An Updated Security Analysis of PFLASH

Ryann Cartor[1] and Daniel Smith-Tone[1,2]

[1]Department of Mathematics, University of Louisville,
Louisville, Kentucky, USA
[2]National Institute of Standards and Technology,
Gaithersburg, Maryland, USA

ryann.cartor@louisville.edu, daniel.smith@nist.gov

Abstract. One application in post-quantum cryptography that appears especially difficult is security for low-power or no-power devices. One of the early champions in this arena was SFLASH, which was recommended by NESSIE for implementation in smart cards due to its extreme speed, low power requirements, and the ease of resistance to side-channel attacks. This heroship swiftly ended with the attack on SFLASH by Dubois et al. in 2007. Shortly thereafter, an old suggestion re-emerged: fixing the values of some of the input variables. The resulting scheme known as PFLASH is nearly as fast as the original SFLASH and retains many of its desirable properties but without the differential weakness, at least for some parameters.

PFLASH can naturally be considered a form of high degree HFE^- scheme, and as such, is subject to any attack exploiting the low rank of the central map in HFE^-. Recently, a new attack has been presented that affects HFE^- for many practical parameters. This development invites the investigation of the security of PFLASH against these techniques.

In this vein, we expand and update the security analysis of PFLASH by proving that the entropy of the key space is not greatly reduced by choosing parameters that are provably secure against differential adversaries. We further compute the complexity of the new HFE^- attack on instances of PFLASH and conclude that PFLASH is secure against this avenue of attack as well. Thus PFLASH remains a secure and attractive option for implementation in low power environments.

Key words: Multivariate Cryptography, HFE, PFLASH, Discrete Differential, MinRank

1 Introduction

In December of 2016, the National Institute of Standards and Technology (NIST) published an open call for proposals for new post-quantum standards for some of the most critical security applications in digital communication infrastructure, see [1]. The post-quantum technologies this project aspires to vet and standardize

are designed to be secure against adversaries with access to quantum computing devices— machines capable of acheiving exponential speed-up over classical computers on certain problems, see [2].

Many avenues to post-quantum security are developing, including techniques from lattice theory, coding theory and algebraic geometry. Each of these areas enjoy hard computational problems that have been studied extensively and have histories going back many decades. They also share the common trait that the fundamental computational problems in these fields have no known significant speed-up in the quantum paradigm.

One of the hard computational problems on which the security of many post-quantum cryptosystems is based is the problem of solving systems of multivariate equations. Generically, solving systems of multivariate quadratic equations is hard, so a valid technique for constructing a cryptosystem is to find a class of quadratic vector-valued functions on a vector space that is easy to invert, and transform it into a system that appears random.

Both of these tasks present challenges. The standard technique for the second task is computing a morphism of the system in an attempt to remove the properties allowing the system to be inverted. Techniques for the prior task are more varied, and in this work our focus is on a particular big field scheme.

1.1 Prior Work

The progenitor of all "big field" schemes is commonly known as C^*, or the Matsumoto-Imai scheme, see [3]. This scheme exploits the vector space structure of extension fields to provide two versions of a function— a vector-valued version which is quadratic over the base field, and a monomial function whose input and output lie in the extension field. The cryptanalysis of this scheme by Patarin in [4] inspired many big field constructions.

In [5], Patarin introduced the Hidden Field Equations (HFE) cryptosystem, a natural generalization of the monomial based C^* in which the monomial map is replaced with a low degree polynomial. Also described in the above work is the minus modifier— the removal of public equations— which can be applied to both HFE, producing HFE$^-$, and to C^*, creating C^{*-}.

A popular iteration of C^{*-} was SFLASH, see [6], which was very efficient, but unfortunately insecure. An attack by Dubois et al. in [7] broke SFLASH by way of a symmetric differential relation present in the central monomial map.

In [8], a way to resist the attack on SFLASH is presented. The augmentation of the scheme, known as projection, fixes the value of d of the input variables producing a scheme we now call PFLASH. PFLASH is still a very fast signature scheme and is amenable to low-power environments without sacrificing side-channel resistance. This projected C^{*-} system is shown to resist differential cryptanalysis for restricted parameters, that is, when the degree is bounded by $q^{n/2-d}$, in [9] and is fully specified with paractical parameters in [10].

Since the design of PFLASH there have been a number of cryptanalytic developments in the big field venue. The development of differential invariant attacks in [11] and their further application in [12] are examples of advancement

in this active area. Furthermore, the improved efficiency of the Kipnis-Shamir (KS) attack of [13] presented in [14] is directly impactful to PFLASH, as one can consider PFLASH as a possibly high degree but still low rank version of HFE$^-$.

