

On Weak Keys and Forgery Attacks Against Polynomial-Based MAC Schemes

Gordon Procter^(✉) and Carlos Cid

Information Security Group, Royal Holloway,
University of London, London, UK
{gordon.procter.2011, carlos.cid}@rhul.ac.uk

Abstract. Universal hash functions are commonly used primitives for fast and secure message authentication in the form of Message Authentication Codes (MACs) or Authenticated Encryption with Associated Data (AEAD) schemes. These schemes are widely used and standardised, the most well known being McGrew and Viega’s Galois/Counter Mode (GCM). In this paper we identify some properties of hash functions based on polynomial evaluation that arise from the underlying algebraic structure. As a result we are able to describe a general forgery attack, of which Saarinen’s cycling attack from FSE 2012 is a special case. Our attack removes the requirement for long messages and applies regardless of the field in which the hash function is evaluated. Furthermore we provide a common description of all published attacks against GCM, by showing that the existing attacks are the result of these algebraic properties of the polynomial-based hash function. Finally, we greatly expand the number of known weak GCM keys and show that almost every subset of the keyspace is a weak key class.

Keywords: Universal hashing · MAC · Galois/Counter Mode · Cycling attacks · Weak keys

1 Introduction

The study of information-theoretic message authentication codes and universal hashing was initiated by Gilbert et al. [14] and Carter and Wegman [10, 11, 38, 39]. Universal hash functions can be used to construct message authentication codes in both the information-theoretically secure and computationally secure settings (see [9, 39]). Simmons [33] provides a general summary of the theory of unconditionally secure message authentication. Bernstein [2, 3] provides a thorough description of the geneology and more recent literature of unconditionally secure message authentication, including a description of the contributions of

The work described in this paper has been supported in part by the European Commission through the ICT programme under contract ICT-2007-216676 ECRYPT II.

Bierbrauer et al. [5], den Boer [12], and Taylor [36] to polynomial-based hashing. Bernstein [4] also gives an interesting overview of the security of universal hash function based MACs in the computationally secure setting. Shoup [32] describes several methods for realising universal hash function families that are related to polynomials including the evaluation hash [5, 12, 36] which is a variant of the division hash or cryptographic CRC of Krawczyk [21] (itself a variant of Rabin’s fingerprinting codes [27]).

In this paper, we focus on message authentication codes constructed from universal hash functions that are realised by polynomial evaluation. These are widely used and standardised; for examples see [2, 13, 17, 20, 22, 31]. McGrew and Viega’s Galois/Counter Mode (GCM) [26] is the most widely deployed polynomial-based scheme. The algorithm is generally assumed to be secure, with a small number of papers containing attacks against the authentication component via the universal hash function: Ferguson’s attack against truncated GCM tags [15], demonstrating that the security of short tags is significantly lower than would be expected; Joux’s ‘forbidden attack’ [19], illustrating the brittleness of GCM under nonce reuse; Handschuh and Preneel’s [16] extension to Joux’s attack [16]; and Saarinen’s cycling attacks [29], which highlight a weakness due to the underlying algebraic structure of a hash function based on polynomial evaluation. Both Handschuh and Preneel [16] and Saarinen [29] have described classes of weak keys for polynomial evaluation based universal hash functions, with Saarinen particularly focusing on GCM.

Contributions. A motivation of this work was the observation that all existing attacks against GCM are algebraic in nature, and in fact seem to exploit a fundamental underlying algebraic structure of the polynomial-based hash function. The contributions of this paper are to identify and study some of the properties of hash functions based on polynomial evaluation that are the result of this underlying algebraic structure. As a result, we are able to describe a general forgery attack, of which Saarinen’s cycling attack is a special case; our attack can however be used with short messages, applies regardless of the field in which the hash is evaluated, and facilitates length extension attacks against GCM. Furthermore, we provide a common description of all published attacks against GCM by showing that the existing attacks are the result of these algebraic properties of the polynomial-based hash function. Finally, we greatly expand the number of known weak GCM keys, and show that almost every subset of the key space is a weak key class. We note that the attacks presented in this paper do not in any way contradict the security bounds for GCM given by McGrew and Viega [24]. However the algebraic properties (and related attacks) discussed in this paper appear to be an inherent feature of polynomial-based authentication schemes and therefore should be considered in the security assessment of new schemes and extensions of existing ones.

Structure. This paper is structured as follows. In Sect. 2 we introduce the notation that will be used throughout this paper and provide a brief description of the

syntax and security of message authentication codes. In Sect. 3 we give a basic overview of three schemes that use hash functions based on polynomial evaluation for message authentication, including GCM and SGCM. In Sect. 4 we describe the main technique used in this paper for the cryptanalysis of polynomial-based authentication schemes and discuss some features of the resulting attack that make it more interesting than cycling attacks. Section 5 contains a common description of the existing attacks against GCM. In Sect. 6 we show that there are many more weak key classes for hash functions based on polynomial evaluation than have previously been described and suggest a method to realise a key recovery attack against polynomial-based hash function schemes. Section 7 contains a discussion of the consequences of this attack.

2 Preliminaries

2.1 Notation

We consider a message M parsed as $M_1 || \dots || M_m$, where each M_i is n bits long and $||$ represents concatenation of strings. In the syntax of authenticated encryption with associated data [28], this message consists of associated data $A \in \mathcal{A}$ that is authenticated but not encrypted and plaintext $P \in \mathcal{P}$ that will be encrypted and authenticated.

A family of hash functions will be denoted $\mathcal{H} = \{h_H : \{0, 1\}^* \rightarrow \{0, 1\}^n \mid H \in \mathcal{K}_{\mathcal{H}}\}$ with each hash function h_H indexed by a key $H \in \mathcal{K}_{\mathcal{H}}$. A block cipher E is a family of permutations on $\{0, 1\}^n$, with each permutation indexed by a key $k \in \mathcal{K}_E$. The application of a block cipher to input $x \in \{0, 1\}^n$ using key k will be denoted by $E_k(x)$. Where a nonce is used it will be denoted by N .

