Cryptanalysis of SFLASH with Slightly Modified Parameters

Vivien Dubois, Pierre-Alain Fouque, and Jacques Stern

École normale supérieure
DI, 45 rue d'Ulm, 75230 Paris cedex 05, France
{vivien.dubois,pierre-alain.fouque,jacques.stern}@ens.fr

Abstract. SFLASH is a signature scheme which belongs to a family of multivariate schemes proposed by Patarin *et al.* in 1998 [9]. The SFLASH scheme itself has been designed in 2001 [8] and has been selected in 2003 by the NESSIE European Consortium [6] as the best known solution for implementation on low cost smart cards. In this paper, we show that slight modifications of the parameters of SFLASH within the general family initially proposed renders the scheme insecure. The attack uses simple linear algebra, and allows to forge a signature for an arbitrary message in a question of minutes for practical parameters, using only the public key. Although SFLASH itself is not amenable to our attack, it is worrying to observe that no rationale was ever offered for this "lucky" choice of parameters.

1 Introduction

Multivariate Cryptography is an area of research which attempts to build asymmetric primitives, based on hard computational problems related to multivariate quadratic polynomials over a finite field. Multivariate schemes have recently received much attention, for several reasons. First, the hard problems of reference are not known to be polynomial in the quantum model, unlike integer factorization and the discrete logarithm problems. More importantly, Multivariate Cryptography offers a large collection of primitives and problems of a new flavor. In general, multivariate schemes require modest computational resources and can be implemented on low cost smart cards. Moreover, these schemes benefit from several nice properties such as providing very short or very fast signatures. Also, they are quite versatile: a number of generic non-exclusive variations can be derived from a few basic schemes. Even when the original schemes are weak, variations are often considered to avoid structural attacks.

One of the more elaborate outcomes of Multivariate Cryptography is probably the SFLASH signature scheme. Designed by Patarin *et al.* [8], it is among the fastest signatures schemes known, with NTRUSign and TTS [4,11]. Although initial tweaks in the first version of SFLASH were shown inappropriate [3], the second version of SFLASH is currently considered secure, as testified from the recent acceptance of this primitive by the NESSIE European Consortium [6].

The structure of SFLASH is among the simplest in Multivariate Cryptography. Roughly speaking, SFLASH is a truncated C^* scheme. The C^* scheme was invented by Matsumoto and Imai in 1988 [5], and was shown to be insecure by Patarin in 1995 [7]. Later, Patarin *et al.* considered the simple variation of C^* consisting in deleting from the public key a large number of coordinates [9]. Schemes derived from C^* by this principle are called C^{*-} schemes; they are well suited for signature. As soon as the number of deleted coordinates is large enough, C^{*-} schemes are considered secure. SFLASH belongs to the C^{*-} family and has been chosen as a candidate for the NESSIE selection, and finally accepted.

Our Results. We argue that the security of the C^{*-} schemes remains insufficiently understood. In particular, one may rightfully question the reasons for the particular choice of parameters opted for in SFLASH. Might other parameters yield the same security ?

In this paper, we show that many choices of parameters for C^{*-} schemes are insecure. Our approach uses basic properties of the differential as introduced in [2]. Since the differential is bilinear and symmetric, it seems natural to consider skew-symmetric linear maps with respect to this function. This property is so specific and overdefined that the space of skew-symmetric maps is left unchanged when we replace the full public key of C^* by its truncated version C^{*-}, even when the number of deleted coordinates is very large. Skew-symmetric maps can be recovered from their defining equation in terms of the differential of a C^{*-} public key, using only linear algebra. Once this has been achieved, compositions of these maps with the public key can be used to recover a full C^* public key, which can then be inverted using the original attack by Patarin [7].

