A Closer Look at HMAC

Krzysztof Pietrzak

IST Austria

April 12th 2013

Abstract. Bellare, Canetti and Krawczyk [BCK96] show that cascading an ε-secure (fixed input length) PRF gives an $O(\varepsilon nq)$-secure (variable input length) PRF when making at most q prefix-free queries of length n blocks. We observe that this translates to the same bound for NMAC (which is the cascade without the prefix-free requirement but an additional application of the PRF at the end), and give a matching attack, showing this bound is tight. This contradicts the $O(\varepsilon n)$ bound claimed by Koblitz and Menezes [KM12].

Definitions. For a keyed function $\mathsf{F} : \{0,1\}^c \times \{0,1\}^b \to \{0,1\}^c$ we denote with $\mathsf{casc}^\mathsf{F} : \{0,1\}^{2c} \times \{0,1\}^{b*} \to \{0,1\}^c$ (where $\{0,1\}^{b*} = \bigcup_{z\in\mathbb{N}}\{0,1\}^{bz}$) the cascade (aka. Merkle-Damgård) construction build from F as

$$\mathsf{casc}^\mathsf{F}(k, m_1\|\ldots\|m_n) = y_n \text{ where } y_0 = k \text{ and for } i \geq 1 \ : \ y_i = \mathsf{F}(y_{i-1}, m_i)$$

nmac^F is casc^F with an additional application of F at the end (using some padding if $b > c$).

$$\mathsf{nmac}^\mathsf{F}((k_1,k_2),M) = \mathsf{F}(k_2, \mathsf{casc}^\mathsf{F}(k_1,M))$$

A variable input length function $\mathsf{G} : \{0,1\}^{2c} \times \{0,1\}^{b*} \to \{0,1\}^c$ is a (ε,t,q,n)-secure PRF (for fixed input length functions we omit the parameter n) if for any adversary A of size t, making q queries, each of length at most n (in b-bit blocks) and $\mathcal{R}$ denoting a uniformly random function with the same domain

$$\left|\Pr_{k\leftarrow\{0,1\}^c}[\mathsf{A}^{G(k,.)}] \to 1] - \Pr_{\mathcal{R}}[\mathsf{A}^{\mathcal{R}(.)} \to 1]\right| \leq \varepsilon$$

Upper Bound.

Theorem 1 ([BCK96] casc^F is a PRF[1]). *If F is an (ε,t,q)-secure PRF then casc^F is an (ε',t',q,n)-secure PRF with if queried on prefix-free messages*

$$\varepsilon' = O(\varepsilon q n) \qquad t' = t - \tilde{O}(qn)$$

[1] This is Theorem 3.1 in the full version of [BCK96] http://charlotte.ucsd.edu/ mihir/papers/cascade.pdf

As any q-query distinguisher who can find a collision in $\mathsf{casc}^{\mathsf{F}}$ with advantage $\delta \in O(\varepsilon qn)$ can be turned into a distinguisher for $\mathsf{casc}^{\mathsf{F}}$ with advantage $\delta - q^2/2^c$ (as the probability that a random function collides on any q queries is $\leq q^2/2^c$), we get

Corollary 1. *Let F be as in the above theorem. Then for any q distinct messages $M_1, \ldots, M_q$ of length at most n*

$$\Pr_{k \leftarrow \{0,1\}^c}[\exists i \neq j \ : \ \mathsf{casc}^{\mathsf{F}}(k, M_i) = \mathsf{casc}^{\mathsf{F}}(k, M_j)] = O(\varepsilon qn)$$

Note that unlike in Theorem 1, in Corollary 1 we did not require the messages to be prefix-free. The reason we can drop this requirement is that we can make the M_i's prefix free by adding some block $X \in \{0,1\}^b$ (that does not appear in any of the M_i's) at the end of every message. This will make the messages prefix-free, but will no decrease the collision probability.[2]

Proposition 1 ($\mathsf{nmac}^{\mathsf{F}}$ is a PRF). *If F is an (ε, t, q)-secure PRF then $\mathsf{nmac}^{\mathsf{F}}$ is an (ε', t', q, n)-secure PRF with*

$$\varepsilon' = O(\varepsilon qn) \qquad t' = t - \tilde{O}(qn)$$

Proof. Let $\mathsf{nmac}^{\mathsf{F}}_{+}$ denote $\mathsf{nmac}^{\mathsf{F}}$, but where the outer application of $\mathsf{F}(k_2, .)$ is replaced with a random function $\mathcal{R}(.)$. By the security of F, one cannot distinguish $\mathsf{nmac}^{\mathsf{F}}$ from $\mathsf{nmac}^{\mathsf{F}}_{+}$ but with advantage ε (by a reduction of complexity $\tilde{O}(qn)$).

