Optimal PRFs from Blockcipher Designs

Bart Mennink[1,2] and Samuel Neves[3]

[1] Digital Security Group, Radboud University, Nijmegen, The Netherlands
b.mennink@cs.ru.nl
[2] CWI, Amsterdam, The Netherlands
[3] CISUC, Dept. of Informatics Engineering, University of Coimbra, Portugal
sneves@dei.uc.pt

Abstract. Cryptographic modes built on top of a blockcipher usually rely on the assumption that this primitive behaves like a pseudorandom permutation (PRP). For many of these modes, including counter mode and GCM, stronger security guarantees could be derived if they were based on a PRF design. We propose a heuristic method of transforming a dedicated blockcipher design into a dedicated PRF design. Intuitively, the method consists of evaluating the blockcipher once, with one or more intermediate state values fed-forward. It shows strong resemblance with the optimally secure EDMD construction by Mennink and Neves (CRYPTO 2017), but the use of internal state values make their security analysis formally inapplicable. In support of its security, we give the rationale of relying on the EDMD function (as opposed to alternatives), and present analysis of simplified versions of our conversion method applied to the AES. We conjecture that our main proposal AES-PRF, AES with a feed-forward of the middle state, achieves close to optimal security. We apply the design to GCM and GCM-SIV, and demonstrate how it entails significant security improvements. We furthermore demonstrate how the technique extends to tweakable blockciphers and allows for security improvements in, for instance, PMAC1.

Keywords: PRP · PRF · EDMD · AES-PRF · GCM · GCM-SIV · PMAC1

1 Introduction

The conventional approach to cryptographic designs is to evaluate a blockcipher in a certain mode of operation, and undoubtedly the vast majority of MAC functions, encryption schemes, and authenticated encryption schemes follow this paradigm. It allows to reduce the security of the (keyed) construction in a standard model argument to the security of the keyed underlying primitive. The approach is, to a certain extent, a natural one. Ample literature discusses the design [DR02, KR11, DR01, RDP+96, DKR97, Vau03, Mat96, HT96, DPU+16] and analysis [BS93, Mat93, Knu94, LH94, JK97, BBS99, Wag99, BDK01] of blockciphers, and we even have a widely deployed and well-understood standardized blockcipher, the AES [DR02]. For modes that evaluate the cryptographic primitive in the forward as well as inverse direction, one in fact *needs* an invertible primitive, and a blockcipher is the most logical choice.

However, many cryptographic primitives in the literature evaluate the underlying blockcipher in forward direction only. For example, counter mode encryption [BDJR97] (the observation applies equally well to authenticated encryption mode GCM [MV04]) internally uses a blockcipher $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ to encrypt a message $m = m_1 \cdots m_l \in (\{0,1\}^n)^l$ as $c_i = E_k(\mathrm{ctr} + i) \oplus m_i$ for $i = 1, \ldots, l$. Counter mode can be distinguished from a random encryption scheme in about $2^{n/2}$ data blocks: an adversary can keep m_i constant and observe that the c_i never collide whereas they likely collide for a

IACR Transactions on Symmetric Cryptology ISSN 2519-173X, Vol. 2017, No. 3, pp. 228–252
DOI:10.13154/tosc.v2017.i3.228-252

random encryption scheme. Inspired by this, it makes more sense to not use a blockcipher, but rather a dedicated pseudorandom function inside counter mode (and GCM).

This is no longer a theoretical purity concern. Birthday-style attacks on common blockcipher modes of operation have been shown to be feasible in real scenarios [BL16, McG12], and in the case of counter mode they could be entirely avoided by choosing a PRF instead of a PRP. Although [BL16] could be chalked up to legacy ciphers (e.g., DES or Blowfish) being used in modern protocols, this is not necessarily the case in general. Over the last decade, there has been a flurry of ciphers targeting low-end hardware, which overwhelmingly have small block sizes in common. Some, like SIMON [BSS+13], SPECK, SIMECK [YZS+15], KATAN, or KTANTAN [CDK09], go as far as having 32-bit block variants. A birthday bound here renders these ciphers nearly unusable in several relevant modes of operation.

Another prominent example scheme that, to a lesser extent, benefits from using a pseudorandom function over a pseudorandom permutation is Wegman-Carter MAC [WC81, Bra82]:

$$\mathrm{WC}_{k,h}(u, m) = F_k(u) \oplus h(m)\,, \tag{1}$$

where F_k is a PRF and h is a universal hash function. Not only is WC typically used with a PRF, the concept of PRFs was developed (in part) to make Wegman-Carter work with short keys [Bra82, GGM86]. Nevertheless, given that blockciphers are better understood, Shoup suggested to use a blockcipher instead, in what is now known as the Wegman-Carter-Shoup MAC [Sho96],

$$\mathrm{WCS}_{k,h}(u, m) = E_k(u) \oplus h(m)\,, \tag{2}$$

where E_k is a PRP. But despite quantitative improvements from Shoup [Sho96] and Bernstein [Ber05], WCS based on a PRP remains stuck at the birthday bound, unlike WC based on a PRF.

Further examples of schemes that would benefit from the usage of a PRF abound. Unfortunately, unlike the case of blockciphers, dedicated fixed input length pseudorandom function designs are scarce: the only well-known candidate in literature is SURF by Bernstein [Ber97]. In fact, this scarcity was one of the reasons for the introduction of WCS over WC.

1.1 Generic PRP-PRF Conversion Functions

Various methods of generically transforming a PRP into a PRF have appeared in literature. First off, the well-known PRP-PRF switch [Fre77, IR88, BKR94, HWKS98, BR06, CN08] suggests to simply view the PRP as a PRF, which can be done as long as $q \ll 2^{n/2}$. Mennink and Neves [MN17] summarized four main directions in achieving security beyond the birthday bound.

A first direction is in truncating permutations, as first suggested in cryptographic context by Hall et al. [HWKS98]. Bellare and Impagliazzo [BI99] and later Gilboa and Gueron [GG16] proved that truncating an n-bit blockcipher by $m < n$ bits is secure up to about $2^{\frac{m+n}{2}}$ queries.[1] On the downside, truncation decreases the rate at which randomness is generated and hence makes the mode less efficient.

Bellare et al. [BKR98] were the first to suggest the xor of permutations,

$$\mathrm{XoP}_{k_1,k_2}(x) = E_{k_1}(x) \oplus E_{k_2}(x)\,. \tag{3}$$

Following a sequence of analyses by Lucks [Luc00] and Bellare and Impagliazzo [BI99], Patarin achieved $2^n/67$ security [Pat08, Pat13b, Pat10]. The results generalize to more

[1] In a non-cryptographic context, Stam [Sta78] in fact derived this result already in 1978.

permutations [CLP14, MP15], as well as to the single key variant with domain separation [Pat10]. Using XoP in counter mode or Wegman-Carter would yield optimal security, but the resulting scheme is twice as expensive.

Iwata [Iwa06] offered a compromise between the xor of permutations and traditional counter mode with the CENC mode of operation: E_{k_1} can be used as a mask which is used for the encryption of $w \geq 1$ blocks using E_{k_2}. If $w = 1$, CENC constitutes counter mode based on XoP, but also for larger (but still reasonably small) values of w, CENC achieves essentially optimal security [IMV16] and shares most of the benefits of the xor of permutations construction without much of a performance hit. Nevertheless, CENC remains a mode of operation, not a PRF primitive, and thus is not usable as a general replacement for a PRF.

Two novel constructions are EDM by Cogliati and Seurin [CS16] and EDMD by Mennink and Neves [MN17]:

$$\mathrm{EDM}_{k_1,k_2}(x) = E_{k_2}(E_{k_1}(x) \oplus x)\,, \tag{4}$$
$$\mathrm{EDMD}_{k_1,k_2}(x) = E_{k_2}(E_{k_1}(x)) \oplus E_{k_1}(x)\,. \tag{5}$$

Mennink and Neves proved security of EDM up to approximately $2^n/(67n)$ queries and EDMD up to approximately $2^n/67$ queries. However, just like the xor of permutations, these two generic modes are again twice as expensive.

1.2 Towards a Dedicated PRF

None of the above generic methods seem particularly suitable for the design of an efficient PRF. In particular, it is difficult to argue for their practical usage when they entail such a noticeable slowdown. What if, instead, we could design a secure and efficient PRF from scratch?

One could try and design a non-invertible round function, and design a PRF around its iteration. However, non-invertible round functions are hard to get right, as collision probabilities are amplified with each iteration, and the track record of this design approach is not very reassuring [PK14, Dae16].

Instead, our approach is to stick with tried-and-tested designs, namely those of blockciphers. Our key observation is that the EDMD structure of (5) is particularly suited to (heuristically) transform imperfect random permutations into a good PRF. We call this heuristic construction FastPRF. At a high level, the idea of FastPRF is as follows: if a blockcipher E_k consists of r rounds, the blockcipher is evaluated exactly one time, with a predetermined selection of state values fed-forward. Naturally, the strength of FastPRF highly depends on the blockcipher itself, as well as on the choice of states that are fed-forward. For example, if E is any blockcipher, and the 0th state is fed-forward, this effectively corresponds to

$$E_k(x) = E_k(x) \oplus x\,, \tag{6}$$

which can be distinguished from random in about $2^{n/2}$ evaluations (cf., Section 3.3). A more logical choice is to use the middle state. Let E_k^1 and E_k^2 be the first and second $r/2$ rounds of the cipher, for example. Then we can define FastPRF as

$$\mathrm{FastPRF}_k(x) = E_k(x) \oplus E_k^1(x)\,. \tag{7}$$

Closer inspection at SURF reveals that the high-level structure behind (7) matches that of SURF [Ber97], yet FastPRF is more general. We can observe that (7) resembles the structure of EDMD with $E_{k_1}(\cdot) = E_k^1(\cdot)$ and $E_{k_2}(\cdot) = E_k^2(\cdot)$. As a matter of fact, the generalized FastPRF method in Section 2.2 is based on the generalized GEDMD

construction which we introduce in Section 2.1 and which we prove to attain at least the same level of security as EDMD.

Unfortunately, the security guarantees of EDMD and GEDMD do not transfer to FastPRF. For one, the same key is used for both permutations. Additionally, the underlying permutations are neither ideal nor independently drawn. Thus, we cannot claim "provable" security of FastPRF designs. One could argue the security of FastPRF using the "prove-then-prune" approach, introduced by Hoang et al. [HKR15] to argue the security of AEZ, but invoking "prove-then-prune" still requires both a solid heuristic argument for the instantiation, as well as cryptanalytic results. We discuss the rationale of FastPRF in Section 2.3.