The schemes under attack are those for which the internal C^* parameter and the number of variables are not coprime. Such parameters are perfectly acceptable for practical realizations of C^{*-} schemes in the current state of cryptanalysis. SFLASH with the recommended set of parameters escapes this attack. However, this shows that the elements underlying the security of C^{*-} schemes and their relations with parameters are not well identified. To illustrate this point, we show that changing the parameters of SFLASH by one renders the scheme breakable in a few minutes.

Organization of the Paper. In Section 2, we recall the definition of C^* and C^{*-} schemes. Then, in Section 3, we characterize skew-symmetric maps with respect to the differential of C^*. In Section 4, we show that the same maps can be recovered from a truncated version C^{*-} of the original C^* public key. Finally, in section 5, we show how their use allows us to restore a full C^* public key.

2 C^* and C^{*-}

Before we describe the C^* and C^{*-} schemes, we recall the generic construction of multivariate schemes.

2.1 The Generic Construction of Multivariate Schemes

We denote by $\mathbb{F}_q^n$ the n-dimensional vector space over the finite field $\mathbb{F}_q$ with q elements. A function from $\mathbb{F}_q^n$ to $\mathbb{F}_q^m$ is defined by m coordinate-polynomials in n variables. When these polynomials have multivariate degree 2, the function is termed quadratic. Finding a preimage of a quadratic function involves solving a multivariate quadratic system of equations, an NP-hard problem in general. Nevertheless, some classes of easily invertible quadratic functions are known and can form the basis of a multivariate asymmetric scheme. More precisely, the generic construction of multivariate schemes is the following. The key generation algorithm hides an easily invertible quadratic function F by two linear (or affine) changes of coordinates U and T into a function P defined by

$$P = T \circ F \circ U$$

P is the public key and U, T are the secret key. The proponents of multivariate cryptography argue that the function P is a random-looking quadratic function, which is expected to be hard to invert by general purpose techniques. An encrypted message can be decrypted by using the secret key (T, U) to undo the hiding process and by solving the easy internal quadratic system.

2.2 The C^* scheme

The C^* scheme was proposed by Matsumoto and Imai in 1988 [5]. In the C^* scheme, the internal easy-to-invert function is defined from a monomial over the degree n extension field of $\mathbb{F}_q$, denoted $\mathbb{F}_{q^n}$, of the form

$$F(x) = x^{1+q^\theta}$$

where θ is a positive integer. The function F is isomorphic to a quadratic function from $\mathbb{F}_q^n$ into itself and provided q is even, the integer θ can be chosen so that F is a permutation. This happens if and only if $\gcd(q^\theta + 1, q^n - 1) = 1$. In Appendix A, we show that, denoting by d the gcd of θ and n, this is equivalent to the condition that $\frac{n}{d}$ is odd.

The C^* scheme, as previously described, was shown to be insecure by Patarin [7]. It was observed that, for any x, y such that $F(x) = y$, we have

$$y^{q^\theta}.x - y.x^{q^{2\theta}} = 0$$

It follows that there exist n bilinear relations between a ciphertext and the corresponding plaintext. These bilinear relations can be found from the public key with a few plaintext-ciphertext pairs. Using these bilinear relations allows us to recover the plaintext from any ciphertext, by linear algebra.

Several ways to withstand the attack by Patarin were later considered. Among the most promising, are the C^{*-} schemes. In the next section, we recall these schemes in detail.

2.3 C^{*-} schemes

A C^{*-} scheme is derived from a C^* scheme by simply deleting from the C^* public key some of the quadratic polynomials. More precisely, for some additional parameter r, the key generation builds a C^* scheme and then deletes from the public key the last r coordinates. In the sequel, Π will denote the projection on the first $(n-r)$ coordinates, P the C^* public key, and P_Π the resulting C^{*-} public key.