1.2 Our Contribution

We expand and update the analysis in [9] and [10] proving resistance to differential and rank techniques for the vast majority of parameters, and verifying that the provably secure key spaces are not as severely limited as the previous works suggest. This improvement is directly impactful, providing further assurance that attacks based on equivalent keys cannot weaken PFLASH.

The degree bound restriction in [9] reduces the dimension of possible private keys by a factor of more than two. Our updated differential analysis verifies the security of the scheme when the central map has no degree bound, and thus assures us that very little entropy is lost in the key space when restricting to parameters that are provably secure against differential adversaries.

In [10], an argument for the resistance of PFLASH to the technique of [14, Section 8.2] when PFLASH is considered as a low degree projected HFE$^-$ scheme is provided. We make this assessment more robust by also considering the possibility of an adversary attempting to remove the projection modifier from PFLASH considering it to be a higher rank HFE$^-$ scheme. Whereas in the former case, the attack is impossible, in the latter case, the algebraic structure allows the possibility that the attack can succeed; however, the complexity of the attack is directly computed and shown to be infeasible.

1.3 Organization

The paper is organized as follows. The next section introduces the notion of big field schemes and provides the description of those schemes relevant to this work, namely, C^*, $PFLASH$ and HFE. In the following section, we review the cryptanalytic techniques that have proven most successful in attacking big field schemes. The subsequent two sections provide a new proof of security against differential attacks for PFLASH, first by analyzing the projected C^* primitive and then by extending these results to the full scheme. We then conclude, noting parameter choices for PFLASH and discussing applications of the scheme.

2 Big Field Schemes

Many multivariate cryptosystems utilize the structure of a degree n extension $\mathbb{K}$ of a finite field $\mathbb{F}_q$ as an $\mathbb{F}_q$-algebra. Such cryptosystems are collectively known as "big field" schemes. To emphasize a choice of basis, one chooses an $\mathbb{F}_q$-vector space isomorphism $\phi : \mathbb{F}_q^n \to \mathbb{K}$. There is then an equivalence between systems F of n quadratic polynomials in n variables over $\mathbb{F}$ and univariate polynomials of the form

$$f(x) = \sum_{0 \le i \le j < n} \alpha_{ij} x^{q^i + q^j}$$

over $\mathbb{K}$ given by $F = \phi^{-1} \circ f \circ \phi$.

To hide the structure of an easily invertible map, the standard technique is to apply an isomorphism of polynomials to mask the choice of basis for the input and output of f.

Definition 1 *A* polynomial morphism *between two systems of polynomials is a pair of affine maps* (T, U) *such that* $G = T \circ F \circ U$. *If both* T *and* U *are invertible, then the morphism is said to be an isomorphism and* F *and* G *are said to be isomorphic.*

Thus, for big field schemes, the construction of a public key can be summarized with the following diagram.

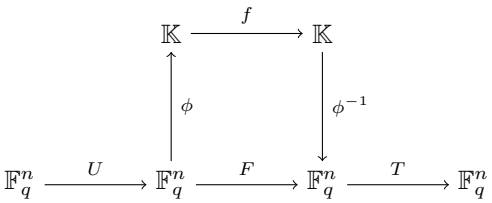

2.1 *C**

Matsumoto and Imai discovered massively multivariate cryptography, introducing the scheme now known as C^* at Eurocrypt '88. The $C^*(q, n)$ scheme is a big field construction in which the vector-valued representation of a quadratic monomial map $f(x) = x^{q^\theta+1}$ is hidden by an isomorphism. Thus the public key is given by $P = T \circ \phi^{-1} \circ f \circ \phi \circ U$.

The C^* scheme was originally envisioned for encryption, but could quite apparently be applied in either encryption or digital signatures. To encrypt (or to verify a signature), one simply computes the output of the public function P. To decrypt (or to sign), the preimage must be determined successively for each of the components of the private key, all of which can be computed efficiently. The interesting step, the inversion of f can be accomplished by noticing that if $b(q^\theta + 1) = 1 \ (mod\ q^n - 1)$, then $(x^{q^\theta+1})^b = x$.

2.2 PFLASH

The PFLASH scheme is a particular parametrization of a projected C^{*-} scheme. The projection and minus modifiers were both originally suggested in reference to C^* in [15]. The idea of projection is to fix the value of d input variables to change the simplicity of the central map. Thus the composition of the projection and the affine map U form a projection onto a codimension d hyperplane. The minus modification removes r equations from the public key. Thus the composition of this projection with T has corank r. The public key of PFLASH(q, n, r, d) is given by $P = \pi_r \circ T \circ \phi^{-1} \circ f \circ \phi \circ U \circ \pi_d$.

We note that the public key is no longer isomorphic to the private monomial function. Instead there is merely a polynomial morphism between the central map and the public key. Since it is well-known that the morphism of polynomials problem is NP-hard, see [16], there is some hope that the information lost to the public key may secure the scheme.