We see FastPRF as a potentially fruitful concrete object of study and analysis, but also as an opportunity for blockcipher designers—particularly in the lightweight space—to define a PRF along with their designs, in order to widen their applicability. As initial concrete target, we propose AES-PRF: an instantiation of FastPRF based on the AES standard. We introduce and analyze this instantiation in Section 3.

We demonstrate the applicability of our scheme in Section 4, by instantiating GCM and GCM-SIV with it. In more detail, whereas the original GCM [MV04] achieves birthday bound security only, GCM instantiated with a dedicated PRF is optimally secure. Likewise, using a dedicated PRF inside GCM-SIV [GL15, GLL17, LLG17] yields significant security and efficiency improvements, both in the subkey derivation and in the internal evaluation of GCM. In Section 5, we briefly elaborate on the neat and almost immediate extension of our technique to tweakable blockciphers such as SKINNY [BJK+16]. Extending FastPRF to tweakable blockciphers effectively results in a compressing fixed-input-length PRF. The resulting construction can contribute to a removal of the length parameter in security bounds, as we exemplify for PMAC1. The extension to tweakable blockciphers finds further applications in MAC functions based on compressing fixed-input-length PRFs such as Yasuda's [Yas08] construction and NI$^+$ [DNP16].

2 Optimal PRFs from Blockciphers

In this section we describe a general method of transforming an iterative blockcipher into a PRF. We begin with a generalization of Mennink and Neves's EDMD construction [MN17], and demonstrate that it achieves at least the same level of security (Section 2.1). Then, in Section 2.2 we show how to use this construction to design native PRFs. We elaborate on its rationale in Section 2.3.

2.1 Generalized EDMD

Let $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$. Let $d \geq 2$, and define generalized EDMD, or GEDMD$^d : \{0,1\}^{\kappa \cdot d} \times \{0,1\}^n \to \{0,1\}^n$, as

$$\text{GEDMD}^d_{k_1,\ldots,k_d}(x) = (E_{k_d} \circ \cdots \circ E_{k_1})(x) \oplus (E_{k_{d-1}} \circ \cdots \circ E_{k_1})(x) \oplus \ldots \oplus E_{k_1}(x)\,. \quad (8)$$

We will demonstrate that GEDMD^d has at least the same level of security as EDMD = GEDMD^2.

2.1.1 Security Model

Denote by $\mathsf{perm}(n)$ the set of all permutations on $\{0,1\}^n$, and by $\mathsf{func}(m,n)$ the set of all functions from $\{0,1\}^m \to \{0,1\}^n$. For a blockcipher $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$, we denote its PRP security against distinguisher $\mathcal{D}$ by

$$\mathbf{Adv}^{\text{prp}}_E(\mathcal{D}) = \left|\mathbf{Pr}\left[\mathcal{D}^{E_k} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}^{p} \Rightarrow 1\right]\right|, \quad (9)$$

where the probabilities are taken over uniform random drawings $k \xleftarrow{\$} \{0,1\}^\kappa$ and $p \xleftarrow{\$} \mathsf{perm}(n)$. For a function $F : \{0,1\}^\kappa \times \{0,1\}^m \to \{0,1\}^n$, we denote its PRF security against distinguisher $\mathcal{D}$ by

$$\mathbf{Adv}^{\mathrm{prf}}_{F}(\mathcal{D}) = \left|\mathbf{Pr}\left[\mathcal{D}^{F_k} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}^{f} \Rightarrow 1\right]\right| , \tag{10}$$

where the probabilities are taken over uniform random drawings $k \xleftarrow{\$} \{0,1\}^\kappa$ and $f \xleftarrow{\$} \mathsf{func}(m,n)$.

2.1.2 Security of GEDMD

Mennink and Neves [MN17] proved that EDMD is secure up to $2^n/67$ evaluations.

Lemma 1 (EDMD [MN17])**.** *Let $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher. For any distinguisher $\mathcal{D}$ with query complexity at most $q \leq 2^n/67$, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{EDMD}}(\mathcal{D}) \leq q/2^n + \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathcal{D}') + \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathcal{D}'') , \tag{11}$$

for some distinguishers $\mathcal{D}'$ and $\mathcal{D}''$ with the same query and time complexity as $\mathcal{D}$.

One can easily observe that, for any $d \geq 3$, GEDMD^d is at least as secure as GEDMD^{d-1}. As such, the bound of Lemma 1 is inherited by GEDMD^d for any $d \geq 2$.[2]

Theorem 1. *Let $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher. Let $d \geq 2$. For any distinguisher $\mathcal{D}$ with query complexity at most $q \leq 2^n/67$, we have*

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^d}(\mathcal{D}) \leq q/2^n + \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathcal{D}') + \mathbf{Adv}^{\mathrm{prp}}_{E}(\mathcal{D}'') . \tag{12}$$

for some distinguishers $\mathcal{D}'$ and $\mathcal{D}''$ with the same query and time complexity as $\mathcal{D}$.

Proof (Proof). Consider any distinguisher $\mathcal{D}$ whose goal is to distinguish $\mathrm{GEDMD}^d \in \mathsf{func}(n,n)$ from $f \xleftarrow{\$} \mathsf{func}(n,n)$. As a first step, note that

$$\mathrm{GEDMD}^d_{k_1,\ldots,k_d}(x) = (\mathrm{GEDMD}^{d-1}_{k_2,\ldots,k_d} \circ E_{k_1})(x) \oplus E_{k_1}(x) .$$

Let $g \xleftarrow{\$} \mathsf{func}(n,n)$, and define

$$F_{k_1,g}(x) = (g \circ E_{k_1})(x) \oplus E_{k_1}(x) .$$

By a hybrid argument, we have

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^d}(\mathcal{D}) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^{d-1}}(\mathcal{D}) + \left|\mathbf{Pr}\left[\mathcal{D}^{F_{k_1,g}} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}^{f} \Rightarrow 1\right]\right| .$$

However, as E_{k_1} is a permutation, the distributions of $F_{k_1,g}$ and f are identical. Inductive application yields, for any $d \geq 2$,

$$\mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^d}(\mathcal{D}) \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^{d-1}}(\mathcal{D}) \leq \cdots \leq \mathbf{Adv}^{\mathrm{prf}}_{\mathrm{GEDMD}^2}(\mathcal{D}) ,$$

where $\mathrm{GEDMD}^2 = \mathrm{EDMD}$. The proof is completed using Lemma 1. □

[2] Alternatively, in the same vein as how Mennink and Neves reduced the security of EDMD to XoP, one can prove that GEDMD^d is at least as secure as the xor of d permutations XoP^d.

2.2 FastPRF Design

To prevent some classes of attacks (e.g., slide attacks [BW99]), most blockciphers consist of distinct round functions. This might be accomplished by a complex key schedule, by adding round constants, or some other way. In effect, the result is that most blockciphers resemble the following structure, in one way or another:

$$E_k(x) = (E_k^r \circ E_k^{r-1} \circ \cdots \circ E_k^1)(x)\,.$$

While this immediately suggests the use of GEDMD with each state value fed-forward, this is not a good idea: round functions are often too weak and too simple for each of them to realistically resemble a random permutation.

Because round functions are individually weak, blockciphers tend to have a lot of them. Instead of looking at the cipher in terms of round functions, we can see it in terms of *groups* of round functions, as follows:

$$E_k(x) = (E_k^d \circ E_k^{d-1} \circ \cdots \circ E_k^1)(x)\,,$$

where each E_k^i function is comprised of a number of rounds, i.e., $d < r$. We can now define a PRF out of this representation by applying GEDMD:

$$\mathrm{FastPRF}_k(x) = (E_k^d \circ E_k^{d-1} \circ \cdots \circ E_k^1)(x) \oplus (E_k^{d-1} \circ \cdots \circ E_k^1)(x) \oplus \ldots \oplus E_k^1(x)\,. \quad (13)$$

Concrete blockciphers are not ideal permutations, and even though a good blockcipher can be considered as a pseudorandom permutation, we cannot directly apply the results of Section 2.1 to FastPRF: one requires the individual groups of round functions E_k^i to be mutually independent and sufficiently random. In general, however, this is the nature of concrete ciphers. There is a long history of concrete designs, such as Feistel networks or key-alternating ciphers, taking provably-secure structures and instantiating them with weaker—but efficiently computable—round functions.

2.3 Rationale

The goal behind FastPRF of (13) is to achieve an optimal or quasi-optimal PRF at the same cost as a regular blockcipher. We are convinced by the plausibility of this quest, given that there is no complexity-theoretic reason that a PRF should be significantly slower to compute than a PRP. To achieve our goal, we looked at the currently known PRP to PRF conversion methods (cf., Section 1.1), and investigated which of them suits our purposes best.

Truncation necessarily entails a significant slowdown compared to the original blockcipher. To salvage this slowdown, one could, hypothetically, split the blockcipher E_k into two halves, $E_k = E_k^2 \circ E_k^1$, and output the concatenation of the truncation of both:

$$F_k(x) = \mathrm{trunc}(E_k^1(x)) \,\|\, \mathrm{trunc}(E_k^2(x))\,,$$

where trunc truncates by $n/2$ bits. This construction has two strong drawbacks:

1. An attacker has direct access to the output of half the cipher, making the function significantly riskier to use than the original blockcipher;
2. Even assuming that both halves are ideal, one would still be far from obtaining optimal security: distinguishing this construction from random can be done in approximately $2^{3n/4}$ queries [GG16].

The xor of permutations can be used likewise with the same cost as a single E_k evaluation as

$$F_k(x) = E_k^1(x) \oplus E_k^2(x)\,.$$

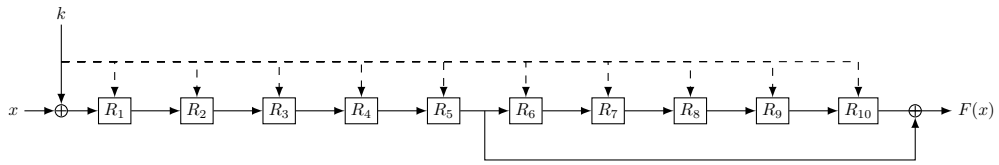

Figure 1: AES-PRF-128.

However, similar to the case of truncation, the attacker has more or less direct access to half the cipher; in particular, many useful properties may survive the xor of two weaker primitives. This, once again, makes this primitive riskier to use than the original cipher.