To find a preimage under P_Π of a string y of $(\mathbb{F}_q)^{n-r}$, the user first has to pad y with some string k of $(\mathbb{F}_q)^r$, and then has to find the preimage of (y,k) by P using its secret key U,T. Using a C^{*-} scheme for encryption is therefore quite awkward: to recover the plaintext, the user has to review all possible paddings k, compute for each k the corresponding preimage, and identify the plaintext among these preimages by using some message redundancy. However, C^{*-} schemes are well-suited for signature, even for large q and r, since in this setting any of the q^r preimages of y by P_Π is a valid signature of y. To sign the message y, the user chooses an arbitrary k and the signature consists in the preimage of (y,k) by P. In the sequel, we only consider the C^{*-} scheme in the signature setting.

C^{*-} schemes were first introduced by Patarin *et al.* [9], but the idea of enhancing the security of multivariate schemes by deleting a few coordinates from the public key first appeared in Shamir [10]. In [9], Patarin *et al.* describe a technique for reconstructing a C^* public key from a C^{*-} public key with complexity of the order of q^r. Accordingly, the parameters q and r must be chosen such that $q^r \simeq 2^{80}$ for practical instantiations of C^{*-} schemes. The illustrative notation C^{*--} is sometimes used in this case. No condition is specified in the literature for choosing the parameter θ besides the obvious requirement that the corresponding monomial should be invertible and, as seen before, all values of θ whose gcd d with n is such that $\frac{n}{d}$ is odd can be chosen. In fact, choosing a large d allows a faster inversion of the C^* monomial, as observed by Ding [1], and can be an attractive choice.

SFLASH. Practical instantiations of C^{*-} schemes were proposed by Patarin *et al.* as candidates to the European call for primitives NESSIE in 2001. These instantiations were called FLASH and SFLASH. Initially, some tweak was added to SFLASH to decrease the size of the public key, however this tweak rendered the scheme insecure, as shown by Gilbert and Minier in 2002 [3], and discarded. Without this tweak, FLASH and SFLASH are very similar and therefore, only SFLASH was later considered by the NESSIE evaluation process, and finally accepted in 2003. The recommended parameters of SFLASH are $q = 2^7$, $n = 37$, $\theta = 11$ and $r = 11$; signatures are 239 bits long. Until now, no weakness was reported in either SFLASH or the general design principle of C^{*-} schemes.

In the sequel, we will show that many C^{*-} schemes are insecure. The C^{*-} schemes under attack are those for which the gcd d of θ and n is not 1. Note that this is different than the condition that $\frac{n}{d}$ is odd, which is needed to make the mapping invertible. The attack makes it possible to forge a signature in a matter of minutes for practical parameters. The attack does not apply to SFLASH for which the recommended parameters θ and n are coprime. However this "lucky" choice appears to be accidental since no rationale was offered for it.

3 Skew-Symmetric Maps w.r.t the Differential of C^*

In this section, we consider some properties of the differential of the internal C^* monomial. Implications of these properties to C^{*-} schemes will be addressed in the next section.

The differential, defined as follows, can be considered for any quadratic function F. For any element a, the difference function $x \mapsto F(x+a) - F(x)$ is affine and its constant term is $F(a) - F(0)$. Its linear part is called the differential of F at a and is denoted $DF(a,x)$:

$$DF(a,x) = F(x+a) - F(x) - F(a) + F(0)$$

$DF(a,x)$ is actually bilinear and symmetric when considered as a function of a and x. Our attack is based on considering skew-symmetric maps with respect to this bilinear function *i.e.* linear maps M such that for all choices of x and a

$$DF(a, M(x)) + DF(M(a), x) = 0$$

This is a very strong condition, and when F is defined by a random collection of quadratic polynomials, only trivial solutions M are expected to satisfy this condition. However, when $F(x) = x^{1+q^\theta}$, its differential is

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

The skew-symmetric maps with respect to the differential of such a C^* monomial are given by the following theorem.