A finite field will be denoted by \mathbb{K} unless the order of the field has particular relevance, in which case it will be denoted by \mathbb{F}_{p^r} with $|\mathbb{F}| = p^r$. The multiplicative group of a field \mathbb{K} will be denoted by \mathbb{K}^* .

2.2 Universal Hash Functions

A family of hash functions is said to be ϵ -almost \oplus universal if for every $M, M' \in \{0, 1\}^*$ with $M \neq M'$ and for every $c \in \{0, 1\}^n$, $\Pr_{H \in \mathcal{K}_{\mathcal{H}}} [h_H(M) \oplus h_H(M') = c] < \epsilon$. Throughout this paper ϵ -almost \oplus universal will be abbreviated to ϵ -AXU. This condition was introduced by Krawczyk [21] under the name ϵ -OTP-Secure as it is a necessary and sufficient condition for unconditional MAC security when the output of the hash function is encrypted with the one time pad in a field of characteristic 2. In this paper we will generally refer to ϵ -AXU hash function families; however any remark made that requires an ϵ -AXU hash function family in characteristic 2 will also hold for an ϵ -almost strongly universal [34] or ϵ -almost Δ universal [35] hash function family in any finite field.

A polynomial based hash function family is a common way to realise an ϵ -AXU hash function family. Shoup [32] describes several examples of this type

of construction; the main example of interest to this paper is the evaluation hash. In the case of the evaluation hash the message M determines a polynomial $g_M = \sum_{i=1}^m M_i x^i \in \mathbb{K}[x]$, where $M = M_1 || \dots || M_m$ with each $M_i \in \mathbb{K}$. The hash key is an element $H \in \mathbb{K}$ and we define the hash function by $h_H(M) = g_M(H)$.

There are several methods for turning a universal hash function into a message authentication code (see [9, 39] for early examples). The two most common methods are $E_k(N) + h_H(M)$ and $E_k(h_H(M))$.

2.3 Syntax

We will follow Black et al. [8] for a description of the syntax of nonce-based message authentication schemes. A message authentication scheme is a pair of algorithms, Gen and MAC , with four associated sets: \mathcal{K} , the set of possible keys; \mathcal{M} , the message space; \mathcal{N} , the set of nonces and \mathcal{T} , the set of possible authentication tags.

The key generation algorithm Gen takes as input the security parameter and probabilistically outputs the shared key $k \in \mathcal{K}$. The algorithm MAC takes as input a key $k \in \mathcal{K}$, a nonce $N \in \mathcal{N}$, and a message $M \in \mathcal{M}$ and outputs a tag $T \in \mathcal{T}$. The authenticity of a tuple (N, M, T) is verified by computing $\text{MAC}(k, N, M)$: if $T = \text{MAC}(k, N, M)$ then the tag is valid, otherwise it is invalid.

2.4 Security

An adversary attacking a message authentication scheme is given access to two oracles: a tag generation oracle \mathcal{S} and a verification oracle \mathcal{V} . At the beginning of the experiment Gen is run to obtain k , then MAC takes queries (N, M) and returns $\text{MAC}(k, N, M)$. The verification oracle takes queries (N, M, T) and returns 1 if $T = \text{MAC}(k, N, M)$ or 0 otherwise. An adversary is said to successfully forge an authentication tag if they can produce a verification query (N, M, T) so that \mathcal{V} returns 1 when (N, M) was not previously queried to \mathcal{S} .

A common restriction of this security notion is to nonce-respecting adversaries where, although the adversary can control the nonce, they never query \mathcal{S} for (N, M') if they have previously queried \mathcal{S} for (N, M) .

McGrew and Viega [24], Ferguson [15], and Handschuh and Preneel [16] all assert that the probability of creating a valid (non-truncated) tag having seen a single valid (message, tag) pair is approximately $m/|\mathbb{K}|$ where the polynomial is evaluated in \mathbb{K} and m is the length of message that the construction operates on. It is worth emphasising that in this context, m is the *maximum* permissible message length. This is included in the original paper [24] but is not made explicitly clear in the later papers [15, 16]. In this paper we will demonstrate the importance of this distinction via a method of forging GCM tags using a longer message than the one that was given in the valid (message, tag) pair from the tag generation oracle.

Throughout this paper we will focus on GCM for concreteness however the majority of the comments apply equally to any other hash function based on

polynomial evaluation. Most of the results in this paper apply equally to both common constructions of MACs from universal hash functions, either $T = E_k(N) + h_H(M)$ or $T = E_k(h_H(M))$, as our results are based on collisions in the hash function. Where necessary it will be made clear that a remark is dependent on one of these general constructions or the specific structure of GCM.

3 Polynomial-Based Authentication Schemes

We present below a brief description of some of the main authentication schemes based on polynomial evaluation hash functions that are of relevance to our work.

3.1 Galois/Counter Mode

Galois/Counter Mode (GCM) is an AEAD scheme submitted to NIST by McGrew and Viega in 2004, with the specification slightly revised in 2005 [26] (although the revision contained ‘*no normative changes [from the 2004 specification]*’). GCM combines counter mode encryption with a polynomial evaluation based MAC following the Encrypt–then–MAC paradigm, although the authentication key is derived from the block cipher key.

AES–GCM encryption takes as input: a key k , an initialisation vector IV (the nonce), plaintext $P = P_1 || \dots || P_p$ and additional data $A = A_1 || \dots || A_a$. The key is 128, 192 or 256 bits long; the IV should preferably be 96 bits long although any length is supported (see [18]); and for each i , $|P_i| = |A_i| = 128$ except for perhaps a partial final block. With this input, AES–GCM returns a ciphertext $C = C_1 || \dots || C_p$ (the same length as the plaintext) and an authentication tag T .

