.

6 Framework for ideal-model cryptography

In Section 5 we reduced the (mu) security of the augmented cascade tightly to the assumed mu prf security of the compression function under leakage. To complete the story, we will, in Section 7, bound the mu prf security of an ideal compression function under leakage and thence obtain concrete bounds for the mu security of the augmented cascade in the same model. Additionally, we will consider the same questions when the compression function is not directly ideal but obtained via

<p><u>Game $\text{PRF}_{\mathbf{F},\mathbf{P}}(\mathcal{A})$</u></p> <p>$v \leftarrow 0$ $c \leftarrow_{\\$} \{0, 1\}$; $\mathbf{P} \leftarrow_{\\$} \mathbf{P}$; $c' \leftarrow_{\\$} \mathcal{A}^{\text{NEW},\text{FN},\text{PRIM}}$ Return ($c = c'$)</p> <p><u>NEW()</u></p> <p>$v \leftarrow v + 1$ If ($c = 1$) then $F_v \leftarrow_{\\$} \mathbf{F}^{\text{PRIM}}$ Else $F_v \leftarrow_{\\$} \mathbf{A}$</p> <p><u>FN($i, x$)</u></p> <p>Return $F_i(x)$</p> <p><u>PRIM(x)</u></p> <p>$y \leftarrow \mathbf{P}(x)$; Return y</p>	<p><u>Game $\text{PRF}_{\mathbf{F},\text{Out},\mathbf{P}}(\mathcal{A})$</u></p> <p>$v \leftarrow 0$ $c \leftarrow_{\\$} \{0, 1\}$; $\mathbf{P} \leftarrow_{\\$} \mathbf{P}$; $c' \leftarrow_{\\$} \mathcal{A}^{\text{NEW},\text{FN},\text{PRIM}}$ Return ($c = c'$)</p> <p><u>NEW()</u></p> <p>$v \leftarrow v + 1$; $K_v \leftarrow_{\\$} \mathbf{F.K}$ If ($c = 1$) then $F_v \leftarrow \mathbf{F}^{\text{PRIM}}(K_v, \cdot)$ Else $F_v \leftarrow_{\\$} \mathbf{A}$ Return $\text{Out}(K_v)$</p> <p><u>FN(i, x)</u></p> <p>Return $F_i(x)$</p> <p><u>PRIM(x)</u></p> <p>$y \leftarrow \mathbf{P}(x)$; Return y</p>
--	--

Figure 3: **Games defining prf security of function family \mathbf{F} in the presence of an ideal primitive \mathbf{P} .** In the basic (left) case there is no leakage, while in the extended (right) case there is leakage represented by Out .

the Davies-Meyer transform on an ideal blockcipher, reflecting the design in popular hash functions. If we gave separate, ad hoc definitions for all these different constructions in different ideal models for different goals, it would be a lot of definitions. Accordingly we introduce a general definition of an ideal primitive (that may be of independent interest) and give a general definition of PRF security of a function family with access to an instance of an ideal primitive, both for the basic setting and the setting with leakage. A reader interested in our results on the mu prf security of ideal primitives can jump ahead to Section 7 and refer back here as necessary.

IDEALIZED CRYPTOGRAPHY. We define an *ideal primitive* to simply be a function family \mathbf{P} : $\mathbf{P.K} \times \mathbf{P.D} \rightarrow \mathbf{P.R}$. Below we will provide some examples but first let us show how to lift security notions to idealized models using this definition by considering the cases of interest to us, namely PRFs and PRFs under leakage.

An *oracle function family* \mathbf{F} specifies for each function \mathbf{P} in its *oracle space* $\mathbf{F.O}$ a function family $\mathbf{F}^{\mathbf{P}}$: $\mathbf{F.K} \times \mathbf{F.D} \rightarrow \mathbf{F.R}$. We say \mathbf{F} and ideal primitive \mathbf{P} are *compatible* if $\{\mathbf{P}(\mathbf{K}\mathbf{K}, \cdot) : \mathbf{K}\mathbf{K} \in \mathbf{P.K}\} \subseteq \mathbf{F.O}$, meaning instances of \mathbf{P} are legitimate oracles for \mathbf{F} . These represent constructs whose security we want to measure in an idealized model represented by \mathbf{P} .

We associate to \mathbf{F}, \mathbf{P} and adversary \mathcal{A} the game PRF in the left of Fig. 3. In this game, \mathbf{A} is the family of all functions with domain $\mathbf{F.D}$ and range $\mathbf{F.R}$. The game begins by picking an instance \mathbf{P} : $\mathbf{P.D} \rightarrow \mathbf{P.R}$ of \mathbf{P} at random. The function \mathbf{P} is provided as oracle to \mathbf{F} and to \mathcal{A} via procedure PRIM . The game is in the multi-user setting, and when $c = 1$ it selects a new instance F_v at random from the function family $\mathbf{F}^{\mathbf{P}}$. Otherwise it selects F_v to be a random function from $\mathbf{F.D}$ to $\mathbf{F.R}$. As usual a query i, x to FN must satisfy $1 \leq i \leq v$ and $x \in \mathbf{F.D}$. A query to PRIM must be in the set $\mathbf{P.D}$. We let $\text{Adv}_{\mathbf{F},\mathbf{P}}^{\text{prf}}(\mathcal{A}) = 2 \Pr[\text{PRF}_{\mathbf{F},\mathbf{P}}(\mathcal{A})] - 1$ be the advantage of \mathcal{A} .

We now extend this to allow leakage on the key. Let Out : $\mathbf{F.K} \rightarrow \text{Out.R}$ be a function with domain $\mathbf{F.K}$ and range Out.R . Game PRF on the right of Fig. 3 is now associated not only to \mathbf{F}, \mathbf{P} and an adversary \mathcal{A} but also to Out . The advantage of \mathcal{A} is $\text{Adv}_{\mathbf{F},\text{Out},\mathbf{P}}^{\text{prf}}(\mathcal{A}) = 2 \Pr[\text{PRF}_{\mathbf{F},\text{Out},\mathbf{P}}(\mathcal{A})] - 1$.

CAPTURING PARTICULAR IDEAL MODELS. The above framework allows us to capture the random oracle model, ideal cipher model and many others as different choices of the ideal primitive \mathbf{P} . Not all of these are relevant to our paper but we discuss them to illustrate how the framework captures known settings.

Let \mathcal{Y} be a non-empty set. Let $\mathbf{P.K}$ be the set of all functions $\mathbf{P}: \{0, 1\}^* \rightarrow \mathcal{Y}$. (Each function is represented in some canonical way, in this case for example as a vector over \mathcal{Y} of infinite length.) Let $\mathbf{P}: \mathbf{P.K} \times \{0, 1\}^* \rightarrow \mathcal{Y}$ be defined by $\mathbf{P}(\mathbf{P}, x) = \mathbf{P}(x)$. Then $\mathbf{P} \leftarrow^s \mathbf{P}$ is a random oracle with domain $\{0, 1\}^*$ and range \mathcal{Y} . In this case, an oracle function family compatible with \mathbf{P} is simply a function family in the random oracle model, and its prf security in the random oracle model is measured by $\text{Adv}_{\mathbf{F}, \mathbf{P}}^{\text{prf}}(\mathcal{A})$.

Similarly let $\mathbf{P.K}$ be the set of all functions $\mathbf{P}: \{0, 1\}^* \times \mathbb{N} \rightarrow \{0, 1\}^*$ with the property that $|\mathbf{P}(x, l)| = l$ for all $(x, l) \in \{0, 1\}^* \times \mathbb{N}$. Let $\mathbf{P}: \mathbf{P.K} \times (\{0, 1\}^* \times \mathbb{N}) \rightarrow \{0, 1\}^*$ be defined by $\mathbf{P}(\mathbf{P}, (x, l)) = \mathbf{P}(x, l)$. Then $\mathbf{P} \leftarrow^s \mathbf{P}$ is a variable output length random oracle with domain $\{0, 1\}^*$ and range $\{0, 1\}^*$.

Let \mathcal{D} be a non-empty set. To capture the single random permutation model, let $\mathbf{P.K}$ be the set of all permutations $\pi: \mathcal{D} \rightarrow \mathcal{D}$. Let $\mathbf{P.D} = \mathcal{D} \times \{+, -\}$. Let $\mathbf{P.R} = \mathcal{D}$. Define $\mathbf{P}(\pi, (x, +)) = \pi(x)$ and $\mathbf{P}(\pi, (y, -)) = \pi^{-1}(y)$ for all $\pi \in \mathbf{P.K}$ and all $x, y \in \mathcal{D}$. An oracle for an instance $\mathbf{P} = \mathbf{P}(\pi, \cdot)$ of \mathbf{P} thus allows evaluation of both π and π^{-1} on inputs of the caller's choice.

Finally we show how to capture the ideal cipher model. If \mathcal{K}, \mathcal{D} are non-empty sets, a function family $E: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{D}$ is a blockcipher if $E(K, \cdot)$ is a permutation on \mathcal{D} for every $K \in \mathcal{K}$, in which case $E^{-1}: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{D}$ denotes the blockcipher in which $E^{-1}(K, \cdot)$ is the inverse of the permutation $E(K, \cdot)$ for all $K \in \mathcal{K}$. Let $\mathbf{P.K}$ be the set of all block ciphers $E: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{D}$. Let $\mathbf{P.D} = \mathcal{K} \times \mathcal{D} \times \{+, -\}$. Let $\mathbf{P.R} = \mathcal{D}$. Define $\mathbf{P}(E, (K, X, +)) = E(K, X)$ and $\mathbf{P}(E, (K, Y, -)) = E^{-1}(K, Y)$ for all $E \in \mathbf{P.K}$ and all $X, Y \in \mathcal{D}$. An oracle for an instance $\mathbf{P} = \mathbf{P}(E, \cdot)$ of \mathbf{P} thus allows evaluation of both E and E^{-1} on inputs of the caller's choice.

