Practical Cryptanalysis of SFLASH

Vivien Dubois[1], Pierre-Alain Fouque[1], Adi Shamir[1,2], and Jacques Stern[1]

[1] École normale supérieure
Département d'Informatique 45, rue d'Ulm
75230 Paris cedex 05, France
Vivien.Dubois@ens.fr,
Pierre-Alain.Fouque@ens.fr, Jacques.Stern@ens.fr
[2] Weizmann Institute of Science Adi.Shamir@weizmann.ac.il

Abstract. In this paper, we present a practical attack on the signature scheme SFLASH proposed by Patarin, Goubin and Courtois in 2001 following a design they had introduced in 1998. The attack only needs the public key and requires about one second to forge a signature for any message, after a one-time computation of several minutes. It can be applied to both SFLASH^{v2} which was accepted by NESSIE, as well as to SFLASH^{v3} which is a higher security version.

1 Introduction

In the last twenty years, multivariate cryptography has emerged as a potential alternative to RSA or DLOG [12, 2] schemes. Many schemes have been proposed whose security appears somehow related to the problem of deciding whether or not a quadratic system of equations is solvable, which is known to be NP-complete [5]. An attractive feature of such schemes is that they have efficient implementations on smart cards, although the public and secret keys are rather large. Contrary to RSA or DLOG schemes, no polynomial quantum algorithm is known to solve this problem.

The SFLASH Scheme. SFLASH is based on the Matsumoto-Imai scheme (MI) [7], also called the C^* scheme. It uses the exponentiation $x \mapsto x^{q^\theta+1}$ in a finite field $\mathbb{F}_{q^n}$ of dimension n over a binary field $\mathbb{F}_q$, and two affine maps on the input and output variables. The MI scheme was broken by Patarin in 1995 [8]. However, based on an idea of Shamir [13], Patarin *et al.* proposed at CT-RSA 2001 [10] to remove some equations from the MI public key and called the resulting scheme C^{*-}. This completely avoids the previous attack and, although not appropriate for an encryption scheme, it is well-suited for a signature scheme. The scheme was selected in 2003 by the NESSIE European Consortium as one of the three recommended public key signature schemes, and as the best known solution for low cost smart cards.

Previous Attacks on SFLASH. The first version of SFLASH, called SFLASH^{v1}, is a more efficient variant of C^{*-} using a small subfield. It has been attacked by Gilbert and Minier in [6]. However, the later versions (SFLASH^{v2} and SFLASH^{v3}) were immune to this attack.

Recently, Dubois, Fouque and Stern in [1] proposed an attack on a special class of SFLASH-like signatures. They show that when the kernel of the linear map $x \mapsto x + x^{q^\theta}$ is non-trivial, the C^{*-} scheme is not secure. The attack is very efficient in this case, but relies on some specific properties which are not met by the NESSIE proposals and which make the scheme look less secure.

Our Results. In this paper, we achieve a total break of the NESSIE standard with the actual parameters suggested by the designers: given only the public key, a signature for any message can be forged in about one second after a one time computation of several minutes. The asymptotic running time of the attack is $O(\log^2(q)n^6)$ since it only needs standard linear algebra algorithms on $O(n^2)$ variables, and n is typically very small. As in [1], the basic strategy of the attack is to recover additional independent equations in order to apply Patarin's attack [8]. To this end, both attacks use the differential of the public key. However, the attacks differ in the way the invariants related to the differential are found. The differential of the public key, also called its polar form, is very important since it transforms quadratic equations into linear ones. Hence, it can be used to find some linear relations that involve the secret keys. Its cryptanalytic significance had been demonstrated in [4].

Organization of the Paper. In section 2, we describe the SFLASH signature scheme and the practical parameters recommended by Patarin *et al.* and approved by NESSIE. Then, in section 3 we present the multiplicative property of the differential that we need. Next, in section 4 we describe how to recover linear maps related to multiplications in the finite field from the public key. In section 5, we show how to break the NESSIE proposal given only the public key. In section 6, we extend the attack to cover the case when up to half of the equations are removed, and finally in section 7, we compare our method with the technique of [1] before we conclude.

