Could SFLASH be repaired ?

(full version*)

Jintai Ding[1], Vivien Dubois[2], Bo-Yin Yang[3*],
Owen Chia-Hsin Chen[3], and Chen-Mou Cheng[4]

[1] Dept. of Mathematics and Computer Sciences, University of Cincinnati
[2] CELAR, France
[3] Institute of Information Sciences, Academia Sinica, Taiwan
[4] Dept. of Electrical Engineering, National Taiwan University

Abstract. The SFLASH signature scheme stood for a decade as the most successful cryptosystem based on multivariate polynomials, before an efficient attack was finally found in 2007. In this paper, we review its recent cryptanalysis and we notice that its weaknesses can all be linked to the fact that the cryptosystem is built on the structure of a large field. As the attack demonstrates, this richer structure can be accessed by an attacker by using the specific symmetry of the core function being used. Then, we investigate the effect of restricting this large field to a purely linear subset and we find that the symmetries exploited by the attack are no longer present. At a purely defensive level, this defines a countermeasure which can be used at a moderate overhead. On the theoretical side, this informs us of limitations of the recent attack and raises interesting remarks about the design itself of multivariate schemes.

Keywords: multivariate cryptography, signature, SFLASH, differential.

1 Introduction

Multivariate schemes are asymmetric primitives based on hard computational problems involving multivariate polynomials. Reference problems are for instance solving a system of multivariate polynomial equations, or deciding whether two sequences of multivariate polynomials are isomorphic. The research for such schemes originates from Matsumoto and Imai's work in the early 80s, but has really been active for a decade. The practical interest for considering such schemes, besides the obvious diversification effort, comes from their usual high performances which make

* An extended abstract of this paper appears in the proceedings of ICALP 2008. An earlier version of this work can also be found at http://eprint.iacr.org/2007/366. Correspondence to BY at by@moscito.org.

them well-suited for implementation on small devices. On the other side, the area is young and much cryptanalytic effort is still to be done to understand well what their security might rely on.

Multivariate schemes are all based on a construction method inspired from McEliece [11]: an easy-to-invert multivariate vectorial function is transformed into a random-looking one by applying secret linear bijections on both variables and coordinates. Of course, such a linear hiding has the nice feature to be very easy to undo by the legitimate user, but it also has the drawback of leaking the invariant properties of the internal function. Whenever such invariant properties can be used in order to devise a cryptanalytic attack (e.g. elimination properties enhancing Gröbner basis computation), one uses additional transformations to destroy them.

SFLASH is a signature scheme proposed by Patarin, Goubin and Courtois [16], following a design they had introduced at Asiacrypt'98 [14]. The easy-to-invert internal function of SFLASH is defined from a single variable polynomial over some field extension $\mathbb{F}_{q^n}$ and turned into a function from $(\mathbb{F}_q)^n$ to itself by using the linear structure of $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$. To allow efficient inversion, this function has a specific shape as a polynomial over $\mathbb{F}_{q^n}$, namely this is a *monomial* which is inverted by raising to the inverse exponent, like in RSA. The basic McEliece-type hiding, *i.e.* using two linear bijections, of such a function was the initial proposal – known as the C* cryptosystem – of Matsumoto and Imai [10], but it was later seen by Patarin [13] that the hidden monomial structure implies some algebraic properties of the public function which can be exploited for an attack. However, Patarin, Goubin and Courtois later showed [14] that algebraic attacks can be very easily avoided by using an additional transformation initially used by Shamir [15] which consists in simply *deleting a few coordinates* of the public function. Schemes obtained from the application of *minus* to C* are termed C^{*-} schemes; they are suitable for signature. SFLASH is a C^{*-} scheme chosen as a candidate for the selection organized by the NESSIE European consortium [1], and accepted in 2003 [12].

Recently, Dubois, Fouque, Shamir and Stern discovered a new property of C* monomials which is almost not affected by the *minus* transformation, and which can be used to recover missing coordinates of the public function [4,3]. As a consequence, all practical parameters choices for C^{*-} schemes, including those of SFLASH, were shown insecure. The attack found by Dubois *et al.* is the most effective development of a new kind of cryptanalysis which targets geometrical properties of multivariate functions. Consequences of this attack are of course a reevaluation of related cryptosystems and a more careful study of the properties of the internal functions being used. However it seems that the mere design principle of multivariate schemes is here in question : can we effectively hide a particular function such as a C* monomial using linear maps ?