Theorem 1. *Let M be a linear map; M is skew-symmetric with respect to $DF(a,x)$ if and only if M is the multiplication by some element ξ satisfying $\xi^{q^\theta} + \xi = 0$.*

Proof. A linear map M over $\mathbb{F}_{q^n}$ is a sum of q-powerings : $M(x) = \sum_{i=0}^{n-1} \lambda_i x^{q^i}$. When DF is the differential of the C^* monomial given by (1), we get for any elements a, x in $\mathbb{F}_{q^n}$

$$\sum_{i=0}^{n-1} \lambda_i a^{q^\theta} x^{q^i} + \sum_{i=0}^{n-1} \lambda_i^{q^\theta} a\, x^{q^{i+\theta}} + \sum_{i=0}^{n-1} \lambda_i a^{q^i} x^{q^\theta} + \sum_{i=0}^{n-1} \lambda_i^{q^\theta} a^{q^{i+\theta}} x = 0$$

Since the monomials $a^{q^u} x^{q^v}$ are a basis of the space of bilinear maps over $\mathbb{F}_{q^n}$, we obtain the following equations corresponding to the various elements of the basis

$$\lambda_0 + \lambda_0^{q^\theta} = 0 \qquad \text{(coefficient of } ax^{q^\theta}\text{)}$$

$$\lambda_i = 0,\ i \neq 0, \theta \qquad \text{(coefficient of } a^{q^i} x^{q^\theta},\ i \neq 0, \theta\text{)}$$

$$(\lambda_\theta)^{q^\theta} = 0 \qquad \text{(coefficient of } ax^{q^{2\theta}}\text{)}$$

Conversely, it is straightforward to see that multiplications by an element ξ satisfying $\xi^{q^\theta} + \xi = 0$ are skew-symmetric with respect to DF :

$$DF(a, \xi.x) + DF(\xi.a, x) = \xi^{q^\theta} a^{q^\theta} x + \xi a x^{q^\theta} + a^{q^\theta} \xi x + a \xi^{q^\theta} x^{q^\theta} = 0$$

which concludes the proof. □

We denote by $\mathcal{K}_\theta$ the set of the elements ξ such that $\xi^{q^\theta} + \xi = 0$. By the linearity of q-powerings, this is a linear space. The non-zero elements of $\mathcal{K}_\theta$ are the $(q^\theta - 1)$-th roots of the unity and the number of these elements is $\gcd(q^\theta - 1, q^n - 1) = q^d - 1$ where d is the gcd of θ and n. Consequently, $\mathcal{K}_\theta$ is a linear space of dimension d.

For any element ξ in $\mathcal{K}_\theta$, we denote by M_ξ multiplication by ξ. As stated by the theorem, the maps M_ξ are the skew-symmetric applications with respect to the differential of the C^* monomial. They form a linear space isomorphic to $\mathcal{K}_\theta$. When $d = 1$, $\mathcal{K}_\theta$ is generated by 1, and all the maps M_ξ are colinear to the identity. This case is trivial since scalar multiples of the identity are skew-symmetric with respect to any bilinear product. Accordingly, non-trivial maps M_ξ only exist when $d > 1$.

4 Recovering the Skew-Symmetric Maps from a C^{*-} Public Key

Let P be a C^* public key and let P_Π be the C^{*-} public key obtained from P by deleting the last r coordinates. Since P is a composition $T \circ F \circ U$ where F is the internal C^* monomial and U, T are secret changes of coordinates, P_Π is the composition $T_\Pi \circ F \circ U$ where T_Π is obtained from T by removing the last r rows. The differential of P_Π is

$$DP_\Pi(a, x) = T_\Pi \left(DF(U(a), U(x)) \right)$$

Since $DF(U(a), U(x))$ is isomorphic by U to $DF(a, x)$, the skew-symmetric maps with respect to $DF(U(a), U(x))$ are the maps denoted N_ξ defined by

$$N_\xi = U^{-1} \circ M_\xi \circ U$$

By the linearity of T_Π, all the maps N_ξ are also skew-symmetric with respect to the truncated DP_Π:

$$DP_\Pi(a, N_\xi(x)) + DP_\Pi(N_\xi(a), x) = 0$$

We argue that they are likely to be the only ones, even when the number r of deleted coordinates is very close to n.

