The Exact PRF-Security of NMAC and HMAC*

Peter Gaži, Krzysztof Pietrzak, and Michal Rybár

IST Austria

August 2014

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: If f is an ε-secure PRF (against q queries) and a δ-*non-adaptively* secure PRF (against q queries), then $\mathsf{NMAC}^{\mathsf{f}}$ 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 Koblitz 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]. We also analyze a variant of NI that does not include the message length in the last call to the compression function, proving a $\ell^{1+o(1)} q^2/2^c$ bound in this case.

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 $\mathsf{f} : \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ (with $b \geq c$) into a keyed variable input-length function $\mathsf{NMAC}^{\mathsf{f}} : \{0,1\}^{2c} \times \{0,1\}^{b*} \to \{0,1\}^c$ (where $\{0,1\}^{b*}$ denotes all bit strings whose length is a multiple of b) as

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

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

$$\mathsf{Casc}^{\mathsf{f}}(K_1, m_1 \| \ldots \| m_\ell) := \mathsf{f}(\ldots \mathsf{f}(\mathsf{f}(K_1, m_1), m_2) \ldots m_\ell) \; .$$

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

* A preliminary version of this paper appears in the proceedings of CRYPTO 2014, this is the full version. This work was partly funded by the European Research Council under an ERC Starting Grant (259668-PSPC).

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 [19]. 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) $\mathsf{Casc}^{\mathsf{f}}$ 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 $\mathsf{Casc}^{\mathsf{f}}(K, M) = \mathsf{Casc}^{\mathsf{f}}(K, M')$ under a random key K, given oracle access to $\mathsf{Casc}^{\mathsf{f}}(K, .)$ (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 $\mathsf{Casc}^{\mathsf{f}}$. However, both of these have been found not to satisfy the WCR notion [32,33], 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 [17]).

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 $\mathsf{NMAC}^{\mathsf{f}}$ 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]

1.1 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 $\mathsf{NMAC}^{\mathsf{f}}$ assuming only that f is a PRF: If f is an ε-secure PRF against q queries, then $\mathsf{NMAC}^{\mathsf{f}}$ 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 $\mathsf{NMAC}^{\mathsf{f}}$ 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 $\mathsf{Casc}^{\mathsf{f}}$ 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 $\mathsf{NMAC}^{\mathsf{f}}$ against arbitrary adaptive queries to the security of $\mathsf{Casc}^{\mathsf{f}}$ against non-adaptive prefix-free queries.

[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 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.

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 $\mathsf{NMAC}^{\mathsf{f}}$ 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 [29] who shows that cascading non-adaptively secure PRFs does not give an adaptively secure PRF in general, to construct an ε-secure f where $\mathsf{NMAC}^{\mathsf{f}}$ 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 [25,31].

PRF-Security Proof for NI. The main difficulty in security analyses of $\mathsf{NMAC}^{\mathsf{f}}$ and $\mathsf{HMAC}^{\mathsf{f}}$ based on the PRF-security of the underlying compression function f is that both these constructions are constantly rekeying f during the evaluation of $\mathsf{Casc}^{\mathsf{f}}$, 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.

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 $\mathsf{NMAC}^{\mathsf{f}}$, 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 indifferentiability [24,10] of HMAC from a (keyed) random oracle. This result is incomparable to ours: While the stronger notion of indifferentiability covers the settings where HMAC is not used as a PRF, the bound achieved in [12] is understandably much weaker, being $\Theta(\ell^2 q^2/2^c)$.

Generic Attacks. There is also a recent line of work investigating generic attacks against iterated hash-based MACs [27,20,26,28]. These works present various attacks against MACs (e.g. related-key attack, universal forgeries, state recovery) that do not exploit the inner structure and potential weaknesses of the compression function, instead they rely solely on the iterative structure of the MACs.

Another look at [18]. As already mentioned, Bellare [3] proved that $\mathsf{NMAC}^{\mathsf{f}}$ 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 [18], Koblitz 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 community[2] 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 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 [18] also claimed that HMAC is an $\varepsilon\ell$-secure PRF, a bound that is falsified by an attack given in this paper. In response, [18] 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 do denote the empty string. We use $\{0,1\}^{b*} := \bigcup_{z\geq 0}\{0,1\}^{bz}$ to 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 $\mathsf{Pow}(\mathcal{S})$ 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

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

is the maximum, over all positive integers $n' \leq n$, of the number of positive divisors of n'. More precisely, we have $\forall \varepsilon > 0\ \exists n_0\ \forall n > n_0 \colon d(n) < n^{(1+\varepsilon)/\ln\ln n}$ [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 $X, Y, \ldots$ and lower-case letters $x, y, \ldots$, 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 $\mathsf{P}_{UA|VB}$ 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 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.

corresponding conditional probability distribution, seen as a (partial) function $\mathcal{U} \times \mathcal{V} \to [0,1]$. The value $\mathsf{P}_{UA|VB}(u,v) = \mathsf{P}[U = u \wedge A | V = v \wedge B]$ is well-defined for all $u \in \mathcal{U}$ and $v \in \mathcal{V}$ such that $\mathsf{P}_{VB}(v) > 0$ and undefined otherwise. Two probability distributions P_U and $\mathsf{P}_{U'}$ on the same set $\mathcal{U}$ are equal, denoted $\mathsf{P}_U = \mathsf{P}_{U'}$, if $\mathsf{P}_U(u) = \mathsf{P}_{U'}(u)$ for all $u \in \mathcal{U}$. Conditional 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. $\mathsf{P}^{\mathcal{E}}_{U|V}(u,v)$. If the distribution of a random variable U is clear from the context, we also sometimes write 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 [23], 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 $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}} \colon \mathcal{Y} \times \mathcal{X}^i \times \mathcal{Y}^{i-1} \to [0,1]$ such that $\sum_{y_i} \mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}(y_i, x^i, y^{i-1}) = 1$ for all $i \geq 1$, $x^i \in \mathcal{X}^i$ and $y^{i-1} \in \mathcal{Y}^{i-1}$. Note that $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$ 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., $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}} = \mathsf{p}^{\mathbf{G}}_{Y_i|X^iY^{i-1}}$ 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 $\mathsf{p}^{\mathbf{D}}_{X_i|X^{i-1}Y^{i-1}}$ for all $1 \leq i \leq q$. (Its first query is determined by $\mathsf{p}^{\mathbf{D}}_{X_1}$.) 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 $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) = |\mathsf{P}^{\mathbf{DF}}_{W_q}(1) - \mathsf{P}^{\mathbf{DG}}_{W_q}(1)|$ 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. 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 $\mathcal{A} = A_0, A_1, \ldots$ 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 $\nu^{\mathbf{D}}(\mathbf{F}, \overline{A_q}) = \mathsf{P}^{\mathbf{DF}}(\overline{A_q})$ 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 $\mu(\mathbf{F}, \overline{A_q}) = \max_{x^q} \mathsf{P}^{\mathbf{F}}_{\overline{A_q}|X^q}(x^q)$ 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 $\mathcal{A} = A_0, A_1, \ldots$ 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 $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}A_i} = \mathsf{p}^{\mathbf{G}}_{Y_i|X^iY^{i-1}}$ for $i \geq 1$, for all arguments for which $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}A_i}$ 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 [23], see also [16] for the proof of claim (ii) and [22] 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) [23, 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) [23, Theorem 2] If $\mathsf{p}^{\mathbf{F}}_{A_i|X^iY^{i-1}A_{i-1}} = \mathsf{p}^{\mathbf{F}}_{A_i|X^iA_{i-1}}$ *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 $f\colon \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ under a key $k \in \mathcal{K}$ we often write $f_k(\cdot)$ instead of $f(k, \cdot)$. A variable input-length keyed function $\mathsf{G} : \{0,1\}^c \times \{0,1\}^{b*} \to \{0,1\}^c$ is an:

– $(\varepsilon, t, q, \ell)$*-secure PRF*, if for any adversary A running in time t and making at most q queries, each of length at most ℓ (in b-bit blocks), a URF $\mathbf{R}\colon \{0,1\}^{b*} \to \{0,1\}^c$ and a uniformly random key $K \leftarrow \{0,1\}^c$, we have $\Delta^{\mathsf{A}}(\mathsf{G}_K, \mathbf{R}) \leq \varepsilon$.

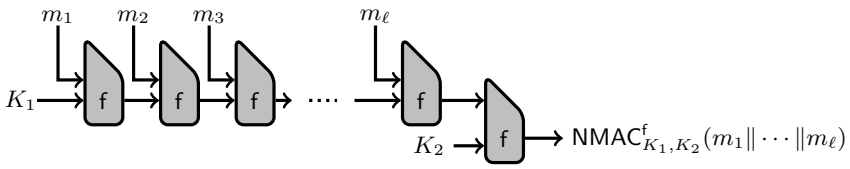

Fig. 1. The construction $\mathsf{NMAC}^{\mathsf{f}}_{K_1,K_2}$.

- $(\varepsilon, t, q, \ell)$-*NA-secure PRF*, if the above is true for all adversaries A that choose their queries non-adaptively (i.e., 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 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 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 A as an $(\varepsilon, t, q, \ell)$-PRF adversary against G if it runs in time t, asks at most q queries each consisting of at most ℓ blocks, and achieves the advantage $\Delta^{\mathsf{A}}(\mathsf{G}_K, \mathbf{R}) = \varepsilon$. We refer analogously to adversaries for the other PRF-notions defined above.