The output of $\mathsf{nmac}^{\mathsf{F}}_{+}(.) = \mathcal{R}(\mathsf{casc}(k_1, .))$ is uniformly random, as long as all the outputs of the inner $\mathsf{casc}(k_1, .)$ function are distinct. This implies that distinguishing $\mathsf{nmac}^{\mathsf{F}}_{+}$ from random is at most as hard as provoking a collision on the inner function (by Theorem 1.(i) [Mau02]), and moreover adaptive strategies do not help (by Theorem 2 from [Mau02]). By Corollary 1 we can upper bound this advantage by $O(\varepsilon qn)$. $\square$

Note that the reduction we just gave is non-uniform as Corollary 1 does not specify how to actually find the messages M_i. To get a uniform reduction we use the fact from any adversary A who can distinguish $\mathsf{nmac}^{\mathsf{F}}_{+}$ from random with advantage δ one can actually extract messages $M_1, \ldots, M_q$ on which $\mathsf{nmac}^{\mathsf{F}}_{+}$ collides with expected probability at least δ by simply invoking A and collecting its queries, while answering them with uniformly random values. We then can make these M_i's prefix-free (if they are not already) by adding some block X to all of them, and now can use these to distinguish $\mathsf{casc}^{\mathsf{F}}$ from random with probability δ.

[2] As for any X, $\mathsf{casc}^{\mathsf{F}}(k, M_i) = \mathsf{casc}^{\mathsf{F}}(k, M_j)] \Rightarrow \mathsf{casc}^{\mathsf{F}}(k, M_i \| X) = \mathsf{casc}^{\mathsf{F}}(k, M_j \| X)]$

Lower Bound. We show that Proposition 1 is tight.

Proposition 2. *If PRFs exist, there exists an (ε, t, q)-secure PRF* F *where* $\mathsf{nmac}^{\mathsf{F}}$ *can be very efficiently (in time $\tilde{O}(qn)$) distinguished from random with advantage $\Omega(\varepsilon qn)$.*

Proof. We start with any $(\varepsilon/2, t, q)$-secure PRF F' from which we construct a (ε, t, q)-secure F by considering any set of "weak keys" $\mathcal{K}$ of size $2^c(\varepsilon/2)$, say the keys where the first $c - \log \varepsilon - 1$ bits are 0. We then define F as

$$\mathsf{F}(k, .) = \mathsf{F}'(k0, .) \text{ if } k \notin \mathcal{K} \text{ and } \mathsf{F}(k, .) = 0^c \text{ otherwise}$$

So, F behaves as F', except for weak keys where it's constantly 0^c (we can replace 0^c with any other weak key). It's not hard to show that F is a (ε, t, q)-secure PRF, i.e. compared to F' we loose at most an $\varepsilon/2$ term in distinguishing advantage by redefining it on an $\varepsilon/2$ fraction of the keys.

Assume we make two queries M_0, M_1 to $\mathsf{nmac}^{\mathsf{F}}(k = (k_1, k_2), .)$, which are sampled by first sampling an $n-1$ block long query $M = m_1 \| \ldots \| m_{n-1} \in \{0,1\}^{b(n-1}$ at random and then setting $M_0 = M \| x_0, M_1 = M \| x_1$ for any $x_0 \neq x_1$.

If one of the $n-1$ intermediate values in the evaluation of the inner function $\mathsf{casc}^{\mathsf{F}}(k_1, M)$ is in $\mathcal{K}$, then the output of $\mathsf{casc}^{\mathsf{F}}(k_1, M \| x)$ is 0^n. As this happens with probability $\approx (n-1)\varepsilon/2$

$$\Pr_{k_1, k_2}[\mathsf{nmac}^{\mathsf{F}}((k_1, k_2), M_0) = \mathsf{nmac}^{\mathsf{F}}((k_1, k_2), M_1) = \mathsf{F}(k_2, 0^c)] = \Theta(n\varepsilon)$$

If we query $\mathsf{nmac}^{\mathsf{F}}$ on $q/2$ such random and independently sampled message pairs M_0, M_1, the probability to observe a collision for at least one such pair is $\Theta(n\varepsilon q)$. As we expect to see a collision for such a pair when querying a random function with probability only $O(q/2^c)$ we get a distinguishing advantage of $\Theta(n\varepsilon q)$ as claimed.

References

[BCK96] Mihir Bellare, Ran Canetti, and Hugo Krawczyk. Pseudorandom functions revisited: The cascade construction and its concrete security. In *37th Annual Symposium on Foundations of Computer Science*, pages 514–523. IEEE Computer Society Press, October 1996.

[KM12] Neal Koblitz and Alfred Menezes. Another look at hmac. Cryptology ePrint Archive, Report 2012/074, 2012. We refer to the 14th revision of http://eprint.iacr.org/cgi-bin/versions.pl?entry=2012/074 (posted on Jan. 6th).

[Mau02] Ueli M. Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 110–132. Springer, April / May 2002.