Mechanically, the scheme works as a digital signature primitive as follows. Verification is accomplished by evaluating the public polynomials at the signature. Signing is done by finding preimages of each of the private maps. To find a preimage of $\pi_r \circ T\phi^{-1}$, randomly append r values to the message, then apply T^{-1} and ϕ. Once f is inverted, an element in the preimage of $\phi \circ U$ and in the image of π_d is selected as the signature.

2.3 HFE

Hidden Field Equation (HFE) scheme of [5] is a generalization of the C^* construction, in which the monomial map is replaced by a more general polynomial with a degree bound D. Given the degree n extension $\mathbb{F} \subseteq \mathbb{K}$ we choose a quadratic polynomial $f : \mathbb{K} \to \mathbb{K}$ of degree bound D. Thus f has the form:

$$f(x) = \sum\nolimits_{q^i+q^j \leq D}^{i \leq j} \alpha_{i,j} x^{q^i+q^j} + \sum_{\substack{i \\ q^i \leq D}} \beta_i x^{q^i} + \gamma,$$

where $\alpha_{i,j}, \beta_i, \gamma \in \mathbb{K}$. The public key is then constructed via the isomorphism:

$$P = T \circ \phi^{-1} \circ f \circ \phi \circ U.$$

Inversion is accomplished by first taking a ciphertext $y = P(x)$, computing $v = T^{-1}(y)$, solving $v = f(u)$ for u via the Berlekamp algorithm, see [17], and then recovering $x = U^{-1}(u)$.

3 Cryptanalyses of Big Field Schemes

The big field multivariate cryptosystems have an extensive history in cryptanalysis. Several techniques have been developed that illustrate that it is very difficult to hide efficient inversion of a system. These techniques can largely be grouped into two categories: those based on differential propertys and those based on rank properties.

3.1 Differential Techniques

By breaking "big field" schemes and also inspiring modifiers, differential attacks have been instrumental in the development and analysis of multivariate public key cryptography. Given a field map f, the discrete differential is defined by $Df(a, x) = f(a + x) - f(a) - f(x) + f(0)$. As an operator on $\mathbb{K}$, D is $\mathbb{K}$-linear and reduces the complexity while increasing the dimension of a function. For

example, the differential of an affine map is zero, the differential of a quadratic map is bilinear, the differential of a cubic map is bi-quadratic, etc.

Patarin's linearization equations attack of [4] can be viewed as a differential attack as follows. The differential of the C^* monomial $f(x) = x^{q^\theta+1}$ is symmetric in characteristic two; hence, it is zero on the diagonal, $Df(x,x) = 0$. Therefore setting $v = f(u)$ we have

$$\begin{aligned} 0 = Df(v, f(u)) &= vu^{q^{2\theta}+q^\theta} + v^{q^\theta}u^{q^\theta+1} \\ &= u^{q^\theta}(vu^{q^{2\theta}} + uv^{q^\theta}), \end{aligned}$$

and whether or not $u = 0$, the right factor must be zero; thus, we obtain a bilinear relation between u and v. Setting $u = Ux$ and $v = T^{-1}y$, we obtain a bilinear relation between plaintext and ciphertext pairs: the linearization equations. Indeed even the higher order linearization equations (HOLEs) attacks pioneered in [18] can similarly be derived via differentials.

Another notable application of symmetric differential techniques in cryptanalysis is the attack on SFLASH of [7]. This attack exploits the fact that C^* polynomials are multiplicative. Specifically, $f(x) = x^{q^\theta+1}$ exhibits a differential symmetry.

Definition 2 *A function $f : \mathbb{K} \to \mathbb{K}$ has a differential symmetry if there exists a pair of $\mathbb{F}$-linear maps $L, \Lambda_L : \mathbb{K} \to \mathbb{K}$ such that*

$$Df(La, x) + Df(a, Lx) = \Lambda_L Df(a, x).$$

The attack uses the fact that left-multiplication maps of elements in $\mathbb{K}$ satisfy the above relation. This equality provides a criterion for the derivation of such maps, and via a linear algebra distillation technique, such a map can be efficiently recovered, and a full rank key derived.

It is important to note that once such a symmetry inducing linear map is discovered, there is no need to recover a full rank private key; an attack can be mounted directly with the recovered representation of the extension field multiplicative structure. Thus, even if a central map does not have a differential symmetry, it is possible that a minus-modified version of the scheme might; thus, an attack may be mounted directly on the choice of representation of the big field. This fact is the basis for the direct analysis of minus-modified schemes of [19] and [20].

It was shown in [21] that a quadratic map can only have the symmetry of Definition 2 with L a representation of left-multiplication by a field element when f is multiplicative; that is, when f has only one quadratic monomial. Later it was shown in [9] that the only linear maps L satisfying the above relation for C^* are the multiplication maps.

