

A preliminary version of this paper appears in *Advances in Cryptology – CRYPTO '06*, Lecture Notes in Computer Science Vol. , ed., Springer-Verlag, 2006. This is the full version.

New Proofs for NMAC and HMAC: Security without Collision-Resistance

MIHIR BELLARE*

June 2006

Abstract

HMAC was proved in [3] to be a PRF assuming that (1) the underlying compression function is a PRF, and (2) the iterated hash function is weakly collision-resistant. However, recent attacks show that assumption (2) is false for MD5 and SHA-1, removing the proof-based support for HMAC in these cases. This paper proves that HMAC is a PRF under the sole assumption that the compression function is a PRF. This recovers a proof based guarantee since no known attacks compromise the pseudorandomness of the compression function, and it also helps explain the resistance-to-attack that HMAC has shown even when implemented with hash functions whose (weak) collision resistance is compromised. We also show that an even weaker-than-PRF condition on the compression function, namely that it is a privacy-preserving MAC, suffices to establish HMAC is a secure MAC as long as the hash function meets the very weak requirement of being computationally almost universal, where again the value lies in the fact that known attacks do not invalidate the assumptions made.

Keywords: Message authentication, hash functions, Pseudorandom Functions, Carter-Wegman.

*Dept. of Computer Science & Engineering 0404, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0404, USA. Email: mihir@cs.ucsd.edu. URL: <http://www-cse.ucsd.edu/users/mihir>. Supported in part by NSF grant CNS-0524765, an IBM Faculty Partnership Development Award, and a gift from INTEL Corporation.

Contents

1	Introduction	3
2	Definitions	5
3	Security of NMAC	7
3.1	The results	7
3.2	Tightness of bound	9
3.3	Proof of Lemma 3.1	9
3.4	Proof of Lemma 3.2	15
3.5	Proof of Theorem 3.4	18
4	MAC-security of NMAC under weaker assumptions	19
4.1	Privacy preserving MACs	19
4.2	Results	20
4.3	Proof of Lemma 4.2	21
5	Security of HMAC	24
5.1	Security of GHMAC	25
5.2	Single-keyed HMAC	26
5.3	Lifting the results of Section 4	27
5.4	Remarks	27
	References	27
	A Attacking weak collision-resistance	29
	B The reduction-from-pf-PRF proof	30

1 Introduction

HMAC [3] is a popular cryptographic-hash-function-based MAC. The basic construct is actually NMAC, of which HMAC can be viewed as a derivative.

THE CONSTRUCTIONS. Succinctly:

$$\begin{aligned}\text{NMAC}(K_{\text{out}}\|K_{\text{in}}, M) &= H^*(K_{\text{out}}, H^*(K_{\text{in}}, M)) \\ \text{HMAC}(K_{\text{out}}\|K_{\text{in}}, M) &= H(K_{\text{out}}\|H(K_{\text{in}}\|M)).\end{aligned}$$

Here H is a cryptographic hash function, eg. MD5 [28], SHA-1 [26], or RIPEMD-160 [17]. Let $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ denote the underlying compression function. (Here $b = 512$ while c is 128 or 160.) Let h^* be the iterated compression function which on input $K \in \{0, 1\}^c$ and a message $x = x[1] \dots x[n]$ consisting of b -bit blocks, lets $a[0] = K$ and $a[i] = h(a[i-1], x[i])$ for $i = 1, \dots, n$, and finally returns $a[n]$. Then $H^*(K, M) = h^*(K, M^*)$ and $H(M) = H^*(\text{IV}, M)$, where M^* denotes M padded appropriately to a length that is a positive multiple of b and IV is a public c -bit initial vector that is fixed as part of the description of H . Both NMAC and HMAC use two keys, which in the case of NMAC are of length c bits each, and in the case of HMAC of length b bits each and derived from a single b -bit key. HMAC is a non-intrusive version of NMAC in the sense that it uses the cryptographic hash function only as a black box, making it easier to implement.

USAGE. HMAC is standardized via an IETF RFC [22], a NIST FIPS [25] and ANSI X9.71 [1], and implemented in SSL, SSH, IPsec and TLS amongst other places. It is often used as a PRF (pseudorandom function [19]) rather than merely as a MAC. In particular this is the case when it is used for key-derivation, as in TLS [16] and IKE (the Internet Key Exchange protocol of IPsec) [20]. HMAC is also used as a PRF in a proposed standard for one-time passwords [24].

WHAT'S KNOWN. The results are for NMAC but can be lifted to HMAC. It is shown in [3] that NMAC is a secure PRF if (1) the underlying compression function h is a secure PRF, and also (2) that the hash function H is *weakly collision resistant* (WCR). The latter, introduced in [3], is a relaxation of collision resistance (CR) that asks that it be computationally infeasible for an adversary, given an oracle for $H^*(K, \cdot)$ under a hidden key K , to find a collision, meaning distinct inputs M_1, M_2 such that $H^*(K, M_1) = H^*(K, M_2)$.

THE PROBLEM. HMAC is usually implemented with MD5 or SHA-1. But, due to recent attacks [32, 31], these functions are *not* WCR. Thus the assumption on which the proof of [3] is based is not true. This does not reflect any actual weaknesses in the NMAC or HMAC constructs, on which no attacks are known. (Being iterated MACs, the generic birthday based forgery attacks of [27] always break NMAC and HMAC in time $2^{c/2}$, but we mean no better-than-birthday attacks are known.) But it means that we have lost the proof-based guarantees. We are interested in recovering them.

LOSS OF WCR. First we pause to expand on the claim above that our main hash functions are not WCR. Although WCR appears to be a weaker requirement than CR due to the hidden key, in fact, for iterated hash functions, it ends up not usually being so. The reason is that collision-finding attacks such as those on MD5 [32] and SHA-1 [31] extend to find collisions in $H^*(\text{IV}, \cdot)$ for an arbitrary but given IV, and, any such attack, via a further extension attack, can be used to compromise WCR, meaning to find a collision in $H^*(K, \cdot)$, given an oracle for this function, even with K hidden. This was pointed out in [3, 21], and, for the curious, we recall the attack in Appendix A.

MAIN RESULT. We show (Theorem 3.3) that NMAC is a PRF under the *sole* assumption that the underlying compression function h is itself a PRF. In other words, the additional assumption that

the hash function is WCR is dropped. (And, in particular, as long as h is a PRF, the conclusion is true even if H is *not* WCR, let alone CR.)

The main advantage of our result is that it is based on an assumption that is not refuted by any known attacks. (There are to date no attacks that compromise the pseudorandomness of the compression functions of MD5 or SHA-1.) Another feature of our result is that it is the first proof for NMAC that is based solely on an assumption about the compression function rather than also assuming something about the entire iterated hash function.

TECHNIQUES. We show (Lemma 3.1) that if a compression function h is a PRF then the iterated compression function h^* is *computationally almost universal* (cAU), a computational relaxation of the standard information-theoretic notion of almost-universality (AU) of [13, 33, 30]. We then conclude with Lemma 3.2 which says that the composition of a PRF and a cAU function is a PRF. (This can be viewed as a computational relaxation of the Carter-Wegman paradigm [13, 33].)

RELATED WORK. If h^* were a PRF, it would imply it is cAU, but h^* is not a PRF due to the extension attack. It is however shown by [4] that if h is a PRF then h^* (which they call the cascade) is a “pf-PRF” (prefix-free PRF), meaning a PRF as long as no query of the adversary is a prefix of another query. It was pointed out to us by Victor Shoup after seeing an early draft of our paper that it is possible to apply this in a black-box way to show that h^* is cAU. Lemma 3.1 however obtains a better bound. Although this is by only a constant factor in the context of the lemma itself, Theorem 3.4 shows that it can translate into an appreciably better security guarantee in the context of NMAC itself. The proof of Lemma 3.1 exploits and extends the ideas of [4] while also adding some new ones. For comparison, we do present the indirect proof in Appendix B and contrast the bounds obtained with ours.

Dodis, Gennaro, Håstad, Krawczyk and Rabin show [18, Lemma 4] that the cascade over a family of *random* functions is AU as long as the two messages whose collision probability one considers have the same length. (In this model, $h(K, \cdot)$ is a random function for each $K \in \{0, 1\}^c$. That is, it is like Shannon’s ideal cipher model, except the component maps are functions not permutations.) This does not imply Lemma 3.1 (showing the cascade h^* is cAU if h is a PRF), because we need to allow the two messages to have different lengths, and also because it is not clear what implication their result has for the case when h is a PRF. (A PRF does *not* permit one to securely instantiate a *family* of random functions.) A second result [18, Lemma 5] in the same paper says that if $h^*(K, M)$ is close to uniformly distributed then so is $h^*(K, M||X)$. (Here M is chosen from some distribution, K is a random but known key, and X is a fixed block.) This result only assumes h is a PRF, but again we are not able to discern any implications for the problems we consider, because in our case the last block of the input is not fixed, we are interested in the cAU property rather than randomness, and our inputs are not drawn from a distribution.

ANOTHER RESULT. The fact that compression functions are underlain by block ciphers, together with the fact that no known attacks compromise the pseudorandomness of the compression functions of MD5, SHA-1, may give us some confidence that it is ok to assume these are PRFs, but it still behooves us to be cautious. What can we prove about NMAC without assuming the compression function is a PRF? We would not expect to be able to prove it is a PRF, but what about just a secure MAC? (Any PRF is a secure MAC [6, 9], so our main result implies NMAC is a secure MAC, but we are interested in seeing whether this can be proved under weaker assumptions.) We show (Theorem 4.3) that NMAC is a secure MAC if h is a *privacy-preserving MAC* (PP-MAC) [8] and h^* (equivalently, H^*) is cAU. A PP-MAC (the definition is provided in Section 4) is stronger than a MAC but weaker than a PRF. This result reverts to the paradigm of [3] of making assumptions both about the compression function and its iteration, but the point is that cAU is a very weak

assumption compared to WCR and PP-MAC is a weaker assumption than PRF.

FROM NMAC TO HMAC. The formal results (both previous and new) we have discussed so far pertain to NMAC. However, discussions (above and in the literature) tend to identify NMAC and HMAC security-wise. This is explained by an observation of [3] which says that HMAC inherits the security of NMAC as long as the compression function is a PRF when keyed via the data input. (So far when we have talked of it being a PRF, it is keyed via the chaining variable.) In our case this means that HMAC is a PRF if the compression function is a “dual-PRF,” meaning a PRF when keyed by either of its two inputs.

However, the analysis above assumes that the two keys $K_{\text{out}}, K_{\text{in}}$ of HMAC are chosen *independently* at random, while in truth they are equal to $K \oplus \text{opad}$ and $K \oplus \text{ipad}$ respectively, where K is a random b -bit key and opad, ipad are fixed, distinct constants. We apply the theory of PRFs under related-key attacks [7] to extend the observation of [3] to this single-key version of HMAC, showing it inherits the security of NMAC as long as the data-input-keyed compression function is a PRF under an appropriate (and small) class of related key attacks. Assuming additionally that the compression function is a PRF in the usual sense, we obtain a (in fact, the first) security proof of the single-key version of HMAC. These results are in Section 5.

VERSIONS. A preliminary version of this paper appears in the proceedings of CRYPTO’06 [2]. The full version you are now reading contains a different and simpler proof of Claim 3.8. It also includes an additional result, namely Theorem 3.4. It contains numerous proofs omitted from the preliminary version due to lack of space. The appendices here were also omitted from the preliminary version for lack of space.

2 Definitions

NOTATION. We denote by $s_1 \| s_2$ the concatenation of strings s_1, s_2 , and by $|s|$ the length of string s . Let b be a positive integer representing a block length, and let $B = \{0, 1\}^b$. Let B^+ denote the set of all strings of length a positive multiple of b bits. Whenever we speak of blocks we mean b -bit ones. If $M \in B^+$ then $\|M\|_b = |M|/b$ is the number of blocks in M , and $M[i]$ denotes its i -th b -bit block, meaning $M = M[1] \dots M[n]$ where $n = \|M\|_b$. If $M_1, M_2 \in B^+$, then M_1 is a prefix of M_2 , written $M_1 \subseteq M_2$, if $M_2 = M_1 \| A$ for some $A \in B^*$. If S is a set then $s \stackrel{\$}{\leftarrow} S$ denotes the operation of selecting s uniformly at random from S . An adversary is a (possibly randomized) algorithm that may have access to one or more oracles. We let

$$A^{\mathcal{O}_1, \dots}(x_1, \dots) \Rightarrow 1 \quad \text{and} \quad y \stackrel{\$}{\leftarrow} A^{\mathcal{O}_1, \dots}(x_1, \dots)$$

denote, respectively, the event that A with the indicated oracles and inputs outputs 1, and the experiment of running A with the indicated oracles and inputs and letting y be the value returned. (This value is a random variable depending on the random choices made by A and its oracles.) Either the oracles or the inputs (or both) may be absent, and often will be.

A family of functions is a two-argument map $f: \text{Keys} \times \text{Dom} \rightarrow \text{Rng}$ whose first argument is regarded as a key. We fix one such family $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ to model a compression function that we regard as being keyed via its c -bit chaining variable. Typical values are $b = 512$ and $c = 128$ or 160 . The *iteration* of family $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ is the family of functions $h^*: \{0, 1\}^c \times B^+ \rightarrow \{0, 1\}^c$ defined via:

Function $h^*(K, M)$ // $K \in \{0, 1\}^c, M \in B^+$
 $n \leftarrow \|M\|_b; a[0] \leftarrow K$
 For $i = 1, \dots, n$ do $a[i] \leftarrow h(a[i-1], M[i])$

Return $a[n]$

This represents the Merkle-Damgård [23, 14] iteration method used in all the popular hash functions but without the “strengthening,” meaning that there is no $|M|$ -based message padding.

PRFs. A prf-adversary A against a family of functions $f: Keys \times Dom \rightarrow Rng$ takes as oracle a function $g: Dom \rightarrow Rng$ and returns a bit. The prf-advantage of A against f is the difference between the probability that it outputs 1 when its oracle is $g = f(K, \cdot)$ for a random key $K \xleftarrow{\$} Keys$, and the probability that it outputs 1 when its oracle g is chosen at random from the set $Maps(Dom, Rng)$ of all functions mapping Dom to Rng , succinctly written as

$$\mathbf{Adv}_f^{\text{prf}}(A) = \Pr \left[A^{f(K, \cdot)} \Rightarrow 1 \right] - \Pr \left[A^{\$} \Rightarrow 1 \right]. \quad (1)$$

In both cases the probability is over the choice of oracle and the coins of A .

cAU AND COLLISION-PROBABILITY. Let $F: \{0, 1\}^k \times Dom \rightarrow Rng$ be a family of functions. cAU is measured by considering an almost-universal (au) adversary A against F . It (takes no inputs and) returns a pair of messages in Dom . Its au-advantage is

$$\mathbf{Adv}_F^{\text{au}}(A) = \Pr \left[F(K, M_1) = F(K, M_2) \wedge M_1 \neq M_2 : (M_1, M_2) \xleftarrow{\$} A; K \xleftarrow{\$} Keys \right].$$

This represents a very weak form of collision-resistance since A must produce M_1, M_2 without being given any information about K . WCR [3] is a stronger notion because here A gets an oracle for $F(K, \cdot)$ and can query this in its search for M_1, M_2 .

