The Exact Security of PMAC

Peter Gaži, Krzysztof Pietrzak and Michal Rybár *

IST Austria
{peter.gazi,pietrzak,michal.rybar}@ist.ac.at

Abstract. PMAC is a simple and parallel block-cipher mode of operation, which was introduced by Black and Rogaway at Eurocrypt 2002. If instantiated with a (pseudo)random permutation over n-bit strings, PMAC constitutes a provably secure variable input-length (pseudo)random function. For adversaries making q queries, each of length at most ℓ (in n-bit blocks), and of total length $\sigma \leq q\ell$, the original paper proves an upper bound on the distinguishing advantage of $O(\sigma^2/2^n)$, while the currently best bound is $O(q\sigma/2^n)$. In this work we show that this bound is tight by giving an attack with advantage $\Omega(q^2\ell/2^n)$.
In the PMAC construction one initially XORs a mask to every message block, where the mask for the ith block is computed as $\tau_i := \gamma_i \cdot L$, where L is a (secret) random value, and γ_i is the i-th codeword of the Gray code. Our attack applies more generally to any sequence of γ_i's which contains a large coset of a subgroup of $GF(2^n)$.
We then investigate if the security of PMAC can be further improved by using τ_i's that are k-wise independent, for $k > 1$ (the original distribution is only 1-wise independent). We observe that the security of PMAC will not increase in general, even if the masks are chosen from a 2-wise independent distribution, and then prove that the security increases to $O(q^2/2^n)$, if the τ_i are 4-wise independent. Due to simple extension attacks, this is the best bound one can hope for, using any distribution on the masks. Whether 3-wise independence is already sufficient to get this level of security is left as an open problem.

Keywords: Message Authentication Codes · PMAC · Attack · Masks

1 Introduction

PMAC (for Parallelizable Message Authentication Code) is a block-cipher mode of operation, introduced by Black and Rogaway at Eurocrypt 2002 [BR02]. The mode, when instantiated with a block-cipher over $\{0,1\}^n$, constitutes a variable input-length pseudorandom function $\{0,1\}^* \to \{0,1\}^n$ (which is then typically used for message authentication, hence the name). PMAC is slightly less efficient than, for example, modes based on CBC MAC, but its main advantage is that unlike CBC-based MACs, it allows to process the message blocks fully in parallel.

The secret key of PMAC specifies two permutations π, π' over $\{0,1\}^n$, and a function $\tau : \mathbb{N} \to \{0,1\}^n$ for determining the masks. On input a message $M = m_1 \| \ldots \| m_\ell, m_i \in \{0,1\}^n$, the output is computed as

$$\mathsf{PMAC}_{\pi,\pi',\tau}(M) = \pi'\left(\bigoplus_{i=1}^{\ell} \pi(m_i \oplus \tau(i))\right) . \tag{1}$$

In [BR02], the key is just a single key $K \in \mathcal{K}$ for a block-cipher $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$, π, π' are instantiated both with $E(K,.)$, and the mask function is defined as $\tau(i) = \gamma_i \cdot L$,

*Research supported by the European Research Council, ERC consolidator grant (682815-TOCNeT)

IACR Transactions on Symmetric Cryptology ISSN 2519-173X, Vol. 2016, No. 2, pp. 145–161
DOI:10.13154/tosc.v2016.i2.145-161

where γ_i is the ith Gray codeword[1] and $L = E(K, 0)$. This is a slightly idealized version of PMAC, we will discuss all the simplifications we make in greater detail in Section 1.3.

1.1 Security of PMAC in the Random Permutation Model

The security of a block-cipher mode of operation is usually analyzed assuming the underlying block-cipher under a random secret key realizes a uniformly random permutation. A bound in this model then implies security when instantiated with a block-cipher, we just have to add an extra term which bounds the advantage of distinguishing the block-cipher from a random permutation (i.e., the PRP security of the block-cipher, cf. Eq.(4) in this paper).

[BR02] proved an upper bound of $\sigma^2/2^n$ on the distinguishing advantage against PMAC for any adversary making a total of q queries, each of length at most ℓ blocks (of n bits), and a total of $\sigma \leq \ell q$ blocks. This was later improved to $q^2\ell/2^n$ by Minematsu and Matsushima at FSE'07 [MM07], and then to $q\sigma/2^n$ by Nandi at FSE'10 [Nan10] (note that $q\sigma$ can be much less than $q^2\ell$, if the message lengths vary a lot).

In this work we show that this bound is tight by giving an attack with advantage $\Omega(q^2\ell/2^n)$. For this, we show that it is possible to construct q messages $M_1, \ldots, M_q$ ($M_a = m_1^{(a)} \| m_2^{(a)} \| \ldots \| m_\ell^{(a)}$), such that for any pair of messages (M_a, M_b),

$$\bigoplus_{i=1}^{\ell} \pi(m_i^{(a)} \oplus \tau(i)) = \bigoplus_{i=1}^{\ell} \pi(m_i^{(b)} \oplus \tau(i))$$

for $\ell - 1$ different choices of L (where $\tau(i) = \gamma_i \cdot L$). Thus, also the PMAC of M_a, M_b (which additionally permutes this value) will collide. This directly gives a distinguishing attack, and even a forgery as now $M_a \| X$ and $M_b \| X$ will collide for any string X. Moreover, the set of L's for which two messages collide will be mostly disjoint for the $\binom{q}{2}$ pairs of messages, so with q messages of length ℓ we will observe a collision with probability in the order of $q^2\ell/2^n$.

Recently, Luykx *et al.* [LPSY16] showed that one can construct a pair of messages which will collide with probability roughly $\ell/2^n$, leading to an attack with advantage $\ell/2^n$ for $q = 2$ messages. However, their attack does not generalize to q messages. In contrast, our attack obtains this high collision probability for every of the $\binom{q}{2}$ message pairs.

1.2 k-wise Independent Masks

Several works show that by somewhat changing the construction, one can boost the security of PMAC [Yas11, Yas12, Zha15] even beyond the $q^2/2^n$ birthday bound. We investigate whether one can make the original construction more secure by just changing the distribution of the masks.

As a warm-up, in Section 4 we prove that if the masks $\tau_1, \tau_2, \ldots$ are uniform and independent, then the security indeed increases to $O(q^2/2^n)$. This is the best we can hope for under any distribution of masks: One can always query on random messages, and if a collision is found (which occurs with probability $q^2/2^n$), add the same block to both colliding messages, which will also lead to the same output.

The original distribution of masks in PMAC is only 1-wise independent, we investigate if the security increases when using k-wise independent distributions for $k > 1$. In Section 6 we show that 2-wise independence in general does not increase security by constructing a 2-wise independent distribution which, for any set of messages, gives us exactly the same collision probability as the original distribution. In Section 5 we show that using any 4-wise independent distribution on masks[2] will boost the security to the optimal $O(q^2/2^n)$. Whether 3-wise independence is sufficient is left as an open problem.

[1] This encoding is chosen to allow for efficient sequential computation of the values $\gamma_1 \cdot L, \gamma_2, \cdot L, \ldots$.

[2] For example computed as $\tau_i = \sum_{j=0}^{3} L_j \cdot i^j$ for random $L_j \in GF(2^n)$.

1.3 Variants of the PMAC Construction

The construction that we analyze is a somewhat simplified version of the actual original PMAC as proposed in [BR02]. We now discuss the existing differences and the applicability of our results to other variants.

One difference is that [BR02] specifies a padding which allows it to take as inputs messages whose length is not a multiple of n, moreover, the last block is not permuted. Additionally, a final mask (which is fixed and independent of ℓ) is XORed to the state before the outer permutation is applied. Our attacks and security proofs can be easily adapted to take these things into account, we chose not to do so for the sake of conceptual and notational simplicity. In particular, for our attack we choose q messages for the "simplified" PMAC as in Eq. (1) in a way that maximises the probability of seeing a collision. XORing a fixed value to the state before applying the outer permutation does not affect this collision probability. To handle the fact that the last block is not permuted we can simply add an arbitrary dummy message block to every message. Again, this will not affect the collision probability.

Another difference is that for our security proofs we assume that the value L used for the masks is sampled uniformly at random, while in the original construction $L := \pi(0)$. This distinction does not matter as long as $\ell q \ll 2^n$ (as then whp. none of the internal queries made is 0), which is satisfied for our main security result (Lemma 4) using 4-wise independent masks, as there we must assume $\ell \leq 2^{n/2}$ anyway. For our "warm-up" proof (Lemma 3) using independent random masks we don't have to make such an assumption, so here it's not clear if the result still applies with this difference for very large ℓ. This distinction also doesn't affect the success probability of our attack, which works for any distribution on L.

