

The Exact PRF-Security of NMAC and HMAC

Peter Gaži, Krzysztof Pietrzak, and Michal Rybár

IST Austria

Abstract. NMAC is a mode of operation which turns a fixed input-length keyed hash function f into a variable input-length function. A practical single-key variant of NMAC called HMAC is a very popular and widely deployed message authentication code (MAC). Security proofs and attacks for NMAC can typically be lifted to HMAC.

NMAC was introduced by Bellare, Canetti and Krawczyk [Crypto'96], who proved it to be a secure pseudorandom function (PRF), and thus also a MAC, assuming that (1) f is a PRF and (2) the function we get when cascading f is weakly collision-resistant. Unfortunately, HMAC is typically instantiated with cryptographic hash functions like MD5 or SHA-1 for which (2) has been found to be wrong. To restore the provable guarantees for NMAC, Bellare [Crypto'06] showed its security based solely on the assumption that f is a PRF, albeit via a non-uniform reduction.

- Our first contribution is a simpler and uniform proof for this fact: If f is an ε -secure PRF (against q queries) and a δ -*non-adaptively* secure PRF (against q queries), then NMAC^f is an $(\varepsilon + \ell q \delta)$ -secure PRF against q queries of length at most ℓ blocks each.
- We then show that this $\varepsilon + \ell q \delta$ bound is basically tight. For the most interesting case where $\ell q \delta \geq \varepsilon$ we prove this by constructing an f for which an attack with advantage $\ell q \delta$ exists. This also violates the bound $O(\ell \varepsilon)$ on the PRF-security of NMAC recently claimed by Kobitz and Menezes.
- Finally, we analyze the PRF-security of a modification of NMAC called NI [An and Bellare, Crypto'99] that differs mainly by using a compression function with an additional keying input. This avoids the constant rekeying on multi-block messages in NMAC and allows for a security proof starting by the standard switch from a PRF to a random function, followed by an information-theoretic analysis. We carry out such an analysis, obtaining a tight $\ell q^2/2^c$ bound for this step, improving over the trivial bound of $\ell^2 q^2/2^c$. The proof borrows combinatorial techniques originally developed for proving the security of CBC-MAC [Bellare et al., Crypto'05].

Keywords: Message authentication codes, pseudorandom functions, NMAC, HMAC, NI.

1 Introduction

NMAC is a mode of operation which transforms a keyed fixed input-length function $f : \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ (with $b \geq c$) into a keyed variable input-length

function $\text{NMAC}^f : \{0, 1\}^{2c} \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ (where $\{0, 1\}^{b^*}$ denotes all bit strings whose length is a multiple of b) as

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

where $\text{Casc}^f : \{0, 1\}^c \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ is the cascade (also known as Merkle-Damgård) construction

$$\text{Casc}^f(K_1, m_1 \| \dots \| m_\ell) := f(\dots f(f(K_1, m_1), m_2) \dots m_\ell).$$

HMAC is a variant of NMAC (we postpone its exact definition to Section 2.2) tweaked for applicability in practice. As security proofs for NMAC can typically be lifted to HMAC, it is usually sufficient to analyse the security of the cleaner NMAC construction, we will discuss this point further in Section 1.2.

NMAC and HMAC were introduced by Bellare, Canetti and Krawczyk in 1996 [4] and later standardized [18]. HMAC has also become very popular and widely used, being implemented in SSL, SSH, IPsec and TLS amongst other places. Although originally designed as a MAC, it is also often employed more broadly, as a pseudorandom function (PRF). This is the case for example when used for key-derivation in TLS and IKE (the Internet Key Exchange protocol of IPsec). This proliferation into practice motivates the need for a good understanding of the exact security guarantees provided by NMAC and HMAC when used as a PRF.

PRF-SECURITY OF NMAC. Bellare *et al.* [4] prove that NMAC is a secure PRF if (1) f is a PRF and (2) Casc^f is weakly collision-resistant (WCR). This is a relaxed notion of collision resistance, where one requires that it is hard to find a pair of messages $M \neq M'$ such that $\text{Casc}^f(K, M) = \text{Casc}^f(K, M')$ under a random key K , given oracle access to $\text{Casc}^f(K, \cdot)$ (but not K , as in the standard definition of collision resistance).

HMAC is typically instantiated with cryptographic hash functions like MD5 or SHA-1 playing the role of Casc^f . However, both of these have been found not to satisfy the WCR notion [26, 27], which renders the security proof from [4] irrelevant for this case. Despite that, no attacks (better than standard birthday attacks) are known for NMAC or HMAC when instantiated with MD5 or SHA-1 (though attacks on reduced round versions exist [16]).

SECURITY WITHOUT COLLISION-RESISTANCE. To restore the provable security of NMAC, Bellare [3] investigates the security of NMAC dropping assumption (2), that is, assuming only that f is a secure PRF. The exact security statement from [3] is a bit technical, but it roughly states that if f is an ε -secure PRF (against an adversary running in time t and asking q queries) and a γ -secure PRF (against time $O(\ell)$ and 2 queries), then NMAC^f is an $(\varepsilon + \ell q^2 \gamma)$ -secure PRF against time t and q queries of length at most ℓ (in b -bit blocks). The security reduction is non-uniform, which means one has to be careful when deducing what this bound exactly means when instantiated in practice, we will discuss this further in Section 1.2.¹

¹ We note that in a very recent update of the ePrint version of [3], Bellare observes that the proof in [3] can also give a uniform reduction, differing from the non-uniform

1.1 Our Contributions

PRF-SECURITY PROOF FOR NMAC. Our first contribution is a simpler, uniform, and as we will show, basically tight proof for the PRF-security of NMAC^f assuming only that f is a PRF: If f is an ε -secure PRF against q queries, then NMAC^f is roughly $\ell q \varepsilon$ -secure against q queries of length at most ℓ blocks each.

Our actual result is more fine-grained, and expresses the security in terms of both the adaptive and non-adaptive security of f . Let δ denote the PRF-security of f against q *non-adaptive* queries. Then our Theorem 1 states that NMAC^f is roughly $(\varepsilon + \ell q \delta)$ -secure (against q queries, each at most ℓ blocks). As non-adaptive adversaries are a subset of adaptive ones we have $\delta \leq \varepsilon$, and if $\delta \ll \varepsilon$, then our fine-grained bound is much better than the simpler $\ell q \varepsilon$ bound. The reduction works in the best running time one could hope for, its overhead being $\tilde{O}(\ell q)$.

The main technical part of our proof closely follows a proof by Bellare *et al.* [5] who show that if f is a secure fixed input-length PRF, then Casc^f is a secure PRF if queried on prefix-free queries. We first observe that their proof also holds in the non-adaptive setting. Then we reduce the security of NMAC^f against arbitrary adaptive queries to the security of Casc^f against non-adaptive prefix-free queries.

MATCHING ATTACK FOR NMAC. In Section 3.2 we prove that the above lower bound is basically tight. From any PRF, we construct another PRF f for which NMAC^f can be broken with advantage $\Theta(\ell q \delta)$. This shows that our bound is tight for the practically most important case when $\ell q \delta$ is larger (or at least comparable) to ε .

We also consider the case where $\varepsilon \gg \ell q \delta$, that is, when the PRF has much better security against non-adaptive than adaptive distinguishers. We observe that for any ε , we can use a result due to Pietrzak [23] who shows that cascading non-adaptively secure PRFs does not give an adaptively secure PRF in general, to construct an ε -secure f where NMAC^f can be broken with advantage $\Theta(\varepsilon^2)$. This only shows the ε term is necessary if ε is constant as then $\Theta(\varepsilon) = \Theta(\varepsilon^2) = \Theta(1)$. We conjecture that $\Theta(\varepsilon^2)$ is the correct value, and the ε term in the lower bound can be improved to $\Theta(\varepsilon^2)$ using security amplification techniques along the lines of [22, 25].