For $M_1, M_2 \in Dom$ it is useful to let $\text{Coll}_F(M_1, M_2) = \Pr[F(K, M_1) = F(K, M_2)]$, the probability being over $K \xleftarrow{\$} \{0, 1\}^k$.

SECURITY AND RESOURCES. As usual with the concrete security approach [6], there is no formal notion of security of a primitive (eg. a PRF) but informally, when we talk of, say, f being a PRF, we mean that the prf-advantage of any (no input) prf-adversary of “practical” resources is “low.” Similarly H is computationally au (cAU) if the au-advantage of any adversary of “practical” resources is “low”. Formal results state the concrete security of reductions, bounding the advantage and resources of one adversary as a function of those of others.

The following conventions will be adopted in measuring resource usage. The time-complexity of an adversary is defined as the total execution time of an overlying experiment (meaning, includes not only the running time of the adversary but the time to compute replies to oracle queries and the time to perform any initializations or to test whether the adversary was successful) plus the size of the code of the adversary, in some fixed model of computation. (In cases like PRFs, where equation (1) defining the advantage involves two experiments, we consider the maximum of the two execution times, with the convention that the picking of a random function is done by building an on-line table while responding to oracle queries. More details on these conventions will appear when they are used.) When we say that a resource measure (such as the time-complexity, number of oracle queries, or their lengths) is at most a certain value, we mean this holds for all coin tosses of the adversary and regardless of how its oracle queries are answered. With regard to time-complexity we will permit ourselves the use of big-oh notation, even though there are no asymptotics here, with the intent that it hides some constant depending only on the model of computation.

MACs. Any PRF is a MAC. (This was established for the basic notion of MAC security in [6], but holds even for the most stringent notions and with tight security reductions [9].) Accordingly, we do not explicitly discuss MACs until Section 4.1 where we consider PP-MACs.

3 Security of NMAC

Let $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ be a family of functions that represents the compression function, here assumed to be a PRF. Let pad denote a padding function such that $s^* = s \parallel \text{pad}(|s|) \in B^+$ for any string s . (Such padding functions are part of the description of current hash functions. Note the pad depends only on the length of s .) Then the family NMAC: $\{0, 1\}^{2c} \times D \rightarrow \{0, 1\}^c$ is defined by $\text{NMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = h(K_{\text{out}}, h^*(K_{\text{in}}, M^*) \parallel \text{fpad})$ where $\text{fpad} = \text{pad}(c) \in \{0, 1\}^{b-c}$ and h^* is the iterated compression function as defined in Section 2. The domain D is the set of all strings up to some maximum length, which is 2^{64} for current hash functions.

It turns out that our security proof for NMAC does not rely on any properties of pad beyond the fact that $M^* = M \parallel \text{pad}(|M|) \in B^+$. (In particular, the Merkle-Damgård strengthening, namely inclusion of the message length in the padding, that is used in current hash functions and is crucial to collision resistance of the hash function, is not important to the security of NMAC.) Accordingly, we will actually prove the security of a more general construct that we call generalized NMAC. The family GNMAC: $\{0, 1\}^{2c} \times B^+ \rightarrow \{0, 1\}^c$ is defined by $\text{GNMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = h(K_{\text{out}}, h^*(K_{\text{in}}, M) \parallel \text{fpad})$ where fpad is any (fixed) $b-c$ bit string. Note the domain is B^+ , meaning inputs have to have a length that is a positive multiple of b bits. (But can be of any length.) NMAC is nonetheless a special case of GNMAC via $\text{NMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = \text{GNMAC}(K_{\text{out}} \parallel K_{\text{in}}, M^*)$ and thus the security of NMAC is implied by that of GNMAC. (Security as a PRF or a MAC, respectively, for both.)

3.1 The results

MAIN LEMMA. The following says that if h is a PRF then its iteration h^* is cAU.

Lemma 3.1 *Let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions and let $n \geq 1$ be an integer. Then there exists a prf-adversary A against h such that the following is true. Let A^* be any au-adversary against h^* such that the two messages output by A^* are at most n_1, n_2 blocks long, respectively, where $1 \leq n_1, n_2 \leq n$. Then*

$$\mathbf{Adv}_{h^*}^{\text{au}}(A^*) \leq (n_1 + n_2 - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A) + \frac{1}{2^c}. \quad (2)$$

Furthermore, A has time-complexity at most $O(nT_h)$, where T_h is the time for one evaluation of h , and makes at most 2 oracle queries. ■

Note that A depends only on (h and) the integer n but not on A^* . (The inequality (2) is true for all A^* the block length of whose output messages satisfy the stated constraints.) In particular the time-complexity of A is small and independent of the time-complexity of A^* . Furthermore A makes only two oracle queries. The proof, which is in Section 3.3, shows how all this is possible.

One might ask whether stronger results hold. For example, assuming h is a PRF, (1) Is h^* a PRF? (2) Is h^* WCR? (Either would imply that h^* is cAU.) But the answer is NO to both questions. The function h^* is *never* a PRF due to the extension attack. On the other hand it is easy to give an example of a PRF h such that h^* is not WCR. Also MD5 and SHA-1 are candidate counter-examples, since their compression functions appear to be PRFs but their iterations are not WCR.

THE PRF(CAU)=PRF LEMMA. The *composition* of families $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ and $F: \{0, 1\}^k \times D \rightarrow \{0, 1\}^b$ is the family $hF: \{0, 1\}^{c+k} \times D \rightarrow \{0, 1\}^c$ defined by $hF(K_{\text{out}} \parallel K_{\text{in}}, M) = h(K_{\text{out}}, F(K_{\text{in}}, M))$. The following lemma says that if h is a PRF and F is cAU then hF is a PRF.

Lemma 3.2 Let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ and $F: \{0, 1\}^k \times D \rightarrow B$ be families of functions, and let $hF: \{0, 1\}^{c+k} \times D \rightarrow \{0, 1\}^c$ be defined by

$$hF(K_{\text{out}} \| K_{\text{in}}, M) = h(K_{\text{out}}, F(K_{\text{in}}, M))$$

for all $K_{\text{out}} \in \{0, 1\}^c$, $K_{\text{in}} \in \{0, 1\}^k$ and $M \in D$. Let A_{hF} be a prf-adversary against hF that makes at most $q \geq 2$ oracle queries, each of length at most n , and has time-complexity at most t . Then there exists a prf-adversary A_h against h and an au-adversary A_F against F such that

$$\mathbf{Adv}_{hF}^{\text{prf}}(A_{hF}) \leq \mathbf{Adv}_h^{\text{prf}}(A_h) + \binom{q}{2} \cdot \mathbf{Adv}_F^{\text{au}}(A_F). \quad (3)$$

Furthermore, A_h has time-complexity at most t and makes at most q oracle queries, while A_F has time-complexity $O(T_F(n))$ and the two messages it outputs have length at most n , where $T_F(n)$ is the time to compute F on an n -bit input. ■

This extends the analogous PRF(AU)=PRF lemma by relaxing the condition on F from AU to cAU. (The PRF(AU)=PRF lemma is alluded to in [4, 11], and variants are in [11, 12].) A simple proof of Lemma 3.2, using games [10, 29], is in Section 3.4.

The reduction of Lemma 3.2 may look loose due to the $\binom{q}{2}$ factor. (And in some settings this is indeed the case and has led to the use of alternative constructs [11].) However in our case this factor turns out only to reflect the existing birthday attack on GNMAC [27] and thus does not reflect a loose reduction. Note that the time-complexity of A_F is small and in particular independent of the time-complexity of A_{hF} .

GNMAC IS A PRF. We now combine the two lemmas above to conclude that if h is a PRF then so is GNMAC.

Theorem 3.3 Assume $b \geq c$ and let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions and let $\text{fpad} \in \{0, 1\}^{b-c}$ be a fixed padding string. Let GNMAC: $\{0, 1\}^{2c} \times B^+ \rightarrow \{0, 1\}^c$ be defined by

$$\text{GNMAC}(K_{\text{out}} \| K_{\text{in}}, M) = h(K_{\text{out}}, h^*(K_{\text{in}}, M) \| \text{fpad})$$

for all $K_{\text{out}}, K_{\text{in}} \in \{0, 1\}^c$ and $M \in B^+$. Let A_{GNMAC} be a prf-adversary against GNMAC that makes at most $q \geq 2$ oracle queries, each of at most m blocks, and has time-complexity at most t . Then there exist prf-adversaries A_1, A_2 against h such that

$$\mathbf{Adv}_{\text{GNMAC}}^{\text{prf}}(A_{\text{GNMAC}}) \leq \mathbf{Adv}_h^{\text{prf}}(A_1) + \binom{q}{2} \left[2m \cdot \mathbf{Adv}_h^{\text{prf}}(A_2) + \frac{1}{2^c} \right]. \quad (4)$$

Furthermore, A_1 has time-complexity at most t and makes at most q oracle queries, while A_2 has time-complexity at most $O(mT_h)$ and makes at most 2 oracle queries, where T_h is the time for one computation of h . ■

Proof of Theorem 3.3: Define $F: \{0, 1\}^c \times B^+ \rightarrow \{0, 1\}^b$ by $F(K_{\text{in}}, M) = h^*(K_{\text{in}}, M) \| \text{fpad}$. Then GNMAC = hF . Apply Lemma 3.2 (with $k = c$, $D = B^+$ and $A_{hF} = A_{\text{GNMAC}}$) to get prf-adversary A_1 and au-adversary A_F with the properties stated in the lemma. Note that $\mathbf{Adv}_F^{\text{au}}(A_F) = \mathbf{Adv}_{h^*}^{\text{au}}(A_F)$. (Because a pair of messages is a collision for $h^*(K_{\text{in}}, \cdot) \| \text{fpad}$ iff it is a collision for $h^*(K_{\text{in}}, \cdot)$.) Let $n = m$ and let A_2 be the prf-adversary A of Lemma 3.1. Apply the lemma to $A^* = A_F$. ■

The upper bound on the advantage of A_{GNMAC} in (4) is a function of a single assumed bound m on the block lengths of all messages queried by A_{GNMAC} . The following more general result expresses the bound in terms of assumed bounds on the block lengths of each query taken individually:

Theorem 3.4 *Assume $b \geq c$ and let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions and let $\text{fpad} \in \{0, 1\}^{b-c}$ be a fixed padding string. Let $\text{GNMAC}: \{0, 1\}^{2c} \times B^+ \rightarrow \{0, 1\}^c$ be defined as in Theorem 3.3. Let A_{GNMAC} be a prf-adversary against GNMAC that makes at most $q \geq 2$ oracle queries, where the i -th query is of at most $m_i \geq 1$ blocks for $1 \leq i \leq q$. Let $n = m_1 + \dots + m_q$ and $m = \max(m_1, \dots, m_q)$. Assume A_{GNMAC} has time-complexity at most t . Then there exist prf-adversaries A_1, A_2 against h such that*

$$\mathbf{Adv}_{\text{GNMAC}}^{\text{prf}}(A_{\text{GNMAC}}) \leq \mathbf{Adv}_h^{\text{prf}}(A_1) + (q - 1) \cdot (n - q/2) \cdot \mathbf{Adv}_h^{\text{prf}}(A_2) + \binom{q}{2} \frac{1}{2^c}. \quad (5)$$

Furthermore, A_1 has time-complexity at most t and makes at most q oracle queries, while A_2 has time-complexity at most $O(mT_h)$ and makes at most 2 oracle queries, where T_h is the time for one computation of h . ■

Theorem 3.4 is more general than Theorem 3.3 in the sense that the former implies the latter. Indeed, if we set $m_i = m$ for all $1 \leq i \leq q$ then we get $n = mq$ and (5) implies (4). On the other hand, Theorem 3.4 can give significantly better security guarantees than Theorem 3.3 in cases where the messages being authenticated are of varying lengths. For example suppose A_{GNMAC} makes a single query of block length m followed by $q-1$ queries, each a single block. Then the bound from (4) multiplies $\mathbf{Adv}_h^{\text{prf}}(A_2)$ by $q(q-1)m$ whereas the bound from (5) multiplies $\mathbf{Adv}_h^{\text{prf}}(A_2)$ by $(q-1)(m+q/2-1)$. These are thus different by a $\Theta(q)$ factor if, say, $m = \Theta(q)$. The proof of Theorem 3.4 can be found in Section 3.5.

3.2 Tightness of bound

The best known attack on GNMAC is the birthday one of [27]. They show that it is possible to break GNMAC with about $2^{c/2}/\sqrt{m}$ queries of at most m blocks each. We now want to assess how close to this is the guarantee we can get from Theorem 3.3 gets. If t is a time-complexity then let $\bar{t} = t/T_h$. Assume that the best attack against h as a PRF is exhaustive key search. (Birthday attacks do not apply since h is not a family of permutations.) This means that $\mathbf{Adv}_h^{\text{prf}}(A) \leq \bar{t} \cdot 2^{-c}$ for any prf-adversary A of time complexity t making $q \leq \bar{t}$ queries. Plugging this into (4) and simplifying, the upper bound on the prf-advantage of any adversary against GNMAC who has time-complexity t and makes at most q queries is $O(\bar{t} + m^2 q^2 T_h) \cdot 2^{-c}$. If we ignore the T_h term, then this hits 1 when $q \approx 2^{c/2}/m$. This means that the bound justifies NMAC up to roughly $2^{c/2}/m$ queries, off from the number in the above-mentioned attack by a factor of \sqrt{m} . It is an interesting open problem to improve our analysis and fill the gap.

3.3 Proof of Lemma 3.1

SOME DEFINITIONS. In this proof it will be convenient to consider prf-adversaries that take inputs. The advantage of A against h on inputs x_1, \dots is defined as

$$\mathbf{Adv}_h^{\text{prf}}(A(x_1, \dots)) = \Pr \left[A^{h(K, \cdot)}(x_1, \dots) \Rightarrow 1 \right] - \Pr \left[A^{\$}(x_1, \dots) \Rightarrow 1 \right],$$

where in the first case $K \xleftarrow{\$} \{0, 1\}^c$ and in the second case the notation means that A is given as oracle a map chosen at random from $\text{Maps}(\{0, 1\}^b, \{0, 1\}^c)$.

OVERVIEW. To start with, we ignore A_{hF} and upper bound $\text{Coll}_F(M_1, M_2)$ as some appropriate function of the prf-advantage of a prf-adversary against h that takes M_1, M_2 as input. We consider first the case that $M_1 \subseteq M_2$ (M_1 is a prefix of M_2) and then the case that $M_1 \not\subseteq M_2$, building in each case a different adversary.