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

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

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

The construction $\mathsf{NMAC}^{\mathsf{f}} : (\{0,1\}^c)^2 \times \{0,1\}^{b*} \to \{0,1\}^c$ is derived from $\mathsf{Casc}^{\mathsf{f}}$ 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,

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

or $\mathsf{NMAC}^{\mathsf{f}}_{K_1,K_2} := \mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathsf{f}_{K_2}$, see Figure 1. 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 of NMAC above (resp. HMAC below) actually corresponds to the generalized constructions denoted as GNMAC (resp. GHMAC) in [3] where this step is also justified in detail.

$\mathsf{HMAC}^{\mathsf{f}}$ is a practice-oriented version of $\mathsf{NMAC}^{\mathsf{f}}$, 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:

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

and fpad is a fixed $(b-c)$-bit padding not affecting the security analysis. (Technically, [19] allows for arbitrary length of the key K: a key shorter than b bits is padded with zeroes before applying

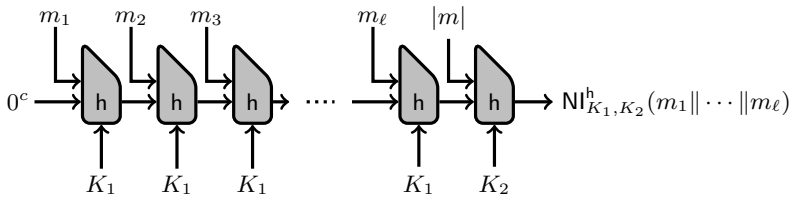

Fig. 2. The construction $\mathsf{NI}^{\mathsf{h}}_{K_1,K_2}$.

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 $\mathsf{h}\colon \{0,1\}^k \times \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$. 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 $\mathsf{NI}^{\mathsf{h}}\colon (\{0,1\}^k)^2 \times \{0,1\}^{b*} \to \{0,1\}^c$ as

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

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 $\mathsf{f}\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ and a key K we will denote by $\mathsf{LenCasc}^{\mathsf{f}}_K$ a system that given a message m outputs the pair $(\mathsf{Casc}^{\mathsf{f}}_K(m), |m|)$. This allows us to describe NI equivalently as $\mathsf{NI}^{\mathsf{h}}_{K_1,K_2} := \mathsf{LenCasc}^{\mathsf{h}_{K_1}}_{\mathbf{0}} \triangleright \mathsf{h}_{K_2}$, see also Figure 2. For a detailed discussion of the relationship of NI to NMAC, see [2]. For completeness, we also consider the modified version of NI that replaces the message length $|m|$ in the last (outer) call of the compression function by the constant bitstring 0^b, we denote this variant as NI2. Formally, we have

$$\mathsf{NI2}^{\mathsf{h}}_{K_1,K_2}(m) := \mathsf{h}_{K_2}(\mathsf{Casc}^{\mathsf{h}_{K_1}}_{\mathbf{0}}(m), 0^b)$$

or $\mathsf{NI2}^{\mathsf{h}}_{K_1,K_2} := \mathsf{ZCasc}^{\mathsf{h}_{K_1}}_{\mathbf{0}} \triangleright \mathsf{h}_{K_2}$, where $\mathsf{ZCasc}^{\mathsf{f}}_K$ a system that given a message m outputs the pair $(\mathsf{Casc}^{\mathsf{f}}_K(m), 0^b)$.

3 PRF-Security of NMAC

In this section we analyze the PRF security of $\mathsf{NMAC}^{\mathsf{f}}$ in terms of the PRF-security of the underlying function f.

3.1 Security Lower Bound

Before moving to the $\mathsf{NMAC}^{\mathsf{f}}$ construction, we start by stating a lower bound on the security of the cascade $\mathsf{Casc}^{\mathsf{f}}$ 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 given in Appendix A.

Proposition 1 ($\mathsf{Casc}^{\mathsf{f}}$ as a NA-PF-PRF). *Let* $\mathsf{f}\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ *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* $\mathsf{Casc}^{\mathsf{f}}$*,* T^{A} *is an* $(\varepsilon_{\mathsf{na}}, t, q)$*-NA-PRF adversary against* f *such that*

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

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

Theorem 1 ($\mathsf{NMAC}^{\mathsf{f}}$ as a PRF). *If* $\mathsf{f}\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ *is an* (ε, t, q)*-secure PRF and an* $(\varepsilon_{\mathsf{na}}, t, q)$*-NA-secure PRF, then* $\mathsf{NMAC}^{\mathsf{f}}$ *is an* $(\varepsilon', t', q, \ell)$*-secure PRF with*

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

The reduction is uniform. Concretely, 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* $\mathsf{NMAC}^{\mathsf{f}}$*,*

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

and their parameters satisfy equations (1).

Proof. Let A be a PRF-adversary running in time t' and asking q queries, each of length at most ℓ blocks. Let $\mathbf{r}\colon \{0,1\}^b \to \{0,1\}^c$, $\mathbf{R}\colon \{0,1\}^{b*} \to \{0,1\}^c$ 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 $\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright g$. Finally, output the decision bit of A. Clearly we have $\Delta^{\mathsf{A}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathsf{f}_{K_2}, \mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}) = \Delta^{\mathsf{T}_1^{\mathsf{A}}}(\mathsf{f}_K, \mathbf{r})$ and if we denote $\Delta^{\mathsf{T}_1^{\mathsf{A}}}(\mathsf{f}_K, \mathbf{r})$ by ε then using triangle inequality we get

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

In the experiment where A interacts with $\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}$, let C_i denote the event that during the first i queries to $\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}$, for any two distinct queries M and M' the values $\mathsf{Casc}^{\mathsf{f}}_{K_1}(M)$ and $\mathsf{Casc}^{\mathsf{f}}_{K_1}(M')$ (inputs to the final $\mathbf{r}$-call) are also distinct. As long as the monotone condition $\mathcal{C} = C_0, C_1, \ldots$ remains satisfied, the responses of $\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}$ to distinct queries are equivalent to outputs of $\mathbf{r}$ on distinct inputs, and thus independent, uniformly random values, in particular $(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}) | \mathcal{C} \equiv \mathbf{R}$. We can therefore apply Lemma 1(i) to conclude that distinguishing $\mathsf{Casc}^{\mathsf{f}} \triangleright \mathbf{r}$ from a URF $\mathbf{R}$ is at least as hard as making the condition $\mathcal{C}$ fail, i.e.,

$$\Delta^{\mathsf{A}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}, \mathbf{R}) \le \nu^{\mathsf{A}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}, \overline{C_q})\ .$$

Below we explain how to use the adversary A to construct[4] a *non-adaptive* adversary A_{na} such that

$$\nu^{\mathsf{A}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}, \overline{C_q}) = \nu^{\mathsf{A}_{\mathsf{na}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}, \overline{C_q})\ . \tag{2}$$

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

[4] One could use a lemma from the random system framework [23] 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.

during this simulated interaction. The equation (2) 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 $\mathsf{Casc}^{\mathsf{f}}_{K_1}$ and A_{na} is non-adaptive, we additionally have

$$\nu^{\mathsf{A}_{\mathsf{na}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1} \triangleright \mathbf{r}, \overline{C_q}) = \nu^{\mathsf{A}_{\mathsf{na}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1}, \overline{C_q}) \ .$$

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 $\mathsf{Casc}^{\mathsf{f}}_{K_1}$) with at least 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

$$\nu^{\mathsf{A}_{\mathsf{na}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1}, \overline{C_q}) \leq \nu^{\mathsf{A}_{\mathsf{pf}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1}, \overline{C_q})$$

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 $\mathsf{Casc}^{\mathsf{f}}_{K_1}$ outputs 1 with probability $\nu^{\mathsf{A}_{\mathsf{pf}}}(\mathsf{Casc}^{\mathsf{f}}_{K_1}, \overline{C_q})$, 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 $\Delta^{\mathsf{A}^*}(\mathsf{Casc}^{\mathsf{f}}_{K_1}, \mathbf{R})$, we have

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

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 $\mathsf{T}^{\mathsf{A}^*}$ against f such that

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

Putting $\mathsf{T}^{\mathsf{A}}_2 := \mathsf{T}^{\mathsf{A}^*}$ hence concludes the proof of Theorem 1. □

3.2 Matching Attacks

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

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

- *f is $(\varepsilon_{\mathsf{na}}, t, q)$-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*

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

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

Roughly speaking, we construct an $(\varepsilon_{\mathsf{na}}, t, q)$-NA-secure PRF f that behaves pseudo-randomly for all keys except for a small, $\varepsilon_{\mathsf{na}}/2$-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 $\mathsf{f}(k,\cdot)$ outputs some constant value $w \in \mathcal{K}$ irrespective of its input.

