On the Differential Security of Multivariate Public Key Cryptosystems

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

daniel.smith@nist.gov

Abstract. Since the discovery of an algorithm for factoring and computing discrete logarithms in polynomial time on a quantum computer, the cryptographic community has been searching for an alternative for security in the approaching post-quantum world. One excellent candidate is multivariate public key cryptography. Though the speed and parameterizable nature of such schemes is desirable, a standard metric for determining the security of a multivariate cryptosystem has been lacking. We present a reasonable measure for security against the common differential attacks and derive this measurement for several modern multivariate public key cryptosystems.

Key words: Matsumoto-Imai, multivariate public key cryptography, differential, symmetry

1 Introduction

In recent years a great deal of focus has been directed towards post-quantum cryptology. This increased attention is indicative of a paradigm shift which has been occurring since, in [1], Peter Shor developed algorithms for factoring and computing discrete logarithms in polynomial time on a quantum computing device. In the face of mounting evidence that quantum computing is not a physical impossibility but merely an engineering challenge, it is more important than ever that we develop secure systems relying on problems of greater difficulty than the classical number theoretic schemes.

Multivariate Public Key Cryptography(MPKC) has emerged as one of a few serious candidates for security in the post-quantum world. This emergence is due to several facts. First, the problem of solving a system of quadratic equations is known to be NP-hard, and seems to be hard even in the average case. No great reduction of the complexity of this problem has been found in the quantum model of computing, and, indeed, if this problem is discovered to be solvable in the quantum model, we can solve all NP problems, which seems particularly wishful. Second, multivariate systems are very efficient, often having

speeds dozens of times faster than RSA, [2–4]. Finally, it is easy to parameterize many multivariate systems in such a way that vastly different schemes are derived with potentially vastly different resistances to specialized attacks.

One of the great challenges facing MPKC is the task of deriving security proofs. In fact, there currently is no widely accepted quantification for indistinguishability between systems of multivariate equations. One reason for the absence of such a quantification is the fact that even with a great deal of structure in the construction of a multivariate cryptosystem, the coefficients can appear to have a uniform distribution. In fact, history has shown that once a way to distinguish a system of structured multivariate equations from a collection of random equations is discovered, a method of solving this system is often quickly developed.

Recently, several cryptanalyses of various multivariate cryptosystems have pointed out weaknesses in the predominant philosophy for the construction of multivariate public key cryptosystems. Several systems, SFLASH, Square, for example, which are based on simple modifications of the prototypical Matsumoto-Imai public key cryptosystem, have been broken by very similar differential attacks exploiting some symmetry which is inherent to the field structure these systems utilize. See [5–8]. In fact, even various attacks on other multivariate schemes, for example the oil-vinegar attack, see [9], can be viewed as a dual attack, finding a differential invariant.

In [10], a classification of field maps exhibiting the multiplicative symmetry was presented. In this article we are interested in the dual problem, that is, identifying all possible initial general linear differential symmetries a field map can possess. Such a characterization will lead to a fuller understanding of the theory, potentially establish a foundation for modeling more general security proofs, and establish a reasonable and quantitative criterion for the development of future multivariate schemes.which we may model

The paper is organized as follows. The next section illustrates the ubiquitous nature of the differential attack by recasting the attack on the balanced oil and vinegar scheme in the differential setting. In the following section, we focus on differential symmetry, presenting the general linear symmetry and discussing the general structure of the space of linear maps exhibiting this symmetry. The subsequent section restricts the analysis of this space to the case in which the hidden field map of the cryptosystem is a C^* monomial. Next the specific case of the squaring map used in Square is analyzed. The space of linear maps is then determined for projected systems such as the projected SFLASH analogue, pSFLASH. Finally, we review these results and analyze the dimension of this space of linear maps as a metric for determining differential security.

2 Differential Symmetries and Invariants

Differential attacks play a crucial role in multivariate public key cryptography. Such attacks have not only broken many of the so called “big field” schemes, they have directed the further development of the field by inspiring modifiers —

Plus (+), Minus (-), Projection (p), Perturbation (P), Vinegar (v) — and the creation of newer more robust techniques.

The differential of a field map, f, is defined by $Df(a,x) = f(a+x) - f(a) - f(x) + f(0)$. The use of this discrete differential appears to occur in very many cryptanalyses of post-quantum multivariate schemes. In fact, we can even consider Patarin's initial attack, in [11], on Imai and Matsumoto's C^* scheme, see [12], as the exploitation of a trivial differential symmetry. Suppose $f(x) = x^{q^\theta+1}$ and let $y = f(x)$. Since the differential of f, Df, is a symmetric bilinear function, $0 = Df(y,y) = Df(y, x^{q^\theta+1}) = yx^{q^{2\theta}+q^\theta} + y^{q^\theta}x^{q^\theta+1} = x^{q^\theta}(yx^{q^{2\theta}} + y^{q^\theta}x)$. Dividing by x^{q^θ} we have Patarin's linear relation, $yx^{q^{2\theta}} = y^{q^\theta}x$; see [11] for details.