<p>Game $G_1(M_1, M_2, l) \ // \ 0 \leq l \leq \ M_1\ _b$</p> <p>$m_1 \leftarrow \ M_1\ _b ; m_2 \leftarrow \ M_2\ _b$</p> <p>$a[l] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>For $i = l + 1$ to m_2 do</p> <p style="padding-left: 20px;">$a[i] \leftarrow h(a[i - 1], M_2[i])$</p> <p>If $a[m_1] = a[m_2]$ then return 1</p> <p style="padding-left: 20px;">else return 0</p> <p>Adversary $A_3^g(M_1, M_2)$</p> <p>$m_1 \leftarrow \ M_1\ _b ; m_2 \leftarrow \ M_2\ _b$</p> <p>$l \xleftarrow{\\$} \{1, \dots, m_1 + 1\}$</p> <p>If $l = m_1 + 1$ then return $A_2^g(M_1, M_2)$</p> <p>Else return $A_1^g(M_1, M_2, l)$</p>	<p>Adversary $A_1^g(M_1, M_2, l) \ // \ 1 \leq l \leq \ M_1\ _b$</p> <p>$m_1 \leftarrow \ M_1\ _b ; m_2 \leftarrow \ M_2\ _b$</p> <p>$a[l] \leftarrow g(M_2[l])$</p> <p>For $i = l + 1$ to m_2 do</p> <p style="padding-left: 20px;">$a[i] \leftarrow h(a[i - 1], M_2[i])$</p> <p>If $a[m_1] = a[m_2]$ then return 1</p> <p style="padding-left: 20px;">else return 0</p> <p>Adversary $A_2^g(M_1, M_2)$</p> <p>$m_1 \leftarrow \ M_1\ _b ; m_2 \leftarrow \ M_2\ _b$</p> <p>$a[m_1 + 1] \leftarrow g(M_2[m_1 + 1])$</p> <p>For $i = m_1 + 2$ to m_2 do</p> <p style="padding-left: 20px;">$a[i] \leftarrow h(a[i - 1], M_2[i])$</p> <p>$y \xleftarrow{\\$} B \setminus \{M_2[m_1 + 1]\}$</p> <p>If $h(a[m_2], y) = g(y)$ then return 1</p> <p style="padding-left: 20px;">else return 0</p>
---	--

Figure 1: Games and adversaries taking input distinct messages M_1, M_2 such that $M_1 \subseteq M_2$. The adversaries take an oracle $g: \{0, 1\}^b \rightarrow \{0, 1\}^c$.

THE CASE $M_1 \subseteq M_2$. We begin with some high-level intuition. Suppose $M_1 \subseteq M_2$ with $m_2 = \|M_2\|_b \geq 1 + m_1$, where $m_1 = \|M_1\|_b$. The argument to upper bound $\text{Coll}_{h^*}(M_1, M_2)$ has two parts. First, a hybrid argument is used to show that $a[m_1] = h^*(K, M_1)$ is computationally close to random when K is drawn at random. Next, we imagine a game in which $a[m_1]$ functions as a key to h . Let $a[m_1 + 1] = h(a[m_1], M_2[m_1 + 1])$ and $a[m_2] = h^*(a[m_1 + 1], M_2[m_1 + 2] \dots M_2[m_2])$. Now, if $a[m_2] = a[m_1]$ then we effectively have a way to recover the “key” $a[m_1]$ given $a[m_1 + 1]$, amounting to a key-recovery attack on $h(a[m_1], \cdot)$ based on one input-output example of this function. But being a PRF, h is also secure against key-recovery.

In the full proof that follows, we use the games and adversaries specified in Figure 1. Adversaries A_1, A_2 represent, respectively, the first and second parts of the argument outlined above, while A_3 integrates the two.

Claim 3.5 Let $M_1, M_2 \in B^+$ with $M_1 \subseteq M_2$ and $1 + \|M_1\|_b \leq \|M_2\|_b$. Suppose $1 \leq l \leq \|M_1\|_b$. Then

$$\begin{aligned} \Pr \left[A_1^{\$}(M_1, M_2, l) \Rightarrow 1 \right] &= \Pr \left[G_1(M_1, M_2, l) \Rightarrow 1 \right] \\ \Pr \left[A_1^{h(K, \cdot)}(M_1, M_2, l) \Rightarrow 1 \right] &= \Pr \left[G_1(M_1, M_2, l - 1) \Rightarrow 1 \right]. \blacksquare \end{aligned}$$

Recall the notation means that in the first case A_1 gets as oracle $g \xleftarrow{\$} \text{Maps}(\{0, 1\}^b, \{0, 1\}^c)$ and in the second case $K \xleftarrow{\$} \{0, 1\}^c$.

Proof of Claim 3.5: $A_1^g(M_1, M_2, l)$ sets $a[l] = g(M_2[l])$. If g is chosen at random then this is equivalent to the $a[l] \xleftarrow{\$} \{0, 1\}^c$ assignment in $G_1(M_1, M_2, l)$. On the other hand if $g = h(K, \cdot)$ for a random K , then K plays the role of $a[l - 1]$ in $G_1(M_1, M_2, l - 1)$. \blacksquare

Claim 3.6 Let $M_1, M_2 \in B^+$ with $M_1 \subseteq M_2$ and $1 + \|M_1\|_b \leq \|M_2\|_b$. Then

$$\begin{aligned} \Pr \left[A_2^{\$}(M_1, M_2) \Rightarrow 1 \right] &= 2^{-c} \\ \Pr \left[A_2^{h(K, \cdot)}(M_1, M_2) \Rightarrow 1 \right] &\geq \Pr \left[G_1(M_1, M_2, m_1) \Rightarrow 1 \right]. \blacksquare \end{aligned}$$

Proof of Claim 3.6: Suppose g is chosen at random. Since $y \neq M_2[m_1 + 1]$, the quantity $g(y)$ is not defined until the last line of the code of A_2 , at which point $h(a[m_2], y)$ is fixed, and thus the probability that the two are equal is 2^{-c} due to the randomness of $g(y)$. Now suppose $g = h(K, \cdot)$ for a random K . Think of K as playing the role of $a[m_1]$ in $G_1(M_1, M_2, m_1)$. Then $a[m_2] = K$ in $A_2^g(M_1, M_2)$ exactly when $a[m_1] = a[m_2]$ in $G_1(M_1, M_2, m_1)$, meaning exactly when the latter game returns 1. But if $a[m_2] = K$ then certainly $h(a[m_2], y) = h(K, y)$, and the latter is $g(y)$, so $A_1^g(M_1, M_2)$ will return 1. (However, it could be that $h(a[m_2], y) = h(K, y)$ even if $a[m_2] \neq K$, which is why we have an inequality rather than an equality in the claim.) \blacksquare

Claim 3.7 Let $M_1, M_2 \in B^+$ with $M_1 \subseteq M_2$ and $1 + \|M_1\|_b \leq \|M_2\|_b$. Let $m_1 = \|M_1\|_b$. Then

$$\mathbf{Adv}_h^{\text{prf}}(A_3(M_1, M_2)) \geq \frac{1}{m_1 + 1} (\text{Coll}_{h^*}(M_1, M_2) - 2^{-c}). \blacksquare$$

Proof of Claim 3.7: From the description of A_3 , whether $g = \$$ or $g = h(K, \cdot)$,

$$\Pr \left[A_3^g(M_1, M_2) \Rightarrow 1 \right] = \frac{1}{m_1 + 1} \left(\Pr \left[A_2^g(M_1, M_2) \Rightarrow 1 \right] + \sum_{l=1}^{m_1} \Pr \left[A_1^g(M_1, M_2, l) \Rightarrow 1 \right] \right).$$

Now Claims 3.6 and 3.5 imply that $\Pr \left[A_3^{h(K, \cdot)}(M_1, M_2) \Rightarrow 1 \right]$ is

$$\begin{aligned} &\geq \frac{1}{m_1 + 1} \left(\Pr \left[G_1(M_1, M_2, m_1) \Rightarrow 1 \right] + \sum_{l=1}^{m_1} \Pr \left[G_1(M_1, M_2, l - 1) \Rightarrow 1 \right] \right) \\ &= \frac{1}{m_1 + 1} \cdot \sum_{l=0}^{m_1} \Pr \left[G_1(M_1, M_2, l) \Rightarrow 1 \right]. \end{aligned} \tag{6}$$

On the other hand, Claims 3.6 and 3.5 also imply that $\Pr \left[A_3^{\$}(M_1, M_2) \Rightarrow 1 \right]$ is

$$= \frac{1}{m_1 + 1} \left(2^{-c} + \sum_{l=1}^{m_1} \Pr \left[G_1(M_1, M_2, l) \Rightarrow 1 \right] \right). \tag{7}$$

Subtracting (7) from (6) and exploiting the cancellation of terms, we get

$$\mathbf{Adv}_h^{\text{prf}}(A_3(M_1, M_2)) \geq \frac{1}{m_1 + 1} (\Pr \left[G_1(M_1, M_2, 0) \Rightarrow 1 \right] - 2^{-c}).$$

Now examination of Game $G_1(M_1, M_2, 0)$ shows that that in this game, $a[m_1] = h^*(a[0], M_1)$, $a[m_2] = h^*(a[0], M_2)$, and $a[0]$ is selected at random. Since the game returns 1 iff $a[m_1] = a[m_2]$, the probability that it returns 1 is exactly $\text{Coll}_{h^*}(M_1, M_2)$. \blacksquare

THE CASE $M_1 \not\subseteq M_2$. For $M_1, M_2 \in B^+$ with $\|M_1\|_b \leq \|M_2\|_b$ and $M_1 \not\subseteq M_2$, we let $\text{LCP}(M_1, M_2)$ denote the length of the longest common blockwise prefix of M_1, M_2 , meaning the largest integer p

<p>Game $G_2(M_1, M_2, l_1, l_2)$ // $(l_1, l_2) \in I(M_1, M_2)$</p> <p>200 $m_1 \leftarrow \ M_1\ _b$; $m_2 \leftarrow \ M_2\ _b$</p> <p>210 $a_1[l_1] \stackrel{\\$}{\leftarrow} \{0, 1\}^c$</p> <p>220 If $(l_1, l_2) \in I_1(M_1, M_2)$ then $a_2[l_2] \leftarrow a_1[l_1]$ else $a_2[l_2] \stackrel{\\$}{\leftarrow} \{0, 1\}^c$</p> <p>230 For $i = l_1 + 1$ to m_1 do $a_1[i] \leftarrow h(a_1[i - 1], M_1[i])$</p> <p>240 For $i = l_2 + 1$ to m_2 do $a_2[i] \leftarrow h(a_2[i - 1], M_2[i])$</p> <p>250 If $a_1[m_1] = a_2[m_2]$ then return 1 else return 0</p> <p>Adversary $A_4^g(M_1, M_2, l_1, l_2)$ // $(l_1, l_2) \in I^*(M_1, M_2)$</p> <p>a00 $m_1 \leftarrow \ M_1\ _b$; $m_2 \leftarrow \ M_2\ _b$; $p \leftarrow \text{LCP}(M_1, M_2)$</p> <p>a10 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p + 1, p + 1)\} \cup I_2(M_1, M_2)$ then $a_1[l_1] \leftarrow g(M_1[l_1])$ else $a_1[l_1] \stackrel{\\$}{\leftarrow} \{0, 1\}^c$</p> <p>a20 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p + 1, p + 1)\} \cup I_3(M_1, M_2)$ then $a_2[l_2] \leftarrow g(M_2[l_2])$ else $a_2[l_2] \stackrel{\\$}{\leftarrow} \{0, 1\}^c$</p> <p>a30 For $i = l_1 + 1$ to m_1 do $a_1[i] \leftarrow h(a_1[i - 1], M_1[i])$</p> <p>a40 For $i = l_2 + 1$ to m_2 do $a_2[i] \leftarrow h(a_2[i - 1], M_2[i])$</p> <p>a50 If $a_1[m_1] = a_2[m_2]$ then return 1 else return 0</p>

Figure 2: Games and adversaries taking input distinct messages $M_1, M_2 \in B^+$ such that $M_1 \not\subseteq M_2$ and $\|M_1\|_b \leq \|M_2\|_b$.

such that $M_1[1] \dots M_1[p] = M_2[1] \dots M_2[p]$ but $M_1[p + 1] \neq M_2[p + 1]$. Letting $p = \text{LCP}(M_1, M_2)$, $m_1 = \|M_1\|_b$ and $m_2 = \|M_2\|_b$ we consider the following sequence of pairs:

$$\underbrace{(0, 0), \dots, (p, p)}_{I_1(M_1, M_2)}, (p + 1, p + 1), \underbrace{(p + 2, p + 1), \dots, (m_1, p + 1)}_{I_2(M_1, M_2)}, \underbrace{(m_1, p + 2), \dots, (m_1, m_2)}_{I_3(M_1, M_2)}. \quad (8)$$

For $j = 1, 2, 3$ we let $I_j(M_1, M_2)$ be the set of pairs indicated above. We then let

$$\begin{aligned} I(M_1, M_2) &= I_1(M_1, M_2) \cup \{(p + 1, p + 1)\} \cup I_2(M_1, M_2) \cup I_3(M_1, M_2) \\ I^*(M_1, M_2) &= I(M_1, M_2) - \{(0, 0)\} \\ I_1^*(M_1, M_2) &= I_1(M_1, M_2) - \{(0, 0)\}. \end{aligned}$$

We consider the pairs to be ordered as per (8), and for $(l_1, l_2) \in I^*(M_1, M_2)$ we let $\text{PD}(l_1, l_2)$ denote the predecessor of (l_1, l_2) , meaning the pair immediately preceding (l_1, l_2) in the sequence of (8). Note that $|I(M_1, M_2)| = m_1 + m_2 - p$. We consider the games and adversary of Figure 2.

Claim 3.8 Let $M_1, M_2 \in B^+$ with $M_1 \not\subseteq M_2$, and $\|M_1\|_b \leq \|M_2\|_b$. Suppose $(l_1, l_2) \in I^*(M_1, M_2)$. Let $(l'_1, l'_2) = \text{PD}(l_1, l_2)$. Then

$$\begin{aligned} \Pr \left[A_4^{\$}(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] &= \Pr \left[G_2(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] \\ \Pr \left[A_4^{h(K, \cdot)}(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] &= \Pr \left[G_2(M_1, M_2, l'_1, l'_2) \Rightarrow 1 \right]. \quad \blacksquare \end{aligned}$$

Proof of Claim 3.8: We begin by justifying the first equality, namely the one where $g \stackrel{\$}{\leftarrow} \text{Maps}(\{0, 1\}^b, \{0, 1\}^c)$. Let us compare the code of $G_2(M_1, M_2, l_1, l_2)$ and $A_4^g(M_1, M_2, l_1, l_2)$. Line a10 is equivalent to line 210 because g is random. Now consider line a20. If $(l_1, l_2) \in I_1^*(M_1, M_2)$