To attack the NA-PRF security of $\mathsf{NMAC}^{\mathsf{f}}_{K=(K_1,K_2)}$, 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 $\mathsf{Casc}^{\mathsf{f}}(K_1, M)$ is in $\mathcal{K}$, then all following intermediate values are w, and in particular we have $\mathsf{Casc}^{\mathsf{f}}(K_1, M_i) = w$ for both $i \in \{1,2\}$, and hence also $\mathsf{NMAC}^{\mathsf{f}}(K, M_1) = \mathsf{NMAC}^{\mathsf{f}}(K, M_2) = \mathsf{f}_{K_2}(w)$. This implies that it is much more likely to get a collision for a pair of messages as described above for $\mathsf{NMAC}^{\mathsf{f}}_K$ than for $\mathbf{R}$. Our adversary A simply choses $q/2$ 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 $\varepsilon_{\mathsf{na}}$. By the union bound this gives us a total probability of $\Theta(\ell q \varepsilon_{\mathsf{na}})$ for observing a collision when querying $\mathsf{NMAC}^{\mathsf{f}}_K$. On the other hand the probability of observing a colliding pair in $\mathbf{R}$ is only $O(q/2^c)$. □

We emphasize that the above attack only uses messages of one particular length and hence works equally well also if the hash function applies some length-dependent padding such as the MD-strengthening.

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

In [29] a pair $\mathsf{g}_1, \mathsf{g}_2$ of PRFs is constructed such that g_1 and g_2 are $\varepsilon_{\mathsf{na}}$-secure *non-adaptive* PRFs for some negligible $\varepsilon_{\mathsf{na}}$, and the serial composition $\mathsf{g}_1 \triangleright \mathsf{g}_2$ with independent keys can be broken by an *adaptive* attack (in a constant number of queries) with advantage almost 1.[5] From such $\mathsf{g}_1, \mathsf{g}_2$ we can get a single PRF f which is an $\varepsilon_{\mathsf{na}}$-secure NA-PRF for a negligible $\varepsilon_{\mathsf{na}}$, an ε-secure PRF for any ε of our choice, and where $\mathsf{f} \triangleright \mathsf{f}$ is not $\Theta(\varepsilon^2)$-secure, by setting $\mathsf{f} := \mathsf{g}_1$ and $\mathsf{f} := \mathsf{g}_2$ with probability $\varepsilon/2$, respectively, and some strong standard PRF with probability $1-\varepsilon$ (over the choice of the key). We now observe that $\mathsf{NMAC}^{\mathsf{f}}_K$ computed on single-block messages is simply a cascade of two f's with independent keys. Thus, when using the above ε-secure PRF f, we can break $\mathsf{NMAC}^{\mathsf{f}}_K$ 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 [25,31].

4 PRF-Security of the NI Construction

In this section we analyze the PRF-security of the constructions NI^{h} and $\mathsf{NI2}^{\mathsf{h}}$ under the assumption that the keyed compression function h is a PRF (when keyed via its k-bit input).

Recall that $d'(n)$ denotes the maximum, over all positive integers $n' \le n$, of the number of positive divisors of n'; i.e., $d'(n) := \max_{n' \in \{1,\ldots,n\}} |\{d \in \mathbb{N} : d \mid n'\}|$.

Theorem 2. *If* $\mathsf{h}\colon \{0,1\}^k \times \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ *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*

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

[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 [30].

and $\mathsf{NI2}^{\mathsf{h}}$ *is an* $(\varepsilon'', t'', q, \ell)$*-secure PRF with*

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

Proof. We first prove Theorem 2 for the case of $\mathsf{NI2}^{\mathsf{h}}$ and then derive the simpler case NI^{h} from it. The proof proceeds in four consecutive steps. First, we use the PRF-security of 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.

From a PRF to a Random Function. Let A be a PRF-adversary against $\mathsf{NI2}^{\mathsf{h}}$ 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

$$\Delta^{\mathsf{A}}(\mathsf{NI2}^{\mathsf{h}}_K, \mathbf{R}) = \Delta^{\mathsf{A}}\left(\mathsf{ZCasc}_{\mathbf{0}}^{\mathsf{h}_{K_1}} \triangleright \mathsf{h}_{K_2}, \mathbf{R}\right) \leq \varepsilon_1 + \varepsilon_2 + \Delta^{\mathsf{A}}\left(\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \mathbf{R}\right) \tag{3}$$

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 $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2$ is very similar to NMAC with an ideal compression function and keys fixed to zeroes.

Bound via Collision Probability. Let $\mathsf{CColl}(\ell)$ denote the probability that a random choice of the compression function $\mathbf{f}_1$ results in a collision in $\mathsf{Casc}_{\mathbf{0}}^{\mathbf{f}_1}$, maximized over the choice of the two distinct inputs to the cascade m_1, m_2 consisting of at most ℓ blocks each. (Note that this implies a collision also for $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1}$.) Formally, for uniformly random $\mathbf{f}_1 \leftarrow \mathcal{F}(c+b, c)$ we define

$$\mathsf{CColl}(\ell) := \max_{\substack{m_1 \neq m_2 \\ |m_1|, |m_2| \leq \ell b}} \mathsf{P}^{\mathbf{f}_1}\left[\mathsf{Casc}_{\mathbf{0}}^{\mathbf{f}_1}(m_1) = \mathsf{Casc}_{\mathbf{0}}^{\mathbf{f}_1}(m_2)\right] \; . \tag{4}$$

In the experiment where A interacts with $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2$, let E_i denote the event that during the first i queries to $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2$, for any two distinct queries M and M' the values $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1}(M)$ and $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1}(M')$ (inputs to the final $\mathbf{f}_2$-call) were also distinct. As long as the monotone condition $\mathcal{E} = E_0, E_1, \ldots$ remains satisfied, the responses of $\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2$ to distinct queries are clearly independent, uniformly random values thanks to $\mathbf{f}_2$. Hence, we have $(\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2)|\mathcal{E} \equiv \mathbf{R}$ and $\mathsf{p}^{\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2}_{E_i | X^i Y^{i-1} E_{i-1}} = \mathsf{p}^{\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2}_{E_i | X^i E_{i-1}}$ and can therefore consecutively apply Lemma 1(i), Lemma 1(ii), and finally the union bound to get

$$\Delta^{\mathsf{A}}(\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \mathbf{R}) \leq \nu(\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \overline{E_q}) \leq \mu(\mathsf{ZCasc}_{\mathbf{0}}^{\mathbf{f}_1} \triangleright \mathbf{f}_2, \overline{E_q}) \leq q^2 \cdot \mathsf{CColl}(\ell) \; . \tag{5}$$

Graph-Based Representation of Casc. The probability $\mathsf{CColl}(\ell)$ 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 $\mathsf{CColl}(\ell)$ is the central part of our proof. To this end, we use an approach inspired by [7] and represent the computation of $\mathsf{Casc}_{\mathbf{0}}^{\mathbf{f}_1}$ on various inputs by directed graphs.

Let m_1 and m_2 be two distinct messages that can be parsed into b-bit blocks as $m_i = m_i^1 \| \cdots \| m_i^{\ell_i}$ for some $\ell_1, \ell_2 \leq \ell$, and let $\Lambda := \ell_1 + \ell_2$. For convenience, we use the notation $m^{(i)}$ as a reference to the block m_1^i if $i \leq \ell_1$, otherwise it denotes the block $m_2^{i-\ell_1}$. 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

$$s_i := \begin{cases} \mathbf{0} & \text{for } i = 0 \\ f(s_{i-1}, m_1^i) & \text{for } 1 \le i \le \ell_1 \\ f(\mathbf{0}, m_2^1) & \text{for } i = \ell_1 + 1 \\ f(s_{i-1}, m_2^{i-\ell_1}) & \text{for } \ell_1 + 2 \le i \le \Lambda \end{cases} \tag{6}$$

and consider the mappings $[\cdot]_G$ and $[\cdot]'_G$ defined on $\{0, \ldots, \Lambda\}$ 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 $i \neq \ell_1$, while $[\ell_1]'_G := 0$. Now we let

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

– $\mathcal{L}\colon \mathcal{V}^2 \to \mathsf{Pow}(\{0,1\}^b)$ 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 $\mathsf{Casc}_{\mathbf{0}}^{\mathbf{f}_1}$ 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 $\mathcal{M}' := (m_1^1 \| \cdots \| m_1^i, \lambda)$ if $i \le \ell_1$ and $\mathcal{M}' := (m_1, m_2^1 \| \cdots \| m_2^{i-\ell_1})$ otherwise. Building on this notion, we call $\mathsf{fColl}(G)$ the *set of f-collisions* that occurred in G:

$$\mathsf{fColl}(G) := \left\{ (i, [i]_G) : [i]_G < i \wedge m^{(i)} \notin \mathcal{L}_{i-1}([i-1]'_G, [i]_G) \right\} . \tag{7}$$

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 $i = 1, \ldots, \Lambda$. The pair $(i, [i]_G)$ belongs to $\mathsf{fColl}(G)$ (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 and its proof is inspired by Lemma 8 from [7].

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*

$$\mathsf{P}^F\left[G_F^{\mathcal{M}} = g\right] \le 2^{-c \cdot |\mathsf{fColl}(g)|} .$$

Proof (of Lemma 2). Let $\mathcal{M} = \{m_1, m_2\}$, $\Lambda = \ell_1 + \ell_2$ and let $m^{(i)}$ denote the i-th block of $m_1 \| m_2$ as before. First, we introduce the notion of consistency. Assume we sample $F \leftarrow \mathcal{F}(c+b, c)$ and the values $S_1, \ldots, S_\Lambda$ belonging to $G = G_F^{\mathcal{M}}$ are revealed to us stepwise. (Recall that S_i is the

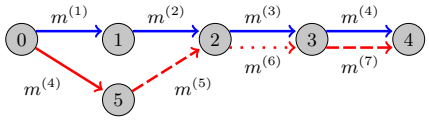

$\mathcal{M} = \{m_1, m_2\}$, $|m_1| = |m_2| = 4b$
$m^{(4)} \neq m^{(1)}, m^{(6)} = m^{(3)}, m^{(7)} \neq m^{(4)}$
- m_1 — fresh
- m_2 - - - collision
. . . . . determined

Fig. 3. Illustration of the three cases from Lemma 2.

random variable representing the chaining variable of the cascade defined in (6) and determined by the choice of F. In turn, the values $S_1, \ldots, S_\Lambda$ completely determine the shape of the structure graph G.) We say that G is *consistent* with the given graph g after step $i \leq \Lambda$, denoted Cons_i, if the structure graphs G_i and g_i are equal as triples $(\mathcal{V}, \mathcal{E}, \mathcal{L})$ (as before, G_i denotes the part of graph G obtained after the first i blocks are processed, and g_i is defined analogously from g).

