

Extension Field Cancellation: A New Central Trapdoor for Multivariate Quadratic Systems

Alan Szepieniec^{1,2(✉)}, Jintai Ding³, and Bart Preneel^{1,2}

¹ Department of Electrical Engineering, ESAT/COSIC, KU Leuven, Leuven, Belgium
`alan.szepieniec@esat.kuleuven.be`

² iMinds, Ghent, Belgium

³ University of Cincinnati, Cincinnati, OH, USA

Abstract. This paper introduces a new central trapdoor for multivariate quadratic (MQ) public-key cryptosystems that allows for encryption, in contrast to time-tested MQ primitives such as Unbalanced Oil and Vinegar or Hidden Field Equations which only allow for signatures. Our construction is a mixed-field scheme that exploits the commutativity of the extension field to dramatically reduce the complexity of the extension field polynomial implicitly present in the public key. However, this reduction can only be performed by the user who knows concise descriptions of two simple polynomials, which constitute the private key. After applying this transformation, the plaintext can be recovered by solving a linear system. We use the minus and projection modifiers to inoculate our scheme against known attacks. A straightforward C++ implementation confirms the efficient operation of the public key algorithms.

Keywords: MQ · Multivariate · Quadratic · Public-key · Post-quantum · Encryption · Mixed-field · Trapdoor

1 Introduction

Since the inception of public-key cryptography, cryptographers have made a huge effort to find new and better computational problems that feature the elusive *trapdoor* — a small piece of information that can turn an otherwise hard to invert function into one that can easily be inverted. This on-going search effort has led to a tremendous diversification of the computational problems that underpin public-key cryptography. This diversification is a good thing: by keeping all the eggs in separate baskets, a breakthrough in one area is unlikely to spill over to other areas, thus limiting the catastrophic potential of scientific advances.

Of particular interest to this paper is the class of problems known as multivariate quadratic (MQ) systems of equations. Not only do cryptosystems based on this primitive offer performance advantages over well-established ones such as RSA or systems based on elliptic curves, MQ cryptography is also conjectured to be post-quantum — that is to say, it holds promise of resisting attacks on

quantum computers. From this point of view, MQ cryptography is certainly a promising line of research.

The key challenge in the design of MQ cryptosystems is to find a suitable central mapping $\mathcal{F} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ which should be easily invertible in addition to being expressible in terms of multivariate quadratic polynomials. The trapdoor information cannot be recovered efficiently from the public key as it is hidden by two affine transformations. Many central mappings have been proposed, most of which fall in two main categories [32]: single field schemes, such as UOV [17], Rainbow [7] and the triangular variants [31], where the central polynomial system is chosen to have a particular structure that enables efficient inversion; and mixed field schemes, such as C* [19], HFE [22] and Multi-HFE [3], where arithmetic in the base field is mixed with arithmetic in an extension field. However, despite the abundance of proposals, MQ cryptography has an awful track record as most of these proposals have been broken [2, 14, 18, 28, 29, 32].

Consequently, much research in the area of MQ cryptography has been devoted to patchwork — finding small modifications to existing systems that render specific attacks infeasible. A few examples among many that fall into this category are the minus modifier (“−”) [25], which inoculates HFE-type systems against Gröbner basis attacks and linearization attacks; vinegar variables (“v”) [17], which combines elements from different trapdoors and like “minus” is capable of making a Gröbner basis attack prohibitively expensive; and projection (“p”) [9] which appears to successfully thwart the Dubois *et al.* differential attack [10, 11] on SFLASH.

However, the search for modifications to fix broken systems has an equally bad track record. Many of the MQ systems that were supposedly inoculated against some attack by the introduction of a modification, were broken by minor variants of that same attack. For example, both the multivariate generalization and the odd field characteristic variant of HFE were introduced and designed specifically to thwart the algebraic attack on HFE [14]; however, neither variant has managed to withstand cryptanalysis [2]. Another example is given by the fate of SFLASH, one of the three recommended signature schemes of the NESSIE project [1]. The addition of the minus modifier to the basic C* construction did not save the scheme from a new type of differential attack [10, 11]. The rapid spawn of attacks that break the inoculated systems seems to suggest the need for a more prudent design strategy: searching for fundamentally different basic principles for MQ trapdoors, rather than tinkering on the edges of existing ones.