2 Description of SFLASH

In 1988, Matsumoto and Imai [7] proposed the C^* scheme for encryption and signature. The basic idea is to hide a quadratic easily invertible mapping F in some large finite field $\mathbb{F}_{q^n}$ by two secret invertible linear (or affine) maps U and T which mix together the n coordinates of F over the small field $\mathbb{F}_q$:

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

where $F(x) = x^{q^\theta+1}$ in $\mathbb{F}_{q^n}$. This particular form was chosen since its representation as a multivariate mapping over the small field is quadratic, and thus the size of the public key is relatively small.

The secret key consists of the maps U and T; the public key P is formed by the n quadratic expressions, whose inputs and outputs are mixed by U and T, respectively. It can be seen that F and P are invertible whenever $\gcd(q^\theta + 1, q^n - 1) = 1$, which implies that q has to be a power of 2 since q is a prime power.

This scheme was successfully attacked by Patarin [8] in 1996. To avoid this attack and restore security Patarin *et al.* proposed in [11] to remove from the public key the last r quadratic expressions (out of the initial n), and called this variant of C^* schemes, C^{*-}. Furthermore, if the value of r is chosen such that $q^r \geq 2^{80}$, then the variant is termed C^{*--}. If we denote by Π the projection of n variables over $\mathbb{F}_q$ onto the first $n - r$ coordinates, we can represent the public key by the composition :

$$P_\Pi = \Pi \circ T \circ F \circ U = T_\Pi \circ F \circ U.$$

In the sequel, P denotes the public key of a C^* scheme whereas P_Π denotes a C^{*-} or C^{*--} public key. In both cases the secret key consists of the two linear maps T and U.

To sign a message m, the last r coordinates are chosen at random, and the signer recovers s such that $P_\Pi(s) = m$ by inverting T, U and F. A signature (m, s) can be checked by computing $P_\Pi(s)$ with the public key, which is extremely fast since it only involves the evaluation of a small number of quadratic expressions over the small finite field $\mathbb{F}_q$.

For the NESSIE project and in [10], Patarin *et al.* proposed two particular recommended choices for the parameters of C^{*--} :

- for SFLASH^{v2} : $q = 2^7$, $n = 37$, $\theta = 11$ and $r = 11$
- for SFLASH^{v3} : $q = 2^7$, $n = 67$, $\theta = 33$ and $r = 11$

SFLASH^{v3} was actually proposed to provide an even more conservative level of security than SFLASH^{v2} [10]. However, the designers made clear that they viewed SFLASH^{v2} as providing adequate security, and no attack on these two choices of parameters had been reported so far.

The important fact to notice here is that in both cases $\gcd(n, \theta) = 1$ and thus the attack described in [1] on a modified version of SFLASH in which $\gcd(n, \theta) > 1$ cannot be applied. The attack described in this paper shares with [1] the basic observation about the multiplicative property of C^{*-} schemes which is described in section 3, but proceeds in a completely different way. More discussion about the relationships between the two attacks can be found in section 7.

3 The Multiplicative Property of the Differential

The attack uses a specific multiplicative property of the differential of the public key of a C^{*-} scheme.