Our results. In this paper, we review the recent cryptanalysis of SFLASH and we notice that its weaknesses can all be linked to the fact that the cryptosystem is built on the structure of a large field. As the attack demonstrates, this richer structure can be accessed by an attacker by using the specific symmetry of the internal C* function that can be perceived from even a small number of public polynomials. Note that, since the large field structure is only necessary to perform the secret operations, it needs not be encapsulated in the public key. Then, we study the effect of restricting this large field to a purely linear subset, and we find that the symmetries exploited by the attack are no longer present. Indeed the symmetries of the C* monomial are fundamentally linked to the large field multiplication and do not hold when restricted to a non-multiplicative subset; we provide mathematical proofs for the target cases explaining this phenomenon in detail. As we will see, this result conveys additional perspective on the general design of multivariate schemes.

Organization of the Paper. In Section 2, we give a brief introduction to SFLASH. In Section 3, we review its recent cryptanalysis [4,3]. In Section 4, we show that the geometrical properties which are exploited by the attack do not hold when restricting the internal function to a proper subspace of the large field. In Section 5, we define a modified family of schemes which resist the attack. We discuss our results in Section 6.

2 The SFLASH Scheme

2.1 The C* scheme

The C* scheme was proposed by Matsumoto and Imai in 1988. It uses a *monomial* over $\mathbb{F}_{q^n}$: $F(x) = x^{1+q^\theta}, x \in \mathbb{F}_{q^n}$, where x can be identified with an n coordinates vector over $\mathbb{F}_q$ by fixing some basis of $\mathbb{F}_{q^n}$. The exponent $1 + q^\theta$ is chosen invertible modulo $q^n - 1$ and raising to its inverse is inverting F. Since $1 + q^\theta$ has q-weight 2, F corresponds to a multivariate function from $(\mathbb{F}_q)^n$ into itself of degree 2. On the other hand, the inverse of $1+q^\theta$ has very high q-weight $\mathcal{O}(n)$ for prescribed values of θ [10], and the inverse of F then corresponds to a multivariate function from $(\mathbb{F}_q)^n$ into itself with very high degree $\mathcal{O}(n)$. A C* scheme is built by transforming F with randomly chosen linear bijections S and T : $\boldsymbol{P} = T \circ F \circ S$.

The resulting function $\boldsymbol{P}$ has the same *multivariate* properties as F, but the twisting provided by S and T hides the *single variable* representation which allows fast inversion. Unfortunately, Patarin showed in 1995 [13] that although the plaintext x is a high degree function in term of the ciphertext y, the pairs (x, y) satisfy many low degree algebraic relations, whose degree is independent of the security parameter n. This implies vulnerability to algebraic attacks.

2.2 SFLASH

To avoid an attacker to possibly reconstruct existing algebraic relations on the pairs (x, y), a simple idea is not to provide the entire description of how these variables are related. The most easy way to realize this was used by Shamir in 1993 [15] and consists in simply removing a few coordinate-polynomials of the public key, say the last r ones where r is an additional parameter. Furthermore, Patarin, Goubin and Courtois showed in 1998 [14] that for a C* scheme, the degree of algebraic relations between x and the partial y is quickly growing with the parameter r. Of course, the resulting scheme is no longer bijective but it can still be used for signature at no performance loss. These schemes were introduced as C^{*-} by Patarin, Goubin and Courtois [14]. A public key consists of the $n - r$ first coordinates of an initial C* public key $\boldsymbol{P} = T \circ F \circ S$ with T and S as the secret key. A rationale for the parameter r is provided in [14]; choosing r with $q^r \geq 2^{80}$ is then required for a 2^{80} security level. Besides, no algebraic attack is expected to succeed when r is not too small in regards to n, the initial number of polynomials. SFLASH is a C^{*-} scheme chosen by Patarin *et al.* for the NESSIE selection. For the recommended parameters $q = 2^7$, $n = 37$, $\theta = 11$ and $r = 11$, the signature length is 239 bits and the public key size is 15 Kbytes.

3 The Symmetry in SFLASH

The design of SFLASH was aimed at resisting algebraic attacks and stood challenging for almost ten years. However, in the last four years, a new kind of cryptanalysis for multivariate schemes has been developed based on geometrical properties of the so-called differential [8,5,6]. As defined in the initial paper by Fouque, Granboulan and Stern [8], the differential transforms a *quadratic* function $\boldsymbol{P}(x)$ into its *bilinear symmetric* associate, denoted $\boldsymbol{DP}(a, b)$. The differential of $\boldsymbol{P}$ can be obtained by substituting monomials $x_i x_j$ by $a_i b_j + a_j b_i$ in the expression of $\boldsymbol{P}$ (if $\boldsymbol{P}$ is not homogeneous, terms of degree 1 and 0 are discarded). The interest of doing so is that $\boldsymbol{DP}$ is linear separately in a and b and its properties relatively to these variables can then be described in terms of linear algebra. Furthermore, when considering a multivariate scheme $\boldsymbol{P} = T \circ F \circ S$, these properties are isomorphic to those of F since S and T are linear bijections.