Related work. Encryption schemes have been the bane of multivariate quadratic cryptography. No MQ encryption scheme has withstood the test of time, while several MQ *signature* schemes have. However, some very recent results and proposals in this area pose new and interesting challenges for cryptanalysts.

Porras *et al.* proposed a new central trapdoor which they call ZHFE [24]. Up until this point, the extension field polynomial in HFE-based cryptosystem required the number of nonzero coefficients to be small and its degree to be relatively low, so as to allow efficient root calculation. The idea of Porras *et al.* exchanges this single low-degree polynomial for a pair of high-degree polynomials

that make up the central map. Additionally, these polynomials are chosen such that there exists a third polynomial, $\Psi(\mathcal{X})$, which is a function of the first two and yet has low degree. In order to invert a given image, it suffices to factorize this third polynomial. As the degree of the polynomials increases, so does the degree of regularity of the system. This increase in the degree of regularity, in turn, renders a direct algebraic attack infeasible, even though the very same attack broke the regular HFE cryptosystem.

Tao *et al.* proposed a multivariate quadratic encryption scheme called Simple Matrix Encryption, or simply ABC Encryption [27]. Their construction is based on a fundamentally new idea: embedding polynomial matrix arithmetic inside the central trapdoor function. The trapdoor can be inverted with high probability because the matrix, albeit evaluated in a single point, can be reconstructed from the output. With high probability this matrix can be inverted, giving rise to a system of linear equations which describe the input.

Our contributions. We introduce a new central trapdoor for multivariate quadratic encryption schemes. Our proposal is a mixed-field scheme — similar to the C* and HFE string of proposals because we use an embedding function to pretend as though a vector of variables in the base field were actually a single variable in the extension field. However, our proposal is notably different from its predecessors, where the restriction on the degree of this embedded polynomial was key both to their efficiency and to their demise; our proposal allows for a high-degree embedded polynomial and undoes this complexity by exploiting the commutative property of the extension field. Our proposal allows for encryption, in stark contrast to most other members of the HFE family.

Like the ABC Encryption Scheme, decryption of a ciphertext consists of essentially solving linear systems. This linear system is parameterized by the particular ciphertext or message: every possible ciphertext or message implicitly defines a unique linear system. Knowledge of the private key allows the user to obtain the linear system efficiently, while the adversary who attacks the system without this crucial information has no advantage to solve the quadratic system.

Like ZHFE, the central map consists of two high-degree extension field polynomials that satisfy a special relation which is obviously hidden from the adversary. The decryption algorithm exploits this relation to turn the otherwise hard inversion problem into an easy one.

Another important similarity between our map and both ABC and ZHFE is that all three are expanding maps, *i.e.*, $\mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ where $m = 2n$. This commonality is no accident, because in order to allow unique decryption, the map must be injective. However, if $m \approx n$, the differential of this nearly-bijective map is readily differentiable from that of a random one — not a desirable property for multivariate quadratic maps to have.

Despite these similarities, the main advantage of our scheme is that its construction is notably *different* from ABC and ZHFE. Consequently, as-yet undiscovered weaknesses or even attacks that affect ABC or ZHFE may leave our scheme intact. Furthermore, this diversification opens the door for a combination of strategies whose end result reaps the benefits of both worlds. Certainly

the case of HFEv proves that such a combination may indeed increase both security and performance.

In line with a common theme throughout MQ cryptography, we are unable to prove the security of our scheme or even to reduce it to a plausible computational assumption. An exhaustive list of all known attacks on MQ systems and why they fail against our system is beyond the scope of this paper. Nevertheless, we identify several pertinent attacks that may be launched against a naïve implementation of our scheme, and we propose strategies to thwart them. Patarin’s linearization attack [21] is foiled by the minus modifier and repeated applications of the same modifier make the extended MinRank attack [4, 18] as well as the direct algebraic attack [14] prohibitively inefficient. The scheme seems naturally resistant to Dubois *et al.*’s differential attack [10, 11], but we nevertheless recommend to use the projection modifier, which is the proper countermeasure against this attack.