Differential methods provide powerful tools for decomposing a multivariate scheme. To illustrate the nearly universal nature of differential attacks, we review the attack of Kipnis and Shamir, see [9], on a non-big-field system, the oil and vinegar scheme. Though they use differing terminology, the attack exploits a symmetry hidden in the differential structure of the scheme.

Recall that the oil and vinegar scheme is based on a hidden quadratic system of equations, $f : k^n \to k^o$, in two types of variables, $x_1, ..., x_o$, the oil variables, and $x_{o+1}, ..., x_{o+v=n}$, the vinegar variables. We focus on the balanced oil and vinegar scheme, in which $o = v$. Let $c_1, ..., c_v$ be random constants. The map f has the property that $f(x_1, ..., x_v, c_1, ..., c_v)$ is affine in $x_1, ..., x_v$. The encryption map, $\overline{f}$ is the composition of f with an n-dimensional invertible affine map, L.

Let O represent the subspace generated by the first v basis vectors, and let V denote the cosummand of O. Notice that the discrete differential given by $Df(a,x) = f(x+a) - f(x) - f(a) + f(0)$ has the property that for all a and x in O, $Df(a,x) = 0$. Thus for each coordinate, i, the differential coordinate form Df_i can be represented:

$$Df_i = \begin{bmatrix} 0 & Df_{i1} \\ Df_{i1}^T & Df_{i2} \end{bmatrix} \Bigg($$

Let M_1 and M_2 be two invertible matrices in the span of the Df_i. Then $M_1^{-1}M_2$ is an O-invariant transformation of the form:

$$\begin{bmatrix} A & B \\ 0 & C \end{bmatrix} \Bigg($$

Now the Df_i are not known, but $D(f \circ L)_i = L^T Df_i L$, so the $L^T Df_i L$ are known. Notice that if M is in the span of the Df_i, then $L^T ML$ is in the span of the $L^T Df_i L$. Also, since $(L^T M_1 L)^{-1}(L^T M_2 L) = L^{-1}M_1^{-1}M_2 L$, there is a large space of matrices leaving $L^{-1}O$ invariant, which Kipnis and Shamir are able to exploit to effect an attack against the balanced oil and vinegar scheme; see [9] for details. Making the oil and vinegar scheme unbalanced, see [13], corrects this problem by making any subspace which is invariant under a general product $M_1^{-1}M_2$ very small, see [14].

While the differential analysis of the oil and vinegar systems is a very specific case of utilizing an invariant related to the differential structure of the hidden

map, several general attacks on big field schemes rely on a type of linear symmetry. The following sections focus on a systematic study of this type of symmetry, and conditions in which such a symmetry can be utilized for a differential attack.

3 Properties of General Linear Symmetries

Let k be an extension field of $\mathbb{F}_q$, the field with q elements. Dubois et al. completed a successful attack against the SFLASH signature scheme, see [8], by utilizing a multiplicative symmetry of the form:

$$Df(\sigma a, x) + Df(a, \sigma x) = (\sigma^{q^\theta} + \sigma)Df(a, x), \tag{1}$$

where $f : k \to k$ is a C^* monomial map, and $\sigma \in k$.

Consider the more general initial linear symmetric relation as suggested by Dubois et al., in [8], of the form:

$$Df(La, x) + Df(a, Lx) = \Lambda_L Df(a, x), \tag{2}$$

where $f : k \to k$ is a function, and $L, \Lambda_L : k \to k$ are $\mathbb{F}_q$-linear. This definition is perfectly appropriate, since we are guaranteed a solution space of dimension at least n for C^* monomial maps, f. In addition, it is clear that we have additive closure, in general. Let S_G denote the set of all linear maps, L, satisfying (2). Notice:

$$\begin{aligned} Df((L+M)a, x) + Df(a, (L+M)x) &= Df(La, x) + Df(a, Lx) \\ &\quad + Df(Ma, x) + Df(a, Mx) \\ &= \Lambda_L Df(a, x) + \Lambda_M Df(a, x) \\ &= (\Lambda_L + \Lambda_M)\, Df(a, x). \end{aligned} \tag{3}$$

For a more general function, f, however, we have no guarantee of such a large space of solutions as possessed by C^* monomials; however, in characteristic two, the discovery of one such symmetric relation allows the generation of a space of maps satisfying the symmetry which has both an additive and square structure. It is worth exploring to see how much structure such a symmetry holds.