Let us assume that Cons_i is true for some i and then bound the probability $\mathsf{P}[\mathsf{Cons}_{i+1}|\mathsf{Cons}_i]$. To this end, we inspect the $(i+1)$-th step in g where there are the following 3 possibilities how the next edge corresponding to $m^{(i+1)}$ might look (see also Fig. 3):

Fresh: It arrives at a new vertex not present in g_i (i.e., $[i+1]_g = i+1$).
Determined: It follows an already existing edge (i.e., $[i+1]_g \leq i$ and $m^{(i+1)}$ is already in the label set of the edge $([i]_g, [i+1]_g)$ in g_i).
Collision: It causes an f-collision (i.e., $[i+1]_g \leq i$ and $m^{(i+1)}$ is not in the label set of the edge $([i]_g, [i+1]_g)$ in g_i). In this case, G_{i+1} will stay consistent if and only if its $(i+1)$-th edge lands on precisely the same vertex as in g_{i+1}, in other words, if $S_{i+1} = s_{i+1}$. The probability of this event (conditioned on Cons_i) is 2^{-c}, as S_{i+1} is uniformly random over $\{0,1\}^n$ and not determined in the first i steps.

Since the third case occurs exactly $|\mathsf{fColl}(g)|$ times, if we trivially upper-bound the probabilities $\mathsf{P}[\mathsf{Cons}_{i+1}|\mathsf{Cons}_i]$ in the other two cases by 1, we obtain the final bound $\mathsf{P}[G = g] = \mathsf{P}[\mathsf{Cons}_\Lambda] \leq 2^{-c\cdot|\mathsf{fColl}(g)|}$ as desired. □

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*

$$\mathsf{P}^F\left[\left|\mathsf{fColl}\left(G_F^{\mathcal{M}}\right)\right| \geq 2\right] \leq \frac{4\Lambda^4}{2^{2c}}\,.$$

Proof (of Lemma 3). Denote by $\mathcal{G}^r(\mathcal{M}) := \{G \in \mathcal{G}(\mathcal{M}) : |\mathsf{fColl}(G)| = r\}$ the set of all structure graphs for $\mathcal{M}$ containing exactly r f-collisions. Then (using Lemma 2 in the last step) we have

$$\mathsf{P}\left[\left|\mathsf{fColl}\left(G_F^{\mathcal{M}}\right)\right| \geq 2\right] = \sum_{r=2}^{\infty} \mathsf{P}\left[\left|\mathsf{fColl}\left(G_F^{\mathcal{M}}\right)\right| = r\right] = \sum_{r=2}^{\infty} \sum_{g\in\mathcal{G}^r(\mathcal{M})} \mathsf{P}\left[G_F^{\mathcal{M}} = g\right] \leq \sum_{r=2}^{\infty} \frac{|\mathcal{G}^r(\mathcal{M})|}{(2^c)^r}\,.$$

Since one can verify that any $G \in \mathcal{G}(\mathcal{M})$ is completely determined by the set of its f-collisions $\mathsf{fColl}(G) \subseteq \{(i,j) : 0 \leq j < i \leq \Lambda\}$ and the latter set has $\Lambda(\Lambda+1)/2$ elements, we have $|\mathcal{G}^r(\mathcal{M})| \leq (\Lambda(\Lambda+1)/2)^r$ and hence

$$\mathsf{P}\left[\left|\mathsf{fColl}\left(G_F^{\mathcal{M}}\right)\right| \geq 2\right] \leq \sum_{r=2}^{\infty} \left(\frac{\Lambda(\Lambda+1)}{2\cdot 2^c}\right)^r \leq \frac{4\Lambda^4}{2^{2c}}\,.$$

In the last step we used that $1 \leq \Lambda \leq 2^{c/2}$ and $c \geq 2$ which can be safely assumed, since otherwise the statement of the lemma is trivially true (as 1 upper-bounds any probability). □

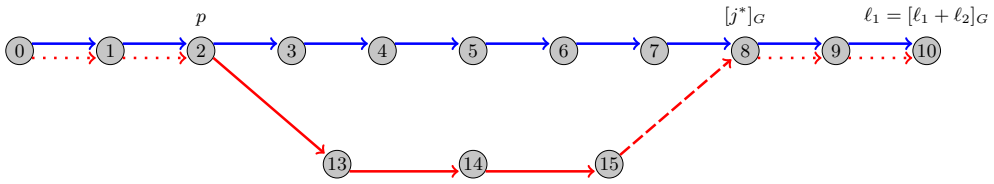

Fig. 4. A sample graph from the set $\mathcal{H}_1$ in the proof of Lemma 4, with $p = 2$ and $j^* = 16$.

From Collision Probability to Counting Graphs. We can now proceed to upper-bounding the value $\mathsf{CColl}(\ell)$. Let $\mathcal{M} := (m_1, m_2)$ be the two distinct messages of length at most ℓ blocks that maximize the probability $\mathsf{CColl}(\ell) := \max_{m_1 \neq m_2} \mathsf{P}^F\left[\mathsf{Casc}_{\mathbf{0}}^F(m_1) = \mathsf{Casc}_{\mathbf{0}}^F(m_2)\right]$. 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 $V_2^i := [\ell_1 + i]'_G$. We also refer to the path $V_j^0, \ldots, V_j^{\ell_j}$ as the *m_j-path*. Using this notation, we have $\mathsf{CColl}(\ell) = \mathsf{P}[V_1^{\ell_1} = V_2^{\ell_2}]$. Since $m_1 \neq m_2$, $V_1^{\ell_1} = V_2^{\ell_2}$ cannot occur without any f-collision, hence we can split $\mathsf{CColl}(\ell)$ into

$$\mathsf{P}\left[V_1^{\ell_1} = V_2^{\ell_2} \wedge |\mathsf{fColl}(G_F^{\mathcal{M}})| = 1\right] + \mathsf{P}\left[V_1^{\ell_1} = V_2^{\ell_2} \wedge |\mathsf{fColl}(G_F^{\mathcal{M}})| \geq 2\right] . \tag{8}$$

The latter probability can be readily upper-bounded by $4\Lambda^4/2^{2c}$ using Lemma 3. As for the former, let us denote by $\mathcal{H}(\mathcal{M})$ the set

$$\mathcal{H}(\mathcal{M}) := \left\{G \in \mathcal{G}^1(\mathcal{M}) : V_1^{\ell_1} = V_2^{\ell_2}\right\}$$

of structure graphs for $\mathcal{M}$ that contain exactly one f-collision and where the vertices $V_1^{\ell_1}$ and $V_2^{\ell_2}$ coincide. The first term in (8) 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. Recall that $d'(n)$ denotes the maximum, over all positive integers $n' \leq n$, of the number of positive divisors of n'; i.e., $d'(n) := \max_{n' \in \{1,\ldots,n\}} |\{d \in \mathbb{N} : d \mid n'\}|$.

Lemma 4. *For two distinct messages $\mathcal{M} = \{m_1, m_2\}$ each of length at most ℓ blocks we have $|\mathcal{H}(\mathcal{M})| \leq \ell d'(\ell)$. If the messages in $\mathcal{M}$ are of the same length then we have $|\mathcal{H}(\mathcal{M})| \leq \ell$.*

Proof (of Lemma 4). Let us first consider the general case where we allow the messages m_1 and m_2 to have different lengths, let us denote them by ℓ_1 and ℓ_2 as before. Without loss of generality let us assume that $\ell_1 \geq \ell_2$. We split the set $\mathcal{H}(\mathcal{M})$ into two partitions: Let $\mathcal{H}_1$ contain all the structure graphs from $\mathcal{H}(\mathcal{M})$ such that the m_1-path does not contain a loop, and let $\mathcal{H}_2$ contain all the rest. Formally, $\mathcal{H}_1 := \{G \in \mathcal{H}(\mathcal{M}); \forall i \in \{1, \ldots, \ell_1\} : [i]_G = i\}$ and $\mathcal{H}_2 := \mathcal{H}(\mathcal{M}) \setminus \mathcal{H}_1$. We now upper-bound the size of both partitions in two separate claims, which together conclude the proof of the first part of Lemma 4.

Claim 1: $|\mathcal{H}_1| \leq \ell$.