Recently, Dubois, Fouque, Shamir and Stern showed a very efficient cryptanalysis of C*$^-$ schemes based on a class of geometrical invariants of the differential of C* [4,3]. We summarize it below.

3.1 Skew-symmetric Maps with respect to the Differential

The differential of the internal C* function is $DF(a,b) = a\,b^{q^\theta} + a^{q^\theta} b$ for $a, b \in \mathbb{F}_{q^n}$. When a and b are identified with n coordinates vectors over $\mathbb{F}_q$, DF is a bilinear symmetric function from $(\mathbb{F}_q)^n \times (\mathbb{F}_q)^n$ to $(\mathbb{F}_q)^n$. Each of the n coordinates of DF is a multivariate polynomial in the coordinates $a_1, \ldots, a_n$ and $b_1, \ldots, b_n$ of a and b respectively, which is linear separately in a and b, and where a and b play symmetric roles. Each such polynomial is written on the basis of terms $a_i b_j + a_j b_i$ so it has $n(n-1)/2$ coefficients. Now, it is observed in [4] that linear maps consisting of *multiplications* by some element ξ of $\mathbb{F}_{q^n}$ have a specific action on DF. Indeed, we have

$$DF(\xi.a, b) + DF(a, \xi.b) = (\xi + \xi^{q^\theta}).DF(a,b) \tag{1}$$

For the particular elements ξ such that $\xi + \xi^{q^\theta} = 0$ (at least 1 is solution), the associated multiplication maps M_ξ satisfy

$$DF(M_\xi(a), b) + DF(a, M_\xi(b)) = 0$$

that is, they are the *skew-symmetric* maps with respect to DF. The existence of non-trivial (*i.e.* not colinear to the identity) such maps is of course very unusual and even for a C* monomial it does not happen for all parameters. However, even when it does not happen, the initial identity can also be interpreted as a skew-symmetry property. Let us indeed define for any linear map M, the skew-symmetric action of M over DF as the bilinear and symmetric function

$$\Sigma[M](a,b) = DF(M(a), b) + DF(a, M(b))$$

Our basic identity infers that in the special case of multiplication maps,

$$\Sigma[M_\xi](a,b) = M_\zeta \circ DF(a,b)$$

where M_ζ is the multiplication by $\xi + \xi^{q^\theta}$. As a consequence, for any element ξ of $\mathbb{F}_{q^n}$, the coordinates of the bilinear and symmetric function $\Sigma[M_\xi](a,b)$ are linear combinations of the coordinates of DF. Therefore, expressed in geometrical terms, multiplication maps have the specific property to leave unchanged under skew-symmetric action the subspace spanned by the coordinates of DF. Note that this property is very strong because the subspace spanned by the n coordinates of DF has dimension at most n while for a random linear map M, the coordinates of $\Sigma[M]$ might be any polynomials in the whole space of bilinear symmetric polynomials of dimension $n(n-1)/2$ and are very unlikely to all be confined in the tiny subspace spanned by the coordinates of DF.

The public key $\boldsymbol{P}$ of a C* scheme inherits of the above properties; the only difference is that the linear maps that play with regards to $\boldsymbol{P}$ the role of multiplications with regards to F are the conjugates $S^{-1} \circ M_\xi \circ S$. Now, a crucial point is : although the latter maps depend on the secret bijection S, they can be computed from their characteristic property with regards to the public key $\boldsymbol{P}$. For instance, considering the simple skew-symmetry condition, $\boldsymbol{DP}(\boldsymbol{M}(a), b) + \boldsymbol{DP}(a, \boldsymbol{M}(b)) = 0$, we see that this equation is linear in $\boldsymbol{M}$. It can be seen [4] that each coordinate of $\boldsymbol{DP}$ provides us with $n(n-1)/2$ linear conditions on the n^2 coefficients of $\boldsymbol{M}$. Then, even a marginal number of coordinates of the public key allows to solve the space of skew-symmetric maps. Solving the more general skew-symmetry condition follows similar principles although more theory is involved; we refer the reader to the original paper [3] for the details.

3.2 Consequences

The properties described above allow an attacker to compute from a C^{*-} public key conjugates $S^{-1} \circ M_\xi \circ S$ of multiplications maps M_ξ. This of course is very annoying because these maps depend on the secret bijection S and were initially considered as secret information. Furthermore, it is shown in [4] that the nature of these maps is an additional problem. We do not consider these aspects here and focus on the initial breach *i.e.* the existence of linear maps which can be computed from the public key although they contain secret information. In the sequel, we investigate the possibility to destroy the skew-symmetry property of C^{*-} schemes.