7 Security of the compression function under leakage

In Section 5 we reduced the (multi-user) security of the augmented cascade tightly to the assumed multi-user prf security of the compression function under leakage. To complete the story, we now study (bound) the multi-user prf security of the compression function under leakage. This will be done assuming the compression function is ideal. Combining these results with those of Section 5 we will get concrete bounds for the security of the augmented cascade for use in applications, discussed in Section 8.

In the (leak-free) multi-user setting, it is well known that prf security of a compression function decreases linearly in the number of users. We will show that this is an extreme case, and as the amount of leakage increases, the multi-user prf security degrades far more gracefully in the number of users (Theorem 7.2). This (perhaps counterintuitive) phenomenon will turn out to be essential to obtain good bounds on augmented cascades. We begin below with an informal overview of the bounds and why this phenomenon occurs.

OVERVIEW OF BOUNDS. The setting of an ideal compression function mapping $\mathcal{K} \times \mathcal{X} \rightarrow \mathcal{D}$ is formally captured, in the framework of Section 6, by the ideal primitive $\mathbf{F}: \mathbf{F.K} \times (\mathcal{K} \times \mathcal{X}) \rightarrow \mathcal{K}$ defined as follows. Let $\mathbf{F.K}$ be the set of all functions mapping $\mathcal{K} \times \mathcal{X} \rightarrow \mathcal{K}$ and let $\mathbf{F}(f, (K, X)) = f(K, X)$. Now, the construction we are interested in is the simplest possible, namely the compression function itself. Formally, again as per Section 6, this means we consider the oracle function family CF whose oracle space CF.O consists of all functions $f: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{K}$, and with $\text{CF}^f = f$.

For this overview we let $\mathcal{K} = \{0, 1\}^c$. We contrast the prf security of an ideal compression

	$\text{Adv}_{\text{CF},\mathbf{F}}^{\text{prf}}(\mathcal{B})$	$\text{Adv}_{\text{CF},\text{Out},\mathbf{F}}^{\text{prf}}(\mathcal{B})$
su	$\frac{q_{\mathbf{F}}}{2^c}$	$\frac{q_{\mathbf{F}}}{2^{c-r}}$
mu , trivial	$\frac{u(q + q_{\mathbf{F}})}{2^c}$	$\frac{u(q + q_{\mathbf{F}})}{2^{c-r}}$
mu , dedicated	$\frac{u^2 + 2uq_{\mathbf{F}}}{2^{c+1}}$	$\frac{u^2 + 2uq_{\mathbf{F}} + 1}{2^c} + \frac{3crq_{\mathbf{F}}}{2^{c-r}}$

Figure 4: **Upper bounds on prf advantage of an adversary \mathcal{B} attacking an ideal compression function mapping $\{0, 1\}^c \times \mathcal{X}$ to $\{0, 1\}^c$. Left:** Basic case, without leakage. **Right:** With leakage Out being the truncation function that returns the first $r \leq c$ bits of its output. **First row:** Single user security, $q_{\mathbf{F}}$ is the number of queries to the ideal compression function. **Second row:** Multi-user security as obtained trivially by applying Lemma 4.1 to the su bound, u is the number of users. **Third row:** Multi-user security as obtained by a dedicated analysis, with the bound in the leakage case being from Theorem 7.2.

function along two dimensions: (1) Number of users, meaning su or mu, and (2) basic (no leakage) or with leakage. The bounds are summarized in Fig. 4 and discussed below. When we say the (i, j) table entry we mean the row i , column j entry of the table of Fig. 4.

First consider the basic (no leakage) case. We want to upper bound $\text{Adv}_{\text{CF},\mathbf{F}}^{\text{prf}}(\mathcal{B})$ for an adversary \mathcal{B} making $q_{\mathbf{F}}$ queries to the ideal compression function (oracle PRIM) and q queries to oracle FN. In the su setting (one NEW query) it is easy to see that the bound is the (1, 1) table entry. This is because a fairly standard argument bounds the advantage by the probability that \mathcal{B} makes a PRIM query containing the actual secret key K used to answer FN queries. We refer to issuing such a query as *guessing the secret key K* . Note that this probability is actually independent of the number q of FN queries and q does not figure in the bound. Now move to the mu setting, and let \mathcal{B} make u queries to its NEW oracle. Entry (2,1) of the table is the trivial bound obtained via Lemma 4.1 applied with \mathbf{F}_1 being our ideal compression function and \mathbf{F}_0 a family of all functions, but one has to be careful in applying the lemma. The subtle point is that adversary \mathcal{A}_1 built in Lemma 4.1 runs \mathcal{B} but makes an additional q queries to PRIM to compute the function \mathbf{F}_1 , so its advantage is the (1, 1) table entry with $q_{\mathbf{F}}$ replaced by $q_{\mathbf{F}} + q$. This term gets multiplied by u according to Equation (6), resulting in our (2, 1) table entry. A closer look shows one can do a tad better: the bound of the (1, 1) table entry extends with the caveat that a collisions between two different keys also allows the adversary to distinguish. In other words, the advantage is now bounded by the probability that \mathcal{B} guesses *any* of the u keys K_1, \dots, K_u , or that any two of these keys collide. This yields the (3, 1) entry of the table. Either way, the (well known) salient point here is that the advantage in the mu case is effectively u times the one in the su case.

We show that the growth of the advantage as a function of the number of users becomes far more favorable when the adversary obtains some leakage about the secret key under some function Out . For concreteness we take the leakage function to be truncation to r bits, meaning $\text{Out} = \text{TRUNC}_r$ is the function that returns the first $r \leq c$ bits of its input. (Theorem 7.2 will consider a general Out .) Now we seek to bound $\text{Adv}_{\text{CF},\text{Out},\mathbf{F}}^{\text{prf}}(\mathcal{B})$. Now, given only $\text{TRUNC}_r(K)$ for a secret key K , then there are only 2^{c-r} candidate secret keys consistent with this leakage, thus increasing the probability that

the adversary can guess the secret key. Consequently, the leakage-free bound from of the (1,1) entry generalizes to the bound of the (1,2) entry. Moving to multiple users, the (2,2) entry represents the naive bound obtained by applying Lemma 4.1. It is perhaps natural to expect that this is best possible as in the no-leakage case. We however observe that this is overly pessimistic. To this end, we exploit the following simple fact: *Every PRIM query (K, X) made by \mathcal{B} to the ideal compression function can only help in guessing a key K_i such that $\text{Out}(K) = \text{Out}(K_i)$.* In particular, every PRIM query (K, X) has only roughly $m \cdot 2^{-(c-r)}$ chance of guessing one of the u keys, where m is the number of generated keys K_i such that $\text{Out}(K_i) = K$. A standard balls-into-bins arguments (Lemma 7.1) can be used to infer that except with small probability (e.g., 2^{-c}), we always have $m \leq 2u/2^r + 3cr$ for any K . Combining these two facts yields our bound, which is the (3,2) entry of the table. Theorem 7.2 gives a more general result and the full proof. Note that if $r = 0$, i.e., nothing is leaked, this is close to the bound of the (1,3) entry and the bound does grow linearly with the number of users, but as r grows, the $3crq_F \cdot 2^{-(c-r)}$ term becomes the leading one, and does *not* grow with u . We now proceed to the detailed proof of the (3,2) entry.

COMBINATORIAL PRELIMINARIES. Our statements below will depend on an appropriate multi-collision probability of the output function $\text{Out}: \text{Out.D} \rightarrow \text{Out.R}$. In particular, for any $X_1, \dots, X_u \in \text{Out.R}$, we first define

$$\mu(X_1, \dots, X_u) = \max_{Y \in \text{Out.R}} |\{i : X_i = Y\}| ,$$

i.e., the number of occurrences of the most frequent value amongst X_1, \dots, X_u . In particular, this is an integer between 1 and u , and $\mu(X_1, \dots, X_u) = 1$ if all elements are distinct, whereas $\mu(X_1, \dots, X_u) = u$ if they are all equal. (Note when $u = 1$ the function has value 1.) Then, the m -collision probability of Out for u users is defined as

$$P_{\text{Out}}^{\text{coll}}(u, m) = \Pr_{K_1, \dots, K_u \leftarrow \text{Out.D}} [\mu(\text{Out}(K_1), \dots, \text{Out}(K_u)) \geq m] . \quad (26)$$

We provide a bound on $P_{\text{Out}}^{\text{coll}}(u, m)$ for the case where $\text{Out}(K)$, for a random K , is close enough to uniform. (We stress that a combinatorial restriction on Out is necessary for this probability to be small – it would be one if Out is the constant function, for example.) To this end, denote

$$\delta(\text{Out}) = \text{SD}(\text{Out}(K), R) = \frac{1}{2} \sum_{y \in \text{Out.R}} \left| \Pr[\text{Out}(K) = y] - \frac{1}{|\text{Out.R}|} \right| , \quad (27)$$

i.e., the statistical distance between $\text{Out}(K)$, where K is uniform on Out.D , and a random variable R uniform on Out.R .