Moreover, in the security proofs we also assume that the inner and outer permutations π, π' are independent, while in the original construction π and π' are the same. If one aims for security in the order of $q^2\ell/2^n$ (or more generally $\sigma\ell/2^n$), this can be handled: informally, as there are q queries to π' and $q\ell$ queries to π, we expect them to overlap only with probability $q^2\ell/2^n$, and as long as they do not overlap, we can treat them as if they were independent. As we aim for $q^2/2^n$ security, it is not clear whether assuming that π and π' are independent is without loss of generality. Again, for our attack this distinction does not matter, the collision probability is the same no matter what π' is.

Let us also mention that there exists a later variant of PMAC due to Rogaway [Rog04] called PMAC1, which for efficiency reasons deviates slightly from PMAC by using a different sequence for the γ_i values. It is not clear if our attack can be adapted to this case. Informally, we require the sequence of $\gamma_1, \ldots, \gamma_\ell$ to contain a large coset of a subgroup of $GF(2^n)$, and it's not clear if the sequence from [Rog04] contains such a set. Let us mention that for similar reasons the attack from [LPSY16] does not apply to the [Rog04] construction either.

Newer variations of PMAC include PMAC+ [Yas11], PMAC with parity [Yas12], and PMACX [Zha15]. These introduce major modifications to the original constructions, therefore we do not discuss them in more detail. Lastly, LightMAC [LPTY16] can be considered a PMAC-like construction.

2 Preliminaries

Basic Definitions. For $n \in \mathbb{N}$ we define $[n] := \{1, \ldots, n\}$, and $\{0,1\}^{n*} := \bigcup_{z\in\mathbb{N}} \{0,1\}^{nz}$ denotes the set of all bitstrings whose length is a multiple of n. In a slight abuse of notation, we interchangeably view strings from $\{0,1\}^{n*}$ also as finite sequences of blocks from $\{0,1\}^n$, i.e., for $s \in \{0,1\}^{nz}$ we also write $s = (s_1, \ldots, s_z)$ for $s_i \in \{0,1\}^n$. The (bit)length of a string w is $|w|$, and if $|w|$ is a multiple of n, $|w|_n = |w|/n$ denotes the

length in n bit blocks. $w^\ell := w\|w\|\dots\|w$ denotes the ℓ-fold concatenation of w. We usually denote sets by calligraphic letters like $\mathcal{X}$. $\mathcal{F}_{b,c}$ (resp. $\mathcal{F}_{b*,c}$) denotes the set of all functions from $\{0,1\}^b$ to $\{0,1\}^c$ (resp. from $\{0,1\}^{b*}$ to $\{0,1\}^c$), $\mathcal{F}_{\mathbb{N},b}$ is the set of all functions $\mathbb{N} \to \{0,1\}^b$ and $\mathcal{P}_n$ the set of all permutations on $\{0,1\}^n$. If P is a (finite or infinite) progression, then by $P_{[\ell]}$ we denote a tuple containing the first ℓ elements of P. A *partition* of a set S is a collection of non-empty subsets A_i, such that if $A_i \neq A_j$, then $A_i \bigcap A_j = \emptyset$, and $\bigcup A_i = S$.

Multisets. We denote with $\mathsf{mult}(x, \mathcal{X})$ the multiplicity of an element x in a multiset $\mathcal{X}$. $\mathcal{X}^\downarrow$ is the subset of $\mathcal{X}$ that contains only the elements of odd multiplicity, i.e.,

$$\mathcal{X}^\downarrow = \{x \in \mathcal{X} : \mathsf{mult}(x, \mathcal{X}) \mod 2 = 1\} .$$

Groups and Cosets. For a definition of a commutative group and a discussion of the notions introduced below, see e.g. [Jor94]. All the groups that we consider in this paper will be commutative, and we will use additive notation for groups. A *subgroup* of G is any subset H that is a group by itself. The *order* of G, denoted $|G|$ is the number of its elements. Lagrange's theorem states that if H is a subgroup of G, then $|H|$ divides $|G|$.

Let G be a group, and H its subgroup. Take $g \in G$. Then the set $g + H := \{g + h : h \in H\}$ is called a *coset of H in G*. Note that trivially any group G is a coset (of G in G), we call a coset *proper* if it is not a group. The set of different cosets of H in G forms a partition of G; and moreover, H itself appears in it as the coset $0 + H$, where 0 is the neutral element of G (and H). The size of a coset is again referred to as its *order*. Finally, the order of G is equal to the product of the order of H and the number of different cosets of H.

Random Variables and Experiments. Random variables and concrete values they can take are usually denoted by upper-case letters $X, Y, \dots$ and lower-case letters $x, y, \dots$ respectively.

If $\mathcal{M}$ is a distribution (respectively, a set), then we denote by $X \xleftarrow{\$} \mathcal{M}$ sampling the random variable X according to $\mathcal{M}$ (respectively, choosing it uniformly at random from $\mathcal{M}$). By X^ℓ we denote ℓ independent and identically distributed copies of a random variable X. A joint probability distribution of q random variables $(X_1, \dots, X_q)$ is *k-wise independent*, if its restriction to any k coordinates is uniform over its domain, e.g., if all X_i have domain $\{0,1\}^n$

$$\forall i_1, \dots, i_k \ , \ 1 \le i_1 < \dots < i_k \le q \ ; \ \forall x_1, \dots, x_k \in \{0,1\}^n \ : \\ \mathsf{Pr}_{(X_1,\dots,X_q)} \left((X_{i_1}, \dots, X_{i_k}) = (x_1, \dots, x_k)\right) = \left(2^{-n}\right)^k \ .$$

More generally, let $\mathcal{M}_n$ be a probability distribution over $\mathcal{F}_{\mathbb{N},n}$. In this case, we call $\mathcal{M}_n$ *k-wise independent*, if any k outputs of $f(\cdot)$ sampled from $\mathcal{M}_n$ are independent. Formally, $\mathcal{M}_n$ is k-wise independent, if:

$$\forall i_1, \dots, i_k \ , \ 1 \le i_1 < \dots < i_k \ ; \ \forall x_1, \dots, x_k \in \{0,1\}^n \ : \\ \mathop{\mathsf{Pr}}_{f \xleftarrow{\$} \mathcal{M}_n} \Big(\big(f(i_1), \dots, f(i_k)\big) = (x_1, \dots, x_k)\Big) = \left(2^{-n}\right)^k \ .$$

Adversaries. In this work an adversary is a probabilistic (polynomial time or computationally unbounded) algorithm, sometimes with access to an oracle $\mathcal{O}(\cdot)$. We use sans-serif letters for adversaries, e.g., $\mathsf{A}^{\mathcal{O}(\cdot)}$, and will only consider "distinguishers", which are adversaries, whose final output is just one bit.

Pseudorandom functions and permutations. We call a function $f : \mathcal{K} \times \mathcal{D} \to \mathcal{R}$ *keyed*, where the first part of the input is referred to as the key (and $\mathcal{K}$ being called the *keyspace* of f). We often write $f_k(\cdot)$ instead of $f(k, \cdot)$. Given a variable input-length keyed function $f\colon \mathcal{K} \times \{0,1\}^{n*} \to \{0,1\}^n$, the PRF-advantage of an adversary A against f is defined as

$$\mathbf{Adv}^{\mathrm{prf}}_f(\mathsf{A}) := \mathsf{Pr}[K \xleftarrow{\$} \mathcal{K} \,:\, \mathsf{A}^{f_K(\cdot)} = 1] - \mathsf{Pr}[\mathsf{R} \xleftarrow{\$} \mathcal{F}_{n*,n} \,:\, \mathsf{A}^{\mathsf{R}(\cdot)} = 1] \,.$$

We also define

$$\mathbf{Adv}^{\mathrm{prf}}_f(q, \ell, t) := \max_{\mathsf{A}} \mathbf{Adv}^{\mathrm{prf}}_f(\mathsf{A})$$

where the maximum goes over all adversaries that run in time at most t, and ask at most q queries, each of length at most ℓ (in n-bit blocks). If we consider computationally unbounded adversaries, we drop the last argument, i.e., $\mathbf{Adv}^{\mathrm{prf}}_f(q, \ell) := \mathbf{Adv}^{\mathrm{prf}}_f(q, \ell, \infty)$.

Pseudorandom permutations (PRPs) and their security notions are defined analogously. Given a keyed permutation (i.e., a block-cipher) $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$, the PRP-advantage of an adversary A against E is defined as

$$\mathbf{Adv}^{\mathrm{prp}}_E(\mathsf{A}) := \mathsf{Pr}[K \xleftarrow{\$} \mathcal{K} \,:\, \mathsf{A}^{E_K(\cdot)} = 1] - \mathsf{Pr}[\mathsf{P} \xleftarrow{\$} \mathcal{P}_n \,:\, \mathsf{A}^{\mathsf{P}(\cdot)} = 1] \,.$$

and

$$\mathbf{Adv}^{\mathrm{prp}}_E(q, t) := \max_{\mathsf{A}} \mathbf{Adv}^{\mathrm{prp}}_E(\mathsf{A})$$

where the maximum goes over all adversaries that run in time at most t and ask at most q queries.