This famous cryptanalysis incited a more careful analysis of a technique originally proposed at ASIACRYPT 1998 in [15] and further suggested after the attack in [8]. The idea is to use projection, that is, to fix some of the input values, to make U singular. PFLASH, whose parameters are defined in [10], is

a particular parametrization of this structure. This change nullifies the basis of the differential symmetric attack as proven in [9] for a certain parameter set. In the resulting scheme, a pC^{*-} scheme, the central map can be made to no longer admit any symmetry. The parameter set which is provably secure against a differential adversary appears quite small, however, and considering the fact that such a scheme can be considered a special case of HFE$^-$ with perhaps a larger degree bound but an even smaller rank, it is necessary to review the rank structure of such schemes as well.

3.2 Rank Techniques

The first significant cryptanalysis of HFE was the Kipnis-Shamir (KS) attack of [13]. The attack is based on the fact that as a quadratic form over the extension field, the public key has low rank. This attack was significantly improved in [14], where minors modeling, instead of the original modeling of the rank property by Kipnis and Shamir, and Gröbner basis techniques are employed. The result is that the security of HFE is polynomial in the degree of the extension $\mathbb{K}$ over $\mathbb{F}_q$.

PFLASH can easily be characterized as an HFE$^-$ scheme with a more efficient inversion process. This characterization is possible by absorbing the projection into the central monomial map to make a more general polynomial. As an HFE$^-$ scheme, the rank of the central map is still 2, thus the central map has a very strong property. The minus modifier, however, provably increases the rank of the public key.

One may even consider PFLASH to be an HFE instance if we append zero polynomials to the public key. In this case, one should suspect that the rank of the central map would be quite high, rendering attacks such as [13] and [14] infeasible. Still, a theoretical verification of this intuition is absent in the literature.

4 Updated Differential Analysis of Projected Primitive

As discussed in [9], we may assume that the projection mapping is tied to f and consider differential symmetries of $f \circ \pi$ where π is chosen in a basis such that $deg(\pi) = q^d$. Clearly, if $f \circ \pi$ has a differential symmetry then the equation $Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$ is satisfied for some M. We can express this relation with matrix multiplication, namely

$$a^\top(\Pi^\top \mathbf{Df} M)x + a^\top(M^\top \mathbf{Df}\Pi)x = \Lambda_M[a^\top(\Pi^\top \mathbf{Df}\Pi)x],$$

where $\mathbf{Df}$ is the matrix representing Df as a bilinear form over $\mathbb{K}$, having one in the $(0, \theta)$ and $(\theta, 0)$ coordinates and zero elsewhere, where $\Pi x = \sum_{i=0}^{d} \beta_i x^{q^i}$ and where $Mx = \sum_{i=0}^{n-1} m_i x^{q^i}$

Examining this equation, we see that $a^\top(\Pi^\top \mathbf{Df} M)x + a^\top(M^\top \mathbf{Df}\Pi)x$ will have nonzero entries restricted to certain coordinates depending only on d and

θ, see Figure 1. Similarly, the right hand side of the equation, $\Pi^\top \mathbf{Df} \Pi$, has a structure dependent upon d and θ, see Figure 2. Notice, the graphs may look different depending on the choice of θ and d.

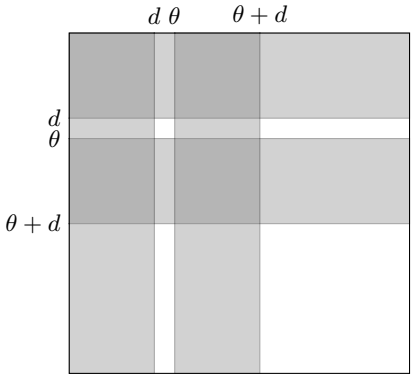

Fig. 1. The shape of the matrix representation over $\mathbb{K}$ of $Df(Ma, \pi x) + Df(\pi a, Mx)$. Shaded regions correspond to possibly nonzero values.

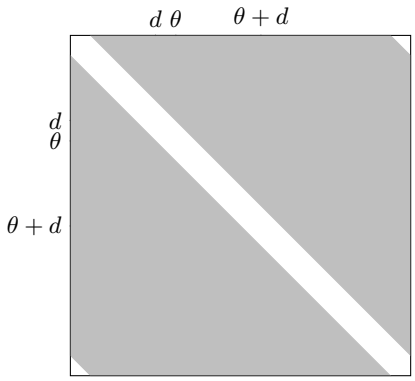

Fig. 2. The shape of the matrix representation of $\Lambda_M Df(\pi a, \pi x)$ over $\mathbb{K}$. Shaded regions correspond to possibly nonzero values.