The differential of the internal quadratic system $F(x) = x^{q^\theta+1}$ is a symmetric bilinear function in $\mathbb{F}_{q^n}$, called DF, and it is defined for all $a, x \in \mathbb{F}_{q^n}$ by the linear operator :

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

When $F(x) = x^{q^\theta+1}$, we get for all $a, x \in \mathbb{F}_{q^n}$

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

Note that this expression is bilinear since exponentiation by q^θ is a linear operation. This map has a very specific multiplicative property: for all $\xi \in \mathbb{F}_{q^n}$

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

We now explain how this identity on the internal polynomial induces a similar one on the differential of the public keys in C^* and C^{*-}. Due to the linearity of the DP operator, we can combine it with the linear maps T and U to get that the differential of any C^* public key P is $DP(a,x) = T \circ DF(U(a), U(x))$. Then, equation (1) becomes for any $\xi \in \mathbb{F}_{q^n}$:

$$\begin{aligned} T \circ DF(\xi \cdot U(a), U(x)) + T \circ DF(U(a), \xi \cdot U(x)) \\ &= T \circ (\xi + \xi^{q^\theta}) \cdot DF(U(a), U(x)) \\ &= T \circ (\xi + \xi^{q^\theta}) \cdot T^{-1}(DP(a,x)). \end{aligned}$$

We denote by M_ξ and $M_{L(\xi)}$ respectively the multiplications by ξ and by $L(\xi) = \xi + \xi^{q^\theta}$. Also, we let N_ξ denote the linear map $U^{-1} \circ M_\xi \circ U$ which depends on the secret key. We still use the word "multiplication" for N_ξ, even though this wording is not actually accurate since this is not the standard multiplication in $\mathbb{F}_{q^n}$, due to the action of the input transformation U. With these notations :

$$DP(N_\xi(a), x) + DP(a, N_\xi(x)) = T \circ M_{L(\xi)} \circ T^{-1}(DP(a,x)).$$

Finally, if DP_Π is the differential of a C^{*-} public key P_Π, then :

$$DP_\Pi(N_\xi(a), x) + DP_\Pi(a, N_\xi(x)) = T_\Pi \circ M_{L(\xi)} \circ T^{-1}(DP(a,x)).$$

Let $\Lambda(L(\xi))$ denote the linear map $T_\Pi \circ M_{L(\xi)} \circ T^{-1}$, then

$$DP_\Pi(N_\xi(a), x) + DP_\Pi(a, N_\xi(x)) = \Lambda(L(\xi))\,(DP(a,x))\,. \tag{2}$$

This last equation is interesting since each coordinate of the left hand side is linear in the unknown coefficients of N_ξ and each coordinate of the right hand side is a linear combination by the unknown coefficients of $\Lambda(L(\xi))$ of the symmetric bilinear coordinate forms of the original DP, which are partially known since their first $(n-r)$ coordinates are public.

The heart of the attack consists in identifying some N_ξ, given the public key and equation (2), and then using its mixing effect on the n coordinates to recover the r missing quadratic forms from the $(n-r)$ known quadratic forms of the public key. In the next section, we will see how to recover some non-trivial multiplication N_ξ, in which ξ can be any value in $\mathbb{F}_{q^n} \setminus \mathbb{F}_q$.

4 Recovering Multiplications from the Public Key

Any linear mapping can be represented by an $n \times n$ matrix with n^2 entries from $\mathbb{F}_q$. Note that the multiplications N_ξ form a tiny subspace of dimension n within the space of all linear maps whose dimension is n^2.

The coordinates of DP_Π are known symmetric bilinear forms that can be seen as $n(n-1)/2$-dimensional vectors. They generate a $(n-r)$-dimensional subspace V_Π which is contained in the n-dimensional space V, generated by the full set of coordinates of DP in the original C^* public key.

Consider now the expression :

$$S_M(a,x) = DP_\Pi(M(a),x) + DP_\Pi(a,M(x))$$

where S_M is defined for any linear mapping M as a $(n-r)$-tuple of symmetric bilinear forms. Most choices of M do not correspond to any multiplication by a large field element ξ, and thus we do not expect them to satisfy the multiplicative property described in section 3. Due to relation (2), when M is a multiplication N_ξ, the $(n-r)$ coordinates of S_{N_ξ} are in V. It is unlikely that they are all in the subspace V_Π. However, there is a huge number of possible values for ξ, and it can be expected that for some choices of $\xi \in \mathbb{F}_{q^n} \setminus \mathbb{F}_q$, some of the bilinear forms in $S_M(a,x)$ will be contained in the known subspace V_Π. Our goal now is to detect such special multiplications.