Collision security. For a keyed function $f\colon \mathcal{K} \times \{0,1\}^{n*} \to \{0,1\}$, we define

$$\mathbf{Adv}^{\mathrm{col}}_f(q, \ell) := \max_{M_1,\ldots,M_q} \mathsf{Pr}_{K \leftarrow \mathcal{K}}[\exists\, i \neq j \,:\, f_K(M_i) = f_K(M_j)] \,,$$

where the maximum goes over all q tuples of distinct messages of length at most ℓ blocks.

3 PMAC and Simplified PMAC

We define the simplified PMAC, $\mathsf{sPMAC}\colon \mathcal{P}_n \times \mathcal{F}_{\mathbb{N},n} \times \{0,1\}^{n*} \to \{0,1\}^n$ as

$$\mathsf{sPMAC}(\pi, \tau, m_1 \| \ldots \| m_\ell) := \bigoplus_{i=1}^{\ell} \pi(m_i \oplus \tau(i)) \,.$$

$\mathsf{PMAC} : \mathcal{P}_n \times \mathcal{P}_n \times \mathcal{F}_{\mathbb{N},n} \times \{0,1\}^{n*} \to \{0,1\}^n$ is derived from sPMAC by additionally encrypting the final output using an independent permutation π':

$$\mathsf{PMAC}(\pi, \pi', \tau, M) = \pi'(\mathsf{sPMAC}(\pi, \tau, M))$$

To save on notation, we will sometimes write τ_i instead $\tau(i)$ and e.g., $\mathsf{PMAC}_{\pi,\pi',\tau}(M)$ instead of $\mathsf{PMAC}(\pi, \pi', \tau, M)$, or, if π, π', τ are clear from the context, simply $\mathsf{PMAC}(M)$.

The first three (two) arguments of PMAC (sPMAC) are the key; consider distributions Π, Π' over $\mathcal{P}_n$, and $\mathtt{T}_n$ over $\mathcal{F}_{\mathbb{N},n}$, then $\mathsf{PMAC}_{\Pi,\Pi',\mathtt{T}_n}(.)$ denotes a keyed function, where the key is sampled according to $(\pi, \pi', \tau) \leftarrow \Pi \times \Pi' \times \mathtt{T}_n$, and then defines the function $\mathsf{PMAC}_{\pi,\pi',\tau}(.)$. If Π is instantiated by a block-cipher $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$, we think of it as the uniform distribution over the multiset of functions $\{E(k, \cdot) \,:\, k \in \mathcal{K}\}$.

For an input message $M = m_1 \| \ldots \| m_\ell$, it will be convenient to define the following variables

$$\begin{aligned} x_i &:= m_i \oplus \tau_i,\ \forall i \quad ; \quad \mathcal{X} := (x_1, \ldots, x_\ell) \\ y_i &:= \pi(x_i),\ \forall i \quad ; \quad \mathcal{Y} := (y_1, \ldots, y_\ell) \end{aligned} \tag{2}$$

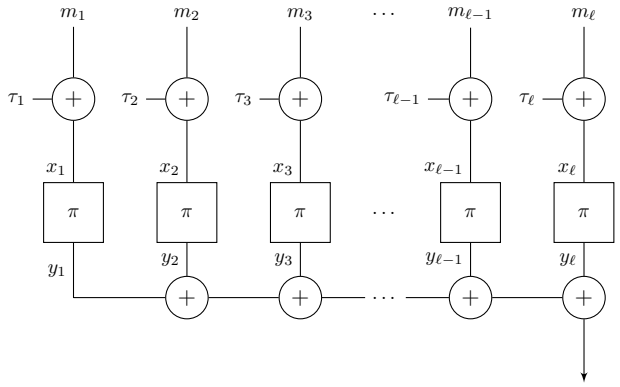

Figure 1: The evaluation of $\mathsf{sPMAC}(\pi, \tau, m_1 \| \ldots \| m_\ell)$, where $\tau_i = \tau(i)$.

We often consider pairs of messages $M = m_1 \| \ldots m_s, M' = m'_1 \| \ldots \| m'_{s'}$, then $\mathcal{X}^*$ denotes the multiset

$$x_i := m_i \oplus \tau_i \ , \ x'_i := m'_i \oplus \tau_i, \ \forall i \qquad ; \qquad \mathcal{X}^* := (x_1, \ldots, x_s, x'_1, \ldots, x'_{s'}) \tag{3}$$

We start by reducing the PRP-security of PMAC with a block-cipher E to the collision security of sPMAC with a random permutation. The argument is fairly standard and allows us to perform the rest of our analysis in the information-theoretic setting.

Lemma 1 (PRF security of PMAC from collision security of sPMAC)**.** *For a block-cipher $E : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$, and for any distribution $\mathtt{T}_n$ over $\mathcal{F}_{\mathbb{N},n}$, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{E,E,\mathtt{T}_n}}(q,\ell,t) \leq 2 \cdot \mathbf{Adv}^{\mathrm{prp}}_{E}(\ell q, t') + \mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n,\mathtt{T}_n}}(q,\ell) + \frac{q^2}{2^n} \ ,$$

where $t' \leq t + O(\ell q)$.

Proof. We first replace the block-cipher E with uniformly random permutations, by a straightforward reduction:

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{E,E,\mathtt{T}_n}}(q,\ell,t) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\mathcal{P}_n\,\mathcal{P}_n,\mathtt{T}_n}}(q,\ell,t) + 2 \cdot \mathbf{Adv}^{\mathrm{prp}}_{E}(\ell q, t') \ . \tag{4}$$

We can now consider computationally unbounded distinguishers (first step below), and replace the outer permutation by a uniformly random function, using the PRF/PRP switching lemma [BR06] in the second step:

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\mathcal{P}_n,\mathcal{P}_n,\mathtt{T}_n}}(q,\ell,t) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\mathcal{P}_n,\mathcal{P}_n,\mathtt{T}_n}}(q,\ell) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\mathcal{P}_n,\mathcal{F}_{n,n},\mathtt{T}_n}}(q,\ell) + \frac{q^2}{2^n} \tag{5}$$

Finally, we claim that distinguishing $\mathsf{PMAC}_{\mathcal{P}_n,\mathcal{F}_{n,n},\mathtt{T}_n}$ from a random function is upper bounded by the collision security of $\mathsf{sPMAC}_{\pi,\tau}$, i.e.,

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\mathcal{P}_n,\mathcal{F}_{n,n},\mathtt{T}_n}}(q,\ell) \leq \mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n,\mathtt{T}_n}}(q,\ell) \tag{6}$$

The statement of the Lemma follows from Eq.(4)-Eq.(6). It remains to prove Eq.(6). As the outer function $f \leftarrow \mathcal{F}_{n,n}$ is uniformly random, the output of $\mathsf{PMAC}_{\pi,f,\tau}(.) \equiv f(\mathsf{sPMAC}_{\pi,\tau}(.))$ is uniformly random, conditioned on not having any collisions on the inner function $\mathsf{sPMAC}_{\pi,\tau}(.)$. By Theorem 1.(i) from [Mau02], this implies that distinguishing $\mathsf{PMAC}_{\pi,f,\tau}(.)$ from a random function is at least as hard as provoking a collision on $\mathsf{sPMAC}_{\pi,\tau}(.)$, and further by Theorem 2 from [Mau02], adaptivity does not help in provoking this condition. This concludes the proof of Eq.(6). □

A *self-cancellation* for a message M (denoted $\mathsf{seCan}(M)$) occurs, if for its corresponding $\mathcal{X}$, we have $\mathcal{X}^{\downarrow} = \emptyset$. A *cross-cancellation* for two messages M, M' (denoted $\mathsf{crCan}(M, M')$) occurs, if for their corresponding $\mathcal{X}^{*\downarrow}$, we have $\mathcal{X}^{*\downarrow} = \emptyset$. A *PMAC-collision* for two messages M, M' (denoted $\mathsf{pCol}(M, M')$) occurs, if $\mathsf{PMAC}(M) = \mathsf{PMAC}(M')$. We define sPMAC-*collision* ($\mathsf{spCol}(M, M')$) analogously. Note that $\mathsf{crCan}(M, M')$ implies $\mathsf{spCol}(M, M')$, and $\mathsf{spCol}(M, M')$ implies $\mathsf{pCol}(M, M')$.