The strategy for finding conditions on π, M and Λ_M for the existence of such a symmetry is then to find coordinates in which one side of this matrix equation is zero while the other side involves only a single unknown coefficient of M or Λ_M. While this system of equations is nonlinear in the coefficients of π, it is linear in both the unknown coefficients of M and those of Λ_M.

The system contains many more equations than variables, but certainly generates a positive dimensional ideal. The reason is that for any fixed π, $M = a\pi$ for any $a \in \mathbb{F}_q$ generates a solution. On the other hand, for a fixed π and a fixed θ, the above system becomes linear with the number of nonzero equations depending on both d and θ. Even in the best case, the number of equations is far larger than the number of variables. Since the coefficients of π are the only source of randomness for this system of linear equations, the great number of equations are not independent in a probabilistic sense. Therefore, probabilistic arguments are difficult, though extensive experiments show that the solution space is generally one dimensional.

Luckily, we can do better by bootstrapping the result of [9]. Specifically, we examine the case when $\theta > \frac{n}{2}$.

Lemma 1. *$f(x^{q^\rho}) = f(x)^{q^\rho}$ when $f(x) = x^{q^\theta+1}$*

Proof. $f(x^{q^\rho}) = (x^{q^\rho})^{q^\theta+1} = x^{(q^\theta+1)q^\rho} = \left(x^{q^\theta+1}\right)^{q^\rho} = f(x)^{q^\rho}$

Consider the special case of Lemma 1 when $\rho = -\theta$. After applying this map to the output of $\mathbf{Df}$, the nonzero terms, originally in the $(\theta, 0)$ and $(0, \theta)$

coordinates, are transported to the $(0, -\theta)$ and $(-\theta, 0)$ coordinates, respectively. This observation leads to the following theorem, revealing that most parameters of PFLASH are provably secure against a differential adversary.

Theorem 1. *Let $f(x) = x^{q^\theta+1}$ be a C^* map, and let M and $\pi x := \sum_{i=0}^{d} x^{q^i}$ be linear. Suppose that f satisfies the symmetric relation:*

$$Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x).$$

If $d < min\{\frac{n}{2} - \theta, |n - 3\theta|, \theta - 1\}$, or if $d < \{\theta - \frac{n}{2}, |2n - 3\theta|, n - \theta - 1\}$, then $M = M_\sigma \circ \pi$ for some $\sigma \in k$.

Proof. Assume $Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$ holds true. Then, we have two cases.

1.) $\theta < \frac{n}{2}$
By [9, Theorem 3], we are done.

2.) $\theta > \frac{n}{2}$
Let $\widetilde{f}(x) = f(x)^{q^{-\theta}} = f\left(x^{q^{-\theta}}\right)$
We have,

$$Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$$

$$[Df(Ma, \pi x) + Df(\pi a, Mx)]^{q^{-\theta}} = [\Lambda_M Df(\pi a, \pi x)]^{q^{-\theta}}$$

$$[Df(Ma, \pi x) + Df(\pi a, Mx)]^{q^{-\theta}} = L_\theta^{-1} \Lambda_M Df(\pi a, \pi x)$$

Let L_θ represent the map that raises terms to the θ^{th} power. We can use the definition of the discrete differential to expand the left hand side of the equation. By linearity, we can distribute the exponent $q^{-\theta}$ to each term. After applying our lemma we get the following,

$$\tilde{f}(Ma+\pi x)+\tilde{f}(Ma)+\tilde{f}(\pi x)+\tilde{f}(\pi a+Mx)+\tilde{f}(\pi a)+\tilde{f}(Mx) = L_\theta^{-1}\Lambda_M Df(\pi a, \pi x)$$

By adding $0 = 2\tilde{f}(0)$ to the left and applying $I = L_\theta L_\theta^{-1}$ to the right we get,

$$D\tilde{f}(Ma, \pi x) + D\tilde{f}(\pi a, Mx) = L_\theta^{-1}\Lambda_M(L_\theta L_\theta^{-1})Df(\pi a, \pi x)$$

And by the lemma we have,

$$D\tilde{f}(Ma, \pi x) + D\tilde{f}(\pi a, Mx) = L_\theta^{-1}\Lambda_M L_\theta D\tilde{f}(\pi a, \pi x)$$

We now have a relation on $\widetilde{f}(x)$ where $-\theta + d < \frac{n}{2}$. Now we can apply [9, Theorem 3] to conclude that $M = M_\sigma \circ \pi$ for some $\sigma \in k$.

We note that the existence of a differential symmetry on $f \circ \pi$ implies a solution of the equation in Theorem 1 as well as the commutativity of M_σ and π. Since the commutativity of M_σ and π requires that π is L-linear, where $\mathbb{F}_q \subseteq L \subseteq k$ and $\sigma \in L$, for any nontrivial differential symmetry to exist,

$(d, n) > 1$. Thus, there is a most desirable value of d from an efficiency and security standpoint: $d = 1$.

