

Generic Security of NMAC and HMAC with Input Whitening

Peter Gazi¹, Krzysztof Pietrzak¹, and Stefano Tessaro²

¹ IST Austria

{peter.gazi,pietrzak}@ist.ac.at

² UC Santa Barbara

tessaro@cs.ucsb.edu

Abstract. HMAC and its variant NMAC are the most popular approaches to deriving a MAC (and more generally, a PRF) from a cryptographic hash function. Despite nearly two decades of research, their exact security still remains far from understood in many different contexts. Indeed, recent works have re-surfaced interest for *generic* attacks, i.e., attacks that treat the compression function of the underlying hash function as a black box.

Generic security can be proved in a model where the underlying compression function is modeled as a random function – yet, to date, the question of proving tight, non-trivial bounds on the generic security of HMAC/NMAC even as a PRF remains a challenging open question.

In this paper, we ask the question of whether a small modification to HMAC and NMAC can allow us to exactly characterize the security of the resulting constructions, while only incurring little penalty with respect to efficiency. To this end, we present simple variants of NMAC and HMAC, for which we prove tight bounds on the generic PRF security, expressed in terms of numbers of construction and compression function queries necessary to break the construction. All of our constructions are obtained via a (near) *black-box* modification of NMAC and HMAC, which can be interpreted as an initial step of key-dependent message pre-processing.

While our focus is on PRF security, a further attractive feature of our new constructions is that they clearly defeat all recent generic attacks against properties such as state recovery and universal forgery. These exploit properties of the so-called “functional graph” which are not directly accessible in our new constructions.

Keywords. message authentication codes, HMAC, generic attacks, provable security

1 Introduction

This paper presents new variants of the HMAC/NMAC constructions of message authentication codes which enjoy *provable* security as a pseudorandom function (PRF) against generic distinguishing attacks, i.e., attacks which treat the compression function of the underlying hash function as a black-box. In particular, we prove concrete *tight* bounds in terms of the number of queries to the construction *and* to the compression function necessary to distinguishing our construction from a random function. Our constructions are the first HMAC/NMAC variants to enjoy such a tight analysis, and we see this as an important stepping stone towards the understanding of the generic security of hash-based message authentication codes.

HASH-BASED MACs. HMAC [3] is the most widely used approach to key a hash function H to obtain a PRF or a MAC. It computes the output on message M and a key K as

$$\text{HMAC}(K, M) = H(K \oplus \text{opad} \parallel H(K \oplus \text{ipad} \parallel M)),$$

where $\text{opad} \neq \text{ipad}$ are constants.³ Usually, H is a hash function like SHA-1, SHA-256 or MD5, in particular following the Merkle-Damgård paradigm [17,5]. That is, it extends a compression function $f : \{0,1\}^c \times \{0,1\}^b \rightarrow \{0,1\}^c$ into a hash function MD_V^f by first padding M into b -bit blocks $M[1], \dots, M[\ell]$, and then producing the output $H(M) = S_\ell$, where

$$S_0 \leftarrow IV, \quad S_i \leftarrow f(S_{i-1} \parallel M[i]) \text{ for all } i = 1, \dots, \ell. \quad (1)$$

³ Some details such as padding and arbitrary key length are addressed in Section 2.

starting with the c -bit initialization value IV. A cleaner yet slightly less practical variant of HMAC is NMAC, which instead outputs

$$\text{NMAC}_{K_{\text{in}}, K_{\text{out}}}(M) = \text{MD}_{K_{\text{out}}}^f(\text{MD}_{K_{\text{in}}}^f(M)),$$

where $K_{\text{in}}, K_{\text{out}} \in \{0, 1\}^c$ are key values.

SECURITY OF HMAC/NMAC. The security of both constructions has been studied extensively, both by obtaining security proofs and proposing attacks. On the former side, NMAC and HMAC were proven to be secure *pseudorandom functions* (PRFs) in the standard model [3], later also using weaker assumptions [2] and via a tight bound in the uniform setting [8]. However, as argued in [8], this standard-model bound might be overly pessimistic, covering also very unnatural constructions of the underlying compression function f (for example the one used in their tightness proof). The authors hence argue for the need of an analysis of the PRF security of HMAC in the so-called *ideal compression function model* where the compression function is modelled as an ideal random function and the adversary is allowed to query it. This model was previously used by Dodis *et al.* [7] to study *indifferentiability* of HMAC, which however only holds for certain key lengths.

This is also the model implicitly underlying many of the recently proposed attacks on hash-based MACs [20,18,16,21,11,6,23]. These attacks are termed *generic*, meaning they can be mounted for any underlying hash function as long as it follows the Merkle-Damgård (MD) paradigm. The complexity of such a generic attack is then expressed in the number of key-dependent queries to the construction (denoted q_C) as well as the number of queries to the underlying compression function (denoted q_f). These two classes of queries are also often referred to as *online* and *offline*, respectively.

All iterated MACs are subject to the long-known Preneel and van Oorschot’s attack [22] which implies a forgery (and hence also distinguishing) attack against HMAC/NMAC making $q_C = 2^{c/2}$ construction queries (consisting of constant-length messages) and no direct compression function queries (i.e., $q_f = 0$). This immediately raises two questions:

How does the security of HMAC and NMAC degrade (in terms of tolerable q_C) by increasing (1) the length ℓ of the messages and (2) the number q_f of compression-function evaluations?

The first question has been partially addressed in [8]. Their result⁴ can be interpreted as giving tight bounds on the PRF security of NMAC against an attacker making q_C key-dependent construction queries (of length at most $\ell < 2^{c/3}$ b -bit blocks) but *no* queries to the compression function. They show that both constructions can only be distinguished from random function with advantage roughly $\epsilon(q_C, \ell) \approx \ell^{1+o(1)} q_C^2 / 2^c$, improving significantly on the bound $\epsilon(q_C, \ell) \approx \ell^2 q_C^2 / 2^c$ provable using standard folklore techniques. From our perspective, this bound can be read as a smooth trade-off: with increasing maximum allowed query length ℓ it tells us how many queries q_C can be tolerated for any acceptable upper bound on advantage.

Still, it is not clear how this trade-off changes when allowing extremely long messages ($\ell > 2^{c/3}$) and/or some queries to the compression function ($q_f > 0$). Note that while huge ℓ can be prevented by standards, in practical settings q_f is very likely to be much higher than q_C , as it represents cheap local (offline) computation of the attacker. We therefore focus on capturing the trade-off between q_C and q_f for values of q_C that do not allow to mount the attack from [22]. However, as we argue below, getting such a tight trade-off for NMAC/HMAC seems to be out of reach for now, we hence relax the problem by allowing for slight modifications to the vanilla NMAC/HMAC construction.

OUR CONTRIBUTIONS. We ask the following question here, and answer it positively:

Can we devise variants of HMAC/NMAC whose security provably degrades gracefully with an increasing number of compression function queries q_f , possibly retaining security for q_f being much larger than 2^c ?

The main contribution of this paper is the introduction and analysis of a variant of NMAC (which we then adapt to the HMAC setting, as described below) which uses additional key material to “whiten”

⁴ Here we refer to Theorem 2 in [8] that formally considers a related construction NI in the standard model. However, its proof starts by a transition to the ideal-model analysis of a construction very closely related to NMAC, while disallowing compression-function queries.

message blocks before being processed by the compression function. Concretely, our construction – termed WNMAC (for “whitened NMAC”) uses an additional extra b -bit key K_w , and given a message M padded as $M[1], \dots, M[\ell]$, operates as NMAC on input padded to blocks $M'[i] = M[i] \oplus K_b$, i.e., every message block is whitened with the *same* key (see also Fig. 1).

The rationale behind WNMAC is two-fold. First, from the security viewpoint, the justification comes from the rich line of research on generic attacks on hash-based MACs. Most recent attacks [20,16,21,11] exploit the so-called “functional graph” of the compression function f , i.e., the graph capturing the structure of f when repeatedly invoked with its b -bit input fixed to some constant (say 0^b). Since our whitening denies the adversary the knowledge of b -bit inputs on which f is invoked during construction queries, intuitively it seems to be the right way to foil such attacks. Moreover, a recent work by Sasaki and Wang [23] suggests that keying *every* invocation of f is necessary in order to prevent suboptimal security against generic state recovery attacks. WNMAC arguably provides the simplest and most natural such keying. Second, from the practical perspective, WNMAC can be implemented on top of an existing implementation of NMAC, using it as a black-box.

PRF-SECURITY OF WNMAC. Our main result shows that WNMAC is a secure PRF; more precisely, no attacker making at most q_C construction queries (for messages padded into at most ℓ blocks) and q_f primitive queries can distinguish WNMAC from a random function, except with distinguishing advantage

$$\epsilon_{\text{WNMAC}}(q_C, q_f, \ell) \leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 1 \right).$$

Here, $d'(\ell)$ is the maximum, over all positive integers $\ell' \leq \ell$, of the number of positive divisors of ℓ' , and grows very slowly, i.e., $d'(\ell) \approx \ell^{1/\ln \ln \ell}$. We also prove that this bound is essentially tight. Namely, we give an attack that achieves advantage roughly $q_C q_f / 2^{2c}$, showing the first term above to be necessary. Additionally, we know from [8] that the third term is tight for $\ell \leq 2^{c/3}$.

Note that in the case of $q_f = 0$, the bound matches exactly the bound from [8]. Moreover, observe that under the realistic assumption that $\ell < \min\{2^{c/3}, 2^{b-c}\}$, the bound simplifies to

$$\epsilon_{\text{WNMAC}}(q_C, q_f, \ell) \leq 3 \frac{q_f q_C}{2^{2c}} + (d'(\ell) + 2) \cdot \frac{\ell q_C^2}{2^c}.$$

Ignoring $d'(\ell)$ for simplicity, we see that we can tolerate up to $q_C \approx 2^{c/2}/\sqrt{\ell}$ construction queries and up to $q_f \approx 2^{1.5c}$ primitive queries. This corresponds to the security threshold ranging from 2^{192} f -queries for MD5 up to 2^{768} f -queries for SHA-512. The first term also clearly characterizes the complete trade-off curve between $q_C < 2^{c/2}/\sqrt{\ell}$ and q_f for any reasonable upper bound on the message length and acceptable distinguishing advantage.

OTHER SECURITY PROPERTIES. Additionally, we also analyze the security level WNMAC achieves with respect to other security notions frequently considered in the attacks literature. By a series of reductions, we show that, roughly speaking, ϵ_{WNMAC} also upper-bounds the adversary’s advantage for *distinguishing- H* and *state recovery*. We believe that addressing these cryptanalytic notions also using the traditional toolbox of provable security is important and see this paper as taking the first step on that path.

LIFTING TO HMAC. We then move our attention from NMAC to HMAC and propose two analogous modifications to it. The first one, called WHMAC, is obtained from HMAC in the same way WNMAC is obtained from NMAC: by whitening the padded message blocks with an independent key (see Fig. 5). The second one, termed WHMAC⁺, additionally processes a fresh key K^+ instead of the first block of the message (see Fig. 6). Both variants can be implemented given only black-box access to HMAC, and we prove that they maintain the same security level as WNMAC as long as the parameters b, c of f satisfy $b \gg 2c$ (for WHMAC) or $b \gg c$ (for WHMAC⁺). Note that for existing hash functions, the former condition is satisfied for both MD5 and SHA-1, while the latter holds also for SHA-256 and SHA-512.