For a given n, ℓ and a distribution $\mathtt{T}_n$, we define the following quantity with the x_i's as defined in Eq.(2):

$$\theta(\ell, n, \mathtt{T}_n) = \max_{\substack{M \neq M' \\ |M|_n, |M'|_n \leq \ell}} \Pr_{\tau \xleftarrow{\$} \mathtt{T}_n} \left[\left\{ x_1, x_2, \ldots, x_{|M|_n}, x'_1, x'_2, \ldots, x'_{|M'|_n} \right\}^{\downarrow} = \emptyset \right] . \tag{7}$$

The quantity $\theta(\ell, n, \mathtt{T}_n)$ bounds the maximum probability over all pairs of distinct messages M, M' of maximum length ℓ, that their reduced set $\mathcal{X}^{*\downarrow}$ is empty, and hence a cross-cancellation occurs. This probability is taken over the sampling of the mask according to the distribution $\mathtt{T}_n$.

The following lemma states that a cross-cancellation is indeed the dominant reason for an sPMAC-collision to occur.

Lemma 2. *For any $n, \mathtt{T}_n$, and $\ell \leq 2^{n-2}$*

$$\mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n, \mathtt{T}_n}}(q, \ell) \leq \theta(\ell, n, \mathtt{T}_n) \cdot q^2 + \frac{q^2}{2^{n-1}} .$$

Proof. By taking a union bound over all q messages, we can upper bound the probability of a collision amongst the q messages by the probability of any pair colliding as:

$$\mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n, \mathtt{T}_n}}(q, \ell) \leq \mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n, \mathtt{T}_n}}(2, \ell) \cdot \binom{q}{2} \leq \mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n, \mathtt{T}_n}}(2, \ell) \cdot q^2 .$$

We upper bound $\mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n, \mathtt{T}_n}}(2, \ell)$ by showing that for any $M \neq M', |M|_n, |M'|_n \leq \ell$ we have:

$$\begin{aligned}
&\Pr_{(\pi,\tau) \leftarrow \mathcal{P}_n \times \mathtt{T}_n}[\mathsf{sPMAC}_{\pi,\tau}(M) = \mathsf{sPMAC}_{\pi,\tau}(M')] \\
&\quad = \Pr_{(\pi,\tau) \leftarrow \mathcal{P}_n \times \mathtt{T}_n}[\mathsf{sPMAC}_{\pi,\tau}(M) = \mathsf{sPMAC}_{\pi,\tau}(M') \wedge \mathsf{crCan}(M, M')] && (8) \\
&\quad + \Pr_{(\pi,\tau) \leftarrow \mathcal{P}_n \times \mathtt{T}_n}[\mathsf{sPMAC}_{\pi,\tau}(M) = \mathsf{sPMAC}_{\pi,\tau}(M') \wedge \overline{\mathsf{crCan}(M, M')}] && (9) \\
&\quad \leq \theta(\ell, n, \mathtt{T}_n) + \frac{1}{2^n - 2\ell} \\
&\quad \leq \theta(\ell, n, \mathtt{T}_n) + \frac{1}{2^{n-1}} && (10)
\end{aligned}$$

Note that this proves the statement of the Lemma. In eq. (10), we have used $\ell \leq 2^{n-2}$. The term (8) can be upper bounded as (using that for any events E_0, E_1,$\Pr[E_0 \wedge E_1] \leq \Pr[E_0]$)

$$(8) \leq \Pr_{(\pi,\tau) \leftarrow \mathcal{P}_n \times \mathtt{T}_n}[\mathsf{crCan}(M, M')] \leq \theta(\ell, n, \mathtt{T}_n) ,$$

where the 2nd step follows by definition.

It remains to upper bound the term (9) by $1/(2^n - 2\ell)$. We first upper bound (9) by fixing τ to the "worst case", and condition on crCan (using $\Pr[E_0 \wedge E_1] \leq \Pr[E_0|E_1]$)

$$(9) \leq \max_{\tau} \Pr_{\pi \leftarrow \mathcal{P}_n}[\mathsf{sPMAC}_{\pi,\tau}(M) = \mathsf{sPMAC}_{\pi,\tau}(M') \mid \overline{\mathsf{crCan}(M, M')}]$$

As $\overline{\mathsf{crCan}(M, M')}$, the set $\mathcal{X}^{*\downarrow} = \{a_1, \ldots, a_s\}$ is non-empty and $s \leq 2\ell$. A necessary (albeit not sufficient) condition to have a collision is that $\bigoplus_{i=1}^{s} \pi(a_i) = 0$. We claim that

$$\Pr_{\pi \leftarrow \mathcal{P}_n}\left[\bigoplus_{i=1}^{s} \pi(a_i) = 0\right] = \Pr_{\pi \leftarrow \mathcal{P}_n}\left[\bigoplus_{i=1}^{s-1} \pi(a_i) = \pi(a_s)\right] \leq \frac{1}{2^n - s + 1}$$

The first equality follows as $A \oplus B = 0$, if and only if $A = B$. To see the second step, assume the output of the random π is defined in a lazy way (sampling a random image without repetition for every fresh input), starting with inputs $a_1, \ldots, a_{s-1}$. Once these have been defined, we know that $\pi(a_s)$ will be uniform over a set of size $2^n - s + 1$, but at most one value, namely $\pi(a_s) = \bigoplus_{i=1}^{s-1} \pi(a_i)$, will satisfy the required condition. □

4 Independent Random Masks

In this section, as a warm-up, we look at the setting where the masks are chosen independently and uniformly at random.

Lemma 3. *For any* $n, \ell \in \mathbb{N}$

$$\theta(n, \ell, \mathcal{F}_{\mathbb{N},n}) \leq \frac{2}{2^n} .$$

Before we prove the lemma, we note that this upper bound is tight: consider the messages $M = 0^n \| 0^n$ and $M' = 1^n \| 1^n$, then for any choice of τ_1, we'll have $\mathcal{X}^{*\downarrow} = \{\tau_1, \tau_2, 1^n \oplus \tau_1, 1^n \oplus \tau_2\} = \emptyset$, if either $\tau_2 = \tau_1$, or $\tau_2 = \tau_1 \oplus 1^n$. As the τ_i are uniform, $\mathsf{Pr}[(\tau_1 = \tau_2) \vee \tau_1 = \tau_2 \oplus 1^n] = 2/2^n$.

Proof. Recall that $\theta(n, \ell, \mathcal{F}_{\mathbb{N},n})$ is the probability, maximized over all $M \neq M'$ of length $|M|_n, |M'|_n \leq \ell$, that $\mathcal{X}^{*\downarrow} = \emptyset$. Let $M = m_1 \| \ldots \| m_s$, $M' = m'_1 \| \ldots \| m'_{s'}$ denote the messages maximising this probability.

If $s = s'$, i.e., the messages are of same length, then let i be the smallest index where $m_i \neq m'_i$ (this i exists as $M \neq M'$). Assume $\tau \leftarrow \mathcal{F}_{\mathbb{N},n}$ is sampled on all inputs, except i. Consider the multiset $\mathcal{Z}_i = \mathcal{X}^* - (x_i, x'_i)$. Now, we claim $\mathcal{X}^{*\downarrow}$ will be empty, if and only if two conditions are satisfied. Firstly, $\mathcal{Z}_i^{\downarrow}$ must contain exactly two elements, let's call them $\{a, b\}$. If this is not the case, then $\mathcal{X}^{*\downarrow}$ will not be empty with probability 1.[3]

If it has exactly two elements $\{a, b\}$, then secondly, $\tau(i)$ must be chosen such that $\{m_i \oplus \tau(i), m'_i \oplus \tau(i)\} = \{x_i, x'_i\} = \{a, b\}$. There are at most two possible values for $\tau(i)$, which satisfy this condition (two, not just one, as the sets are not ordered). As $\tau(i)$ is uniform, the probability it hits one of those two values is $2/2^n$.

Now let us consider the case where $s \neq s'$ (without loss of generality, $s > s'$). We can't use the above argument here as an index i with $m_i \neq m'_i$ will not exist if M is a prefix of M'. In this case we assume $\tau(i)$ is given to us on all inputs except $s + 1$. Following a similar argument as with $\mathcal{Z}_i^{\downarrow}$ above, there will be at most one value for $\tau(s + 1)$ that will cause $\mathcal{X}^{*\downarrow}$ to be empty, which upper bounds the probability of $\mathcal{X}^{*\downarrow} = \emptyset$ to $1/2^n$. □

Lemma 3, in combination with Lemmas 1 and 2, directly give us the following statement.

Theorem 1 (PMAC security with uniform masks)**.** *For any* q, t, n, ℓ*, where* $\ell \leq 2^{n-2}$*, and block-cipher* E *with block-size* n*, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{E,E,\mathcal{F}_{\mathbb{N},n}}}(q, \ell, t) \leq \frac{5q^2}{2^n} + 2 \cdot \mathbf{Adv}^{\mathrm{prp}}_{E}(\ell q, t') ,$$

where $t' \leq t + O(\ell q)$.