The plaintext is encrypted using AES in counter mode, under key k with counter value starting at CTR_1 . If the IV is 96 bits long the initial counter value (CTR_0) is $IV || 0^{31}1$, otherwise it is a polynomial evaluation based hash of IV after zero padding (using the hash key described below). For each i , $CTR_i = \text{inc}(CTR_{i-1})$, where $\text{inc}(\cdot)$ increments the last 32 bits of its argument (modulo 2^{32}).

The authentication tag is computed from a polynomial evaluation hash (in $\mathbb{F}_{2^{128}}$). The message M is parsed as 128-bit blocks (with partial final blocks zero padded) and each block is interpreted as an element of $\mathbb{F}_{2^{128}}$. The first block M_1 encodes the length of the (unpadded) plaintext and additional data and will be referred to as the ‘length field’ throughout this paper. This is followed by blocks of additional data $M_2, \dots, M_{a+1} = A_a, \dots, A_1$ and then the encrypted plaintext $M_{a+2}, \dots, M_{a+p+1} = C_p, \dots, C_1$. Note that in this description the labelling of the blocks M_i are reversed from those given in the original GCM specification as this gives a neater description of the polynomial used in evaluating the hash function. The hash key H is derived from the block cipher key: $H = E_k(0^{128})$. The hash function is then computed as $h_H(M) = \sum_{i=1}^{a+p+1} M_i H^i$ (where all operations are in $\mathbb{F}_{2^{128}}$). The authentication tag is given by:

$$T_M = E_k(CTR_0) \oplus h_H(M).$$

3.2 Sophie Germain Counter Mode

In 2012, Saarinen [29] observed cycling attacks against GCM and other polynomial MACs and hashes. Following this Saarinen proposed SGCM [30] as a variant of GCM; SGCM differs from GCM only by the choice of field in which the hash is computed. SGCM uses \mathbb{F}_q , where $q = 2^{128} + 12451$, rather than $\mathbb{F}_{2^{128}}$, as \mathbb{F}_q^* has significantly fewer subgroups than $\mathbb{F}_{2^{128}}^*$. It was claimed that SGCM offers increased resistance to cycling attacks as a result of this change.

3.3 Poly1305–AES

Bernstein proposed Poly1305–AES in 2005 [2]¹. Poly1305–AES takes as input two 128-bit keys, one for AES and one for the hash (with some specific bits set to zero); a 128 bit nonce; and a message (a byte string). The output of Poly1305–AES is a 128-bit authentication tag.

The hash of a message is computed by evaluating a message-dependant polynomial at the secret key (in $\mathbb{F}_{2^{130-5}}$), and encrypting this by adding (in $\mathbb{F}_{2^{130-5}}$) the output of $\text{AES}_k(N)$ before reducing modulo 2^{128} .

4 Algebraic Structure of Polynomial-Based Authentication Schemes

Let \mathcal{H} be a family of hash functions $\mathcal{H} = \{h_H : \{0, 1\}^* \rightarrow \{0, 1\}^n \mid H \in \mathcal{K}_{\mathcal{H}}\}$ based on polynomial evaluation and let M be an input string. Let $h_H(M) = g_M(H)$, where $g_M(x) = \sum_{i=1}^m M_i x^i \in \mathbb{K}[x]$ and $H \in \mathbb{K}$. Now let $q(x) = \sum_{i=1}^r q_i x^i \in \mathbb{K}[x]$ be a polynomial with constant term zero, such that $q(H) = 0$. Then it follows that

$$h_H(M) = g_M(H) = g_M(H) + q(H) = g_{M+Q}(H) = h_H(M + Q),$$

where $Q = q_1 || q_2 || \dots || q_r$ and the addition $M + Q$ is done block-wise (the shorter is zero-padded if required). Thus given a polynomial $q(x)$ satisfying these properties, it is straightforward to construct collisions for the hash function. It is trivial to see that one can use any polynomial $q(x) \in \langle x^2 - Hx \rangle \subseteq \mathbb{K}[x]$.

Collisions in the hash function correspond to MAC forgeries by substituting the original message for the one that yields a collision in the hash function. These forgeries arise from collisions in the hash function and hence the messages can be substituted without any dependence on the method or key used to encrypt the output of the hash function. This method allows an adversary to create forgeries when he has seen a tuple of (nonce, message, tag) by only modifying the message.

It should be noted that the polynomial defined by the message will always have a zero constant term and therefore the polynomial $q(x)$ that is used to forge will always have x as a factor. If this term were non-zero and the hash

¹ There is a preliminary version from 2004 on his website: <http://cr.yp.to/mac.html>.

of a message was encrypted additively (i.e. $T = E_k(N) + h_H(M)$), it would be possible to flip bits in the first message block and flip the same bits in the authentication tag to create a valid forgery. This is the major difference between Shoup’s Cryptographic CRC [32] and Rabin’s fingerprinting codes [27].

The main observation of this paper is that by working with polynomials in the ideal $\langle x^2 - Hx \rangle$, it is straightforward to produce forgeries for polynomial evaluation based authentication schemes. In [29], Saarinen proposed cycling attacks by working with particular polynomials, namely $x^n - x$ (for more detail, see Sect. 5.4). The forgery is successful if $(x - H)|(x^n - x)$ and therefore if $x^n - x \in \langle x^2 - Hx \rangle$. However, the forgery will be successful if *any* polynomial in this ideal is used to mount a similar attack. Furthermore, use of these polynomials also makes it possible to test for membership of large subsets of the keyspace with a single valid (message, tag) pair and a single verification query (see Sect. 6).

4.1 Malleability

In [29], Saarinen also describes ‘targeted bit forgeries’ against GCM where, rather than swapping the full blocks M_i and M_{i+jt} , corresponding bits in each ciphertext block are flipped. This can also be described by the more general attack, by using a multiple of $q(x)$.