4 Breaking the Symmetry

As we have seen, for C^{*-} schemes, the linear maps which are associated to the skew-symmetry property are connected to the internal field structure, namely they are multiplications by elements of $\mathbb{F}_{q^n}$. In principle, this means that the existence of these maps is tied to the internal field structure. A natural question is: would skew-symmetric maps exist if the internal field structure were truncated, *i.e.* restricted to a subspace of it?

4.1 Projection Breaks the Skew-Symmetry Property of C*$^-$ schemes

Suppose we consider the internal function F restricted to some proper subspace H of $\mathbb{F}_{q^n}$. We denote F_H this restriction. The skew-symmetric maps with respect to the differential DF_H of F_H are by definition the linear maps M_H from H to itself which satisfy :

$$DF_H(M_H(h), k) + DF_H(h, M_H(k)) = 0\ , \quad h, k \in H \tag{2}$$

We expect the solutions M_H to this condition to be the restrictions to H of the skew-symmetric maps w.r.t DF *which map* H *to itself*. When H is an arbitrary subspace, we do not expect non-trivial multiplications M_ξ to map H into itself. Then, the only solutions to our condition should be the scalar multiples of the Identity: $M_H = \lambda.Id_H, \lambda \in \mathbb{F}_q$. Let us now show that our expectation is correct using mathematical arguments. First, we characterize the linear maps M_H which are skew-symmetric with respect to DF_H by transforming the above condition (2) in a condition with respect to DF. That is, we embed the above condition over H in a condition over $\mathbb{F}_{q^n}$. We can embed M_H into a linear map $\bar{M}_H$ which is M_H over H and zero elsewhere. The same way, we can embed the Identity over H into the projection map to H, denoted π_H. Then, (2) is equivalent to:

$$DF(\bar{M}_H(a), \pi_H(b)) + DF(\pi_H(a), \bar{M}_H(b)) = 0\ , \quad a, b \in \mathbb{F}_{q^n}$$

Therefore, the linear maps $\bar{M}_H$ are special solutions to the condition

$$DF(M(a), \pi_H(b)) + DF(\pi_H(a), M(b)) = 0\ , \quad a, b \in \mathbb{F}_{q^n} \tag{3}$$

They are those solutions M left unchanged by composition with π_H :

$$M = M \circ \pi_H = \pi_H \circ M$$

Our method to determine the linear maps $\bar{M}_H$ is then clear : we first find the solutions M to the condition (3), and then find those which are left unchanged by composition with π_H.

Before we do this, let us note an alternative characterization of the linear maps $\bar{M}_H$: they are the common solutions of the two conditions

$$\begin{aligned} &DF(M \circ \pi_H(a), \pi_H(b)) + DF(\pi_H(a), M \circ \pi_H(b)) = 0\ , \\ &DF(\pi_H \circ M(a), \pi_H(b)) + DF(\pi_H(a), \pi_H \circ M(b)) = 0\ , \end{aligned} \qquad a, b \in \mathbb{F}_{q^n} \tag{4}$$

The first condition is the skew-symmetry condition with respect to DF only considered for elements of H. The second condition is the skew-symmetry with respect to $DF(\pi_H, \pi_H)$. Both conditions have additional degrees of freedom compared to the skew-symmetry with respect to DF, and are simultaneously satisfied by the only linear maps $\bar{M}_H$.

The Solutions to Condition 3. As we can see, obvious solutions to Condition 3 are the maps $M_\xi \circ \pi_H$ where M_ξ is skew-symmetric with respect to DF. Since our condition is greatly overdetermined, we do not expect any other solutions. This is confirmed experimentally. In the most simple case when H is a hyperplane, we can give it a mathematical proof.

Lemma 1. *Let H be a hyperplane of $\mathbb{F}_{q^n}$ and DF be the differential of a bijective C^* monomial. The linear maps M which satisfy the condition*

$$DF(M(a), \pi_H(b)) + DF(\pi_H(a), M(b)) = 0\ , \quad a, b \in \mathbb{F}_{q^n}$$

are of the form $M_\xi \circ \pi_H$ where M_ξ is skew-symmetric with respect to DF.