This risk is technically eliminated using Cogliati and Seurin's EDM construction:

$$F_k(x) = E_k^2(E_k^1(x) \oplus x)\,.$$

Indeed, unlike truncation and xor of permutations, EDM does not expose the results of half the cipher directly to an adversary. However, there is still some risk involved with this construction: the attacker has control over the intermediate state. For example, differential collisions are easy to mount in this construction. If a high-probability differential for E_k^1 exists, it is easy to obtain collisions for $F_k(x)$ and its security crumbles. On the other hand, a high-probability differential for E_k^1 does not necessarily spell doom for the blockcipher E_k itself. Additionally, it is not easy to generalize EDM as in Section 2.1 to multiple intermediate state values and preserve its security proof.

This leaves us with the (generalized) EDMD construction, as

1. it generalizes easily to multiple state values fed-forward;
2. it does not give an attacker control over intermediate state values, beyond the control it has over the input;
3. it does not give an attacker access to half the cipher: instead, the state values that are fed-forward are always masked by a full application of the cipher E_k, effectively randomizing it.

3 Concrete Instantiation: AES-PRF

We present a concrete instantiation of FastPRF based on the AES [DR02]. We adopt the most straightforward choice:

- For 128-bit keys and 10 rounds, we define AES-PRF-128 to be AES xored with the internal state after 5 rounds (cf., Figure 1);
- For 192-bit keys and 12 rounds, we define AES-PRF-192 to be AES xored with the internal state after 6 rounds;
- For 256-bit keys and 14 rounds, we define AES-PRF-256 to be AES xored with the internal state after 7 rounds.

Alternative choices naturally exist. One could, for example, split the 192-bit key case into three 4-round permutations, resulting in an arguably stronger PRF. Note that, by design, each of the proposals consists of one full AES xored with an intermediate state. Throughout, we assume that the full AES is a secure pseudorandom permutation. In addition, we denote the first $t \geq 0$ rounds of AES by AES_t, where the instance of AES (128-, 192-, or 256-bit keys) is usually clear from the context. For convention, we define the special case $\mathrm{AES}_0(x)$ to be x instead of $x \oplus k$.

Table 1: Observed speed, in cycles per byte, of AES-128 in counter mode versus AES-PRF-128 in counter mode.

Microarchitecture	AES-128	AES-PRF-128
Intel Sandy Bridge	0.72	0.75
Intel Haswell	0.63	0.63
Intel Skylake	0.63	0.63
AMD Ryzen	0.37	0.37

3.1 Efficiency

It is easy to verify that AES-PRF is essentially as fast as the AES. The additional overhead is composed of

- An extra xor;
- An additional 128 bits to store the intermediate state(s) to xor the output with.

As far as performance goes, the xor is negligible compared to the cost of the full cipher. The extra state might be more of a problem for implementers in heavily constrained environments, but we argue that in most cases, this is not a problem. For example, in counter mode, we can xor the intermediate state directly with the message block to encrypt, thus eliminating the need for another separate state.

We have implemented AES-PRF-128 in counter mode for some Intel and AMD `x86_64` processors, using the AES-NI instruction set, and its performance is nearly indistinguishable from using the AES directly, as can be verified in Table 1. The only noteworthy case is Sandy Bridge, which is particularly sensitive to instruction scheduling. However, for Sandy Bridge, the overhead of incrementing the counter is far more noticeable—0.72 cycles per byte against the optimal 0.63—than the overhead incurred by using AES-PRF over AES. Of course, the same implementation precautions must be taken with AES-PRF as with AES. In particular, implementations that make use of table lookups are susceptible to cache-timing attacks [TOS10]. In most consumer hardware this is no longer an insurmountable problem, with AES-NI being present in the majority of Intel and AMD processors, and with ARMv8-A, SPARC T4, POWER8, and others also having dedicated constant-time AES instructions available.

3.2 Security Analysis

To assess the concrete security of the AES-PRF construction, it is necessary to dive into the details of the AES. In particular, the 5-round AES is the weakest link in the construction.

There are well-established bounds for the maximum expected differential and linear probabilities of 4-round AES [KMT01, PSLL03, KS07]. Therefore, we do not expect differential or linear attacks to have any meaningful success against AES-PRF.

While there are several attacks that do efficiently break AES reduced to 5 rounds, e.g., [DKR97, BK00, Bir04, Tun12], these attacks appear to be inapplicable here, as the 5-round output is masked by a full AES application. Moreover, most such attacks rely on 3- or 4-round distinguishers, followed by key recovery; direct distinguishers for 5 rounds have only recently been discovered, and have massive data and time requirements [SLG+16, GRR16]. The most promising attack known to date, by Grassi et al. [GRR17], belongs to the subspace trail family of attacks [GRR16], and expects the number of output differences of a certain kind of input difference to be a multiple of 8. In AES-PRF, it is not possible to directly access this information, so it seems unlikely that this kind of distinguisher is feasible in our setting.

More generally, attacks that rely on observing relations between tuples of outputs seem unlikely to be successful—the attacker is only able to observe these relations masked by a strong PRP. Suppose, for the sake of argument, that a high-probability differential $\Delta \xrightarrow{\mathrm{AES}_5} \Delta'$ existed. What the attacker would be able to see, in effect, would be

$$\mathrm{AES}_{10}(x_i) \oplus \mathrm{AES}_{10}(x_i \oplus \Delta) \oplus \mathrm{AES}_5(x_i) \oplus \mathrm{AES}_5(x_i \oplus \Delta)\,.$$

Even if $\mathrm{AES}_5(x_i) \oplus \mathrm{AES}_5(x_i \oplus \Delta)$ were effectively constant, $\mathrm{AES}_{10}(x_i) \oplus \mathrm{AES}_{10}(x_i \oplus \Delta)$ is not, and is in itself essentially indistinguishable from the distribution of differentials of a random function [DR07] (assuming, again, security of the full AES). This masking becomes even more pronounced for higher-order attacks, which are the most dangerous attack class for reduced-round AES.

AES-PRF is also, of course, vulnerable to any attack that does not rely on particular properties of AES_5, but only on the high-level structure of EDMD. The very costly distinguishers enabled by this structure are, for the sake of completeness, described in Appendix A. We conjecture that AES-PRF cannot be distinguished from random significantly faster than by either bruteforcing the key *or* by the generic attacks of Appendix A.

3.3 Unbalanced Variants

To gauge AES-PRF's resistance against attacks, one may consider "reduced-round variants" of the PRF. In our case, we propose unbalanced variants of it: instead of $\text{AES-PRF}(x) = \mathrm{AES}_{10}(x) \oplus \mathrm{AES}_5(x)$, we suggest looking at

$$\text{AES-PRF}_t(x) = \mathrm{AES}_{10}(x) \oplus \mathrm{AES}_t(x)\,,$$

for $t \in \{0, \ldots, 9\}$.

3.3.1 AES-PRF$_0$

AES-PRF_0 is exactly the Davies-Meyer construction. It is known to be no less distinguishable than the underlying blockcipher. In more detail, distinguishing AES-PRF_0 from a random function f is equally hard as distinguishing $\text{AES-PRF}_0 \oplus \mathrm{id}$ from random. This function, however, does not expose collisions, and can be distinguished from random with $2^{n/2}$ queries [CS16].

3.3.2 AES-PRF$_1$

AES-PRF_1 is the first nontrivial case, and the above attack no longer works. However, AES-PRF_1 is still not more secure than AES-PRF_0: one can perform a key-recovery attack with 2^{67} queries.

We may rewrite $\text{AES-PRF}_1(x)$ as

$$\text{AES-PRF}_1(x) = \mathrm{AES}_{10}(x) \oplus P(x \oplus k) \oplus k_1\,,$$

where P is the non-keyed portion of the AES round, i.e., $\mathsf{MixColumns} \circ \mathsf{ShiftRows} \circ \mathsf{SubBytes}$. We rely here in the following property of the AES round:

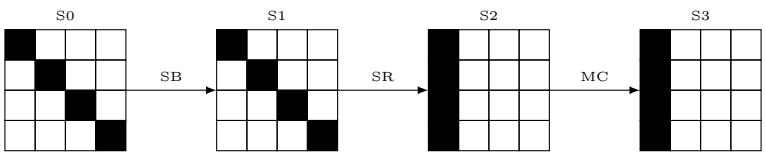

In other words, 4 well-chosen bytes will only affect some other 4 bytes of the output. In particular, input bytes s_0, s_5, s_{10}, s_{15} only affect output bytes s_0, s_1, s_2, s_3; input bytes s_4, s_9, s_{14}, s_{13} only affect output bytes s_4, s_5, s_6, s_7; input bytes s_8, s_{13}, s_2, s_7 only affect output bytes s_8, s_9, s_{10}, s_{11}; and input bytes s_{12}, s_1, s_6, s_{11} only affect output bytes $s_{12}, s_{13}, s_{14}, s_{15}$.

Thus, if we correctly guess 4 bytes of the key on one of those input positions, this will cancel out the contribution of $P(x \oplus k)$ on the corresponding 4 output bytes, in which case we obtain $\mathrm{AES}_{10}(x) \oplus k_1$ verbatim, which can be distinguished from random by its lack of collisions.

This observation leads to the following simple key-recovery attack:

1. Initialize the 16-byte candidate key $k' = 0$;
2. Accumulate a large number $q \approx 2^{67}$ of queries $(x_i, y_i = \text{AES-PRF}_1(x_i))$, yielding approximately $\binom{2^{67}}{2}/2^{128} \approx 32$ collisions among the y_i;
3. Let $(A \to B)$ be any pair of corresponding input-output affected bytes for one round of AES, i.e., $(0, 5, 10, 15) \to (0, 1, 2, 3)$, $(4, 9, 14, 3) \to (4, 5, 6, 7)$, $(8, 13, 2, 7) \to (8, 9, 10, 11)$, and $(12, 1, 6, 11) \to (12, 13, 14, 15)$. Let s_A (resp. s_B) denote the bytes of s at the positions defined by A (resp. B);
4. For each $(A \to B)$ pair:
   (a) For all values a of A:
      i. Compute $S = \{y_i \oplus P(x_i \oplus a)\}$ for all q queries;
      ii. If there are no collisions in S, we either succeeded in finding the correct key or did not collect enough queries;
      iii. If there are collisions in S, but $S_{i\,B} = S_{j\,B}$ for all $i \neq j$, the current value a is likely the correct choice for the value of k'_A. Set $k'_A = a$, and move on to the next $(A \to B)$ pair.