For any pair (a, x), the equation

$$DP_{\Pi}(a, L(x)) + DP_{\Pi}(L(a), x) = 0 \tag{2}$$

gives us $n - r$ linear equations in the n^2 coefficients of the unknown L. Since Equation (2) is bilinear and symmetric in (a, x) and trivial when $a = x$, taking n^2 linearly independent choices for a and x, we construct a system of $(n-r)n(n-1)/2$ linear equations in the n^2 coefficients of L. The kernel of these equations must contain the d-dimensional space formed by the maps N_ξ. Assuming that all the generated linear equations are otherwise independent, the kernel does not contain other solutions up to r satisfying

$$(n-r)\frac{n(n-1)}{2} \geq n^2 - d$$

According to this heuristic, the maps N_ξ are likely to be the only solutions of our greatly overdefined system of linear equations provided that $r \leq r_{max}$ where

$$r_{max} = n - \left\lceil 2\frac{n^2 - d}{n(n-1)} \right\rceil = n - 3$$

which is very close to n. Consequently, we expect to find the same linear subspace of solutions even if we delete from the C^* public key almost all the quadratic polynomials, in order to generate the C^{*-} public key.

Though this analysis is rather naive, it provides a good estimate of the actual value of r_{max} as observed from some computer experiments. In the table below, we report on the actual value of r_{max} found for several parameters, to be compared with the heuristic value $n - 3$.

n	36	36	38	39	39	40	42	42	44
θ	8	12	10	13	9	8	12	14	12
d	4	12	2	13	3	8	6	14	4
r_{max}	33	32	35	35	36	37	39	38	41

The parameters chosen for these experiments are very close to the recommended parameters $n = 37$ and $\theta = 11$ for SFLASH, with the same value of $q = 2^7$. Note that in practice r would be chosen to be much smaller than n – about $\frac{n}{3}$ in SFLASH – and thus our approach could be easily applied even if not all the equations happen to be sufficiently independent.

Using Equation (2) with n^2 independent choices for a and x, we find all maps N_ξ by linear algebra. This takes a few seconds for practical parameters.

5 Recomposing a C^* Public Key using Skew-Symmetric Maps

At this point, we assume that the linear space of skew-symmetric maps N_ξ has been computed. Non-trivial N_ξ are those which are not colinear to the identity.

For any non-trivial N_ξ, we can now generate two C^{*-} public keys P_Π and $P_\Pi \circ N_\xi$. We next show that, provided r is at most $\frac{n}{2}$, completing P_Π with r arbitrary polynomials from $P_\Pi \circ N_\xi$ creates a valid C^* public key with high probability. Though higher values of r are not of practical interest, the technique can be generalized to $r \leq n(1 - \frac{1}{d})$ using $d - 1$ linearly independent non-trivial maps N_ξ, as shown in Appendix B.

Let us recall that the function P_Π is a composition $T_\Pi \circ F \circ U$, where U is a secret isomorphism, F is the C^* monomial and T_Π consists of $n-r$ linearly independent rows. Besides, N_ξ equals $U^{-1} \circ M_\xi \circ U$, where M_ξ denotes multiplication by ξ. The composition of P_Π and N_ξ is

$$P_\Pi \circ N_\xi = T_\Pi \circ F \circ M_\xi \circ U$$

Since F is multiplicative, multiplying the input by ξ results in multiplying the output by $F(\xi)$. Therefore

$$P_\Pi \circ N_\xi = T_\Pi \circ M_{F(\xi)} \circ F \circ U$$

where $M_{F(\xi)}$ denotes the multiplication by $F(\xi)$. Since N_ξ is non-trivial, ξ is not colinear to 1, and since the inverse of F is a power function, $F(\xi)$ is not colinear to 1 either. Hence, $M_{F(\xi)}$ is non-trivial and the matrices T_Π and $T_\Pi \circ M_{F(\xi)}$ are distinct.