PRF-SECURITY PROOF FOR NI. The main difficulty in security analyses of NMAC^f and HMAC^f based on the PRF-security of the underlying compression function f is that both these constructions are constantly rekeying f during the evaluation of Casc^f , using the output from the last invocation as the key for the next one. This prevents the proof approach typically applied to constructions that use a PRF f under a fixed random secret key, where the analysis starts by replacing the PRF with an ideal random function (introducing an error that is upper-bounded by the PRF-security of f) and proceeds by a fully information-theoretic argument.

case only in the running time of the 2-query adversary which then becomes t . The uniform bound given in this paper is better for most reasonable parameters.

To circumvent this issue, as our third contribution we investigate the PRF-security of the nested iterated (NI) construction introduced in [2]. The construction NI^{h} is very similar to NMAC^{f} , but is based on a compression function h that (compared to f) takes an additional k -bit input which is used for keying instead of the chaining input: NI^{h} uses h under the same key throughout the whole cascade. Additionally, it includes the length of the message in the input to the final, outer h -call. The modified keying allows for the simple switching argument from PRF to a random function. We focus on enhancing the information-theoretic analysis that follows this switch and prove an essentially tight $\ell q^2/2^c$ bound for this step, improving significantly over the trivial bound of $\ell^2 q^2/2^c$. For completeness, we also consider the modification of NI that does not include the message length in the last h -call and show a security bound of $\ell d'(\ell)q^2/2^c$ for this case, where $d'(\ell) \approx \ell^{1/\ln \ln \ell}$ denotes the maximum number of divisors of any positive integer not greater than ℓ . Our proofs employ combinatorial techniques originally developed for proving the security of CBC-MAC [7], considerably adapted for our setting.

1.2 More Related Work

INDIFFERENTIABILITY. In practice, the HMAC construction is sometimes used in a setting where stronger guarantees than PRF-security are needed. Motivated by this, recent work [12] investigates the indifferenciability [21, 10] of HMAC from a (keyed) random oracle. This result is incomparable to ours: While the stronger notion of indifferenciability covers the settings where HMAC is not used as a PRF, the bound achieved in [12] is understandably much weaker, being $\Theta(\ell^2 q^2/2^c)$.

ANOTHER LOOK AT [17]. As already mentioned, Bellare [3] proved that NMAC^{f} is an $(\varepsilon + \ell q^2 \gamma)$ -secure PRF against q queries if f is ε -secure against q queries, and γ -secure against 2 queries. In a recent paper [17], Kobitz and Menezes present a criticism of the way [3] discusses the practical implications of this result. In a nutshell, Bellare estimates that for a well-designed PRF the γ term is roughly $t/2^c$ (for a 2-query adversary running in time t), but as this γ is derived in a non-uniform way, it is in the order of $2^{-c/2}$ already for constant t .

At the time when [3] appeared, the fact that non-uniform attacks can distinguish any pseudorandom object generated using a c -bit key with advantage $2^{-c/2}$ in constant time was not widely known in the crypto community² and overoptimistic estimates for the exact security implied by non-uniform reductions have appeared in numerous papers.³ This changed at the latest with the

² Let us stress that this only holds for pseudorandom objects which do not require additional *public* randomness, such as PRFs. This does not extend to weak PRFs, which are defined like PRFs but the adversary only sees the output on random inputs.

³ This should not be confused with the (less trivial, but in the crypto community long well-known) fact that non-uniform generic attacks beating simple brute-force key search exist for “large” running times, as shown in a classical result by Hellman [14]. Hellman’s result for example implies that there almost certainly exist key-recovery

Crypto 2010 paper [11], who discuss this issue in detail and attribute such generic non-uniform attacks to the 1992 paper by Alon *et al.* [1].

The paper [17] also claimed that HMAC is an $\varepsilon\ell$ -secure PRF, a bound that is falsified by an attack given in this paper. In response, [17] was updated to take account of this by employing a non-standard definition of a PRF for the underlying compression function. We believe that the updated claim can be obtained via a simpler proof from [5].

HMAC vs NMAC. The proofs in this paper consider NMAC. There is a standard reduction of HMAC-to-NMAC PRF-security given by Bellare [3], albeit under some additional requirements on the underlying compression function f . Informally, one needs to assume that f is a PRF even when keyed through the b -bit data input, as opposed to being keyed by the c -bit chaining variable. Moreover, security of the single-key version of HMAC requires the PRF to be secure under a specific class of related-key attacks. Formally, the reductions are given in Lemmas 5.1 and 5.2 in the full version of [3] for the case of double- and single-keyed HMAC, respectively. Since these reductions only relate to NMAC via its PRF-security, they apply to our result in a blackbox way, thus giving clear statements also for HMAC.

2 Preliminaries

BASIC DEFINITIONS. We reserve the letter λ to denote the empty string. With $\{0, 1\}^{b*} := \bigcup_{z \geq 0} \{0, 1\}^{bz}$ we denote the set of all bitstrings whose length is a multiple of b . $\mathcal{F}(b, c)$ (resp. $\mathcal{F}(b*, c)$) denotes the sets of all functions from $\{0, 1\}^b$ to $\{0, 1\}^c$ (resp. from $\{0, 1\}^{b*}$ to $\{0, 1\}^c$). We denote by $\text{Pow}(\mathcal{S})$ the power set of the set \mathcal{S} . 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 $\forall \varepsilon > 0 \exists n_0 \forall n > n_0 : d(n) < n^{(1+\varepsilon)/\ln \ln n}$ [13]. All logarithms considered in the paper are base 2 unless indicated otherwise.

RANDOM VARIABLES AND EXPERIMENTS. Random variables and concrete values they can take are usually denoted by upper-case letters X, Y, \dots and lower-case letters x, y, \dots , respectively. If \mathcal{M} is a distribution (respectively, a set), then we denote by $X \leftarrow \mathcal{M}$ sampling the random variable X according to \mathcal{M} (respectively, choosing it uniformly at random from \mathcal{M}). For events A and B and random variables U and V with ranges \mathcal{U} and \mathcal{V} , respectively, we denote by $P_{U|V}$ the corresponding conditional probability distribution, seen as a (partial) function $\mathcal{U} \times \mathcal{V} \rightarrow [0, 1]$. The value $P_{U|V}(u, v) = P[U = u \wedge A | V = v \wedge B]$ is well-defined for all $u \in \mathcal{U}$ and $v \in \mathcal{V}$ such that $P_V(v) > 0$ and undefined otherwise. Two probability distributions P_U and $P_{U'}$ on the same set \mathcal{U} are equal, denoted $P_U = P_{U'}$, if $P_U(u) = P_{U'}(u)$ for all $u \in \mathcal{U}$. Conditional

attacks against AES with a k bit key (k being 128, 192 or 256) which succeed with probability at least $1/2$ and run in time $\approx 2^{2k/3}$, and in particular much less than 2^k required for brute-force key search.

probability distributions are equal if the equality holds for all arguments for which both of them are defined. To emphasize the random experiment \mathcal{E} in consideration, we sometimes write it in the superscript, e.g. $\mathbf{P}_{U|V}^{\mathcal{E}}(u, v)$. If the distribution of a random variable U is clear from the context, we also sometimes write \mathbf{P}^U to refer to the random experiment where U is chosen according to its distribution.