We will use the following lemma, which we prove using standard balls-into-bins techniques.

Lemma 7.1 (Multi-collision probability) *Let $\text{Out} : \text{Out.D} \rightarrow \text{Out.R}$, $u \geq 1$, and $\lambda \geq 0$. Then, for any $m \leq u$ such that*

$$m \geq \frac{2u}{|\text{Out.R}|} + \lambda \ln |\text{Out.R}| , \quad (28)$$

we have

$$P_{\text{Out}}^{\text{coll}}(u, m) \leq u \cdot \delta(\text{Out}) + \exp(-\lambda/3) .$$

We stress that the factor 2 in Equation (28) can be omitted (one can use an additive Chernoff bound when u is sufficiently large in the proof given below, rather than a multiplicative one) at the cost of a less compact statement. As this factor will not be crucial in the following, we keep this simpler variant.

Proof of Lemma 7.1: To start with, note that the probability of having at least m values colliding increases only by at most $u \cdot \delta(\text{Out})$ if in the definition of P^{coll} , we replace the u outputs of Out under random inputs K_1, \dots, K_u by uniform elements of Out.R . Thus, we are going to consider the uniform case only, at the cost of the additive factor $u \cdot \delta(\text{Out})$. Throughout this proof, denote $R = |\text{Out.R}|$ for notational convenience.

To this end, let R_1, \dots, R_u be each uniform over Out.R , and let us fix $Y \in \text{Out.R}$. Let $T_i \in \{0, 1\}$ be 1 if and only if $R_i = Y$, and 0 otherwise. We are interested in bounding $T = \sum_{i=1}^u T_i$, i.e., showing that the probability it is at least m is at most $\frac{1}{R} \cdot \exp(-\lambda/3)$. Note that $\mu = \mathbf{E}[T] = \frac{u}{R}$, and recall that by the Chernoff bound,

$$\Pr[T \geq (1 + \epsilon) \cdot \mu] \leq \exp\left(-\frac{\epsilon^2}{2 + \epsilon} \mu\right).$$

Note that $m \geq 2\mu + \lambda \ln R$, and thus $\Pr[T \geq m] \leq \Pr[T \geq 2\mu + \lambda \ln R]$. To bound the latter, we consider two cases here: First, assume that $u \geq \lambda R \ln R$. Then,

$$\Pr[T \geq 2\mu + \lambda \ln R] \leq \Pr[T \geq 2\mu] \leq \exp(-\mu/3) \leq \frac{1}{R} \exp(-\lambda/3).$$

Second, assume instead that $u \leq \lambda R \ln R$. Then, let $\epsilon = \lambda R \ln R / u$. We have in particular,

$$\frac{\epsilon^2}{2 + \epsilon} \mu = \frac{\lambda^2 R (\ln R)^2}{2u + \lambda R \ln R} \geq \frac{1}{3} \lambda \ln R.$$

Therefore, again by the above Chernoff bound,

$$\Pr[T \geq 2\mu + \lambda \ln R] \leq \Pr[T \geq \mu(1 + \epsilon)] \leq \frac{1}{R} \exp(-\lambda/3).$$

So far, we have computed the probability for a specific and arbitrary $Y \in \text{Out.R}$. The final bound follows by the union bound over all R elements of Out.R . ■

For the analysis below, we also need to use a lower bound the number of potential preimages of a given output. To this end, given $\text{Out}: \text{Out.D} \rightarrow \text{Out.R}$, we define

$$\rho(\text{Out}) = \min_{y \in \text{Out.R}} |\text{Out}^{-1}(y)|.$$

SECURITY OF IDEAL COMPRESSION FUNCTIONS. The following theorem establishes the multi-user security under key-leakage of a random compression function. We stress that the bound here does *not* depend on the number of queries the adversary \mathcal{B} makes to oracle FN . Also, the parameter m can be set arbitrarily in the theorem statement for better flexibility, even though our applications below will mostly use the parameters from Lemma 7.1.

Theorem 7.2 *Let $\text{Out}: \mathcal{K} \rightarrow \text{Out.R}$. Then, for all $m \geq 1$, and all adversaries \mathcal{B} making u queries to NEW , and q_{F} queries to PRIM ,*

$$\text{Adv}_{\text{CF}, \text{Out}, \text{F}}^{\text{prf}}(\mathcal{B}) \leq \frac{u^2}{2|\mathcal{K}|} + \text{P}_{\text{Out}}^{\text{coll}}(u, m) + \frac{(m-1) \cdot q_{\text{F}}}{\rho(\text{Out})}.$$

The statement could be rendered useless whenever $\rho(\text{Out}) = 1$ because a single point has a single pre-image. We note here that Theorem 7.2 can easily be generalized to use a “soft” version of $\rho(\text{Out})$ guaranteeing that the number of preimages of a point is bounded from below by $\rho(\text{Out})$, except with some small probability ϵ , at the cost of an extra additive term $u \cdot \epsilon$. This more general version will not be necessary for our applications. We also note that it is unclear how to use the *average* number of preimages of $\text{Out}(\mathcal{K})$ in our proof.

<p><u>Game G_0, G_1</u></p> <p>$v \leftarrow 0$ $c' \leftarrow_s \mathcal{B}^{\text{NEW, FN, PRIM}}$ Return ($c' = 1$)</p> <p><u>NEW()</u> $v \leftarrow v + 1; K_v \leftarrow_s \mathcal{K}$ Return $\text{Out}(K_v)$</p> <p><u>PRIM(k, x)</u> if $T_F[k, x] = \perp$ then $T_F[k, x] \leftarrow_s \mathcal{K}$ If $\exists j : k = K_j$ and $T_{\text{FN}}[j, x] \neq \perp$ then $\text{bad}_1 \leftarrow \text{true}$ $T_F[k, x] \leftarrow T_{\text{FN}}[j, x]$ Return $T_F[k, x]$</p>	<p><u>FN(i, x)</u> If $T_{\text{FN}}[i, x] = \perp$ then $T_{\text{FN}}[i, x] \leftarrow_s \mathcal{K}$ If $T_F[K_i, x] \neq \perp$ then $\text{bad}_1 \leftarrow \text{true}$ $T_{\text{FN}}[i, x] \leftarrow T_F[K_i, x]$ else if $\exists j \neq i : K_j = K_i$ and $T_{\text{FN}}[j, x] \neq \perp$ then $\text{bad}_2 \leftarrow \text{true}$ $T_{\text{FN}}[i, x] \leftarrow T_{\text{FN}}[j, x]$ Return $T_{\text{FN}}[i, x]$</p>
---	---

Figure 5: **Games G_0 and G_1 in the proof of Theorem 7.2.** The boxed assignment statements are only executed in Game G_1 , but not in Game G_0 .

Proof of Theorem 7.2: The first step of the proof involves two games, G_0 and G_1 , given in Fig. 5. Game G_1 is semantically equivalent to $\text{PRF}_{\text{CF, Out, F}}$ with challenge bit $c = 1$, except that we have modified the concrete syntax of the oracles. In particular, the randomly sampled function $f \leftarrow_s \mathbf{F}$ is now implemented via lazy sampling, and the table entry $T_F[k, x]$ contains the value of $f(k, x)$ if it has been queried. Otherwise, T_F is \perp on all entries which have not been set. Also, the game keeps another table T_{FN} such that $T_{\text{FN}}[i, x]$ contains the value returned upon a query $\text{FN}(i, x)$. Note that the game enforces that any point in time, if $T_{\text{FN}}[i, x]$ and $T_F[K_i, x]$ are both set (i.e., they are not equal \perp), then we also have $T_{\text{FN}}[i, x] = T_F[K_i, x]$ and that, moreover, if $K_i = K_j$, then $T_{\text{FN}}[i, x] = T_{\text{FN}}[j, x]$ whenever both are not \perp . Finally, whenever any of these entries is set for the first time, then it is set to a fresh random value from \mathcal{K} . This guarantees that the combined behavior of the FN and the PRIM oracles are the same as in $\text{PRF}_{\text{CF, Out, F}}$ for the case $c = 1$. Thus,

$$\Pr[G_1] = \Pr[\text{PRF}_{\text{CF, Out, F}} | c = 1].$$

It is easier to see that in game G_0 , in contrast, the PRIM and FN oracles always return random values, and thus, since we are checking whether c' equals 1, rather than c , we get $\Pr[G_0] = 1 - \Pr[\text{PRF}_{\text{CF, Out, F}} | c = 0]$, and consequently,

$$\text{Adv}_{\text{CF, Out, F}}^{\text{prf}}(\mathcal{B}) = \Pr[G_1] - \Pr[G_0].$$

Both games G_0 and G_1 also include two flags bad_1 and bad_2 , initially false, which can be set to true when specific events occur. In particular, bad_1 is set whenever one of the following two events happens: Either \mathcal{B} queries $\text{FN}(i, x)$ after querying $\text{PRIM}(K_i, x)$, or \mathcal{B} queries $\text{PRIM}(K_i, x)$ after querying $\text{FN}(i, x)$. Moreover, bad_2 is set whenever \mathcal{B} queries $\text{FN}(i, x)$ after $\text{FN}(j, x)$, $K_i = K_j$, and $\text{PRIM}(K_i, x) = \text{PRIM}(K_j, x)$ was not queried earlier. (Note that if the latter condition is not true, then bad_1 has been set already.) It is immediate to see that G_0 and G_1 are identical until