THE DUAL CONSTRUCTION. Motivated by the most restrictive term $q_C q_f / 2^{2c}$ in ϵ_{WNMAC} , the final construction we propose in this paper is a “dual” version of WNMAC denoted DWNMAC, that differs in the final, outer f -call. Instead of $f(K_2, s \parallel 0^{b-c})$ for a c -bit key K_2 and a c -bit state s padded with zeroes, the outer

call in DWNMAC computes $f(s, K_2)$ for a longer, b -bit key. As expected, we prove that this tweak removes the need for the $q_C q_f / 2^{2c}$ term and replaces it by the strictly favourable term $q_C q_f / 2^{b+c}$, proving that the zero-padding in the outer call of WNMAC was actually responsible for the “bottle-neck” term in its security bound.

OUR TECHNIQUES. In our information-theoretic analysis of WNMAC we employ the H-coefficient technique by Patarin [19], partially inheriting the notational framework from the recent analysis of keyed sponges by Gaži, Pietrzak, and Tessaro [9]. On a high level, the heart of our proof is a careful analysis of the probability that two sets intersect in the ideal experiment: (1) the set of adversarial queries to f , and (2) the set of inputs on which f is invoked when answering the adversary’s queries to WNMAC. Obtaining a bound on the probability of this event then allows us to exclude it and use the result from [8] that considers $q_f = 0$, properly adapted to the WNMAC setting.

RELATED WORK. As mentioned above, the motivation for our work partially stems from the recent line of work on generic attacks against iterated hash-based MACs [20,18,16,21,11,6,23]. While our security bound for WNMAC does not exclude attacks of the complexity (in terms of numbers of queries and message lengths) considered in these papers, the design of WNMAC was partially guided by the structure of these attacks and seems to prevent them. We find in particular the work [23] to be a good justification for investigating the security of WNMAC and related constructions. Iterated MAC that uses keying in every f -invocation was already considered by An and Bellare [1], their construction NI was later subject to analysis [8] that we adapt and reuse. One can see WNMAC as a conceptual simplification of NI where the key is simply used to whiten the b -bit input to the compression function. Finally, our dual construction considered in Section 5 bears resemblance to the Sandwich MAC analyzed by Yasuda [24], we believe that our methods could be easily adapted to cover this construction as well.

PERSPECTIVE AND OPEN PROBLEMS. We stress that the reader should not conclude from this work that NMAC and HMAC are necessarily less secure than the constructions proposed in this paper, specifically with respect to PRF security. In fact, we are not aware of any attacks showing a separation between the PRF security of our constructions and that of the original NMAC/HMAC constructions, finding one is an interesting open problem.

While obtaining a non-tight birthday-type bound for NMAC/HMAC is feasible (for most key-length values, a bound follow directly from the indifferentiability analysis of [7]), proving *tight* bounds in terms of compression function and construction queries on the generic PRF security of NMAC/HMAC is a challenging open problem, on which little progress has been made. The main challenge is to understand how partial information in form of f -queries can help the attacker to break security (i.e., distinguish) in settings with $q_C \ll 2^{c/2} / \sqrt{\ell}$, when the attack from [8] does not apply. This will require in particular developing a better understanding of the functional graph defined by queries to the function f . Some of its properties have been indeed exploited in existing generic attacks, but proving security appears to require a much deeper understanding: Most of the recent attacks, which are probably still not tight, do not come with rigorous proofs but instead rely on conjectures on the structure of these graphs [11]. The difficulty of this question for NMAC/HMAC is also well documented by the fact that even proving security of the whitened constructions presented in this paper required some novel tricks and considerable effort.

Similarly, it remains equally challenging to prove that for the properties considered by the recent attacks against HMAC/NMAC (such as distinguishing-H, state recovery or various types of forgeries), the security of WNMAC/WHMAC is provably superior. Yet, we note that our construction invalidates direct application of all existing attacks, and hence we feel confident conjecturing that its security is much higher.

BLACK-BOX INSTANTIATIONS. Throughout the paper we implicitly assume we can add a key to each b -bit input block, even though we aim for a black-box instantiation. For many MD-based hash functions, such fine-grained control of the input to the compression function is generally not possible via a black-box message pre-processing. Concretely, the functions from the SHA-family with 512-bit blocks only allow to effectively control (via alterations of the message) the first 447 bits of the last block, since the remaining 65 bits are reserved for the 64-bit length, and an additional 1-bit. Our analysis can be easily modified to take this into account. The resulting bound will change very little, and will result in the term $\ell q_C q_f / 2^{b+c}$ being replaced by the term $(\ell - 1 + 2^d) \cdot q_C \cdot q_f / 2^{b+c}$, where d is the length of the non-controllable part of the input (for

SHA-functions, $d = 65$). Note that since $d \ll b - c$, this will not affect the tightness of the bounds for concrete parameters.

2 Preliminaries

BASIC NOTATION. We denote $[n] := \{1, \dots, n\}$. Moreover, for a finite set \mathcal{S} (e.g., $\mathcal{S} = \{0, 1\}$), we let \mathcal{S}^n , \mathcal{S}^+ and \mathcal{S}^* be the sets of sequences of elements of \mathcal{S} of length n , of arbitrary (but non-zero) length, and of arbitrary length, respectively (with ε denoting the empty sequence, as opposed to ϵ which is a small quantity). As a shorthand, let $\{0, 1\}^{b*}$ denote $(\{0, 1\}^b)^*$. We denote by $S[i]$ the i -th element of $S \in \mathcal{S}^n$ for all $i \in [n]$. Similarly, we denote by $S[i \dots j]$, for every $1 \leq i \leq j \leq n$, the sub-sequence consisting of $S[i], S[i+1], \dots, S[j]$, with the convention that $S[i \dots i] = S[i]$. Moreover, we denote by $S \parallel S'$ the concatenation of two sequences in \mathcal{S}^* , and also, we let $S \mid T$ be the usual prefix-of relation: $S \mid T \Leftrightarrow (\exists S' \in \mathcal{S}^* : S \parallel S' = T)$.

For an integer n , $d(n) = |\{i \in \mathbb{N} : i \mid n\}|$ is the number of its positive divisors and

$$d'(n) := \max_{n' \in \{1, \dots, n\}} |\{d \in \mathbb{N} : d \mid n'\}| \approx n^{1/\ln \ln n}$$

is the maximum, over all positive integers $n' \leq n$, of the number of positive divisors of n' . More precisely, we have $\forall \varepsilon > 0 \exists n_0 \forall n > n_0 : d(n) < n^{(1+\varepsilon)/\ln \ln n}$ [12].

We also let $\mathcal{F}(\mathcal{D}, \mathcal{R})$ be the set of all functions from \mathcal{D} to \mathcal{R} ; and with a slight abuse of notation we sometimes write $\mathcal{F}(m, n)$ (resp. $\mathcal{F}(*, n)$) to denote the set of functions mapping m -bit strings to n -bit strings (resp. from $\{0, 1\}^*$ to $\{0, 1\}^n$). We denote by $x \stackrel{\$}{\leftarrow} \mathcal{X}$ the act of sampling x uniformly at random from \mathcal{X} . Finally, we denote the event that an adversary \mathbf{A} , given access to an oracle \mathbf{O} , outputs a value y , as $\mathbf{A}^{\mathbf{O}} \Rightarrow y$. To emphasize the random experiment considered, we sometimes denote the probability of an event A in a random experiment \mathbf{E} by $\mathbf{P}^{\mathbf{E}}[A]$. Finally, the min-entropy $H_\infty(X)$ of a random variable X with range \mathcal{X} is defined as $-\log(\max_{x \in \mathcal{X}} \mathbf{P}_X(x))$.

PSEUDORANDOM FUNCTIONS. We consider *keyed* functions $\mathbf{F} : \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}$ taking a κ -bit key (i.e., $\mathcal{K} = \{0, 1\}^\kappa$), a message $M \in \mathcal{D}$ as input, and returning an output from \mathcal{R} . For a keyed function \mathbf{F} under a key $k \in \mathcal{K}$ we often write $\mathbf{F}_k(\cdot)$ instead of $\mathbf{F}(k, \cdot)$. One often considers the security of \mathbf{F} as a *pseudorandom function* (or PRF, for short) [10]. This is defined via the following advantage measure, involving an adversary \mathbf{A} :

$$\text{Adv}_{\mathbf{F}}^{\text{prf}}(\mathbf{A}) := \left| \mathbf{P} \left[K \stackrel{\$}{\leftarrow} \{0, 1\}^\kappa : \mathbf{A}^{\mathbf{F}^K} \Rightarrow 1 \right] - \mathbf{P} \left[f \stackrel{\$}{\leftarrow} \mathcal{F}(\mathcal{D}, \mathcal{R}) : \mathbf{A}^f \Rightarrow 1 \right] \right|.$$

Informally, we say that \mathbf{F} is a PRF if this advantage is “negligible” for all “efficient” adversaries \mathbf{A} .

PRFS IN THE IDEAL COMPRESSION FUNCTION MODEL. For our analysis below, we are going to consider keyed constructions $\mathbf{C}[f] : \{0, 1\}^\kappa \times \mathcal{D} \rightarrow \mathcal{R}$ which make queries to a randomly chosen compression function $f \stackrel{\$}{\leftarrow} \mathcal{F}(c + b, c)$ which can also be evaluated by the adversary (we sometimes write \mathbf{C}^f instead of $\mathbf{C}[f]$). For this case, we use the following notation to express the PRF advantage of \mathbf{A} :

$$\text{Adv}_{\mathbf{C}[f]}^{\text{prf}}(\mathbf{A}) := \left| \mathbf{P} \left[K \stackrel{\$}{\leftarrow} \{0, 1\}^\kappa, f \stackrel{\$}{\leftarrow} \mathcal{F}(c + b, c) : \mathbf{A}^{\mathbf{C}_K^f} \Rightarrow 1 \right] - \mathbf{P} \left[\mathbf{R} \stackrel{\$}{\leftarrow} \mathcal{F}(\mathcal{D}, \mathcal{R}), f \stackrel{\$}{\leftarrow} \mathcal{F}(c + b, c) : \mathbf{A}^{\mathbf{R}, f} \Rightarrow 1 \right] \right|.$$

We call \mathbf{A} 's queries to its first oracle *construction queries* (or C-queries) and its queries to the second oracle as *primitive queries* (or f-queries).

Note that the notion of PRF-security is identical to the notion of *distinguishing-R*, first defined in [14] and often used in the cryptanalytic literature on hash-based MACs.