Outline. We introduce notation and recall basic properties of MQ systems as well as of extension field embeddings in Sect. 2. Next, Sect. 3 defines the trapdoor proposed in this paper as well as several necessary modifiers. We recommend parameters for 80 bits of security in the first part of Sect. 4 and afterwards discuss the efficiency of our scheme, both from a theoretical point of view and by referencing timing results from a software implementation. Section 5 concludes the text.

2 Preliminaries

2.1 Notation and Definitions

We use small case letters (s) to denote scalars in the base field; extension field elements are denoted by calligraphic capital letters (\mathcal{C}); small case bold letters (\mathbf{v}) denote column vectors; and regular capital letters are used for matrices (M).

Let \mathbb{F}_q denote the finite field with q elements, which we call the *base field*. With any combination of a finite field \mathbb{F}_q with a polynomial $f(x) \in \mathbb{F}_q[x]$ one can associate a finite ring $\mathbb{E} = \mathbb{F}_q[x]/\langle f(x) \rangle$ of residue classes after division by $f(x)$. If f is irreducible over \mathbb{F}_q and has degree n , then $\mathbb{E} = \mathbb{F}_{q^n}$ is a finite field we call the *extension field*. There exists a natural homomorphism $\varphi: (\mathbb{F}_q)^n \rightarrow \mathbb{F}_{q^n}$ that maps a vector $\mathbf{v} = (v_1, \dots, v_n)^\top \in \mathbb{F}_q^n$ onto an element $v \in \mathbb{F}_{q^n}$ of the extension field. We can apply this embedding function to the vector of indeterminates \mathbf{x} in order to get the extension field indeterminate $\mathcal{X} = \varphi(\mathbf{x})$.

2.2 Multivariate Quadratic Systems

The public key of an MQ cryptosystem is a system of quadratic polynomials mapping n input variables to m output variables: $\mathcal{P}: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$; the public operation consists of evaluating this system of polynomials in a point. The secret key consists of a pair of invertible affine mappings on the input and output

variables, S and T , and an alternate quadratic system of polynomials, $\mathcal{F} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$, such that $\mathcal{P} = T \circ \mathcal{F} \circ S$. The affine transformations are trivially inverted; the central system \mathcal{F} is constructed in such a way that it is also easy to invert. However, the attacker cannot efficiently recover \mathcal{F} from \mathcal{P} and calculate the inverse as \mathcal{F} is hidden by the affine transformations. A schematic overview is given in Fig. 1.

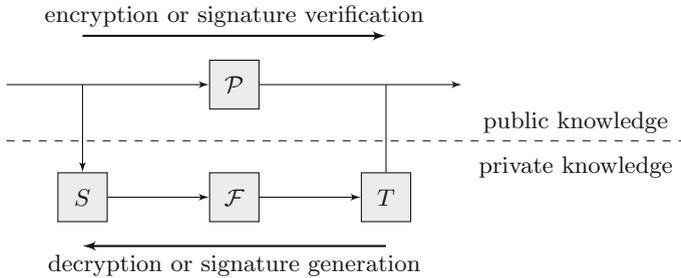


Fig. 1. Schematic representation of multivariate quadratic cryptosystems.

Given a central trapdoor \mathcal{F} it is easy to construct a multivariate quadratic cryptosystem by composing it with two affine transformations. This process is out of the scope of the present paper. Rather, we restrict our attention to the construction of the central trapdoors.

3 Central Map

3.1 The Basic Construction

Let $A \in \mathbb{F}_q^{n \times n}$ be a random matrix over the base field. Then $A\mathbf{x} \in (\mathbb{F}_q[\mathbf{x}])^n$ represents a vector where each element is a linear polynomial in \mathbf{x} . And then $\alpha(\mathbf{x}) = \varphi(A\mathbf{x})$ is an extension field element. The square matrix that represents multiplication by $\alpha(\mathbf{x})$ is denoted by $\alpha_m(\mathbf{x}) \in \mathbb{F}_q^{n \times n}$. We use $\alpha(\mathcal{X})$ to stress the fact that α may also be considered as a univariate polynomial in \mathcal{X} over the extension field, regardless of its representation, although the degree of this polynomial is larger than one.