<p>Game $G_3(M_1, M_2, l_1, l_2)$ // $(l_1, l_2) \in I^*(M_1, M_2)$</p> <p>300 $m_1 \leftarrow \ M_1\ _b$; $m_2 \leftarrow \ M_2\ _b$; $p \leftarrow \text{LCP}(M_1, M_2)$; $K \xleftarrow{\\$} \{0, 1\}^c$</p> <p>310 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_2(M_1, M_2)$</p> <p style="padding-left: 2em;">then $a_1[l'_1] \leftarrow K$; $a_1[l_1] \leftarrow h(a_1[l'_1], M_1[l_1])$ else $a_1[l_1] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>320 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_3(M_1, M_2)$</p> <p style="padding-left: 2em;">then $a_2[l'_2] \leftarrow K$; $a_2[l_2] \leftarrow h(a_2[l'_2], M_2[l_2])$ else $a_2[l_2] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>330 For $i = l_1 + 1$ to m_1 do $a_1[i] \leftarrow h(a_1[i-1], M_1[i])$</p> <p>340 For $i = l_2 + 1$ to m_2 do $a_2[i] \leftarrow h(a_2[i-1], M_2[i])$</p> <p>350 If $a_1[m_1] = a_2[m_2]$ then return 1 else return 0</p> <p>Game $G_4(M_1, M_2, l_1, l_2)$ // $(l_1, l_2) \in I^*(M_1, M_2)$</p> <p>400 $m_1 \leftarrow \ M_1\ _b$; $m_2 \leftarrow \ M_2\ _b$; $p \leftarrow \text{LCP}(M_1, M_2)$; $K \xleftarrow{\\$} \{0, 1\}^c$</p> <p>410 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_2(M_1, M_2)$</p> <p style="padding-left: 2em;">then $a_1[l'_1] \leftarrow K$ else $a_1[l'_1] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>420 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_3(M_1, M_2)$</p> <p style="padding-left: 2em;">then $a_2[l'_2] \leftarrow K$ else $a_2[l'_2] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>430 For $i = l'_1 + 1$ to m_1 do $a_1[i] \leftarrow h(a_1[i-1], M_1[i])$</p> <p>440 For $i = l'_2 + 1$ to m_2 do $a_2[i] \leftarrow h(a_2[i-1], M_2[i])$</p> <p>450 If $a_1[m_1] = a_2[m_2]$ then return 1 else return 0</p> <p>Game $G_5(M_1, M_2, l_1, l_2)$ // $(l_1, l_2) \in I^*(M_1, M_2)$</p> <p>500 $m_1 \leftarrow \ M_1\ _b$; $m_2 \leftarrow \ M_2\ _b$; $p \leftarrow \text{LCP}(M_1, M_2)$</p> <p>510 $a_1[l'_1] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>520 If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\}$ then $a_2[l'_2] \leftarrow a_1[l'_1]$ else $a_2[l'_2] \xleftarrow{\\$} \{0, 1\}^c$</p> <p>530 For $i = l'_1 + 1$ to m_1 do $a_1[i] \leftarrow h(a_1[i-1], M_1[i])$</p> <p>540 For $i = l'_2 + 1$ to m_2 do $a_2[i] \leftarrow h(a_2[i-1], M_2[i])$</p> <p>550 If $a_1[m_1] = a_2[m_2]$ then return 1 else return 0</p>
--

Figure 3: Games equivalent to $G_2(M_1, M_2, l'_1, l'_2)$, where $M_1, M_2 \in B^+$ are such that $M_1 \not\subseteq M_2$ and $\|M_1\|_b \leq \|M_2\|_b$, and $(l_1, l_2) \in I^*(M_1, M_2)$.

then we have $a_2[l_2] = a_1[l_1]$ just as per line 220. This is true because, in this case, we have $M_1[l_1] = M_2[l_2]$ and hence $a_2[l_2] = g(M_2[l_2]) = g(M_1[l_1]) = a_1[l_1]$. If $(l_1, l_2) = (p+1, p+1)$ then setting $a_2[l_2] = g(M_2[l_2])$ is equivalent to choosing $a_2[l_2]$ at random because g is random and has not previously been invoked on $M_2[l_2]$. (We use here in a crucial way that the point $M_1[p+1]$ on which g has previously been invoked is different from $M_2[p+1]$.) If $(l_1, l_2) \in I_3(M_1, M_2)$ then setting $a_2[l_2] = g(M_2[l_2])$ is equivalent to choosing $a_2[l_2]$ at random because g is random and has not previously been invoked anywhere. (In this case, $a_1[l_1]$ is chosen at random and not via g .) Otherwise, $a_2[l_2]$ is chosen at random. We have shown that line a20 is equivalent to line 220. The rest of the code is the same in both cases.

Now we justify the second equality, namely the one where K is selected at random from $\{0, 1\}^c$

and $g = h(K, \cdot)$. We will consider some auxiliary games, shown in Figure 3. We claim that

$$\Pr \left[A_4^{h(K, \cdot)}(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] = \Pr [G_3(M_1, M_2, l_1, l_2) \Rightarrow 1] \quad (9)$$

$$= \Pr [G_4(M_1, M_2, l_1, l_2) \Rightarrow 1] \quad (10)$$

$$= \Pr [G_5(M_1, M_2, l_1, l_2) \Rightarrow 1] \quad (11)$$

$$= \Pr [G_2(M_1, M_2, l'_1, l'_2) \Rightarrow 1] \quad (12)$$

We now justify the above relations.

Since $a_1[l'_1]$ and $a_2[l'_2]$ in $G_3(M_1, M_2, l_1, l_2)$ both equal the key K , the output of this game is the same as that of $A_4^{h(K, \cdot)}(M_1, M_2, l_1, l_2)$, justifying (9).

Game $G_4(M_1, M_2, l_1, l_2)$ was obtained by dropping the application of $h(K, \cdot)$ in 310, 320 and beginning the loops of 330, 340 at l'_1, l'_2 rather than l_1, l_2 , respectively. To justify (10) we consider some cases. If $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_2(M_1, M_2)$ then $l'_1 = l_1 - 1$, so dropping the application of $h(K, \cdot)$ in 310 is compensated for by beginning the loop at 430 one step earlier. On the other hand if $(l_1, l_2) \in I_3(M_1, M_2)$ then $l'_1 = l_1 = m_1$, so starting the loop of 430 at l'_1 effectively makes no change, and we can also replace l_1 by l'_1 in the “else” clause as shown in 410. Similarly if $(l_1, l_2) \in I_1^*(M_1, M_2) \cup \{(p+1, p+1)\} \cup I_3(M_1, M_2)$ then $l'_2 = l_2 - 1$, so dropping the application of $h(K, \cdot)$ in 320 is compensated for by beginning the loop at 440 one step earlier. On the other hand if $(l_1, l_2) \in I_2(M_1, M_2)$ then $l'_2 = l_2 = p+1$, so starting the loop of 440 at l'_2 effectively makes no change, and we can also replace l_2 by l'_2 in the “else” clause as shown in 420.

Rather than picking K upfront and assigning it to $a_1[l'_1]$, Game $G_5(M_1, M_2, l_1, l_2)$ picks $a_1[l'_1]$ directly at random and then adjusts its choice of $a_2[l'_2]$ to ensure that it equals $a_1[l'_1]$ whenever they were both set to K . This justifies (11). Finally the test at 520 is equivalent to testing whether $(l'_1, l'_2) \in I_1(M_1, M_2)$, justifying (12). ■

We now define prf adversary $A_5^g(M_1, M_2)$ against h as follows. It picks $(l_1, l_2) \stackrel{\$}{\leftarrow} I^*(M_1, M_2)$ and returns $A_4^g(M_1, M_2, l_1, l_2)$.

Claim 3.9 Let $M_1, M_2 \in B^+$ with $M_1 \not\subseteq M_2$ and $\|M_1\|_b \leq \|M_2\|_b$. Let $m = \|M_1\|_b + \|M_2\|_b - \text{LCP}(M_1, M_2) - 1$. Then

$$\mathbf{Adv}_h^{\text{prf}}(A_5) \geq \frac{1}{m} \cdot (\text{Coll}_{h^*}(M_1, M_2) - 2^{-c}) . \quad \blacksquare$$

Proof of Claim 3.9: Note that $m = I^*(M_1, M_2)$. By Claim 3.8 we have the following, where the sums are over $(l_1, l_2) \in I^*(M_1, M_2)$ and we are letting $\text{PD}^j(l_1, l_2)$ be the j -th component of $\text{PD}(l_1, l_2)$ for $j = 1, 2$:

$$\begin{aligned} & \mathbf{Adv}_h^{\text{prf}}(A_5) \\ &= \frac{1}{m} \cdot \sum \Pr \left[A_4^{h(K, \cdot)}(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] - \frac{1}{m} \cdot \sum \Pr \left[A_4^{\$}(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] \\ &= \frac{1}{m} \cdot \sum \Pr \left[G_2(M_1, M_2, \text{PD}^1(l_1, l_2), \text{PD}^2(l_1, l_2)) \Rightarrow 1 \right] - \frac{1}{m} \cdot \sum \Pr \left[G_2(M_1, M_2, l_1, l_2) \Rightarrow 1 \right] \\ &= \frac{1}{m} \cdot (\Pr [G_2(M_1, M_2, 0, 0) \Rightarrow 1] - \Pr [G_2(M_1, M_2, m_1, m_2) \Rightarrow 1]) , \end{aligned}$$

where $m_1 = \|M_1\|_b$ and $m_2 = \|M_2\|_b$. Examination of Game $G_2(M_1, M_2, 0, 0)$ shows that that, in this game, $a_1[m_1] = h^*(a_1[0], M_1)$, $a_2[m_2] = h^*(a_2[0], M_2)$, and $a_1[0] = a_2[0]$ is selected at

random. Since the game returns 1 iff $a_1[m_1] = a_2[m_2]$, the probability that it returns 1 is exactly $\text{Coll}_{h^*}(M_1, M_2)$. On the other hand, the values $a_1[m_1]$ and $a_2[m_2]$ in $G_2(M_1, M_2, m_1, m_2)$ are chosen independently at random, and so the probability that they are equal, which is the probability this game returns 1, is 2^{-c} . ■

PUTTING IT TOGETHER. We consider the following adversary A_6 that takes input any distinct $M_1, M_2 \in B^+$ such that $\|M_1\|_b \leq \|M_2\|_b$:

Adversary $A_6^g(M_1, M_2)$
 If $M_1 \subseteq M_2$ then return $A_3^g(M_1, M_2)$
 Else return $A_5^g(M_1, M_2)$

Let $M_1^*, M_2^* \in B^+$ be distinct messages such that $\|M_1^*\|_b \leq \|M_2^*\|_b \leq n$ and

$$\mathbf{Adv}_h^{\text{prf}}(A_6(M_1, M_2)) \leq \mathbf{Adv}_h^{\text{prf}}(A_6(M_1^*, M_2^*))$$

for all distinct $M_1, M_2 \in B^+$ with $\|M_1\|_b \leq \|M_2\|_b \leq n$, where n is as in the Lemma statement. Now let A be the adversary that has M_1^*, M_2^* hardwired in its code and, given oracle g , returns $A_6^g(M_1^*, M_2^*)$. The adversary A has time-complexity as claimed in the Lemma statement. Now let A^* be any au-adversary against h^* such that the two messages output by A^* are at most n_1, n_2 blocks long, respectively, where $1 \leq n_1, n_2 \leq n$. We show that (2) holds. We assume wlog that the two messages M_1, M_2 output by A^* are always distinct, in B^+ , and satisfy $\|M_1\|_b \leq \|M_2\|_b$. Then we have

$$\mathbf{Adv}_{h^*}^{\text{au}}(A^*) = \sum_{M_1 \subseteq M_2} \text{Coll}_{h^*}(M_1, M_2) \cdot \Pr[M_1, M_2] + \sum_{M_1 \not\subseteq M_2} \text{Coll}_{h^*}(M_1, M_2) \cdot \Pr[M_1, M_2]$$

where $\Pr[M_1, M_2]$ denotes the probability that A^* outputs (M_1, M_2) . Now use Claims 3.7 and 3.9, and also the assumptions $\|M_1\|_b \leq n_1$, $\|M_2\|_b \leq n_2$ and $n_2 \geq 1$ from the lemma statement, to get

$$\begin{aligned} \mathbf{Adv}_{h^*}^{\text{au}}(A^*) &\leq \sum_{M_1 \subseteq M_2} \left[(n_1 + n_2 - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A_3(M_1, M_2)) + 2^{-c} \right] \cdot \Pr[M_1, M_2] \\ &\quad + \sum_{M_1 \not\subseteq M_2} \left[(n_1 + n_2 - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A_5(M_1, M_2)) + 2^{-c} \right] \cdot \Pr[M_1, M_2] \\ &= \sum_{M_1, M_2} \left[(n_1 + n_2 - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A_6(M_1, M_2)) + 2^{-c} \right] \cdot \Pr[M_1, M_2] \\ &\leq 2^{-c} + (n_1 + n_2 - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A), \end{aligned}$$

where in the last line we used the definition of prf-adversary A .

3.4 Proof of Lemma 3.2

Game G_0 of Figure 4 implements an oracle for $hF(K_{\text{out}} \| K_{\text{in}}, \cdot)$ with the keys chosen at random, while Game G_2 implements an oracle for a random function. So

$$\begin{aligned} \mathbf{Adv}_{hF}^{\text{prf}}(A_{hF}) &= \Pr[A_{hF}^{G_0} \Rightarrow 1] - \Pr[A_{hF}^{G_2} \Rightarrow 1] \\ &= (\Pr[A_{hF}^{G_0} \Rightarrow 1] - \Pr[A_{hF}^{G_1} \Rightarrow 1]) + (\Pr[A_{hF}^{G_1} \Rightarrow 1] - \Pr[A_{hF}^{G_2} \Rightarrow 1]) \end{aligned} \quad (13)$$

where in the last step we simply subtracted and then added back in the value $\Pr[A_{hF}^{G_1} \Rightarrow 1]$.

Game G_0 $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k$ $K_{\text{out}} \xleftarrow{\$} \{0, 1\}^c$ On query M: Reply $h(K_{\text{out}}, F(K_{\text{in}}, M))$	Game G_1 $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k$ $f \xleftarrow{\$} \text{Maps}(\{0, 1\}^b, \{0, 1\}^c)$ On query M: Reply $f(F(K_{\text{in}}, M))$	Game G_2 $g \xleftarrow{\$} \text{Maps}(D, \{0, 1\}^c)$ On query M: Reply $g(M)$
--	--	---

Figure 4: Games G_0, G_1, G_2 for the proof of Lemma 3.2.

It is easy to construct a prf-adversary A_h against h such that

$$\mathbf{Adv}_h^{\text{prf}}(A_h) = \Pr [A_{hF}^{G_0} \Rightarrow 1] - \Pr [A_{hF}^{G_1} \Rightarrow 1]. \quad (14)$$

(Namely, A_h , given an oracle for a function $f: B \rightarrow \{0, 1\}^c$, picks $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k$. It then runs A_{hF} , replying to oracle query M by $f(F(K_{\text{in}}, M))$, and returns whatever output A_{hF} returns. We omit the simple analysis that establishes (14).) The main part of the proof is to construct au-adversary A'_F against F such that

$$\Pr [A_{hF}^{G_1} \Rightarrow 1] - \Pr [A_{hF}^{G_2} \Rightarrow 1] \leq \binom{q}{2} \cdot \mathbf{Adv}_F^{\text{au}}(A'_F). \quad (15)$$

This adversary, however, will have time-complexity t (and output messages of at most n bits). A standard “coin-fixing” argument will then be used to derive from A'_F an au-adversary A_F that has time-complexity $O(T_F(n))$ (and also outputs messages of n bits) such that

$$\mathbf{Adv}_F^{\text{au}}(A'_F) \leq \mathbf{Adv}_F^{\text{au}}(A_F). \quad (16)$$

Combining (13), (14), (15) and (16) we get (3), completing the proof of the lemma.

Towards constructing A'_F , consider Games G_3, G_4, G_5 of Figure 5. (Game G_3 is defined by the code on the left of the Figure. Game G_4 is the same except that the boxed code-statement is omitted.) We will assume now that a prf-adversary never repeats an oracle query. This is wlog, and is used below without explicit mention.