DISTINGUISHING-H. A further security notion defined in [14] is the so-called *distinguishing-H* security. Here, the goal of the adversary is to distinguish the hash-based MAC construction $\mathbf{C}_K[f]$ using its underlying compression function f (say SHA-1) and a random key K , from the same construction $\mathbf{C}_K[g]$ built on top of an independent random compression function g . In the ideal compression function model, where we model

already the initial compression function f as ideal, this corresponds to distinguishing a pair of oracles $(C_K[f], f)$ from $(C_K[f], g)$. Formally,

$$\text{Adv}_C^{\text{dist-H}}(A) := \left| \mathbb{P} \left[K \xleftarrow{\$} \{0, 1\}^\kappa, f \xleftarrow{\$} \mathcal{F}(c+b, c) : A^{C_K^f, f} \Rightarrow 1 \right] - \mathbb{P} \left[K \xleftarrow{\$} \{0, 1\}^\kappa, f, g \xleftarrow{\$} \mathcal{F}(c+b, c) : A^{C_K^f, g} \Rightarrow 1 \right] \right|.$$

STATE RECOVERY. An additional notion considered in the literature is security against *state recovery*. Since the definition of this notion needs to be tailored for the concrete construction it is applied to, we postpone the formal definition of security against state recovery to Section 3.10.

MACS AND UNPREDICTABILITY. It is well known that a good PRF also yields a good message-authentication code (MAC). A concrete security bound for unforgeability can be obtained from the PRF bound via a standard argument.

ITERATED MACS. For a keyed function $f : \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ we denote with $\text{Casc}^f : \{0, 1\}^c \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ the cascade construction (also known as Merkle-Damgård) built from f as

$$\text{Casc}^f(K, m_1 \parallel \dots \parallel m_\ell) := y_\ell \text{ where } y_0 := K \text{ and for } i \geq 1 : y_i := f(y_{i-1}, m_i),$$

in particular $\text{Casc}^f(K, \varepsilon) := K$.

The construction $\text{NMAC}^f : (\{0, 1\}^c)^2 \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ is derived from Casc^f by adding an additional, independently keyed application of f at the end. It assumes that the domain sizes of f satisfy $b \geq c$ and the output of the cascade is padded with zeroes before the last f -call. Formally,

$$\text{NMAC}^f((K_1, K_2), M) := f(K_2, \text{Casc}^f(K_1, M) \parallel 0^{b-c}).$$

Note that practical MD-based hash functions take as input arbitrary-length bitstrings and then pad them to a multiple of the block length, often including the message length in the so-called MD-strengthening. This padding then also appears in NMAC (and HMAC) but here we take the customary shortcut and our definition of NMAC above (resp. HMAC below) actually corresponds to the generalized constructions denoted as GNMAC (resp. GHMAC) in [2] where this step is also justified in detail.

HMAC^f is a practice-oriented version of NMAC^f, where the two keys (K_1, K_2) are derived from a single key $K \in \{0, 1\}^b$ by xor-ing it with two fixed b -bit strings *ipad* and *opad*. In addition, the keys are not given through the key-input of the compression function f , but are prepended to the message instead. This allows for the usage of existing implementations of hash functions that contain a hard-coded initialization vector *IV*. Formally:

$$\begin{aligned} \text{HMAC}^f(K, m) &:= \text{Casc}^f(\text{IV}, K_2 \parallel \text{Casc}^f(\text{IV}, K_1 \parallel m) \parallel \text{fpad}) \\ &\text{where } (K_1, K_2) := (K \oplus \text{ipad}, K \oplus \text{opad}) \end{aligned}$$

and *fpad* is a fixed $(b-c)$ -bit padding not affecting the security analysis. (Technically, [15] allows for arbitrary length of the key K : a key shorter than b bits is padded with zeroes before applying the xor transformations, a longer key is first hashed.)

3 The Whitened NMAC Construction

We now present our main construction called *Whitened NMAC* (or *WNMAC* for short). To that end, let us first consider a modification of the cascade construction Casc called *whitened cascade* and denoted WCasc . For a keyed function $f : \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ we denote with $\text{WCasc}^f : (\{0, 1\}^c \times \{0, 1\}^b) \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ the whitened cascade construction built from f as

$$\begin{aligned} \text{WCasc}^f((K_1, K_w), m_1 \parallel \dots \parallel m_\ell) &:= y_\ell \\ \text{where } y_0 &:= K_1 \text{ and for } i \geq 1 : y_i := f(y_{i-1}, m_i \oplus K_w), \end{aligned}$$




Fig. 1. The construction $\text{WNMAC}[f]_{K_1, K_2, K_w}$.

in particular $\text{WCasc}^f((K_1, K_w), \varepsilon) := K_1$.

The construction WNMAC is derived from NMAC , the only difference being that the inner cascade Casc is replaced by the whitened cascade WCasc . More precisely,

$$\text{WNMAC}^f((K_1, K_2, K_w), M) := f(K_2, \text{WCasc}^f((K_1, K_w), M) || 0^{b-c}).$$

For a graphical depiction of WNMAC , see Figure 1. We devote most of this section to the proof of the following theorem that quantifies the PRF-security of WNMAC .

Theorem 1 (PRF-Security of WNMAC). *Let A be an adversary making at most q_f queries to the compression function f and at most q_C construction queries, each of length at most ℓ b -bit blocks. Let $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ be a tuple of random keys. Then we have*

$$\text{Adv}_{\text{WNMAC}_K}^{\text{prf}}(A) \leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 1 \right). \quad (2)$$

Note that as observed in Section 2, this also covers the so-called distinguishing-R security of WNMAC . Moreover, our analysis also implies security bounds for distinguishing-H and state recovery, as we discuss later.

3.1 Basic Notation, Message Trees and Repetition Patterns

Let us fix an adversary A . We assume that A is deterministic, it makes *exactly* q_f queries to f and q_C construction queries, and it never repeats the same query twice. All these assumptions are without loss of generality for an information-theoretic indistinguishability analysis, since an arbitrary (possibly randomized) adversary making at most this many queries can be transformed into one satisfying the above constraints and achieving advantage which is at least as large.

Let $\mathcal{Q}_C \subseteq (\{0, 1\}^b)^*$ be any non-empty set of messages (later this will represent the set of A 's C-queries). Based on it, we now introduce the *message tree* and its labeled version, which capture the inherent combinatorial structure of the messages \mathcal{Q}_C , as well as the internal values computed while these messages are processed by WCasc^f inside of WNMAC^f . The message tree $T(\mathcal{Q}_C) = (V, E)$ for \mathcal{Q}_C is defined as follows:

- The vertex set is $V := \left\{ M' \in (\{0, 1\}^b)^* : \exists M \in \mathcal{Q}_C : M' \mid M \right\}$, where \mid is the prefix-of partial ordering of strings. In particular, note that the empty string ε is a vertex and that $\mathcal{Q}_C \subseteq V$.
- The set $E \subseteq V \times V$ of (directed) edges is

$$E := \left\{ (M, M') : \exists m \in \{0, 1\}^b : M' = M \parallel m \right\}.$$

To simplify our exposition, we also define the following two mappings based on $T(\mathcal{Q}_C)$.

- The mapping $\pi(v) : V \setminus \{\varepsilon\} \rightarrow V$ returns the unique parent node of $v \in V \setminus \{\varepsilon\}$; i.e., the unique node u such that $(u, v) \in E$.




Fig. 2. Labeled message tree. Example of a labeled message tree $T_K^f(\mathcal{Q}_C)$ for four messages $\mathcal{Q}_C = \{0, 0 \parallel 0, 0 \parallel 1 \parallel 1, 1\}$, where $r = r^b$ for $r \in \{0, 1\}$. The gray vertices correspond to these four messages. Next to each vertex v and edge (u, v) , we give the label $\lambda(v)$ and the value $\mu(v)$, respectively.

- The mapping $\mu(v): V \setminus \{\varepsilon\} \rightarrow \{0, 1\}^b$ returns the unique message block $m \in \{0, 1\}^b$ such that $\pi(v) \parallel \mu(v) = v$ (intuitively, this will be the message block that is processed when “arriving” in vertex v).

Alternatively, with a slight abuse of notation we will also refer to the vertices in V as $v_1, \dots, v_{|V|}$ which is an arbitrary ordering of them such that for all $1 \leq i, j \leq |V|$ it satisfies $v_i \mid v_j \Rightarrow i \leq j$. Note that one obtains such an ordering for example if one, intuitively speaking, processes the messages in \mathcal{Q}_C block-wise and labels the vertices by their “first appearance”: in particular $v_1 = \varepsilon$ is the tree root.

Additionally, for a mapping $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and a key tuple $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ we also consider an extended version of $T(\mathcal{Q}_C)$ which we call the *labeled message tree* and denote $T_K^f(\mathcal{Q}_C) = (V, E, \lambda)$, and which is defined as follows:

- The set of vertices V and edges E are defined exactly as for $T(\mathcal{Q}_C)$ above.
- The vertex-labeling function $\lambda: V \rightarrow \{0, 1\}^c$ is defined iteratively: $\lambda(\varepsilon) := K_1$ and for each non-root vertex $v \in V \setminus \{\varepsilon\}$ we put $\lambda(v) := f(\lambda(\pi(v)), \mu(v) \oplus K_w)$.

An example of a labeled message tree is given in Figure 2. Note that each vertex label $\lambda(v)$ is exactly the output of the inner, whitened cascade $\text{WCasc}_{K_1, K_w}^f(v)$ in WNMAC_K^f (recall that v is actually a message from $\{0, 1\}^{b^*}$).

For any message tree $T(\mathcal{Q}_C) = (V, E)$, a *repetition pattern* is any equivalence relation ρ on V . For a labeled message tree $T_K^f(\mathcal{Q}_C) = (V, E, \lambda)$ we say that a repetition pattern ρ is *induced* by it if it satisfies

$$\forall u, v \in V : \lambda(u) = \lambda(v) \Leftrightarrow \rho(u, v) .$$

3.2 Interactions and Transcripts

Let \mathcal{QR}_C denote the set of q_C pairs (x, r) such that $x \in \{0, 1\}^{b^*}$ is a construction query and $r \in \{0, 1\}^c$ is a potential response to it (what we mean by “potential” will be clear from below). Similarly let \mathcal{QR}_f denote the set of q_f pairs (x, r) such that $x \in \{0, 1\}^c \times \{0, 1\}^b$ is an f -query and $r \in \{0, 1\}^c$ is a potential response to it. Let $\mathcal{Q}_C \subseteq \{0, 1\}^{b^*}$ and $\mathcal{Q}_f \subseteq \{0, 1\}^c \times \{0, 1\}^b$ denote the sets of first coordinates (i.e., the queries) in \mathcal{QR}_C and \mathcal{QR}_f , respectively; we have $|\mathcal{Q}_C| = q_C$ and $|\mathcal{Q}_f| = q_f$.

We call the pair of sets $(\mathcal{QR}_C, \mathcal{QR}_f)$ *valid* if the adversary A would indeed ask these queries throughout the experiment, assuming that each of her queries would be replied by the respective response in \mathcal{QR}_C or \mathcal{QR}_f (note that once a deterministic A is fixed, this determines whether a given pair $(\mathcal{QR}_C, \mathcal{QR}_f)$ is valid).

We then define a *valid transcript* to be of the form

$$\tau = (\mathcal{QR}_C, \mathcal{QR}_f, K = (K_1, K_2, K_w), T_K^f(\mathcal{Q}_C)) ,$$

where $(\mathcal{QR}_C, \mathcal{QR}_f)$ is valid, $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ is a function and $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ is a key tuple.