Similarly, let $\beta(\mathbf{x}) = \varphi(B\mathbf{x})$ for a random $n \times n$ matrix $B \in \mathbb{F}_q^{n \times n}$. With these polynomials α and β , we define the central trapdoor as follows:

$$\mathcal{F} : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^{2n} : \mathbf{x} \mapsto \begin{pmatrix} \alpha_m(\mathbf{x})\mathbf{x} \\ \beta_m(\mathbf{x})\mathbf{x} \end{pmatrix} . \tag{1}$$

To see how we are able to invert $\mathcal{F}(\mathbf{x}) = \begin{pmatrix} \mathbf{d}_1 \\ \mathbf{d}_2 \end{pmatrix}$, consider first the equality $\alpha(\mathbf{x})\beta(\mathbf{x}) = \beta(\mathbf{x})\alpha(\mathbf{x})$ which holds due to the commutativity of the extension field. We can proceed to construct a system of linear equations in \mathbf{x} :

$$\beta_m(\mathbf{x})\mathbf{d}_1 - \alpha_m(\mathbf{x})\mathbf{d}_2 = 0 . \tag{2}$$

While Gaussian elimination is in this case guaranteed to find a solution, this solution need not be unique. Nevertheless, this set of solutions is expected to be small, in accordance with the number of solutions to random linear systems. Moreover, this set can be pruned by iteratively plugging the potential solution into the function \mathcal{F} and verifying that the correct output image $(\mathbf{d}_1; \mathbf{d}_2)$ is produced.

3.2 Modifiers

The trapdoor as described above is insecure. In particular, it is broken by the bilinear attack, the MinRank attack, as well as an algebraic attack using fast Gröbner basis algorithms. We apply the “minus” to inoculate basic EFC against these attacks. While not strictly necessary, “projection” may guard against new differential attacks at very little cost whereas “Frobenius tail” drastically drops the cost of decryption.

Minus. Although Patarin’s linearization attack [21] was originally conceived to attack C^* , it also applies to unprotected EFC. Indeed, Eq. 2 describes a bilinear polynomial in the plaintext and ciphertext, whose coefficients can be calculated using linear algebra after obtaining enough plaintext-ciphertext pairs. Once these coefficients are known, obtaining a plaintext that matches a given ciphertext is easy. However, dropping just one polynomial from the public key is enough to foil this attack. In this case, the attacker must guess the missing information for every plaintext-ciphertext pair, making them useless for exact linear algebra.

This “minus” modifier, which consists of removing one or more polynomials from the public key [23], is more than just a countermeasure against Patarin’s attack. A pair of important results by Ding *et al.* [6, 8] indicates that this modifier is much better thought of as a fundamental building block of multivariate quadratic cryptosystems rather than a mere patch. Indeed, not only does the first application of this modifier block Patarin’s linearization attack; every repeated application increments by one the rank of the quadratic form associated with the extension field polynomial, rendering the MinRank attack due to Kipnis and Shamir [18] as well as its subsequent improvement by Courtois [4] that much more infeasible. Furthermore, this rank increase in turn increases the degree of regularity of the system, resulting in a similarly infeasible algebraic attack.

The use of this modifier does come at the cost of a performance penalty. In particular, the decryption algorithm must first guess the values of the missing polynomials before undoing the output transformation T . Under this guess, it can proceed to the linear system in Eq. 2 and compute the potential matching plaintext \mathbf{x} . If indeed $\mathcal{F}(\mathbf{x}) = (\mathbf{d}_1; \mathbf{d}_2)$, then the correct plaintext was found. If not, then the guess was wrong and the algorithm must start all over again with a new one.

Fortunately, as long as the number of dropped polynomials a is small enough, the correct plaintext will still be found with overwhelming probability. In order for the decryption algorithm to produce the wrong plaintext \mathbf{x} upon decrypting the ciphertext \mathbf{y} , there must exist at least two guesses $\mathbf{g}_1 \in \mathbb{F}_q^a$ and $\mathbf{g}_2 \in \mathbb{F}_q^a$ such

that both $(\mathbf{y}; \mathbf{g}_1)$ and $(\mathbf{y}; \mathbf{g}_2)$ are in the range of \mathcal{P} . If \mathcal{P} is to be modeled as a random function $\mathbb{F}_q^n \rightarrow \mathbb{F}_q^{2n-a}$, then its range is a uniform subset of \mathbb{F}_q^{2n-a} of size q^n , and then the probability of this event is approximately $q^n \times q^{-2n+a} = q^{-n+a}$. Consequently, as long as $a \ll n$, the probability of decryption error remains astronomically small.