[3] Intuitively, we need $\{a, b\}$ and $\{x_i, x'_i\}$ to cancel each other out for $\mathcal{X}^{*\downarrow} = \emptyset$. If $\mathcal{Z}_i^{\downarrow}$ had more, or less than precisely two elements, this would not be possible.

5 4-wise Independent Masks

In this section we investigate the security of PMAC if the mask distribution $\mathtt{T}_n$ is 4-wise independent. By the following lemma this assumption is sufficient to prove a bound of order $q^2/2^n$ (i.e., independent of the message length ℓ) on the PRF-security of PMAC assuming an exponential upper bound on ℓ.

Lemma 4. *For any $n, \ell \in \mathbb{N}$ where $\ell \leq 2^{n/2}$ and any 4-wise independent distribution $\mathtt{T}_n$, we have*

$$\theta(\ell, n, \mathtt{T}_n) \leq \frac{4}{2^n} \,. \tag{11}$$

Proof. Let $M = m_1 \| \ldots \| m_s, M' = m'_1 \| \ldots \| m'_{s'}, s, s' \leq \ell$ be messages maximizing the probability in the definition of $\theta(\ell, n, \mathtt{T}_n)$ (cf. Eq.(7)) for the 4-wise independent mask distribution $\mathtt{T}_n$.

We start with the case where $s = s'$. Let $I = \{i : m_i \neq m'_i\}$ be the indices of message blocks where the two messages differ, and $\mathcal{X}^*_I \subseteq \mathcal{X}^* = \{x_1, \ldots, x_s, x'_1, \ldots, x'_s\}$ the multiset containing only x_i, x'_i for $i \in I$. Note that $\mathcal{X}^{*\downarrow} = \mathcal{X}^{*\downarrow}_I$, since $m_i = m'_i$ implies $x_i = x'_i$ and for any multiset $\mathcal{S}$ and any x such that $\mathsf{mult}(\mathcal{S}, x) \geq 2$, we have $\mathcal{S}^{\downarrow} = (\mathcal{S} \setminus \{x, x\})^{\downarrow}$. In order to bound the probability that $\mathcal{X}^{*\downarrow} = \emptyset$, it suffices to bound the probability that $\mathcal{X}^{*\downarrow}_I = \emptyset$, let us denote this event with E_{col}.

If E_{col} holds, then we can find a complete matching (i.e., a subgraph where every vertex has degree exactly 1) in a graph, whose vertices are the elements of $\mathcal{X}^*_I$ and two vertices are connected by an edge if and only if they have the same value. Note that if E_{col} holds then every value appears with even multiplicity, so this graph consists of cliques of even size.

We will define a set of events $\{E_{\alpha,\beta} : (\alpha, \beta) \in (I \times \{0,1\})^2\}$ and prove that if $|I| \geq 4$ (we will discuss the cases where $|I| < 4$, and $s \neq s'$ at the end), then:

i. For any α, β, $\mathsf{Pr}[E_{\alpha,\beta}] \leq 2^{-2n}$.

ii. E_{col} implies that for some α, β the event $E_{\alpha,\beta}$ holds.

The above two points then imply $\mathsf{Pr}[E_{col}] \leq \sum_{\alpha,\beta} \mathsf{Pr}[E_{\alpha,\beta}] \leq 2^2 |I|^2 / 2^{2n} \leq 4\ell^2 / 2^{2n}$, which is upper bounded by $4/2^n$ if $\ell \leq 2^{n/2}$, as claimed in the statement of the lemma.

It will be convenient to define the index of a message as $\alpha = (\alpha_i, \alpha_b) \in I \times \{0,1\}$, where $m_\alpha = m_{\alpha_i}$ if $\alpha_b = 0$ and $m_\alpha = m'_{\alpha_i}$ if $\alpha_b = 1$ (similarly for x_α), so the part α_i identifies the block number and the bit α_b indicates whether we consider M or M'.

Let $I = \{i_1, i_2, \ldots\}$ and $\gamma = (\gamma_i, \gamma_b) = (i_1, 0)$, now the event $E_{\alpha,\beta}$ is defined as follows. Let

$$\delta = (\delta_i, \delta_b) = (\min\{I \setminus \{\gamma_i, \alpha_i, \beta_i\}\}, 0)$$

Note that above $\min\{I \setminus \{\gamma_i, \alpha_i, \beta_i\}$ is non-empty, as $|I| \geq 4$. If $\gamma_i = \alpha_i$, or $\alpha = \beta$, the event $E_{\alpha,\beta}$ is defined to never hold, so from now on we assume this is not the case. Then $E_{\alpha,\beta}$ is defined as

$$E_{\alpha,\beta} \iff (x_\gamma = x_\alpha) \wedge (x_\delta = x_\beta) \,.$$

We first prove that

$$\begin{aligned} \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[E_{\alpha,\beta}] &= \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[(x_\gamma = x_\alpha) \wedge (x_\delta = x_\beta)] \\ &= \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[x_\gamma = x_\alpha] \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[x_\delta = x_\beta | x_\gamma = x_\alpha] \\ &= 2^{-n} \cdot 2^{-n} \,. \end{aligned}$$

To see the last step above, note that

$$\mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[x_\gamma = x_\alpha] = \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[m_\gamma \oplus \tau_{\gamma_i} = m_\alpha \oplus \tau_{\alpha_i}] = 2^{-n}$$

holds as $\tau_{\gamma_i}, \tau_{\alpha_i}$, coming from a 4-wise independent distribution, are uniformly random and independent (recall we assume $\alpha_i \neq \gamma_i$). To show

$$\mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[x_\delta = x_\beta | x_\gamma = x_\alpha] = \mathsf{Pr}_{\tau \leftarrow \mathtt{T}_n}[m_\delta \oplus \tau_{\delta_i} = m_\beta \oplus \tau_{\beta_i} | m_\gamma \oplus \tau_{\gamma_i} = m_\alpha \oplus \tau_{\alpha_i}] = 2^{-n} \quad (12)$$

we note that, as the τ_i are 4-wise independent and $\delta_i \notin \{\alpha_i, \beta_i, \gamma_i\}$, the τ_{δ_i} is uniformly random even given all the other masks $\tau_{\alpha_i}, \tau_{\beta_i}, \tau_{\gamma_i}$. This concludes the proof of the condition (i), establishing that $\mathsf{Pr}[E_{\alpha,\beta}] \leq 2^{-2n}$.

It remains to show condition (ii), claiming that E_{col} implies that for some α, β the event $E_{\alpha,\beta}$ holds. For this we simply note that if E_{col} holds, then x_γ (with γ as defined above) must collide with at least some value x_α, and then the value x_δ (with δ as defined above) must collide with some x_β, thus $E_{\alpha,\beta}$ holds.

We have so far assumed that $|I| \geq 4$ and $s = s'$. If $|I| < 4$ (but we still assume $s = s'$), then there are at most $2(|I| - 1) = 4$ possible values x_γ can collide with, this probability is easily upper bonded by $4/2^n$ (2-wise independence of the τ_i is sufficient here). As x_γ colliding with another value x_α (where $\alpha_i \in I$) is a necessary condition for E_{col} to hold, the same upper bound holds of E_{col}.

We now shortly describe how to adapt the proof if the messages have different lengths, say $s > s'$. Let I again denote the set of indices $i \in \{1, \ldots, s'\}$ such that $m_i \neq m'_i$.

If $2|I| + (s' - s) \leq 6$ then we use basically the same argument as for the $|I| < 4$ case above; To have the event E_{col} the value x_γ (where $\gamma = (\min\{I\}, 0)$, or if $|I| = 0$, $\gamma = (s' + 1, 0)$) must collide with some x_α, and as there are at most 4 possibilities for α this probability is at most $4/2^n$. If $2|I| + (s' - s) > 6$ then we have at least 4 indices (namely I and $s' + 1, \ldots, s$) with correspond to x's that must collide, and for this one can use a slight generalisation of the argument for $|I| \geq 4$ from above. □

Again, combining Lemma 4 with Lemmas 1 and 2 give us the following statement.

Theorem 2 (PMAC security with 4-wise independent masks)**.** *For any q, t, n and $\ell \leq 2^{n/2}$, any block-cipher E with block-size n, and any 4-wise independent distribution $\mathtt{T}_n$ over $\mathcal{F}_{\mathbb{N},n}$, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{E,E,\mathtt{T}_n}}(q, \ell, t) \leq \frac{7q^2}{2^n} + 2 \cdot \mathbf{Adv}^{\mathrm{prp}}_{E}(\ell q, t') \ ,$$

where $t' \leq t + O(\ell q)$.