We differentiate between the ways in which such valid transcripts are generated in the real and in the ideal worlds (or experiments), respectively, by defining corresponding distributions \mathbb{T}_{real} and $\mathbb{T}_{\text{ideal}}$ over the set of valid transcripts:

Real world. The transcript \mathbb{T}_{real} for the adversary \mathbf{A} is obtained by sampling $f \xleftarrow{\$} \mathcal{F}(c+b, c)$ and $K = (K_1, K_2, K_w) \leftarrow \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$, and letting \mathbb{T}_{real} denote

$$(\mathcal{QR}_C = \{(M_i, Y_i)\}_{i=1}^{q_C}, \mathcal{QR}_f = \{(X_i, R_i)\}_{i=1}^{q_f}, K = (K_1, K_2, K_w), T_K^f(\mathcal{Q}_C)),$$

where we execute \mathbf{A} , which asks construction queries M_1, \dots, M_{q_C} answered with $Y_i := \text{WNMAC}[f]_K(M_i)$ for all $i \in [q_C]$; and f -queries X_1, \dots, X_{q_f} answered with $R_i := f(X_i)$ for all $i \in [q_f]$ (note that the C -queries and f -queries may in general be interleaved adaptively, depending on \mathbf{A}). Finally, we let $T_K^f(\mathcal{Q}_C)$ be the labeled message tree corresponding to \mathcal{Q}_C , f and K .

Ideal world. The transcript $\mathbb{T}_{\text{ideal}}$ for the adversary \mathbf{A} is obtained similarly to the above, but here, together with the random function $f \xleftarrow{\$} \mathcal{F}(c+b, c)$ and the key tuple $K = (K_1, K_2, K_w) \leftarrow \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$, we also sample q_C independent random values $Y_1, \dots, Y_{q_C} \in \{0, 1\}^r$. Then we let $\mathbb{T}_{\text{ideal}}$ denote

$$(\mathcal{QR}_C = \{(M_i, Y_i)\}_{i=1}^{q_C}, \mathcal{QR}_f = \{(X_i, R_i)\}_{i=1}^{q_f}, K = (K_1, K_2, K_w), T_K^f(\mathcal{Q}_C)),$$

where we execute \mathbf{A} , answer each its C -query M_i with Y_i for all $i \in [q_C]$ and each its f -query X_i with $R_i := f(X_i)$ for all $i \in [q_f]$. Then we let $T_K^f(\mathcal{Q}_C)$ be the labeled message tree corresponding to \mathcal{Q}_C , f and K .

Later we refer to the above two random experiments as *real* and *ideal*, respectively. Note that the range of \mathbb{T}_{real} is included in the range of $\mathbb{T}_{\text{ideal}}$ by definition, and that the range of $\mathbb{T}_{\text{ideal}}$ is easily seen to contain all valid transcripts.

3.3 The H-Coefficient Method

We upper-bound the advantage \mathbf{A} in distinguishing $\text{WNMAC}[f]_K$ for $f \xleftarrow{\$} \mathcal{F}(c+b, c)$ from a random function in terms of the statistical distance of the transcripts, i.e.,

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(\mathbf{A}) \leq \text{SD}(\mathbb{T}_{\text{real}}, \mathbb{T}_{\text{ideal}}) = \frac{1}{2} \sum_{\tau} |\mathbb{P}[\mathbb{T}_{\text{real}} = \tau] - \mathbb{P}[\mathbb{T}_{\text{ideal}} = \tau]|, \quad (3)$$

where the sum is over all valid transcripts. This is because an adversary for \mathbb{T}_{real} and $\mathbb{T}_{\text{ideal}}$, whose optimal advantage is exactly $\text{SD}(\mathbb{T}_{\text{real}}, \mathbb{T}_{\text{ideal}})$, can always output the same decision bit as \mathbf{A} , ignoring any extra information provided by the transcript.

We are going to use Patarin's H-coefficient method [19], as recently revisited in [4]. This means that we need to partition the set of valid transcripts into *good* transcripts GT and *bad* transcripts BT and then apply the following lemma, whose proof is given for completeness in Appendix A.

Lemma 1 (The H-Coefficient Method [19]). *Let $\delta, \epsilon \in [0, 1]$ be such that:*

- (a) $\mathbb{P}[\mathbb{T}_{\text{ideal}} \in \text{BT}] \leq \delta$.
- (b) For all $\tau \in \text{GT}$,

$$\frac{\mathbb{P}[\mathbb{T}_{\text{real}} = \tau]}{\mathbb{P}[\mathbb{T}_{\text{ideal}} = \tau]} \geq 1 - \epsilon.$$

Then,

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(\mathbf{A}) \leq \text{SD}(\mathbb{T}_{\text{real}}, \mathbb{T}_{\text{ideal}}) \leq \epsilon + \delta.$$

More verbally, we want a set of good transcripts GT such that with very high probability (i.e., $1 - \delta$) a generated transcript *in the ideal world* is going to be in this set, and moreover, for each such good transcript, the probabilities that it occurs in the real and in the ideal worlds are *roughly* the same, i.e., at most a multiplicative factor $1 - \epsilon$ apart.

3.4 Good and Bad Transcripts

Given a valid transcript τ we define the sets $\mathcal{L}_{\text{in}}, \mathcal{L}_{\text{out}} \subseteq \{0, 1\}^c \times \{0, 1\}^b$ as

$$\begin{aligned}\mathcal{L}_{\text{in}} &:= \{(\lambda(\pi(v)), \mu(v) \oplus K_w) : v \in V \setminus \{\varepsilon\}\} \\ \mathcal{L}_{\text{out}} &:= \{(K_2, \lambda(v) \parallel 0^{b-c}) : v \in \mathcal{Q}_C\},\end{aligned}$$

and let $\mathcal{L} = \mathcal{L}_{\text{in}} \cup \mathcal{L}_{\text{out}}$. Intuitively, \mathcal{L} represents the set of inputs on which f is evaluated while processing A 's construction queries in the real experiment. This set is also well-defined in the ideal experiment by the above equations, and in both experiments it is determined by the transcript. We refer to \mathcal{L}_{in} as the set of *inner f-invocations*, i.e., those invocations of f that were required to evaluate the inner, whitened cascade WCasc^f in WNMAC ; and similarly, \mathcal{L}_{out} denotes the *outer invocations*.

If there is an intersection between the adversary's f -queries and the inputs in \mathcal{L}_{in} (resp. \mathcal{L}_{out}), we call this an *inner (resp., outer) C-f-collision*. We then denote by $\text{C-f-coll}_{\text{in}}$ (resp., $\text{C-f-coll}_{\text{out}}$) the event that any inner (resp., outer) C-f-collision occurs. Formally,

$$\text{C-f-coll}_{\text{in}} := (\mathcal{Q}_f \cap \mathcal{L}_{\text{in}} \neq \emptyset) \quad \text{and} \quad \text{C-f-coll}_{\text{out}} := (\mathcal{Q}_f \cap \mathcal{L}_{\text{out}} \neq \emptyset)$$

and let $\text{C-f-coll} := \text{C-f-coll}_{\text{in}} \cup \text{C-f-coll}_{\text{out}}$. Furthermore, if the vertex labels $\lambda(M)$ collide for two messages $M, M' \in \mathcal{Q}_C$, we call this a C-collision and denote such an event by

$$\text{C-coll} := (\exists M, M' \in \mathcal{Q}_C : \lambda(M) = \lambda(M')) .$$

Definition 1 (Good Transcripts). *Let*

$$\tau = (\mathcal{QR}_C, \mathcal{QR}_f, K = (K_1, K_2, K_w), T_K^f(\mathcal{Q}_C) = (V, E, \lambda))$$

be a valid transcript. We say that the transcript is good (and thus $\tau \in \text{GT}$) if the following properties are true:

- (1) *The event $\text{C-f-coll}_{\text{out}}$ has not occurred.*
- (2) *The event C-coll has not occurred.*
- (3) *For any $v \in V$ we have $\lambda(v) \neq K_2$.*

We denote as GT the set of all good transcripts, and BT the set of all *bad* transcripts, i.e., transcripts which can possibly occur (i.e., they are in the range of T_{ideal}) and are not good. More specifically, we denote by BT_i the set of all bad transcripts that do not satisfy the i -th property in the definition of a good transcript above, hence we have $\text{BT} = \bigcup_{i=1}^3 \text{BT}_i$.

3.5 Probability of a C-f-collision

In this section we upper-bound the probability of C-f-coll by considering inner and outer C-f-collisions separately.

Lemma 2. *We have $\text{P}^{\text{ideal}}[\text{C-f-coll}_{\text{in}}] \leq \ell q_C q_f / 2^{b+c}$.*

Proof. We start by modifying the ideal experiment to obtain an experiment denoted ideal' and the corresponding transcript distribution $T_{\text{ideal}'}$. The experiment ideal' is given in Figure 3. Clearly, ideal' differs from the ideal experiment only in the way the vertex labeling function $\lambda(\cdot)$ is determined.

We now argue that $\text{P}^{\text{ideal}}[\text{C-f-coll}_{\text{in}}] = \text{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{in}}]$. To see this, consider an intermediate experiment ideal'' that is defined exactly as ideal except that it uses a separate ideal compression function g to generate the vertex labels of the tree contained in the transcript, where g is completely independent of f queried by the adversary (i.e., the adversary queries f and the transcript contains \mathcal{QR}_f and $T_K^g(\mathcal{Q}_C)$). It is now clear that $\text{P}^{\text{ideal}}[\text{C-f-coll}_{\text{in}}] = \text{P}^{\text{ideal}''}[\text{C-f-coll}_{\text{in}}]$ since as long as no inner C-f-collision happens, the experiments are identical.

1. **The adversary asks its C-queries and f-queries and these are answered by independent random values.** Once the q_C queries in \mathcal{Q}_C are fixed, they also determine the message tree $T(\mathcal{Q}_C)$ and mappings μ and π as defined in Section 3.1 (the labeling λ is so far undefined).
2. **Sample a repetition pattern ρ .** The equivalence relation ρ is determined indirectly by first iteratively defining a mapping $\hat{\rho}: V \rightarrow [|V|]$. Recall the vertex ordering $v_1, \dots, v_{|V|}$ defined in Section 3.1. First, set $\hat{\rho}(v_1) := 1$. Then, for i taking values $2, \dots, |V|$, determine $\hat{\rho}(v_i)$ as follows. If there exists $j \in [i-1]$ such that $\mu(v_j) = \mu(v_i)$ and $\hat{\rho}(\pi(v_j)) = \hat{\rho}(\pi(v_i))$ then let $\hat{\rho}(v_i) := \hat{\rho}(v_j)$ for the minimal such j . Otherwise let $z := \max_{j \in [i-1]} \{\hat{\rho}(v_j)\}$ and sample $\hat{\rho}(v_i)$ as

$$\hat{\rho}(v_i) := \begin{cases} 1 & \text{with probability } 2^{-c} \\ \vdots & \vdots \\ z & \text{with probability } 2^{-c} \\ z+1 & \text{with probability } 1 - z \cdot 2^{-c}. \end{cases}$$

Finally, for all $i, j \in [|V|]$ let $\rho(v_i, v_j) \Leftrightarrow (\hat{\rho}(v_i) = \hat{\rho}(v_j))$.