Note that if L is in S_G:

$$\begin{aligned} Df(L^2a, x) + Df(a, L^2x) &= Df(L^2a, x) + Df(La, Lx) \\ &\quad + Df(La, Lx) + Df(a, L^2x) \\ &= \Lambda_L Df(La, x) + \Lambda_L Df(a, Lx) \\ &= \Lambda_L \left(Df(La, x) + Df(a, Lx)\right) \\ &= \Lambda_L^2 Df(a, x). \end{aligned} \tag{4}$$

Notice that for odd characteristic, there is no way to add the needed terms of the form $Df(La, Lx)$. We do not have, in general, multiplicative closure, but for any polynomial function, p, with terms of degree zero or a power of two, if

$L \in S_G$ then $p(L) \in S_G$. Thus, the existence of a single linear map L satisfying the initial general linear symmetry guarantees the existence of a relatively large space of maps satisfying the symmetry.

Therefore S_G is the $\mathbb{F}_q$-vector space sum of rings of the form $\mathbb{F}_q$ or $\mathbb{F}_q\left[L^{2^i}\right]$. Given just a few elements of S_G, we can potentially generate a large subspace of S_G, which is a very appealing situation for an adversary.

This situation is exactly the scenario which has resulted in the breaking of SFLASH and other C^* variants. In [8], it was shown that $k < S_G$ when f is a C^* monomial, and thus S_G is so large that an element can be detected using the relation (2) even when up to one half of the public equations are removed.

Thus the task of constructing a differentially secure multivariate cryptosystem must necessarily include an analysis of the space of linear maps, S_G, illustrating the symmetry. If S_G is very small, then recovering an element from this subspace may be an infeasible task, and the differential attack is doomed.

4 Properties Relative to C^* Monomials

If we restrict our attention to the case in which f is a C^* monomial map of the form $f(x) = x^{q^\theta+1}$, we can derive some additional properties of S_G indicating why so many C^* variants have fallen to differential attacks. Immediately, we know that there is an injective map $g : k \to S_G$, since f has the multiplicative symmetry. Furthermore, by considering the linearized polynomial form of an arbitrary linear map, $L \in GL(\mathbb{F}_q, n)$, we can continue, revealing the exact multiplicative structure of S_G.

Theorem 1 *If f is a C^* monomial, then S_G, equipped with standard multiplication is a k-algebra, and consequently has a large dimension as an $\mathbb{F}_q$-vector space. Furthermore, if $3\theta \neq n$, $S_G \cong k$.*

Proof. Consider the linearized polynomial form of $M \in S_G$, $Mx = \sum_{i=0}^{n-1} m_i x^{q^i}$. We will find conditions on the coefficients, m_i of this linearized polynomial form. For the generic C^* monomial map, $f(x) = x^{q^\theta+1}$, we have that the discrete differential, $Df(a, x) = a^{q^\theta}x + ax^{q^\theta}$. Thus:

$$\begin{aligned} Df(Ma, x) + Df(a, Mx) &= \sum_{i=0}^{n-1}\left(m_i^{q^\theta} a^{q^{i+\theta}} x + m_i a^{q^i} x^{q^\theta}\right)\Bigg(\\ &\quad + \sum_{i=0}^{n-1}\Bigg(\left(m_i a^{q^\theta} x^{q^i} + m_i^{q^\theta} a x^{q^{i+\theta}}\right)\Bigg(\\ &= \sum_{i=0}^{n-1}\left(m_i^{q^\theta}\left(a^{q^{i+\theta}} x + a x^{q^{i+\theta}}\right)\right)\Bigg(\\ &\quad + \sum_{i=0}^{n-1}\left(m_i\left(a^{q^i} x^{q^\theta} + a^{q^\theta} x^{q^i}\right)\right)\Bigg(\end{aligned} \tag{5}$$

Since $M \in S_G$, there is an $\mathbb{F}_q$-linear map, $\Lambda_M(x) = \sum_{i=0}^{n-1} \lambda_i x^{q^i}$ such that the equation $Df(Ma, x) + Df(a, Mx) = \Lambda_M Df(a, x)$ holds. Therefore we have:

$$\begin{aligned}\sum_{i=0}^{n-1} \Bigg(m_i^{q^\theta} \left(a^{q^{i+\theta}} x + a x^{q^{i+\theta}} \right) + m_i \left(a^{q^i} x^{q^\theta} + a^{q^\theta} x^{q^i} \right) &= \sum_{i=0}^{n-1} \Bigg(\lambda_i \left(a^{q^\theta} x + a x^{q^\theta} \right)^{q^i} \\ &= \sum_{i=0}^{n-1} \Bigg(\lambda_i \left(a^{q^{i+\theta}} x^{q^i} + a^{q^i} x^{q^{i+\theta}} \right) \Bigg(\end{aligned} \tag{6}$$

We can collect the coefficients of each monomial, $a^i x^j$, and set each to zero, obtaining relations on the coefficients of the linearized form of M and Λ_M.