6 2-wise Independent Masks

In Section 5, we showed that the security of PMAC with 4-wise independent masks is $q^2/2^n$. On the other hand, in Section 7 we will show that when using the original distribution on masks from [BR02], which is only 1-wise independent, the security is just $\ell q^2/2^n$. This leaves open the question, whether we can get $q^2/2^n$ security already using any 2-wise or 3-wise independent distribution on masks. Below, we show that using a 2-wise independent distribution will in general not improve security: We slightly change the original distribution to make it 2-wise independent, and observe that this does not change the collision probability of sPMAC, and thus also attacker's distinguishing advantage of PMAC in the ideal permutation model, at all. Whether 3-wise independence is sufficient is left as an open problem.

Recall that in [BR02] the masks are computed by means of a function chosen at random from the following family

$$\{i \to a \cdot p_i \mid a \in GF(2^n)\} \ ,$$

where p_i is the i-th Gray codeword. For the following argument $P = (p_1, p_2 \ldots, p_{2^n})$ can be any progression without repetitions. Let $\mathtt{T}_n$ denote this distribution, and note that it is

1-wise, but not 2-wise, independent. Let $\mathtt{T}_n^+$ denote the uniform distribution over

$$\{i \to a \cdot p_i \oplus b \mid a, b \in GF(2^n)\} ,$$

which is 2-wise independent.

By the following lemma, the collision security of sPMAC is exactly the same for $\mathtt{T}_n$ and $\mathtt{T}_n^+$, thus also the security of PMAC implied by Lemma 1 will be the same for both distributions.

Lemma 5. *Let $\mathtt{T}_n$ and $\mathtt{T}_n^+$ be distributions as defined above. Then we have*

$$\mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n,\mathtt{T}_n}}(q,\ell) = \mathbf{Adv}^{\mathrm{col}}_{\mathsf{sPMAC}_{\mathcal{P}_n,\mathtt{T}_n^+}}(q,\ell) \ .$$

Proof. Consider any messages M, M' and $\mathcal{X}^* = (x_1, \ldots, x_{|M|_n}, x'_1, \ldots, x'_{|M'|_n})$ where $x_i = m_i \oplus a \cdot p_i \oplus b, x'_i = m'_i \oplus a \cdot p_i \oplus b$ for random a, b, i.e., according to mask distribution $\mathtt{T}_n^+$. To prove the lemma, it is sufficient to observe that if $\mathcal{X}^{*\downarrow} = \emptyset$, then we will still have $\mathcal{X}^{*\downarrow} = \emptyset$ even if we replace b with any other element of the field, in particular, we can assume $b = 0$ in which case we get mask distribution $\mathtt{T}_n$. □

7 1-wise Independent Masks: PMAC with a Gray Code

In this section we analyse the PRF-security of PMAC with a one-wise independent mask distribution.

The Gray Code. The original PMAC construction uses a mask distribution based on a Gray code, which is an example of a one-wise independent distribution. A Gray code is an ordering $\gamma^\ell = \gamma_0^\ell \gamma_1^\ell \ldots \gamma_{2^\ell-1}^\ell$ of $\{0,1\}^\ell$, for any $\ell \geq 1$, such that successive points differ in precisely one bit. The canonical Gray code from [BR02] is defined as follows:

$$\begin{aligned}
\gamma^1 &= (\gamma_0^1, \gamma_1^1) := (0, 1) \\
\gamma^2 &= (\gamma_0^2, \gamma_1^2, \gamma_2^2, \gamma_3^2) := (00, 01, 11, 10) \\
&\vdots \\
\gamma^{\ell+1} &= (0\gamma_0^\ell, 0\gamma_1^\ell, \cdots, 0\gamma_{2^\ell-2}^\ell, 0\gamma_{2^\ell-1}^\ell, 1\gamma_{2^\ell-1}^\ell, 1\gamma_{2^\ell-2}^\ell, \cdots, 1\gamma_1^\ell, 1\gamma_0^\ell)
\end{aligned}$$

In PMAC the sequence $\tau_1, \tau_2, \ldots$ of masks is defined as $\tau_i := \gamma_i^n \cdot L$ for a pseudorandom $L = E_K(0)$. Let us stress that the first mask is τ_1, so the first codeword $\gamma_0^n = 0^n$ is omitted. This fact makes our attack somewhat more complicated, as the lack of the zero element in the progression $\gamma_1^n, \gamma_2^n, \ldots$ will force us to argue over cosets of subgroups, instead of subgroups directly.

The [LPSY16] Attack. [LPSY16] show an attack on PMAC using two messages of length ℓ (for ℓ being any power of 2) with advantage roughly $\ell/2^n$. This attack exploits the fact that the first 2^w codewords of the *canonical* Gray code form a subgroup of the additive group of the finite field $GF(2^n)$. Hence, this two-query attack improves linearly with the increasing message length ℓ.

However, it is unclear whether this length-dependent attack can be generalized to a larger number of queries q. This is because the two attack queries are derived from the Gray code codewords being used, and are fully determined by them. Therefore, having more available message queries does not increase the success probability of the attack. Moreover, the set of L values that cause the two messages to collide on PMAC output is also predetermined by these codewords. Hence, there is a simple countermeasure against the attack: the user could simply avoid these "weak" keys.

7.1 Our Attack on PMAC

In this section we present an attack which scales with q, achieving success probability roughly $\ell q^2/2^n$ against PMAC. Moreover, this attack is randomized, so no "weak" keys exist, therefore a countermeasure against the [LPSY16] attack as mentioned above no longer applies.

Our attack can be mounted against PMAC using a similar class of 1-wise independent mask distributions as the attack in [LPSY16]. Namely, we assume that the masks are derived as $\tau_i := p_i \cdot R$ for some progression $P = (p_1, \ldots, p_{2^n})$ where every $p_i \in \{0,1\}^n$, and a value $R \xleftarrow{\$} \{0,1\}^n$ which we model as sampled uniformly at random.[4] We assume that all elements of P are distinct (any Gray code satisfies this property by definition). Our attack differs from the one in [LPSY16] in the message construction, and the type of collisions that it is aiming for. While in [LPSY16] the authors construct a pair of messages M, M' such that $\mathsf{seCan}(M)$ and $\mathsf{seCan}(M')$ occur with probability $\ell/2^n$ (over the choice of R), we choose q messages $M_1, \ldots, M_q$ such that for every pair M_i, M_j of them, $\mathsf{crCan}(M_i, M_j)$ occurs with probability $\ell/2^n$.

7.1.1 Description

We will use the following notation: given messages (i.e., attack queries) $M_1, \ldots, M_q$ of length ℓ each, we denote the i-th block of the a-th message by $m_i^{(a)}$. We also analogously define $x_i^{(a)} := m_i^{(a)} \oplus p_i \cdot R$.

The adversary $\mathsf{A} := \mathsf{A}_{\ell,q,n}^{\mathcal{O}(\cdot)}$ we present is parametrized by variables ℓ, q, n (maximal length of messages, number of messages, size of message blocks), and expects to interact with an oracle $\mathcal{O}(\cdot)$ that is either PMAC, or a random function. Its pseudocode is given as Algorithm 1.

The adversary A first identifies the largest possible subset $S \subseteq P_{[\ell]} = (p_1, \ldots, p_\ell)$ that is an additive subgroup of $GF(2^n)$; or more generally, a coset of any group H in G, where both H and G are additive subgroups of $GF(2^n)$ and do not need to be subsets of $P_{[\ell]}$. We denote the order of S by ℓ_S and the indices of S within $P_{[\ell]}$ by I_S. Additionally, we choose an arbitrary fixed element e in S. If S is a group then for notational convenience we choose $e := 0$, but this is of no significance to the attack, or its proof. Then, we denote by I'_S the indices of $S \setminus \{e\}$ within $P_{[\ell]}$.

Having identified S, the adversary samples q message blocks $\hat{m}^{(1)}, \ldots, \hat{m}^{(q)} \xleftarrow{\$} \{0,1\}^n$ one by one, using a form of rejection sampling. Namely, it maintains a set

$$\mathcal{U}_{a-1} = \left\{ \frac{\hat{m}^{(b)} \oplus \hat{m}^{(c)}}{e \oplus p_i} \;:\; b, c \in [a-1], b \neq c, i \in I'_S \right\} ,$$

where a is the index of $\hat{m}^{(a)}$ currently sampled (intuitively, all $u \in \mathcal{U}_{a-1}$ have the property that if $R = u$ then $\mathsf{crCan}(M_b, M_c)$ for some $b \neq c \leq [a-1]$ occurs). A random value sampled for $\hat{m}^{(a)}$ is then accepted, only if the intersection

$$\left\{ \frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{e \oplus p_i} \;:\; b \in [a-1], i \in I'_S \right\} \cap \mathcal{U}_{a-1}$$

is not too large (more precisely, if it is not larger than roughly twice its expected value).