2.1 Random Systems

To present our results we make use of Maurer’s random systems framework [20], which we now introduce in a self-contained exposition sufficient to follow the rest of the paper. This choice is a matter of authors’ taste, we believe that the results could also be obtained using the game-playing framework [8].

We start by observing that the input-output behavior of any kind of reactive discrete system with inputs in \mathcal{X} and outputs in \mathcal{Y} can be described by an infinite family of functions specifying, for each $i \geq 1$, the probability distribution of the system’s i -th output $Y_i \in \mathcal{Y}$, given the values of the first i inputs $X^i \in \mathcal{X}^i$ and the previous $i - 1$ outputs $Y^{i-1} \in \mathcal{Y}^{i-1}$. Using this viewpoint, we say that an $(\mathcal{X}, \mathcal{Y})$ -*(random) system* \mathbf{F} is an infinite sequence of functions $\mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}}: \mathcal{Y} \times \mathcal{X}^i \times \mathcal{Y}^{i-1} \rightarrow [0, 1]$ such that $\sum_{y_i} \mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}}(y_i, x^i, y^{i-1}) = 1$ for all $i \geq 1$, $x^i \in \mathcal{X}^i$ and $y^{i-1} \in \mathcal{Y}^{i-1}$. Note that $\mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}}$ by itself does not represent a (conditional) probability distribution in any particular random experiment with well-defined random variables Y_i, X^i, Y^{i-1} until the system is connected to a distinguisher (see below), in which case these random variables will exist and take the role of the transcript. We shall typically define discrete systems by a high level description, as long as the resulting conditional probability distributions could be derived easily from this description. Two systems \mathbf{F} and \mathbf{G} are called *equivalent* (denoted $\mathbf{F} \equiv \mathbf{G}$) if their input-output behaviors are the same, i.e., $\mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}} = \mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{G}}$ for all $i \geq 1$.

A system \mathbf{F} might often be used as a component (subsystem) in a construction $\mathbf{C}^{(\cdot)}$, resulting in the composed system $\mathbf{C}^{\mathbf{F}}$. $\mathbf{F} \triangleright \mathbf{G}$ denotes the serial composition of systems: every input to $\mathbf{F} \triangleright \mathbf{G}$ is fed to \mathbf{F} , its output is fed to \mathbf{G} and the output of \mathbf{G} is used as the output of $\mathbf{F} \triangleright \mathbf{G}$. In case \mathbf{G} takes as inputs longer bitstrings than \mathbf{F} outputs (as will be the case in the definition of NMAC), the construction $\mathbf{F} \triangleright \mathbf{G}$ pads the outputs of \mathbf{F} with trailing zeroes before passing them to \mathbf{G} .

EXAMPLES. We denote by \mathbf{R} a system that provides access to a function chosen uniformly at random from the set of all functions with domain $\{0, 1\}^{b^*}$ and range $\{0, 1\}^c$. (This unusual domain slightly deviates from the standard definition of \mathbf{R} in the random-systems literature, but will be advantageous for our exposition.) Similarly, for a finite domain $\{0, 1\}^b$ we denote by \mathbf{r} a system realizing a function chosen uniformly from $\mathcal{F}(b, c)$. Finally, we also consider a system \mathbf{f} realizing a function chosen uniformly from $\mathcal{F}(c+b, c)$. We refer to \mathbf{R} , \mathbf{r} and \mathbf{f} as a uniformly random function (URF), a fixed input-length URF, and an ideal compression

function, respectively. In each case the parameters b and c will be clear from the context.

DISTINGUISHERS AND ADVERSARIES. A *distinguisher* \mathbf{D} for an $(\mathcal{X}, \mathcal{Y})$ -random system asking q queries is a $(\mathcal{Y}, \mathcal{X})$ -random system which is “one query ahead:” its input-output behavior is defined by the conditional probability distributions of its queries $\mathbf{p}_{\mathcal{X}_i | \mathcal{X}^{i-1} \mathcal{Y}^{i-1}}^{\mathbf{D}}$ for all $1 \leq i \leq q$. (Its first query is determined by $\mathbf{p}_{\mathcal{X}_1}^{\mathbf{D}}$.) After the distinguisher asks all q queries, it outputs a bit W_q depending on the transcript (X^q, Y^q) . Given a random system \mathbf{F} and a distinguisher \mathbf{D} , we denote by \mathbf{DF} the random experiment where \mathbf{D} interacts with \mathbf{F} , with the distributions of the transcript (X^q, Y^q) and of the bit W_q being uniquely defined by their conditional probability distributions. For two $(\mathcal{X}, \mathcal{Y})$ -random systems \mathbf{F} and \mathbf{G} , the *distinguishing advantage* of \mathbf{D} in distinguishing systems \mathbf{F} and \mathbf{G} by q queries is the quantity $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) = |\mathbf{P}_{W_q}^{\mathbf{DF}}(1) - \mathbf{P}_{W_q}^{\mathbf{DG}}(1)|$ and the maximal distinguishing advantage over all distinguishers asking q queries is denoted by $\Delta_q(\mathbf{F}, \mathbf{G}) = \max_{\mathbf{D}} \Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G})$ (with \mathbf{D} ranging over all such distinguishers).

As opposed to the information-theoretic notion of a distinguisher, we often need to consider an attacker with restricted computational resources. Although such an attacker also participates in a distinguishing experiment, to emphasize this restriction we call it an *adversary* and denote using a sans-serif symbol (e.g. \mathbf{A}). Note that a computationally restricted adversary implicitly defines a random system by its input-output behavior and hence any notation defined for information-theoretic distinguishers is also well-defined for such an adversary. We often restrict the computational power of an adversary by its running time, for this we assume some reasonable fixed model of computation.

MONOTONE CONDITIONS. For a random system \mathbf{F} , we often consider an internal *monotone condition* defined on it. Such a condition is initially satisfied (true), but once it gets violated, it cannot become true again (hence the name monotone). We use such conditions to capture whether the behavior of the system meets some additional requirement (e.g. distinct outputs, consistent outputs) or this was already violated during the interaction that occurred so far. A monotone condition is formalized by a sequence of events $\mathcal{A} = A_0, A_1, \dots$ such that A_0 always holds, and A_i holds if the condition holds after answering the i -th query. The probability that a distinguisher \mathbf{D} issuing q queries to \mathbf{F} makes a monotone condition \mathcal{A} fail in the random experiment \mathbf{DF} is denoted by $\nu^{\mathbf{D}}(\mathbf{F}, \overline{A}_q) = \mathbf{P}^{\mathbf{DF}}(\overline{A}_q)$ and maximum over all such distinguishers is denoted by $\nu(\mathbf{F}, \overline{A}_q) = \max_{\mathbf{D}} \nu^{\mathbf{D}}(\mathbf{F}, \overline{A}_q)$. We also define $\mu(\mathbf{F}, \overline{A}_q) = \max_{x^q} \mathbf{P}_{\overline{A}_q | \mathcal{X}^q}^{\mathbf{F}}(x^q)$ to be the maximal probability of violating the condition \mathcal{A} by a sequence of q non-adaptive queries.