Figure 2 offers empirical validation of this argument. It shows the probability of decryption error for various even values for a as a function of n . Only when a and n are on the same order of magnitude, is this probability noticeable; when n rises to practical values, this probability does indeed drop to zero.

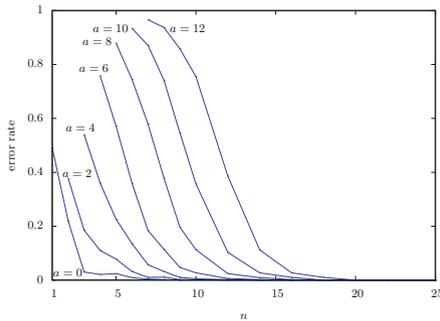


Fig. 2. Observed decryption error rate.

In similar fashion to C^{*-} and HFE^- , this modifier will be denoted by the superscript “-”, *i.e.*, EFC^- . The number of dropped polynomials will be denoted by a .

Projection. The differential symmetry attacks by Dubois *et al.* [10,11] on SFLASH, a C^* variant, show that the minus operator is not enough to secure it. Dubois *et al.* identify a symmetry in the differential of the C^* map \mathcal{F} :

$$DF(L\mathbf{x}, \mathbf{y}) + DF(\mathbf{x}, L\mathbf{y}) = \Lambda\mathcal{F}(\mathbf{x}, \mathbf{y})$$

for some matrices L and Λ . The presence of this symmetry proved fatal.

Fortunately, Ding *et al.* [9] show experimentally that a small tweak by the name of “projection” completely foils this line of attack. In particular, pSFLASH projects the input vector \mathbf{x} onto a lower-dimensional space before passing it through the central map. Smith-Tone [26] has since offered a theoretical basis for the efficacy of this modifier. At the core of Smith-Tone’s argument is the following theorem:

Theorem 1 (Smith-Tone, [26]). *A polynomial $f : \mathbb{F}_{q^n} \rightarrow \mathbb{F}_{q^n}$ with a bilinear differential has the multiplicative symmetry if and only if it has one quadratic monomial summand.*

While the components of EFC do have bilinear differentials, they do not consist of a single quadratic monomial but of a sum of them. For example, the first component is described by $\alpha(\mathcal{X})\mathcal{X} = \sum_{i=0}^{n-1} \mathcal{A}_i \mathcal{X}^{q^i+1}$ where the coefficients \mathcal{A}_i are with overwhelming probability not all but one equal to zero. Therefore, by Smith-Tone’s theorem, the differential multiplicative symmetry is absent with overwhelming probability.

Nevertheless, in anticipation of more general attacks using a similar differential invariant, we follow a perspective offered at the conclusion Smith-Tone’s paper: *projection does not destroy the differential symmetry, but pushes it down to a subfield*. Since this modifier is cheap in terms of performance and cannot degrade security, we choose to err on the side of safety and ensure that no such subfield can exist. In particular, we guarantee that the matrices A and B have rank $n - 1$, and that n is a prime number. Moreover, the kernels of A and B do not intersect except at the origin. This modifier will be denoted by the subscript p , e.g. EFC_p .

Frobenius Tail in Characteristic Two (or Three). The trapdoor as described so far can be implemented over any base field and unless the minus operator is applied, the rank of the quadratic forms associated with the extension field is two. However, if we restrict to characteristic two, we can naturally increase this rank by adding an extra “tail” term to both expressions. In turn, we must drop fewer equations to ensure the same level of security, and this results in a significant speedup of the decryption algorithm. We will use the subscript t^2 to denote the use of this technique, e.g. EFC_{t^2} .

This trick exploits the following property of fields of characteristic two. Let $f(\mathcal{X})$ be a linear function, then $f(\mathcal{X})^3$ is a quadratic function and multiplication by $f(\mathcal{X})$ gives $f(\mathcal{X})^4$ which is once again a linear function.