Proof. The idea of the proof is to replace M and π_H by their expressions as sums of q-powerings, and to express our condition as the vanishing of a polynomial in a, b over $\mathbb{F}_{q^n}$. We have $M(a) = \sum_{i=0}^{n-1} m_i\, a^{q^i}$. The hyperplane H is the kernel of a linear form. On the other hand, any linear form has the form $a \mapsto tr(ua)$ for some element u in $\mathbb{F}_{q^n}$ and where tr denotes the trace operator, $tr(a) = \sum_{i=0}^{n-1} a^{q^i}$. Since multiplication by u is a linear change of coordinates, one can suppose $u = 1$. Since $a \mapsto a^q - a$ maps $\mathbb{F}_{q^n}$ to the kernel of tr, we can give the explicit form $a^q - a$ to the elements of H. We can rewrite our condition : $A(a, b) - B(a, b) = 0$, where

$$\begin{aligned} A(a, b) &= DF(M(a), b) + DF(a, M(b)) \\ B(a, b) &= DF(M(a), b^q) + DF(a^q, M(b)) \end{aligned}$$

Both expressions are written on the basis of symmetric terms of the form $a^{q^i} b^{q^j} + a^{q^j} b^{q^i}$ and their respective coefficients are :

$$\begin{aligned} &A(a, b) : \text{coefficient}\{i, 0\} = m_{i-\theta}^{q^\theta}\ \ ;\ \text{coefficient}\{i, \theta\} = m_i \\ &B(a, b) : \text{coefficient}\{i, 1\} = m_{i-\theta}^{q^\theta}\ \ ;\ \text{coefficient}\{i, \theta + 1\} = m_i \end{aligned}$$

From the coefficient of $a^{q^\theta}b + ab^{q^\theta}$, we find $m_0 + m_0^{q^\theta} = 0$. From the coefficient of $a^{q^{\theta+1}}b + ab^{q^{\theta+1}}$, we find $m_1^{q^\theta} = m_0$. From the coefficient of $a^{q^{\theta+i}}b + ab^{q^{\theta+i}}, i \notin \{0, 1\}$, we find $m_i = 0$. Denoting $m_0 = m_1$ by ξ we have $M(a) = \xi(a^q - a)$ where $\xi^{q^\theta} + \xi = 0$. □

Solutions which are Left Unchanged by Composition with the Projection. As we have shown, the linear maps $\bar{M}_H$ which correspond to the skew-symmetric maps with respect to DF_H, are the solutions to Condition 3 which are left unchanged by composition with π_H. As argued in the previous section, the solutions to this condition are $M_\xi \circ \pi_H$ where M_ξ is multiplication by some element ξ. These maps are unchanged by composition with π_H if and only if M_ξ commutes with π_H, *i.e.* if and only if M_ξ maps H to itself. Then, since for any ξ, M_ξ is bijective, we have $\xi.H = H$. Our goal is to show that, except for specific choices of H which are very sparse, the only ξ satisfying this property are the scalar multiples of 1. As a first step, we notice that these elements ξ form a multiplicative group, independently of the choice of H. Therefore, they actually form a subfield of $\mathbb{F}_{q^n}$ and H is a linear space over this subfield. Finally, the subspaces H for which our property is satisfied by non-trivial elements ξ are subspaces over intermediate subfields of $\mathbb{F}_{q^n}$. As a second step, we upperbound the probability that a random subspace H of a prescribed dimension s is a subspace over an intermediate subfield of $\mathbb{F}_{q^n}$. (In this case, we say that H is degenerate). We show that this probability is negligible in terms of q and n.

Lemma 2. *Degenerate subspaces of $\mathbb{F}_{q^n}$ only exist at dimensions s not coprime with n. In particular, degenerate hyperplanes never exist. The proportion of degenerate subspaces in $\mathbb{F}_{q^n}$ of a prescribed dimension is always at most $\mathcal{O}(q^2 q^{-n})$.*

Proof. When H is a subspace over $\mathbb{F}_{q^r}$, its dimension over $\mathbb{F}_q$ is a multiple of r. Since r must itself be a divisor of n, degenerate subspaces only exist at dimensions s not coprime with n. For instance, degenerate hyperplanes never exist since $n-1$ is always coprime with n. Let r be a common divisor of s and n. It can be shown that the number of subspaces of dimension s in a vector space of dimension n is of the order of $q^{s(n-s)}$ [9]. Then, the number of $\mathbb{F}_{q^r}$-subspaces of dimension s/r in $\mathbb{F}_{q^n}$ is of the order of $q^{s(n-s)/r}$. The number of degenerate subspaces of dimension s in $\mathbb{F}_{q^n}$ is dominated by the latter quantity considered for the smallest common factor r of n and s. Since the smallest possible value of r is 2, the proportion of degenerate subspaces of dimension s in $\mathbb{F}_{q^n}$ is at most of the order of $q^{-s(n-s)/2}$. Since $s(n-s)$ is minimal for $s = 2$ (2 is a common factor of s and n), the searched proportion is dominated by $q^{-(n-2)}$. $\square$

Application to the General Skew-Symmetry Property of C*$^-$ schemes. In the preceding paragraphs, we have shown that restricting the internal function F to some proper subspace H of $\mathbb{F}_{q^n}$ destroys the simple skew-symmetry property (2). In this paragraph, we consider the general skew-symmetry property of C*$^-$ schemes. This property expresses that there exists non-trivial linear maps which leave the space spanned by the coordinates of DF unchanged under skew-symmetric action. The linear maps satisfying this condition are the whole space of multiplications. Using similar techniques as before, we can show that this property considered for the restricted function F_H admits only trivial solutions. We refer the reader to the appendix for the details.