For a random system \mathbf{F} with a monotone condition $\mathcal{A} = A_0, A_1, \dots$ and a random system \mathbf{G} , we say that \mathbf{F} *conditioned on \mathcal{A} is equivalent to \mathbf{G}* , denoted $\mathbf{F} | \mathcal{A} \equiv \mathbf{G}$, if $\mathbf{p}_{\mathcal{Y}_i | \mathcal{X}^i \mathcal{Y}^{i-1} A_i}^{\mathbf{F}} = \mathbf{p}_{\mathcal{Y}_i | \mathcal{X}^i \mathcal{Y}^{i-1}}^{\mathbf{G}}$ for $i \geq 1$, for all arguments for which $\mathbf{p}_{\mathcal{Y}_i | \mathcal{X}^i \mathcal{Y}^{i-1} A_i}^{\mathbf{F}}$ is defined. Intuitively, this captures the fact that as long as the condition \mathcal{A} holds in \mathbf{F} , it behaves the same as \mathbf{G} . The following useful claims

were given in [20], see also [15] for the proof of claim (ii) and [19] for further discussion.

Lemma 1. *Let \mathbf{F} and \mathbf{G} be random systems, let \mathcal{A} be a monotone condition defined on \mathbf{F} , let \mathbf{D} be a distinguisher asking q queries. Then:*

- (i) [20, Lemma 7] *If $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$ then $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) \leq \nu^{\mathbf{D}}(\mathbf{F}, \overline{A_q})$.*
- (ii) [20, Theorem 2] *If $\mathbb{P}_{A_i|X^i Y^{i-1} A_{i-1}}^{\mathbf{F}} = \mathbb{P}_{A_i|X^i A_{i-1}}^{\mathbf{F}}$ for all $i \geq 1$, then $\nu(\mathbf{F}, \overline{A_q}) = \mu(\mathbf{F}, \overline{A_q})$.*

2.2 Message Authentication Codes and PRFs

The standard security requirement for a MAC is *unforgeability under chosen-message attack*. However, it is well-known that any PRF attains this property [6], hence in this paper we focus on PRF-security of the analyzed constructions.

If the first component of the input to a function f is to be seen as a key, we sometimes call f a *keyed* function to emphasize this. For a keyed function $f: \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}$ under a key $k \in \mathcal{K}$ we often write $f_k(\cdot)$ instead of $f(k, \cdot)$. A variable input-length keyed function $\mathbf{G}: \{0, 1\}^c \times \{0, 1\}^{b*} \rightarrow \{0, 1\}^c$ is an:

- $(\varepsilon, t, q, \ell)$ -*secure PRF*, if for any adversary \mathbf{A} running in time t and making at most q queries, each of length at most ℓ (in b -bit blocks), a URF $\mathbf{R}: \{0, 1\}^{b*} \rightarrow \{0, 1\}^c$ and a uniformly random key $K \leftarrow \{0, 1\}^c$, we have $\Delta^{\mathbf{A}}(\mathbf{G}_K, \mathbf{R}) \leq \varepsilon$.
- $(\varepsilon, t, q, \ell)$ -*NA-secure PRF*, if the above is true for all adversaries \mathbf{A} that choose their queries non-adaptively (i.e., \mathbf{A} has to choose its q queries before seeing any of the outputs).
- $(\varepsilon, t, q, \ell)$ -*PF-secure PRF*, if the above is true for all adversaries \mathbf{A} that choose their queries to be prefix-free (i.e., no query is a prefix of another query).
- $(\varepsilon, t, q, \ell)$ -*NA-PF-secure PRF*, if the above is true for all adversaries \mathbf{A} that choose queries *both* non-adaptively and prefix-free.

For fixed input-length functions, we define analogous notions by omitting the parameter ℓ and distinguishing from \mathbf{r} instead of \mathbf{R} . Moreover, we refer to an adversary \mathbf{A} as an $(\varepsilon, t, q, \ell)$ -PRF adversary against \mathbf{G} if it runs in time t , asks at most q queries each consisting of at most ℓ blocks, and achieves the advantage $\Delta^{\mathbf{A}}(\mathbf{G}_K, \mathbf{R}) = \varepsilon$. We refer analogously to adversaries for the other PRF-notions defined above.

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 $\text{Casc}^f(K, m_1 \| \dots \| m_\ell) := y_\ell$ where $y_0 := K$ and for $i \geq 1$ we have $y_i := f(y_{i-1}, m_i)$, in particular $\text{Casc}^f(K, \lambda) := K$.

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,

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

or $\text{NMAC}_{K_1, K_2}^f := \text{Casc}_{K_1}^f \triangleright f_{K_2}$. 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 since it does not affect any of our arguments, we take the customary shortcut and our definition above actually corresponds to the generalized construction denoted as GNMAC in [3] 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:

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

and fpad is a fixed $(b - c)$ -bit padding not affecting the security analysis. (Technically, [18] 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.) As discussed in Section 1.2, we can focus on the PRF-security of NMAC as it translates to analogous results for HMAC under the assumptions stated in [3].

Finally, we also introduce the nested iterated (NI) construction defined in [2]. For this, we consider a keyed compression function $h: \{0, 1\}^k \times \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$. When such h is used in a cascading construction, its c -bit and b -bit inputs are used for the chaining value and the next block, respectively. In contrast to the function f considered above, h has an additional k -bit input that is used for keying. Formally, for such h we define the *nested iterated* construction $\text{NI}^h: (\{0, 1\}^k)^2 \times \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ as

$$\text{NI}_{K_1, K_2}^h(m) := h_{K_2}(\text{Casc}_{\mathbf{0}}^{h_{K_1}}(m), |m|)$$

where $\mathbf{0}$ denotes the all zero bitstring 0^c and $|m|$ is the length of m encoded as a b -bit string. Alternatively, for a function $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and a key K we will denote by LenCasc_K^f a system that given a message m outputs the pair $(\text{Casc}_K^f(m), |m|)$. This allows us to describe NI equivalently as $\text{NI}_{K_1, K_2}^h := \text{LenCasc}_{\mathbf{0}}^{h_{K_1}} \triangleright h_{K_2}$. For a detailed discussion of the relationship of NI to NMAC, see [2].

3 PRF-Security of NMAC

In this section we analyze the PRF security of NMAC^f in terms of the PRF-security of the underlying function f .

3.1 Security Lower Bound

Before moving to the NMAC^f construction, we start by stating a lower bound on the security of the cascade Casc^f when queried on prefix-free inputs. A similar statement has already been proven in [5], and we follow their proof, modifying it where necessary to obtain security against *non-adaptive* adversaries, assuming only *non-adaptive security* of the underlying compression function f . The proof of Proposition 1 is postponed to the full version due to space constraints.

Proposition 1 (Casc^f as a NA-PF-PRF). *Let $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ be a compression function. There exists an explicit reduction T (described in the proof) such that for any $(\varepsilon', t', q, \ell)$ -NA-PF-PRF adversary A against Casc^f , T^{A} is an $(\varepsilon_{\text{na}}, t, q)$ -NA-PRF adversary against f such that*

$$\varepsilon' \leq \ell q \varepsilon_{\text{na}} \quad \text{and} \quad t = t' + \tilde{O}(\ell q) .$$

This allows us to present our main result in this section, which relates the adaptive PRF-security of the construction NMAC^f to both the adaptive and non-adaptive PRF-security of f .

Theorem 1 (NMAC^f as a PRF). *Let $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ be a compression function. There exist explicit reductions T_1 and T_2 (described in the proof) such that for any $(\varepsilon', t', q, \ell)$ -PRF adversary A against NMAC^f ,*

1. $\mathsf{T}_1^{\mathsf{A}}$ is an (ε, t, q) -PRF adversary against f ,
2. $\mathsf{T}_2^{\mathsf{A}}$ is an $(\varepsilon_{\text{na}}, t, q)$ -NA-PRF adversary against f ,

and their parameters satisfy

$$\varepsilon' \leq \varepsilon + (\ell + 1)q\varepsilon_{\text{na}} + \frac{q^2}{2^c} \quad \text{and} \quad t = t' + \tilde{O}(\ell q) .$$

Proof. Let A be a PRF-adversary running in time t' and asking q queries, each of length at most ℓ blocks. Let $\mathbf{r}: \{0, 1\}^b \rightarrow \{0, 1\}^c$, $\mathbf{R}: \{0, 1\}^{b^*} \rightarrow \{0, 1\}^c$ and $K = (K_1, K_2) \leftarrow \{0, 1\}^c \times \{0, 1\}^c$ denote a fixed input-length URF, a URF and a key pair chosen independently at random, respectively.