<p><u>Game H₀</u></p> <p>$v \leftarrow 0$ $c' \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{NEW, FN, PRIM}}$ Return $(\exists j, x: T_{\text{F}}[K_j, x] \neq \perp)$</p> <p><u>Game H₁</u></p> <p>$v \leftarrow 0$ $c' \leftarrow_{\mathcal{B}} \mathcal{B}^{\text{NEW, FN, PRIM}}$ for $i = 0$ to $v - 1$ do $K'_i \leftarrow_{\mathcal{S}} \{ k' : \text{Out}(k') = Y_i \}$ Return $(\exists j, x: T_{\text{F}}[K'_j, x] \neq \perp)$</p>	<p><u>NEW()</u></p> <p>$v \leftarrow v + 1; K_v \leftarrow_{\mathcal{S}} \mathcal{K}; Y_v \leftarrow \text{Out}(K_v)$ Return Y_v</p> <p><u>PRIM(k, x)</u></p> <p>if $T_{\text{F}}[k, x] = \perp$ then $T_{\text{F}}[k, x] \leftarrow_{\mathcal{S}} \mathcal{K}$ Return $T_{\text{F}}[k, x]$</p> <p><u>FN(i, x)</u></p> <p>If $T_{\text{FN}}[i, x] = \perp$ then $T_{\text{FN}}[i, x] \leftarrow_{\mathcal{S}} \mathcal{K}$ Return $T_{\text{FN}}[i, x]$</p>
---	--

Figure 6: **Games H₀ and H₁ in the proof of Theorem 7.2.** Both games share the same NEW, PRIM, and FN oracles, the only difference being the additional re-sampling of the secret keys K'_i in the main procedure of H₁.

$\text{bad}_1 \vee \text{bad}_2$ is set. Therefore, by the fundamental lemma of game playing [5],

$$\text{Adv}_{\text{CF, Out, F}}^{\text{prf}}(\mathcal{B}) = \Pr[\text{G}_1] - \Pr[\text{G}_0] \leq \Pr[\text{G}_0 \text{ sets } \text{bad}_1] + \Pr[\text{G}_0 \text{ sets } \text{bad}_2]. \quad (29)$$

We immediately note that in order for bad_2 to be set in G_0 , we *must* have $K_i = K_j$ for distinct $i \neq j$, i.e., two keys must collide. Since we know that at most u calls are made to NEW, a simple Birthday bound yields

$$\Pr[\text{G}_0 \text{ sets } \text{bad}_2] \leq \frac{u^2}{2 \cdot |\mathcal{K}|}. \quad (30)$$

The rest of the proof thus deals with the more difficult problem of bounding $\Pr[\text{G}_0 \text{ sets } \text{bad}_1]$. To simplify this task, we first introduce a new game, called H₀ (cf. Fig. 6), which behaves as G_0 , except that it only checks at the end of the game whether the bad event triggering bad_1 has occurred during the interaction, in which case the game outputs true. Note that we are relaxing this check a bit further compared with G_0 , allowing it to succeed as long as a query to PRIM of form (K_j, x) for some j and some x was made, even if $\text{FN}(j, x)$ was never queried before. Therefore,

$$\Pr[\text{G}_0 \text{ sets } \text{bad}_1] \leq \Pr[\text{H}_0]. \quad (31)$$

Note that in H₀, the replies to all oracle calls made by \mathcal{B} do not depend on the keys K_1, K_2, \dots anymore, *except* for the leaked values $\text{Out}(K_1), \text{Out}(K_2), \dots$ returned by calls to NEW. We introduce a new and final game H₁ which modifies H₀ by pushing the sampling of the actual key values as far as possible in the game: That is, we first only gives values to \mathcal{B} with the correct leakage *distribution*, and in the final phase of H₁, when computing the game output, we sample keys that are consistent with this leakage. In other words, in the final check we replace the keys K_1, K_2, \dots with *freshly* sampled key K'_1, K'_2, \dots , which are uniform, under the condition that $\text{Out}(K_i) = \text{Out}(K'_i) = Y_i$.

It is not hard to see that $\Pr[\text{H}_0] = \Pr[\text{H}_1]$. This follows from two observations: First, for every i , the joint distribution of $(K_i, Y_i = \text{Out}(K_i))$ is identical to that of $(K'_i, Y_i = \text{Out}(K_i))$, since given Y_i , both K_i and K'_i are uniformly distributed over the set of pre-images of Y_i . Second, the behavior of both H₀ and H₁, before the final check to decide their outputs, only depends on values

$Y_i = \text{Out}(K_i)$, and *not* on the K_i 's. The actual keys K_i are only used for the final check, and since the probability distributions of K_i and K'_i conditioned on $\text{Out}(Y_i)$ are identical, then so are the probabilities of outputting true in games H_0 and H_1 .

Thus, combining Equation (29), Equation (30), and Equation (31), we have

$$\text{Adv}_{\text{CF,Out,F}}^{\text{prf}}(\mathcal{B}) \leq \frac{u^2}{2 \cdot |\mathcal{K}|} + \Pr[H_1]. \quad (32)$$

We are left with computing an upper bound on $\Pr[H_1]$. For this purpose, denote by \mathcal{S} the set of pairs (k, x) on which $T_F[k, x] \neq \perp$ after \mathcal{B} outputs its bit c' in H_1 . Also, let \mathcal{Y} be the multi-set $\{Y_0, Y_1, \dots, Y_{u-1}\}$ of values output by NEW to \mathcal{B} , and denote $\bar{\mathcal{Y}}$ the resulting set obtained by removing repetitions. Note that $|\mathcal{S}| \leq q_F$ and $|\bar{\mathcal{Y}}| \leq |\mathcal{Y}| \leq u$, and the first inequality may be strict, since some elements can be repeated due to collisions $\text{Out}(K_i) = \text{Out}(K_j)$.

Assume that now \mathcal{S} and \mathcal{Y} are given and fixed. We proceed to compute the probability that H_1 outputs true conditioned on the event that \mathcal{S} and \mathcal{Y} have been generated. For notational help, for every $y \in \bar{\mathcal{Y}}$, also denote

$$\mathcal{S}_y = \{ (k, x) \in \mathcal{S} : \text{Out}(k) = y \},$$

and let $q_y = |\mathcal{S}_y|$. Also, let n_y be the number of occurrence of $y \in \bar{\mathcal{Y}}$ in \mathcal{Y} . Note that except with probability $\text{P}_{\text{Out}}^{\text{coll}}(u, m)$, we have $n_y \leq m - 1$ for all $y \in \bar{\mathcal{Y}}$, and thus

$$\begin{aligned} \Pr[H_1] &\leq \Pr[\exists y \in \bar{\mathcal{Y}} : n_y \geq m] + \Pr[H_1 | \forall y \in \bar{\mathcal{Y}} : n_y < m] \\ &= \text{P}_{\text{Out}}^{\text{coll}}(u, m) + \Pr[H_1 | \forall y \in \bar{\mathcal{Y}} : n_y < m]. \end{aligned} \quad (33)$$

Therefore, let us assume we are given \mathcal{S} and \mathcal{Y} such that $n_y \leq m - 1$ for all $y \in \bar{\mathcal{Y}}$. Denote by $\Pr[H_1 | \mathcal{S}, \mathcal{Y}]$ the probability that H_1 outputs true conditioned on the fact that this \mathcal{S} and \mathcal{Y} has been generated. Using the fact that the keys $K'_0, K'_1, \dots, K'_{u-1}$ are sampled independently of \mathcal{S} , we compute

$$\begin{aligned} \Pr[H_1 | \mathcal{S}, \mathcal{Y}] &= \Pr[\exists j, x : (K'_j, x) \in \mathcal{S}] \leq \sum_{y \in \bar{\mathcal{Y}}} \frac{q_y \cdot n_y}{|\text{Out}^{-1}(y)|} \\ &\leq (m - 1) \cdot \sum_{y \in \bar{\mathcal{Y}}} \frac{q_y}{|\text{Out}^{-1}(y)|} \leq \frac{m - 1}{\rho(\text{Out})} \sum_{y \in \bar{\mathcal{Y}}} q_y \leq \frac{(m - 1)q_F}{\rho(\text{Out})}. \end{aligned}$$

Since the bound holds for all such \mathcal{S} and \mathcal{Y} , we also have

$$\Pr[H_1 | \forall y \in \bar{\mathcal{Y}} : n_y < m] \leq \frac{(m - 1)q_F}{\rho(\text{Out})}. \quad (34)$$

The final bound follows by combining Equation (32), Equation (33), and Equation (34). \blacksquare

8 Quantitative bounds for augmented cascades and AMAC

We consider two instantiations of augmented cascades, one using bit truncation, the other using modular reduction. We give concrete bounds on the mu prf security of these constructions in the ideal compression function model, combining results from above. This will give us good guidelines for a comparison with existing constructions – such as NMAC and sponges – in Section 9.

BIT TRUNCATION. Let $\mathcal{K} = \{0, 1\}^c$, and $\text{Out} = \text{TRUNC}_r : \{0, 1\}^c \rightarrow \{0, 1\}^r$, for $r \leq c$, outputs the first r bits of its inputs, i.e., $\text{TRUNC}_r(X) = X[1..r]$. Note that $\delta(\text{TRUNC}_r) = 0$, since omitting

$c - r$ bits does not affect uniformity, and $\rho(\text{TRUNC}_r) = 2^{c-r}$, since every r -bit strings has 2^{c-r} preimages. Then, combining Lemma 7.1 with Theorem 7.2, using $m = 2u/2^r + 3cr$, we obtain the following corollary, denoting with \mathbf{F}_c the ideal compression function for $\mathcal{K} = \{0, 1\}^c$. (We do not specify \mathcal{X} further, as it does not influence the statement.)