Towards bounding $|\mathcal{H}_1|$, note that if m_2 is a prefix of m_1 then clearly $|\mathcal{H}_1| = 0$, therefore we assume that this is not the case. Let $m_1^1 \| \cdots \| m_1^p$ be the blocks forming the longest common prefix of m_1 and m_2; i.e., let $p \in \mathbb{N}$ be the smallest index such that $m_1^{p+1} \neq m_2^{p+1}$ (for illustration see Fig. 4). Since f is a function, we clearly have $V_1^i = V_2^i$ for all $i \leq p$. Let us now consider $j^* := \min\{j > \ell_1 + p : [j]_G \leq \ell_1\}$. Such a j^* is well-defined, since at least the value $\ell_1 + \ell_2$ belongs to the considered set (we have $\ell_1 + \ell_2 > \ell_1 + p$ and $[\ell_1 + \ell_2]_G = \ell_1$).

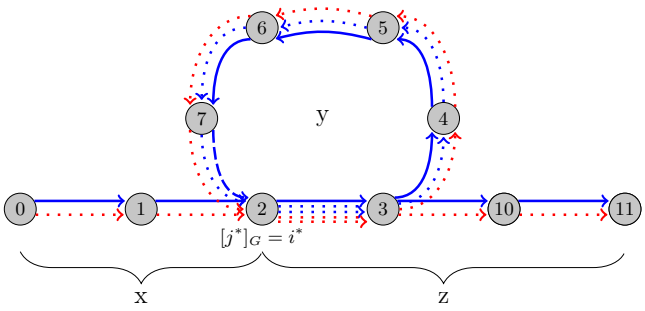

Fig. 5. A sample graph from the set $\mathcal{H}_2$ in the proof of Lemma 4, with $i^* = 2$ and $j^* = 8$.

We now prove that the j^*-th edge $([j^*-1]'_G, [j^*]_G)$ in G must create an f-collision, i.e., that $(j^*, [j^*]_G) \in \mathsf{fColl}(G)$. We have $[j^*]_G \in \mathcal{V}_{j^*-1}$ by definition of j^* and to also see that $m^{(j^*)} \notin \mathcal{L}_{j^*-1}([j^*-1]'_G, [j^*]_G)$ we consider two cases:

1. If $[j^*]_G \geq 1$ and $[j^*]_G - 1 = [j^*-1]'_G$ (the vertices directly preceding the vertex $V_1^{[j^*]_G}$ on m_1-path and m_2-path coincide), then we must have $j^* = p+1$, otherwise this would contradict the minimality of j^*. However, this implies that $m^{([j^*]_G)} \neq m^{(j^*)}$ (as otherwise the common prefix would be longer than p blocks) and hence $m^{(j^*)} \notin \mathcal{L}_{j^*-1}([j^*-1]'_G, [j^*]_G) = \{m^{([j^*]_G)}\}$.
2. On the other hand, if $[j^*]_G - 1 \neq [j^*-1]'_G$, then we claim that there was no edge $([j^*-1]'_G, [j^*]_G)$ in G_{j^*-1} and hence $m^{(j^*)} \notin \mathcal{L}_{j^*-1}([j^*-1]'_G, [j^*]_G) = \emptyset$. Indeed, the only edge leading into the vertex $[j^*]_G$ in G_{j^*-1} can be from $[j^*]_G - 1$, as anything else would contradict either the absence of cycles within the m_1-path, or the minimality of j^*.

Given the j^*-th edge causes an f-collision and $|\mathsf{fColl}(G)| = 1$, no f-collision in G occurs beyond the j^*-th edge. However, we have $[\ell_1]_G = [\ell_1 + \ell_2]_G$ and to achieve this without any additional collision, clearly, we need that $m^{([j^*]_G+1)} \| \cdots \| m^{(\ell_1)} = m^{(j^*+1)} \| \cdots \| m^{(\ell_1+\ell_2)}$, i.e., the suffixes of m_1 and m_2 after the collision are the same. This, however, implies that the value j^* completely determines the structure graph within $\mathcal{H}_1$ and hence we arrive at $|\mathcal{H}_1| \leq \ell$.

Claim 2: $|\mathcal{H}_2| \leq \ell \cdot d'(\ell)$.
For this part, let $j^* := \min\{j : [j]_G < j\}$ and $i^* := [j^*]_G$, where such a $j^* \leq \ell_1$ exists by definition of $\mathcal{H}_2$. Moreover, it creates an f-collision (i.e., $(j^*, i^*) \in \mathsf{fColl}(G)$) by an argument similar to the one from Claim 1. We now split m_1 into $x := m_1^1 \| \cdots \| m_1^{i^*}$, $y := m_1^{i^*+1} \| \cdots \| m_1^{j^*}$ and some z that is chosen to be the shortest string possible such that $m_1 = x \| y^k \| z$ holds for some $k \geq 1$ (note that such z always exists and is unique, possibly empty). This situation is illustrated in Fig. 5.

We claim that in any $G \in \mathcal{H}_2$ the m_2-path is a subgraph of the m_1-path (ignoring the labels for now). Indeed, if the m_2-path contained any edges not contained in the m_1-path, then (since $V_1^{\ell_1} = V_2^{\ell_2}$) the last such "outlying" edge would create an f-collision. To see this, observe that since this is the last edge not in the path of m_1 its end vertex will be contained in the path of both messages, which causes an f-collision when this edge is added (see (7)). However, the m_1-path already created one f-collision and hence creating another one would violate the definition of $\mathcal{H}(\mathcal{M})$.

Moreover, for the same reason the m_2-path cannot introduce new labels to the edges in m_1-path, as this would cause another f-collision. This implies that m_2 has to be of the form $m_2 = x \| y^{k'} \| z$ for some $k' < k$. To achieve this, the number of blocks in y (i.e., $j^* - i^*$) must divide $\ell_1 - \ell_2$.

For any fixed $\mathcal{M}$, a structure graph in $\mathcal{H}_2$ is fully determined by the choice of $j^* \in \{1, \ldots, \ell_1\}$ and $i^* \in \{0, \ldots, j^* - 1\}$, such that $(j^* - i^*) \mid \ell_1 - \ell_2$. There are at most ℓ ways to choose such a

j^* and at most $d'(\ell)$ ways to choose a consistent i^*. Consequently, we obtain $|\mathcal{H}_2| \leq \ell \cdot d'(\ell)$, which concludes the proof for the case of distinct-length messages.

For the second part of the claim it now suffices to observe that if $|m_1| = |m_2|$ then $|\mathcal{H}_2| = 0$. This is because in $\mathcal{H}_2$ the m_1-path already contains an f-collision, and since only one such f-collision is allowed to occur, the only way to achieve $V_1^{\ell_1} = V_2^{\ell_2}$ would hence be if $m_1 = m_2$. This however contradicts the assumption that the messages are distinct. □

In Appendix C we also show that Lemma 4 is tight, and discuss the implications for the tightness of Theorem 2.

Finally, combining the equations (3), (5), (8), and the bounds obtained in Lemma 3 and Lemma 4, we get

$$\Delta^{\mathsf{A}}(\mathsf{NI2}^{\mathsf{h}}_K, \mathbf{R}) \leq \varepsilon_1 + \varepsilon_2 + q^2 \cdot \left(\frac{\ell \cdot d'(\ell)}{2^c} + \frac{4\Lambda^4}{2^{2c}} \right) \leq \varepsilon_1 + \varepsilon_2 + \frac{q^2}{2^c} \cdot \left(\ell \cdot d'(\ell) + \frac{64\ell^4}{2^c} \right)$$

and conclude the proof of Theorem 2 for $\mathsf{NI2}^{\mathsf{h}}$.

The case of NI is handled in the same way as $\mathsf{NI2}$, with the only difference being that it contains $\mathsf{LenCasc}$ instead of ZCasc. Hence, to imply a collision for $\mathsf{LenCasc}$, we require the messages m_1 and m_2 in the definition of $\mathsf{CColl}(\ell)$ to be of the same length. This leads to the use of the second part of Lemma 4 that assumes equal-length messages, arriving at the claimed bound. □

Acknowledgements. We thank the anonymous CRYPTO 2014 reviewers for useful comments and suggestions.

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. Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *Advances in Cryptology – CRYPTO'88*, volume 403 of *Lecture Notes in Computer Science*, pages 8–26. Springer, August 1990.
16. Dimitar Jetchev, Onur Özen, and Martijn Stam. Understanding adaptivity: Random systems revisited. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 313–330. Springer, 2012.
17. 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.
18. Neal Koblitz and Alfred Menezes. Another look at HMAC. Cryptology ePrint Archive, Report 2012/074, 2012.
19. Hugo Krawczyk, Mihir Bellare, and Ran Canetti. HMAC: Keyed-Hashing for Message Authentication. IETF Internet Request for Comments 2104, February 1997.
20. Gatan Leurent, Thomas Peyrin, and Lei Wang. New Generic Attacks against Hash-Based MACs. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology - ASIACRYPT 2013*, volume 8270 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2013.
21. Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal on Computing*, 17(2):373–386, 1988.
22. 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.
23. 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.
24. 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.
25. 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.
26. 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, *Advances in Information and Computer Security*, volume 8231 of *Lecture Notes in Computer Science*, pages 83–98. Springer, 2013.
27. Thomas Peyrin, Yu Sasaki, and Lei Wang. Generic Related-Key Attacks for HMAC. In Xiaoyun Wang and Kazue Sako, editors, *Advances in Cryptology - ASIACRYPT 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 580–597. Springer, 2012.
28. Thomas Peyrin and Lei Wang. Generic Universal Forgery Attack on Iterative Hash-Based MACs. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 147–164. Springer, 2014.
29. 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.
30. 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.
31. 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.
32. 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.