We turn A into an adversary $\mathsf{T}_1^{\mathsf{A}}$ against the PRF-security of f_K as follows: Given access to g (which is either f_K or \mathbf{r}), sample some key K_1 at random, and then invoke A , answering its queries with $\text{Casc}_{K_1}^f \triangleright g$. Finally, output the decision bit of A . Clearly we have $\Delta^{\mathsf{A}}(\text{Casc}_{K_1}^f \triangleright f_{K_2}, \text{Casc}_{K_1}^f \triangleright \mathbf{r}) = \Delta^{\mathsf{T}_1^{\mathsf{A}}}(\mathbf{r}, \mathbf{r})$ and if we denote $\Delta^{\mathsf{T}_1^{\mathsf{A}}}(\mathbf{r}, \mathbf{r})$ by ε then using triangle inequality we get

$$\Delta^{\mathsf{A}}(\text{NMAC}_K^f, \mathbf{R}) = \Delta^{\mathsf{A}}(\text{Casc}_{K_1}^f \triangleright f_{K_2}, \mathbf{R}) \leq \varepsilon + \Delta^{\mathsf{A}}(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \mathbf{R}) .$$

In the experiment where A interacts with $\text{Casc}_{K_1}^f \triangleright \mathbf{r}$, let C_i denote the event that during the first i queries to $\text{Casc}_{K_1}^f \triangleright \mathbf{r}$, for any two distinct queries M and M' the values $\text{Casc}_{K_1}^f(M)$ and $\text{Casc}_{K_1}^f(M')$ (inputs to the final \mathbf{r} -call) are also distinct. As long as the monotone condition $\mathcal{C} = C_0, C_1, \dots$ remains satisfied, the

responses of $\text{Casc}_{K_1}^f \triangleright \mathbf{r}$ to distinct queries are equivalent to outputs of \mathbf{r} on distinct inputs, and thus independent, uniformly random values, in particular $(\text{Casc}_{K_1}^f \triangleright \mathbf{r}) | \mathcal{C} \equiv \mathbf{R}$. We can therefore apply Lemma 1(i) to conclude that distinguishing $\text{Casc}^f \triangleright \mathbf{r}$ from a URF \mathbf{R} is at least as hard as making the condition \mathcal{C} fail, i.e.,

$$\Delta^A(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \mathbf{R}) \leq \nu^A(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \overline{\mathcal{C}}) .$$

Below we explain how to use the adversary A to construct⁴ a *non-adaptive* adversary A_{na} such that

$$\nu^A(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \overline{\mathcal{C}}) = \nu^{A_{\text{na}}}(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \overline{\mathcal{C}}) . \quad (1)$$

A_{na} simply runs A and responds to all its fresh queries by fresh random values, while answering repeated queries consistently. In the end, A_{na} (non-adaptively) asks all the queries that A asked during this simulated interaction. The equation (1) follows from the fact that the simulation for A is perfect as long as its queries do not violate \mathcal{C} . Since \mathcal{C} is defined on $\text{Casc}_{K_1}^f$ and A_{na} is non-adaptive, we additionally have

$$\nu^{A_{\text{na}}}(\text{Casc}_{K_1}^f \triangleright \mathbf{r}, \overline{\mathcal{C}}) = \nu^{A_{\text{na}}}(\text{Casc}_{K_1}^f, \overline{\mathcal{C}}) .$$

Next, for A_{na} we can construct another non-adaptive adversary A_{pf} that violates the condition \mathcal{C} (i.e., creates a collision in the outputs of $\text{Casc}_{K_1}^f$) with the same probability as A_{na} , but all its queries are *prefix-free*. This can be done, for example, by simply appending an additional block to all queries asked by A_{na} , such that this block does not appear in the original queries. Hence we have

$$\nu^{A_{\text{na}}}(\text{Casc}_{K_1}^f, \overline{\mathcal{C}}) = \nu^{A_{\text{pf}}}(\text{Casc}_{K_1}^f, \overline{\mathcal{C}})$$

for a non-adaptive adversary A_{pf} asking prefix-free queries of length at most $\ell+1$.

Finally, consider the non-adaptive adversary A^* that simply asks the same prefix-free queries as A_{pf} and then outputs 1 if and only if the responses to these queries contain a collision. Then A^* interacting with $\text{Casc}_{K_1}^f$ outputs 1 with probability $\nu^{A_{\text{pf}}}(\text{Casc}_{K_1}^f, \overline{\mathcal{C}})$, while in an interaction with \mathbf{R} it outputs 1 with probability at most $q^2/2^c$ via the well-known birthday bound. Hence, by the definition of $\Delta^{A^*}(\text{Casc}_{K_1}^f, \mathbf{R})$, we have

$$\nu^{A_{\text{pf}}}(\text{Casc}_{K_1}^f, \overline{\mathcal{C}}) \leq \Delta^{A^*}(\text{Casc}_{K_1}^f, \mathbf{R}) + \frac{q^2}{2^c} .$$

Since A^* is non-adaptive and prefix-free, we can now employ the reduction T guaranteed by Proposition 1 to obtain an NA-PRF adversary T^{A^*} against f such that

$$\Delta^{A^*}(\text{Casc}_{K_1}^f, \mathbf{R}) \leq (\ell+1)q \cdot \Delta^{T^{A^*}}(f, \mathbf{r}) .$$

Putting $T_2^A := T^{A^*}$ hence concludes the proof of Theorem 1. \square

⁴ One could use a lemma from the random system framework [20] in the spirit of Lemma 1(ii) to switch to non-adaptivity. We prefer to spell out the actual construction to emphasize the uniformity of our reduction.

Corollary 1. *If $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ is an (ε, t, q) -secure PRF and an $(\varepsilon_{\text{na}}, t, q)$ -NA-secure PRF, then NMAC^f is an $(\varepsilon', t', q, \ell)$ -secure PRF with*

$$\varepsilon' = \varepsilon + (\ell + 1)q\varepsilon_{\text{na}} + \frac{q^2}{2^c} \quad \text{and} \quad t = t' + \tilde{O}(\ell q) .$$

3.2 Matching Attacks

We now argue that the bound obtained in Theorem 1 is essentially tight. First, we show that the term $\ell q\varepsilon_{\text{na}}$ is unavoidable (up to a constant factor) by constructing a particular compression function f , which is an $(\varepsilon_{\text{na}}, t, q)$ -NA-secure PRF, yet there is a simple attack against the PRF-security of NMAC^f achieving advantage roughly $\ell q\varepsilon_{\text{na}}$.

Proposition 2. *Let b, c, ℓ be positive integers such that $b \geq c$, let $\varepsilon_{\text{na}} \in (0, 1)$, and moreover, assume that pseudo-random functions exist. Then there exists a function $f: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and an adversary A against NMAC^f such that for any q that satisfies $\varepsilon_{\text{na}} = \omega(q^2 2^{-b}, 2^{-c})$, we have:*

- f is $(\varepsilon_{\text{na}}, t, q)$ -NA-secure PRF;
- the adversary A , when asking q queries of length ℓ blocks each, runs in time $\tilde{O}(\ell q)$ and achieves distinguishing advantage

$$\Delta^A(\text{NMAC}_K^f, \mathbf{R}) = \Theta(\ell q\varepsilon_{\text{na}}) .$$

In particular, NMAC^f is not an $(o(\ell q\varepsilon_{\text{na}}), \tilde{O}(\ell q), q, \ell)$ -secure PRF.