Corollary 8.1 *For any $c \leq r$, and all adversaries \mathcal{B} making u queries to NEW and q_F queries to PRIM,*

$$\text{Adv}_{\text{CF,TRUNC}_r,\mathbf{F}_c}^{\text{prf}}(\mathcal{B}) \leq \frac{u^2}{2^{c+1}} + \frac{2u \cdot q_F}{2^c} + \frac{3cr \cdot q_F}{2^{c-r}} + \exp(-c) . \blacksquare$$

We can then use this result to obtain our bounds for the augmented cascade $\mathbf{ACSC}[\text{CF}, \text{CF}, \text{TRUNC}_r]$ when using an ideal compression function $\{0, 1\}^c \times \mathcal{X} \rightarrow \{0, 1\}^c$.

Theorem 8.2 (mu prf security for r -bit truncation) *For any $r \leq n$, and all adversaries \mathcal{A} making q queries to FN consisting of vectors from \mathcal{X}^* of length at most ℓ , q_F queries to PRIM, and $u \leq q$ queries to NEW,*

$$\text{Adv}_{\mathbf{ACSC}[\text{CF},\text{CF},\text{TRUNC}_r],\mathbf{F}_c}^{\text{prf}}(\mathcal{A}) \leq \frac{5\ell^2 q^2 + 3\ell q q_F}{2^c} + \frac{3cr\ell \cdot (q\ell + q_F)}{2^{c-r}} + \ell \exp(-c) \blacksquare$$

Proof of Theorem 8.2: Let \mathbf{A}_{c-r} be the set of all functions with domain \mathcal{X}^* and $c - r$ -bit outputs, and let \mathbf{A}_c be set of all functions with domain \mathcal{X}^* and c -bit outputs. First note that by Theorem 5.3,

$$\begin{aligned} \text{Adv}_{\text{TRUNC}_r \circ \mathbf{A}_c, \text{TRUNC}_r \circ \mathbf{CSC}[\text{CF}, \text{CF}], \mathbf{F}_c}^{\text{dist}}(\mathcal{A}) \\ \leq \ell \cdot \text{Adv}_{\text{CF,TRUNC}_r,\mathbf{F}_c}^{\text{prf}}(\mathcal{A}_h) + 2 \text{Adv}_{\text{CF},\mathbf{F}_c}^{\text{prf}}(\mathcal{A}_g) . \end{aligned} \quad (35)$$

Note that $\text{TRUNC}_r \circ \mathbf{A}_c$ and \mathbf{A}_{c-r} have the same distribution, and thus

$$\text{Adv}_{\mathbf{ACSC}[\text{CF},\text{CF},\text{TRUNC}_r],\mathbf{F}_c}^{\text{prf}}(\mathcal{A}) = \text{Adv}_{\text{TRUNC}_r \circ \mathbf{A}_c, \text{TRUNC}_r \circ \mathbf{CSC}[\text{CF}, \text{CF}], \mathbf{F}_c}^{\text{dist}}(\mathcal{A}) . \quad (36)$$

To upper bound the two advantages in Equation (35), recall that adversary \mathcal{A}_h makes at most q queries to its NEW oracle and at most q queries to its FN oracle. Adversary \mathcal{A}_g makes u queries to its NEW oracle and at most q queries to its FN oracle. The running time of both constructed adversaries is about that of \mathcal{A} plus the time for $q\ell$ computations of CF, which in particular means that both adversaries make $q_F + q \cdot \ell$ queries to PRIM. Therefore, by Corollary 8.1,

$$\text{Adv}_{\text{CF,TRUNC}_r,\mathbf{F}_c}^{\text{prf}}(\mathcal{A}_h) \leq \frac{q^2}{2^{c+1}} + \frac{2q \cdot (q\ell + q_F)}{2^c} + \frac{3cr \cdot (q\ell + q_F)}{2^{c-r}} + \exp(-c) , \quad (37)$$

as well as

$$\text{Adv}_{\text{CF,TRUNC}_r,\mathbf{F}_c}^{\text{prf}}(\mathcal{A}_g) \leq \frac{u^2}{2^{c+1}} + \frac{u \cdot (q\ell + q_F)}{2^c} . \quad (38)$$

The theorem statement then follow by combining Equation (35), Equation (36), and Equation (37). \blacksquare

MODULAR REDUCTION. Our second example becomes particularly important for the application to the Ed25519 signature scheme.

Here, we let $\mathcal{K} = \mathbb{Z}_N$, and consider the output function $\text{Out} = \text{MOD}_M : \mathbb{Z}_N \rightarrow \mathbb{Z}_M$ for $M \leq N$ is such that $\text{MOD}_M(X) = X \bmod M$. (Note that as a special case, we think of $\mathcal{K} = \{0, 1\}^c$ here as \mathbb{Z}_{2^c} .) We need the following two properties of MOD_M .

Lemma 8.3 For all $M \leq N$: (1) $\rho(\text{MOD}_M) \geq \frac{N}{M} - 1$, (2) $\delta(\text{MOD}_M) \leq M/N$.

Proof of Lemma 8.3: For every $a \in \mathbb{Z}_M$, we start by counting the number of $x \in \mathbb{Z}_N$ with $x \bmod M = a$. In particular, the $\lfloor N/M \rfloor$ integers $a_i = M \cdot i + a$ for $i \in [0 \dots \lfloor N/M \rfloor - 1]$ are all those with this property if $a \geq N \bmod M$. Otherwise, if $a < N \bmod M$, also $a_{\lfloor N/M \rfloor} = M \cdot \lfloor N/M \rfloor + a$ additionally has this property. Thus

$$\rho(\text{MOD}_M) \geq \left\lfloor \frac{N}{M} \right\rfloor \geq \frac{N}{M} - 1.$$

Now, for the statistical distance, $\delta(\text{MOD}_M)$, note that by the above, for all $a \geq N \bmod M$, and U_N uniformly distributed on \mathbb{Z}_N ,

$$\Pr[\text{MOD}_M(U_N) = a] = \frac{1}{N} \left\lfloor \frac{N}{M} \right\rfloor \leq \frac{1}{M},$$

whereas for all $a < N \bmod M$ (if they exist),

$$\Pr[\text{MOD}_M(U_N) = a] = \frac{1}{N} \left\lceil \frac{N}{M} \right\rceil \geq \frac{1}{M}.$$

Thus,

$$\begin{aligned} \delta(\text{MOD}_M) &= \sum_{a: a < N \bmod M} \left(\Pr[\text{MOD}_M(U_N) = a] - \frac{1}{M} \right) \\ &\leq M \cdot \left(\frac{1}{N} \left\lceil \frac{N}{M} \right\rceil - \frac{1}{M} \right) \leq M \cdot \left(\frac{1}{N} \frac{N}{M} + \frac{1}{N} - \frac{1}{M} \right) \leq \frac{M}{N}. \end{aligned}$$

■

Then, combining Lemma 7.1 and Lemma 8.3 with Theorem 7.2, using $m = 2u/M + 3 \ln N \ln M$, we obtain the following corollary, denoting with \mathbf{F}_N the ideal compression function with $\mathcal{K} = \mathbb{Z}_N$. (As above, we do not specify \mathcal{X} further, as it does not influence the statement.)

Corollary 8.4 For any $M \leq N/2$, and all adversaries \mathcal{B} making u queries to NEW and q_F queries to PRIM,

$$\text{Adv}_{\text{CF}, \text{MOD}_M, \mathbf{F}_N}^{\text{prf}}(\mathcal{B}) \leq \frac{u^2}{2N} + \frac{uM}{N} + \frac{4u \cdot q_F}{N} + \frac{6M \ln N \ln M \cdot q_F}{N} + \frac{1}{N}. \blacksquare$$

This can once again be used to obtain the final analysis of the augmented cascade using modular reduction. The proof is similar to that of Theorem 8.2.

Theorem 8.5 (mu prf security for modular reduction) For any $M \leq N/2$, and all adversaries \mathcal{A} making q queries to FN consisting of vectors from \mathcal{X}^* of length at most ℓ , q_F queries to PRIM, and $u \leq q$ queries to NEW,

$$\begin{aligned} \text{Adv}_{\text{ACSC}[\text{CF}, \text{CF}, \text{MOD}_M], \mathbf{F}_N}^{\text{prf}}(\mathcal{A}) &\leq \frac{5\ell^2 q^2 + 3\ell q q_F}{N} \\ &\quad + \frac{7M \ln N \ln M (\ell^2 q + \ell q_F)}{N} + \frac{\ell}{N}. \blacksquare \end{aligned}$$

Proof of Theorem 8.5: Let \mathbf{A}_M be the set of all functions with domain \mathcal{X}^* and outputs in \mathbb{Z}_M , and let \mathbf{A}_N be set of all functions with domain \mathcal{X}^* and outputs in \mathbb{Z}_N . First note that by

Theorem 5.3,

$$\begin{aligned} \text{Adv}_{\text{MOD}_M \circ A_N, \text{MOD}_M \circ \text{CSC}[\text{CF}, \text{CF}], \mathbf{F}_N}^{\text{dist}}(\mathcal{A}) \\ \leq \ell \cdot \text{Adv}_{\text{CF}, \text{MOD}_M, \mathbf{F}_N}^{\text{prf}}(\mathcal{A}_h) + 2 \text{Adv}_{\text{CF}, \mathbf{F}_N}^{\text{prf}}(\mathcal{A}_g). \end{aligned} \quad (39)$$