From the blocks $\hat{m}^{(1)}, \ldots, \hat{m}^{(q)}$, A constructs a set of queries by repeating the same

[4] Note that this is not completely true for the value L described above, but we can afford this imprecision when modelling an *attack*, as it obviously does not significantly affect its performance.

Algorithm 1: Attacker $\mathsf{A}^{\mathcal{O}(\cdot)}_{\ell,q,n}$ against PMAC, where $P = (p_1, \ldots, p_{2^n-1})$

1 $I_S :=$ indices in $P_{[\ell]}$ of a coset $S \subseteq P_{[\ell]}$ of a subgroup H in an additive group $G \subseteq GF(2^n)$
2 $\ell_S := |S|$
3 fix arbitrary $e \in S$ (if S is a group, set $e := 0$)
4 $I'_S :=$ indices in $P_{[\ell]}$ of $S \setminus \{e\}$
5 $\mathcal{U}_0 := \emptyset$
6 **for** $a := 1 \ldots q$ **do**
7 **repeat**
8 $\hat{m}^{(a)} \xleftarrow{\$} \{0,1\}^n$
9 **until** $\left|\left\{\frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{e \oplus p_i} \ : \ b \in [a-1], i \in I'_S\right\} \cap \mathcal{U}_{a-1}\right| \leq \frac{2(a-1)^3(\ell_S-1)^2}{2^n}$
10 $\mathcal{U}_a := \mathcal{U}_{a-1} \cup \left\{\frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{e \oplus p_i} \ : \ b \in [a-1], i \in I'_S\right\}$
11 $M_a := \emptyset$
12 **for** $i := 1 \ldots \ell$ **do**
13 **for** $a := 1 \ldots q$ **do**
14 **if** $i \in I_S$ **then**
15 $M_a := M_a || \hat{m}^{(a)}$
16 **else**
17 $M_a := M_a || 0^n$
18 **for** $i := 1 \ldots q$ **do**
19 $\mathsf{Tag}_i := \mathcal{O}(M_i)$
20 **for** $i := 1 \ldots (q-1)$ **do**
21 **for** $j := (i+1) \ldots q$ **do**
22 **if** $\mathsf{Tag}_i = \mathsf{Tag}_j$ **then**
23 **return** 1
24 **return** 0

block ℓ times:

$$\begin{aligned} M_1 &= (\hat{m}^{(1)})^\ell = \hat{m}^{(1)} || \hat{m}^{(1)} || \ldots || \hat{m}^{(1)} \\ M_2 &= (\hat{m}^{(2)})^\ell = \hat{m}^{(2)} || \hat{m}^{(2)} || \ldots || \hat{m}^{(2)} \\ &\ldots \\ M_q &= (\hat{m}^{(q)})^\ell = \hat{m}^{(q)} || \hat{m}^{(q)} || \ldots || \hat{m}^{(q)} \end{aligned}$$

and then replaces all the blocks of these newly created messages that correspond to indices not in I_S by an all-zero block (in fact, any block with fixed value would do).

From this point on, the attack is simple: A submits the messages constructed above as the attack queries; if there is a collision among the outputs of the oracle it outputs 1, otherwise it outputs 0.

7.1.2 Analysis

We first look at the running time of A. The only nontrivial part of it that is worth consideration is the loop on lines 7–9, which might potentially never terminate. However,

note that the expected size of the set

$$\mathop{\mathbb{E}}_{\hat{m}^{(a)} \xleftarrow{\$} \{0,1\}^n} \left[\left| \left\{ \frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{e \oplus p_i} \; : \; b \in [a-1], i \in I'_S \right\} \cap \mathcal{U}_{a-1} \right| \right] \leq \frac{(a-1)^3(\ell_S - 1)^2}{2^n} \; ,$$

since each of the $(a-1)(\ell_S - 1)$ elements of the set intersected with $\mathcal{U}_{a-1}$ is individually uniform over $\{0,1\}^n$, and we are assessing the probability that it hits the set $\mathcal{U}_{a-1}$, where $|\mathcal{U}_{a-1}| \leq (a-1)^2(\ell_S - 1)$. Hence, the probability that a single iteration of the loop fails to satisfy the condition on line 9 is at most $1/2$ by Markov's inequality. Since every sampling on line 8 is independent, the probability (for a fixed a) that the loop is executed more than k times is upper bounded by 2^{-k}.

Now we move on to analyze the advantage achieved by our attack.

Theorem 3. *Let $P = (p_1, \ldots, p_{2^n}) \in GF(2^n)$ be a progression as defined above, and let $\mathtt{T}_n$ be the mask distribution defined as $\tau_i = p_i \cdot R$ for a random $R \xleftarrow{\$} \{0,1\}^n$. Let Π, Π' be any distributions over $\mathcal{P}_n$ and assume that $\ell q^2 \leq 2^{n-1}$. The adversary $\mathsf{A}_{\ell,q,n}$ given in Algorithm 1 achieves*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\Pi,\Pi',\mathtt{T}_n}}(\mathsf{A}_{\ell,q,n}) \geq \frac{(\ell_S - 1)(q-1)^2}{2^{n+2}} - \frac{q^2}{2^n} \; ,$$

where ℓ_S is the order of the largest coset S of some subgroup H in an additive subgroup G of $GF(2^n)$, such that the coset S is fully contained in $P_{[\ell]} = (p_1, \ldots, p_\ell)$.

Note that as a special case, we can have $S = G$ and hence ℓ_S may be the order of the largest additive group contained in $P_{[\ell]}$.

Proof. We start by investigating the probability of $\mathsf{crCan}(M_a, M_b)$ for two distinct indices $a, b \in \{1, \ldots, q\}$.

Lemma 6. *Let a, b be any two distinct indices from $\{1, \ldots, q\}$. Then, we have*

$$\mathsf{Pr}\left[\mathsf{crCan}(M_a, M_b)\right] \geq \frac{\ell_S - 1}{2^n} \; .$$

Proof. (of Lemma 6) To slightly simplify the notation, we first prove the theorem for the case where S is a group and then describe the straightforward extensions needed to handle the case where S is a proper coset.

Let us hence assume that S is a group and therefore $e = 0$. We will denote by z the index of e in $P_{[\ell]}$, i.e., $p_z = e = 0$. For $i \in I'_S$, let r_i denote the value

$$r_i := \frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{p_i} \; , \tag{13}$$

where the division occurs in $GF(2^n)$ (recall that $i \in I'_S$, and hence $p_i \neq 0$). We observe that if R is sampled to equal r_i, we obtain

$$\hat{m}^{(a)} \oplus \hat{m}^{(b)} = R \cdot p_i = R \cdot (p_z \oplus p_i) \; , \tag{14}$$

and hence

$$\hat{m}^{(a)} \oplus p_z \cdot R = \hat{m}^{(b)} \oplus p_i \cdot R \; , \tag{15}$$

which is equivalent to $x_z^{(a)} = x_i^{(b)}$.

Moreover, we claim that if $R = r_i$, we obtain a complete cross-cancellation for M_a and M_b. To observe this, first note that the equation (15) also trivially implies $x_i^{(a)} = x_z^{(b)}$.

Additionally, recall that we work in a field of characteristic 2, and hence the set $\{0 = p_z, p_i\}$ is a subgroup of S. Consequently, it induces a partition of S into $\ell_S/2$ cosets of the form $\{p_j, p_j \oplus p_i\}$, for $j \in I_S$. For a fixed j, let $k \in I_S$ be an index, such that $p_k = p_i \oplus p_j$ (there is a unique such index, since S is a group). For each coset $\{p_j, p_k\}$, we then obtain equalities $x_j^{(a)} = x_k^{(b)}$ and $x_k^{(a)} = x_j^{(b)}$, since (14) also implies

$$\hat{m}^{(a)} \oplus \hat{m}^{(b)} = R \cdot p_i = R \cdot (p_j \oplus p_i \oplus p_j) = R \cdot (p_j \oplus p_k) \ .$$

This is true for any $j \in I_S$ (hence, for all the $\ell_S/2$ cosets), implying a cross-cancellation.

Finally, note that for any $i \neq j$, we have $r_i \neq r_j$. This follows from equation (13), and the fact that S is a group. Hence, whenever R is sampled to take any of the $\ell_S - 1$ distinct values $\{r_i : i \in I'_S\}$, the event $\mathsf{crCan}(M_a, M_b)$ occurs, which concludes the proof for the case where S is a group.