If $q^\theta + 1$ shares a nontrivial factor with $q^n - 1$, then f is not strictly speaking a C^* monomial, since it is not a permutation polynomial. Thus we treat the case $\theta \notin \{0, \frac{n}{2}, \frac{n}{4}\}$, encompassing all C^* monomials, as well as many functions which are not C^* monomials. If we collect the coefficients of the monomial ax^{q^θ}, we get the relation $\lambda_0 = m_0 + m_0^{q^\theta}$. The coefficients of monomials of the form ax^{q^i}, for $i \notin \{0, \pm\theta\}$, generate the relations $m_{i-\theta} = 0$. Thus $m_i = 0$ for all $i \notin \{0, -\theta, -2\theta\}$. Collecting the coefficients of the monomials of the form $a^{q^\theta} x^{q^i}$ for $i \notin \{0, \theta, 2\theta\}$, we have $m_i = 0$.

Therefore, if a nonzero coefficient exists other than m_0, then either $-\theta = \theta$, which implies $\theta = \frac{n}{2}$, $-\theta = 2\theta$, implying $3\theta = n$, or $-2\theta = 2\theta$, which implies $\theta = \frac{n}{4}$. Of these cases, only $3\theta = n$ represents a possible C^* monomial. Thus, if $3\theta \neq n$, then for all $i \neq 0$, $m_i = 0$, and in this case, $Mx = m_0 x$ is multiplication by an element in k; consequently, $S_G \cong k$.

If $3\theta = n$, then m_0, $m_{\frac{n}{3}}$, and $m_{\frac{2n}{3}}$ can possibly be nonzero. To prove that S_G is still a ring in this case, notice that given two linear maps, M and L, each with all coefficients zero except possibly m_0, $m_{\frac{n}{3}}$, $m_{\frac{2n}{3}}$, l_0, $l_{\frac{n}{3}}$, and $l_{\frac{2n}{3}}$, we have:

$$\begin{aligned} LMx = &(l_0 m_0 + l_{\frac{n}{3}} m_{\frac{2n}{3}}^{q^{\frac{n}{3}}} + l_{\frac{2n}{3}} m_{\frac{n}{3}}^{q^{\frac{2n}{3}}}) x \\ &+ (l_0 m_{\frac{n}{3}} + l_{\frac{n}{3}} m_0^{q^{\frac{n}{3}}} + l_{\frac{2n}{3}} m_{\frac{2n}{3}}^{q^{\frac{2n}{3}}}) x^{q^{\frac{n}{3}}} \\ &+ (l_0 m_{\frac{2n}{3}} + l_{\frac{n}{3}} m_{\frac{n}{3}}^{q^{\frac{n}{3}}} + l_{\frac{2n}{3}} m_0^{q^{\frac{2n}{3}}}) x^{q^{\frac{2n}{3}}}, \end{aligned} \tag{7}$$

which is, again a linear map with all coefficients zero except for the 0-th, $\frac{n}{3}$-th, and $\frac{2n}{3}$-th. Thus S_G has multiplicative closure, and is a 3-dimensional k-algebra.

In the above theorem we didn't mention anything about characteristic. Strictly speaking, a C^* monomial is linearly equivalent to a quadratic permutation polynomial of the form $f(x) = x^{q^\theta + 1}$. This is only possible, however, when q is even, since trivially, $2 | (q^\theta + 1, q^n - 1)$. Some cryptosystems, however, do use this form of core map in odd characteristic, choosing a map which is 2-to-1, or few-to-1. Such systems never use $\theta \in \{\frac{n}{2}, \frac{n}{4}\}$, since such maps would have exponential collisions. It is for this reason that in the above theorem we relaxed the constraints

and allowed any map with $\theta \notin \{0, \frac{n}{2}, \frac{n}{4}\}$. We have completely characterized the symmetries in these cases.

5 Symmetries for Non-permutation Polynomials

In [5, 6], two notable systems, Square and Square-Vinegar, introduced the idea of utilizing a quadratic map over a field of odd characteristic. The C^* form of the core map of Square is $f(x) = x^{q^\theta+1}$ where $\theta = 0$. The theorem of the preceding section doesn't apply to the case $\theta = 0$, therefore we will treat this case separately, and completely characterize S_S, the space of linear maps, L, satisfying (2).