3. **Sample a vertex labeling $\lambda(\cdot)$ according to ρ .** Namely, sample $|\rho|$ distinct uniformly random values $s_1, \dots, s_{|\rho|} \in \{0, 1\}^c$ where $|\rho|$ is the number of equivalence classes of ρ , and let $\lambda(v_i) := s_{\hat{\rho}(v_i)}$ for all $i \in [|V|]$. Also let $K_1 := \lambda(\varepsilon)$.
4. **Sample random keys $(K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^b$.**

Fig. 3. The random experiment ideal' for the proofs of Lemmas 2 and 3.

The remaining equality $\text{P}^{\text{ideal}''}[\text{C-f-coll}_{\text{in}}] = \text{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{in}}]$ follows from the definition of ideal' . It is easy to see that the distribution of vertex labels sampled in steps 2 and 3 of ideal' and by labeling the tree $T_K^{\mathcal{E}}(\mathcal{Q}_C)$ in ideal'' are the same. In both cases, repeated inputs to the compression function lead to consistent outputs, while fresh inputs lead to independent random outputs. The two experiments only differ in the order of sampling: ideal'' first samples \mathbf{g} and then performs the labeling, while ideal' starts by sampling the repetition pattern, and then chooses the actual labels correspondingly. The same distribution of vertex labels in these two experiments then implies the same probability of $\text{C-f-coll}_{\text{in}}$ occurring.

Finally, we upper-bound the probability $\text{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{in}}]$. Conditioned on the repetition pattern ρ taking some fixed value rp , in step 2, we have

$$\begin{aligned} \text{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{in}} \mid \rho = rp] &\leq \sum_{v \in V \setminus \{\varepsilon\}} \text{P}^{\text{ideal}'}[(\lambda(\pi(v)), \mu(v) \oplus K_w) \in \mathcal{Q}_f \mid \rho = rp] \\ &= \sum_{v \in V \setminus \{\varepsilon\}} \text{P}^{\text{ideal}'}[(s_{\hat{\rho}(\pi(v))}, \mu(v) \oplus K_w) \in \mathcal{Q}_f \mid \rho = rp] \\ &= \sum_{v \in V \setminus \{\varepsilon\}} q_f / 2^{b+c} \leq \ell_{q_C} q_f / 2^{b+c} \end{aligned}$$

because the random variables s_i and K_w sampled in steps 3 and 4 are uniformly distributed and independent of \mathcal{Q}_f . Since this bound holds conditioned on ρ being any fixed repetition pattern rp , it remains valid also without conditioning on it, hence concluding the proof. \square

We proceed by upper-bounding the probability of an outer C-f-collision.

Lemma 3. *We have*

$$\text{P}^{\text{ideal}}[\text{C-f-coll}_{\text{out}}] \leq \frac{\ell_{q_C} q_f}{2^{b+c}} + \frac{q_C q_f}{2^{2c}}.$$

Proof. Let us again consider the experiments ideal' and ideal'' defined in the proof of Lemma 2. We start by the simple observation that for any event A we have

$$\begin{aligned} \text{P}^{\text{ideal}}[A] &= \text{P}^{\text{ideal}}[A \wedge \text{C-f-coll}_{\text{in}}] + \text{P}^{\text{ideal}}[A \wedge \neg \text{C-f-coll}_{\text{in}}] \\ &\leq \frac{\ell_{q_C} q_f}{2^{b+c}} + \text{P}^{\text{ideal}''}[A \wedge \neg \text{C-f-coll}_{\text{in}}] \leq \frac{\ell_{q_C} q_f}{2^{b+c}} + \text{P}^{\text{ideal}''}[A], \end{aligned} \quad (4)$$

which follows from Lemma 2 and the observation that ideal and ideal'' only differ if $\text{C-f-coll}_{\text{in}}$ occurs.

Applying (4) to the event $\text{C-f-coll}_{\text{out}}$ as A , it remains to bound the probability $\mathbb{P}^{\text{ideal}''}[\text{C-f-coll}_{\text{out}}]$; for this we observe that $\mathbb{P}^{\text{ideal}''}[\text{C-f-coll}_{\text{out}}] = \mathbb{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{out}}]$ similarly as before: the repetition pattern ρ sampled in step 2 of ideal' has the same distribution as the repetition pattern induced by the tree $T_K^g(\mathcal{Q}_C)$ in ideal'' , and this together with the sampling performed in step 3 results in the same distribution of vertex labels in ideal'' and ideal' and hence also in the same probability of $\text{C-f-coll}_{\text{out}}$ in both experiments.

Finally, to upper-bound the probability $\mathbb{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{out}}]$, again conditioned on the repetition pattern ρ sampled in step 2 taking some fixed value rp , we have

$$\begin{aligned} \mathbb{P}^{\text{ideal}'}[\text{C-f-coll}_{\text{out}} \mid \rho = rp] &\leq \sum_{v \in \mathcal{Q}_C} \mathbb{P}^{\text{ideal}'}[(K_2, \lambda(v) \parallel 0^{b-c}) \in \mathcal{Q}_f \mid \rho = rp] \\ &\leq \sum_{v \in \mathcal{Q}_C} \mathbb{P}^{\text{ideal}'}[(K_2, s_{\hat{\rho}}(v) \parallel 0^{b-c}) \in \mathcal{Q}_f \mid \rho = rp] \\ &= \sum_{v \in \mathcal{Q}_C} q_f / 2^{2c} \leq q_C q_f / 2^{2c} \end{aligned}$$

because the random variables s_i and K_2 sampled in steps 3 and 4 are uniformly distributed and independent of \mathcal{Q}_f . Since this bound holds conditioned on ρ being any fixed repetition pattern rp , it remains valid also without conditioning on it. \square

3.6 Probability of Repeated Outer Invocations

In this section we analyze the probability that any of the outer f -invocations in the ideal experiment will not be fresh, in particular we upper-bound both $\mathbb{P}[\text{T}_{\text{ideal}} \in \text{BT}_2]$ and $\mathbb{P}[\text{T}_{\text{ideal}} \in \text{BT}_3]$.

Lemma 4. *We have*

$$\mathbb{P}^{\text{ideal}}[\text{C-coll}] \leq \frac{\ell q_C q_f}{2^{b+c}} + \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} \right).$$

Proof. Applying (4) to the event C-coll , we have $\mathbb{P}^{\text{ideal}}[\text{C-coll}] \leq \ell q_C q_f / 2^{b+c} + \mathbb{P}^{\text{ideal}''}[\text{C-coll}]$. Since the queries \mathcal{Q}_C in the experiment ideal'' are chosen non-adaptively (with respect to the keys K_1 , K_w and the function g used to later compute the tree labeling), we can obtain via a union bound that

$$\mathbb{P}^{\text{ideal}''}[\text{C-coll}] \leq q_C^2 \cdot \max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \mathbb{P}^{g, K_1, K_w} [\text{WCasc}_{K_1, K_w}^g(M_1) = \text{WCasc}_{K_1, K_w}^g(M_2)].$$

Moreover, we have

$$\begin{aligned} &\max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \mathbb{P}^{g, K_1, K_w} [\text{WCasc}_{K_1, K_w}^g(M_1) = \text{WCasc}_{K_1, K_w}^g(M_2)] \\ &= \max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \sum_{\substack{K_1 \in \{0,1\}^c \\ K_w \in \{0,1\}^b}} \frac{1}{2^{c+b}} \cdot \mathbb{P}^g [\text{WCasc}_{K_1, K_w}^g(M_1) = \text{WCasc}_{K_1, K_w}^g(M_2)] \\ &\leq \sum_{\substack{K_1 \in \{0,1\}^c \\ K_w \in \{0,1\}^b}} \frac{1}{2^{c+b}} \cdot \max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \mathbb{P}^g [\text{WCasc}_{K_1, K_w}^g(M_1) = \text{WCasc}_{K_1, K_w}^g(M_2)] \\ &= \sum_{\substack{K_1 \in \{0,1\}^c \\ K_w \in \{0,1\}^b}} \frac{1}{2^{c+b}} \cdot \max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \mathbb{P}^g [\text{Casc}_{K_1}^g(M_1 \oplus K_w) = \text{Casc}_{K_1}^g(M_2 \oplus K_w)] \\ &= \sum_{\substack{K_1 \in \{0,1\}^c \\ K_w \in \{0,1\}^b}} \frac{1}{2^{c+b}} \cdot \underbrace{\max_{\substack{M_1 \neq M_2 \\ |M_1|, |M_2| \leq \ell b}} \mathbb{P}^g [\text{Casc}_{K_1}^g(M_1) = \text{Casc}_{K_1}^g(M_2)]}_{\text{CascColl}(\ell)}, \end{aligned}$$

where the notation $M_i \oplus K_w$ denotes XOR-ing the key K_w to each of the blocks of M_i .

The last maximization term above was already studied in the context of the construction NI2 in [8], where it was denoted as $\text{CColl}(\ell)$, but we will refer to it as $\text{CascColl}(\ell)$ to avoid confusion with the event C-coll considered here. It was shown in [8] that

$$\text{CascColl}(\ell) \leq \frac{\ell \cdot d'(\ell)}{2^c} + \frac{64\ell^4}{2^{2c}} . \quad (5)$$

We give an overview of the approach used in [8] to obtain this bound in Appendix B.

Putting all the above bounds together concludes the proof of Lemma 4. \square

Lemma 5. *We have*

$$P^{\text{ideal}}[\exists v \in V : \lambda(v) = K_2] \leq \frac{\ell q_C}{2^c} .$$

Proof. As is clear from the description of the ideal experiment, the key K_2 is chosen uniformly at random and independently of the rest of the experiment, in particular of the labels $\lambda(v)$. The lemma hence follows by a simple union bound over all ℓq_C vertices $v \in V$. \square

3.7 Good Transcripts and Putting Pieces Together

Let us consider a good transcript τ . First, since $\tau \notin \text{BT}_1$, there is no overlap between the outer f-involutions and the f-queries issued by the adversary. Second, since $\tau \notin \text{BT}_2$, there is also no repetition between the outer f-involutions themselves. Finally, since $\tau \notin \text{BT}_3$, there is also no overlap between the outer f-involutions and the inner f-involutions (all the outer involutions contain K_2 as their first component). Altogether, this means that each outer f-involution in *real* is fresh and hence its outcome can be seen as freshly uniformly sampled (since f is an ideal random function). Therefore, the distribution of these outcomes will be the same as in *ideal*, where they correspond to the independent random values Y_i . Hence, for all $\tau \in \text{GT}$, we have

$$\frac{P[\text{T}_{\text{real}} = \tau]}{P[\text{T}_{\text{ideal}} = \tau]} = 1 .$$

Plugging this into Lemma 1, together with the bounds from Lemmas 3, 4 and 5, we obtain

$$\begin{aligned} \text{Adv}_{\text{WNMAC}}^{\text{prf}}(A) &\leq \sum_{i=1}^3 P[\text{T}_{\text{ideal}} \in \text{BT}_i] \\ &\leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} \right) + \frac{\ell q_C}{2^c} \\ &\leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 1 \right) , \end{aligned}$$

which concludes the proof of Theorem 1. \square

3.8 Tightness

We now argue that the $q_C q_f / 2^{2c}$ term in our bound on the security of WNMAC as given in (2) is tight, by giving a matching attack (up to a linear factor $O(c)$). For most practical parameters, this will be the dominating term in (2), and thus for those parameters Theorem 1 gives a tight bound. For simpler exposition, here we only describe an attack for the case where $q_C = \Theta(c)$ is very small, we provide the description and analysis of the general attack in Appendix C.