Claim 3.10 Game G_4 is equivalent to Game G_2 while Game G_3 is equivalent to Game G_1 . ■

Proof of Claim 3.10: In Game G_4 , the “If” statement does nothing beyond setting the `bad` flag, and the reply to query M_s is always the random value Z_s . Thus, Game G_4 implements a random function just like Game G_2 . Game G_3 returns random values except that it also ensures that if $F(K_{\text{in}}, M_i) = F(K_{\text{in}}, M_j)$ then the answers to queries M_i, M_j are the same. Thus, it is equivalent to Game G_1 . ■

Now we have:

$$\Pr [A_{hF}^{G_1} \Rightarrow 1] - \Pr [A_{hF}^{G_2} \Rightarrow 1] = \Pr [A_{hF}^{G_3} \Rightarrow 1] - \Pr [A_{hF}^{G_4} \Rightarrow 1] \quad (17)$$

$$\leq \Pr [A_{hF}^{G_4} \text{ sets bad}]. \quad (18)$$

Above, (17) is by Claim 3.10. Since G_3, G_4 differ only in statements that follow the setting of `bad`, (18) follows from the Fundamental Lemma of Game Playing [10].

We define au-adversary A'_F against F , as follows: It runs $A_{hF}^{G_5}$, then picks at random i, j subject to $1 \leq i < j \leq q$, and finally outputs the messages M_i, M_j . In other words, it runs A_{hF} , replying to

<p>Games $G3, G4$</p> <p>$K_{\text{in}} \xleftarrow{\\$} \{0, 1\}^k ; s \leftarrow 0$ $Z_1, \dots, Z_q \xleftarrow{\\$} \{0, 1\}^c$</p> <p>On query M:</p> <p>$s \leftarrow s + 1 ; M_s \leftarrow M$ $Y_s \leftarrow F(K_{\text{in}}, M_s)$ If $(\exists r < s : Y_r = Y_s)$ then $\text{bad} \leftarrow \text{true} ; \boxed{Z_s \leftarrow Z_r}$ Reply Z_s</p>	<p>Game $G5$</p> <p>$s \leftarrow 0$ $Z_1, \dots, Z_q \xleftarrow{\\$} \{0, 1\}^c$</p> <p>On query M:</p> <p>$s \leftarrow s + 1 ; M_s \leftarrow M$ Reply Z_s</p>
--	--

Figure 5: Games $G3, G4$ are defined by the code on the left, where Game $G3$ includes the boxed statement while Game $G4$ does not.

the oracle queries of the latter with random values, and then outputs a random pair of messages that A_{hF} queries to its oracle. (In order for M_i, M_j to always be defined, we assume A_{hF} always makes *exactly* q oracle queries rather than at most q where by “always” we mean no matter how its oracle queries are replied to. This is wlog.) We claim that

$$\Pr [A_{hF}^{G4} \text{ sets bad}] \leq \binom{q}{2} \cdot \text{Adv}_F^{\text{au}}(A'_F). \quad (19)$$

Combining (18) and (19) yields (15). We now justify (19). Intuitively, it is true because i, j are chosen at random after the execution of A_{hF} is complete, so A_{hF} has no information about them. A rigorous proof however needs a bit more work. Consider the experiment defining the au-advantage of A'_F . (Namely, we run A_{hF}^{G5} , pick i, j at random subject to $1 \leq i < j \leq q$, and then pick $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k$.) In this experiment, consider the following events defined for $1 \leq \alpha < \beta \leq q$:

$$\begin{aligned} C_{\alpha, \beta} &: F(K_{\text{in}}, M_\alpha) = F(K_{\text{in}}, M_\beta) \\ C &: \bigvee_{1 \leq \alpha < \beta \leq q} C_{\alpha, \beta}. \end{aligned}$$

Notice that the events “ $C_{\alpha, \beta} \wedge (i, j) = (\alpha, \beta)$ ” ($1 \leq \alpha < \beta \leq q$) are disjoint. (Even though the events $C_{\alpha, \beta}$ for $1 \leq \alpha < \beta \leq q$ are not.) Thus:

$$\begin{aligned} \text{Adv}_F^{\text{au}}(A'_F) &= \Pr \left[\bigvee_{1 \leq \alpha < \beta \leq q} (C_{\alpha, \beta} \wedge (i, j) = (\alpha, \beta)) \right] \\ &= \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha, \beta} \wedge (i, j) = (\alpha, \beta)]. \end{aligned}$$

Since i, j are chosen at random after the execution of A_{hF}^{G5} is complete, the events “ $(i, j) = (\alpha, \beta)$ ” and $C_{\alpha, \beta}$ are independent. Thus the above equals

$$\begin{aligned} \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha, \beta}] \cdot \Pr [(i, j) = (\alpha, \beta)] &= \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha, \beta}] \cdot \frac{1}{\binom{q}{2}} \\ &= \frac{1}{\binom{q}{2}} \cdot \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha, \beta}] \\ &\geq \frac{1}{\binom{q}{2}} \cdot \Pr [C]. \end{aligned}$$

The proof of (19) is concluded by noting that $\Pr[C]$ equals $\Pr[G4 \text{ sets bad}]$.

Finally, note that

$$\mathbf{Adv}_F^{\text{au}}(A'_F) = \mathbf{E}_{M_1, M_2} [\text{Coll}_F(M_1, M_2)]$$

where the expectation is over $(M_1, M_2) \stackrel{s}{\leftarrow} A'_F$. Thus there must exist $M_1, M_2 \in B^+$ such that $\text{Coll}_F(M_1, M_2) \geq \mathbf{Adv}_F^{\text{au}}(A'_F)$. (And these messages are distinct because A_{hF} never repeats an oracle query.) Let A_F be the au-adversary that has M_1, M_2 hardwired as part of its code and, when run, simply outputs these messages and halts. Then (16) follows. Furthermore the time complexity of A_F is $O(mT_h)$. (Remember that by our convention the time-complexity is that of the overlying experiment, so includes the time to compute T_F on the messages that A_F outputs.)

3.5 Proof of Theorem 3.4

Define $F: \{0, 1\}^c \times B^+ \rightarrow \{0, 1\}^b$ by $F(K_{\text{in}}, M) = h^*(K_{\text{in}}, M) \parallel \text{fpad}$. Note $\text{GNMAC} = hF$. Consider the proof of Lemma 3.2 with A_{GNMAC} playing the role of A_{hF} . We assume as usual that A_{hF} does not repeat an oracle query and makes exactly q oracle queries. Let A_1 be the adversary A_h from that proof. Referring also to the games in that proof, we have

$$\mathbf{Adv}_{\text{GNMAC}}^{\text{prf}}(A_{hF}) \leq \mathbf{Adv}_h^{\text{prf}}(A_1) + \Pr[A_{hF}^{G4} \text{ sets bad}] . \quad (20)$$

Lemma 3.1 tells us that there is an adversary A_2 , with resource bounds as in the theorem statement, such that

$$\text{Coll}_F(M, M') \leq (\|M\|_b + \|M'\|_b - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A_2) + \frac{1}{2^c} \quad (21)$$

for all distinct $M, M' \in B^+$ such that $\|M\|_b, \|M'\|_b \leq m$. (We get this by applying the lemma to the adversary $A_{M, M'}^*$ that has M, M' hardwired in its code, outputs them, and halts.) Let M_1, \dots, M_q denote the queries made by A_{hF} in its execution with $G4$. Regarding these as random variables over the coins of this execution, and also taking the probability over the choice of K_{in} , we have

$$\Pr[A_{hF}^{G4} \text{ sets bad}] = \Pr[\exists i < j : F(K_{\text{in}}, M_i) = F(K_{\text{in}}, M_j)] \quad (22)$$

$$\begin{aligned} &\leq \sum_{i < j} \Pr[F(K_{\text{in}}, M_i) = F(K_{\text{in}}, M_j)] \\ &\leq \sum_{i < j} (n_i + n_j - 1) \cdot \mathbf{Adv}_h^{\text{prf}}(A_2) + \frac{1}{2^c} \end{aligned} \quad (23)$$

$$\begin{aligned} &= \mathbf{Adv}_h^{\text{prf}}(A_2) \cdot \sum_{i < j} (n_i + n_j - 1) + \binom{q}{2} \frac{1}{2^c} \\ &= \mathbf{Adv}_h^{\text{prf}}(A_2) \cdot (q-1) \cdot (n - q/2) + \binom{q}{2} \frac{1}{2^c} . \end{aligned}$$

Above, (22) is by the definition of game $G4$. We obtained (23) by applying (21). To do this, one can first write each term of the sum as itself a sum, over all distinct messages M, M' , of $\text{Coll}_F(M, M')$, weighted by the probability that $M_i = M$ and $M_j = M'$, the latter taken only over the execution of A_{hF} with $G4$. Note we use our convention that resource bounds assumed on an adversary (in this case, the block lengths of its queries) hold not only in the game defining security but across all possible answers to queries, and hence in particular in its execution with $G4$.

4 MAC-security of NMAC under weaker assumptions

Since any PRF is a secure MAC [6, 9], Theorem 3.3 implies that NMAC is a secure MAC if the compression function is a PRF. Here we show that NMAC is a secure MAC under a weaker-than-PRF assumption on the compression function —namely that it is a privacy-preserving MAC— coupled with the assumption that the hash function is cAU. This is of interest given the still numerous usages of HMAC as a MAC (rather than as a PRF). This result can be viewed as attempting to formalize the intuition given in [3, Remark 4.4].

4.1 Privacy preserving MACs

MAC FORGERY. Recall that the mac-advantage of mac-adversary A against a family of functions $f: Keys \times Dom \rightarrow Rng$ is

$$\mathbf{Adv}_f^{\text{mac}}(A) = \Pr \left[A^{f(K, \cdot), \text{VF}_f(K, \cdot, \cdot)} \text{ forges} : K \stackrel{\$}{\leftarrow} Keys \right] .$$

The verification oracle $\text{VF}_f(K, \cdot, \cdot)$ associated to f takes input M, T , returning 1 if $f(K, M) = T$ and 0 otherwise. Queries to the first oracle are called mac queries, and ones to the second are called verification queries. A is said to forge if it makes a verification query M, T the response to which is 1 but M was not previously a mac query. Note we allow multiple verification queries [9].

PRIVACY-PRESERVING MACs. The privacy notion for MACs that we use adapts the notion of left-or-right indistinguishability of encryption [5] to functions that are deterministic, and was first introduced by [8] who called it indistinguishability under distinct chosen-plaintext attack. An oracle query of an ind-adversary A against family $f: \{0, 1\}^k \times \{0, 1\}^l \rightarrow \{0, 1\}^L$ is a pair of l -bit strings. The reply is provided by one or the other of the following games:

Game Left	Game Right
$K \stackrel{\$}{\leftarrow} \{0, 1\}^c$	$K \stackrel{\$}{\leftarrow} \{0, 1\}^c$
On query (x_0, x_1) :	On query (x_0, x_1) :
Reply $f(K, x_0)$	Reply $f(K, x_1)$

Each game has an initialization step in which it picks a key; it then uses this key in computing replies to all the queries made by A . The ind-advantage of A is

$$\mathbf{Adv}_f^{\text{ind}}(A) = \Pr \left[A^{\text{Right}} \Rightarrow 1 \right] - \Pr \left[A^{\text{Left}} \Rightarrow 1 \right] .$$

However, unlike for encryption, the oracles here are deterministic. So A can easily win (meaning, obtain a high advantage), for example by making a pair of queries of the form $(x, z), (y, z)$, where x, y, z are distinct, and then returning 1 iff the replies returned are the same. (We expect that $h(K, x) \neq h(K, y)$ with high probability over K for functions h of interest, for example compression functions.) We fix this by simply outlawing such behavior. To be precise, let us say that A is *legitimate* if for any sequence $(x_0^1, x_1^1), \dots, (x_0^q, x_1^q)$ of oracle queries that it makes, x_0^1, \dots, x_0^q are all distinct l -bit strings, and also x_1^1, \dots, x_1^q are all distinct l -bit strings. (As a test, notice that the adversary who queried $(x, z), (y, z)$ was not legitimate.) It is to be understood henceforth that an ind-adversary means a legitimate one. When we say that f is privacy-preserving, we mean that the ind-advantage of any (legitimate) practical ind-adversary is low.

Privacy-preservation is not, by itself, a demanding property. For example, it is achieved by a constant family such as the one defined by $f(K, x) = 0^L$ for all K, x . We will however want the property for families that are also secure MACs.

PP-MAC $<$ PRF. We claim that a privacy-preserving MAC (PP-MAC) is strictly weaker than a PRF, in the sense that any PRF is (a secure MAC [6, 9] and) privacy-preserving, but not vice-versa. This means that when (below) we assume that a compression function h is a PP-MAC, we are indeed assuming less of it than that it is a PRF. Let us now provide some details about the claims made above. First, the following is the formal statement corresponding to the claim that any PRF is privacy-preserving:

Proposition 4.1 *Let $f: \{0, 1\}^k \times \{0, 1\}^l \rightarrow \{0, 1\}^L$ be a family of functions, and A_{ind} an ind-adversary against it that makes at most q oracle queries and has time-complexity at most t . Then there is a prf-adversary A_{prf} against f such that $\mathbf{Adv}_f^{\text{ind}}(A_{\text{ind}}) \leq 2 \cdot \mathbf{Adv}_f^{\text{prf}}(A_{\text{prf}})$. Furthermore, A_{prf} makes at most q oracle queries and has time-complexity at most t . ■*

The proof is a simple exercise and is omitted. Next we explain why a PP-MAC need not be a PRF. The reason (or one reason) is that if the output of a family of functions has some structure, for example always ending in a 0 bit, it would disqualify the family as a PRF but need not preclude its being a PP-MAC. To make this more precise, let $f: \{0, 1\}^k \times \{0, 1\}^l \rightarrow \{0, 1\}^L$ be a PP-MAC. Define $g: \{0, 1\}^k \times \{0, 1\}^l \rightarrow \{0, 1\}^{L+1}$ by $g(K, x) = f(K, x)\|0$ for all $K \in \{0, 1\}^k$ and $x \in \{0, 1\}^l$. Then g is also a PP-MAC, but is clearly not a PRF.

4.2 Results

The following implies that if h is a PP-MAC and F is cAU then their composition hF is a secure MAC.

Lemma 4.2 *Let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ and $F: \{0, 1\}^k \times D \rightarrow B$ be families of functions, and let $hF: \{0, 1\}^{c+k} \times D \rightarrow \{0, 1\}^c$ be defined by*

$$hF(K_{\text{out}}\|K_{\text{in}}, M) = h(K_{\text{out}}, F(K_{\text{in}}, M))$$

for all $K_{\text{out}} \in \{0, 1\}^c, K_{\text{in}} \in \{0, 1\}^k$ and $M \in D$. Let A_{hF} be a mac-adversary against hF that makes at most q_{mac} mac queries and at most q_{vf} verification queries, with the messages in each of these queries being of length at most n . Suppose A_{hF} has time-complexity at most t . Let $q = q_{\text{mac}} + q_{\text{vf}}$ and assume $2 \leq q < 2^b$. Then there exists a mac-adversary A_1 against h , an ind-adversary A_2 against h , and an au-adversary A_F against F such that

$$\mathbf{Adv}_{hF}^{\text{mac}}(A_{hF}) \leq \mathbf{Adv}_h^{\text{mac}}(A_1) + \mathbf{Adv}_h^{\text{ind}}(A_2) + \binom{q}{2} \cdot \mathbf{Adv}_F^{\text{au}}(A_F). \quad (24)$$

Furthermore, A_1 makes at most q_{mac} mac queries and at most q_{vf} verification queries and has time-complexity at most t ; A_2 makes at most q oracle queries and has time-complexity at most t ; and A_F outputs messages of length at most n , makes 2 oracle queries, and has time-complexity $O(T_F(n))$, where $T_F(n)$ is the time to compute F on an n -bit input. ■

The proof is in Section 4.3. As a corollary we have the following, which says that if h is a PP-MAC and h^* is cAU then GNMAC is a secure MAC.