Note that $\text{MOD}_M \circ A_N$ and A_M do *not* have the same distribution. Still, by the triangle inequality,

$$\begin{aligned} \text{Adv}_{\text{ACSC}[\text{CF}, \text{CF}, \text{MOD}_M], \mathbf{F}_N}^{\text{prf}}(\mathcal{A}) \\ \leq \text{Adv}_{A_M, \text{MOD}_M \circ A_N, \mathbf{F}_N}^{\text{dist}}(\mathcal{A}) + \text{Adv}_{\text{MOD}_M \circ A_N, \text{MOD}_M \circ \text{CSC}[\text{CF}, \text{CF}], \mathbf{F}_N}^{\text{dist}}(\mathcal{A}), \end{aligned} \quad (40)$$

and Lemma 8.3 directly yields

$$\text{Adv}_{A, \text{MOD}_M \circ A_N, \mathbf{F}_N}^{\text{dist}}(\mathcal{A}) \leq q \cdot \delta(\text{MOD}_M) \leq q \cdot \frac{M}{N}. \quad (41)$$

To upper bound the two advantages in Equation (39), recall that adversary \mathcal{A}_h makes at most q queries to its NEW oracle and at most q queries to its FN oracle. Adversary \mathcal{A}_g makes u queries to its NEW oracle and at most q queries to its FN oracle. The running time of both constructed adversaries is about that of \mathcal{A} plus the time for $q\ell$ computations of CF, which in particular means that both adversaries make $q_F + q \cdot \ell$ queries to PRIM. Therefore, by Corollary 8.4,

$$\begin{aligned} \text{Adv}_{\text{CF}, \text{MOD}_M, \mathbf{F}_N}^{\text{prf}}(\mathcal{A}_h) \\ \leq \frac{q^2}{2N} + \frac{qM}{N} + \frac{4q \cdot (q_F + q \cdot \ell)}{N} + \frac{6Mc \ln M \cdot (q_F + q \cdot \ell)}{N} + \exp(-c). \end{aligned} \quad (42)$$

as well as

$$\text{Adv}_{\text{CF}, \text{MOD}_N, \mathbf{F}_N}^{\text{prf}}(\mathcal{A}_g) \leq \frac{u^2}{2N} + \frac{u \cdot (q\ell + q_F)}{N}. \quad (43)$$

The theorem statement then follow by combining all of the above equations. \blacksquare

BOUNDS FOR AMAC. The above bounds are for augmented cascades, but they can easily be adapted to AMAC, at the cost of adding an extra additive term, which we now discuss. Recall that $\text{AMAC}(K, M) = \text{Out}(H(K||M))$, where the iterated hash function H is derived from a compression function h . We only consider here the special case where the key K is completely handled by the first compression function call of H (and is exactly a random element of \mathcal{X}), and the message is processed from the second call onwards. In other words, AMAC is the 2-tier cascade with the first tier being the dual of h , meaning the key and input roles are swapped. In particular, we can use Theorem 5.3, which would give us a modified version of the above bounds with an additional additive term, accounting $2 \text{Adv}_g^{\text{prf}}(\mathcal{A}_g)$ for \mathcal{A}_g as given in the reduction. This can easily be upper bounded (using the dedicated mu bound from Fig. 4) as

$$2 \cdot \text{Adv}_g^{\text{prf}}(\mathcal{A}_g) \leq \frac{u^2 + u(q_F + q\ell)}{|\mathcal{X}|} \leq \frac{q^2 + q(q_F + q\ell)}{|\mathcal{X}|}.$$

9 Comparisons

We compare the quantitative bounds obtained for augmented cascades above with those from NMAC and sponges. In particular, we show that the security of augmented cascades is comparable to that of NMAC when its output is truncated, and superior to that of keyed sponges.

COMPARISON WITH NMAC. We start with concrete bounds for the ideal compression function security of NMAC, for both cases where the output is processed by TRUNC_r or MOD_m . In particular, for a compression function $f : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{K}$, we consider the construction NMAC such that

$$\text{NMAC}[f]((K^{\text{in}}, K^{\text{out}}), \mathbf{X}) = f(K^{\text{out}}, \text{CSC}[f, f](K^{\text{in}}, \mathbf{X}) \parallel \text{pad}) .$$

We are interested in the concrete security of NMAC when f is replaced by CF using the ideal compression function \mathbf{F} as in Section 7. To the best of our knowledge, the following concrete bounds for NMAC have not appeared anywhere, but the statement is somewhat folklore in the single-user case. (For some parameter regime, similar bounds can be inferred from [14]) The proof follows by a fairly standard argument, as sketched below.

Theorem 9.1 (NMAC with r -bit truncation) *For any $r \leq n$, and all adversaries \mathcal{A} making q queries to FN consisting of vectors from \mathcal{X}^* of length at most ℓ , $q_{\mathbf{F}}$ queries to PRIM, and u queries to NEW,*

$$\text{Adv}_{\text{TRUNC}_r \circ \text{NMAC}[\text{CF}, \mathbf{F}_c]}^{\text{prf}}(\mathcal{A}) \leq \frac{u^2}{2^{c+1}} + \frac{2\ell^2 q^2}{2^c} + \frac{2q\ell q_{\mathbf{F}}}{2^c} + \frac{2uq_{\mathbf{F}}}{2^c} . \blacksquare$$

Theorem 9.2 (NMAC with modular truncation) *For any $M \leq N$, and all adversaries \mathcal{A} making q queries to FN consisting of vectors from \mathcal{X}^* of length at most ℓ , $q_{\mathbf{F}}$ queries to PRIM, and u queries to NEW,*

$$\text{Adv}_{\text{MOD}_M \circ \text{NMAC}[\text{CF}, \mathbf{F}_N]}^{\text{prf}}(\mathcal{A}) \leq q \cdot \frac{M}{N} + \frac{u^2}{2N} + \frac{2\ell^2 q^2}{N} + \frac{2q\ell q_{\mathbf{F}}}{N} + \frac{2uq_{\mathbf{F}}}{2^c} . \blacksquare$$

PROOF SKETCH. Let us look at Theorem 9.1 for the case $r = n$ first.

Each of the at most q queries (u_i, \mathbf{X}_i) to FN makes internally up to $\ell + 1$ queries to \mathbf{F}_c . A pair (u_i, \mathbf{X}_i) and (u_j, \mathbf{X}_j) of such FN queries may be forced to invoke \mathbf{F}_c on some common inputs trivially, in particular when these two queries are for (1) the same user, and (2) the corresponding messages share a common prefix (this includes the case that \mathbf{X}_i is itself a prefix of \mathbf{X}_j , or vice versa).

A sufficient condition for the outputs of these FN queries to be random looking is that every internal query (K, X) to \mathbf{F}_c is *unique*, i.e., no other FN or PRIM query results in invoking \mathbf{F}_c on the same input (K, X) . The only exception are those unavoidable collisions as above, i.e., the same query (K, X) to \mathbf{F}_c is made when processing another query which needs to evaluate some common queries. It is not hard to show that the probability that this condition is not true is at most

$$\frac{u^2}{2^{c+1}} + \frac{2\ell^2 q^2}{2^c} + \frac{2q\ell q_{\mathbf{F}}}{2^c} + \frac{2uq_{\mathbf{F}}}{2^c} ,$$

which is also an upper bound on the advantage. This also implies the general case for $r \leq n$, as truncation does not increase the adversary's advantage. As for modular reduction, one needs to additionally take into account the additional bias of the outputs, which results in an additive $q \cdot \delta(\text{MOD}_M) = q \cdot \frac{M}{N}$ term.

As a side remark, we note that the bound is somewhat generous and not tight, yet improving upon these bounds remains an open problem. A recent paper by Gaži, Pietrzak, and Tessaro [18] improves this bound by considering a modification of NMAC using block whitening, but it is open to verify whether their bound holds for NMAC, too. Still, we note that the main issue with respect to tightness is the growth of the bound with respect to the message length. For short message length, the above bound is essentially tight.

NMAC VS AUGMENTED CASCADES. We note that the bounds of Theorem 8.2 and Theorem 9.1 are similar in many respects, showing that the concrete security for both constructions is comparable. The main difference between the two constructions is the presence of additional terms with denominator 2^{c-r} in Theorem 8.2 which are not present in Theorem 9.1. These have order $q\ell^2/2^{c-r}$ and $q_{\mathbb{F}}\ell/2^{c-r}$ (ignoring small multiplicative factors), respectively. The crucial point here is that *the dependencies on q and $q_{\mathbb{F}}$ are linear*. For example, in the typical case that $c = 2r$ (e.g., $r = 256$), for small ℓ (which is common in many applications) these terms become large roughly when q and $q_{\mathbb{F}}$ approach $2^{c/2}$. In this regime, NMAC is similarly insecure.

SPONGES VS AUGMENTED CASCADES. In contrast to NMAC, the situation is inverted with respect to keyed sponges, which also use truncation as a means to prevent extension attacks and achieve prf security. Indeed, GPT [17] show that when keyed sponges use a c -bit state and output r bits of this state as their output, then there exist attacks with advantage $q^2/2^{c-r}$ and $qq_{\mathbb{F}}/2^{c-r}$, even in the single-user case. For $c = 2r = 512$, for example, keyed sponges can be distinguished with $q = 2^{128}$ queries to FN. In contrast, no terms of such order appear in Theorem 8.2, thus showing that the security of augmented cascades is superior to that of sponges. Needless to say this is not a problem for practical instantiations (like those based on SHA-3), where $c - r$ is generally at least 512, but it shows theoretical gains of augmented cascades over sponges when setting equal parameters.