THE $q_C = \Theta(c)$ CASE. We must define an adversary $A^{\mathcal{O},f}$ who can distinguish the case where the first oracle \mathcal{O} implements a random function R from the case where it implements $\text{WNMAC}^f((K_1, K_2, K_w), \cdot)$ with random keys K_1, K_2, K_w using the random function $f : \{0, 1\}^{b+c} \rightarrow \{0, 1\}^c$ which is given as the second oracle.

$A^{\mathcal{O},f}$ first picks $t := q_f/2^c$ keys $\tilde{K}_1, \dots, \tilde{K}_t$ arbitrarily, and then uses its q_f function queries to learn the outputs

$$\mathcal{Z}_i = \{f(\tilde{K}_i, x) \mid 0^{b-c}\} : x \in \{0, 1\}^c\}$$

for all the keys. When throwing 2^c balls randomly into 2^c bins, we expect a $1 - 1/e \approx 0.63$ fraction of the bins to be non-empty (and the value is strongly concentrated around this expectation). We can think of evaluating the random function $f(\tilde{K}_i, \cdot) \mid 0^{b-c} : \{0, 1\}^c \rightarrow \{0, 1\}^c$ as throwing 2^c balls (the inputs) to random bins (the outputs), and thus have $|\mathcal{Z}_i| \approx 0.63 \cdot 2^c$. Then $A^{\mathcal{O},f}$ queries \mathcal{O} on $\Theta(c)$ random inputs, let \mathcal{Q}_c denote the corresponding outputs. Now $A^{\mathcal{O},f}$ outputs 1 if and only if for some i we have $\mathcal{Q}_c \subset \mathcal{Z}_i$. If $\mathcal{O}(\cdot) = \text{WNMAC}^f((K_1, K_2, K_w), \cdot) = f(K_2, \text{WCasc}^f((K_1, K_w), \cdot) \mid 0^{b-c})$ and moreover $K_2 = \tilde{K}_i$ for some i – which happens with probability $t/2^c$ – then all the outputs of $\mathcal{O}(\cdot)$ are in the range of $f(\tilde{K}_i, \cdot) \mid 0^{b-c}$ and thus $A^{\mathcal{O},f}$ outputs 1.

On the other hand, if $\mathcal{O}(\cdot)$ is a random function, then every single query will miss the set \mathcal{Z}_i with constant probability 0.37. Using this, we get by a Chernoff bound (and the union bound over all t keys) that

$$\mathbb{P}[\exists i : \mathcal{Q}_c \subset \mathcal{Z}_i] \leq \frac{t}{2^{\Theta(q_C)}}.$$

Summing up we get for $q_C = \Theta(c)$ and $t = q_f/2^c$

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(A_{q_C, t}) \geq \left| \frac{t}{2^c} - \frac{t}{2^{\Theta(q_C)}} \right| \geq \frac{t}{2^{c-1}} \geq \frac{q_f}{2^{2c-1}} = \frac{q_f q_C}{2^{2c} \cdot \Theta(c)}$$

which matches our term $q_f q_C / 2^{2c}$ from the lower bound up to a $\Theta(c)$ factor.

3.9 Distinguishing-H Security of WNMAC

The above results also imply a bound on the distinguishing-H security of WNMAC. To capture this, we first introduce the notion of distinguishing-C, which corresponds to PRF-security with the restriction that the distinguisher only uses construction queries.

Definition 2 (Distinguishing-C). Let $C[f] : \{0, 1\}^\kappa \times \mathcal{D} \rightarrow \mathcal{R}$ be a keyed construction making queries to a randomly chosen compression function $f \xleftarrow{\$} \mathcal{F}(c + b, c)$. The distinguishing-C advantage of an adversary A is defined as

$$\text{Adv}_{C[f]}^{\text{dist-C}}(A) := \left| \mathbb{P} \left[K \xleftarrow{\$} \{0, 1\}^\kappa, f \xleftarrow{\$} \mathcal{F}(c + b, c) : A^{C_K^f} \Rightarrow 1 \right] - \mathbb{P} \left[R \xleftarrow{\$} \mathcal{F}(\mathcal{D}, \mathcal{R}) : A^R \Rightarrow 1 \right] \right|.$$

The notion of distinguishing-C is useful for bridging distinguishing-H and PRF-security, as the following simple lemma shows.

Lemma 6. For every adversary A asking q_C and q_f construction and primitive queries, respectively, there exists an adversary A' asking q_C queries to its single oracle such that

$$\text{Adv}_C^{\text{dist-H}}(A) \leq \text{Adv}_{C[f]}^{\text{prf}}(A) + \text{Adv}_{C[f]}^{\text{dist-C}}(A')$$

and

$$\text{Adv}_{C[f]}^{\text{prf}}(A) \leq \text{Adv}_C^{\text{dist-H}}(A) + \text{Adv}_{C[f]}^{\text{dist-C}}(A').$$

Proof (sketch). Both statements follow from the triangle inequality for distinguishing advantage and from the observation that having access to an additional oracle that is independent from the rest of the experiment and is the same in both distinguished experiments cannot increase the advantage of the adversary. \square

1. **The adversary asks its C-queries.** For each of them, only the repetition pattern for the state values belonging to this query is sampled (as in the experiment *ideal'* in Figure 3) and the query is answered with a fresh random value, unless the outer f-invocation happens on a repeated value, in which case the query is answered consistently. After answering all queries, we have a complete repetition pattern ρ for all state values.
2. **Let A output its guess** (M, s) .
3. **Sample a vertex labeling** $\lambda(\cdot)$ **according to** ρ , **let** $K_1 := \lambda(\varepsilon)$.
4. **Sample random keys** $(K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^b$.

Fig. 4. The random experiment \mathcal{E}' for the proof of Theorem 3.

One can readily obtain a bound on the distinguishing-C security of WNM_{AC} using Theorem 1 with $q_f = 0$.

Lemma 7 (Distinguishing-C Security of WNM_{AC}). *Let A be an adversary making at most q_C construction queries, each of length at most ℓ b-bit blocks. Let $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ be a tuple of random keys. Then we have*

$$\text{Adv}_{\text{WNMAC}_K}^{\text{dist-C}}(\mathbf{A}) \leq \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 1 \right).$$

By combining Theorem 1 and Lemmas 6 and 7, we get the following theorem.

Theorem 2 (Distinguishing-H Security of WNM_{AC}). *Let A be an adversary making at most q_f queries to the compression function and at most q_C construction queries, each of length at most ℓ b-bit blocks. Let $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ be a tuple of random keys. Then we have*

$$\text{Adv}_{\text{WNMAC}_K}^{\text{dist-H}}(\mathbf{A}) \leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + 2 \cdot \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 1 \right).$$

3.10 State Recovery for WNM_{AC}

We now formally define the notion of security against state recovery for WNM_{AC}. We consider the strong notion where the goal of the adversary is to output a pair (M, s) such that the state s occurs *at any point* during the evaluation of WCasc on M . Formally, we define $\text{Adv}_{\text{WNMAC}_K^f}^{\text{sr}}(\mathbf{A})$ to be

$$\mathbb{P} \left[K \stackrel{\$}{\leftarrow} \mathcal{K}, f \stackrel{\$}{\leftarrow} \mathcal{F}, \mathbf{A}^{\text{WNMAC}_{K,f}^f} \Rightarrow (M, s) : \exists M' \in \{0, 1\}^{b*} \text{ s.t. } M' \mid M \wedge \text{WCasc}_{K_1, K_w}^f(M') = s \right]$$

where $\mathcal{K} = \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$, $K = (K_1, K_2, K_w)$ and $\mathcal{F} := \mathcal{F}(c + b, c)$.

Theorem 3 (State-Recovery Security of WNM_{AC}). *Let A be an adversary making at most q_f queries to the compression function and at most q_C construction queries, each of length at most ℓ b-bit blocks. Let $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^c \times \{0, 1\}^b$ be a tuple of random keys. Then we have*

$$\text{Adv}_{\text{WNMAC}_K^f}^{\text{sr}}(\mathbf{A}) \leq \frac{q_f q_C}{2^{2c}} + 2 \cdot \frac{\ell q_C q_f}{2^{b+c}} + 2 \cdot \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 2 \right).$$

Proof (sketch). First, we replace the compression function oracle f by an independent random function g completely unrelated to WNM_{AC}^f. The error introduced by this is upper-bounded by Theorem 2 and now, compression-function queries are useless to the adversary, hence we can disregard them.

Let us denote by \mathcal{E} the experiment where A interacts with WNM_{AC}^f (without direct access to f). Consider an alternative experiment \mathcal{E}' given in Figure 4. As long as the key K_2 chosen in step 4 does not hit any of the internal states that occurred during the query evaluation, the experiment \mathcal{E}' is identical to \mathcal{E} . Moreover, since K_2 is chosen independently at random, such a hit can only occur with probability at most $\ell q_C / 2^c$. Since the vertex labels are only sampled after the adversary makes its guess for the state, the probability that the guess will be correct is at most $\ell / 2^c$. \square




Fig. 5. The construction $\text{WHMAC}[f]_{K, K_w}$.




Fig. 6. The construction $\text{WHMAC}^+[f]_{K, K_w, K^+}$.

4 Whitening HMAC

HMAC is a “practice-oriented” variant of NMAC, see Section 2 for its definition. In this section we consider a “whitened” variant WHMAC of HMAC which is derived from HMAC in the same way as WNMAC was derived from NMAC, i.e., by XORing a random key K_w to every message block. We also consider a variant WHMAC^+ where the first message block is a fresh key $K^+ \in \{0, 1\}^b$. More precisely,

$$\text{WHMAC}_{K, K_w}[f](m) := f\left(K'_2, \text{WCasc}_{K'_1, K_w}^f(m) \parallel \text{fpad}\right)$$

where

$$K'_1 := f(\text{IV}, K \oplus \text{ipad}) \quad \text{and} \quad K'_2 := f(\text{IV}, K \oplus \text{opad}) \quad (6)$$

and fpad is some fixed padding; and

$$\text{WHMAC}^+_{K, K_w, K^+}[f](m) := f\left(K'_2, \text{WCasc}_{K'_1, K_w}^f(m) \parallel \text{fpad}\right),$$

where this time

$$Z := f(\text{IV}, K \oplus \text{ipad}) \quad \text{and} \quad K'_1 := f(Z, K^+) \quad \text{and} \quad K'_2 := f(\text{IV}, K \oplus \text{opad})$$

and fpad is again some padding. For graphical representations of WHMAC and WHMAC^+ see Figures 5 and 6, respectively. Note that both variants, WHMAC and WHMAC^+ , can be implemented given just black-box access to an implementation of HMAC.

The theorem below relates the security of WHMAC and WHMAC^+ to the security of WNMAC.