Proof (sketch). Here we only describe the high-level idea for constructing f and A and defer the discussion of the technical obstacles in implementing this idea to the full version.

Roughly speaking, we construct an $(\varepsilon_{\text{na}}, t, q)$ -NA-secure PRF f that behaves pseudo-randomly for all keys except for a small, $\varepsilon_{\text{na}}/2$ -fraction of them. We denote the set of these keys by \mathcal{K} and refer to them as the *weak keys*. Under any weak key k , the function $f(k, \cdot)$ outputs some constant value $w \in \mathcal{K}$ irrespective of its input.

To attack the NA-PRF security of $\text{NMAC}_{K=(K_1, K_2)}^f$, consider a pair of messages M_1, M_2 chosen by sampling $M \leftarrow \{0, 1\}^{b(\ell-1)}$ at random and then setting $M_1 = M \| x_1$ and $M_2 = M \| x_2$ for some distinct blocks $x_1, x_2 \in \{0, 1\}^b$. If some of the $\ell - 1$ intermediate values in the evaluation of the inner function $\text{Casc}^f(K_1, M)$ is in \mathcal{K} , then all following intermediate values are w , and in particular we have $\text{Casc}^f(K_1, M_i) = w$ for both $i \in \{1, 2\}$, and hence also $\text{NMAC}^f(K, M_1) = \text{NMAC}^f(K, M_2) = f_{K_2}(w)$. This implies that it is much more likely to get a collision for a pair of messages as described above for NMAC_K^f than for \mathbf{R} . Our adversary A simply chooses $q/2$ message pairs at random as above, and it outputs 1 if it observes a collision for at least one of those pairs. As there are $q/2$ message pairs, each of length ℓ , we have a total of $\ell q/2$ possibilities to “hit” a weak key, each having probability ε_{na} . By the union bound this gives us

a total probability of $\Theta(\ell q \varepsilon_{\text{na}})$ for observing a collision when querying NMAC_K^f . On the other hand the probability of observing a colliding pair in \mathbf{R} is only $O(q/2^c)$. \square

We now consider the tightness of the bound in Theorem 1 when $\varepsilon \gg \ell q \varepsilon_{\text{na}}$ is the dominating term. This is the case when the best adaptive attack against \mathbf{f} is by more than a factor ℓq better than any non-adaptive attack.

In [23] a pair $\mathbf{g}_1, \mathbf{g}_2$ of PRFs is constructed such that \mathbf{g}_1 and \mathbf{g}_2 are ε_{na} -secure *non-adaptive* PRFs for some negligible ε_{na} , and the serial composition $\mathbf{g}_1 \triangleright \mathbf{g}_2$ with independent keys can be broken by an *adaptive* attack (in a constant number of queries) with advantage almost 1.⁵ From such $\mathbf{g}_1, \mathbf{g}_2$ we can get a single PRF \mathbf{f} which is an ε_{na} -secure NA-PRF for a negligible ε_{na} , an ε -secure PRF for any ε of our choice, and where $\mathbf{f} \triangleright \mathbf{f}$ is not $\Theta(\varepsilon^2)$ -secure, by setting $\mathbf{f} := \mathbf{g}_1$ and $\mathbf{f} := \mathbf{g}_2$ with probability $\varepsilon/2$, respectively, and some strong standard PRF with probability $1 - \varepsilon$ (over the choice of the key). We now observe that NMAC_K^f computed on single-block messages is simply a cascade of two \mathbf{f} 's with independent keys. Thus, when using the above ε -secure PRF \mathbf{f} , we can break NMAC_K^f with advantage $\Theta(\varepsilon^2)$. This shows that the ε term in Theorem 1 is necessary if ε is constant as then $\Theta(\varepsilon) = \Theta(\varepsilon^2) = \Theta(1)$. We conjecture that $\Theta(\varepsilon^2)$ is the correct value, and the ε term in the lower bound can be improved to $\Theta(\varepsilon^2)$ using security amplification techniques along the lines of [22, 25].

4 PRF-Security of the NI Construction

In this section we analyze the PRF-security of the NI^h construction under the assumption that the keyed compression function \mathbf{h} is a PRF (when keyed via its k -bit input).

Theorem 2. *If $\mathbf{h}: \{0, 1\}^k \times \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ is an (ε_1, t, q) -secure PRF and an $(\varepsilon_2, t, \ell q)$ -secure PRF, then NI^h is an $(\varepsilon', t', q, \ell)$ -secure PRF with*

$$\varepsilon' = \varepsilon_1 + \varepsilon_2 + \frac{q^2}{2^c} \cdot \left(\ell + \frac{64\ell^4}{2^c} \right) \quad \text{and} \quad t = t' + \tilde{O}(\ell q).$$

Proof. We prove Theorem 2 in four consecutive steps. First, we use the PRF-security of \mathbf{h} to replace it by an ideal compression function, making the rest of our analysis information-theoretic. Second, we observe that the resulting system behaves identically to \mathbf{R} as long as no non-trivial collision occurs in the outputs of the initial cascade. Third, we reduce estimating the probability of such a collision to a counting problem of upper-bounding the number of graphs satisfying certain properties (modeling the computation of the cascade). Finally, we give a bound on the number of these graphs, hence concluding the argument.

⁵ The NA-PRF security of this construction relies on the DDH assumption, [9] construct such a PRF under the weaker assumption that “uniform transcript key-agreement” exists, and this assumption is necessary [24].

FROM A PRF TO A RANDOM FUNCTION. Let \mathbf{A} be a PRF-adversary against NI^h running in time t and asking q queries, each of length at most ℓ blocks. To simplify the notation let $\mathbf{0} := 0^c$. By a standard argument as in the proof of Theorem 1, we have

$$\Delta^A(\text{NI}_K^h, \mathbf{R}) = \Delta^A\left(\text{LenCasc}_0^{h_{K_1}} \triangleright h_{K_2}, \mathbf{R}\right) \leq \varepsilon_1 + \varepsilon_2 + \Delta^A\left(\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \mathbf{R}\right) \quad (2)$$

where $K = (K_1, K_2) \leftarrow (\{0, 1\}^k)^2$ is a uniformly random key and \mathbf{f}_1 and \mathbf{f}_2 are two independent ideal compression functions. Interestingly, the system $\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2$ is very similar to NMAC with an ideal compression function and $(\mathbf{0}, |m|)$ being used instead of the key pair.