In total, we require approximately 2^{67} queries, 2^{101} computations, and 2^{67} memory. The number of queries is justified by the following criteria:

- The probability that there is no collision after q queries is approximately $e^{-\binom{q}{2}/2^{128}}$. With $q = 2^{67}$, this probability is suitably small: approximately 2^{-46};
- The probability that every collision misses the currently active B is 2^{-32} per collision.

The key observation here is that, under the wrong key randomization hypothesis, only the correct key guess k' results in the set $\{\text{AES-PRF}_1(x_i) \oplus P(x_i \oplus k')\}$ to have no collisions with high probability. We take advantage of the fact that 1 round of AES does not have full diffusion and allows us to guess 32 bits of the key at a time.

This attack may improve its time complexity by noticing that once bytes of k are available, so are the corresponding bytes of $k_1 = \mathrm{AES}_{10}(x) \oplus P(x \oplus k)$. Exploiting the key schedule may accelerate the filtering of incorrect key guesses.

3.3.3 AES-PRF$_2$

The same attack strategy used for AES-PRF$_1$ no longer works for two rounds. Any single byte affects every output byte, so detecting where collisions happen is no longer reliable as a means to verify correct keys. However, we do believe that AES-PRF$_2$ is still within reach of an efficient attack, and leave it as an open problem.

Table 2: Distribution of output values of $x \oplus S(x)$ by number of preimages.

Preimages	Count
1	91
2	54
3	15
4	3

3.3.4 AES-PRF$_9$

We can also look at the other end of unbalanced variants of AES-PRF. The first thing to notice is that the last round does not include MixColumns.[3] This means that AES-PRF$_9$ can be written, for the first row of the state, as

$$\begin{aligned} x_0 \oplus S(x_0) &\oplus k_0 \ , \\ x_4 \oplus S(x_4) &\oplus k_4 \ , \\ x_8 \oplus S(x_8) &\oplus k_8 \ , \\ x_{12} \oplus S(x_{12}) &\oplus k_{12} \, . \end{aligned}$$

By observing the distribution of $x \oplus S(x)$ (see Table 2), we see that only 3 outputs have maximal probability 1/64:

$$\begin{aligned} &\texttt{0x7E} \oplus S(\texttt{0x7E}) = \texttt{0x81} \oplus S(\texttt{0x81}) = \texttt{0xDA} \oplus S(\texttt{0xDA}) = \texttt{0xE4} \oplus S(\texttt{0xE4}) = \texttt{0x8D} \, , \\ &\texttt{0x1D} \oplus S(\texttt{0x1D}) = \texttt{0xC1} \oplus S(\texttt{0xC1}) = \texttt{0xD8} \oplus S(\texttt{0xD8}) = \texttt{0xF8} \oplus S(\texttt{0xF8}) = \texttt{0xB9} \, , \\ &\texttt{0x17} \oplus S(\texttt{0x17}) = \texttt{0x47} \oplus S(\texttt{0x47}) = \texttt{0x56} \oplus S(\texttt{0x56}) = \texttt{0xC2} \oplus S(\texttt{0xC2}) = \texttt{0xE7} \, . \end{aligned}$$

Therefore, by observing the frequency of bytes $0, 4, 8, 12$ of AES-PRF$_9$ for sufficiently many outputs ($q \gg 2^8$), we are able to derive each of k_0, k_4, k_8, k_{12} as one of 3 possibilities with high confidence. For the other rows the same principle does not apply, since their bytes are of the form $x_i \oplus S(x_j)$, whose output is balanced. Nevertheless, recovering the first 32 bits of the key cheaply is still an attack, and the distribution of $x \oplus S(x)$ acts as very efficient distinguisher here.

4 Application to GCM and GCM-SIV

We discuss the security of GCM by McGrew and Viega [MV04] and GCM-SIV by Gueron and Lindell [GL15], in case they are instantiated using FastPRF. The reasoning below directly applies to counter mode encryption, as GCM uses this mode internally.

4.1 Security Model

Formally, an authenticated encryption scheme AE consists of two algorithms Enc, Dec. The encryption algorithm Enc gets as input a key k, a nonce n, associated data ad, a message m, and outputs a ciphertext c, and tag t. The decryption algorithm Dec gets as input a key k, a nonce n, associated data ad, a ciphertext c, and a tag t, and outputs either a message m or a dedicated $\perp$-sign, where $\mathrm{Dec}(k, n, ad, \mathrm{Enc}(k, n, ad, m)) = m$ is required to hold for any (k, n, ad, m).

[3] Consequences of the lack of MixColumns in the last round have been previously studied by Dunkelman and Keller [DK10].

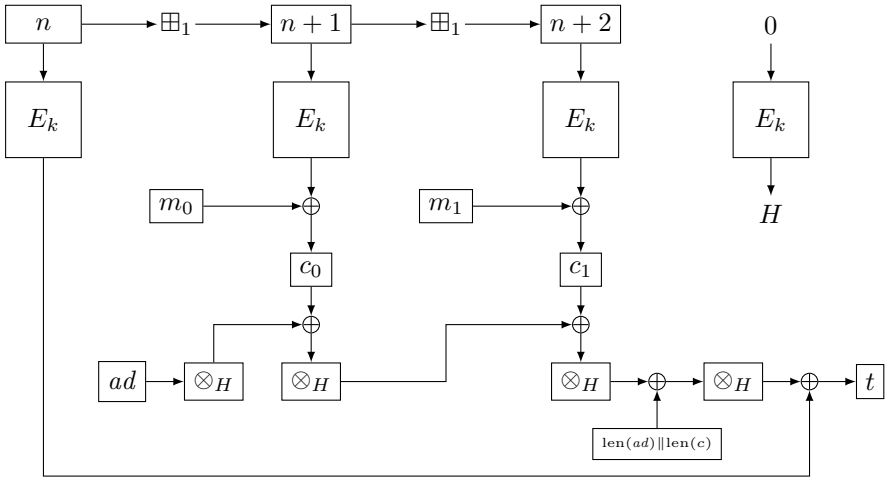

Figure 2: Overview of the GCM mode with a 96-bit nonce, a single block of associated data, and two blocks of plaintext (resp. ciphertext). $\otimes_H$ is multiplication by H in $\mathbb{F}_{2^{128}}$, whereas $\boxplus_1$ is addition modulo 2^{32} of the least significant bytes of the state.

Security of $\mathrm{AE} = (\mathrm{Enc}, \mathrm{Dec})$ is usually measured via its confidentiality and authenticity. We denote the confidentiality of AE against a distinguisher $\mathcal{D}$ by

$$\mathbf{Adv}^{\mathrm{conf}}_{\mathrm{AE}}(\mathcal{D}) = \left|\mathbf{Pr}\left[\mathcal{D}^{\mathrm{Enc}_k} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}^{\$} \Rightarrow 1\right]\right| , \tag{14}$$

where $\$(n, a, m)$ always returns a random $(c, t) \xleftarrow{\$} \{0,1\}^{|m|+\tau}$ (where τ is the tag size), and where the probabilities are taken over uniform random drawings $k \xleftarrow{\$} \{0,1\}^{\kappa}$ and $\$$. The authenticity of AE against a distinguisher $\mathcal{D}$ is denoted by

$$\mathbf{Adv}^{\mathrm{auth}}_{\mathrm{AE}}(\mathcal{D}) = \left|\mathbf{Pr}\left[\mathcal{D}^{\mathrm{Enc}_k,\mathrm{Dec}_k} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}^{\mathrm{Enc}_k,\perp} \Rightarrow 1\right]\right| , \tag{15}$$

where $\perp$ always returns the $\perp$-sign, and where the probabilities are taken over the uniform random drawing $k \xleftarrow{\$} \{0,1\}^{\kappa}$. For authenticity, the distinguisher is not allowed to relay a response from its first oracle to its second oracle. Unless explicitly stated otherwise, we will consider the case where $\mathcal{D}$ is required to be nonce-respecting: it is not allowed to repeat a nonce in an encryption query (it may reuse a nonce in a decryption query).

4.2 AES-PRF-GCM

GCM is an authenticated encryption scheme by McGrew and Viega [MV04]. It internally uses counter mode on top of a blockcipher (see Figure 2).

McGrew and Viega [MV04] and later Iwata et al. [IOM12] proved the following result for GCM with 96-bit nonce. (We express the result in terms of AES as underlying primitive for convenience.)

Theorem 2 (GCM [MV04, IOM12])**.** *Let* $\mathrm{AES} : \{0,1\}^{\kappa} \times \{0,1\}^{n} \to \{0,1\}^{n}$ *be the* AES *blockcipher, and* τ *be the tag length. For any distinguisher* $\mathcal{D}$ *with encryption query complexity at most* q*, decryption query complexity at most* q' *(*$= 0$ *for confidentiality), per-query length at most* ℓ*, and total complexity at most* σ*, we have*

$$\mathbf{Adv}^{\mathrm{conf}}_{\mathrm{GCM}[\mathrm{AES},\tau]}(\mathcal{D}) \leq \binom{q+\sigma+1}{2}/2^{n} + \mathbf{Adv}^{\mathrm{prp}}_{\mathrm{AES}}(\mathcal{D}') , \tag{16}$$

$$\mathbf{Adv}^{\mathrm{auth}}_{\mathrm{GCM}[\mathrm{AES},\tau]}(\mathcal{D}) \leq \frac{q'(\ell+1)}{2^{\tau}} + \binom{q+q'+\sigma+1}{2}/2^{n} + \mathbf{Adv}^{\mathrm{prp}}_{\mathrm{AES}}(\mathcal{D}') , \tag{17}$$

for some distinguisher $\mathcal{D}'$ with the same time complexity as $\mathcal{D}$ and making at most $q + q' + \sigma + 1$ queries.