Let us specifically consider this most desired value $d = 1$. Then the only restriction on θ for provable differential security is

$$\theta \in \left(2, \frac{n-1}{3}\right) \cup \left(\frac{n+1}{3}, \frac{n}{2} - 1\right) \cup \left(\frac{n}{2} + 1, \frac{2n-1}{3}\right) \cup \left(\frac{2n+1}{3}, n-2\right).$$

Furthermore, since $\theta = \frac{n}{2}$ always produces a many-to-one map in any characteristic, the restriction to provably secure parameters for PFLASH eliminates at most four possible values for θ for all extension degrees n.

5 Extension to PFLASH

We now generalize the analysis of the previous section in application to PFLASH. First we derive a heuristic argument for bootstrapping the provable security of the composition $f \circ \pi$ to statistical security for the projected primitive. We then clarify the resistance of PFLASH to analysis as an HFE$^-$ scheme. Finally, we derive security bounds for various PFLASH parameters.

5.1 Differential Analysis

As mentioned in Section 3, proof that differential symmetries do not exist for the central map of a scheme verifies that a differential adversary cannot recover a full rank key. Such a proof does not, however, verify that a differential adversary cannot find a symmetry revealing the extension field multiplicative structure and directly attack the scheme.

To illustrate this principal, imagine a high degree variant of HFE in which the central map has the form $f(x) = x^{q^\theta+1} + \pi_2(Q(x))$ over an extension of degree $2n$, where π_2 is a rank n projection onto the complement of the subfield of size q^n and Q is an arbitrary quadratic. Then any minus variant in which the image of π_2 is the kernel of T is a C^{*-} public key, but one with multiplicative symmetry. In particular, any map L representing muliplication by an element in the intermediate extension of degree n would satisfy

$$D(T \circ f \circ U)(U^{-1}La, x) + D(T \circ f \circ U)(a, U^{-1}Lx) = (L^{q^\theta} + L)D(T \circ f \circ U)(a, x).$$

Thus the minus scheme has a multiplicative symmetry even though the original scheme provably does not. In fact, even more strongly, we have computed functions of the form of f above over a degree 6 extension of $GF(2)$ for which no linear differential symmetry of any form exists, but under projection onto the degree 3 subfield, the *multiplicative* symmetry is exhibited.

In the case of PFLASH, we may attempt the strategy of the previous section for proving security. We may always model the removal of r equations as the application of a polynomial $\pi(x) = \sum_{i=0}^{r} a_i x^{q^i}$ to the central map. If only a few equations are removed, then the analysis proceeds just like in [19], because $f \circ \pi$

is a low rank albeit high degree polynomial. Since no parameters suggested for PSFLASH are near this range, however, this analysis does not apply. When we perform this analysis with $r \approx \frac{n}{3}$ and $f \circ \pi$, however, the methods of the previous section fail to generate a provably secure class of private keys.

Fortunately, there is an easy heuristic argument revealing a simple relationship between symmetries of the central map and symmetries of a map with the minus modifier that shows that symmetry should be statistically no more likely for any minus modified scheme than for the original. Let T' be the minus projection composed with the inclusion mapping with domain $\mathbb{F}_q^{n-r}$ and codomain $\mathbb{K}$. Suppose that $T' \circ f \circ \pi$ has a differential symmetry. Then

$$D(T' \circ f)(\pi a, Mx) + D(T' \circ f)(Ma, \pi x) = \Lambda_M D(T' \circ f)(\pi a, \pi x)$$
$$T' \left[Df(\pi a, Mx) + Df(Ma, \pi x) \right] = \Lambda_M T' Df(\pi a, \pi x).$$

Since the left is clearly in $T'\mathbb{K}$, the right must be as well. Thus, with high probability, that is, when $Span_{a,x}(Df(\pi a, \pi x)) = \mathbb{K}$, we have that $\Lambda_M T'\mathbb{K} = T'\mathbb{K}$. We know from linear algebra that in this case there exists at least one invertible transformation Λ'_M such that $\Lambda_M T' = T'\Lambda'_M$. Therefore, we obtain the relation

$$Df(\pi a, Mx) + Df(Ma, \pi x) = \Lambda'_M Df(\pi a, \pi x) \pmod{ker(T')}. \tag{1}$$

Clearly, this argument is not reversible for any Λ'_M satisfying (1); therefore, we cannot in general conclude that the scheme with the minus modifier inherits any differential symmetry from the central map. On the other hand, satisfying (1) imposes $n - r$ constraints on Λ_M, while the "commuting" of Λ_M with T' imposes another r constraints. Thus, the existence of a symmetry in the minus case imposes the same number of constraints on Λ_M as for the central map and so we expect the probability of the existence of a differential symmetry to be no higher than for the central map.