The $n-r$ quadratic polynomials defining P_Π are linear combinations encoded by the rows of T_Π of the n quadratic polynomials defining $F \circ U$, whereas the $n - r$ quadratic polynomials defining $P_\Pi \circ N_\xi$ are different linear combinations encoded by the rows of $T_\Pi \circ M_{F(\xi)}$ of the same n quadratic polynomials defining $F \circ U$. Adding r polynomials of $P_\Pi \circ N_\xi$ to P_Π recomposes a valid C^* public key if and only if the corresponding rows of $T_\Pi \circ M_{F(\xi)}$ added to the rows of T_Π form a full rank system. Let us select, for instance, the r first rows of $T_\Pi \circ M_{F(\xi)}$. The rows of T_Π generate a subspace of dimension $n-r$ of $(\mathbb{F}_q)^n$. A random vector lies in a subspace of dimension $n - k$ of $(\mathbb{F}_q)^n$ with probability q^{-k}. Therefore, if we assume that the selected rows of $T_\Pi \circ M_{F(\xi)}$ are random vectors, the probability that they form with the rows of T_Π a full rank system is

$$\left(1 - \frac{1}{q^r}\right)\left(1 - \frac{1}{q^{r-1}}\right) \ldots \left(1 - \frac{1}{q}\right) \simeq 1 - \frac{1}{q}$$

With this probability, adding the r first polynomials of $P_\Pi \circ N_\xi$ to P_Π will recover a valid C^* public key (which is not necessarily identical to the C^* key we started with). This public key corresponds to a secret key T obtained by adding to T_Π the first r rows of $T_\Pi \circ M_{F(\xi)}$. We then apply Patarin's attack and recover n message-signature bilinear relations. If adding the r first polynomials fails to recover a C^* public key (which can be detected by the failure of Patarin's attack), we can retry with a different set of r polynomials of $P_\Pi \circ N_\xi$, or try a different value of ξ. The probability of success in t independent trials is expected to be $q^{-(t-1)}$.

The table below provides some timings (in seconds) for an actual implementation of our attack on a single PC. We successfully recovered a C^* public key from a C^{*-} public key for all the listed values of the parameters n, θ which are close to those of SFLASH and with the same value of $q = 2^7$.

n	36	36	38	39	39	40	42	42	44
θ	8	12	10	13	9	8	12	14	12
d	4	12	2	13	3	8	6	14	4
r	11	11	11	12	12	12	13	13	13
$C^{*-} \mapsto C^*$	$57s$	$57s$	$94s$	$105s$	$90s$	$105s$	$141s$	$155s$	$155s$

6 Forging Signatures using Patarin's attacks

Our attack makes it possible to recover a C^* public key from a C^{*-} public key in a few seconds for practical parameters. Then, it remains to apply Patarin's attack to this public key and this is the "expensive" step of the attack.

As shown in [7], once Patarin's bilinear relations have been computed, we get for any message a subspace of dimension d containing at least one valid signature. Finding this signature requires trying all the q^d elements of this subspace. When d is large, additional linear equations can be generated to avoid exhaustive search using another attack also described in [7] which takes advantage of a small value of $\frac{n}{d}$.

The first attack, involving a precomputation in time $(\log_2 q)^2 n^6$ and then $q^d (\log_2 q)^2 n^3$ for each signature, is efficient when d is small. The second attack, involving a precomputation in time $(\log_2 q)^2 n^{3\frac{k+1}{2}}$ where $k = \frac{n}{d}$ and then $(\log_2 q)^2 n^3$ for each signature, is efficient when $\frac{n}{d}$ is small.