33. 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.

A Non-Adaptive Security of the Cascade

Here we prove Proposition 1 that states the PRF-security of the construction $\mathsf{Casc}^{\mathsf{f}}$ against non-adaptive prefix-free adversaries, assuming that f itself is a non-adaptively secure PRF. Our argument follows the proof for the adaptive case in [5] with minor modifications and we include it here for completeness.

Given a compression function $\mathsf{f}\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ and a tuple of independent random keys $\overline{K} = (K_1, \ldots, K_q) \in (\{0,1\}^c)^q$, let $\mathsf{qf}_{\overline{K}} = (\mathsf{f}_{K_1}, \ldots, \mathsf{f}_{K_q})$ denote the q-tuple of oracles providing access to q copies of f, each one being assigned a different key from $\overline{K}$. Moreover, let $\mathbf{qr} = (\mathbf{r}_1, \ldots, \mathbf{r}_q)$ denote the q-tuple of independent, uniformly random functions $\mathbf{r}_i\colon \{0,1\}^b \to \{0,1\}^c$. Following [5], we say that f is (ε, t, q)-NA-PRF^q-secure, if for any non-adaptive adversary A running in time t and asking at most q queries, we have $\Delta^{\mathsf{A}}(\mathsf{qf}_{\overline{K}}, \mathbf{qr}) \leq \varepsilon$.

Proposition 1 (restated). *Let $\mathsf{f}\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ 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 $\mathsf{Casc}^{\mathsf{f}}$, T^{A} is an $(\varepsilon_{\mathsf{na}}, t, q)$-NA-PRF adversary against f such that*

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

Proof. The proof consists of two consecutive reductions. First, out of an assumed attacker against the NA-PF-PRF security of $\mathsf{Casc}^{\mathsf{f}}$ we construct an attacker against the NA-PRF^q security of f. Second, we use the latter to construct an attacker against the NA-PRF-security of f. In each of these two steps the success probabilities of the two attackers are related by a hybrid argument. We describe and analyze each of these two steps in a separate lemma below.

Lemma 5. *There exists an explicit reduction T_1 (described in the proof) such that for any non-adaptive adversary A_1 against the NA-PF-PRF security of $\mathsf{Casc}^{\mathsf{f}}$, running in time t' and asking q prefix-free queries of length at most ℓ blocks each, $\mathsf{A}_2 := \mathsf{T}_1^{\mathsf{A}_1}$ is a non-adaptive adversary against the NA-PRF^q-security of f running in time $t' + O(\ell q)$ and asking at most q queries, such that $\Delta^{\mathsf{A}_1}(\mathsf{Casc}^{\mathsf{f}}_K, \mathbf{R}) \leq \ell \cdot \Delta^{\mathsf{A}_2}(\mathsf{qf}_{\overline{K}}, \mathbf{qr})$.*

Proof (of Lemma 5). We start by describing a sequence of adversaries $\mathsf{A}_2^{(i)}$ for $i \in \{1, \ldots, \ell\}$. Given access to oracles $(\mathsf{g}_1, \ldots, \mathsf{g}_q)$ which are either $\mathsf{qf}_{\overline{K}} = (\mathsf{f}_{K_1}, \ldots, \mathsf{f}_{K_q})$ (for independent random keys $K_1, \ldots, K_q$), or q independent random functions $\mathbf{qr} = (\mathbf{r}_1, \ldots, \mathbf{r}_q)$, $\mathsf{A}_2^{(i)}$ works as follows:

1. It runs A_1 to obtain its q non-adaptive prefix-free queries $x_1, \ldots, x_q$, each of length at most ℓ blocks (without loss of generality we assume that $x_1, \ldots, x_q \in \{0,1\}^{b*}$ are distinct). Each query x_j is parsed into blocks as $x_j = x_j^1 \| \cdots \| x_j^{\ell_j}$, where each $x_j^z \in \{0,1\}^b$.
2. The response r_j to each query x_j is determined: If $\ell_j < i$, then r_j is chosen independently and uniformly at random. Otherwise, an index $c_j \in \{1, \ldots, q\}$ is determined consecutively for all queries of length at least i in an arbitrary way, given that two queries x_j and $x_{j'}$ share the same index (i.e., $c_j = c_{j'}$) if and only if their first $i-1$ blocks are identical (i.e., $x_j^1 \| \cdots \| x_j^{i-1} = x_{j'}^1 \| \cdots \| x_{j'}^{i-1}$). The response r_j is then computed as

$$r_j \leftarrow \mathsf{Casc}^{\mathsf{f}}_{\mathsf{g}_{c_j}(x_j^i)} \left(x_j^{i+1} \| \cdots \| x_j^{\ell_j} \right)\ .$$

All g-values required for this computation are obtained by querying the g-oracles; note that this can be done non-adaptively. The tuple of responses $(r_1, \ldots, r_q)$ is given to A_1.

3. $\mathsf{A}_2^{(i)}$ outputs the same bit that A_1 does.

A straightforward analysis using the definition of $\mathsf{A}_2^{(i)}$ allows one to establish the following three facts:

(i) $\mathsf{A}_1(\mathsf{Casc}_K^{\mathsf{f}}) = \mathsf{A}_2^{(1)}(\mathsf{qf}_{\overline{K}})$,
(ii) $\mathsf{A}_1(\mathbf{R}) = \mathsf{A}_2^{(\ell)}(\mathbf{qr})$,
(iii) $\mathsf{A}_2^{(i+1)}(\mathsf{qf}_{\overline{K}}) = \mathsf{A}_2^{(i)}(\mathbf{qr})$ for all $i \in \{1, \ldots, \ell\}$,

where the equalities represent equal distributions of the output bits. Combining these facts, we get

$$\begin{aligned}\Delta^{\mathsf{A}_1}(\mathsf{Casc}_K^{\mathsf{f}}, \mathbf{R}) &= \left|\mathsf{P}[\mathsf{A}_1(\mathsf{Casc}_K^{\mathsf{f}}) = 1] - \mathsf{P}[\mathsf{A}_1(\mathbf{R}) = 1]\right| \overset{(i),(ii)}{=} \left|\mathsf{P}[\mathsf{A}_2^{(1)}(\mathsf{qf}_{\overline{K}}) = 1] - \mathsf{P}[\mathsf{A}_2^{(\ell)}(\mathbf{qr}) = 1]\right| \\ &\overset{(iii)}{\leq} \sum_{i=1}^{\ell} \left|\mathsf{P}[\mathsf{A}_2^{(i)}(\mathsf{qf}_{\overline{K}}) = 1] - \mathsf{P}[\mathsf{A}_2^{(i)}(\mathbf{qr}) = 1]\right| = \sum_{i=1}^{\ell} \Delta^{\mathsf{A}_2^{(i)}}(\mathsf{qf}_{\overline{K}}, \mathbf{qr}) \,. \end{aligned} \tag{9}$$

Now we define A_2 to initially choose an index $i \in \{1, \ldots, \ell\}$ uniformly at random and then act as $\mathsf{A}_2^{(i)}$. This implies

$$\Delta^{\mathsf{A}_2}(\mathsf{qf}_{\overline{K}}, \mathbf{qr}) = \frac{1}{\ell} \cdot \sum_{i=1}^{\ell} \Delta^{\mathsf{A}_2^{(i)}}(\mathsf{qf}_{\overline{K}}, \mathbf{qr})$$

and hence concludes the proof of Lemma 5. □

Lemma 6. *There exists an explicit reduction T_2 (described in the proof) such that for any non-adaptive adversary A_2 against the NA-PRF$^{\mathsf{q}}$-security of f, running in time $t' + O(\ell q)$ and asking at most q queries, $\mathsf{A}_3 := \mathsf{T}_2^{\mathsf{A}_2}$ is a non-adaptive adversary against the NA-PRF-security of f running in time $t' + O(\ell q)$ and asking at most q queries, such that $\Delta^{\mathsf{A}_2}(\mathsf{qf}_{\overline{K}}, \mathbf{qr}) \leq q \cdot \Delta^{\mathsf{A}_3}(\mathsf{f}_K, \mathbf{r})$.*

Proof (of Lemma 6). Let us again describe a sequence of adversaries $\mathsf{A}_3^{(i)}$ for $i \in \{1, \ldots, q\}$. Given access to an oracle g, which is either f_K (for an independent random key K), or an independent random function $\mathbf{r}$, $\mathsf{A}_3^{(i)}$ works as follows:

1. It runs A_2 to obtain its q non-adaptive queries $(o_1, x_1), \ldots, (o_q, x_q)$, each consisting of a pair (o, x) representing a query x to A_2's o-th oracle.
2. $\mathsf{A}_3^{(i)}$ chooses $i - 1$ independent random keys $K_1, \ldots, K_{i-1} \in \{0,1\}^c$. Then, it determines the response r_j to each query (o_j, x_j) as

$$r_j \leftarrow \begin{cases} \mathsf{f}_{K_{o_j}}(x_j) & \text{if } o_j < i \\ \mathsf{g}(x_j) & \text{if } o_j = i \\ \mathbf{r}_{o_j}(x_j) & \text{if } o_j > i, \end{cases}$$

   where $\mathbf{r}_{i+1}, \ldots, \mathbf{r}_q$ are independent uniformly random functions, sampled internally by $\mathsf{A}_3^{(i)}$ (using lazy sampling to maintain efficiency). All g-values required for this computation are obtained by querying the g-oracle and once again this can be done non-adaptively. The tuple of responses $(r_1, \ldots, r_q)$ is given to A_2.
3. $\mathsf{A}_3^{(i)}$ outputs the same bit that A_2 does.