If $q(H) = 0$, then $\alpha \cdot q(H) = 0$ for any $\alpha \in \mathbb{K}$ and

$$\begin{aligned} T_M &= E_k(N) + h_H(M) \\ &= E_k(N) + M_1 \cdot H + \dots + M_m \cdot H^m \\ &= E_k(N) + (M_1 + \alpha q_1) \cdot H + \dots + (M_m + \alpha q_m) \cdot H^m \\ &= T_{M+\alpha Q} \end{aligned}$$

where $T_{M+\alpha Q}$ is the authentication tag for the message $M_1 \oplus \alpha \cdot q_1 || \dots || M_m \oplus \alpha \cdot q_m$ (recall that M contains the associated data, encrypted plaintext and the length of both).

If the plaintext is encrypted using a stream cipher (or a block cipher in counter mode) flipping bits in the ciphertext causes the same bits in the plaintext to be flipped. This allows us to predict relations between the original plaintext and the forged plaintext (as $C_i \oplus \alpha q_i$ decrypts to $P_i \oplus \alpha q_i$). Because α can be chosen so as to set $C_i \oplus \alpha q_i$ equal to any value chosen by the adversary (for a single i), an adversary can choose a differential (in a single block) between the original message and the forged message.

If further control over the underlying plaintext is required, several forgery polynomials could be used. In the best case, using t polynomials permits the adversary control over t message blocks. The cost of this extra malleability is that the forgery is only successful if the authentication key is a root of the greatest common divisor of the two polynomials. This can be extended to give as much control over the plaintext as required, but for every extra malleable block the success probability is reduced by at least $\frac{1}{|\mathcal{K}_{\mathcal{K}}|}$.

If the plaintext were encrypted using a block cipher (not in counter mode) then an adversary would not have this fine control over the plaintext, but would still be able to manipulate the ciphertext in this way.

This property also permits an adversary to create as many forgeries as there are non-zero elements in the field (see [7, 25] for further discussion of multiple forgeries).

4.2 Length Extension

In the GCM specification, the last block input to the hash function (corresponding to the term $M_1 \cdot H$ in the MAC calculation) describes the length of the plaintext and additional data. The more general attack described in this paper allows an adversary to manipulate the length field (even though it does not explicitly appear in the sent message). If an adversary is given a valid tag for a message then the content of the length field is known, as it correctly encodes the length of the plaintext and additional data. It is therefore possible to choose a differential in the length field so that it corresponds to the length of the new message. In particular, forgeries can be created using high degree polynomial $q(x)$ regardless of the size of the message in the initial (message, tag) pair.

This is an important remark as it removes one significant limitation on the effectiveness of cycling attacks against GCM [29], which is the length of the message necessary to launch an attack. For a cycling attack to be attempted, an adversary requires as many blocks of correctly authenticated data as there are elements in the subgroup with which he wishes to forge, in order to swap the first and last blocks. By manipulating the length field any forgery probability can be realised starting with a valid authentication tag on a single message block.

A common criticism of GCM is that the maximum message length may be restrictive in the future as data rates increase [15]. However, it follows from our work (and the original security proofs [24]) that increasing the maximum permissible length would significantly decrease the security of the scheme.

4.3 Key Recovery

Saarinen suggests that once a weak key has been identified (by a successful cycling attack), the adversary would create many forgeries by further cycling attacks [29, Sect. 9]. Translating this to the more general polynomial root description: once a successful forgery occurs, the authentication key is known to be one of the roots of the ‘forgery polynomial’ q . Therefore rather than making repeated ‘cycling forgeries’ with guaranteed success but limited control of the plaintext, the adversary can aim to recover the authentication key and forge authentication tags for arbitrary messages. By attempting to forge using a subset of the roots of the forgery polynomial (and reducing the number of roots in the subset after each successful attempt), an adversary can gradually recover the authentication key using a method that is independent of encryption method or key used. This would give a forgery probability less than 1 at each stage, however the adversary can choose a trade-off between the forgery probability and the speed of

recovering the authentication key. This is analogous to the key recovery attack described by Handschuh and Preneel [16] (where the subsets are chosen to realise a binary search of the keyspace). Note that in the case of GCM, recovery of the hash key H does not lead to the recovery of the encryption key k as $H = E_k(0)$.

4.4 Choosing Polynomials

To maximise the probability of a successful forgery it is important that the polynomial used to attempt a forgery has many distinct roots, as a root with multiplicities increases the degree of the polynomial (and hence the length of the attempted forgery) without increasing the probability of success. The naïve way to achieve this is to compute $q(x) = \prod_i (x - H_i)$ for as many H_i as is required to give the desired forgery probability.

Alternatively, if the polynomial defined by the hash function is evaluated in \mathbb{F}_{p^r} and the irreducible factorisation of $x^{p^r} - x$ is computed in a subfield \mathbb{F}_{p^d} , a subset of these factors can be multiplied together (in \mathbb{F}_{p^d}). By choosing distinct irreducible factors, the roots of the product polynomial will be distinct. Cycling attacks [29] employ a variation on this method. The factorisation

$$2^{2^n} - 1 = \prod_{i=1}^n 2^{2^{i-1}} + 1$$

allows Saarinen to find factors of $x^{2^{128}} - x$ in $\mathbb{F}_2[x]$ which can be used in a cycling attack (although they are not necessarily irreducible):

$$\begin{aligned} x^{2^{128}} - x &= x(x-1) \frac{(x^3-1)}{x-1} \frac{(x^5-1)}{x-1} \frac{(x^{17}-1)}{x-1} \dots \\ &= x(x-1)(1+x+x^2)(1+x+\dots+x^4)(1+x+\dots+x^{16}) \dots \end{aligned}$$

To carry out a cycling attack using a subgroup of order t , the factors x , $(x-1)$ and $\frac{(x^t-1)}{x-1}$ are multiplied together to obtain the polynomial $x^{t+1} - x$. In general there is no requirement to select $(x-1)$ or to use only three factors, for example the polynomial $x(1+x+x^2)(1+x+\dots+x^{16})$ could be used to give a forgery probability of $\frac{19}{2^{128}}$. This is not a cycling attack, as the polynomial used contains more than two terms so the forgery does not involve simply swapping two message blocks, but it does rely on the same underlying algebraic structure.