4.2 Experimental verifications

We checked experimentally, for various C* parameters n and θ, the effect of restricting the internal function to a randomly chosen subspace H of various dimensions s. For instance, for parameters $n = 36$ and $\theta = 4$ (which are more interesting than those of SFLASH since they are not prime numbers), we obtain the table below for the dimension of the solution space of the general skew-symmetry condition as the number of coordinate-wise conditions grows.

# conditions	$s = 0$	$s = 1$	$s = 2$	$s = 3$	$s = 4$	$s = 9$	$s = 18$
1	1296	1225	1156	1089	1024	769	324
2	708	669	632	598	564	414	207
3	168	145	124	109	104	99	90
4	36	1	1	1	1	1	1
6	36	1	1	1	1	1	1
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

5 Projected C*$^-$ schemes

Based on the previous results, we are led to define a new family of schemes that we call *projected* C^{*-} schemes. As we will see, these schemes actually consists in hiding a C* monomial using non-bijective linear maps. We next define the (ad-hoc) computational problems on which the security of these schemes is based. Finally, we discuss possible choices of parameters and suggest one concrete choice with performances comparable to SFLASH.

Description. A projected C*$^-$ scheme is defined as follows. Start from a C* scheme $F(x) = x^{1+q^\theta}$ with secret linear maps S and T. Let r and s be two integers between 0 and n. Let T^- be the projection of T on the last r coordinates and S^- be the restriction of S on the last s coordinates. Compute $\hat{\boldsymbol{P}} = T^- \circ F \circ S^-$. The generated function $\hat{\boldsymbol{P}}$ is used as the public key and the secret linear bijections S and T are used as the secret key. Note that $\hat{\boldsymbol{P}}$ is a quadratic function from $(\mathbb{F}_q)^{n-s}$ to $(\mathbb{F}_q)^{n-r}$. To find a preimage by the public function of a given message m, the legitimate user first pads m with a random vector m' of $(\mathbb{F}_q)^r$ and compute the preimage of (m, m') by $T \circ F \circ S$. If this element has its last s coordinates to 0, then its $n - s$ first coordinates are a valid signature for m. Otherwise, he discards this element and tries with an other random padding m'. When $r > s$, the process ends with probability 1 and costs on average q^s inversions of F. In practice, r is chosen a significant fraction of n to make the public key resistant to algebraic attacks; s can be chosen as small as 1 to destroy symmetries arising from the internal field structure. As for C*$^-$ schemes, the significant value of r makes projected C*$^-$ schemes only suitable for signature, since reviewing all possible paddings m' is not efficient. Finally, we mention that projection already appeared in the literature as a possible modifier [17] but was never considered as a useful measure let alone a defensive measure.

Possible Angles of Analysis. As usual for multivariate schemes, the security relies on several ad-hoc computational problems. The first problem is solving the public system of quadratic equations. Since s is chosen small, this is about as hard as solving the initial C*$^-$ system. The second problem is recovering the functional decomposition of the public key or at least some information on the secret maps S^-, T^-. There is no efficient strategy to solve this problem in general [7], and the attack by Dubois *et al.* which falls into this category for C*$^-$ schemes is here prevented by the projection. Remains the strategy consisting in recovering the public key into a valid C*$^-$ public key. Showing this to be possible is actually the new challenge opened by the new family of schemes.

Parameters. n, θ, r are chosen following the rationales for C^{*-} schemes. We choose $s = 1$ as it induces the minimal factor q on the secret operations. The value of q can be chosen small but, at constant blocksize, this requires a larger value of n and therefore a larger public key. As a possible trade-off, we propose pFLASH with $q = 2^4$, $n = 74$, $\theta = 11$, $r = 22$ and $s = 1$. Our tests have pFLASH signing at $\lesssim 1$ million K8/C2 cycles, in line with expectations of $\sim 16\times$ time of SFLASH [2]; private key size is $2\times$ at 5.4kB. These are still attractive features for small device implementation.

6 Conclusion

In this paper, we provide additional insight on the recent cryptanalysis of SFLASH by exhibiting a simple modification which provably avoids the attack. Our study shows that the attack against SFLASH has deeper roots than the mere fact that it is based on a C* monomial : the attack is made possible because the large field structure is embedded in the public key and is stopped when it is no more the case. Then, we realize that, indeed, one might not hope to hide effectively a particular function defined on a large field using linear bijections; this might at most be achievable in some security range using compressive linear maps. But then, is it still possible to build a practical cryptosystem in this setting ? At the present state, we can still define a modified family of C* -based schemes which is of practical interest. Analysis of this most simple case would probably yield additional understanding of the ways to distinguish a specifically-built multivariate function and would provide further insight on the very possibility to obfuscate such a function using linear maps.