This time it is easy to observe that we have

(iv) $\mathsf{A}_2(\mathsf{qf}_{\overline{K}}) = \mathsf{A}_3^{(q)}(\mathsf{f}_K)$
(v) $\mathsf{A}_2(\mathbf{qr}) = \mathsf{A}_3^{(1)}(\mathbf{r})$
(vi) $\mathsf{A}_3^{(i)}(\mathsf{f}_K) = \mathsf{A}_3^{(i+1)}(\mathbf{r})$ for all $i \in \{1, \ldots, q\}$

and hence, similarly as in (9), we get

$$\Delta^{\mathsf{A}_2}(\mathsf{qf}_{\overline{K}}, \mathbf{qr}) \overset{(iv),(v)}{=} \left|\mathsf{P}[\mathsf{A}_3^{(q)}(\mathsf{f}_K) = 1] - \mathsf{P}[\mathsf{A}_3^{(1)}(\mathbf{r}) = 1]\right| \overset{(vi)}{\leq} \sum_{i=1}^{q} \Delta^{\mathsf{A}_3^{(i)}}(\mathsf{f}_K, \mathbf{r}) \,. \tag{10}$$

Again, letting A_3 be an adversary that chooses a random index $i \in \{1, \ldots, q\}$ and then simulates $\mathsf{A}_3^{(i)}$ gives us

$$\Delta^{\mathsf{A}_3}(\mathsf{f}_K, \mathbf{r}) = \frac{1}{q} \cdot \sum_{i=1}^{q} \Delta^{\mathsf{A}_3^{(i)}}(\mathsf{f}_K, \mathbf{r}) \,,$$

thus proving Lemma 6. □

The proof of Proposition 1 is now concluded by combining the two reductions described above. For any $(\varepsilon', t', q, \ell)$-NA-PF-PRF adversary A against $\mathsf{Casc}^{\mathsf{f}}$, we let $\mathsf{T}^{\mathsf{A}} := \mathsf{T}_2^{\mathsf{T}_1^{\mathsf{A}}}$ and observe that $\Delta^{\mathsf{A}}(\mathsf{Casc}^{\mathsf{f}}_K, \mathbf{R}) \leq \ell q \cdot \Delta^{\mathsf{T}^{\mathsf{A}}}(\mathsf{f}_K, \mathbf{r})$ while T^{A} runs in time $t' + \tilde{O}(\ell q)$ and asks at most q queries as desired. □

B Proof of Proposition 2

In this appendix we fill in the details omitted in the sketch of the proof of Proposition 2 in Section 3.2.

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

- *f is $(\varepsilon_{\mathsf{na}}, t, q)$-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*

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

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

Proof. We start by showing how to construct the $(\varepsilon_{\mathsf{na}}, t, q)$-NA-secure PRF f. To simplify our technical arguments later, we design f in such a way that besides having weak keys as sketched in Section 3, it also satisfies the additional property that for any key $k \in \{0,1\}^c$ and a uniformly distributed input $U \in \{0,1\}^b$, the value $\mathsf{f}(k, U)$ is also uniformly distributed. Having this goal in mind, we construct f starting from a pseudo-random permutation (which exists by our assumption and the result [21]). Consider any $(\varepsilon_{\mathsf{na}}/4, t, q)$-NA-secure PRP $\pi\colon \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^b$ and a set of "weak keys" $\mathcal{K} \subseteq \{0,1\}^c$ of size $2^c(\varepsilon_{\mathsf{na}}/2)$, defined as $\mathcal{K} := 0^{1-\log \varepsilon_{\mathsf{na}}} \| \{0,1\}^{c+\log \varepsilon_{\mathsf{na}} - 1}$ (the set of keys where the first $1 - \log \varepsilon_{\mathsf{na}}$ bits are 0). Let $[\cdot]_c$ represent the truncation of a longer bitstring to its first c bits. We fix a value $w \in \mathcal{K}$ (say $w = 0^c$) and define f as

$$\mathsf{f}(k, x) := \begin{cases} w & \text{for } k \in \mathcal{K}, \\ [\pi(k, x)]_c & \text{for } k \notin \mathcal{K}. \end{cases}$$

Hence, f behaves as a truncated version of π except when a weak key from $\mathcal{K}$ is used, in this case $\mathsf{f}(k,\cdot)$ always outputs w. By the well-known PRF/PRP switching lemma [15] we obtain that π is also an $(\varepsilon_{\mathsf{na}}/4 + q^2/2^b, t, q)$-NA-secure PRF and by assumption $\varepsilon_{\mathsf{na}}/4 + q^2/2^b \leq \varepsilon_{\mathsf{na}}/2$. It is easy to see that this implies that also $[\pi(\cdot)]_c$ is an $(\varepsilon_{\mathsf{na}}/2, t, q)$-NA-secure PRF. By redefining $[\pi(\cdot)]_c$ on an $\varepsilon_{\mathsf{na}}/2$-fraction of the keys at most an $\varepsilon_{\mathsf{na}}/2$ term in the PRF-distinguishing advantage is lost, hence the function f is an $(\varepsilon_{\mathsf{na}}, t, q)$-NA-secure PRF.

Now consider two queries M_1, M_2 to $\mathsf{NMAC}^{\mathsf{f}}(K = (K_1, K_2), \cdot)$ which are determined by first sampling an $(\ell-1)$-block message $M = m_1 \| \cdots \| m_{\ell-1} \in \{0,1\}^{b(\ell-1)}$ at random and then setting $M_1 = M \| x_1$ and $M_2 = M \| x_2$ for some distinct blocks $x_1, x_2 \in \{0,1\}^b$. Let $Z_0 := K_1$ and $Z_i := \mathsf{f}(Z_{i-1}, m_i)$ for $i \in \{1, \ldots, \ell-1\}$. If any of the $\ell-1$ intermediate values $Z_1, \ldots, Z_{\ell-1}$ in the evaluation of the inner function $\mathsf{Casc}^{\mathsf{f}}(K_1, M)$ is in $\mathcal{K}$, then $\mathsf{Casc}^{\mathsf{f}}(K_1, M_i) = w$ for both $i \in \{1,2\}$ and hence also $\mathsf{NMAC}^{\mathsf{f}}(K, M_1) = \mathsf{NMAC}^{\mathsf{f}}(K, M_2)$. We now lower-bound the probability of this event occurring. Since M is chosen independently and uniformly at random, the construction of f from a permutation implies that each value Z_i will also be distributed uniformly at random and independently of $\mathcal{K}$, as long as $Z_{i-1} \notin \mathcal{K}$. Therefore, we obtain

$$\begin{aligned}
\mathsf{P}^{K,M}\left[\{Z_1, \ldots, Z_{\ell-1}\} \cap \mathcal{K} \neq \emptyset\right] &= 1 - \mathsf{P}^{K,M}\left[\{Z_1, \ldots, Z_{\ell-1}\} \cap \mathcal{K} = \emptyset\right] \\
&= 1 - \left(\mathsf{P}^{K,M}\left[Z_0 \notin \mathcal{K}\right] \cdot \prod_{i=1}^{\ell-1} \mathsf{P}^{K,M}\left[Z_i \notin \mathcal{K} | Z_{i-1} \notin \mathcal{K}\right] \right) \\
&= 1 - \left(1 - \frac{\varepsilon_{\mathsf{na}}}{2}\right)^{\ell} \geq \ell \varepsilon_{\mathsf{na}}/4 \,.
\end{aligned}$$

As explained above, this also lower-bounds the probability of a collision between $\mathsf{NMAC}^{\mathsf{f}}(K, M_1)$ and $\mathsf{NMAC}^{\mathsf{f}}(K, M_2)$.

Now, consider an adversary A that queries $\mathsf{NMAC}^{\mathsf{f}}_K$ on $q/2$ such random and independently sampled message pairs M_1, M_2 and outputs 1 if and only if it observes a collision for at least one such pair. A interacting with $\mathsf{NMAC}^{\mathsf{f}}_K$ outputs 1 with probability

$$1 - \left(1 - \frac{\ell \varepsilon_{\mathsf{na}}}{4}\right)^{q/2} \geq \frac{\ell q \varepsilon_{\mathsf{na}}}{16} = \Theta(\ell q \varepsilon_{\mathsf{na}}) \,.$$

However, in the interaction with the random function $\mathbf{R}$, A clearly outputs 1 with probability only $O(q/2^c)$. By our assumption on $\varepsilon_{\mathsf{na}}$, we get $q/2^c = o(\ell q \varepsilon_{\mathsf{na}})$ and hence also $\Delta^{\mathsf{A}}(\mathsf{NMAC}^{\mathsf{f}}_K, \mathbf{R}) = \Omega(\ell q \varepsilon_{\mathsf{na}})$ as desired. □

C Tightness of Lemma 4 and Theorem 2

In this appendix we prove a lower bound, for a particular pair of messages $\mathcal{M} = \{m_1, m_2\}$, on the number of structure graphs that contain exactly one f-collision and where the final vertices $V_1^{\ell_1}$ and $V_2^{\ell_2}$ of their message paths coincide. As in Lemma 4 we consider both the case where the messages are required to have the same length, and the case without this requirement. Recall that $\mathcal{H}(\mathcal{M})$ denotes the set

$$\mathcal{H}(\mathcal{M}) := \left\{ G \in \mathcal{G}^1(\mathcal{M}) : V_1^{\ell_1} = V_2^{\ell_2} \right\}$$

of such graphs and $d'(n) := \max_{n' \in \{1,\ldots,n\}} |\{d \in \mathbb{N} : d \mid n'\}|$. Proposition 3 below shows that Lemma 4 is tight (up to a constant factor 4).