Theorem 2 *Let q be odd. Then $S_S \cong k$.*

Proof. First, $Df(a, x) = 2ax$. Therefore, by the symmetric application of the linear function $M(x) = \sum_{i=0}^{n-1} m_i x^{q^i}$, we have:

$$Df(Ma, x) + Df(a, Mx) = 2 \left(\sum_{i=0}^{n-1} m_i a^{q^i} \right) x + 2a \left(\sum_{i=0}^{n-1} m_i x^{q^i} \right) \left(\right. \quad (8)$$

Setting this quantity equal to $\Lambda_M Df(a, x)$ we have:

$$2 \left(\sum_{i=0}^{n-1} \left(m_i a^{q^i} \right) x + 2a \left(\sum_{i=0}^{n-1} m_i x^{q^i} \right) \right) \left(= \sum_{i=0}^{n-1} \lambda_i 2^{q^i} a^{q^i} x^{q^i}. \right. \quad (9)$$

We can collect the coefficients of each monomial $a^{q^i} x^{q^j}$ and set each equal to zero to determine relations between M and Λ_M. Collecting coefficients for monomials of the form ax^{q^i}, for $i \neq 0$, we get the relations, $2m_i = 0$. Thus $m_i = 0$ for all $i \neq 0$, and M is multiplication by m_0 in k; consequently, $S_S \approx k$.

It is important to note that the Square systems have been broken by a differential attack in [7] which recovers the multiplicative structure of k by utilizing a symmetry Square exhibits under left composition. This method of finding a terminal symmetry under left composition was discovered for two reasons: first, the Square systems did not preclude such an attack by employing the minus modifier or an alternative precaution; and second, the designers were able to mask the initial multiplicative symmetry of the core map of Square by projecting the input of the C^* monomial into a subspace, making an attack using a symmetry of the form (2) infeasible. If we include the minus modifier, i.e. consider Square-, then the attack of [7] fails, and the question of which symmetries exist over a subspace becomes more critical.

6 Symmetries over Subspaces

In [15], Ding et al. began the work of classifying the initial general linear symmetries for C^* monomial maps over subspaces. Their result was imprecisely stated,

but they successfully proved that "almost always" if a field map has an initial general linear symmetry over a subspace then that symmetry is a multiplicative symmetry.

As stated, the claim indicated that for a bijective C^* monomial, f, given any hyperplane, $H = \pi(k)$, if we have:

$$Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x), \tag{10}$$

for all $a, x \in k$, then $M = M_\sigma \circ \pi$ and $\Lambda_M = M_{\sigma + \sigma^{q^\theta}}$, for some $\sigma \in k$.

To prove that the statement as given in [15] is in err, let us define the space saving notation $S_f(A, B)(a, x) = Df(Aa, Bx) + Df(Ba, Ax)$, and take the following example. Let $k = GF(64)$ over $\mathbb{F}_2$, $f(x) = x^5$, $\pi x = x + x^2$, and $Mx = x^4 + x^8$. By a simple calculation,

$$\begin{aligned} S_f(M, \pi)(a, x) &= (a^{16} + a^{32})(x + x^2) + (a + a^2)(x^{16} + x^{32}) \\ &= a^{16}x + ax^{16} + a^{32}x + ax^{32} + a^{16}x^2 + a^2x^{16} + a^{32}x^2 + a^2x^{32} \\ &= \left(ax^4 + a^4x + a^2x^4 + a^4x^2 + ax^8 + a^8x + a^2x^8 + a^8x^2\right)^{16} \\ &= Df(\pi a, \pi x)^{16}. \end{aligned} \tag{11}$$

(Here we note that two terms of the form $(a^4 + a^8)(x^4 + x^8)$ cancelled each other in the first line above.) Thus, we have found a counterexample with $\Lambda_M x = x^{16}$ and $Mx = (\pi x)^4$, which is certainly not the composition of a multiplication map and π. Here the fact that $2\,(codim(H) + \theta) = n$ created some extra symmetries in the relations between the coefficients of M and Λ_M. Informally, θ was an exceptional choice which permits the existence of a linear map allowing collisions between monomials generated from $Df(Ma, \pi x)$ and $Df(\pi a, Mx)$. Since the arithmetic of k has characteristic 2, collision corresponds with annihilation.

We can resolve the minor issues with the result of Ding et al. and generalize the statement somewhat by providing a more detailed analysis of the symmetry:

$$Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x), \tag{12}$$

for more general linear maps, π. In particular, a more precise formulation of the result of Ding et al. is the special case of $d = 1$ in the following theorem.