Theorem 4.3 *Assume $b \geq c$ and let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions and let $\text{fpad} \in \{0, 1\}^{b-c}$ be a fixed padding string. Let GNMAC: $\{0, 1\}^{2c} \times B^+ \rightarrow \{0, 1\}^c$ be defined by*

$$\text{GNMAC}(K_{\text{out}}\|K_{\text{in}}, M) = h(K_{\text{out}}, h^*(K_{\text{in}}, M)\|\text{fpad})$$

for all $K_{\text{out}}, K_{\text{in}} \in \{0, 1\}^c$ and $M \in B^+$. Let A_{GNMAC} be a mac-adversary against GNMAC that makes at most q_{mac} mac queries and at most q_{vf} verification queries, with the messages in each

of these queries being of at most m blocks. Suppose A_{GNMAC} has time-complexity at most t . Let $q = q_{\text{mac}} + q_{\text{vf}}$ and assume $2 \leq q < 2^b$. Then there exists a mac-adversary A_1 against h , an ind-adversary A_2 against h , and an au-adversary A^* against h^* such that

$$\mathbf{Adv}_{\text{GNMAC}}^{\text{mac}}(A_{\text{GNMAC}}) \leq \mathbf{Adv}_h^{\text{mac}}(A_1) + \mathbf{Adv}_h^{\text{ind}}(A_2) + \binom{q}{2} \cdot \mathbf{Adv}_{h^*}^{\text{au}}(A^*). \quad (25)$$

Furthermore, A_1 makes at most q_{mac} mac queries and at most q_{vf} verification queries and has time-complexity at most t ; A_2 makes at most q oracle queries and has time-complexity at most t ; and A^* outputs messages of at most m blocks, makes 2 oracle queries, and has time-complexity $O(mT_h)$, where T_h is the time for one computation of h . \blacksquare

We remark that Lemma 4.2 can be extended to show that hF is not only a MAC but itself privacy-preserving. (This assumes h is privacy-preserving and F is cAU. We do not prove this here.) This implies that GNMAC is privacy-preserving as long as h is privacy-preserving and h^* is cAU. This is potentially useful because it may be possible to show that a PP-MAC is sufficient to ensure security in some applications where HMAC is currently assumed to be a PRF.

4.3 Proof of Lemma 4.2

A mac-adversary against h gets a mac oracle $h(K_{\text{out}}, \cdot)$ and corresponding verification oracle $\text{VF}_h(K_{\text{out}}, \cdot, \cdot)$. By itself picking key K_{in} and invoking these oracles, it can easily simulate the mac oracle $h(K_{\text{out}}, F(K_{\text{in}}, \cdot))$ and verification oracle $\text{VF}_h(K_{\text{out}}, F(K_{\text{in}}, \cdot), \cdot)$ required by a mac-adversary against hF . This leads to the following natural construction of A_1 :

Adversary $A_1^{h(K_{\text{out}}, \cdot), \text{VF}_h(K_{\text{out}}, \cdot, \cdot)}$

$K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k$; $i \leftarrow 0$

Run A_{hF} , replying to its oracle queries as follows:

On mac query M or verification query M, T :

$i \leftarrow i + 1$; $M_i \leftarrow M$; $y_i \leftarrow F(K_{\text{in}}, M)$

If mac query M then reply $h(K_{\text{out}}, y_i)$ to A_{hF}

If verification query M, T then reply $\text{VF}_h(K_{\text{out}}, y_i, T)$ to A_{hF}

Consider the experiment defining the mac-advantage of A_1 . Namely, choose $K_{\text{out}} \xleftarrow{\$} \{0, 1\}^c$ and run A_1 with oracles $h(K_{\text{out}}, \cdot)$ and $\text{VF}_h(K_{\text{out}}, \cdot, \cdot)$. Let Coll (for ‘‘collision’’) be the event that there exist j, l such that $y_j = y_l$ but $M_j \neq M_l$. Then

$$\begin{aligned} \mathbf{Adv}_h^{\text{mac}}(A_1) &= \Pr[A_1 \text{ forges}] \\ &\geq \Pr[A_{hF} \text{ forges} \wedge \overline{\text{Coll}}] \\ &\geq \Pr[A_{hF} \text{ forges}] - \Pr[\text{Coll}] \\ &= \mathbf{Adv}_{hF}^{\text{mac}}(A_{hF}) - \Pr[\text{Coll}]. \end{aligned} \quad (26)$$

The rest of the proof is devoted to upper bounding $\Pr[\text{Coll}]$. Consider the games $G1, G2$ of Figure 6, where we denote by $\langle i \rangle_b$ the representation of integer i as a b -bit string. (The assumption $q < 2^b$ made in the lemma statement means that we can always represent i this way in $G2$.) These games differ only in the manner in which tag T_i is computed. In $G1$, it is equal to $hF(K_{\text{out}} \| K_{\text{in}}, M_i)$. In $G2$, however, it is the result of applying $h(K_{\text{out}}, \cdot)$ to the current value of the counter i , and as such does not depend on K_{in} . Now note that

$$\begin{aligned} \Pr[\text{Coll}] &= \Pr[A_{hF}^{G1} \text{ sets bad}] \\ &= (\Pr[A_{hF}^{G1} \text{ sets bad}] - \Pr[A_{hF}^{G2} \text{ sets bad}]) + \Pr[A_{hF}^{G2} \text{ sets bad}]. \end{aligned} \quad (27)$$

<p>Game $G1$</p> <p style="text-align: center;">$K_{\text{out}} \xleftarrow{\\$} \{0, 1\}^c ; K_{\text{in}} \xleftarrow{\\$} \{0, 1\}^k ; i \leftarrow 0$</p> <p>On mac query M or verification query M, T:</p> <p style="padding-left: 20px;">$i \leftarrow i + 1 ; M_i \leftarrow M ; y_i \leftarrow F(K_{\text{in}}, M)$</p> <p style="padding-left: 20px;">If $\exists j < i : y_i = y_j$ and $M_i \neq M_j$ then $\text{bad} \leftarrow \text{true}$</p> <p style="padding-left: 20px;">If $\exists j < i : M_i = M_j$ then $T_i \leftarrow T_j$</p> <p style="padding-left: 20px;">Else $T_i \leftarrow h(K_{\text{out}}, y_i)$</p> <p style="padding-left: 20px;">If mac query M then reply T_i</p> <p style="padding-left: 20px;">If verification query M, T then</p> <p style="padding-left: 40px;">If $T_i = T$ then reply 1 else reply 0</p> <p>Game $G2$</p> <p style="text-align: center;">$K_{\text{out}} \xleftarrow{\\$} \{0, 1\}^c ; K_{\text{in}} \xleftarrow{\\$} \{0, 1\}^k ; i \leftarrow 0$</p> <p>On mac query M or verification query M, T:</p> <p style="padding-left: 20px;">$i \leftarrow i + 1 ; M_i \leftarrow M ; y_i \leftarrow F(K_{\text{in}}, M)$</p> <p style="padding-left: 20px;">If $\exists j < i : y_i = y_j$ and $M_i \neq M_j$ then $\text{bad} \leftarrow \text{true}$</p> <p style="padding-left: 20px;">If $\exists j < i : M_i = M_j$ then $T_i \leftarrow T_j$</p> <p style="padding-left: 20px;">Else $T_i \leftarrow h(K_{\text{out}}, \langle i \rangle_b)$</p> <p style="padding-left: 20px;">If mac query M then reply T_i</p> <p style="padding-left: 20px;">If verification query M, T then</p> <p style="padding-left: 40px;">If $T_i = T$ then reply 1 else reply 0</p>
--

Figure 6: Games for the proof of Lemma 4.2.

Now consider the adversaries A_2, A'_F described in Figure 7. We claim that

$$\Pr [A_{hF}^{G1} \text{ sets bad}] - \Pr [A_{hF}^{G2} \text{ sets bad}] \leq \mathbf{Adv}_h^{\text{ind}}(A_2) \quad (28)$$

$$\Pr [A_{hF}^{G2} \text{ sets bad}] \leq \binom{q}{2} \cdot \mathbf{Adv}_F^{\text{au}}(A'_F) . \quad (29)$$

Adversary A'_F , however, has time-complexity t (and outputs messages of at most n bits). A standard “coin-fixing” argument will be used to derive from A'_F an au-adversary A_F that has time-complexity $O(T_F(n))$ (and also outputs messages of n bits) such that

$$\mathbf{Adv}_F^{\text{au}}(A'_F) \leq \mathbf{Adv}_F^{\text{au}}(A_F) . \quad (30)$$

Combining (30), (29), (28), (27) and (26) yields (24) and completes the proof of Lemma 4.2. It remains to prove (28), (29) and (30). We begin with the first of these.

Recall that an ind-adversary against h is given an oracle that takes as input a pair of b -bit strings x_0, x_1 . We are denoting this oracle by g . Now it is easy to see that

$$\begin{aligned} \Pr [A_2^{\text{Right}} \Rightarrow 1] &= \Pr [A_{hF}^{G1} \text{ sets bad}] \\ \Pr [A_2^{\text{Left}} \Rightarrow 1] &= \Pr [A_{hF}^{G2} \text{ sets bad}] , \end{aligned}$$

which implies (28). However, there is one important thing we still need to verify, namely that A_2 is legitimate. So consider the sequence $(x_1, y_1), (x_2, y_2), \dots$ of oracle queries it makes. The left halves

Adversary $A_2^{g(\cdot, \cdot)}$
 $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^k ; i \leftarrow 0$
Run A_{hF} , replying to its oracle queries as follows:
On mac query M or verification query M, T :
 $i \leftarrow i + 1 ; M_i \leftarrow M ; y_i \leftarrow F(K_{\text{in}}, M)$
If $\exists j < i : y_i = y_j$ and $M_i \neq M_j$ then return 1
If $\exists j < i : M_i = M_j$ then $T_i \leftarrow T_j$
Else $T_i \leftarrow g(\langle i \rangle_b, y_i)$
If mac query M then reply T_i to A_{hF}
If verification query M, T then
If $T_i = T$ then reply 1 to A_{hF} else reply 0 to A_{hF}
Return 0

Adversary A'_F
 $K_{\text{out}} \xleftarrow{\$} \{0, 1\}^c ; i \leftarrow 0$
Run A_{hF} , replying to its oracle queries as follows:
On mac query M or verification query M, T :
 $i \leftarrow i + 1 ; M_i \leftarrow M$
If $\exists j < i : M_i = M_j$ then $T_i \leftarrow T_j$
Else $T_i \leftarrow h(K_{\text{out}}, \langle i \rangle_b)$
If mac query M then reply T_i to A_{hF}
If verification query M, T then
If $T_i = T$ then reply 1 to A_{hF} else reply 0 to A_{hF}
Pick i, j at random subject to $1 \leq i < j \leq q$
Return M_i, M_j

Figure 7: Ind-adversary A_2 against h , taking an oracle g that on input a pair of b -bit strings returns a c -bit string, and au-adversary A'_F against F , that outputs a pair of strings in D .

x_1, x_2, \dots are values of the counter i in different loop iterations and are thus strictly increasing (although not necessarily successive) and in particular different. On the other hand the right half values y_1, y_2, \dots are distinct because as soon as $y_i = y_j$ for some $j < i$, adversary A_2 halts (and returns 1), never making an oracle query whose right half is y_i .

Next we turn to A'_F . In order for the messages M_i, M_j it returns to always be defined, we assume wlog that A_{hF} always makes exactly, rather than at most, q_{mac} mac queries and exactly, rather than at most, q_{vf} verification queries. Intuitively, (29) is true because i, j are chosen at random after the execution of A_{hF} is complete, so A_{hF} has no information about them. This can be made rigorous just as in the proof of Lemma 3.2, and the details follow. Consider the experiment defining the au-advantage of A'_F . In this experiment, consider the following events defined for $1 \leq \alpha < \beta \leq q$:

$$\begin{aligned}
C_{\alpha, \beta} & : F(K_{\text{in}}, M_\alpha) = F(K_{\text{in}}, M_\beta) \wedge M_\alpha \neq M_\beta \\
C & : \bigvee_{1 \leq \alpha < \beta \leq q} C_{\alpha, \beta} .
\end{aligned}$$

Notice that the events “ $C_{\alpha,\beta} \wedge (i, j) = (\alpha, \beta)$ ” ($1 \leq \alpha < \beta \leq q$) are disjoint. (Even though the events $C_{\alpha,\beta}$ for $1 \leq \alpha < \beta \leq q$ are not.) Thus:

$$\begin{aligned} \text{Adv}_F^{\text{au}}(A'_F) &= \Pr \left[\bigvee_{1 \leq \alpha < \beta \leq q} (C_{\alpha,\beta} \wedge (i, j) = (\alpha, \beta)) \right] \\ &= \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha,\beta} \wedge (i, j) = (\alpha, \beta)]. \end{aligned}$$

Since i, j are chosen at random after the execution of A_{hF} is complete, the events “ $(i, j) = (\alpha, \beta)$ ” and $C_{\alpha,\beta}$ are independent. Thus the above equals

$$\begin{aligned} \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha,\beta}] \cdot \Pr [(i, j) = (\alpha, \beta)] &= \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha,\beta}] \cdot \frac{1}{\binom{q}{2}} \\ &= \frac{1}{\binom{q}{2}} \cdot \sum_{1 \leq \alpha < \beta \leq q} \Pr [C_{\alpha,\beta}] \\ &\geq \frac{1}{\binom{q}{2}} \cdot \Pr [C]. \end{aligned}$$

The proof of (29) is concluded by noting that $\Pr [C]$ equals $\Pr [G2 \text{ sets bad}]$.

Finally, au-adversary A_F of time-complexity $O(T_F(n))$ satisfying (30) can be obtained from A_F just as in the proof of Lemma 3.2.

5 Security of HMAC

In this section we show how our security results about NMAC lift to corresponding ones about HMAC. We begin by recalling the observation of [3] as to how this works for HMAC with two independent keys, and then discuss how to extend this to the single-keyed version of HMAC.

THE CONSTRUCTS. Let $h: \{0, 1\}^c \times \{0, 1\}^b \rightarrow \{0, 1\}^c$ as usual denote the compression function. Let pad be the padding function as described in Section 3, so that $s^* = s \parallel \text{pad}(|s|) \in B^+$ for any string s . Recall that the cryptographic hash function H associated to h is defined by $H(M) = h^*(\text{IV}, M^*)$, where IV is a c -bit initial vector that is fixed as part of the description of H and M is a string of any length up to some maximum length that is related to pad . (This maximum length is 2^{64} for current hash functions.) Then $\text{HMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = H(K_{\text{out}} \parallel H(K_{\text{in}} \parallel M))$, where $K_{\text{out}}, K_{\text{in}} \in \{0, 1\}^b$. If we write this out in terms of h^* alone we get

$$\text{HMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = h^*(\text{IV}, K_{\text{out}} \parallel h^*(\text{IV}, K_{\text{in}} \parallel M \parallel \text{pad}(b + |M|)) \parallel \text{pad}(b + c)).$$

As with NMAC, the details of the padding conventions are not important to the security of HMAC as a PRF, and we will consider the more general construct GHMAC: $\{0, 1\}^{2b} \times B^+ \rightarrow \{0, 1\}^c$ defined by

$$\text{GHMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = h^*(\text{IV}, K_{\text{out}} \parallel h^*(\text{IV}, K_{\text{in}} \parallel M) \parallel \text{fpad}) \quad (31)$$

for all $K_{\text{out}}, K_{\text{in}} \in \{0, 1\}^b$ and all $M \in B^+$. Here $\text{IV} \in \{0, 1\}^c$ and $\text{fpad} \in \{0, 1\}^{b-c}$ are fixed strings. HMAC is a special case, via $\text{HMAC}(K_{\text{out}} \parallel K_{\text{in}}, M) = \text{GHMAC}(M \parallel \text{pad}(b + |M|))$ with $\text{fpad} = \text{pad}(b + c)$, and thus security properties of GHMAC (as a PRF or MAC) are inherited by HMAC, allowing us to focus on the former.