Proof (Proof (sketch)). We only discuss the high-level structure. The proof consists of three steps: (i) replacing AES by a random permutation $\pi \xleftarrow{\$} \mathsf{perm}(n)$, (ii) subsequently replacing π by a random function $f \xleftarrow{\$} \mathsf{func}(n, n)$, and (iii) analyzing (with slight abuse of notation) GCM$[f, \tau]$. In more detail, the following bound for confidentiality/authenticity follows by a hybrid argument:

$$\begin{aligned}\mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[\text{AES},\tau]}(\mathcal{D}) &\leq \mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[\pi,\tau]}(\mathcal{D}) + \mathbf{Adv}^{\text{prp}}_{\text{AES}}(\mathcal{D}') \\ &\leq \mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[f,\tau]}(\mathcal{D}) + \left|\mathbf{Pr}\left[\mathcal{D}''^{f} \Rightarrow 1\right] - \mathbf{Pr}\left[\mathcal{D}''^{\pi} \Rightarrow 1\right]\right| + \mathbf{Adv}^{\text{prp}}_{\text{AES}}(\mathcal{D}') \\ &\leq \mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[f,\tau]}(\mathcal{D}) + \binom{q + q' + \sigma + 1}{2}/2^n + \mathbf{Adv}^{\text{prp}}_{\text{AES}}(\mathcal{D}')\,, \qquad (18)\end{aligned}$$

for some distinguishers $\mathcal{D}', \mathcal{D}''$ that make at most $q + q' + \sigma + 1$ queries. The core part in the analysis of GCM centers around the analysis of GCM based on a random function f, and McGrew and Viega [MV04] and later Iwata et al. [IOM12] proved that

$$\begin{aligned}\mathbf{Adv}^{\text{conf}}_{\text{GCM}[f,\tau]}(\mathcal{D}) &= 0\,, \\ \mathbf{Adv}^{\text{auth}}_{\text{GCM}[f,\tau]}(\mathcal{D}) &\leq \frac{q'(\ell + 1)}{2^\tau}\,,\end{aligned}$$

which completes the proof. □

From high-level inspection of the security analysis of GCM, it becomes clear that FastPRF can be used to improve GCM's security significantly. In more detail, if we use AES-PRF instead of AES, steps (i) and (ii) in the proof merge and become "replacing AES-PRF by a random function $f \xleftarrow{\$} \mathsf{func}(n, n)$." In other words, we simply get

$$\mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[\text{AES-PRF},\tau]}(\mathcal{D}) \leq \mathbf{Adv}^{\text{conf/auth}}_{\text{GCM}[f,\tau]}(\mathcal{D}) + \mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}')$$

instead of (18), and we obtain the following corollary.

Corollary 1. *Let* AES-PRF : $\{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ *be the* AES-PRF *construction of Section 3, and τ be the tag length. For any distinguisher $\mathcal{D}$ with encryption query complexity at most q, decryption query complexity at most q' ($= 0$ for confidentiality), per-query length at most ℓ, and total complexity at most σ, we have*

$$\mathbf{Adv}^{\text{conf}}_{\text{GCM}[\text{AES-PRF},\tau]}(\mathcal{D}) \leq \mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}')\,, \qquad (19)$$

$$\mathbf{Adv}^{\text{auth}}_{\text{GCM}[\text{AES-PRF},\tau]}(\mathcal{D}) \leq \frac{q'(\ell + 1)}{2^\tau} + \mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}')\,, \qquad (20)$$

for some distinguisher $\mathcal{D}'$ with the same time complexity as $\mathcal{D}$ and making at most $q + q' + \sigma + 1$ queries.

4.3 AES-PRF-GCM-SIV

GCM is notoriously sensitive to nonce repeats, which lead to forgeries and even key recovery [Jou06, BZD+16]. GCM-SIV [GL15, GLL17, LLG17] is an authenticated encryption mode based on GCM that aims to be more robust to such usage failures. In particular, GCM-SIV aims for a slightly different security notion than GCM—misuse-resistant authenticated encryption, or mrAE. This notion comprises (14) and (15), with the exception that the requirement that nonces are unique is lifted.

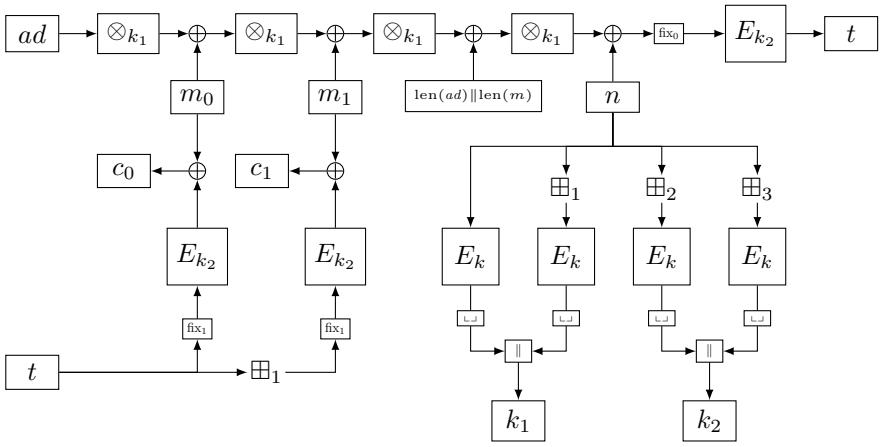

Figure 3: A high-level overview of GCM-SIV with a 128-bit key, one block of associated data, and two blocks of plaintext (resp. ciphertext). Notation matches that of Figure 2. fix_0 sets the most significant bit of a block to 0; fix_1 sets it to 1. $\llcorner\lrcorner$ indicates truncation of the first 64 bits; $\|$ denotes concatenation.

There are several variants of GCM-SIV [GL15, IM16, GLL17]. In this work, we consider the most recent one, [GLL17], which is also being considered as an IETF RFC [LLG17]. It is based on the SIV [RS06] mode, and reuses the individual components of GCM (see Figure 3). The basic GCM-SIV construction uses two keys, one of size n bits, one of size at most $2n$ bits, and GCM-SIV of [GLL17] uses the `DeriveKey` mechanism[4] to derive two subkeys from a single one in the following way:

$$\begin{aligned} k_1 &= \text{trunc}(\text{AES}_k(n\|0)) \,\|\, \text{trunc}(\text{AES}_k(n\|1))\,, \\ k_2 &= \text{trunc}(\text{AES}_k(n\|2)) \,\|\, \cdots \,\|\, \text{trunc}(\text{AES}_k(n\|5))\,, \end{aligned} \tag{21}$$

where trunc truncates by $n/2$ bits (recall the truncation construction of Section 1.1).

An earlier security analysis of GCM-SIV was performed by Gueron et al. [GL15, GLL17]. Iwata and Seurin [IS17] pointed out several shortcomings in the analysis, and performed an improved analysis.

Theorem 3 (GCM-SIV [IS17])**.** *Let* $\text{AES} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ *be the* AES *blockcipher. For any distinguisher* $\mathcal{D}$ *that can make encryption queries for at most* q_u *distinct nonces and at most* r *repeats per nonce, and that can make* q_D *decryption queries with total complexity at most* σ_D*, all with per-query associated data and message length at most* ℓ_a *and* ℓ_m*, we have*

$$\mathbf{Adv}^{\text{mrAE}}_{\text{GCM-SIV[AES]}}(\mathcal{D}) \le \mathbf{Adv}^{\text{prp}}_{\text{AES}}(\mathcal{D}'') + \min\left(\frac{\binom{6(q_u+q_D)}{2}}{2^n}, \frac{6(q_u+q_D)}{2^{3n/4}}\right) + \tag{22}$$

$$(q_u+q_D)\cdot\mathbf{Adv}^{\text{prf}}_{\text{AES}}(\mathcal{D}') + \tag{23}$$

$$\frac{(q_u+q_D)^2}{2^{\kappa+1}} + \frac{q_u r^2(\ell_a+3\ell_m+3)}{2^n} + \frac{2rq_D(\ell_a+\ell_m+1)}{2^n} + \frac{q_D}{2^n}\,, \tag{24}$$

for some distinguishers $\mathcal{D}'$ *making at most* $r(\ell_m+1)+1+\sigma_D$ *queries and* $\mathcal{D}''$ *making at most* $6(q_u+q_D)$ *queries (both with the same time complexity as* $\mathcal{D}$*).*

[4] `DeriveKey` was introduced to GCM-SIV after analysis revealed that the original key derivation mechanism was weaker than expected [NSA17].

Part (22) of the bound comes from how well the `DeriveKey` functionality behaves like random, part (23) reflects the security of AES used in GCM-SIV based on two uniformly randomly generated subkeys, and (24) reflects the security of GCM-SIV based on uniformly randomly generated primitives. We remark that the bound of Iwata and Seurin is slightly stronger, having the multi-user PRF security $\mathbf{Adv}^{\text{mu-prf}}_{\text{AES}}(\mathcal{D}')$ instead of (23), where now distinguisher $\mathcal{D}'$ makes at most $r(\ell_m + 1)$ queries for at most q_u distinct users and an additional amount of $q_D + \sigma_D$ queries freely distributed over all users. We have adopted the slightly simplified bound.

Suppose that, instead of (21), the key is derived using AES-PRF:

$$\begin{aligned} k_1 &= \text{AES-PRF}(n\|0)\,, \\ k_2 &= \text{AES-PRF}(n\|1)\|\text{AES-PRF}(n\|2)\,, \end{aligned} \tag{25}$$

then this subkey derivation function inherits the PRF security of AES-PRF against a distinguisher $\mathcal{D}''$ making at most $3(q_u+q_D)$ queries. In other words, part (22) of Theorem 3 becomes $\mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}'')$, for some distinguisher $\mathcal{D}''$ making at most $3(q_u + q_D)$ queries. In addition, (23) has a term that measures the PRF security of AES, which is at best $\binom{r(\ell_m+1)+1+\sigma_D}{2}/2^n$ due to the PRP-PRF switch. By directly using AES-PRF instead of AES, this implicit birthday term gets eliminated. We thus obtain the following corollary.

Corollary 2. *Let* AES-PRF : $\{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ *be the* AES-PRF *construction of Section 3. For any distinguisher* $\mathcal{D}$ *that can make encryption queries for at most* q_u *distinct nonces and at most* r *repeats per nonce, and that can make* q_D *decryption queries with total complexity at most* σ_D*, all with per-query associated data and message length at most* ℓ_a *and* ℓ_m*, we have*

$$\begin{aligned} \mathbf{Adv}^{\text{mrAE}}_{\text{GCM-SIV[AES-PRF]}}(\mathcal{D}) \le\ & \mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}'') + (q_u + q_D) \cdot \mathbf{Adv}^{\text{prf}}_{\text{AES-PRF}}(\mathcal{D}') \\ & + \frac{(q_u + q_D)^2}{2^{\kappa+1}} + \frac{q_u r^2(\ell_a + 3\ell_m + 3)}{2^n} + \frac{2rq_D(\ell_a + \ell_m + 1)}{2^n} + \frac{q_D}{2^n}\,, \end{aligned} \tag{26}$$

for some distinguishers $\mathcal{D}'$ *making at most* $r(\ell_m + 1) + 1 + \sigma_D$ *queries and* $\mathcal{D}''$ *making at most* $3(q_u + q_D)$ *queries (both with the same time complexity as* $\mathcal{D}$*).*

5 Extension to Tweakable Blockciphers

Tweakable blockciphers are a relatively recent invention formalized by Liskov et al. [LRW02, LRW11]. A tweakable blockcipher, as the name implies, is a blockcipher that takes one additional input beyond the message and key—a tweak. The security of a tweakable blockcipher is then defined as the indistinguishability of the construction against a collection of random permutations, one per each key and tweak.