We summarize in the table below the complexities of Patarin's attacks for several choices of parameters which are close to those of SFLASH (and with the same value of $q = 2^7$). The star symbol at parameter d or $\frac{n}{d}$ specifies which of the two attacks devised by Patarin is considered.

n	36	36	38	39	39	40	42	42	44
θ	8	12	10	13	9	8	12	14	12
d	4^*	12	2^*	13	3^*	8	6^*	14	4^*
n/d	9	3^*	19	3^*	13	5^*	7	3^*	11
r	11	11	11	12	12	12	13	13	13
Precomputation	2^{36}	2^{36}	2^{36}	2^{36}	2^{36}	2^{51}	2^{36}	2^{36}	2^{36}
Signature forgery	2^{49}	2^{21}	2^{35}	2^{21}	2^{36}	2^{21}	2^{57}	2^{21}	2^{49}

7 Conclusion

We have demonstrated a very simple but very powerful attack against a large class of C^{*-} schemes, namely those for which the number of variables n and the C^* parameter θ are not coprime. This attack transforms any such C^{*-} scheme into a full C^* scheme in a few seconds, even when the number of deleted coordinates is much larger than encountered for practical purposes. This is a major discovery since it was currently believed that even a weak scheme such as C^* can be made secure by simply deleting a sufficiently large number of coordinates from the public key. We have shown that this design fails for some choices of parameters. This shows that the security of C^{*-} schemes relies on mechanisms which are more subtle than anticipated, and does not necessarily improve when we increase the parameters. In particular, it is quite worrying to observe that no rationale was ever offered for the parameters recommended for SFLASH. It should be added that further unpublished work performed by the authors together with Adi Shamir has shown that the weakness of C^{*-} schemes was not only a matter of parameter choice, since they were able to mount a practical attack against the actual SFLASH schemes.

Acknowledgements. We are very grateful to Adi Shamir for interesting discussions and helpful remarks. Part of this work is supported by the Commission of the European Communities through the IST program under contract IST-2002-507932 ECRYPT.

References

1. J. Ding. A New Variant of the Matsumoto-Imai Cryptosystem through Perturbation. In *PKC '04*, LNCS 2947, pages 305–318. Springer-Verlag, 2004.
2. P. A. Fouque, L. Granboulan, and J. Stern. Differential Cryptanalysis for Multivariate Schemes. In *Eurocrypt '05*, LNCS 3494, pages 341–353. Springer-Verlag, 2005.
3. H. Gilbert and M. Minier. Cryptanalysis of SFLASH. In *Eurocrypt '02*, LNCS 2332, pages 288–298. Springer-Verlag, 2002.
4. J. Hoffstein, N. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte. NTRUSIGN : Digital Signatures Using the NTRU Lattice. In *CT-RSA '03*, LNCS 2612, pages 122–140. Springer-Verlag, 2003.
5. T. Matsumoto and H. Imai. Public Quadratic Polynomial-tuples for Efficient Signature-Verification and Message-Encryption. In *Eurocrypt '88*, LNCS 330, pages 419–453. Springer-Verlag, 1988.
6. NESSIE. New European Schemes for Signatures Integrity and Encryption. Portfolio of recommended cryptographic primitives. `http://www.nessie.eu.org/index.html`
7. J. Patarin. Cryptanalysis of the Matsumoto and Imai Public Key Scheme of Eurocrypt'88. In *Crypto '95*, LNCS 963, pages 248–261. Springer-Verlag, 1995.
8. J. Patarin, N. Courtois, and L. Goubin. FLASH, a Fast Multivariate Signature Algorithm. In *CT-RSA '01*, LNCS 2020, pages 297–307. Springer-Verlag, 2001.
9. J. Patarin, L. Goubin, and N. Courtois. C^*_{-+} and HM : Variations Around Two Sechemes of T. Matsumoto and H. Imai. In *Asiacrypt '98*, LNCS 1514, pages 35–49. Springer-Verlag, 1998.
10. A. Shamir. Efficient Signature Scheme Based on Birational Permutations. In *Crypto '93*, LNCS 773, pages 1–12. Springer-Verlag, 1993.
11. B. Y. Yang and J. M. Chen. Building Secure Tame-like Multivariate Public-Key Cryptosystems: The New TTS. In *ACISP '05*, LNCS 3574, pages 518–531. Springer-Verlag, 2005.