Proposition 3. *There exist two distinct messages $\mathcal{M} = \{m_1, m_2\}$, each of length at most ℓ blocks, such that $|\mathcal{H}(\mathcal{M})| \geq \frac{\ell \cdot d'(\ell)}{4}$. Moreover, if we additionally require $|m_1| = |m_2|$ then there exist two equal-length messages $\mathcal{M} = \{m_1, m_2\}$ of length at most ℓ blocks such that $|\mathcal{H}(\mathcal{M})| \geq \ell$.*

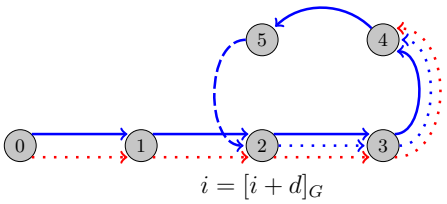

Fig. 6. A sample graph $G_{i,d}$ for the proof of Proposition 3.

Proof. Again, let us first consider the case where $|m_1| = |m_2|$ is *not* required. Given ℓ, let $\ell' \leq \ell/2$ be any positive integer such that $d(\ell') = d'(\ell/2)$ (it exists by the definition of $d'(\cdot)$). We choose $m_1, m_2 \in \{0,1\}^{b*}$ to be the messages consisting of $\ell/2 + \ell'$ and $\ell/2$ equal blocks 0^b, respectively. Now, we describe $\ell \cdot d'(\ell)/4$ distinct structure graphs and show that they are all in $\mathcal{H}(\mathcal{M})$, thus establishing the proof of the first part.

For every $i \in \{0, \ldots, \ell/2\}$ and every d that is a divisor of ℓ', we denote by $G_{i,d}$ the structure graph constructed as follows: Informally, the graph corresponding to m_1 starts with a path of length $i + d - 1$ edges, and the $(i + d)$-th edge returns to vertex i, hence causing a collision. Note that now we have a ρ-shaped graph (where the cycle has length d), and the remaining edges of m_1 must follow the edges along the cycle in that graph. Since m_2 is a prefix of m_1, this also determines the m_2-path (see Figure 6 for a sample $G_{i,d}$). Formally, $G_{i,d} := (\mathcal{V}, \mathcal{E}, \mathcal{L})$ where

$$\begin{aligned} \mathcal{V} &:= \{0, \ldots, i + d - 1\}, \\ \mathcal{E} &:= \{(j-1, j) \mid 1 \leq j \leq i + d - 1\} \cup \{(i + d - 1, i)\} \text{ and} \\ \mathcal{L}(u, v) &:= \{0^b\} \text{ for all } (u, v) \in \mathcal{E}\ . \end{aligned}$$

It is clear from the definition of $G_{i,d}$ that for distinct $(i, d) \neq (i', d')$ we also have $G_{i,d} \neq G_{i',d'}$. Moreover, we claim that for each (i, d) chosen as above, $G_{i,d} \in \mathcal{H}(\mathcal{M})$. To see this, observe that the m_1-path ends in the vertex $i + (\ell/2 + \ell' - i \mod d)$, while the m_2-path ends in the vertex $i + (\ell/2 - i \mod d)$. Since $d | \ell'$, this is actually the same vertex and we have $V_1^{\ell_1} = V_2^{\ell_2}$, establishing $G_{i,d} \in \mathcal{H}(\mathcal{M})$. There are $(\ell/2 + 1) \cdot d(\ell')$ ways to choose a tuple (i, d) with $i \in \{0, \ldots, \ell/2\}$ and d being a divisor of ℓ', and thus $\mathcal{H}(\mathcal{M})$ has at least $(\ell/2 + 1) \cdot d(\ell') \geq \ell/2 \cdot d'(\ell/2) \geq \ell d'(\ell)/4$ distinct elements as claimed.

For the case $|m_1| = |m_2|$, consider the messages $m_1 = 1^b 0^{b(\ell-1)}$ and $m_2 = 01^{b-1}0^{b(\ell-1)}$. These messages are both of length ℓ blocks and differ in their first blocks, while the remaining $\ell - 1$ blocks consist of zeroes in both messages. We again construct ℓ distinct structure graphs and show that they all belong to $\mathcal{H}(\mathcal{M})$.

For every $i \in \{1, \ldots, \ell\}$, we denote by $G(i)$ the structure graph constructed as follows: Informally, the subgraph corresponding to m_1 is a path of length ℓ, not containing any collision itself. Since m_2 differs from m_1 in the first block, the m_2-path will not overlap with the m_1-path as long as no f-collision occurs. In the graph $G(i)$, we let this collision happen for the i-th edge of the m_2-path, hitting the vertex i on the m_1-path. In particular, in the case $i = 1$ the collision occurs by having $V_1^1 = V_2^1$ even though the first blocks of the messages differ. See Figure 7 for a sample

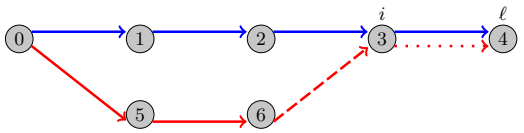

$\ell = 4,\ i = 3$
$m_1 = 1^b 0^{3b},\ m_2 = 01^{b-1}0^{3b}$

Fig. 7. A sample graph $G(i)$ for the proof of Proposition 3.

$G(i)$. Formally, $G(i) := (\mathcal{V}, \mathcal{E}, \mathcal{L})$ where

$$\mathcal{V} := \{0, \ldots, \ell + i - 1\},$$

$$\mathcal{E} := \begin{cases} \{(j-1, j) \mid 1 \le j \le \ell\} & \text{if } i = 1 \\ \{(j-1, j) \mid j \in \{1, \ldots, \ell + i - 1\} \setminus \{\ell + 1\}\} \cup \{(0, \ell + 1)\} \cup \{(\ell + i - 1, i)\} & \text{if } i > 1 \end{cases}$$

$$\mathcal{L}(u, v) := \begin{cases} \{1^b, 01^{b-1}\} & \text{if } (u, v) = (0, 1) \text{ and } i = 1 \\ \{1^b\} & \text{if } (u, v) = (0, 1) \text{ and } i > 1 \\ \{0^b\} & \text{if } (u, v) \in \{(j-1, j) \mid j \in \{1, \ldots, \ell + i - 1\} \setminus \{\ell + 1\}\} \\ \{0^b\} & \text{if } (u, v) = (\ell + i - 1, i) \text{ and } i > 1 \\ \{01^{b-1}\} & \text{if } (u, v) = (0, \ell + 1) \text{ and } i > 1 \\ \emptyset & \text{otherwise.} \end{cases}$$

Again, it is clear from the definition of $G(i)$ that for distinct $i \neq i'$ we also have $G(i) \neq G(i')$. Moreover, it is easy to see that for each $i \in \{1, \ldots, \ell\}$ we have $G(i) \in \mathcal{H}(\mathcal{M})$. This proves that in this case $|\mathcal{H}(\mathcal{M})| \ge \ell$ as desired. □

Finally, the ideas from the proof of Proposition 3 above can be used to give a simple non-adaptive distinguishing attack achieving advantage $\Theta(\ell q^2/2^c)$ against $\mathsf{LenCasc}^{\mathbf{f}_1}_{\mathbf{0}} \triangleright \mathbf{f}_2$, i.e., against the system that we obtain after replacing h in NI^{h} by a random compression function. We sketch this attack below, hence showing that the information-theoretic analysis in Theorem 2 is tight.

The adversary simply chooses q messages $m_1, \ldots, m_q$ of the form $m_i = x_i \| 0^{b(\ell-1)}$ for arbitrary distinct x_i's. For any $1 \le i < j \le q$ and any $1 \le p \le \ell$, we will have a collision $\mathbf{f}_2(\mathsf{LenCasc}^{\mathbf{f}_1}_{\mathbf{0}}(m_i)) = \mathbf{f}_2(\mathsf{LenCasc}^{\mathbf{f}_1}_{\mathbf{0}}(m_j))$ if the outputs after computing the inner cascade $\mathsf{Casc}^{\mathbf{f}_1}_{\mathbf{0}}$ on the p-block prefixes of m_i and m_j collide (as their suffixes and lengths are identical, and thus such a collision implies that also the final values collide). The probability that for any fixed (i, j, p) this happens, conditioned on that this collision is not predetermined (i.e., either $p = 1$ or $\mathsf{Casc}^{\mathbf{f}_1}_{\mathbf{0}}$ applied to the $(p-1)$-block prefixes did not collide) is roughly 2^{-c} as long as $\ell \ll 2^{c/2}$. We can choose triples (p, i, j) in $\ell q(q-1)/2 = \Theta(\ell q^2)$ ways, and as just explained every such triple defines a possible event that leads to a collision and has probability $\approx 2^{-c}$ (and these events are disjoint as we required the collisions not to be predetermined), hence this gives the claimed $\Theta(\ell q^2/2^c)$ bound.