Now assume that the set S is a proper coset of some subgroup H in a group $G \subseteq GF(2^n)$. Observe that $S = e \oplus H$, hence we can rewrite any element $p \in S$ as $p = e \oplus h$ for some $h \in H$ and vice versa, $h = e \oplus p$. For the sake of argument, imagine that the values $p_i \in S$ (note $S \subseteq P_{[\ell]} = (p_1, \ldots, p_\ell)$) on all positions in I_S would be replaced by $h_i := p_i \oplus g \in H$ instead; i.e., we would replace S by H in $P_{[\ell]}$ (recall that $|S| = |H|$). Then the previous analysis (for S being a subgroup) would apply, since H is a group. Now, if a cross-cancellation occurs in this modified setting with S replaced by H in $P_{[\ell]}$, then it also occurs before the replacement, as we have

$$\begin{aligned} \hat{m}^{(a)} \oplus p_i \cdot R = \hat{m}^{(b)} \oplus p_j \cdot R &\Leftrightarrow \hat{m}^{(a)} \oplus (p_i \oplus e) \cdot R = \hat{m}^{(b)} \oplus (p_j \oplus e) \cdot R \\ &\Leftrightarrow \hat{m}^{(a)} \oplus h_i \cdot R = \hat{m}^{(b)} \oplus h_j \cdot R \ . \end{aligned}$$

Hence, all the cancellations occur as before even if we replace H by S in $P_{[\ell]}$, and the rest of the analysis remains the same. □

The above lemma shows that for each M_a, M_b there are at least $\ell_S - 1$ "good" values R can take that would cause a cross-cancellation for M_a and M_b. Interestingly, this holds even if M_a and M_b are constructed from arbitrary distinct fixed values $\hat{m}^{(a)}$ and $\hat{m}^{(b)}$.

Let us refer to these potential values of R as (a, b)-*good*, and let $\mathcal{R}_{a,b}$ denote the set of all (a, b)-good values, formally

$$\mathcal{R}_{a,b} = \{r \in \{0,1\}^n : (R = r) \Rightarrow \mathsf{crCan}(M_a, M_b)\} \ .$$

Let $\mathcal{R} = \bigcup_{a \neq b \in [q]} R_{a,b}$ denote the set of all good values.

We now need to show that when we look at all $\binom{q}{2}$ pairs of A's queries, most of these good values for R will not overlap, giving us $|\mathcal{R}| = \Omega(\ell_S q^2)$ in total. To this end, we leverage the rejection sampling that A used to choose the building blocks $\hat{m}^{(a)}$.

Lemma 7. *Assuming $\ell_S q^2 \leq 2^{n-1}$, we have*

$$|\mathcal{R}| \geq \frac{(\ell_S - 1)(q-1)^2}{4} \ .$$

Proof. (of Lemma 7) For $a \in [q]$, let $\mathcal{V}_a$ denote the set of fresh values that are added to the set $\mathcal{U}_{a-1}$ in the a-th iteration of step 10 of the algorithm $\mathsf{A}^{\mathcal{O}(\cdot)}_{\ell,q,n}$ to form the set $\mathcal{U}_a$, formally $\mathcal{V}_a := \mathcal{U}_a \setminus \mathcal{U}_{a-1}$. By the definition of $\mathcal{U}_a$ on line 10, and the fact that we only count fresh values, we have

$$\mathcal{V}_a = \left\{ \frac{\hat{m}^{(a)} \oplus \hat{m}^{(b)}}{e \oplus p_i} \ : \ b \in [a-1], i \in I'_S \right\} \setminus \mathcal{U}_{a-1} \ .$$

The size of the set above before subtracting $\mathcal{U}_{a-1}$ is $(a-1)(\ell_S-1)$, and by the choice of $\hat{m}^{(a)}$ on lines 7–9, we know that the subtraction removes at most $2(a-1)^3(\ell_S-1)^2/2^n$ elements. Hence, we have

$$\begin{aligned}|\mathcal{V}_a| &\geq (a-1)(\ell_S-1) - \frac{2(a-1)^3(\ell_S-1)^2}{2^n} \\ &\geq (a-1)(\ell_S-1)\left(1 - \frac{2(a-1)^2(\ell_S-1)}{2^n}\right) \\ &\geq \frac{(a-1)(\ell_S-1)}{2}\,,\end{aligned}$$

where the last inequality follows, since $a^2\ell_S \leq q^2\ell_S \leq 2^{n-1}$. Clearly, $\mathcal{U}_q = \bigcup_{a=1}^{q} \mathcal{V}_a$, and by construction the sets $\mathcal{V}_a$ are disjoint. Hence, we obtain

$$|\mathcal{U}_q| = \sum_{a=1}^{q} |\mathcal{V}_a| \geq \sum_{a=1}^{q} \frac{(a-1)(\ell_S-1)}{2} \geq \frac{(\ell_S-1)(q-1)^2}{4}\,.$$

Finally, by observations in the proof of Lemma 6, we have $\mathcal{U}_q \subseteq \mathcal{R}$. Therefore, we can also conclude that $|\mathcal{R}| \geq \frac{(\ell_S-1)(q-1)^2}{4}$. □

To conclude the proof of Theorem 3, note that when $\mathcal{O} = \mathsf{PMAC}$, and if the randomly sampled R takes any value from $\mathcal{R}$, A observes a tag collision and outputs 1. According to Lemma 7, this happens with probability at least $(\ell_S-1)(q-1)^2/2^{n+2}$. On the other hand, if $\mathcal{O}$ is a random function, A observes such a collision (and hence outputs 1) with probability at most $q^2/2^n$. □

Consider the Gray code used in the original PMAC construction. This code does not include the zero element, hence the progression $P = (p_1, \ldots, p_{2^n-1})$ in this case does not contain any additive groups. However, it does contain some proper cosets. To see this, let G_i denote the additive subgroup of $GF(2^n)$ of size 2^i containing elements of the form $0^{n-i}w$ for $w \in \{0,1\}^i$. Then for any $\ell \geq 2^k - 1$ we get that $P_{[\ell]}$ contains the only proper coset of G_{k-1} in G_k, which is of size 2^{k-1}. This gives us the following corollary.

Corollary 1. *Consider the setting from Theorem 3, and let* $\mathtt{T}_n$ *be the mask distribution defined as* $\tau_i = \gamma_i^n \cdot R$ *for a random* $R \xleftarrow{\$} \{0,1\}^n$*, and* γ_i^n *being the* i*-th codeword in the canonical Gray code. Then, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathsf{PMAC}_{\Pi,\Pi',\mathtt{T}_n}}(\mathsf{A}_{\ell,q,n}) = \Omega(\ell q^2/2^n)\,.$$

References

[BR02] John Black and Phillip Rogaway. A block-cipher mode of operation for parallelizable message authentication. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 384–397. Springer, Heidelberg, April / May 2002.

[BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.

[Jor94] C. R. Jordan. *Groups*. Newnes, Oxford, 1994.

[LPSY16] Atul Luykx, Bart Preneel, Alan Szepieniec, and Kan Yasuda. On the influence of message length in PMAC's security bounds. LNCS, pages 596–621. Springer, Heidelberg, 2016.

[LPTY16] Atul Luykx, Bart Preneel, Elmar Tischhauser, and Kan Yasuda. A MAC mode for lightweight block ciphers. In *FSE 2016*, LNCS, pages 43–59. Springer, Heidelberg, 2016.

[Mau02] Ueli M. Maurer. Indistinguishability of random systems. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 110–132. Springer, Heidelberg, April / May 2002.

[MM07] Kazuhiko Minematsu and Toshiyasu Matsushima. New bounds for PMAC, TMAC, and XCBC. In Alex Biryukov, editor, *FSE 2007*, volume 4593 of *LNCS*, pages 434–451. Springer, Heidelberg, March 2007.

[Nan10] Mridul Nandi. A unified method for improving PRF bounds for a class of blockcipher based MACs. In Seokhie Hong and Tetsu Iwata, editors, *FSE 2010*, volume 6147 of *LNCS*, pages 212–229. Springer, Heidelberg, February 2010.

[Rog04] Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In Pil Joong Lee, editor, *ASIACRYPT 2004*, volume 3329 of *LNCS*, pages 16–31. Springer, Heidelberg, December 2004.

[Yas11] Kan Yasuda. A new variant of PMAC: Beyond the birthday bound. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 596–609. Springer, Heidelberg, August 2011.

[Yas12] Kan Yasuda. PMAC with parity: Minimizing the query-length influence. In Orr Dunkelman, editor, *CT-RSA 2012*, volume 7178 of *LNCS*, pages 203–214. Springer, Heidelberg, February / March 2012.

[Zha15] Yusi Zhang. Using an error-correction code for fast, beyond-birthday-bound authentication. In Kaisa Nyberg, editor, *CT-RSA 2015*, volume 9048 of *LNCS*, pages 291–307. Springer, Heidelberg, April 2015.