10 Security of the Davies-Meyer construction

One might object that practical compression functions are not un-structured enough to be treated as random because they are built from blockciphers via the Davies-Meyer construction. Accordingly, we study the mu PRF security under leakage of the Davies-Meyer construction with an ideal blockcipher and show that bounds of the quality we have seen for a random compression function continue to hold. Of course, one could extend this and prove similar bounds for other compression functions, like the PGV [25] ones.

Recall that given a block cipher $E: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{D}$, the Davies-Meyer (DM) construction outputs $E_X(H) \oplus H$ for chaining value $H \in \mathcal{D}$ and message block $X \in \mathcal{K}$. Formally, we can model the (keyed) Davies-Meyer construction as an oracle function family DM where $\text{DM.K} = \mathcal{D}$, $\text{DM.D} = \mathcal{K}$, and $\text{DM.R} = \mathcal{D}$. Its oracle space consist of all oracles allowing both forward and backward access to a block cipher, which is the same as that of the ideal-cipher primitive (which we denote as **IC**) defined above.

Theorem 10.1 (prf security of Davies-Meyer) *Let $\text{Out}: \mathcal{D} \rightarrow \text{Out.R}$. Then, for all $m \geq 2$, and all adversaries \mathcal{B} making u queries to NEW, q_{FN} queries to FN, and $q_{\mathbb{F}}$ queries to PRIM,*

$$\text{Adv}_{\text{DM,Out,IC}}^{\text{prf}}(\mathcal{B}) \leq \frac{u(u + 2q_{\text{FN}})}{2|\mathcal{D}|} + \text{P}_{\text{Out}}^{\text{coll}}(u, m) + \frac{(m-1) \cdot q_{\mathbb{F}}}{\rho(\text{Out})}.$$

Proof of Theorem 10.1: The proof is similar to that of Theorem 7.2, and we assume that the reader is familiar with its proof, as this will allow for somewhat more compact explanations here. The first step of the proof involves two games, G_0 and G_1 , given in Fig. 7.

Game G_1 (where the boxed statements are executed) is semantically equivalent to $\text{PRF}_{\text{DM,Out,IC}}$ with challenge bit $c = 1$, except that we have modified the concrete syntax of the oracles. In particular, the underlying ideal cipher **IC** is implemented via lazy sampling, and the table entries $T_{\text{PRIM}}[K, X, +]$ and $T_{\text{PRIM}}[K, Y, -]$ contain the values $E(K, X)$ and $E^{-1}(K, Y)$, respectively, for the sampled block cipher E , if it has been queried via a PRIM query. (Note that the game always set

<p><u>Game $G_0, \boxed{G_1}$</u></p> <p>$v \leftarrow 0$ $c' \leftarrow \mathcal{B}^{\text{NEW, FN, F}}$ Return ($c' = 1$)</p> <p><u>PRIM($x, k, +$)</u> if $T_{\text{PRIM}}[x, k, +] = \perp$ then $T_{\text{PRIM}}[x, k, +] \leftarrow \mathcal{D} \setminus T_{\text{PRIM}}[x, \cdot, +].\text{R}$ If $\exists j : k = K_j$ and $T_{\text{FN}}[j, x] \neq \perp$ then $\text{bad}_1 \leftarrow \text{true}$ $\boxed{T_{\text{PRIM}}[x, k, +] \leftarrow T_{\text{F}}[j, x]}$ $T_{\text{PRIM}}[x, T_{\text{PRIM}}[x, k, +], -] \leftarrow k$ Return $T_{\text{PRIM}}[x, k, +]$</p> <p><u>PRIM($x, z, -$)</u> if $T_{\text{PRIM}}[x, z, -] = \perp$ then $T_{\text{PRIM}}[x, z, -] \leftarrow \mathcal{D} \setminus T_{\text{PRIM}}[x, \cdot, -].\text{R}$ If $\exists j : T_{\text{FN}}[j, x] = z$ then $\text{bad}_1 \leftarrow \text{true}$ $\boxed{T_{\text{PRIM}}[x, k, +] \leftarrow K_j}$ $T_{\text{PRIM}}[x, T_{\text{PRIM}}[x, z, -], +] \leftarrow z$ Return $T_{\text{PRIM}}[k, z, -]$</p>	<p><u>NEW()</u></p> <p>$v \leftarrow v + 1 ; K_v \leftarrow \mathcal{K}$ Return $\text{Out}(K_v)$</p> <p><u>FN(i, x)</u> If $T_{\text{FN}}[i, x] = \perp$ then $T_{\text{FN}}[i, x] \leftarrow \mathcal{D}$ If $T_{\text{PRIM}}[x, K_i, +] \neq \perp$ then $\text{bad}_1 \leftarrow \text{true}$ $\boxed{T_{\text{FN}}[i, x] \leftarrow T_{\text{PRIM}}[x, K_i, +]}$ else if $\exists j \neq i : K_j = K_i$ and $T_{\text{FN}}[j, x] \neq \perp$ then $\text{bad}_2 \leftarrow \text{true}$ $\boxed{T_{\text{FN}}[i, x] \leftarrow T_{\text{FN}}[j, x]}$ else if $\exists j \neq i : T_{\text{FN}}[j, x] = T_{\text{FN}}[i, x]$ then $\text{bad}_3 \leftarrow \text{true}$ $\boxed{T_{\text{FN}}[i, x] \leftarrow \mathcal{D} \setminus T_{\text{FN}}[\cdot, x].\text{R}}$ Return $T_{\text{FN}}[i, x] \oplus K_i$</p>
--	--

Figure 7: **Games G_0 and G_1 in the proof of Theorem 10.1.** The boxed statements are only executed in Game G_1 , but not in Game G_0 . Here, $T_{\text{PRIM}}[x, \cdot, +].\text{R}$ is the set of all values z such that there exists a k with $T_{\text{FN}}[x, k, +] = z$. The notations $T_{\text{PRIM}}[x, \cdot, -].\text{R}$ and $T_{\text{FN}}[\cdot, x].\text{R}$ are defined analogously.

$T_{\text{PRIM}}[K, X, +]$ and $T_{\text{PRIM}}[K, Y, -]$ jointly in a consistent way.) Otherwise, T_{PRIM} is \perp on all entries which have not been set so far.

Also, the game keeps another table T_{FN} such that $T_{\text{FN}}[i, x] \oplus K_i$ is the value returned upon a query $\text{FN}(i, x)$. Note that the game enforces that any point in time, if $T_{\text{FN}}[i, x]$ and $T_{\text{PRIM}}[x, K_i, +]$ are both set (i.e., they are not equal \perp), then we also have $T_{\text{FN}}[i, x] = T_{\text{PRIM}}[x, K_i, +]$ and that, moreover, if $K_i = K_j$, then $T_{\text{FN}}[i, x] = T_{\text{FN}}[j, x]$ whenever both are not \perp . Finally, whenever any of these entries is set for the first time, then it is set to a fresh random value from \mathcal{D} constrained on not violating the permutation constraint. This guarantees that the combined behavior of the FN and the PRIM oracles are the same as in $\text{PRF}_{\text{DM, Out, IC}}$ for the case $c = 1$. Thus,

$$\Pr[G_1] = \Pr[\text{PRF}_{\text{DM, Out, IC}} \mid c = 1].$$

It is easier to see that in game G_0 , in contrast, the FN oracles always return random values (the fact that K_i is xored to the $T_{\text{FN}}[i, x]$ does not modify the distribution), and PRIM behaves like an

independent ideal cipher. Thus, since we are checking whether c' equals 1, rather than c , we have

$$\Pr [G_0] = 1 - \Pr[\text{PRF}_{\text{DM,Out,IC}} | c = 0] .$$

Consequently,

$$\text{Adv}_{\text{DM,Out,IC}}^{\text{prf}}(\mathcal{B}) = \Pr [G_1] - \Pr [G_0] .$$

Both games G_0 and G_1 also include three flags bad_1 , bad_2 , and bad_3 , initially false, which can be set to true when specific events occur. It is immediate to see that G_0 and G_1 are identical until $\text{bad}_1 \vee \text{bad}_2 \vee \text{bad}_3$ is true. Therefore, by the fundamental lemma of game playing [5],

$$\begin{aligned} \text{Adv}_{\text{DM,Out,IC}}^{\text{prf}}(\mathcal{B}) &= \Pr [G_1] - \Pr [G_0] \\ &\leq \Pr [G_0 \text{ sets } \text{bad}_1] + \Pr [G_0 \text{ sets } \text{bad}_2] + \Pr [G_0 \text{ sets } \text{bad}_3] . \end{aligned} \tag{44}$$

As in the proof of Theorem 7.2,

$$\Pr [G_0 \text{ sets } \text{bad}_2] \leq \frac{u^2}{2 \cdot |\mathcal{D}|} . \tag{45}$$

As for bad_3 , note that for this to happen, one of the random values generated when answering an FN query for (i, x) must hit a previously generated value for (j, x) , for which there are at most u candidates. Thus, by the union bound,

$$\Pr [G_0 \text{ sets } \text{bad}_3] \leq \frac{q_{\text{FN}} \cdot u}{|\mathcal{D}|} . \tag{46}$$

We are left with the problem of upper bounding $\Pr [G_0 \text{ sets } \text{bad}_1]$. We note however that this part of the proof can be carried out exactly as in the proof of Theorem 7.2, and results in the identical bound. It is thus omitted. \blacksquare

Acknowledgments

We thank the Eurocrypt 2016 reviewers for their comments.