BOUND VIA COLLISION PROBABILITY. Let $\text{CColl}(\ell)$ denote the probability that a random choice of the compression function \mathbf{f}_1 results in a collision in $\text{Casc}_0^{\mathbf{f}_1}$, maximized over the choice of the two distinct, equal-length inputs m_1, m_2 consisting of at most ℓ blocks each. (Note that we require length equality $|m_1| = |m_2|$ to obtain a collision also for $\text{LenCasc}_0^{\mathbf{f}_1}$.) Formally, for uniformly random $\mathbf{f}_1 \leftarrow \mathcal{F}(c + b, c)$ we define

$$\text{CColl}(\ell) := \max_{\substack{m_1 \neq m_2 \\ |m_1| = |m_2| \leq \ell b}} \text{P}^{\mathbf{f}_1} \left[\text{Casc}_0^{\mathbf{f}_1}(m_1) = \text{Casc}_0^{\mathbf{f}_1}(m_2) \right]. \quad (3)$$

In the experiment where \mathbf{A} interacts with $\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2$, let E_i denote the event that during the first i queries to $\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2$, for any two distinct queries M and M' the values $\text{LenCasc}_0^{\mathbf{f}_1}(M)$ and $\text{LenCasc}_0^{\mathbf{f}_1}(M')$ (inputs to the final \mathbf{f}_2 -call) were also distinct. As long as the monotone condition $\mathcal{E} = E_0, E_1, \dots$ remains satisfied, the responses of $\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2$ to distinct queries are clearly independent, uniformly random values thanks to \mathbf{f}_2 . Hence, we have $(\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2) | \mathcal{E} \equiv \mathbf{R}$ and $\text{p}_{E_i | X^i Y^{i-1} E_{i-1}}^{\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2} = \text{p}_{E_i | X^i E_{i-1}}^{\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2}$ and can therefore consecutively apply Lemma 1(i), Lemma 1(ii), and finally the union bound to get

$$\Delta^A(\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \mathbf{R}) \leq \nu(\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \overline{E_q}) \leq \mu(\text{LenCasc}_0^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \overline{E_q}) \leq q^2 \cdot \text{CColl}(\ell). \quad (4)$$

GRAPH-BASED REPRESENTATION OF Casc . The probability $\text{CColl}(\ell)$ could trivially be upper-bounded by $O(\ell^2/2^c)$ using a union-bound argument, achieving a non-trivial and significantly better bound on $\text{CColl}(\ell)$ is the central part of our proof. To this end, we use an approach inspired by [7] and represent the computation of $\text{Casc}_0^{\mathbf{f}_1}$ on various inputs by directed graphs.

Let m_1 and m_2 be two distinct, equal-length messages that can be parsed into b -bit blocks as $m_i = m_i^1 \| \dots \| m_i^{\ell'}$ for some $\ell' \leq \ell$, and let $\Lambda := 2\ell'$. For convenience, we use the notation $m^{(i)}$ as a reference to the block m_1^i if $i \leq \ell'$, otherwise it denotes the block $m_2^{i-\ell'}$. For any fixed compression function $f \in \mathcal{F}(c + b, c)$ and a pair of such messages $\mathcal{M} = (m_1, m_2)$, we define the *structure graph* $G_f^{\mathcal{M}}$ to be the triple $G_f^{\mathcal{M}} = (\mathcal{V}, \mathcal{E}, \mathcal{L})$, such that:

– $(\mathcal{V}, \mathcal{E})$ is a directed graph. To describe it, let

$$s_i := \begin{cases} \mathbf{0} & \text{for } i = 0 \\ f(s_{i-1}, m_1^i) & \text{for } 1 \leq i \leq \ell' \\ f(\mathbf{0}, m_2^1) & \text{for } i = \ell' + 1 \\ f(s_{i-1}, m_2^{i-\ell'}) & \text{for } \ell' + 2 \leq i \leq \Lambda \end{cases} \quad (5)$$

and consider the mappings $[\cdot]_G$ and $[\cdot]'_G$ defined on $\{0, \dots, \Lambda\}$ such that $[i]_G := \min\{j : s_i = s_j\}$ (so $[i]_G = i$ if and only if s_i is “fresh”) and $[i]'_G := [i]_G$ for $i \neq \ell'$, while $[\ell']'_G := 0$. Now we let

$$\mathcal{V} := \{[i]_G : 0 \leq i \leq \Lambda\} \quad \text{and} \quad \mathcal{E} := \{([i-1]'_G, [i]_G) : 1 \leq i \leq \Lambda\}.$$

– $\mathcal{L} : \mathcal{V}^2 \rightarrow \text{Pow}(\{0, 1\}^b)$ is a labeling function that labels every edge $(u, v) \in \mathcal{E}$ with the set $\{m^{(i)} : [i-1]'_G = u \wedge [i]_G = v\}$ and every pair of vertices that do not form an edge with the empty set \emptyset (to simplify our notation later).

Intuitively, if all the values s_i are distinct, $G_f^{\mathcal{M}}$ simply consists of two directed paths starting in the root vertex 0, representing the evaluation of $\text{Casc}_{\mathbf{0}}^{f_1}$ on the messages m_1 and m_2 (the edges are labeled by the corresponding blocks). If some collisions among the values s_i occur, one can obtain the graph $G_f^{\mathcal{M}}$ by collapsing every pair of vertices i, j where $s_i = s_j$ into one vertex labeled $\min\{i, j\}$, as well as merging the edge labels in the natural way.

Let $\mathcal{G}(\mathcal{M}) := \{G_f^{\mathcal{M}} : f \in \mathcal{F}(c+b, c)\}$ denote the set of all structure graphs associated with the message pair \mathcal{M} . Note that the uniformly distributed random variable $F \leftarrow \mathcal{F}(c+b, c)$ also induces a distribution on $\mathcal{G}(\mathcal{M})$, therefore we denote by $G_F^{\mathcal{M}}$ the resulting random variable (taking on structure graphs as values). Similarly, F also induces a distribution on the values s_i defined above and we denote the resulting random variables S_i .

For a fixed structure graph $G = G_f^{\mathcal{M}}$ we denote by $G_i = (\mathcal{V}_i, \mathcal{E}_i, \mathcal{L}_i)$ the graph that is obtained after processing only the first i out of Λ blocks of \mathcal{M} . More formally, $G_i := G_f^{\mathcal{M}'}$ where $\mathcal{M}' := (m_1^1 \parallel \dots \parallel m_1^i, \lambda)$ if $i \leq \ell'$ and $\mathcal{M}' := (m_1, m_2^1 \parallel \dots \parallel m_2^{i-\ell'})$ otherwise. Building on this notion, we call $\text{fColl}(G)$ the set of f -collisions that occurred in G :

$$\text{fColl}(G) := \left\{ (i, [i]_G) : [i]_G < i \wedge m^{(i)} \notin \mathcal{L}_{i-1}([i-1]'_G, [i]_G) \right\}. \quad (6)$$

Informally, imagine we reveal the structure graph G step by step, i.e., by a sequence of transitions from G_{i-1} to G_i , for $i = 1, \dots, \Lambda$. The pair $(i, [i]_G)$ belongs to $\text{fColl}(G)$ (and we say that the i -th step caused an f -collision), if during this step, instead of adding a new vertex, we arrive at a vertex already visited, while not following an existing edge already labeled with $m^{(i)}$ (i.e., not repeating a step we have made before).

PROPERTIES OF STRUCTURE GRAPHS. We first upper-bound the probability of $G_F^{\mathcal{M}}$ taking the form of any particular fixed structure graph $g \in \mathcal{G}(\mathcal{M})$. The following result is inspired by Lemma 8 from [7]. Due to space constraints, we postpone the proofs of all technical lemmas below to the full version of this paper.

Lemma 2. *Let $F \leftarrow \mathcal{F}(c + b, c)$ be chosen uniformly at random. For a fixed graph $g \in \mathcal{G}(\mathcal{M})$ we have*

$$\mathbb{P}^F [G_F^{\mathcal{M}} = g] \leq 2^{-c \cdot |\text{fColl}(g)|} .$$

Using Lemma 2, it is easy to see that the event that at least two f -collisions occur in G is highly unlikely.