The FastPRF construction can be generalized to such designs as well, though more care is necessary to ensure that the tweak and the key contribute to each of the individual permutations. More detailed, if $\widetilde{E}$ is a tweakable blockcipher which can be partitioned into groups of round functions as follows:

$$\widetilde{E}_{k,t}(x) = (\widetilde{E}^d_{k,t} \circ \widetilde{E}^{d-1}_{k,t} \circ \cdots \circ \widetilde{E}^1_{k,t})(x)\,,$$

we can define a compressing PRF out of this representation by applying GEDMD:

$$\widetilde{\text{FastPRF}}_k(t,x) = (\widetilde{E}^d_{k,t} \circ \widetilde{E}^{d-1}_{k,t} \circ \cdots \circ \widetilde{E}^1_{k,t})(x) \oplus (\widetilde{E}^{d-1}_{k,t} \circ \cdots \circ \widetilde{E}^1_{k,t})(x) \oplus \ldots \oplus \widetilde{E}^1_{k,t}(x)\,. \tag{27}$$

Once again: this construction only works if the groups of rounds are sufficiently strong individually, but in addition, we require that each of the rounds is sufficiently dependent on t.

Unlike the case of blockciphers, there do not exist many native designs of tweakable blockciphers to draw from. Indeed, most tweakable blockciphers in current use are in fact generic blockcipher-based constructions with far from optimal security (e.g., Rogaway's XEX construction [Rog04] is used in OCB2, a large number of CAESAR submissions, and XTS disk encryption). Some particular examples of dedicated tweakable blockcipher designs to apply $\widetilde{\text{FastPRF}}$ on are Threefish [FLS+10], SCREAM [GLS+15], or the more general TWEAKEY [JNP14] framework which is for instance adopted by the developers of SKINNY [BJK+16]. One may for example take SKINNY-128-256 as tweakable blockcipher with a 128-bit state and 256-bit tweakey: it consists of 48 rounds, and $\widetilde{\text{SKINNY-PRF-128}}$ can be defined to be SKINNY-128-256 xored with the internal state after 24 rounds. Like the AES, SKINNY has solid design principles; 6 rounds are already sufficient for full diffusion, and 24 rounds are already sufficient to withstand several classes of attacks. However, there has not yet been enough cryptanalytic research on SKINNY to confidently claim that the resulting PRF is heuristically secure.

Assuming the existence of $\widetilde{\text{FastPRF}}$, one could use this construction instead of existing tweakable blockciphers in settings where the tweakable blockcipher is not evaluated in inverse direction. For example, consider PMAC1 from Rogaway [Rog04], the tweakable blockcipher based variant of PMAC. Rogaway proved the following result on the PRF security of PMAC1 with tag length n (the definition of Section 2.1.1 generalizes to variable input sizes).

Theorem 4 (PMAC1 [Rog04])**.** *Let $\widetilde{E} : \{0,1\}^\kappa \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be a tweakable blockcipher. For any distinguisher $\mathcal{D}$ with encryption query complexity at most q and total complexity at most σ, we have*

$$\mathbf{Adv}^{\text{prf}}_{\text{PMAC1}[\widetilde{E}]}(\mathcal{D}) \leq \binom{q}{2}/2^n + \binom{\sigma}{2}/2^n + \mathbf{Adv}^{\text{tprp}}_{\widetilde{E}}(\mathcal{D}')\,, \tag{28}$$

for some distinguisher $\mathcal{D}'$ with the same time complexity as $\mathcal{D}$ and making at most σ queries.

Closer inspection of the security analysis reveals that $\binom{\sigma}{2}/2^n$ comes from viewing $\widetilde{E}$ as a random function (one could call this a TPRP-TPRF-switch, although a tweakable PRF is just a compressing fixed-input-length PRF). Following a similar reasoning as in Section 4, one can observe that directly using $\widetilde{\text{FastPRF}}$ in PMAC1 yields the following corollary.

Corollary 3. *Let $\widetilde{\text{FastPRF}} : \{0,1\}^\kappa \times \{0,1\}^\tau \times \{0,1\}^n \to \{0,1\}^n$ be the construction of (27). For any distinguisher $\mathcal{D}$ with encryption query complexity at most q and total complexity at most σ, we have*

$$\mathbf{Adv}^{\text{prf}}_{\text{PMAC1}[\widetilde{\text{FastPRF}}]}(\mathcal{D}) \leq \binom{q}{2}/2^n + \mathbf{Adv}^{\text{prf}}_{\widetilde{\text{FastPRF}}}(\mathcal{D}')\,, \tag{29}$$

for some distinguisher $\mathcal{D}'$ with the same time complexity as $\mathcal{D}$ and making at most σ queries.

In other words, unlike for the original PMAC1, the security bound of PMAC1 based in $\widetilde{\text{FastPRF}}$ does not admit a quadratic security loss on σ, provided $\widetilde{\text{FastPRF}}$ is in turn built on a dedicated tweakable blockcipher.

Acknowledgments

Bart Mennink is supported by a postdoctoral fellowship from the Netherlands Organisation for Scientific Research (NWO) under Veni grant 016.Veni.173.017. The authors are thankful to the anonymous reviewers of FSE 2018 for their useful technical comments, to Pierre Karpman and Atul Luykx for preliminary discussions, to Tetsu Iwata and Yannick Seurin for discussion regarding [IS17], and to Dan Bernstein for discussion regarding SURF [Ber97].

References

[BBS99] Eli Biham, Alex Biryukov, and Adi Shamir. Cryptanalysis of Skipjack reduced to 31 rounds using impossible differentials. In Jacques Stern, editor, *Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding*, volume 1592 of *Lecture Notes in Computer Science*, pages 12–23. Springer, 1999.

[BDJR97] Mihir Bellare, Anand Desai, E. Jokipii, and Phillip Rogaway. A concrete security treatment of symmetric encryption. In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 394–403. IEEE Computer Society, 1997.

[BDK01] Eli Biham, Orr Dunkelman, and Nathan Keller. The rectangle attack - rectangling the serpent. In Pfitzmann [Pfi01], pages 340–357.

[Ber97] Daniel J. Bernstein. SURF: simple unpredictable random function. https://cr.yp.to/papers.html#surf, April 1997.

[Ber05] Daniel J. Bernstein. Stronger security bounds for Wegman-Carter-Shoup authenticators. In Ronald Cramer, editor, *Advances in Cryptology - EUROCRYPT 2005, 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Aarhus, Denmark, May 22-26, 2005, Proceedings*, volume 3494 of *Lecture Notes in Computer Science*, pages 164–180. Springer, 2005.

[BI99] M. Bellare and R. Impagliazzo. A tool for obtaining tighter security analyses of pseudorandom function based constructions, with applications to PRP to PRF conversion. Cryptology ePrint Archive, Report 1999/024, 1999. http://eprint.iacr.org/1999/024.

[Bih97] Eli Biham, editor. *Fast Software Encryption, 4th International Workshop, FSE '97, Haifa, Israel, January 20-22, 1997, Proceedings*, volume 1267 of *Lecture Notes in Computer Science*. Springer, 1997.

[Bir04] Alex Biryukov. The boomerang attack on 5 and 6-round reduced AES. In Hans Dobbertin, Vincent Rijmen, and Aleksandra Sowa, editors, *Advanced Encryption Standard - AES, 4th International Conference, AES 2004, Bonn, Germany, May 10-12, 2004, Revised Selected and Invited Papers*, volume 3373 of *Lecture Notes in Computer Science*, pages 11–15. Springer, 2004.

[BJK+16] Christof Beierle, Jérémy Jean, Stefan Kölbl, Gregor Leander, Amir Moradi, Thomas Peyrin, Yu Sasaki, Pascal Sasdrich, and Siang Meng Sim. The SKINNY family of block ciphers and its low-latency variant MANTIS. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology - CRYPTO*

2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II, volume 9815 of *Lecture Notes in Computer Science*, pages 123–153. Springer, 2016.

[BK00] Eli Biham and Nathan Keller. Cryptanalysis of reduced variants of Rijndael. In *Third AES Candidate Conference (AES3)*, 2000.

[BKR94] Mihir Bellare, Joe Kilian, and Phillip Rogaway. The security of cipher block chaining. In Desmedt [Des94], pages 341–358.

[BKR98] Mihir Bellare, Ted Krovetz, and Phillip Rogaway. Luby-Rackoff backwards: Increasing security by making block ciphers non-invertible. In Kaisa Nyberg, editor, *Advances in Cryptology - EUROCRYPT '98, International Conference on the Theory and Application of Cryptographic Techniques, Espoo, Finland, May 31 - June 4, 1998, Proceeding*, volume 1403 of *Lecture Notes in Computer Science*, pages 266–280. Springer, 1998.

[BL16] Karthikeyan Bhargavan and Gaëtan Leurent. On the practical (in-)security of 64-bit block ciphers: Collision attacks on HTTP over TLS and OpenVPN. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 456–467. ACM, 2016.

[BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Vaudenay [Vau06], pages 409–426.

[Bra82] Gilles Brassard. On computationally secure authentication tags requiring short secret shared keys. In David Chaum, Ronald L. Rivest, and Alan T. Sherman, editors, *Advances in Cryptology: Proceedings of CRYPTO '82, Santa Barbara, California, USA, August 23-25, 1982.*, pages 79–86. Plenum Press, New York, 1982.

[BS93] Eli Biham and Adi Shamir. *Differential Cryptanalysis of the Data Encryption Standard*. Springer, 1993.

[BSS+13] Ray Beaulieu, Douglas Shors, Jason Smith, Stefan Treatman-Clark, Bryan Weeks, and Louis Wingers. The SIMON and SPECK families of lightweight block ciphers. Cryptology ePrint Archive, Report 2013/404, 2013. http://eprint.iacr.org/2013/404.