Dimension of the overall linear maps space. Let us consider k of the published expressions, for instance the first k, and let us study the vector space $E(1,\ldots,k)$ of linear maps M such that the first k coordinates of $S_M(a,x)$ are all contained in V_Π. Since membership in V_Π is expressed by the vanishing of $n(n-1)/2-(n-r)$ linear forms, the elements of this subspace satisfy a system of $k \cdot (n(n-1)/2-(n-r))$ linear equations in the n^2 unknown coefficients of M. If all these equations were independent, the dimension of $E(1,\ldots,k)$ would be $n^2 - k \cdot (n(n-1)/2-(n-r))$ which is clearly impossible as soon as $k \geq 3$. Otherwise, we can only claim that it is lower-bounded by this number. On the other hand, it can be seen that the space $E(1,\ldots,k)$ contains a subspace of multiplications, whose dimension is now to be computed.

Dimension of the multiplications space. For a multiplication N_ξ, thanks to equation (2), the coordinates of S_{N_ξ} are guaranteed to be linear combinations of the coordinates of DP, whose coefficients $\Lambda(L(\xi))$ are linear in $\xi + \xi^{q^\theta}$. Setting $\zeta = \xi + \xi^{q^\theta}$, the first k linear combinations are given by the k linear forms

$$\Lambda_i(\zeta) = \Pi_i \circ T \circ M_\zeta \circ T^{-1}$$

for $i = 1,\ldots,k$ where Π_i is the projection on the ith coordinate. Note that $\Lambda_i : \zeta \mapsto \Lambda_i(\zeta)$ are linear bijections from $\mathbb{F}_{q^n}$ to $(\mathbb{F}_q^n)^*$, the vector space of linear forms over $\mathbb{F}_q^n$. Indeed, the kernel of Λ_i consists of the elements ζ such that the ith row of $T \circ M_\zeta \circ T^{-1}$ is zero. Since $T \circ M_\zeta \circ T^{-1}$ is invertible for $\zeta \neq 0$, the kernel of Λ_i must be trivial. This implies that Λ_i is a linear bijection, and we will use this property. Note that this is the converse of the assumption underlying the attack in [1], and in this sense, our new attack and the old attack can be seen as complementary.

Let us consider the subspace L' of $(\mathbb{F}_q^n)^*$ generated by the first $(n-r)$ coordinate projections. In this case, the k conditions $\Pi_i \circ S_{N_\xi} \in V_\Pi$ become

$$\Lambda_i(L(\xi)) \in L', \ \forall i = 1, \ldots, k \tag{3}$$

which means that $\Lambda_i(L(\xi))$ only depends on the $(n-r)$ first rows of DP, *i.e.* only on the known DP_Π.

Consequently, when searching for a multiplication by ξ for which equation (3) holds, we get the following set of conditions on $\zeta = L(\xi) = \xi + \xi^{q^\theta}$:

(i) $\zeta \in \mathrm{Im}(L)$
(ii) $\Lambda_i(\zeta) \in L'$ for $i = 1, \ldots, k$

Since $\zeta = \xi + \xi^{q^\theta}$ and $\gcd(n, \theta) = 1$, ζ is non-zero unless $\xi = 0$ or 1. This means that the kernel of L has dimension 1, hence ζ ranges over a space of dimension $n-1$. Condition (i) corresponds to a single linear relation over the coordinates of $L(\xi)$, since $\dim \mathrm{Im}(L) = n - 1$. Also, since Λ_i is a linear bijection and L' is of codimension r, each of the conditions in (ii) corresponds to r additional linear relations. Altogether, this means that we have $kr + 1$ linear equations. Furthermore, since we are interested in the space of N_ξ's and not in the space of M_ζ's, the dimension is $n - kr - 1 + 1 = n - kr$ since the kernel of L is of dimension 1. This implies that whenever we add a condition (*i.e.* increase k by 1), we add about $n^2/2$ linear equations on the full space of linear maps, but their effect on the subspace of multiplications is to reduce its dimension only by r. Finally, the space of multiplications in $E(1, \ldots, k)$ includes at least one non-trivial multiplication, *i.e.* a multiplication by an element outside $\mathbb{F}_q$ whenever