Theorem 4 (Relating Security of WHMAC to WNMAC). *Consider any $\text{xxx} \in \{\text{prf}, \text{dist-H}, \text{sr}\}$. Assume that for every adversary A making at most q_f queries to the compression function f and at most q_C construction queries, each of length at most ℓ b -bit blocks, we have*

$$\text{Adv}_{\text{WNMAC}_{K_1, K_2, K_w}[f]}^{\text{xxx}}(A) \leq \epsilon,$$

where here and below, $K_1, K_2 \in \{0, 1\}^c$ and $K, K_w, K^+ \in \{0, 1\}^b$ are uniformly random keys. Then for every such adversary A we have

$$\text{Adv}_{\text{WHMAC}_{K, K_w}}^{\text{xxx}}[f](A) \leq \epsilon + 2^{-\frac{b-2c}{2}} \quad (7)$$

and

$$\text{Adv}_{\text{WHMAC}^+_{K, K_w, K^+}}^{\text{xxx}}[f](A) \leq \epsilon + 2 \cdot 2^{-\frac{b-c}{2}} + 2^{-c}. \quad (8)$$

Proof. Intuitively, for WHMAC one can think of f as an extractor which extracts keys K'_1, K'_2 from K , and the bound then readily follows by the leftover hash lemma. For WNM AC^+ one can roughly think of K'_1 and K'_2 as being extracted from independent keys K^+ and K , respectively. For the latter it is thus sufficient that b (which is the length, and thus also the entropy of the uniform K and K^+) is sufficiently larger than c (the length of K'_1, K'_2), whereas for the former we need b to be sufficiently larger than $2c$. The details of both proofs follow.

In order to prove the bound (7) it is sufficient to show that the statistical distance between the transcripts (as seen by the adversary) when interacting with WNM AC or WHMAC is at most $2^{-\frac{b-2c}{2}}$. As the only difference between WNM AC and WHMAC is that we replace the uniform keys K_1, K_2 with keys K'_1, K'_2 derived according to (6), to bound the distance between the transcripts, it is sufficient to bound the distance between the random and derived keys. As K'_1, K'_2 are not independent of f , it is important to bound the distance when given f , concretely, we must show that

$$\text{SD}((K'_1, K'_2, f), (K_1, K_2, f)) \leq 2^{-\frac{b-2c}{2}}.$$

We will use the leftover hash lemma [13] which states that for any random variable $X \in \{0, 1\}^m$ with min-entropy at least $H_\infty(X) \geq k$ and a hash function $h : \{0, 1\}^m \rightarrow \{0, 1\}^\ell$ chosen from a family of pairwise independent hash functions we have (with U_ℓ being uniform over $\{0, 1\}^\ell$)

$$\text{SD}((h(X), h), (U_\ell, h)) \leq 2^{\frac{\ell - H_\infty(X)}{2}} \leq 2^{\frac{\ell - k}{2}}.$$

Since $f : \{0, 1\}^{b+c} \rightarrow \{0, 1\}^c$ is uniformly random, also the function

$$f'(K) = (f(\text{IV}, K \oplus \text{ipad}), f(\text{IV}, K \oplus \text{opad}))$$

is uniformly random, and thus also pairwise independent. Using $H_\infty(K) = H_\infty(K \oplus \text{ipad}) = b$ and $(K'_1, K'_2) = f'(K)$ we thus get

$$\text{SD}((K'_1, K'_2, f'), (K_1, K_2, f')) = \text{SD}((K'_1, K'_2, f), (K_1, K_2, f)) \leq 2^{-\frac{b-2c}{2}}$$

as required. The first equality above holds as f defines all of f' and vice versa.

To establish the bound (8), consider now the keys K'_1, K'_2 derived according to (4). We will again upper bound the statistical distance of the derived and the random keys, proving

$$\text{SD}((K'_1, K'_2, f), (K_1, K_2, f)) \leq 2 \cdot 2^{-\frac{b-c}{2}} + 2^{-c}.$$

Instead of K'_1 , consider a key K''_1 which is computed as $K''_1 = f(z, K^+)$ for some fixed $z \neq \text{IV}$. By the leftover hash lemma we get

$$\begin{aligned} \text{SD}((K''_1, f(z, \cdot)), (K_1, f(z, \cdot))) &\leq 2^{-\frac{b-c}{2}} \\ \text{SD}((K'_2, f(\text{IV}, \cdot)), (K_2, f(\text{IV}, \cdot))) &\leq 2^{-\frac{b-c}{2}}. \end{aligned}$$

Noting that $f(z, \cdot)$ and $f(\text{IV}, \cdot)$ are independent random functions, and K and K^+ are also independent, we can apply the triangle inequality for statistical distance to obtain

$$\text{SD}((K''_1, K'_2, f), (K_1, K_2, f)) \leq 2 \cdot 2^{-\frac{b-c}{2}}.$$

If we replace $K''_1 = f(z, K^+)$ with $K'_1 = f(Z, K^+)$ above we get an extra 2^{-c} term which accounts for the probability that $Z = \text{IV}$. \square




Fig. 7. The construction $\text{DWNMAC}[f]_{K_1, K_2, K_w}$.

5 The Dual WNMAC Construction

Looking at the security bounds for WNMAC given in Section 3 from a distance, it seems that under reasonable assumptions the most restrictive term in the bounds is $q_f q_C / 2^{2c}$. Intuitively speaking, the reason for this term is the outer f -call in WNMAC that only takes $2c$ bits of actual inputs and adds $b - c$ padding zeroes.

In an attempt to overcome this limitation, we propose a variant of the WNMAC construction that we call *Dual WNMAC* (DWNMAC). We prove the PRF-security of DWNMAC that goes beyond the restrictive term $q_f q_C / 2^{2c}$ and our proof again extends also to distinguishing-H and state-recovery security. The price we pay for this improvement is a slight increase in the key length and the fact that DWNMAC cannot be implemented using only black-box access to NMAC. Similarly, if we apply the same modification to WHMAC, the resulting construction can no longer be implemented using black-box access to HMAC.

The construction DWNMAC is derived from WNMAC, the only difference being that the outer f -call is performed on the c -bit state and a b -bit key K_2 . More precisely, for a key tuple $(K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^b \times \{0, 1\}^b$ and a message $M \in \{0, 1\}^{b^*}$, we define

$$\text{DWNMAC}^f((K_1, K_2, K_w), M) := f(\text{WCasc}_{K_1, K_w}^f(M), K_2).$$

For a graphical depiction of DWNMAC, see Figure 7. Note that DWNMAC is slightly similar to what we would obtain by whitening from the Sandwich MAC construction [24].

We now summarize the security of DWNMAC.

Theorem 5 (Security of DWNMAC). *Let A be an adversary making at most q_f queries to the compression function f and at most q_C construction queries, each of length at most ℓ b -bit blocks. Let $K = (K_1, K_2, K_w) \in \{0, 1\}^c \times \{0, 1\}^b \times \{0, 1\}^b$ be a tuple of random keys. Then we have*

$$\text{Adv}_{\text{DWNMAC}_K^c}^{\text{xxx}}(A) \leq 3 \cdot \frac{\ell q_C q_f}{2^{b+c}} + 2 \cdot \frac{\ell q_C^2}{2^c} \cdot \left(d'(\ell) + \frac{64\ell^3}{2^c} + 2 \right)$$

for all $\text{xxx} \in \{\text{prf}, \text{dist-H}, \text{sr}\}$.

Proof (sketch). The proofs are analogous to the proofs for WNMAC given in Section 3, with the main modification needed in Lemma 3 where the probability of an outer C - f -collision can be upper-bounded by $q_C q_f / 2^{b+c}$. Roughly speaking, this is because the outer call in DWNMAC does not contain the 0^{b-c} padding and instead processes $b + c$ bits of input that are hard to predict for the attacker. \square

Acknowledgments We thank the anonymous reviewers for their helpful comments. Gaži and Pietrzak's work was partly funded by the European Research Council under an ERC Starting Grant (259668-PSPC). Tessaro's research was partially supported by NSF grant CNS-1423566 and by the Glen and Susanne Culler Chair.

References

1. Jee Hea An and Mihir Bellare. Constructing VIL-MACs from FIL-MACs: Message authentication under weakened assumptions. In Michael J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 252–269, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany.
2. Mihir Bellare. New proofs for NMAC and HMAC: Security without collision-resistance. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 602–619, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany.
3. Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 1–15, Santa Barbara, CA, USA, August 18–22, 1996. Springer, Heidelberg, Germany.
4. Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
5. Ivan Damgård. A design principle for hash functions. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 416–427, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany.
6. Itai Dinur and Gaëtan Leurent. Improved generic attacks against hash-based MACs and HAIFA. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 149–168, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
7. Yevgeniy Dodis, Thomas Ristenpart, John P. Steinberger, and Stefano Tessaro. To hash or not to hash again? (in)differentiability results for h^2 and HMAC. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 348–366, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.
8. Peter Gaži, Krzysztof Pietrzak, and Michal Rybár. The exact PRF-security of NMAC and HMAC. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 113–130, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
9. Peter Gaži, Krzysztof Pietrzak, and Stefano Tessaro. The exact PRF security of truncation: Tight bounds for keyed sponges and truncated CBC. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 368–387, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
10. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. On the cryptographic applications of random functions. In G. R. Blakley and David Chaum, editors, *CRYPTO'84*, volume 196 of *LNCS*, pages 276–288, Santa Barbara, CA, USA, August 19–23, 1984. Springer, Heidelberg, Germany.
11. Jian Guo, Thomas Peyrin, Yu Sasaki, and Lei Wang. Updates on generic attacks against HMAC and NMAC. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 131–148, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
12. G. H. Hardy and Edward M. Wright. *An Introduction to the Theory of Numbers (sixth edition)*. Oxford University Press, USA, 2008.
13. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
14. Jongsung Kim, Alex Biryukov, Bart Preneel, and Seokhie Hong. On the security of HMAC and NMAC based on HAVAL, MD4, MD5, SHA-0 and SHA-1 (extended abstract). In Roberto De Prisco and Moti Yung, editors, *SCN 06*, volume 4116 of *LNCS*, pages 242–256, Maiori, Italy, September 6–8, 2006. Springer, Heidelberg, Germany.
15. Hugo Krawczyk, Mihir Bellare, and Ran Canetti. HMAC: Keyed-hashing for message authentication. IETF Internet Request for Comments 2104, February 1997.
16. Gaëtan Leurent, Thomas Peyrin, and Lei Wang. New generic attacks against hash-based MACs. In Kazuo Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 1–20, Bangalore, India, December 1–5, 2013. Springer, Heidelberg, Germany.
17. Ralph C. Merkle. One way hash functions and DES. In Gilles Brassard, editor, *CRYPTO'89*, volume 435 of *LNCS*, pages 428–446, Santa Barbara, CA, USA, August 20–24, 1990. Springer, Heidelberg, Germany.
18. Yusuke Naito, Yu Sasaki, Lei Wang, and Kan Yasuda. Generic state-recovery and forgery attacks on ChopMD-MAC and on NMAC/HMAC. In Kazuo Sakiyama and Masayuki Terada, editors, *IWSEC 13*, volume 8231 of *LNCS*, pages 83–98, Okinawa, Japan, 2013. Springer, Heidelberg, Germany.
19. Jacques Patarin. The “coefficients H” technique (invited talk). In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345, Sackville, New Brunswick, Canada, August 14–15, 2009. Springer, Heidelberg, Germany.