[BW99] Alex Biryukov and David Wagner. Slide attacks. In Knudsen [Knu99], pages 245–259.

[BZD+16] Hanno Böck, Aaron Zauner, Sean Devlin, Juraj Somorovsky, and Philipp Jovanovic. Nonce-disrespecting adversaries: Practical forgery attacks on GCM in TLS. In *10th USENIX Workshop on Offensive Technologies, WOOT 16, Austin, TX, August 8-9, 2016*. USENIX Association, 2016.

[CDK09] Christophe De Cannière, Orr Dunkelman, and Miroslav Knezevic. KATAN and KTANTAN - A family of small and efficient hardware-oriented block ciphers. In Christophe Clavier and Kris Gaj, editors, *Cryptographic Hardware and Embedded Systems - CHES 2009, 11th International Workshop, Lausanne, Switzerland, September 6-9, 2009, Proceedings*, volume 5747 of *Lecture Notes in Computer Science*, pages 272–288. Springer, 2009.

[CLP14] Benoit Cogliati, Rodolphe Lampe, and Jacques Patarin. The indistinguishability of the XOR of k permutations. In Cid and Rechberger [CR15], pages 285–302.

[CN08] Donghoon Chang and Mridul Nandi. A short proof of the PRP/PRF switching lemma. Cryptology ePrint Archive, Report 2008/078, 2008. http://eprint.iacr.org/2008/078.

[CR15] Carlos Cid and Christian Rechberger, editors. *Fast Software Encryption - 21st International Workshop, FSE 2014, London, UK, March 3-5, 2014. Revised Selected Papers*, volume 8540 of *Lecture Notes in Computer Science*. Springer, 2015.

[CS16] Benoît Cogliati and Yannick Seurin. EWCDM: an efficient, beyond-birthday secure, nonce-misuse resistant MAC. In Robshaw and Katz [RK16], pages 121–149.

[Dae16] Joan Daemen. Spectral characterization of iterating lossy mappings. In Claude Carlet, M. Anwar Hasan, and Vishal Saraswat, editors, *Security, Privacy, and Applied Cryptography Engineering - 6th International Conference, SPACE 2016, Hyderabad, India, December 14-18, 2016, Proceedings*, volume 10076 of *Lecture Notes in Computer Science*, pages 159–178. Springer, 2016.

[Des94] Yvo Desmedt, editor. *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*. Springer, 1994.

[DK10] Orr Dunkelman and Nathan Keller. The effects of the omission of last round's MixColumns on AES. *Inf. Process. Lett.*, 110(8-9):304–308, 2010.

[DKR97] Joan Daemen, Lars R. Knudsen, and Vincent Rijmen. The block cipher Square. In Biham [Bih97], pages 149–165.

[DNP16] Avijit Dutta, Mridul Nandi, and Goutam Paul. One-key compression function based MAC with security beyond birthday bound. In Joseph K. Liu and Ron Steinfeld, editors, *Information Security and Privacy - 21st Australasian Conference, ACISP 2016, Melbourne, VIC, Australia, July 4-6, 2016, Proceedings, Part I*, volume 9722 of *Lecture Notes in Computer Science*, pages 343–358. Springer, 2016.

[DPU+16] Daniel Dinu, Léo Perrin, Aleksei Udovenko, Vesselin Velichkov, Johann Großschädl, and Alex Biryukov. Design strategies for ARX with provable bounds: Sparx and LAX. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *Advances in Cryptology - ASIACRYPT 2016 - 22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4-8, 2016, Proceedings, Part I*, volume 10031 of *Lecture Notes in Computer Science*, pages 484–513, 2016.

[DR01] Joan Daemen and Vincent Rijmen. The wide trail design strategy. In Bahram Honary, editor, *Cryptography and Coding, 8th IMA International Conference, Cirencester, UK, December 17-19, 2001, Proceedings*, volume 2260 of *Lecture Notes in Computer Science*, pages 222–238. Springer, 2001.

[DR02] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Information Security and Cryptography. Springer, 2002.

[DR07] Joan Daemen and Vincent Rijmen. Probability distributions of correlation and differentials in block ciphers. *J. Mathematical Cryptology*, 1(3):221–242, 2007.

[Ebe17] Sean Eberhard. More on additive triples of bijections. *CoRR*, abs/1704.02407, 2017.

[FLS+10] Niels Ferguson, Stefan Lucks, Bruce Schneier, Doug Whiting, Mihir Bellare, Tadayoshi Kohno, Jon Callas, and Jesse Walker. The Skein hash function family. https://www.schneier.com/academic/paperfiles/skein1.3.pdf, 2010.

[Fre77] David Freedman. A remark on the difference between sampling with and without replacement. *Journal of the American Statistical Association*, 72(359):681–681, 1977.

[GG16] Shoni Gilboa and Shay Gueron. The advantage of truncated permutations. *CoRR*, abs/1610.02518, 2016.

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.

[GL15] Shay Gueron and Yehuda Lindell. GCM-SIV: full nonce misuse-resistant authenticated encryption at under one cycle per byte. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12-6, 2015*, pages 109–119. ACM, 2015.

[GLL17] Shay Gueron, Adam Langley, and Yehuda Lindell. AES-GCM-SIV: specification and analysis. Cryptology ePrint Archive, Report 2017/168, 2017. http://eprint.iacr.org/2017/168.

[GLS+15] Vincent Grosso, Gaëtan Leurent, François-Xavier Standaert, Kerem Varici, Anthony Journault, François Durvaux, Lubos Gaspar, and Stéphanie Kerckhof. SCREAM: side-channel resistant authenticated encryption with masking. https://perso.uclouvain.be/fstandae/SCREAM/, 2015.

[Gol96] Dieter Gollmann, editor. *Fast Software Encryption, Third International Workshop, Cambridge, UK, February 21-23, 1996, Proceedings*, volume 1039 of *Lecture Notes in Computer Science*. Springer, 1996.

[GRR16] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. Subspace trail cryptanalysis and its applications to AES. *IACR Trans. Symmetric Cryptol.*, 2016(2):192–225, 2016.

[GRR17] Lorenzo Grassi, Christian Rechberger, and Sondre Rønjom. A new structural-differential property of 5-round AES. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 289–317, 2017.

[HKR15] Viet Tung Hoang, Ted Krovetz, and Phillip Rogaway. Robust authenticated-encryption AEZ and the problem that it solves. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 15–44. Springer, 2015.

[HT96] Howard M. Heys and Stafford E. Tavares. Substitution-permutation networks resistant to differential and linear cryptanalysis. *J. Cryptology*, 9(1):1–19, 1996.

[HWKS98] Chris Hall, David Wagner, John Kelsey, and Bruce Schneier. Building PRFs from PRPs. In Hugo Krawczyk, editor, *Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings*, volume 1462 of *Lecture Notes in Computer Science*, pages 370–389. Springer, 1998.

[IM16] Tetsu Iwata and Kazuhiko Minematsu. Stronger security variants of GCM-SIV. *IACR Trans. Symmetric Cryptol.*, 2016(1):134–157, 2016.

[IMV16] Tetsu Iwata, Bart Mennink, and Damian Vizár. CENC is optimally secure. Cryptology ePrint Archive, Report 2016/1087, 2016. http://eprint.iacr.org/2016/1087.

[IOM12] Tetsu Iwata, Keisuke Ohashi, and Kazuhiko Minematsu. Breaking and repairing GCM security proofs. In Reihaneh Safavi-Naini and Ran Canetti, editors, *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*, pages 31–49. Springer, 2012.

[IR88] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In Shafi Goldwasser, editor, *Advances in Cryptology - CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings*, volume 403 of *Lecture Notes in Computer Science*, pages 8–26. Springer, 1988.

[IS17] Tetsu Iwata and Yannick Seurin. Reconsidering the security bound of AES-GCM-SIV. Cryptology ePrint Archive, Report 2017/708, 2017. http://eprint.iacr.org/2017/708.

[Iwa06] Tetsu Iwata. New blockcipher modes of operation with beyond the birthday bound security. In Matthew J. B. Robshaw, editor, *Fast Software Encryption, 13th International Workshop, FSE 2006, Graz, Austria, March 15-17, 2006, Revised Selected Papers*, volume 4047 of *Lecture Notes in Computer Science*, pages 310–327. Springer, 2006.

[JK97] Thomas Jakobsen and Lars R. Knudsen. The interpolation attack on block ciphers. In Biham [Bih97], pages 28–40.

[JNP14] Jérémy Jean, Ivica Nikolic, and Thomas Peyrin. Tweaks and keys for block ciphers: The TWEAKEY framework. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014, Proceedings, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 274–288. Springer, 2014.

[Jou06] Antoine Joux. Authentication failures in NIST version of GCM. http://csrc.nist.gov/groups/ST/toolkit/BCM/documents/Joux_comments.pdf, 2006.

[KMT01] Liam Keliher, Henk Meijer, and Stafford E. Tavares. New method for upper bounding the maximum average linear hull probability for SPNs. In Pfitzmann [Pfi01], pages 420–436.

[Knu94] Lars R. Knudsen. Truncated and higher order differentials. In Bart Preneel, editor, *Fast Software Encryption: Second International Workshop. Leuven, Belgium, 14-16 December 1994, Proceedings*, volume 1008 of *Lecture Notes in Computer Science*, pages 196–211. Springer, 1994.

[Knu99] Lars R. Knudsen, editor. *Fast Software Encryption, 6th International Workshop, FSE '99, Rome, Italy, March 24-26, 1999, Proceedings*, volume 1636 of *Lecture Notes in Computer Science*. Springer, 1999.

[KR11] Lars R. Knudsen and Matthew Robshaw. *The Block Cipher Companion*. Information Security and Cryptography. Springer, 2011.

[KS07] Liam Keliher and Jiayuan Sui. Exact maximum expected differential and linear probability for two-round advanced encryption standard. *IET Information Security*, 1(2):53–57, 2007.

[LH94] Susan K. Langford and Martin E. Hellman. Differential-linear cryptanalysis. In Desmedt [Des94], pages 17–25.

[LLG17] Yehuda Lindell, Adam Langley, and Shay Gueron. AES-GCM-SIV: Nonce Misuse-Resistant Authenticated Encryption. Internet-Draft draft-irtf-cfrg-gcmsiv-05, Internet Engineering Task Force, May 2017. Work in Progress.