$$n \geq kr + 2. \tag{4}$$

Consequently, the dimension of $E(1, \ldots, k)$ is

$$\max\left\{n^2 - k\left(\frac{n(n-1)}{2} - (n-r)\right), n - kr, 1\right\}.$$

Figure 1 describes the expected evolution of the dimension of the space of all linear maps and of the dimension of the subspace of multiplications for two different choices of r. The intuition behind our attack is that initially there are many "useless maps" and few multiplications. However, the number of useless maps drops rapidly as we add more equations, whereas the number of multiplications drops slowly (since many of the equations are linearly related on the subspace of multiplications). This leads to an elimination race, and we hope to

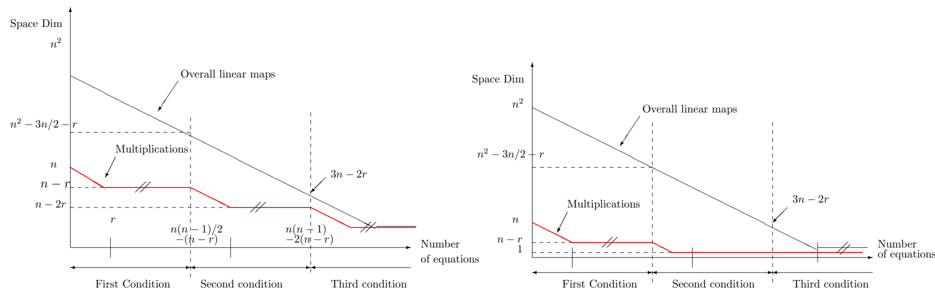

Fig. 1. Evolution of the dimensions of the overall linear maps and their subspace of multiplications when $r < n/3$ (left figure) and when $r \geq n/3$ (right figure), as we add more linear equations.

get rid of all the "bad maps" before we inadvertantly kill off all the "good maps" by imposing too many conditions.

Taking $k = 3$, it can be seen that the first expression of the max is not positive. This seems to indicate that $E(1, \dots, k)$ consists entirely of multiplications. This is demonstrated in the left figure. This subspace contains non-trivial multiplications, whenever $n - 3r > 1$. Therefore, the attack is expected to work for values of r up to $(n-2)/3$. The right figure shows a case in which r is too large, and thus the "good maps" are eliminated before the "bad maps". We will see in section 6 how to improve the attack and deal with values of r up to about $n/2$. Note that even without this improvement, our technique is already sufficient to recover non-trivial multiplications for the recommended parameters of SFLASH^{v2} and SFLASH^{v3}, since $r = 11$ is smaller than both $35/3$ and $65/3$. Of course, the argument that was offered is only heuristic. However, it was confirmed by a large number of experiments, in which the attack always behaved as expected by our heuristic analysis, and signatures were successfully forged.

5 Recovering a Full C^* Public Key

The final part of the attack is to recover a set P'_Π of additional equations which are independent of the first system P_Π. If the rank of the concatenation of the original P_Π and the newly computed r equations of P'_Π is full, then Patarin's attack on MI [8] can be mounted, although we do not necessarily reconstruct the r original equations of the full public key. This idea is the same as in [1].

Recovering a full rank system. To reconstruct a full rank system, we note that the action of the final linear map T is to compute different linear combinations of the full (*i.e.* non-truncated) internal quadratic polynomials $F \circ U$. Consequently, if we were able to mix by some linear mapping the internal quadratic coordinates $F \circ U$ before the action of T_Π, then we will be able to create new quadratic polynomials which could replace the r missing ones.