A third option is to use a randomly selected polynomial in $\mathbb{F}_{p^r}[x]$. One potential issue with this method is the presence of repeated factors. Square-free factorisation has been extensively studied as it is a common first step in many polynomial factorisation algorithms (for example, see [37, Ch. 14]). It may be feasible to sample polynomials from $\mathbb{F}_{p^r}[x]$ randomly and process this polynomial to make it more desirable by removing repeated factors. This method does not appear promising due to the large number of irreducible polynomials of any given degree in $\mathbb{F}_{p^r}[x]$ and the observation that a degree d polynomial that consists of a single linear factor and an irreducible polynomial of degree $d - 1$ is

almost as bad as a degree d irreducible polynomial from a forgery probability perspective. Irreducible polynomials in \mathbb{F}_p that are known to have a root in \mathbb{F}_{p^r} would be good candidates for attempting forgeries as the normality of $\mathbb{F}_{p^r}/\mathbb{F}_p$ guarantees that these polynomials will split into linear factors. Unfortunately this does not appear to be a well studied area. A further disadvantage of choosing random polynomials is that, although the roots of a polynomial in $\mathbb{K}[x]$ can be identified efficiently (see [1] for example), it would be unlikely that a non-intersecting subset of the keyspace would be used for a second forgery attempt if the first was unsuccessful.

5 Existing Attacks Against GCM

We show below that the four known attacks on GCM can be described as special cases of the properties discussed in Sect. 4.

5.1 Ferguson’s Short Tag Attack

Ferguson’s attack against GCM when short tags are used [15] begins by attempting to forge using a particular class of polynomials (linearised polynomials). Linearised polynomials have the property that their roots form a linear subspace of the splitting field of the polynomial (see [23, Chap. 3.4] for an overview). Ferguson uses polynomials in $\mathbb{F}_2[x]$ that split over $\mathbb{F}_{2^{128}}$, so the roots correspond to possible authentication keys and it is possible to describe the roots of a linearised polynomial using a matrix over \mathbb{F}_2 . Multiple successful forgeries reduce the dimension of the subspace of the keyspace that contains the authentication key and eventually an adversary can recover the key.

5.2 Joux’s Forbidden Attack

Joux’s ‘forbidden attack’ against GCM [19] is also a specific case of the properties discussed in this paper. This attack requires two messages, M and M' , that are authenticated with the same (key, IV) pair. Reusing the (key, IV) pair in GCM has the effect of reusing H , k and N :

$$\begin{aligned} T_M \oplus T_{M'} &= (h_H(M) \oplus f_k(N)) \oplus (h_H(M') \oplus f_k(N)) \\ &= h_H(M) \oplus h_H(M') \\ &= h_H(M \oplus M') \end{aligned}$$

The adversary knows T_M , $T_{M'}$ and both messages so is able to derive a polynomial that is satisfied by the hash key. This attack is prevented if we only consider nonce-respecting adversaries.

5.3 Handschuh and Preneel

Handschuh and Preneel [16] describe a key recovery attack and a method to verify a guess for a key. They identify the key recovery attack as an extension of Joux’s ‘forbidden attack’ which does not require nonce reuse. It consists of attempting to create a forgery and then searching through the roots of the polynomial defined by the difference between the original message and the forged message. This was initially identified by Black and Cochran [6], but extended and generalised by Handschuh and Preneel. The method for verifying a key guess H corresponds precisely with attempting to forge using the polynomial $(x^2 - Hx)$.

Handschuh and Preneel consider their attack to be infeasible for GCM due to the blocksize of 128 bits, however it is precisely as feasible as Saarinen’s cycling attacks.

5.4 Saarinen’s Cycling Attacks

In 2012, Saarinen observed cycling attacks against GCM and other polynomial-based MACs and hashes [29]. If a hash key H lies in a subgroup of order t , then $H^t = 1 \in \mathbb{K}$ and (for any i, j) message blocks M_i and M_{i+jt} can be swapped without changing the value of the hash.

For example (ignoring GCM’s length encoding), if $H^4 = H$ then blocks M_1 and M_4 can be swapped without changing the value of the hash:

$$\begin{aligned} h_H(M_1||M_2||M_3||M_4) &= M_1 \cdot H \oplus M_2 \cdot H^2 \oplus M_3 \cdot H^3 \oplus M_4 \cdot H^4 \\ &= M_4 \cdot H \oplus M_2 \cdot H^2 \oplus M_3 \cdot H^3 \oplus M_1 \cdot H^4 \\ &= h_H(M_4||M_2||M_3||M_1). \end{aligned}$$

It is more natural and general to consider the authentication keys that fall in low order subgroups as roots of a low degree polynomial. Cycling attacks correspond to the general attack introduced in this paper, using the polynomial

$$q(x) = (M_i - M_{i+jt})(x^{t+1} - x),$$

noting that in fields of characteristic 2 subtraction is the same as \oplus .

$$\begin{aligned} h_H(M_1||M_2||M_3||M_4) &= M_1 \cdot H \oplus M_2 \cdot H^2 \oplus M_3 \cdot H^3 \oplus M_4 \cdot H^4 \\ &= M_1 \cdot H \oplus M_2 \cdot H^2 \oplus M_3 \cdot H^3 \oplus M_4 \cdot H^4 \\ &\quad \oplus (M_1 \oplus M_4) \cdot H \oplus (M_1 \oplus M_4) \cdot H^4 \\ &= M_4 \cdot H \oplus M_2 \cdot H^2 \oplus M_3 \cdot H^3 \oplus M_1 \cdot H \\ &= h_H(M_4||M_2||M_3||M_1) \end{aligned}$$

Using the more general ‘polynomial roots’ description it is possible to forge using any subset of the keyspace. If the authentication keys that we wish to

attempt to forge with are the elements of a low order subgroup, for example the order three subgroup of $\mathbb{F}_{2^{128}}^*$ (identified by Saarinen [29, Sect. 4.1]) plus the all zero key:

$$\begin{aligned} H_0 &= 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00 \\ H_1 &= 80\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00\ 00 \\ H_2 &= 10\ D0\ 4D\ 25\ F9\ 35\ 56\ E6\ 9F\ 58\ CE\ 2F\ 8D\ 03\ 5A\ 94 \\ H_3 &= 90\ D0\ 4D\ 25\ F9\ 35\ 56\ E6\ 9F\ 58\ CE\ 2F\ 8D\ 03\ 5A\ 94 \end{aligned}$$

then the polynomial that is created corresponds precisely to Saarinen's cycling attack. In this case $(x - H_0)(x - H_1)(x - H_2)(x - H_3) = x^4 - x$.