5.2 Rank Analysis

One can consider PFLASH to be a high degree version of HFE$^-$ by absorbing the projection of the variables into the central map. Notice that the rank of the composition is still only two, thus PFLASH must achieve its security from the minus modifier.

Recently, in [22], a key recovery attack valid for all parameters of HFE$^-$ is presented. For an HFE$^-$ instance with parameters (q, n, D, r), the complexity is noted as $\mathcal{O}(\binom{n+\lceil log_q(D)\rceil+1}{\lceil log_q(D)\rceil+r+1}^{\omega})$.

In application to PFLASH, there are two things to note about this attack. First, the attack produces an equivalent HFE$^-$ key, not a pC^{*-} key. This fact may not limit the attack, because it will still recover a central map of rank two of the form $f \circ \pi$ which we may then attack as a pC^* scheme in the manner of [23]. Second, the quantity $\lceil log_q(D) \rceil$ in the complexity estimate is derived from the rank structure that the degree bound of HFE implies, not directly from the

degree bound itself. Thus, the rank of the C^* monomial, which is two, plays the role of $\lceil log_q(D) \rceil$ in the application of the techniques of [22] to PFLASH.

In fact, instances of PFLASH with quite inappropriate but still large parameters can be broken with this method. In particular we note that for a PFLASH$(256, 44, 3, 1)$ that the complexity of the attack is roughly estimated $44^{(3+2+1)\omega} \sim 2^{78}$. For large values of r, however, such as in all parameter sets in [10], this attack is infeasible. For example, the smallest parameters suggested in [10] still resist this attack to dozens of orders of magnitude beyond brute force. Thus, for sensible parameters with r sufficiently large, PFLASH is secure.

5.3 Security Estimates

Now with a refined security analysis, we can eliminate differential attacks for a larger set of parameters, thus doubling the entropy of the key space for PFLASH. In addition, with the complexity estimate of $\mathcal{O}(n^{(r+3)\omega})$ and practical values of r, PFLASH is quite secure against the new attack on HFE$^-$ schemes. In conjunction with the invariant analysis of [10], we conclude that the security of PFLASH is determined by its resistance to algebraic and brute force attacks.

Viewing PFLASH as an HFE$^-$ scheme, we may use the bound in [24] to estimate the degree of regularity of PFLASH. This upper bound can be computed

$$\frac{(q-1)(R+r)}{2} + 2,$$

where R is the rank of the central map; in the case of PFLASH, this quantity is two. Though this is an upper bound, empirical evidence suggests that it is tight for random systems of rank R. Thus the degree of regularity is far too high for practical schemes to be weakened. Furthermore, direct algebraic attacks for large schemes are impractical even with smaller complexity bounds because the space complexity of the best algorithms are too large to be practical.

Therefore, we corroborate the claims of [10] that brute force collision attacks are the greatest threat to PFLASH schemes. The evidence from our increase of the entropy of the key space and the verification that PFLASH resists recent weaknesses revealed in HFE$^-$ suggest the security levels in Table 1 (all of which are in agreement with [10]).

Scheme	Public Key (Bytes)	Security (bits)
PFLASH$(16, 62, 22, 1)$	39,040	80
PFLASH$(16, 74, 22, 1)$	72,124	104
PFLASH$(16, 94, 30, 1)$	142,848	128

Table 1. Security levels for standard parameters of PFLASH

6 Conclusion

The history of PFLASH intersects with most of the major advances in design and cryptanalysis in asymmetric multivariate cryptography. Interestingly, essentially all of the major cryptanalytic techniques that have proven successful in attacking multivariate schemes are relevant for PFLASH, and so any security metric for the scheme must inherently be complex. In spite of all of the tools available to an adversary, PFLASH remains secure.

Our analysis expands upon and complements previous analysis of PFLASH. We verify that the entropy of the key space is not significantly reduced by selecting parameters for which differential security is provable. We further verify security against new developments in rank analysis relevant to schemes employing the minus modifier. We conclude that any attack that fundamentally reduces the security of PFLASH below the brute force bound must include techniques as of yet undeveloped.

In venues for which speed, digest size, storage and power are severe limitations PFLASH seems to be one of the most performant options. When one considers devices in which no public key needs to be transported, such as some applications of smart cards, PFLASH is a leading candidate. In light of the security assurance this analysis provides, PFLASH appears ready for deployment.