When we compose the multiplication $N_\xi = U^{-1} \circ M_\xi \circ U$ (which was found in the previous part of the attack) with the truncated public key P_Π, the inputs of the internal quadratic mapping $F(x) = x^{q^\theta+1}$ are multiplied by ξ. Indeed,

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

since $P_\Pi \circ N_\xi(x) = T_\Pi \circ F \circ U \circ U^{-1} \circ M_\xi(U(x)) = T_\Pi(F(M_\xi(U(x))))$. Let us denote this new system by P'_Π. We can show that the outputs of the internal quadratic equations $F \circ U$ are multiplied by $\xi^{q^\theta+1}$. Indeed, $T_\Pi \circ F(\xi \cdot U(x)) = T_\Pi((\xi \cdot U(x))^{q^\theta+1}) = T_\Pi(\xi^{q^\theta+1} \cdot F(U(x)))$, and so :

$$P'_\Pi = P_\Pi \circ N_\xi = T_\Pi \circ M_{\xi^{q^\theta+1}} \circ F \circ U$$

Let us consider the special case $\xi \in \mathbb{F}_{q^n} \setminus \mathbb{F}_q$. In this situation, we say that N_ξ is non-trivial. Since F is a permutation and thus $F(\mathbb{F}_q) = \mathbb{F}_q$, $\xi^{q^\theta+1}$ is not in $\mathbb{F}_q$ either. Thus, the multiplication by $M_{\xi^{q^\theta}+1}$ is non-trivial, *i.e.* corresponds in particular to a non-diagonal matrix.

Therefore, in the sets P_Π and P'_Π the internal quadratic coordinates of $F \circ U$ are mixed with two different linear combinations, T_Π and $T_\Pi \circ M_{\xi^{q^\theta}+1}$. We hope that for some value $\xi \in \mathbb{F}_{q^n} \setminus \mathbb{F}_q$, r equations in the set P'_Π together with P_Π will form a full rank system. This special case is not necessary since we could use different values of ξ to add r different quadratic forms to the $(n-r)$ public ones. However, in our experiments it was always sufficient to use one ξ, and then Patarin's attack could be applied to forge actual signatures.

In practice, to determine if the new system of n equations is of full rank, we simply tested whether Patarin's attack succeeded. If not, another set of r equations was chosen amongst the $(n-r)$ equations of P'_Π. For each choice of r equations, the success probability was approximately $1-1/q$, which is close to 1 for $q = 2^7$

If $\xi \in \mathbb{F}_q$ (*i.e.* the multiplication is trivial), P'_Π is simply P_Π where each coordinate has been multiplied by the same element of $\mathbb{F}_q$, since $F(\mathbb{F}_q) = \mathbb{F}_q$ and multiplication by an element of $\mathbb{F}_q$ is a diagonal matrix. Thus, such trivial ξ are not interesting for our attack and this is the reason why they were discarded from our search for appropriate N_ξ in the previous section.

Practical results. We carried our experiments on a 2GHz AMD Opteron PC using different parameters. The following table provides the time to recover a non-trivial multiplication and the time to recover an independent set of equations which form a full rank system. This computation has to be done only once per public key. Then Patarin's attack requires about one second to forge an actual signature for any given message. All these operations can be carried out by solving various systems of linear equations with a relatively small number of variables ($O(n^2)$ or $O(n)$, depending on the operation).

The two columns in bold font represent the time to attack SFLASH^{v2} and SFLASH^{v3}. The notation 's' is for seconds and 'm' is for minutes.

n	37	**37**	67	**67**	131
θ	11	**11**	33	**33**	33
q	2	**128**	2	**128**	2
r	11	**11**	11	**11**	11
N_ξ Recovery	4s	**70s**	1m	**50m**	35m
C^* Recovery	7.5s	**22s**	2m	**10m**	7m
Forgery	0.01s	**0.5s**	0.02s	**2s**	0.1s

6 Breaking SFLASH when the Number of Deleted Quadratic Equations r is up to $n/2$

In this section, we deal with this problem by a technique which we call *distillation*, since it allows to gradually filter additional linear maps which are not multiplications. When $r \leq (n-2)/3$, we can use three conditions to eliminate all the useless linear maps, while retaining at least a two dimensional subspace of multiplications (since we reduce the initial n coordinates three times by r). When $r > (n-2)/3$, this will usually kill all the multiplications along with the useless linear maps.