6 Weak Keys for Polynomial-Based Authentication Schemes

For any cryptographic algorithm, a relevant question for its security assessment is whether it contains *weak keys*. Handschuh and Preneel [16, Sect. 3.1] give the following definition of weak keys:

In symmetric cryptology, a class of keys $[\mathcal{D}]$ is called a *weak key class* if for the members of that class the algorithm behaves in an unexpected way and if it is easy to detect whether a particular unknown key belongs to this class. For a MAC algorithm, the unexpected behavior can be that the forgery probability for this key is substantially larger than average. Moreover, if a weak key class $[\mathcal{D}]$ is of size C , one requires that identifying that a key belongs to this class requires testing fewer than C keys by exhaustive search and fewer than C verification queries.

Handschuh and Preneel [16] identify 0 as a weak authentication key for GCM and other similar constructions as $h_0(M) = 0$ for every message M . Following the definition above and because $|\mathcal{D}| = 1$, an adversary is not allowed to test any key by exhaustive search, nor are they allowed any verification queries. For a single element subset of the key space $\mathcal{D} = \{H^*\}$ to be a weak key class, a nonce-respecting adversary has to be able to identify whether or not $H = H^*$ when they are given a number of (message, tag) pairs of their choosing (each created using a different IV). We note that a nonce-respecting adversary can detect whether $\mathcal{D} = \{0\}$ if $|\text{IV}| \neq 96$: in this case all IVs hash to give the same initial counter value and $h_0(M) = 0$ for every message M so all messages have the same authentication tag (as identified in [24, Sect. 5]). If $|\text{IV}| = 96$ a different initial counter value is used to encrypt the output of the hash function and so although the output of the hash function does not change this cannot be detected given the output of the MAC algorithm.

Saarinen [29] demonstrated that the situation is much worse than described by Handschuh and Preneel, as he was able to find classes of weak keys where the authentication key falls in a low order subgroup of \mathbb{K}^* . It is then possible to

create a valid forgery by swapping two message blocks of a valid (message, tag) pair without changing the authentication tag if the authentication key lies in a subgroup with order dividing the distance between the swapped message blocks.

This forgery will be successful if and only if the key is an element of such a subgroup and therefore this provides a simple method for identifying weak keys which requires one valid (message, tag) pair and one verification query. These classes of weak keys therefore meet Handschuh and Preneel's definition of weak keys.

For example, the subset of authentication keys corresponding zero and the elements of the subgroup of order 3 in $\mathbb{F}_{2^{128}}$ is a weak key class. Membership of this subset can be confirmed by a successful forgery if M_i and M_j are swapped and $i \equiv j \pmod{3}$. This is equivalent to attempting a forgery using (a multiple of) the polynomial $x^4 - x$.

However, it follows from the discussion in Sect. 4 that it is possible to derive comparable statements for any set of authentication keys in $\mathbb{F}_{2^{128}}$, except that rather than 'nice' binary descriptions, the polynomial description will involve elements of $\mathbb{F}_{2^{128}}$. In particular, for any set of authentication keys \mathcal{D} we can use any polynomial in the ideal $\prod_{H \in \mathcal{D}} (x^2 - Hx)$ to test for membership of that subset of the keyspace. It follows that *almost every* subset of the GCM keyspace is weak. We discuss this issue further in Sects. 6.1 and 7.

6.1 Keyspace Search

Based on the properties discussed in Sect. 4 it is possible to test for membership of any subset of the keyspace using at most two verification queries. Membership of subsets that include the zero key can be tested by setting $q(x) = \prod_{H \in \mathcal{D}} (x - H)$. This therefore requires one verification query, independent of the size of \mathcal{D} . To test for membership of a subset \mathcal{D} that does not include zero, first test whether $H \in \mathcal{D} \cup \{0\}$ and then rule out $H = 0$ using the method described below. This therefore requires two verification queries, but again is independent of the size of \mathcal{D} . The distinction between subsets including zero or not including zero is a consequence of the constant term of $g_M(x)$ being zero to avoid predictable changes in the output of the hash from flipping low order bits.

Therefore, using Handschuh and Preneel's definition, a set \mathcal{D} of GCM authentication keys is a weak key class if either: $|\mathcal{D}| \geq 3$ or $|\mathcal{D}| \geq 2$ and $0 \in \mathcal{D}$.

Given one valid (message, tag) pair for a single block message and one verification query it is easy to determine whether or not $H = 0$. If the adversary attempts to forge using any other single block message and the same tag, then the forgery is successful if and only if $H = 0$ as seen below.

If no length encoding is used:

$$\begin{aligned} T &= E(\text{CTR}_0) + (M \cdot H) \\ &= E(\text{CTR}_0) + (M' \cdot H) \\ &\Leftrightarrow (M - M') \cdot H = 0 \\ &\Leftrightarrow M = M' \text{ or } H = 0 \end{aligned}$$

If a GCM style length encoding is used:

$$\begin{aligned}
 T &= E(\text{CTR}_0) + (\text{length} \cdot H) + (M \cdot H^2) \\
 &= E(\text{CTR}_0) + (\text{length} \cdot H) + (M' \cdot H^2) \\
 &\Leftrightarrow (M - M') \cdot H^2 = 0 \\
 &\Leftrightarrow M = M' \text{ or } H = 0
 \end{aligned}$$

By testing for membership of subsets of the keyspace, it is plausible that an adversary could recover one bit of the authentication key with each forgery attempt. If $q(x) = \prod_{H \in \mathcal{Y}} (x - H)$, where \mathcal{Y} is the set of authentication keys for which the first bit is zero, then a successful forgery confirms that the first bit of the authentication key is zero and a failure confirms that the first bit is one. Repeating this for each bit of the authentication key, the whole key could be recovered using 128 verification queries.