References

1. Cryptographic Technology Group: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. NIST CSRC (2016) http://csrc.nist.gov/groups/ST/post-quantum-crypto/documents/call-for-proposals-final-dec-2016.pdf.
2. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Sci. Stat. Comp. **26, 1484** (1997)
3. Matsumoto, T., Imai, H.: Public Quadratic Polynominal-Tuples for Efficient Signature-Verification and Message-Encryption. In: EUROCRYPT. (1988) 419–453
4. Patarin, J.: Cryptoanalysis of the Matsumoto and Imai Public Key Scheme of Eurocrypt'88. In Coppersmith, D., ed.: CRYPTO. Volume 963 of Lecture Notes in Computer Science., Springer (1995) 248–261
5. Patarin, J.: Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): Two New Families of Asymmetric Algorithms. In: EUROCRYPT. (1996) 33–48
6. Patarin, J., Courtois, N., Goubin, L.: Flash, a fast multivariate signature algorithm. CT-RSA 2001, LNCS **2020** (2001) 297–307
7. Dubois, V., Fouque, P.A., Shamir, A., Stern, J.: Practical Cryptanalysis of SFLASH. In Menezes, A., ed.: CRYPTO. Volume 4622 of Lecture Notes in Computer Science., Springer (2007) 1–12
8. Ding, J., Dubois, V., Yang, B.Y., Chen, C.H.O., Cheng, C.M.: Could SFLASH be Repaired? In Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I., eds.: ICALP (2). Volume 5126 of Lecture Notes in Computer Science., Springer (2008) 691–701

9. Smith-Tone, D.: On the differential security of multivariate public key cryptosystems. In Yang, B.Y., ed.: PQCrypto. Volume 7071 of Lecture Notes in Computer Science., Springer (2011) 130–142
10. Chen, M.S., Yang, B.Y., Smith-Tone, D.: Pflash - secure asymmetric signatures on smart cards. Lightweight Cryptography Workshop 2015 (2015) http://csrc.nist.gov/groups/ST/lwc-workshop2015/papers/session3-smith-tone-paper.pdf.
11. Moody, D., Perlner, R.A., Smith-Tone, D.: An asymptotically optimal structural attack on the ABC multivariate encryption scheme. [25] 180–196
12. Moody, D., Perlner, R.A., Smith-Tone, D. In: Key Recovery Attack on the Cubic ABC Simple Matrix Multivariate Encryption Scheme. Springer (2017)
13. Kipnis, A., Shamir, A.: Cryptanalysis of the hfe public key cryptosystem by relinearization. Advances in Cryptology - CRYPTO 1999, Springer **1666** (1999) 788
14. Bettale, L., Faugère, J., Perret, L.: Cryptanalysis of hfe, multi-hfe and variants for odd and even characteristic. Des. Codes Cryptography **69** (2013) 1–52
15. Patarin, J., Goubin, L., Courtois, N.: C^*_{-+} and HM: variations around two schemes of t. matsumoto and h. imai. In Ohta, K., Pei, D., eds.: Advances in Cryptology - ASIACRYPT '98, International Conference on the Theory and Applications of Cryptology and Information Security, Beijing, China, October 18-22, 1998, Proceedings. Volume 1514 of Lecture Notes in Computer Science., Springer (1998) 35–49
16. Patarin, J., Goubin, L., Courtois, N.: Improved algorithms for isomorphisms of polynomials. In Nyberg, K., ed.: 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., Springer (1998) 184–200
17. Berlekamp, E.R.: Factoring polynomials over large finite fields. Mathematics of Computation **24** (1970) pp. 713–735
18. Ding, J., Hu, L., Nie, X., Li, J., Wagner, J. In: High Order Linearization Equation (HOLE) Attack on Multivariate Public Key Cryptosystems. Springer Berlin Heidelberg, Berlin, Heidelberg (2007) 233–248
19. Daniels, T., Smith-Tone, D.: Differential properties of the HFE cryptosystem. [25] 59–75
20. Cartor, R., Gipson, R., Smith-Tone, D., Vates, J.: On the differential security of the hfev- signature primitive. In Takagi, T., ed.: Post-Quantum Cryptography - 7th International Workshop, PQCrypto 2016, Fukuoka, Japan, February 24-26, 2016, Proceedings. Volume 9606 of Lecture Notes in Computer Science., Springer (2016) 162–181
21. Smith-Tone, D.: Properties of the discrete differential with cryptographic applications. In Sendrier, N., ed.: PQCrypto. Volume 6061 of Lecture Notes in Computer Science., Springer (2010) 1–12
22. Vates, J., Smith-Tone, D.: Key recovery attack for all parameters of hfe-. In Current Submission (2017)
23. Billet, O., Macario-Rat, G.: Cryptanalysis of the square cryptosystems. ASIACRYPT 2009, LNCS **5912** (2009) 451–486
24. Ding, J., Kleinjung, T.: Degree of regularity for HFE-. IACR Cryptology ePrint Archive **2011** (2011) 570
25. Mosca, M., ed.: Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014, Waterloo, ON, Canada, October 1-3, 2014. Proceedings. Volume 8772 of Lecture Notes in Computer Science., Springer (2014)