Theorem 3 *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 $Df(Ma, \pi x) + Df(\pi a, Mx) = \Lambda_M Df(\pi a, \pi x)$. If $\theta + d < \frac{n}{2}$, $|n - 3\theta| > d$, and $0 < d < \theta - 1$, then $M = M_\sigma \pi$ for some $\sigma \in k$.*

Proof. Our strategy for the proof will be to determine relations between the coefficients of the linearized polynomial forms of M and Λ_M. We will zigzag back and forth between solving for coefficients of M and of Λ_M, further resolving the relationship between the two maps with each step. We will extensively use the "space of indices," the torus consisting of the pairs (r, s) (mod n) which correspond to monomials of the form $a^{q^r} x^{q^s}$. The geometry of this space of indices will be useful in determining relations on the coefficients of the corresponding monomials in the expansions of (12).

Expanding the right hand side of (12) repeatedly, using the bilinearity of Df, we obtain:

$$\begin{aligned}
\Lambda_M Df(\pi a, \pi x) &= \sum_{i=0}^{n-1} \Big(\lambda_i Df(\pi a, \pi x)^{q^i} \\
&= \sum_{i=0}^{n-1} \Big(\lambda_i Df(\sum_{j=0}^{d} a^{q^j}, \sum_{l=0}^{d} x^{q^l})^{q^i} \\
&= \sum_{i=0}^{n-1} \sum_{j=0}^{d} \sum_{l=0}^{d} \Big(\lambda_i Df(a^{q^j}, x^{q^l})^{q^i} \\
&= \sum_{i=0}^{n-1} \sum_{j=0}^{d} \sum_{l=0}^{d} \Big(\lambda_i \left(a^{q^{\theta+j}} x^{q^l} + a^{q^j} x^{q^{\theta+l}} \right)^{q^i} .
\end{aligned} \tag{13}$$

Notice that for each monomial term in this expression, the difference between the exponent of q in the power of a and the exponent of q in the power of x is $l - \theta - j \pmod n$ or $l + \theta - j \pmod n$. Also, there is the restriction that $0 \leq l, j \leq d$. From these facts we can determine which monomials never occur in the right side of (12).

The monomial $a^{q^r} x^{q^s}$ may only occur in the right side of (12) if the difference between the coordinates, $(s - r) \ (mod \ n) \in [-\theta - d, -\theta + d] \cup [\theta - d, \theta + d]$, where we require $2\theta + 2d < n$, avoiding overlap. Also, implicitly, we have the restriction that for such an interval, (u, v), the positive residues u and v satisfy $0 \leq v - u \leq n - 1$. For all other pairs, (r, s), $a^{q^r} x^{q^s}$ certainly has a coefficient of zero in the right side of (12). Therefore, we will study the set of pairs of indices,

$$E = \{(r, s) | s - r \in (-\theta + d, \theta - d) \cup (\theta + d, -\theta - d)\}.$$

This set is the diagonal band in the space of indices for which the corresponding coefficients have no contribution from the right side of (12); refer to the shaded region in the figure below.

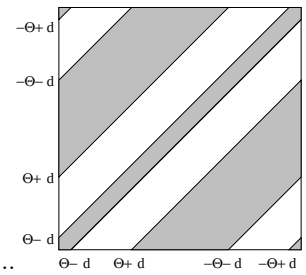

Fig. 1. The space of indices with the shaded region corresponding to monomials which cannot occur on the right side of (12).

Expanding the left hand side of (12), similarly:

$$\begin{aligned}
S_f(M,\pi)(a,x) &= Df(Ma,\pi x) + Df(\pi a, Mx) \\
&= Df(\sum_{i=0}^{n-1} m_i a^{q^i}, \pi x) + Df(\pi a, \sum_{i=0}^{n-1} m_i x^{q^i}) \\
&= \sum_{i=0}^{n-1} \Big(Df(m_i a^{q^i}, \pi x) + Df(\pi a, m_i x^{q^i}) \Big) \Big(\\
&= \sum_{i=0}^{n-1} \Big(Df(m_i a^{q^i}, \sum_{j=0}^{d} x^{q^j}) + Df(\sum_{j=0}^{d} a^{q^j}, m_i x^{q^i}) \Big) \Big(\\
&= \sum_{i=0}^{n-1} \sum_{j=0}^{d} \Big(Df(m_i a^{q^i}, x^{q^j}) + Df(a^{q^j}, m_i x^{q^i}) \Big) \Big(\\
&= \sum_{i=0}^{n-1} \sum_{j=0}^{d} \Big(m_{i-\theta}^{q^\theta} a^{q^i} x^{q^j} + m_i a^{q^i} x^{q^{\theta+j}} + m_i a^{q^{\theta+j}} x^{q^i} + m_{i-\theta}^{q^\theta} a^{q^j} x^{q^i} \Big) \Big(\\
&\qquad (14)
\end{aligned}$$