Acknowledgment

JD and BY are grateful to the Alexander van Humboldt Foundation, the Taft Fund, and TWISC [National Science Council project NSC 96-2219-E-011-008 / NSC 96-2219-E-001-001] without whose valuable support much of this work would've not happened. BY, CC, OC would also like to thank NSC for partial sponsorship via project NSC 96-2221-E-001-031-MY3.

References

1. *European project IST-1999-12324 on New European Schemes for Signature, Integrity and Encryption.* `http://www.cryptonessie.org`.
2. Daniel J. Bernstein. eBATs benchmark results. `http://ebats.cr.yp.to`.

3. Vivien Dubois, Pierre-Alain Fouque, Adi Shamir, and Jacques Stern. Practical Cryptanalysis of SFLASH. In Alfred Menezes, editor, *CRYPTO*, volume 4622 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2007.
4. Vivien Dubois, Pierre-Alain Fouque, and Jacques Stern. Cryptanalysis of SFLASH with Slightly Modified Parameters. In *Proceedings of Eurocrypt 2007*, volume LNCS 4515, pages 264–275, 2007.
5. Vivien Dubois, Louis Granboulan, and Jacques Stern. An Efficient Provable Distinguisher for HFE. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 156–167. Springer, 2006.
6. Vivien Dubois, Louis Granboulan, and Jacques Stern. Cryptanalysis of HFE with Internal Perturbation. In *Proceedings of PKC 2007*, volume LNCS 4450, pages 249–265. Springer, 2007.
7. Jean-Charles Faugère and Ludovic Perret. Polynomial Equivalence Problems: Algorithmic and Theoretical Aspects. In Serge Vaudenay, editor, *EUROCRYPT*, volume 4004 of *Lecture Notes in Computer Science*, pages 30–47. Springer, 2006.
8. Pierre-Alain Fouque, Louis Granboulan, and Jacques Stern. Differential Cryptanalysis for Multivariate Schemes. In Ronald Cramer, editor, *EUROCRYPT*, volume 3494 of *Lecture Notes in Computer Science*, pages 341–353. Springer, 2005.
9. Jay Goldman and Gian-Carlo Rota. The Number of Subspaces of a Vector Space. In W.T.Tutte, editor, *Recent Progress in Combinatorics*, pages 75–83. Academic Press, 1969.
10. Tsutomu Matsumoto and Hideki Imai. Public Quadratic Polynominal-Tuples for Efficient Signature-Verification and Message-Encryption. In *EUROCRYPT*, pages 419–453, 1988.
11. Robert J. McEliece. A Public-Key Cryptosystem based on Algebraic Coding Theory. In *JPL DSN Progress Report*, pages 114–116, California Inst. Technol., Pasadena, 1978.
12. NESSIE, New European Schemes for Signatures, Integrity, and Encryption. *Portfolio of Recommended Cryptographic Primitives.* `http://www.nessie.eu.org`.
13. Jacques Patarin. Cryptanalysis of the Matsumoto and Imai Public Key Scheme of Eurocrypt'88. In Don Coppersmith, editor, *CRYPTO*, volume 963 of *Lecture Notes in Computer Science*, pages 248–261. Springer, 1995.
14. Jacques Patarin, Louis Goubin, and Nicolas Courtois. $C^{*}{}_{-+}$ and HM: Variations Around Two Schemes of T. Matsumoto and H. Imai. In Kazuo Ohta and Dingyi Pei, editors, *ASIACRYPT*, volume 1514 of *Lecture Notes in Computer Science*, pages 35–49. Springer, 1998.
15. Adi Shamir. Efficient Signature Schemes Based on Birational Permutations. In Douglas R. Stinson, editor, *CRYPTO*, volume 773 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 1993.
16. Specifications of SFLASH. *Final Report NESSIE*, pages 669–677. 2004.
17. Christopher Wolf and Bart Preneel. Taxonomy of Public Key Schemes based on the problem of Multivariate Quadratic equations. ePrint Archive Report 2005/077. `http://eprint.iacr.org/2005/077`.