This would require unfeasibly large messages to be used in the forgery attempts in the case of authentication keys corresponding to elements of a field with $|\mathbb{K}| \approx 2^{128}$, but it is a strong argument against using a hash function based on polynomial evaluation in a field with $|\mathbb{K}| \ll 2^{128}$. This may be a direction taken by variants of GCM designed to improve the performance of GCM (see [40] for one such example), however we recommend extreme caution when considering these modifications. In the case of GCM the size of the subsets that can be tested is limited to around 2^{56} as the maximum message length is limited.

One advantage of being able to test for membership of arbitrary subsets is that it allows the adversary to use any partial knowledge of the authentication key that they may have.

7 Discussions and Conclusions

7.1 Choice of Fields

It is true that the security against cycling attacks, as presented in [29], can be increased by evaluating a hash function in a field with a multiplicative group, the order of which does not have many factors. However the attack introduced in this paper (of which cycling attacks is a special case) applies equally well in any finite field, so Saarinen's claim that *'The security of polynomial-evaluation MACs against attacks of this type of attack can be determined from the factorization of the group size in a straightforward manner'* [29, Sect. 8] is somewhat misleading.

Saarinen's claim is valid in the sense that the factorisation of $|\mathbb{K}| - 1$ determines the extent to which the process of computing irreducible factors will succeed; however an attack using $\prod_{H \in \mathcal{D}} (x - H)$ will work equally well in every field. In particular, it follows from our work that the SGCM variant of GCM has the same inherent weaknesses regarding polynomial based forgery attacks.

7.2 Length Extension

It is unfortunate that including the length of the additional authenticated data and plaintext in the input to the hash function is not sufficient to prevent the length extension attack presented in this paper. In schemes that use a GCM-like length encoding, if the value of the length field were encrypted using a block cipher before being input to the hash function, it would not be possible to alter the message length as described in Sect. 4. However, one of the design goals of GCM was to take advantage of AES pipelining, which precludes the use of the block cipher to compute the authentication tag.

7.3 Malleability

Part of the reason that this weakness in the algebraic structure of polynomial hashing is problematic for GCM is that it allows an adversary to choose the changes that are made to the plaintext in a forged message. This is because addition in a field of characteristic 2 is used for both the counter mode encryption and the hash function evaluation.

One way to avoid this issue is to use different operations during encryption and MAC generation. This is one significant advantage that (CTR & Poly1305-AES) [2] has over GCM, as in this scheme the MAC is computed using addition in a prime order field while the message is encrypted using addition in a field of characteristic 2.

An alternative method to increase the difficulty for an adversary attempting to make meaningful manipulations of plaintext is to use a mode of operation other than CTR as this will prevent the ‘targeted bit forgeries’ described by Saarinen [29, Sect. 6] and the analogous forgeries in this paper.

GCM roughly follows the Encrypt-then-MAC paradigm, as is generally perceived to be best practice (although MAC-then-Encrypt has also been proved secure in the nonce-based AEAD setting [28]). Despite going against the perceived best practice, using a MAC-then-Encrypt approach (in addition to the proposed changes described above) would make it harder for an adversary to create ciphertexts that correctly decrypt to a plaintext known to be related to a (plaintext,ciphertext) pair obtained from a query. We note however that the introduction of other weaknesses caused by making these changes has not been ruled out.

7.4 Weak Keys

The weak key classes that are identified in Sect. 6 cause the forgery probability to be higher than expected because an adversary can detect whether the authentication key that is being used is a member of that class and can then forge with probability one.

The broader issue with polynomial evaluation based hashes is that it is possible to test for membership of large subsets of the keyspace with only one or

two verification queries and once an adversary has successfully confirmed membership of a subset he can either continue to forge messages or conduct a search of a much reduced keyspace. This is an unusual and undesirable property of a cryptosystem.

It is interesting that the two-element subsets of the keyspace containing zero are weak key classes, while those that do not contain zero are not, yet any subset of the keyspace containing at least three elements is weak. This perhaps suggests a problem with the definition of a weak key class. In our opinion the definition is correct and the observations made in this paper are unavoidable properties of hash functions based on polynomial evaluation that result from the algebraic structure of the construction, so are not best described in terms of the number of weak keys.

The most important discussion around this issue is whether an algorithm in which almost every subset of the keyspace is a weak key class is a weak algorithm or whether this is a property of the construction that, although highly undesirable, is not considered to reduce the security of the scheme to an unacceptable level. We suggest that in the case of GCM it is the latter; in other polynomial-based MAC schemes with different parameters it may be the former and this property must be considered when designing and evaluating schemes.