Now, to analyze which monomials of the form $a^{q^r} x^{q^s}$, have nontrivial coefficients for the pair of "indices" (r,s), we construct four index sets, A, B, C, and D, relative to the four monomials in the above expression, respectively. We have:

$$\begin{aligned}
A &= [0, n-1] \times [0, d] \\
B &= [0, n-1] \times [\theta, \theta + d] \\
C &= [\theta, \theta + d] \times [0, n-1] \\
D &= [0, d] \times [0, n-1] \,.
\end{aligned} \tag{15}$$

We can see that only the pairs (A, C), (A, D), (B, C), and (B, D) have non trivial intersections. Isolating the index pairs occurring in only one of these index spaces we can find relations on the coefficients of M and Λ_M which involve only one m_i. If, furthermore, the index pair occurs in E, then the corresponding coefficient from Λ_M is zero. Let $*$ denote the operation of taking one of these sets minus the union of the other three. We notice that:

$$\begin{aligned}
A^* &= ([d+1, \theta-1] \cup [\theta+d+1, n-1]) \times [0, d] \\
B^* &= ([d+1, \theta-1] \cup [\theta+d+1, n-1]) \times [\theta, \theta+d]
\end{aligned} \tag{16}$$

if $d < \theta$.

For both A^* and B^*, the first coordinate is the index associated with the coefficient of M in (14); we are, therefore, interested in which values of the first coordinate are possible in $A^* \cap E$ and $B^* \cap E$. Equivalently, we want to discover $\pi_1(A^* \cap E)$ and $\pi_1(B^* \cap E)$, where π_1 is the projection mapping onto the first

coordinate. By a simple calculation, we have:

$$\begin{aligned}\pi_1(A^* \cap E) &= [\theta + d + 1, -\theta - 1] \cup [-\theta + d + 1, n - 1] \cup [d + 1, \theta - 1] \\ \pi_1(B^* \cap E) &= [d + 1, \theta - 1] \cup [\theta + d + 1, 2\theta - 1] \cup [2\theta + d + 1, n - 1]\,,\end{aligned} \tag{17}$$

see the figure below.

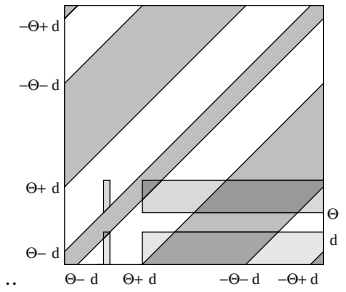

Fig. 2. The intersection of A^* and B^* with E.

Since the coefficient of M associated with (r, s) in A^* is $m_{r-\theta}$, and the coefficient associated with (r, s) in B^* is m_r, we know that $m_r = 0$ for every r in the union, $((\pi_1(A^* \cap E) - \theta) \cup \pi_1(B^* \cap E))$, where:

$$\pi_1(A^* \cap E) - \theta = [d + 1, -2\theta - 1] \cup [-2\theta + d + 1, -\theta - 1] \cup [-\theta + d + 1, n - 1]\,. \tag{18}$$

Notice that $\pi_1(B^* \cap E)$ and $\pi_1(A^* \cap E) - \theta$ are symmetric with respect to $[d + 1, n - 1]$, and therefore their union is $[d + 1, n - 1]$ if and only if the first "gap", $[\theta, \theta + d]$, of $\pi_1(B^* \cap E)$ is contained in the first or second subinterval of $\pi_1(A^* \cap E) - \theta$. This occurs when either $\theta + d \leq n - 2\theta - 1$, which is equivalent to $3\theta + d < n$, or $n - 2\theta + d + 1 \leq \theta$, which is equivalent to $n < 3\theta - d$; thus, since by hypothesis $|n - 3\theta| > d$, we have $m_r = 0$ for all $r \in [d + 1, n - 1]$.

Furthermore, since the boundary of E, ∂E, corresponds to regions at which the coefficient of the right side of (12) is a single λ_i, we can use the complementary technique, checking the coefficients corresponding to $\partial E - (A \cup B \cup C \cup D)$, to reveal that $\lambda_i = 0$ for $i \in [d + 1, \theta - 1] \cup [\theta + 1, n - \theta - 1] \cup [n - \theta + 1, n - d - 1]$. Moreover, we can compare coefficients at the intersection of ∂E and one of A^*, B^*, C^*, or D^*. For $\partial E \cap A^*$, we get the relations $\lambda_i = m_{d+i}$ for $i \in [1, d - 1]$, and for $\partial E \cap B^*$, we get $\lambda_i = m_i$ for $i \in [n - d + 1, n - 1]$. Since we have already shown that such coefficients of M are zero, λ can only be nonzero for the values λ_0, λ_d, λ_θ, $\lambda_{n-\theta}$, and λ_{n-d}.