Distillation is performed by relaxing the constraints, *i.e.* by forcing only two coordinates of S_M to be in V_Π. This will cancel a large fraction of useless linear maps, but not all of them. To clarify the situation, we use in the rest of this section angular brackets to demonstrate the stated number of dimensions for the SFLASH^{v3} parameters of $n = 67$ and $r = 11$.

After forcing the two conditions, the dimension of the space of linear maps is reduced to

$$n^2 - 2(n(n-1)/2 - (n-r)) = 3n - 2r \ \langle 179 \rangle$$

of the n^2 $\langle 4489 \rangle$ at the beginning, while the dimension of the good subspace (*i.e.* the subspace of multipications) is $n - 2r$ $\langle 45 \rangle$. Now, to find at least one non-trivial multiplication, we need to eliminate all the remaining useless linear maps. The new idea is that we can perform this process twice with different pairs of coordinates, *i.e.* coordinates 1 and 2 for the first time and coordinates 3 and 4 for the second, and get two different sets of linear maps, say VS_1 and VS_2, which contain both good and bad linear maps. Two random linear subspaces of dimension m in a linear space of dimension t are likely to have a nonzero intersection if and only if $m > t/2$, and then the dimension of the intersection is expected to be $2m - t$. We can apply this criterion separately to the space of all linear maps (in which $t = n^2$) and to the subspace of multiplications (in which $t = n$). In our example $VS_1 \cap VS_2$ is likely to contain non-trivial multiplications since $\langle 45 \rangle > \langle 67 \rangle/2$, but is not likely to contain other maps since $\langle 179 \rangle < \langle 4489 \rangle/2$. More generally, we may have to replace each one of SV_1 and SV_2 by the sum of several such linear subspaces in order to build up the dimension of the multiplications to more than $n/2$. For example, if each VS_i has only a $\langle 10 \rangle$-dimensional subspace of multiplications, we can replace it by the sum of four such linear subspaces to get the expected dimension up to $\langle 40\rangle$, and the intersection of two such sums will have an expected dimension of $\langle 13\rangle$, and thus many non-trivial multiplications.

Asymptotic Analysis. We now show how to deal with any $r < (1-\varepsilon)n/2$ for a fixed ε and large enough n. Note that our goal here is to simplify the description, rather than to provide the most efficient construction or tightest analysis. Since $n - 2r > \varepsilon n$, we can impose pairs of conditions and create linear subspaces VS_i of total dimension $O(n)$ which contain a subspace of multiplications of dimension $\varepsilon n \geq 2$. If we add $1/\varepsilon$ such subspaces, the dimension of the subspace of multiplications will increase to almost n, while the total dimension will remain n/ε, which is much smaller than n^2. Consequently, the intersection of two such sums is likely to consist entirely of multiplications.

Experimentations. We get the following timing results when r is close to $n/2$ and 's', 'm' and 'h' respectively denotes seconds, minutes and hours.

n	37	37	67	67
θ	11	11	33	33
q	2	128	2	128
r	17	16	32	31
N_ξ Recovery	$8s$	$4m$	$3.5m$	$10h$
C^* Recovery	$7.5s$	$22s$	$3m$	$10m$
Forgery	$0.01s$	$0.4s$	$0.02s$	$2s$

7 Comparison with the Method of Dubois *et al.* [1]

In both attacks, the basic strategy is to recover additional independent equations in order to apply Patarin's attack [8]. They both use the differential of the public key, but differ in the way the invariants of the differential are found. The method of [1] can only deal with schemes where $\gcd(n, \theta) > 1$, which implies that the kernel of $L(\xi) = \xi + \xi^{q^\theta}$ is of dimension strictly larger than 1.