A Constructing a Bijective C^* Monomial

The internal C^* monomial x^{1+q^θ} is bijective in the field $\mathbb{F}_{q^n}$ if and only if $q^\theta+1$ and q^n-1 are coprime.

When q is odd, both $q^\theta+1$ and q^n-1 are even, and their gcd is a multiple of 2. Therefore, q odd never yields a bijective C^* monomial.

When q is even, then $q^\theta-1$ and $q^\theta+1$ are coprime and therefore

$$\gcd(q^{2\theta}-1,q^n-1)=\gcd(q^\theta-1,q^n-1).\gcd(q^\theta+1,q^n-1)$$

We denote by A,B and C the above gcds. We determine A and B and then deduce C. Denoting by d the gcd of θ and n, B equals q^d-1. On the other hand, A equals $q^{\gcd(2\theta,n)}-1$. We have

$$\gcd(2\theta,n)=d.\gcd(2\frac{\theta}{d},\frac{n}{d})$$

and since $\frac{\theta}{d}$ and $\frac{n}{d}$ are coprime, the right-hand gcd is 2 when $\frac{n}{d}$ is even and 1 otherwise. Hence, A equals $q^{2d}-1$ when $\frac{n}{d}$ is even, and q^d-1 when $\frac{n}{d}$ is odd. Finally, C equals q^d+1 when $\frac{n}{d}$ is even, and 1 when $\frac{n}{d}$ is odd.

The choices of θ and n yielding a bijective C^* monomial are therefore those for which $\frac{n}{d}$ is odd.

B Recovering a full C^* when r is over $\frac{n}{2}$

In Section 5, we have shown how to recover a C^{*-} public key into a full C^* public key, using one single non-trivial skew-symmetric map N_ξ, when $r\le\frac{n}{2}$. In this appendix, we show that this technique can be generalized up to

$$r=\min\left\{\,r_{max}\,;\,n\left(1-\tfrac{1}{d}\right)\,\right\}$$

Let us recall that r_{max} is the maximal value of r allowing to recover the d-dimensional space of skew-symmetric maps. This value can be found experimentally and is given in Section 4 for some parameters. For r smaller than r_{max}, let $N_\xi^1,\ldots,N_\xi^{d-1}$ form with the identity a basis of the space of skew-symmetric maps. Aside from P_Π, we get $d-1$ independent C^{*-} public keys $P_\Pi\circ N_\xi^1,\ldots,P_\Pi\circ N_\xi^{d-1}$. We use coordinates of these additional C^{*-} public key to complete P_Π into a full C^* public key. The overall number of coordinates available is $d(n-r)$, so that there is no hope to recover a full C^* if $r>n(1-\frac{1}{d})$. When all coordinates are linearly independent, we can recover a full C^* up to $r=n(1-\frac{1}{d})$. This has never failed to work in practice. The table below provides timings for some parameters and the largest value of r allowing the attack. The star symbol at parameter r indicates that the value considered corresponds to r_{max}.

n	36	36	38	39	39	40	42	42	44
θ	8	12	10	13	9	8	12	14	12
d	4	12	2	13	3	8	6	14	4
$r = \min\{r_{max}, n(1-\frac{1}{d})\}$	27	32^*	19	35^*	26	35	35	38^*	33
$C^{*-} \mapsto C^*$	$65s$	$51s$	$112s$	$79s$	$107s$	$95s$	$134s$	$117s$	$202s$