References

- [1] E. Andreeva, J. Daemen, B. Mennink, and G. V. Assche. Security of keyed sponge constructions using a modular proof approach. In G. Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 364–384. Springer, Heidelberg, Mar. 2015. 3, 7
- [2] M. Bellare. New proofs for NMAC and HMAC: Security without collision-resistance. In C. Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 602–619. Springer, Heidelberg, Aug. 2006. 4, 6, 7, 17
- [3] M. Bellare, R. Canetti, and H. Krawczyk. Keying hash functions for message authentication. In N. Kobitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 1–15. Springer, Heidelberg, Aug. 1996. 3, 6, 7
- [4] M. Bellare, R. Canetti, and H. Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th FOCS*, pages 514–523. IEEE Computer Society Press, Oct. 1996. 3, 5, 7, 10, 11, 12
- [5] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. 8, 24, 32

- [6] D. J. Bernstein. Extending the Salsa20 nonce. In *Symmetric key encryption workshop (SKEW)*, February 2011. URL: <https://cr.yp.to/papers.html#xsalsa>. 11
- [7] D. J. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, Heidelberg, Sept. / Oct. 2011. 3, 34
- [8] G. Bertoni, J. Daemen, M. Peeters, and G. Assche. On the security of the keyed sponge construction. In *Symmetric key encryption workshop (SKEW)*, February 2011. 3, 7
- [9] N. Brown. Things that use Ed25519. <http://ianix.com/pub/ed25519-deployment.html>. 3
- [10] D. Chang, M. Dworkin, S. Hong, J. Kelsey, and M. Nandi. A keyed sponge construction with pseudo-randomness in the standard model. In *The Third SHA-3 Candidate Conference (March 2012)*, 2012. 3, 7
- [11] J.-S. Coron, Y. Dodis, C. Malinaud, and P. Puniya. Merkle-Damgård revisited: How to construct a hash function. In V. Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 430–448. Springer, Heidelberg, Aug. 2005. 6, 7
- [12] I. Damgård. A design principle for hash functions. In G. Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 416–427. Springer, Heidelberg, Aug. 1990. 3
- [13] Y. Dodis and K. Pietrzak. Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In T. Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 21–40. Springer, Heidelberg, Aug. 2010. 8
- [14] Y. Dodis, T. Ristenpart, J. P. Steinberger, and S. Tessaro. To hash or not to hash again? (In)differentiability results for h^2 and HMAC. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 348–366. Springer, Heidelberg, Aug. 2012. 29
- [15] S. Dziembowski and K. Pietrzak. Leakage-resilient cryptography. In *49th FOCS*, pages 293–302. IEEE Computer Society Press, Oct. 2008. 8
- [16] P. Gazi, K. Pietrzak, and M. Rybár. The exact PRF-security of NMAC and HMAC. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 113–130. Springer, Heidelberg, Aug. 2014. 4, 7
- [17] P. Gazi, K. Pietrzak, and S. Tessaro. The exact PRF security of truncation: Tight bounds for keyed sponges and truncated CBC. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 368–387. Springer, Heidelberg, Aug. 2015. 3, 7, 30
- [18] P. Gazi, K. Pietrzak, and S. Tessaro. Generic security of NMAC and HMAC with input whitening. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 85–109. Springer, Heidelberg, Nov. / Dec. 2015. 29
- [19] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, Oct. 1986. 3, 5, 11, 12
- [20] U. M. Maurer, R. Renner, and C. Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In M. Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, Feb. 2004. 7
- [21] B. Mennink, R. Reyhanitabar, and D. Vizár. Security of full-state keyed sponge and duplex: Applications to authenticated encryption. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 465–489. Springer, Heidelberg, Nov. / Dec. 2015. 3, 7
- [22] R. C. Merkle. One way hash functions and DES. In G. Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 428–446. Springer, Heidelberg, Aug. 1990. 3
- [23] N. Mouha and A. Luykx. Multi-key security: The Even-Mansour construction revisited. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 209–223. Springer, Heidelberg, Aug. 2015. 5

- [24] E. D. Mulder, M. Hutter, M. E. Marson, and P. Pearson. Using Bleichenbacher’s solution to the hidden number problem to attack nonce leaks in 384-bit ECDSA. In G. Bertoni and J.-S. Coron, editors, *CHES 2013*, volume 8086 of *LNCS*, pages 435–452. Springer, Heidelberg, Aug. 2013. 8
- [25] B. Preneel, R. Govaerts, and J. Vandewalle. Hash functions based on block ciphers: A synthetic approach. In D. R. Stinson, editor, *CRYPTO’93*, volume 773 of *LNCS*, pages 368–378. Springer, Heidelberg, Aug. 1994. 7, 30
- [26] S. Tessaro. Optimally secure block ciphers from ideal primitives. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 437–462. Springer, Heidelberg, Nov. / Dec. 2015. 5

A Comparing speed of different hash-based MACs

This appendix quantifies the speed advantage of AMAC over HMAC for short messages.

CONTEXT. We note at the outset that applications concerned purely with hashing speed should use neither AMAC nor HMAC: non-hash-based MACs are faster. However, non-hash-based MACs cost extra code size, while hash-based MACs have the advantage of reusing hash implementations. Standards use HMAC rather than NMAC for a similar reason: an HMAC implementation can treat the entire hash function as a black box.

SHA-512 SPEED FOR LONG MESSAGES. For concreteness we focus on the common Intel Haswell line of CPUs, and we focus on SHA-512 as the underlying hash function. Despite its very high target security level, SHA-512 is the fastest standard hash function on Haswell, running at 8 cycles/byte. For comparison, SHA3-256 uses nearly 9 cycles/byte, SHA-256 uses more than 11 cycles/byte, and SHA3-512 uses more than 16 cycles/byte.

The basic reason that SHA-512 outperforms SHA-256 here is that SHA-512 compression handles 128 message bytes, while SHA-256 compression handles only 64 message bytes. SHA-512 compression has only 25% more operations than SHA-256 compression; each operation inside SHA-512 compression handles 64 bits rather than 32 bits, and most of these 64-bit operations run as quickly as 32-bit operations on these CPUs.

SHA-512 SPEED FOR SHORT MESSAGES. An “8 cycles/byte” statement for SHA-512 is obtained from observing that, e.g., hashing a 2048-byte message takes $8 \cdot 1024$ cycles more than hashing a 1024-byte message. This underestimates the cost of hashing short messages. Specifically, an n -byte message actually produces $\lceil (n + 17)/128 \rceil$ compression-function calls. Benchmarks show a constant per-hash-call overhead of approximately 256 cycles, for a total cost of approximately $256 + 1024 \lceil (n + 17)/128 \rceil$ cycles.

SPEED OF SHA-512-BASED MACs. Both AMAC and HMAC take the same 8 cycles/byte for long messages. We now extrapolate from the observed speed of SHA-512 to predict the AMAC overhead and the HMAC overhead for short messages.

The definition of AMAC includes a computation of `Out`: e.g., the specific `Out` in [7] maps a 512-bit integer h to $h \bmod \ell$ for a standard prime $\ell \approx 2^{252}$. We have checked that existing software for this (`sc25519_barrett` in `ed25519/amd64-51`) takes only 100 cycles. More to the point, if HMAC-SHA-512 were used in [7] then one would also want to reduce its 512-bit output modulo ℓ . In other words, the `Out` overhead exists anyway; AMAC takes advantage of this by integrating `Out` into the MAC definition. From now on we ignore the cost of `Out`.

The most obvious advantage of AMAC over HMAC is that it hashes fewer bytes of data. If the key has 32 bytes and the message has n bytes then AMAC hashes $n + 32$ bytes, taking approximately

$256 + 1024\lceil(n + 49)/128\rceil$ cycles; e.g., approximately 1280 cycles for $n \leq 79$, and approximately 9472 cycles for $n = 1024$. For HMAC there are two hashing layers:

- The first hashing layer expands the key to a 128-byte block and hashes this block together with the message. If the first compression call is precomputed then this hashes n bytes, taking approximately $256 + 1024\lceil(n + 17)/128\rceil$ cycles. If the hash function is instead called as a black box then this instead hashes $n + 128$ bytes, taking approximately $1280 + 1024\lceil(n + 17)/128\rceil$ cycles.
- The second hashing layer expands the key to another 128-byte block and hashes this block together with the 64-byte output of the first hashing layer. With precomputation this hashes 64 bytes, taking approximately 1280 cycles. Without precomputation this hashes 192 bytes, taking approximately 2304 cycles.

Overall HMAC takes approximately $2560 + 1024\lceil(n + 17)/128\rceil$ cycles if 128 bytes of compression outputs are precomputed, and $4608 + 1024\lceil(n + 17)/128\rceil$ cycles if not; e.g., for $n \leq 111$, approximately 3584 cycles with precomputation, or 5632 cycles without; for $n = 1024$, approximately 11776 cycles with precomputation, or 13824 cycles without.

In other words: The time for AMAC is approximately the time for $0.25 + \lceil(n + 49)/128\rceil$ compression-function calls. The time for HMAC is approximately the time for $2.5 + \lceil(n + 17)/128\rceil$ compression-function calls with precomputation, or $4.5 + \lceil(n + 17)/128\rceil$ compression-function calls without.