Using this information we can greatly simplify (13), and as a consequence, get further information about the coefficients of M. In particular, from collecting coefficients for monomials with indices (θ, i), for $i \in [0, d]$, we get the relations $m_0^{q^\theta} + m_i = \lambda_0 + \lambda_{-d}$. Thus, $m_i = m_0$ for $i \leq d$, and, finally, we see that $M = m_0 \pi$.

The preceding theorem gives us precise criteria for when the space of linear maps, S_G, consists of only projected multiplication maps. Furthermore, it was stated in [10] and [15] that these multiplication maps satisfy the relation (2) only if the multiplication commutes with the projection, which happens precisely when the image of the projection is a subspace over an intermediate extension field of $\mathbb{F}_q$. Clearly, in the case $d = 1$, $d + \theta < \frac{n}{2}$, and $|n - 3\theta| > 1$, πk is a hyperplane, and thus $S_G \cong \mathbb{F}_q$, which is optimal.

7 Conclusion

Multivariate public key cryptography has several desirable traits as a potential candidate for post-quantum security. Unfortunately, a standard metric by which we can judge the security of a multivariate scheme has yet to be determined. One consequence of this current status of the field is the similar cryptanalyses of several promising ideas.

We offer the size of the space of linear maps, S_G, illustrating the initial differential symmetries of the core map, f, as a benchmark for the judgement of differential security in modern multivariate public key cryptosystems. As evidence of the feasibility and utility of this method as a measurement of differential security, we measure these spaces for several key players in the evolution of the recent big-field schemes. In the cases of schemes which have been broken, we find that these spaces are large, at least as large as the size of the big field. In the cases of currently considered secure variants, such as the projected SFLASH scheme, pSFLASH, we find that we can make this space as small as possible.

References

1. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Sci. Stat. Comp. **26, 1484** (1997)
2. Chen, A.I.T., Chen, M.S., Chen, T.R., Cheng, C.M., Ding, J., Kuo, E.L.H., Lee, F.Y.S., Yang, B.Y.: Sse implementation of multivariate pkcs on modern x86 cpus. CHES 2009, LNCS, Springer, IACR **5747** (2009) 33–48
3. Chen, A.I.T., Chen, C.H.O., Chen, M.S., Cheng, C.M., Yang, B.Y.: Practical-sized instances of multivariate pkcs: Rainbow, tts, and ℓic-derivatives. Post-Quantum Crypto, LNCS **5299** (2008) 95–106
4. Yang, B.Y., Cheng, C.M., Chen, B.R., Chen, J.M.: Implementing minimized multivariate public-key cryptosystems on low-resource embedded systems. 3rd Security of Pervasive Computing Conference, LNCS **3934** (2006) 73–88
5. Clough, C., Baena, J., Ding, J., Yang, B.Y., Chen, M.S.: Square, a New Multivariate Encryption Scheme. In Fischlin, M., ed.: CT-RSA. Volume 5473 of Lecture Notes in Computer Science., Springer (2009) 252–264
6. Baena, J., Clough, C., Ding, J.: Square-vinegar signature scheme. PQCRYPTO 2008, LNCS **5299** (2008) 17–30
7. Billet, O., Macario-Rat, G.: Cryptanalysis of the square cryptosystems. ASIACRYPT 2009, LNCS **5912** (2009) 451–486

8. 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
9. Shamir, A., Kipnis, A.: Cryptanalysis of the oil & vinegar signature scheme. CRYPTO 1998. LNCS **1462** (1998) 257–266
10. Smith-Tone, D.C.: Properties of the discrete differential with cryptographic applications. PQCRYPTO 2010, LNCS **6061** (2010) 1–12
11. Patarin, J.: Cryptanalysis of the Matsumoto and Imai public key scheme of Eurocrypt 88. Crypto 1995, Springer **963** (1995) 248–261
12. Matsumoto, T., Imai, H.: Public quadratic polynomial-tuples for efficient signature verification and message-encryption. Eurocrypt '88, Springer **330** (1988) 419–545
13. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. EUROCRYPT 1999. LNCS **1592** (1999) 206–222
14. Patarin, J.: The oil and vinegar algorithm for signatures. Presented at the Dagsthul Workshop on Cryptography (1997)
15. 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