References

1. Berlekamp, E.R.: Factoring polynomials over large finite fields. *Math. Comput.* **24**(111), 713–735 (1970)
2. Bernstein, D.J.: The Poly1305-AES message-authentication code. In: Gilbert, H., Handschuh, H. (eds.) *FSE 2005*. LNCS, vol. 3557, pp. 32–49. Springer, Heidelberg (2005)
3. Bernstein, D.J.: The Poly1305-AES message-authentication code. Slides from *FSE (2005)*. <http://cr.yp.to/talks/2005.02.21-1/slides.pdf>
4. Bernstein, D.J.: Stronger security bounds for Wegman-Carter-Shoup authenticators. In: Cramer, R. (ed.) *EUROCRYPT 2005*. LNCS, vol. 3494, pp. 164–180. Springer, Heidelberg (2005)
5. Bierbrauer, J., Johansson, T., Kabatianskii, G., Smeets, B.: On families of hash functions via geometric codes and concatenation. In: Stinson, D.R. (ed.) *CRYPTO 1993*. LNCS, vol. 773, pp. 331–342. Springer, Heidelberg (1994)
6. Black, J., Cochran, M.: MAC reforgeability. *Cryptology ePrint Archive*, report 2006/095 (2006)
7. Black, J., Cochran, M.: MAC reforgeability. In: Dunkelman, O. (ed.) *FSE 2009*. LNCS, vol. 5665, pp. 345–362. Springer, Heidelberg (2009)
8. Black, J., Halevi, S., Krawczyk, H., Krovetz, T., Rogaway, P.: UMAC: fast and secure message authentication. In: Wiener, M. (ed.) *CRYPTO 1999*. LNCS, vol. 1666, pp. 216–233. Springer, Heidelberg (1999)
9. Brassard, G.: On computationally secure authentication tags requiring short secret shared keys. In: Chaum, D., Rivest, R.L., Sherman, A.T. (eds.) *CRYPTO*, pp. 79–86. Plenum Press, New York (1982)
10. Carter, L., Wegman, M.N.: Universal classes of hash functions (extended abstract). In: Hopcroft, J.E., Friedman, E.P., Harrison, M.A. (eds.) *STOC*, pp. 106–112. ACM (1977)

11. Carter, L., Wegman, M.N.: Universal classes of hash functions. *J. Comput. Syst. Sci.* **18**(2), 143–154 (1979)
12. den Boer, B.: A simple and key-economical unconditional authentication scheme. *J. Comput. Secur.* **2**, 65–72 (1993)
13. Dworkin, M.: Recommendation for block cipher modes of operation: Galois/Counter Mode (GCM) and GMAC. NIST Special Publication 800–38D, NIST, Nov 2007
14. MacWilliams, F.J., Gilbert, E.N., Sloane, N.J.A.: Codes which detect deception. Technical report 3, Bell Sys. Tech. J., Mar 1974
15. Ferguson, N.: Authentication weaknesses in GCM. Comments submitted to NIST Modes of Operation Process (2005)
16. Handschuh, H., Preneel, B.: Key-recovery attacks on universal hash function based MAC algorithms. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 144–161. Springer, Heidelberg (2008)
17. Igoe, K., Solinas, J.: AES Galois Counter Mode for the secure shell transport layer protocol. IETF Request for Comments 5647 (2009)
18. Iwata, T., Ohashi, K., Minematsu, K.: Breaking and repairing GCM security proofs. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 31–49. Springer, Heidelberg (2012)
19. Joux, A.: Authentication failures in NIST version of GCM. Comments submitted to NIST Modes of Operation Process (2006)
20. Kohno, T., Viega, J., Whiting, D.: CWC: a high-performance conventional authenticated encryption mode. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 408–426. Springer, Heidelberg (2004)
21. Krawczyk, H.: LFSR-based hashing and authentication. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 129–139. Springer, Heidelberg (1994)
22. Law, L., Solinas, J.: Suite B cryptographic suites for IPsec. IETF Request for Comments 6379 (2011)
23. Lidl, R., Neiderreiter, H.: Finite Fields, vol. 20, 2nd edn. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge (1997)
24. McGrew, D.A., Viega, J.: The security and performance of the Galois/Counter Mode (GCM) of operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 343–355. Springer, Heidelberg (2004)
25. McGrew, D.A., Fluhrer, S.R.: Multiple forgery attacks against message authentication codes. Comments submitted to NIST on the Choice Between CWC or GCM (2005)
26. McGrew, D.A., Viega, J.: The Galois/Counter Mode of operation (GCM). Submission to NIST Modes of Operation Process, May 2005
27. Rabin, M.O.: Fingerprinting with random polynomials. Technical report (1981)
28. Rogaway, P.: Authenticated-encryption with associated-data. In: Atluri, V. (ed.) ACM Conference on Computer and Communications Security, pp. 98–107. ACM (2002)
29. Saarinen, M.-J.O.: Cycling attacks on GCM, GHASH and other polynomial MACs and hashes. In: Canteaut, A. (ed.) FSE 2012. LNCS, vol. 7549, pp. 216–225. Springer, Heidelberg (2012)
30. Saarinen, M.-J.O.: SGCM: the Sophie Germain Counter Mode. Cryptology ePrint Archive, report 2012/326 (2012)
31. Salter, M., Housley, R.: Suite B profile for transport layer security (TLS). IETF Request for Comments 6460 (2011)

32. Shoup, V.: On fast and provably secure message authentication based on universal hashing. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 313–328. Springer, Heidelberg (1996)
33. Simmons, G.J.: Contemporary Cryptology: The Science of Information Integrity. Institute of Electrical, and Electronics Engineers. IEEE Press, Piscataway (1992)
34. Stinson, D.R.: Universal hashing and authentication codes. *Des. Codes Crypt.* **4**(3), 369–380 (1994)
35. Stinson, D.R.: On the connections between universal hashing, combinatorial designs and error-correcting codes. *Electron. Colloquium Comput. Complexity (ECCC)* **2**(52), 1–24 (1995)
36. Taylor, R.: Near optimal unconditionally secure authentication. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, pp. 244–253. Springer, Heidelberg (1995)
37. von zur Gathen, J., Gerhard, J.: Modern computer algebra, 2nd edn. Cambridge University Press, Cambridge (2003)
38. Wegman, M.N., Carter, L.: New classes and applications of hash functions. In: FOCS, pp. 175–182. IEEE Computer Society (1979)
39. Wegman, M.N., Carter, L.: New hash functions and their use in authentication and set equality. *J. Comput. Syst. Sci.* **22**(3), 265–279 (1981)
40. Aoki, K., Yasuda, K.: The security and performance of “GCM” when short multiplications are used instead. In: Kutyłowski, M., Yung, M. (eds.) Inscrypt 2012. LNCS, vol. 7763, pp. 225–245. Springer, Heidelberg (2013)