Lemma 3. *Let $F \leftarrow \mathcal{F}(c + b, c)$ be chosen uniformly at random. Then*

$$\mathbb{P}^F [|\text{fColl}(G_F^{\mathcal{M}})| \geq 2] \leq \frac{4A^4}{2^{2c}} .$$

FROM COLLISION PROBABILITY TO COUNTING GRAPHS. We can now proceed to upper-bounding the value $\text{CColl}(\ell)$. Let $\mathcal{M} := (m_1, m_2)$ be the two distinct, equal-length messages of length at most ℓ blocks that maximize the probability $\text{CColl}(\ell) := \max_{m_1 \neq m_2} \mathbb{P}^F [\text{Casc}_0^F(m_1) = \text{Casc}_0^F(m_2)]$. For $j \in \{1, 2\}$ let V_j^i be the random variable denoting the i -th vertex (counting from 0) in the path corresponding to m_j in $G_F^{\mathcal{M}}$ (randomness taken over the uniform choice of F). Formally, $V_1^i := [i]_G$ and $V_2^i := [\ell' + i]_G$. Using this notation, we have $\text{CColl}(\ell) = \mathbb{P}[V_1^{\ell'} = V_2^{\ell'}]$. Since $m_1 \neq m_2$, $V_1^{\ell'} = V_2^{\ell'}$ cannot occur without any f -collision, hence we can split $\text{CColl}(\ell)$ into

$$\mathbb{P} [V_1^{\ell'} = V_2^{\ell'} \wedge |\text{fColl}(G_F^{\mathcal{M}})| = 1] + \mathbb{P} [V_1^{\ell'} = V_2^{\ell'} \wedge |\text{fColl}(G_F^{\mathcal{M}})| \geq 2] . \quad (7)$$

The latter probability can be readily upper-bounded by $4A^4/2^{2c}$ using Lemma 3. As for the former, let us denote by $\mathcal{H}(\mathcal{M})$ the set of structure graphs for \mathcal{M} that contain exactly one f -collision and where the vertices $V_1^{\ell'}$ and $V_2^{\ell'}$ coincide. The first term in (7) can then be upper-bounded by $|\mathcal{H}(\mathcal{M})|/2^c$ using Lemma 2, hence it remains to bound the size of the set $\mathcal{H}(\mathcal{M})$.

COUNTING THE STRUCTURE GRAPHS. We give such a bound in the following lemma, proven in the full version of this paper.

Lemma 4. *For two distinct, equal-length messages $\mathcal{M} = \{m_1, m_2\}$ each of length at most ℓ blocks, we have $|\mathcal{H}(\mathcal{M})| \leq \ell$.*

Finally, combining the equations (2), (4), (7), and the bounds obtained in Lemma 3 and Lemma 4, we get

$$\Delta^A(\text{NI}_K^h, \mathbf{R}) \leq \varepsilon_1 + \varepsilon_2 + q^2 \cdot \left(\frac{\ell}{2^c} + \frac{4A^4}{2^{2c}} \right) \leq \varepsilon_1 + \varepsilon_2 + \frac{q^2}{2^c} \cdot \left(\ell + \frac{64\ell^4}{2^c} \right)$$

and conclude the proof of Theorem 2. \square

In the full version we also show that Lemma 4 is tight, and discuss the implications for the tightness of Theorem 2. Moreover, we show a generalization of Lemma 4 that does not require the messages in \mathcal{M} to have the same length, in which case we prove $|\mathcal{H}(\mathcal{M})| \leq \ell d'(\ell)$. This translates directly into a PRF-security statement for a variant of NI that does not include the message length in its last h -call, giving a bound that is equivalent to Theorem 2 except for the term $\ell q^2/2^c$ that is replaced by $\ell d'(\ell)q^2/2^c$.

Acknowledgements. We thank the anonymous reviewers for useful comments and suggestions. This work was partly funded by the European Research Council under an ERC Starting Grant (259668-PSPC).

References

1. Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple construction of almost k -wise independent random variables. *Random Struct. Algorithms*, 3(3):289–304, 1992.
2. Jee Hea An and Mihir Bellare. Constructing VIL-MACs from FIL-MACs: Message authentication under weakened assumptions. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO’99*, volume 1666 of *Lecture Notes in Computer Science*, pages 252–269. Springer, August 1999.
3. Mihir Bellare. New proofs for NMAC and HMAC: Security without collision-resistance. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 602–619. Springer, August 2006.
4. Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Keying hash functions for message authentication. In Neal Koblitz, editor, *Advances in Cryptology – CRYPTO’96*, volume 1109 of *Lecture Notes in Computer Science*, pages 1–15. Springer, August 1996.
5. Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th Annual Symposium on Foundations of Computer Science*, pages 514–523. IEEE Computer Society Press, October 1996.
6. Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, 61(3):362–399, 2000.
7. Mihir Bellare, Krzysztof Pietrzak, and Phillip Rogaway. Improved security analyses for CBC MACs. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 527–545. Springer, August 2005.
8. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426. Springer, May / June 2006.
9. Chongwon Cho, Chen-Kuei Lee, and Rafail Ostrovsky. Equivalence of uniform key agreement and composition insecurity. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 447–464. Springer, August 2010.
10. Jean-Sébastien Coron, Yevgeniy Dodis, Cécile Malinaud, and Prashant Puniya. Merkle-Damgård revisited: How to construct a hash function. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer, August 2005.
11. Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 649–665. Springer, August 2010.

12. 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, *Advances in Cryptology – CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 348–366. Springer, August 2012.
13. G. H. Hardy and Edward M. Wright. *An Introduction to the Theory of Numbers (sixth edition)*. Oxford University Press, USA, 2008.
14. Martin E. Hellman. A cryptanalytic time-memory trade-off. *IEEE Transactions on Information Theory*, 26(4):401–406, 1980.
15. Dimitar Jetchev, Onur zen, and Martijn Stam. Understanding adaptivity: Random systems revisited. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT12*, volume 7658 of *Lecture Notes in Computer Science*, pages 313–330. Springer, 2012.
16. 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. In Roberto Prisco and Moti Yung, editors, *Security and Cryptography for Networks*, volume 4116 of *Lecture Notes in Computer Science*, pages 242–256. Springer, 2006.
17. Neal Koblitz and Alfred Menezes. Another look at HMAC. Cryptology ePrint Archive, Report 2012/074, 2012.
18. Hugo Krawczyk, Mihir Bellare, and Ran Canetti. HMAC: Keyed-Hashing for Message Authentication. IETF Internet Request for Comments 2104, February 1997.
19. Ueli Maurer. Conditional equivalence of random systems and indistinguishability proofs. In *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 3150–3154, July 2013.
20. Ueli M. Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. Springer, April / May 2002.
21. Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004: 1st Theory of Cryptography Conference*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer, February 2004.
22. Ueli M. Maurer and Stefano Tessaro. Computational indistinguishability amplification: Tight product theorems for system composition. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 355–373. Springer, August 2009.
23. Krzysztof Pietrzak. Composition does not imply adaptive security. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 55–65. Springer, August 2005.
24. Krzysztof Pietrzak. Composition implies adaptive security in minicrypt. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 328–338. Springer, May / June 2006.
25. Stefano Tessaro. Security amplification for the cascade of arbitrarily weak PRPs: Tight bounds via the interactive hardcore lemma. In Yuval Ishai, editor, *TCC 2011: 8th Theory of Cryptography Conference*, volume 6597 of *Lecture Notes in Computer Science*, pages 37–54. Springer, March 2011.
26. Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding collisions in the full SHA-1. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 17–36. Springer, August 2005.
27. Xiaoyun Wang and Hongbo Yu. How to break MD5 and other hash functions. In Ronald Cramer, editor, *Advances in Cryptology – EUROCRYPT 2005*, volume 3494 of *Lecture Notes in Computer Science*, pages 19–35. Springer, May 2005.