To recover non-trivial multiplication in [1], skew-symmetric mappings with respect to a bilinear form B are considered, *i.e.* linear maps M such that $B(M(a), x) = -B(a, M(x))$. In fact, the authors show that skew-symmetric mappings related to the symmetric bilinear forms of a C^* public key are specific multiplications in the extension $\mathbb{F}_{q^n}$ by means of a suitable transformation depending on the secret key, namely $U^{-1} \circ M_\xi \circ U$ where $\xi \in \operatorname{Ker} L$. For such maps, we get $DP(M(a), x) + DP(a, M(x)) = 0$. Since DP can be computed from the public key, this equation defines linear equations in the unknowns of M. However, in the case considered in this paper, *i.e.* when $\dim \operatorname{Ker} L = 1$ or equivalently when $\gcd(n, \theta) = 1$, the only skew-symmetric maps are the trivial multiplications which are useless to recover new independent quadratic equations.

To recover non-trivial multiplications, we introduce here different and more elaborate conditions related to the vector space generated by the various images of the differential in public key coordinates. In this case, we are also able to detect images of multiplications. However, the multiplications to be found are not known in advance but are only shown to exist by counting arguments, and the way we find them is by setting up an elimination race between the multiplications and other linear maps.

8 Conclusion

Multivariate cryptographic schemes are very efficient but have a lot of exploitable mathematical structure. Their security is not fully understood, and new attacks against them are found on a regular basis. It would thus be prudent not to use them in any security-critical applications.

One of the most interesting open problems is whether the new techniques described in this paper can be applied to the HFE cryptosystem [9]. The main attacks discovered so far against HFE are based on Gröbner bases [3], and are very slow. So far, we could not find a way how to detect non-trivial multiplications in HFE, since it lacks the multiplicative property described in section 3, but this is a very promising line of attack which should be pursued further.

Acknowledgements

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. V. Dubois, P. A. Fouque, and J. Stern. Cryptanalysis of SFLASH with Slightly Modified Parameters. In *Eurocrypt '07*, LNCS 4515, pages 264–275. Springer-Verlag, 2007.
2. T. El Gamal. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. In *IEEE Transactions on Information Theory*, volume IT–31, no. 4, pages 469–472, july 1985.
3. J.-C. Faugère and A. Joux. Algebraic Cryptanalysis of Hidden Field Equation (HFE) Cryptosystems Using Gröbner Bases. In *Crypto '03*, LNCS 2729, pages 44–60. Springer-Verlag, 2003.
4. P. A. Fouque, L. Granboulan, and J. Stern. Differential Cryptanalysis for Multivariate Schemes. In *Eurocrypt '05*, LNCS 3494, pages 341–353. Springer-Verlag, 2005.
5. M. R. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of* NP-*Completeness*. Freeman, New-York, 1979.
6. H. Gilbert and M. Minier. Cryptanalysis of SFLASH. In *Eurocrypt '02*, LNCS 2332, pages 288–298. Springer-Verlag, 2002.

7. 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.
8. 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.
9. J. Patarin. Hidden field equations (HFE) and isomorphisms of polynomials (IP): two new families of asymmetric algorithms. In *Eurocrypt' 96*, LNCS 1070, pages 33–48. Springer-Verlag, 1996.
10. 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.
11. 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.
12. R.L. Rivest, A. Shamir, and L.M. Adleman. A method for obtaining digital signatures and public-key cryptosystem. *Communications of the ACM*, 21(2):120–126, 1978.
13. A. Shamir. Efficient Signature Schemes Based on Birational Permutations. In *Crypto '93*, LNCS 773, pages 1–12. Springer-Verlag, 1993.