THE DUAL FAMILY. To state the results, it is useful to define $\bar{h}: \{0, 1\}^b \times \{0, 1\}^c \rightarrow \{0, 1\}^c$, the dual of family h , by $\bar{h}(x, y) = h(y, x)$. The assumption that h is a PRF when keyed by its data input is formally captured by the assumption that \bar{h} is a PRF.

<p>Game G_0 On query M: $Z \xleftarrow{\\$} \{0, 1\}^c$ Reply Z</p>	<p>Game G_1 $K'_{\text{out}}, K'_{\text{in}} \xleftarrow{\\$} \{0, 1\}^c$ On query M: Reply $\text{GNMAC}(K'_{\text{out}} \ K'_{\text{in}}, M)$</p>
<p>Game G_2 $K_{\text{in}} \xleftarrow{\\$} \{0, 1\}^b$ $K'_{\text{out}} \xleftarrow{\\$} \{0, 1\}^c$; $K'_{\text{in}} \leftarrow h(\text{IV}, K_{\text{in}})$ On query M: Reply $\text{GNMAC}(K'_{\text{out}} \ K'_{\text{in}}, M)$</p>	<p>Game G_3 $K_{\text{in}}, K_{\text{out}} \xleftarrow{\\$} \{0, 1\}^b$ $K'_{\text{out}} \leftarrow h(\text{IV}, K_{\text{out}})$; $K'_{\text{in}} \leftarrow h(\text{IV}, K_{\text{in}})$ On query M: Reply $\text{GNMAC}(K'_{\text{out}} \ K'_{\text{in}}, M)$</p>

Figure 8: Games G_0, G_1, G_2, G_3 for the proof of Lemma 5.1.

5.1 Security of GHMAC

Let $K'_{\text{out}} = h(\text{IV}, K_{\text{out}})$ and $K'_{\text{in}} = h(\text{IV}, K_{\text{in}})$. The observation of [3] is that

$$\begin{aligned} \text{GHMAC}(K_{\text{out}} \| K_{\text{in}}, M) &= h(K'_{\text{out}}, h^*(K'_{\text{in}}, M) \| \text{fpad}) \\ &= \text{GNMAC}(K'_{\text{out}} \| K'_{\text{in}}, M). \end{aligned} \quad (32)$$

This effectively reduces the security of GHMAC to GNMAC. Namely, if \bar{h} is a PRF and $K_{\text{out}}, K_{\text{in}}$ are chosen at random, then $K'_{\text{out}}, K'_{\text{in}}$ will be computationally close to random. Now (32) implies that if GNMAC is a PRF then so is GHMAC. The formal statement follows.

Lemma 5.1 *Assume $b \geq c$ and let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions. Let $\text{fpad} \in \{0, 1\}^{b-c}$ be a fixed padding string and $\text{IV} \in \{0, 1\}^c$ a fixed initial vector. Let $\text{GHMAC}: \{0, 1\}^{2b} \times B^+ \rightarrow \{0, 1\}^c$ be defined by (31) above. Let A be a prf-adversary against GHMAC that has time-complexity at most t . Then there exists a prf-adversary $A_{\bar{h}}$ against \bar{h} such that*

$$\text{Adv}_{\text{GHMAC}}^{\text{prf}}(A) \leq 2 \cdot \text{Adv}_{\bar{h}}^{\text{prf}}(A_{\bar{h}}) + \text{Adv}_{\text{GNMAC}}^{\text{prf}}(A).$$

Furthermore, $A_{\bar{h}}$ makes only 1 oracle query, this being IV , and has time-complexity at most t . \blacksquare

Proof of Lemma 5.1: Assume wlog that A never repeats an oracle query. Consider the games in Figure 8, and let $p_i = \Pr[A^{G_i} \Rightarrow 1]$ for $0 \leq i \leq 3$. Then

$$\text{Adv}_{\text{GHMAC}}^{\text{prf}}(A) = p_3 - p_0 = (p_3 - p_1) + (p_1 - p_0).$$

Clearly $p_1 - p_0 = \text{Adv}_{\text{GNMAC}}^{\text{prf}}(A)$. To complete the proof we construct $A_{\bar{h}}$ so that

$$\text{Adv}_{\bar{h}}^{\text{prf}}(A_{\bar{h}}) = \frac{p_3 + p_2}{2} - \frac{p_2 + p_1}{2} = \frac{p_3 - p_1}{2}. \quad (33)$$

Given an oracle $g: \{0, 1\}^c \rightarrow \{0, 1\}^c$, adversary $A_{\bar{h}}$ picks a bit c at random. Then it picks keys via

If $c = 1$ then $K'_{\text{out}} \leftarrow g(\text{IV})$; $K_{\text{in}} \xleftarrow{\$} \{0, 1\}^b$; $K'_{\text{in}} \leftarrow h(\text{IV}, K_{\text{in}})$

Else $K'_{\text{out}} \leftarrow \{0, 1\}^c$; $K'_{\text{in}} \leftarrow g(\text{IV})$.

Finally it runs A with oracle $\text{GNMAC}(K'_{\text{out}} \| K'_{\text{in}}, \cdot)$, and returns whatever A returns. \blacksquare

Combining this with Theorem 3.3 or Theorem 3.4 yields the result that GHMAC is a PRF assuming h, \bar{h} are both PRFs. Note that the PRF assumption on \bar{h} is mild because $A_{\bar{h}}$ makes only one oracle query.

5.2 Single-keyed HMAC

HMAC, GHMAC as described and analyzed above use two keys that are assumed to be chosen independently at random. However, HMAC is in fact usually implemented with these keys derived from a single b -bit key. Here we provide the first security proofs for the actually-implemented single-key version of HMAC.

Specifically, let $\text{opad}, \text{ipad} \in \{0, 1\}^b$ be distinct, fixed and known constants. (Their particular values can be found in [3] and are not important here.) Then the single-key version of HMAC is defined by

$$\text{HMAC-1}(K, M) = \text{HMAC}(K \oplus \text{opad} \| K \oplus \text{ipad}, M).$$

As before, we look at this as a special case of a more general construct, namely GHMAC-1: $\{0, 1\}^b \times B^+ \rightarrow \{0, 1\}^c$, defined by

$$\text{GHMAC-1}(K, M) = \text{GHMAC}(K \oplus \text{opad} \| K \oplus \text{ipad}, M) \quad (34)$$

for all $K \in \{0, 1\}^b$ and all $M \in B^+$. We now focus on GHMAC-1. We will show that GHMAC-1 inherits the security of GNMAC as long as \bar{h} is a PRF against an appropriate class of related key attacks. In such an attack, the adversary can obtain input-output examples of h under keys related to the target key. Let us recall the formal definitions following [7].

A related-key attack on a family of functions $\bar{h}: \{0, 1\}^b \times \{0, 1\}^c \rightarrow \{0, 1\}^c$ is parameterized by a set $\Phi \subseteq \text{Maps}(\{0, 1\}^b, \{0, 1\}^b)$ of *key-derivation* functions. We define the function $\text{RK}: \Phi \times \{0, 1\}^b \rightarrow \{0, 1\}^b$ by $\text{RK}(\phi, K) = \phi(K)$ for all $\phi \in \Phi$ and $K \in \{0, 1\}^b$. A rka-adversary $A_{\bar{h}}$ may make an oracle query of the form ϕ, x where $\phi \in \Phi$ and $x \in \{0, 1\}^c$. Its rka-advantage is defined by

$$\text{Adv}_{\bar{h}, \Phi}^{\text{rka}}(A_{\bar{h}}) = \Pr \left[A_{\bar{h}}^{\bar{h}(\text{RK}(\cdot, K), \cdot)} \Rightarrow 1 \right] - \Pr \left[A_{\bar{h}}^{G(\text{RK}(\cdot, K), \cdot)} \Rightarrow 1 \right].$$

In the first case, K is chosen at random from $\{0, 1\}^b$ and the reply to query ϕ, x of $A_{\bar{h}}$ is $\bar{h}(\phi(K), x)$. In the second case, $G \stackrel{\$}{\leftarrow} \text{Maps}(\{0, 1\}^b \times \{0, 1\}^c, \{0, 1\}^c)$ and $K \stackrel{\$}{\leftarrow} \{0, 1\}^b$, and the reply to query ϕ, x of $A_{\bar{h}}$ is $G(\phi(K), x)$. For any string $s \in \{0, 1\}^b$ let $\Delta_s: \{0, 1\}^b \rightarrow \{0, 1\}^b$ be defined by $\Delta_s(K) = K \oplus s$.

Lemma 5.2 *Assume $b \geq c$ and let $B = \{0, 1\}^b$. Let $h: \{0, 1\}^c \times B \rightarrow \{0, 1\}^c$ be a family of functions. Let $\text{fpad} \in \{0, 1\}^{b-c}$ be a fixed padding string, $\text{IV} \in \{0, 1\}^c$ a fixed initial vector, and $\text{opad}, \text{ipad} \in \{0, 1\}^b$ fixed, distinct strings. Let GHMAC-1: $\{0, 1\}^b \times B^+ \rightarrow \{0, 1\}^c$ be defined by (34) above. Let $\Phi = \{\Delta_{\text{opad}}, \Delta_{\text{ipad}}\}$. Let A be a prf-adversary against GHMAC-1 that has time-complexity at most t . Then there exists a rka-adversary $A_{\bar{h}}$ against \bar{h} such that*

$$\text{Adv}_{\text{GHMAC-1}}^{\text{prf}}(A) \leq \text{Adv}_{\bar{h}, \Phi}^{\text{rka}}(A_{\bar{h}}) + \text{Adv}_{\text{GNMAC}}^{\text{prf}}(A).$$

Furthermore, $A_{\bar{h}}$ makes 2 oracle queries, these being $\Delta_{\text{opad}}, \text{IV}$ and $\Delta_{\text{ipad}}, \text{IV}$, and has time-complexity at most t . ■

Proof of Lemma 5.2: Adversary $A_{\bar{h}}$ queries its oracle with $\Delta_{\text{opad}}, \text{IV}$ and lets K'_{out} denote the value returned. It also queries its oracle with $\Delta_{\text{ipad}}, \text{IV}$ and lets K'_{in} denote the value returned. It then runs A , answering the latter's oracle queries via $\text{GNMAC}(K'_{\text{out}} \| K'_{\text{in}}, \cdot)$, and returns whatever A returns. ■

Combining this with Theorem 3.3 or Theorem 3.4 yields the result that GHMAC-1 is a PRF assuming h is a PRF and \bar{h} is a PRF under Φ -restricted related-key attacks, where Φ is as in Lemma 5.2. We remark that Φ is a small set of simple functions, which is important because it is shown in [7] that if Φ is too rich then no family can be a PRF under Φ -restricted related-key attacks. Furthermore, the assumption on \bar{h} is rendered milder by the fact that $A_{\bar{h}}$ makes only two oracle queries, in both of which the message is the same, namely is the initial vector.

5.3 Lifting the results of Section 4

The procedure above to lift the NMAC results of Section 3 to HMAC applies also to lift the results of Section 4 to HMAC. Specifically, if h is a PP-MAC, h^* is AU and \bar{h} is a PRF then GHMAC is a (privacy-preserving) MAC. Also if h is a PP-MAC, h^* is AU and \bar{h} is a PRF under Φ -restricted related-key attacks, with Φ as in Lemma 5.2, then GHMAC-1 is a (privacy-preserving) MAC. Note that the assumption on \bar{h} continues to be that it is a PRF or PRF against Φ -restricted related-key attacks. (Namely, this has not been reduced to its being a PP-MAC.) This assumption is however mild in this context since (as indicated by Lemmas 5.2 and 5.1) it need only hold with respect to adversaries that make very few queries and these of a very specific type.

5.4 Remarks

Let $h: \{0, 1\}^{128} \times \{0, 1\}^{512} \rightarrow \{0, 1\}^{128}$ denote the compression function of MD5 [28]. An attack by den Boer and Bosselaers [15] finds values x_0, x_1, K such that $h(x_0, K) = h(x_1, K)$ but $x_0 \neq x_1$. In a personal communication, Rijmen has said that it seems possible to extend this to an attack that finds such x_0, x_1 even when K is unknown. If so, this might translate into the following attack showing h is not a PRF when keyed by its data input. (That is, \bar{h} is not a PRF.) Given an oracle $g: \{0, 1\}^{128} \rightarrow \{0, 1\}^{128}$, the attacker would find x_0, x_1 and obtain $y_0 = g(x_0)$ and $y_1 = g(x_1)$ from the oracle. It would return 1 if $y_0 = y_1$ and 0 otherwise. How does this impact the above, where we are assuming \bar{h} is a PRF? Interestingly, the actual PRF assumptions we need on \bar{h} are so weak that even such an attack does not break them. In Lemma 5.1, we need \bar{h} to be a PRF only against adversaries that make just one oracle query. (Because $A_{\bar{h}}$ makes only one query.) But the attack above makes two queries. On the other hand, in Lemma 5.2, we need \bar{h} to be a related-key PRF only against adversaries that make two related-key queries in both of which the 128-bit message for \bar{h} is the same, this value being the initial vector used by the hash function. Furthermore, the related key functions must be $\Delta_{\text{opad}}, \Delta_{\text{ipad}}$. The above-mentioned attack, however, uses two different messages x_0, x_1 and calls the oracle under the original key rather than the related keys. In summary, the attack does not violate the assumptions made in either of the lemmas.

Acknowledgments

Thanks to Dan Brown, Ran Canetti, Hugo Krawczyk, Mridul Nandi, Vincent Rijmen, Phillip Rogaway, Victor Shoup, Paul Van Oorschot and the Crypto 2006 PC for comments and references.

References

- [1] American National Standards Institution. ANSI X9.71, Keyed hash message authentication code, 2000. (Cited on page 3.)