A Projection Also Breaks the General Skew-Symmetry Property

As before, we denote F_H the restriction of the C* function F to a proper subspace H of $\mathbb{F}_{q^n}$. Note that F_H is a function from H to $\mathbb{F}_{q^n}$. A linear map M_H from H to itself satisfies the general skew-symmetry property with respect to DF_H if and only if there exists an associated linear map N_{M_H} from $\mathbb{F}_{q^n}$ to itself such that

$$DF_H(M_H(h), k) + DF_H(h, M_H(k)) = N_{M_H} \circ DF_H(h, k) , \quad h, k \in H \quad (5)$$

As before, we can embed this identity over H into an identity over $\mathbb{F}_{q^n}$. We denote $\bar{M}_H$ the linear map from $\mathbb{F}_{q^n}$ to itself which is M_H over H and zero elsewhere. We denote π_H the projection to H. Identity (5) is equivalent to

$$DF(\bar{M}_H(a), \pi_H(b)) + DF(\pi_H(a), \bar{M}_H(b)) = N_{M_H} \circ DF(\pi_H(a), \pi_H(b))$$

for any a, b in $\mathbb{F}_{q^n}$. The linear maps $\bar{M}_H$ are special solutions to the condition

$$DF(M(a), \pi_H(b)) + DF(\pi_H(a), M(b)) = N_M \circ DF(\pi_H(a), \pi_H(b)) \quad (6)$$

for any a, b in $\mathbb{F}_{q^n}$. They are those solutions M which are left unchanged by composition with π_H:

$$M = M \circ \pi_H = \pi_H \circ M$$

Obvious solutions to Condition (6) are the maps $M_\xi \circ \pi_H$ where M_ξ is multiplication by an element ξ of $\mathbb{F}_{q^n}$. Since Condition (6) is greatly overdetermined, we do not expect any parasitic solutions, and this is confirmed in practice. When H is a hyperplane, we can actually give it a mathematical proof (see below). Then, the maps $M_\xi \circ \pi_H$ left unchanged by composition with π_H are those for which H is closed by multiplication by ξ. We know from Lemma 2 that except for negligibly sparse choices of H, the only elements ξ which satisfy this property are the scalar multiples of 1.

Lemma 3. *Let H be a hyperplane of $\mathbb{F}_{q^n}$ and DF be the differential of a bijective C* monomial. The linear maps M for which there exists a linear map N_M such that, for any a, b in $\mathbb{F}_{q^n}$,*

$$DF(M(a), \pi_H(b)) + DF(\pi_H(a), M(b)) = N_M \circ DF(\pi_H(a), \pi_H(b))$$

are the $M_\xi \circ \pi_H$ where M_ξ is multiplication by an element ξ of $\mathbb{F}_{q^n}$.

Proof. The proof is analogous to the proof of Lemma 1. Recall that, up to a linear change of coordinates, elements of H have the form $a^q - a$ for a in $\mathbb{F}_{q^n}$. Let us also recall that bijective C* monomials only exist in characteristic 2. We rewrite our condition $A(a,b) - B(a,b) = C(a,b) - D(a,b)$ where $A(a,b)$ and $B(a,b)$ are the same as in the proof of Lemma 1, and

$$C(a,b) = N(DF(a^q,b)) + DF(a,b^q))$$
$$D(a,b) = N(DF(a^q,b^q)) - DF(a,b))$$

Both expressions are written on the basis of symmetric terms of the form $a^{q^i}b^{q^j} + a^{q^j}b^{q^i}$ and their respective coefficients are, in terms of the coefficients $m_0,\dots,m_{n-1}$ of M and $n_0,\dots,n_{n-1}$ of N:

$A(a,b)$: coefficient$\{i,0\} = m_{i-\theta}^{q^\theta}$; coefficient$\{i,\theta\} = m_i$
$B(a,b)$: coefficient$\{i,1\} = m_{i-\theta}^{q^\theta}$; coefficient$\{i,\theta+1\} = m_i$
$C(a,b)$: coefficient$\{\theta+1+i,i\} = n_i$; coefficient$\{\theta+i,i+1\} = n_i$
$D(a,b)$: coefficient$\{\theta+i,i\} = n_{i-1} - n_i$

From the coefficient of $a^{q^\theta}b + ab^{q^\theta}$, we find

$$m_0 + m_0^{q^\theta} = n_{-1} - n_0.$$

From the coefficient of $a^{q^{\theta+1}}b + ab^{q^{\theta+1}}$, we find

$$m_1^{q^\theta} - m_0 = n_0.$$

From the coefficient of $a^{q^{\theta+i}}b + ab^{q^{\theta+i}}, i \notin \{0,1\}$, we find

$$m_i = 0.$$

On the other hand, from the coefficient of $a^{q^{\theta+1+i}}b^{q^i} + a^{q^i}b^{q^{\theta+1+i}}, i \neq 0$, we find

$$n_i = 0.$$

Denote $m_0 = m_1$ by ξ. We end up with $M(a) = \xi(a^q - a)$ and $N(a) = (\xi^{q^\theta} + \xi)a$. □