[LRW02] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In Moti Yung, editor, *Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2002.

[LRW11] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. *J. Cryptology*, 24(3):588–613, 2011.

[Luc00] Stefan Lucks. The sum of PRPs is a secure PRF. In Bart Preneel, editor, *Advances in Cryptology - EUROCRYPT 2000, International Conference on the Theory and Application of Cryptographic Techniques, Bruges, Belgium, May 14-18, 2000, Proceeding*, volume 1807 of *Lecture Notes in Computer Science*, pages 470–484. Springer, 2000.

[Mat93] Mitsuru Matsui. Linear cryptanalysis method for DES cipher. In Tor Helleseth, editor, *Advances in Cryptology - EUROCRYPT '93, Workshop on the Theory and Application of of Cryptographic Techniques, Lofthus, Norway, May 23-27, 1993, Proceedings*, volume 765 of *Lecture Notes in Computer Science*, pages 386–397. Springer, 1993.

[Mat96] Mitsuru Matsui. New structure of block ciphers with provable security against differential and linear cryptanalysis. In Gollmann [Gol96], pages 205–218.

[McG12] David McGrew. Impossible plaintext cryptanalysis and probable-plaintext collision attacks of 64-bit block cipher modes. Cryptology ePrint Archive, Report 2012/623, 2012. http://eprint.iacr.org/2012/623.

[MN17] Bart Mennink and Samuel Neves. Encrypted Davies-Meyer and its dual: Towards optimal security using mirror theory. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 556–583. Springer, 2017.

[MP15] Bart Mennink and Bart Preneel. On the XOR of multiple random permutations. In Tal Malkin, Vladimir Kolesnikov, Allison Bishop Lewko, and Michalis Polychronakis, editors, *Applied Cryptography and Network Security - 13th International Conference, ACNS 2015, New York, NY, USA, June 2-5, 2015, Revised Selected Papers*, volume 9092 of *Lecture Notes in Computer Science*, pages 619–634. Springer, 2015.

[MV04] David A. McGrew and John Viega. The security and performance of the Galois/counter mode (GCM) of operation. In Anne Canteaut and Kapalee Viswanathan, editors, *Progress in Cryptology - INDOCRYPT 2004, 5th International Conference on Cryptology in India, Chennai, India, December 20-22, 2004, Proceedings*, volume 3348 of *Lecture Notes in Computer Science*, pages 343–355. Springer, 2004.

[NSA17] NSA IA. Key recovery attacks on AES-GCM-SIV. https://mailarchive.ietf.org/arch/msg/cfrg/k2mpWgod4mbdOxsvN6EtXHb0BAg, 2017.

[Pat08] Jacques Patarin. A proof of security in $O(2^n)$ for the xor of two random permutations. In Reihaneh Safavi-Naini, editor, *Information Theoretic Security, Third International Conference, ICITS 2008, Calgary, Canada, August 10-13, 2008, Proceedings*, volume 5155 of *Lecture Notes in Computer Science*, pages 232–248. Springer, 2008.

[Pat10] Jacques Patarin. Introduction to mirror theory: Analysis of systems of linear equalities and linear non equalities for cryptography. Cryptology ePrint Archive, Report 2010/287, 2010. http://eprint.iacr.org/2010/287.

[Pat13a] Jacques Patarin. Generic attacks for the xor of k random permutations. In Michael J. Jacobson Jr., Michael E. Locasto, Payman Mohassel, and Reihaneh Safavi-Naini, editors, *Applied Cryptography and Network Security - 11th International Conference, ACNS 2013, Banff, AB, Canada, June 25-28, 2013. Proceedings*, volume 7954 of *Lecture Notes in Computer Science*, pages 154–169. Springer, 2013.

[Pat13b] Jacques Patarin. Security in $O(2^n)$ for the xor of two random permutations – proof with the standard H technique–. Cryptology ePrint Archive, Report 2013/368, 2013. http://eprint.iacr.org/2013/368.

[Pfi01] Birgit Pfitzmann, editor. *Advances in Cryptology - EUROCRYPT 2001, International Conference on the Theory and Application of Cryptographic Techniques, Innsbruck, Austria, May 6-10, 2001, Proceeding*, volume 2045 of *Lecture Notes in Computer Science*. Springer, 2001.

[PK14] Léo Perrin and Dmitry Khovratovich. Collision spectrum, entropy loss, t-sponges, and cryptanalysis of GLUON-64. In Cid and Rechberger [CR15], pages 82–103.

[PSLL03] Sangwoo Park, Soo Hak Sung, Sangjin Lee, and Jongin Lim. Improving the upper bound on the maximum differential and the maximum linear hull probability for SPN structures and AES. In Thomas Johansson, editor, *Fast Software Encryption, 10th International Workshop, FSE 2003, Lund, Sweden, February 24-26, 2003, Revised Papers*, volume 2887 of *Lecture Notes in Computer Science*, pages 247–260. Springer, 2003.

[RDP+96] Vincent Rijmen, Joan Daemen, Bart Preneel, Antoon Bosselaers, and Erik De Win. The cipher SHARK. In Gollmann [Gol96], pages 99–111.

[RK16] Matthew Robshaw and Jonathan Katz, editors. *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part I*, volume 9814 of *Lecture Notes in Computer Science*. Springer, 2016.

[Rog04] Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In Pil Joong Lee, editor, *Advances in Cryptology - ASIACRYPT 2004, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5-9, 2004, Proceedings*, volume 3329 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2004.

[RS06] Phillip Rogaway and Thomas Shrimpton. A provable-security treatment of the key-wrap problem. In Vaudenay [Vau06], pages 373–390.

[Sho96] Victor Shoup. On fast and provably secure message authentication based on universal hashing. In Neal Koblitz, editor, *Advances in Cryptology - CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings*, volume 1109 of *Lecture Notes in Computer Science*, pages 313–328. Springer, 1996.

[SLG+16] Bing Sun, Meicheng Liu, Jian Guo, Longjiang Qu, and Vincent Rijmen. New insights on AES-like SPN ciphers. In Robshaw and Katz [RK16], pages 605–624.

[Sta78] A. J. Stam. Distance between sampling with and without replacement. *Statistica Neerlandica*, 32(2):81–91, 1978.

[TOS10] Eran Tromer, Dag Arne Osvik, and Adi Shamir. Efficient cache attacks on AES, and countermeasures. *J. Cryptology*, 23(1):37–71, 2010.

[Tun12] Michael Tunstall. Improved "partial sums"-based Square attack on AES. Cryptology ePrint Archive, Report 2012/280, 2012. http://eprint.iacr.org/2012/280.

[Vau03] Serge Vaudenay. Decorrelation: A theory for block cipher security. *J. Cryptology*, 16(4):249–286, 2003.

[Vau06] Serge Vaudenay, editor. *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, volume 4004 of *Lecture Notes in Computer Science*. Springer, 2006.

[Wag99] David Wagner. The boomerang attack. In Knudsen [Knu99], pages 156–170.

[WC81] Mark N. Wegman and Larry Carter. New hash functions and their use in authentication and set equality. *J. Comput. Syst. Sci.*, 22(3):265–279, 1981.

[Yas08] Kan Yasuda. A one-pass mode of operation for deterministic message authentication- security beyond the birthday barrier. In Kaisa Nyberg, editor, *Fast Software Encryption, 15th International Workshop, FSE 2008, Lausanne, Switzerland, February 10-13, 2008, Revised Selected Papers*, volume 5086 of *Lecture Notes in Computer Science*, pages 316–333. Springer, 2008.

[YZS+15] Gangqiang Yang, Bo Zhu, Valentin Suder, Mark D. Aagaard, and Guang Gong. The Simeck family of lightweight block ciphers. In Tim Güneysu and Helena Handschuh, editors, *Cryptographic Hardware and Embedded Systems - CHES 2015 - 17th International Workshop, Saint-Malo, France, September 13-16, 2015, Proceedings*, volume 9293 of *Lecture Notes in Computer Science*, pages 307–329. Springer, 2015.

A Security Against Generic Attacks

We consider how AES-PRF behaves when its underlying permutations are idealized. Patarin [Pat13a] found several attacks against $\mathrm{XoP}^d_{k_1,\ldots,k_d}$ dependent on q, the number of oracle queries. Here we adapt the attacks to the two-permutation GEDMD construction used in AES-PRF.

A.1 $q = 2^n$

Given access to the full codebook, the following property occurs with probability 1:

$$\bigoplus_{x\in\{0,1\}^n} (E_{k_2}(E_{k_1}(x)) \oplus E_{k_1}(x)) = \bigoplus_{x\in\{0,1\}^n} E_{k_2}(E_{k_1}(x)) \oplus \bigoplus_{x\in\{0,1\}^n} E_{k_1}(x) = 0\,.$$

In a random function, this event has probability 2^{-n}. This yields an attack with advantage of $1 - 2^{-n}$ with running time of 2^n xor operations.

A.2 $q < 2^n$

In this setting, we can distinguish EDMD from a random function by counting the number of collisions. Let $n_{coll}(q)$ be this quantity. In a random function, the expected number of collisions is $n_{coll} = \binom{q}{2}/2^n$; in EDMD it is $\binom{q}{2}/(2^n - 1)$. This distinguisher stems from the fact that given a collision $E_{k_2}(E_{k_1}(x)) \oplus E_{k_1}(x) = E_{k_2}(E_{k_1}(y)) \oplus E_{k_1}(y)$, we equivalently have $E_{k_2}(E_{k_1}(x)) \oplus E_{k_2}(E_{k_1}(y)) = E_{k_1}(x) \oplus E_{k_1}(y)$, in which neither side can be 0.

When $q < 2^{n/2}$, the distinguisher simply outputs 1 when a collision exists, and 0 otherwise. The advantage is given by $\frac{\binom{q}{2}}{2^{2n}-2^n} \approx q^2/2^{2n}$. When $q > 2^{n/2}$, the distinguisher is slightly different: output 1 when $n_{coll} \geq \binom{q}{2}/2^n$, 0 otherwise. The advantage here is more complex to calculate, but Patarin calculates it to be $O(q/2^{3n/2})$. This attack strategy is likely to be optimal, as it matches the recent asymptotic bound on the sum of permutations by Eberhard [Ebe17, Theorem 1.5].

None of these attacks is particularly threatening to the PRF security of AES-PRF, as no amount of extra computation—short of bruteforcing the key—will be of any help. In effect, the advantage remains negligible even when the attacker obtains nearly the entire codebook.