- [2] M. Bellare. New Proofs for NMAC and HMAC: Security without Collision-Resistance. Preliminary version of this paper. *Advances in Cryptology – CRYPTO '06*, Lecture Notes in Computer Science Vol. , ed., Springer-Verlag, 2006. (Cited on page 5.)
- [3] M. Bellare, R. Canetti and H. Krawczyk. Keying hash functions for message authentication. *Advances in Cryptology – CRYPTO '96*, Lecture Notes in Computer Science Vol. 1109, N. Kobitz ed., Springer-Verlag, 1996. (Cited on pages 1, 3, 4, 5, 6, 19, 24, 25, 26, 29 and 30.)
- [4] M. Bellare, R. Canetti and H. Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. <http://www-cse.ucsd.edu/users/mihir>. (Preliminary version in *Proceedings of the 37th Symposium on Foundations of Computer Science*, IEEE, 1996.) (Cited on pages 4, 8, 30 and 31.)
- [5] M. Bellare, A. Desai, E. Jorjipii and P. Rogaway. A concrete security treatment of symmetric encryption. *Proceedings of the 38th Symposium on Foundations of Computer Science*, IEEE, 1997. (Cited on page 19.)
- [6] M. Bellare, J. Kilian and P. Rogaway. The security of the cipher block chaining message authentication code. *Journal of Computer and System Sciences*, Vol. 61, No. 3, Dec 2000, pp. 362–399. (Cited on pages 4, 6, 19 and 20.)
- [7] M. Bellare and T. Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. *Advances in Cryptology – EUROCRYPT '03*, Lecture Notes in Computer Science Vol. 2656, E. Biham ed., Springer-Verlag, 2003. (Cited on pages 5, 26 and 27.)
- [8] M. Bellare, C. Namprempre and T. Kohno. Authenticated Encryption in SSH: Provably Fixing the SSH Binary Packet Protocol. *ACM Transactions on Information and System Security (TISSEC)*, Vol. 7, Iss. 2, May 2004, pp. 206–241. (Cited on pages 4 and 19.)
- [9] M. Bellare, O. Goldreich and A. Mityagin. The power of verification queries in message authentication and authenticated encryption. *Cryptology ePrint Archive: Report 2004/309*, 2004. (Cited on pages 4, 6, 19 and 20.)
- [10] M. Bellare and P. Rogaway. The game-playing technique and its application to triple encryption. *Cryptology ePrint Archive: Report 2004/331*, 2004. (Cited on pages 8 and 16.)
- [11] J. BLACK, S. HALEVI, H. KRAWCZYK, T. KROVETZ AND P. ROGAWAY. UMAC: Fast and Secure Message Authentication. *Advances in Cryptology – CRYPTO '99*, Lecture Notes in Computer Science Vol. 1666, M. Wiener ed., Springer-Verlag, 1999. (Cited on page 8.)
- [12] J. Black and P. Rogaway. CBC MACs for arbitrary-length messages: The three-key constructions. *Advances in Cryptology – CRYPTO '00*, Lecture Notes in Computer Science Vol. 1880, M. Bellare ed., Springer-Verlag, 2000. (Cited on page 8.)
- [13] L. CARTER AND M. WEGMAN. Universal classes of hash functions. *Journal of Computer and System Sciences*, Vol. 18, No. 2, 1979, pp. 143–154. (Cited on page 4.)
- [14] I. Damgård. A design principle for hash functions. *Advances in Cryptology – CRYPTO '89*, Lecture Notes in Computer Science Vol. 435, G. Brassard ed., Springer-Verlag, 1989. (Cited on page 6.)
- [15] B. den Boer and A. Bosselaers. Collisions for the compression function of MD5. *Advances in Cryptology – EUROCRYPT '93*, Lecture Notes in Computer Science Vol. 765, T. Hellesteth ed., Springer-Verlag, 1993. (Cited on page 27.)
- [16] T. Dierks and C. Allen. The TLS protocol. Internet RFC 2246, 1999. (Cited on page 3.)
- [17] H. Dobbertin, A. Bosselaers and B. Preneel. RIPEMD-160: A strengthened version of RIPEMD. *Fast Software Encryption '96*, Lecture Notes in Computer Science Vol. 1039, D. Gollmann ed., Springer-Verlag, 1996. (Cited on page 3.)
- [18] Y. Dodis, R. Gennaro, J. Håstad, H. Krawczyk, and T. Rabin. Randomness extraction and key derivation using the CBC, Cascade, and HMAC modes. *Advances in Cryptology – CRYPTO '04*, Lecture Notes in Computer Science Vol. 3152, M. Franklin ed., Springer-Verlag, 2004. (Cited on page 4.)
- [19] O. GOLDBREICH, S. GOLDWASSER AND S. MICALI. How to construct random functions. *Journal of the ACM*, Vol. 33, No. 4, 1986, pp. 210–217. (Cited on page 3.)
- [20] D. Harkins and D. Carrel. The Internet Key Exchange (IKE). Internet RFC 2409, 1998. (Cited on page 3.)
- [21] S. Hirose. A note on the strength of weak collision resistance. *IEICE Transactions on Fundamentals*, Vol. E87-A, No. 5, May 2004, pp. 1092–1097. (Cited on pages 3 and 29.)

- [22] H. Krawczyk, M. Bellare and R. Canetti. HMAC: Keyed-hashing for message authentication. Internet RFC 2104, 1997. (Cited on page 3.)
- [23] R. Merkle. One-way hash functions and DES. *Advances in Cryptology – CRYPTO '89*, Lecture Notes in Computer Science Vol. 435, G. Brassard ed., Springer-Verlag, 1989. (Based on an unpublished paper from 1979 and the author's Ph. D thesis, Stanford, 1979). (Cited on page 6.)
- [24] D. M'Raihi, M. Bellare, F. Hoornaert, D. Naccache, O. Ranen. HOTP: An HMAC-based one time password algorithm. Internet RFC 4226, December 2005. (Cited on page 3.)
- [25] National Institute of Standards and Technology. The keyed-hash message authentication code (HMAC). FIPS PUB 198, March 2002. (Cited on page 3.)
- [26] National Institute of Standards and Technology. Secure hash standard. FIPS PUB 180-2, August 2000. (Cited on page 3.)
- [27] B. Preneel and P. van Oorschot. On the security of iterated message authentication codes. *IEEE Transactions on Information Theory*, Vol. 45, No. 1, January 1999, pp. 188–199. (Preliminary version, entitled “MD-x MAC and building fast MACs from hash functions,” in CRYPTO 95.) (Cited on pages 3, 8 and 9.)
- [28] R. RIVEST. The MD5 message-digest algorithm. Internet RFC 1321, April 1992. (Cited on pages 3 and 27.)
- [29] V. Shoup. Sequences of games: A tool for taming complexity in security proofs. *Cryptology ePrint Archive: Report 2004/332*, 2004. (Cited on page 8.)
- [30] D. Stinson. Universal hashing and authentication codes. *Designs, Codes and Cryptography*, Vol. 4, 1994, 369–380. (Cited on page 4.)
- [31] X. Wang, Y. L. Yin and H. Yu. Finding collisions in the full SHA-1. *Advances in Cryptology – CRYPTO '05*, Lecture Notes in Computer Science Vol. 3621, V. Shoup ed., Springer-Verlag, 2005. (Cited on pages 3 and 29.)
- [32] X. Wang and H. Yu. How to break MD5 and other hash functions. *Advances in Cryptology – EUROCRYPT '05*, Lecture Notes in Computer Science Vol. 3494, R. Cramer ed., Springer-Verlag, 2005. (Cited on pages 3 and 29.)
- [33] M. WEGMAN AND L. CARTER. New hash functions and their use in authentication and set equality. *Journal of Computer and System Sciences*, Vol. 22, No. 3, 1981, pp. 265–279. (Cited on page 4.)

A Attacking weak collision-resistance

Recall that H represents the cryptographic hash function (eg. MD5, SHA-1) while H^* is the extended hash function, which is the hash function with the initial vector made explicit as an (additional) first input. Let us use the term *general collision-finding attack* to refer to an attack that finds collisions in $H^*(IV, \cdot)$ for an arbitrary but given IV . As we discussed in Section 1, it was noted in [3, 21] that any general collision-finding attack can be used to compromise the weak collision-resistance (WCR) of H . (And since the known collision-finding attacks on MD5 and SHA-1 [32, 31] do extend to general ones, the WCR of these functions is no more than their CR.) Here we recall the argument that shows this. It is a simple extension attack, and works as follows.

To compromise WCR of H , an attacker given an oracle for $H^*(K, \cdot)$ under a hidden key K must output distinct M_1, M_2 such that $H^*(K, M_1) = H^*(K, M_2)$. Our attacker picks some string x and calls its oracle to obtain $IV = H^*(K, x)$. Then it runs the given general collision-finding attack on input IV to obtain a collision X_1, X_2 for $H^*(IV, \cdot)$. (That is, X_1, X_2 are distinct strings such that $H^*(IV, X_1) = H^*(IV, X_2)$.) Now let $M_1 = x \parallel \text{pad}(|x|) \parallel X_1 \parallel \text{pad}(|X_1|)$ and $M_2 = x \parallel \text{pad}(|x|) \parallel X_2 \parallel \text{pad}(|X_2|)$. (Here $\text{pad}(n)$ is a padding string that when appended to a string of length n results in a string whose length is a positive multiple of b bits where b is the block-length of the underlying compression function. The function pad is part of the description of the cryptographic hash function.) Then it follows that $H^*(K, M_1) = H^*(K, M_2)$.

We clarify that these attacks on the WCR of the hash function do not break HMAC. What these attacks show is that the WCR assumption made for the security proof of [3] is not true for MD5 and

SHA-1. This means we lose the proof-based guarantee of [3], but it does not imply any weakness in the construct. Our results show that WCR of the iterated hash function is not necessary for the security of HMAC: pseudorandomness of the compression function suffices. This helps explain why no attacks have emerged on HMAC even when it is implemented with hash functions that are not WCR.

B The reduction-from-pf-PRF proof

We sketch how one can obtain the result that h a PRF implies h^* is cAU by using the result of [4] that says that h a PRF implies h^* is a pf-PRF (a PRF as long as no query of the adversary is a prefix of another query). We then compare this with the direct proof given in Section 3.3 .

THE RESULT OF [4]. A pf-adversary is said to be prefix-free if no query it makes is a prefix of another. The result of [4] is that if D is a prefix-free pf-adversary against h^* that makes at most q queries, each of at most m blocks, then there is a pf-adversary A against h such that

$$\mathbf{Adv}_{h^*}^{\text{prf}}(D) \leq qm \cdot \mathbf{Adv}_h^{\text{prf}}(A) \quad (35)$$

and A has about the same time-complexity as D . (We remark that there is a typo in the statement of Theorem 3.1 of the proceedings version of [4] in this regard: the factor q is missing from the bound. This is however corrected in the on-line version of the paper.)

THE REDUCTION-FROM-PF-PRF. The result obtained via the reduction-from-pf-PRF proof will be slightly worse than the one of Lemma 3.1. Namely we claim that, under the same conditions as in that lemma, (2) is replaced by

$$\mathbf{Adv}_{h^*}^{\text{au}}(A^*) \leq 2 \cdot [\max(n_1, n_2) + 1] \cdot \mathbf{Adv}_h^{\text{prf}}(A) + \frac{1}{2^c}, \quad (36)$$

and the time-complexity of A increases from $(n_1 + n_2)$ computations of h to $2 \max(n_1, n_2)$ computations of h . For some intuition about the proof, imagine h^* were a PRF. (It is not.) Then $\text{Coll}_{h^*}(M_1, M_2)$ would be about the same as the probability that $f(M_1) = f(M_2)$ for a random function f , because otherwise the pf-adversary who queried its oracle with M_1, M_2 and accepted iff the replies were the same would be successful. With h^* in fact only a pf-PRF, the difficulty is the case that $M_1 \subseteq M_2$, which renders the adversary just described not prefix-free. There is a simple (and well-known) observation —we will call it the extension trick— to get around this. Namely, assuming wlog $M_1 \neq M_2$ and $\|M_1\|_b \leq \|M_2\|_b$, let $x \in B$ be a block different from $M_2[\|M_1\|_b + 1]$, and let $M'_1 = M_1 \| x$ and $M'_2 = M_2 \| x$. Then $\text{Coll}_{h^*}(M_1, M_2) \leq \text{Coll}_{h^*}(M'_1, M'_2)$ but M_1 is not a prefix of M_2 . This leads to the prefix-free pf-adversary against h^* below:

Adversary D^f

$(M_1, M_2) \xleftarrow{\$} A^*$

If $M_1 \subseteq M_2$ then $x \xleftarrow{\$} B \setminus \{M_2[\|M_1\|_b + 1]\}$; $M'_1 \leftarrow M_1 \| x$; $M'_2 \leftarrow M_2 \| x$

Else $M'_1 \leftarrow M_1$; $M'_2 \leftarrow M_2$

If $f(M'_1) = f(M'_2)$ then return 1 else return 0

Here $f: B^+ \rightarrow \{0, 1\}^c$ and we assume wlog that $M_1, M_2 \in B^+$ are distinct messages with $\|M_1\|_b \leq \|M_2\|_b$. Now the result of [4] gives us a pf-adversary A against h such that (35) holds. Thus:

$$\begin{aligned} \mathbf{Adv}_{h^*}^{\text{au}}(A^*) - 2^{-c} &\leq \mathbf{Adv}_{h^*}^{\text{prf}}(D) \\ &\leq 2 \cdot [\max(n_1, n_2) + 1] \cdot \mathbf{Adv}_h^{\text{prf}}(A) + 2^{-c}. \end{aligned}$$

Re-arranging terms yields (36).

The time-complexity of A as per [4] is (essentially) the time-complexity t of A^* , rather than being a small quantity independent of t as in Lemma 3.1. It is not clear whether or not the coin-fixing argument of our proof of Section 3.3 can be applied to A to reduce this time-complexity. (One would have to enter into the details of the proof of [4] to check.) However, instead, we can first modify A^* to an adversary that has embedded in its code a pair $M_1, M_2 \in B^+$ of distinct messages that maximize $\text{Coll}_{h^*}(M_1, M_2)$. It just outputs these messages and halts. We then apply the argument above to this modified A^* , and now A will have time-complexity of $2 \max(n_1, n_2)$ computations of h plus minor overhead.

COMPARISONS. For the case that $M_1 \subseteq M_2$, our direct proof (meaning the one of Section 3.3) uses a different (and novel) idea as opposed to the extension trick, which leads to a factor of only $n_1 + 1$ in the bound (Claim 3.7) in this case, as opposed to the $2[\max(n_1, n_2) + 1]$ factor obtained via the reduction-from-pf-PRF proof. The difference can be significant in the case that M_2 is long and M_1 is short. In the case $M_1 \not\subseteq M_2$ our direct proof relies heavily on ideas of [4], but avoids the intermediate reduction to the multi-oracle model they use and exploits the non-adaptive nature of the setting to improve the factor in the bound from $2 \max(n_1, n_2)$ to $n_1 + n_2 - 1$. These improvements (overall, a constant factor) may not seem significant in this context. But they become significant when translated to NMAC, as evidenced by Theorem 3.4. As we saw, the latter gives appreciably better bounds than Theorem 3.3 in the case that authenticated messages have varying lengths. However, our direct proof (meaning, Lemma 3.1) is crucial to obtaining Theorem 3.4. The reduction-from-pf-PRF proof will only yield (a result that is slightly worse than) Theorem 3.3.

The reduction-from-pf-PRF proof is certainly simpler than our direct proof if one is willing to take the result of [4] as given. However, especially for a construct that is as widely standardized as HMAC, we think it is useful to have from-scratch proofs that are as easily verifiable as possible. If the measure of complexity is that of a from-scratch (i.e. self-contained) proof, we contend that our direct one (although not trivial) is simpler than that of [4]. (In particular because we do not use the multi-oracle model.) We remark that if a reader's primary interest is the simplest possible self-contained proof regardless of the quality of the bound, the way to get it is to use our direct proof for the case $M_1 \not\subseteq M_2$ and then the extension trick (as opposed to our direct proof) for the case $M_1 \subseteq M_2$.