20. Thomas Peyrin, Yu Sasaki, and Lei Wang. Generic related-key attacks for HMAC. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 580–597, Beijing, China, December 2–6, 2012. Springer, Heidelberg, Germany.
21. Thomas Peyrin and Lei Wang. Generic universal forgery attack on iterative hash-based MACs. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 147–164, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
22. Bart Preneel and Paul C. van Oorschot. MDx-MAC and building fast MACs from hash functions. In Don Coppersmith, editor, *CRYPTO'95*, volume 963 of *LNCS*, pages 1–14, Santa Barbara, CA, USA, August 27–31, 1995. Springer, Heidelberg, Germany.
23. Yu Sasaki and Lei Wang. Generic attacks on strengthened HMAC: n-bit secure HMAC requires key in all blocks. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 324–339, Amalfi, Italy, September 3–5, 2014. Springer, Heidelberg, Germany.
24. Kan Yasuda. “sandwich” is indeed secure: How to authenticate a message with just one hashing. In Josef Pieprzyk, Hossein Ghodosi, and Ed Dawson, editors, *ACISP 07*, volume 4586 of *LNCS*, pages 355–369, Townsville, Australia, July 2–4, 2007. Springer, Heidelberg, Germany.

A The H -Coefficient Method

In this section we prove the basic lemma underlying Patarin’s H -Coefficient method [19].

Lemma 1 (restated). *Let $\delta, \epsilon \in [0, 1]$ be such that:*

- (a) $\mathbb{P}[\mathbf{T}_{\text{ideal}} \in \text{BT}] \leq \delta$.
- (b) For all $\tau \in \text{GT}$,

$$\frac{\mathbb{P}[\mathbf{T}_{\text{real}} = \tau]}{\mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau]} \geq 1 - \epsilon .$$

Then,

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(\mathbf{A}) \leq \text{SD}(\mathbf{T}_{\text{real}}, \mathbf{T}_{\text{ideal}}) \leq \epsilon + \delta .$$

Proof. Let \mathcal{T} be the set of valid transcripts such that $\mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] \geq \mathbb{P}[\mathbf{T}_{\text{real}} = \tau]$. Then,

$$\text{SD}(\mathbf{T}_{\text{real}}, \mathbf{T}_{\text{ideal}}) = \sum_{\tau \in \mathcal{T}} (\mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] - \mathbb{P}[\mathbf{T}_{\text{real}} = \tau])$$

by the fundamental properties of the statistical distance. Then, note that \mathcal{T} can be partitioned into two blocks $\mathcal{T} \cap \text{BT}$ and $\mathcal{T} \cap \text{GT}$. On the one hand, we can use (a) to upper bound

$$\sum_{\tau \in \mathcal{T} \cap \text{BT}} (\mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] - \mathbb{P}[\mathbf{T}_{\text{real}} = \tau]) \leq \sum_{\tau \in \mathcal{T} \cap \text{BT}} \mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] \leq \sum_{\tau \in \text{BT}} \mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] \leq \delta .$$

On the other hand, (b) implies

$$\sum_{\tau \in \mathcal{T} \cap \text{GT}} (\mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] - \mathbb{P}[\mathbf{T}_{\text{real}} = \tau]) \leq \epsilon \cdot \sum_{\tau \in \mathcal{T} \cap \text{GT}} \mathbb{P}[\mathbf{T}_{\text{ideal}} = \tau] \leq \epsilon .$$

Therefore, $\text{SD}(\mathbf{T}_{\text{real}}, \mathbf{T}_{\text{ideal}}) \leq \epsilon + \delta$. Moreover, every adversary \mathbf{A} can be turned into a distinguisher \mathbf{A}' for \mathbf{T}_{real} and $\mathbf{T}_{\text{ideal}}$, which looks at the first part of the transcript (i.e., the one containing the q message-output pairs $(M_1, Y_1), \dots, (M_q, Y_q)$), and outputs the corresponding decision bit \mathbf{A} would output (this bit is uniquely defined by the fact that \mathbf{A} is deterministic). Then, we clearly have

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(\mathbf{A}) = \mathbb{P}[\mathbf{A}'(\mathbf{T}_{\text{real}}) \Rightarrow 1] - \mathbb{P}[\mathbf{A}'(\mathbf{T}_{\text{ideal}}) \Rightarrow 1] \leq \text{SD}(\mathbf{T}_{\text{real}}, \mathbf{T}_{\text{ideal}}) \leq \epsilon + \delta ,$$

as the statistical distance is the quantity corresponding to the advantage of the best \mathbf{A}' . \square

B Probability of a Cascade Collision

Here we briefly summarize the technique used in [8] to upper-bound the quantity `CascColl` considered in the proof of Lemma 4. For further details, we refer the reader to [8].

Intuitively, the problem of upper-bounding the probability $\mathbb{P}[\text{CascColl}]$ is reduced to a combinatorial counting problem. The objects counted are so-called “structure graphs” that represent the structure of intermediate values obtained when labeling the message tree corresponding to \mathcal{Q}_C using a random compression function f .⁵ On a high level, these structure graphs differ from our labeled message trees by merging the vertices that end up having the same label (and hence resulting in a directed acyclic graph that does not necessarily have to be a tree). It is shown in [8, Lemma 2] that the probability (over the randomness of f) of a fixed structure graph occurring can be upper-bounded using the number of so-called f -collisions in this graph. Then, loosely speaking, [8, Lemma 3] shows that structure graphs containing two or more such f -collisions are too unlikely to matter, while [8, Lemma 4] performs the actual counting and gives an upper bound on the number of distinct structure graphs that contain at most one such f -collision. Overall, this results in the desired bound (5).

C Description of the General Attack

Here we describe how to generalize the attack from Section 3.8 to the case when q_C is large, concretely, for a given q_C we let $0 \leq \alpha \leq 1$ be such that

$$q_C = 2^{\alpha \cdot c} \cdot \Theta(c)$$

and further we set

$$t = q_f / 2^{c(1-\alpha)}$$

Note that for $\alpha = 0$ we get the parameters of the previous attack. $\mathcal{A}^{\mathcal{O},f}$ first picks some (arbitrary) subset $\mathcal{X} \subset \{0, 1\}^c$ of size $|\mathcal{X}| = 2^{c(1-\alpha)}$, and t (arbitrary) keys $\tilde{K}_1, \dots, \tilde{K}_t$, and then uses its q_f function queries to learn the outputs of f on those keys:

$$\mathcal{Z}_i = \{f(\tilde{K}_i, x) \parallel 0^{b-c} : x \in \mathcal{X}\}$$

We now let $\mathcal{A}^{\mathcal{O},f}$ query \mathcal{O} on q_C random inputs and let \mathcal{Q}_c denote the outputs. In the previous attack we simply let $\mathcal{A}^{\mathcal{O},f}$ output 1 if $\mathcal{Q}_c \subset \mathcal{Z}_i$ for some i , as this could only happen (except with extremely tiny probability) if \mathcal{O} implemented WNMAC where the key K_2 for the outer invocation was \tilde{K}_i .

Now we have to be a bit more careful, as \mathcal{Z}_i contains just a (possibly very small) fraction of the outputs of $f(\tilde{K}_i, \cdot \parallel 0^{b-c})$, and thus \mathcal{Q}_c will not be a subset of \mathcal{Z}_i even if $K_c = \tilde{K}_i$. Instead we will look at the size of the intersection

$$\mathcal{I}_i = \mathcal{Q}_c \cap \mathcal{Z}_i .$$

If $\mathcal{O}(\cdot)$ implements a random function, then \mathcal{Q}_c is a set of $q_c = 2^{\alpha \cdot c} \cdot \Theta(c)$ random values from $\{0, 1\}^c$, and \mathcal{Z}_i is a random subset of size $2^{(1-\alpha)c}$.⁶ Let \mathcal{U}_s denote a random subset of $\{0, 1\}^c$ of size s . As $|\mathcal{Q}_c| \cdot |\mathcal{Z}_i| = 2^c \cdot \Theta(c)$ we have for some constant γ_i (this constant is small, and its exact value depends on $|\mathcal{Q}_c|$ and $|\mathcal{Z}_i|$)

$$E[|\mathcal{Q}_c| \cdot |\mathcal{Z}_i|] = E[|\mathcal{U}_{|\mathcal{Q}_c|}| \cdot |\mathcal{U}_{|\mathcal{Z}_i|}|] = \gamma_i \cdot c$$

Now consider the case where the first oracle is

$$\mathcal{O}(\cdot) = \text{WNMAC}^f((K_1, K_2, K_w), \cdot) = f(K_2, \text{WCasc}^f((K_1, K_w), \cdot) \parallel 0^{b-c})$$

⁵ We note that in [8] the initial state is fixed to a c -bit all-zero string $\mathbf{0}$ instead of the value K_1 considered in our Lemma 4, but it is easy to verify that since K_1 is already fixed at that point, this does not introduce any difference.

⁶ The sets \mathcal{Q}_c and \mathcal{Z}_i can be slightly smaller than $2^{(1-\alpha)c}$ and $2^{\alpha \cdot c} \Theta(c)$, respectively, due to collisions. But this can be ignored as even in the extreme case where α is 0 or 1, this will affect the size of the set only by a factor $1/e$.

and $K_2 = \tilde{K}_i$ for some i . We claim that in this case the intersection will be roughly twice as large:

$$E[|\mathcal{Q}_c| \cdot |\mathcal{Z}_i|] = E[|\mathcal{U}_{|\mathcal{Q}_c|}| \mathcal{U}_{|\mathcal{Z}_i|}] \approx 2 \cdot \gamma_i \cdot c. \quad (9)$$

The reason is that now we can have collisions on either the inputs to $f(K_2 = \tilde{K}_i, \cdot)$, or on the outputs, thus doubling our chances to see a collision.

The strategy of $A^{\mathcal{O},f}$ is now to check, for all $i = 1, \dots, t$, if $|\mathcal{Q}_c \cap \mathcal{Z}_i|$ is much bigger than it should be assuming \mathcal{O} is a random function, say

$$|\mathcal{Q}_c \cap \mathcal{Z}_i| \geq 1.5 \cdot \gamma_i$$

and output 1 if this is the case. In the case where $\mathcal{O}(\cdot)$ is WNMAC and $K_2 = \tilde{K}_i$ for some i , we'll almost certainly output 1. If $\mathcal{O}(\cdot)$ is a random function, we expect to pass the above check (and thus get a false positive) with probability only $2^{-\Theta(c)}$ for every key (this again follows by the Chernoff bound), and thus by a union bound with probability $t/2^{\Theta(c)}$ overall. If the hidden constant in the previous Θ too small to make $t/2^{\Theta(c)}$ small ($< 1/2$ is sufficient), we can let the adversary save up some of its q_f queries (at the prize of a slightly smaller t) for extra queries used to get additional outputs for keys \tilde{K}_i where the above check passes, in order to make the probability of a false positive extremely small.

Summing up we get an advantage of roughly (recall that $q_C = 2^{\alpha \cdot c} \cdot \Theta(c)$ and $t = q_f/2^{c(1-\alpha)}$)

$$\text{Adv}_{\text{WNMAC}}^{\text{prf}}(A_{q_C,t}) \approx \frac{t}{2^c} = \frac{q_f}{2^{c(1-\alpha)}2^c} = \frac{q_f}{2^{-c\alpha}2^{2c}} = \frac{q_f q_C}{2^{2c} \cdot \Theta(c)}$$

which again matches our term $q_f q_C / 2^{2c}$ from the lower bound up to a factor $O(c)$.