Let α and β be defined as earlier. Then this enhancement adds the quadratic terms $\alpha(\mathcal{X})^3$ and $\beta(\mathcal{X})^3$ as follows:

$$\mathcal{F} : \mathbb{F}_{2^n} \rightarrow \mathbb{F}_{2^n} : \mathcal{X} \mapsto \begin{pmatrix} \alpha(\mathcal{X})\mathcal{X} + \beta(\mathcal{X})^3 \\ \beta(\mathcal{X})\mathcal{X} + \alpha(\mathcal{X})^3 \end{pmatrix}. \tag{3}$$

In order to decrypt $\mathcal{F}(\mathcal{X}) = (\mathcal{D}_1; \mathcal{D}_2)$, the user solves the linear system

$$\alpha(\mathcal{X})\mathcal{D}_2 - \beta(\mathcal{X})\mathcal{D}_1 = \alpha(\mathcal{X})^4 - \beta(\mathcal{X})^4. \tag{4}$$

Afterwards, the set of solutions is pruned based on $\mathcal{F}(\mathcal{X}) = (\mathcal{D}_1; \mathcal{D}_2)$.

A similar trick is possible in fields of characteristic three. For linear functions $f(\mathcal{X})$ the term $f(\mathcal{X})^2$ is quadratic and multiplication by $f(\mathcal{X})$ gives $f(\mathcal{X})^3$ which is once again a linear function. Although this particular Frobenius tail does destroy the common factor in the two polynomials, it merely increases the rank of the quadratic form to three. The use of this trick will be denoted by the subscript t^3 .

4 Efficiency

4.1 Recommended Parameters

We predict that the most efficient attack on our system is the algebraic attack using efficient Gröbner basis algorithms such as Faugère's F_4 or F_5 [12, 13]. Taking this attack into account, we propose parameters to ensure at least 80 bits of security.

We follow the argument due to Ding *et al.* [5, 8], who develop an upper bound for the degree of regularity of HFE⁻ systems. In this line of reasoning, the degree of regularity D_{reg} is intricately linked to the rank r of the quadratic form associated with the extension field polynomial. Moreover, a applications of the minus modifier effectively increases this rank by a . Especially for small base fields, the degree of regularity is expected to lie near its upper bound:

$$D_{\text{reg}} \leq \frac{(q-1)(r+a)}{2} + 2. \quad (5)$$

This argument applies to a single quadratic form. However, the central map of EFC consists of *two* quadratic forms. Nevertheless, we argue that the effect of minus is replicated across both quadratic forms. The polynomials are dropped *after* the output transformation T is applied, meaning that the effect of the missing information passes through T^{-1} and is not isolated to one quadratic form but spread across both. Although this reasoning underscores the following parameter recommendations, we note it is not perfectly rigorous and warrants further study.

Considering the two components of our central map separately, we see that their rank is $r = 2$. If the Frobenius tail modifiers are applied, this is increased to $r = 4$ and $r = 3$ for characteristics 2 and 3, respectively. For a security level of 80 bits, we recommend to ensure this adjusted rank is at least 12 for \mathbb{F}_2 and 8 for \mathbb{F}_3 .

$$a = \begin{cases} 10 & q = 2, n = 83, \text{EFC}_p^- \\ 8 & q = 2, n = 83, \text{EFC}_{pt^2}^- \\ 6 & q = 3, n = 59, \text{EFC}_p^- \end{cases}. \quad (6)$$

Then we can estimate the degrees of regularity for these base fields:

$$D_{\text{reg}} \leq \frac{(q-1)(r+a)}{2} + 2 = \begin{cases} 8 & q = 2 \\ 10 & q = 3 \end{cases}. \quad (7)$$

The running time of efficient Gröbner basis algorithms is dominated by Gaussian elimination in the matrix of coefficients associated with the monomials of degree D_{reg} . We can use this bottleneck to estimate the algorithm's total complexity. In particular, the number of monomials of this degree is given by $T = \binom{n}{D_{\text{reg}}} \approx 2^{35}$ both for $n = 83, q = 2$ as well as $n = 59, q = 3$. Moreover, the number of nonzero monomials is on the order of $\tau = \binom{n}{2} \geq 2^{10}$. Assuming a Wiedemann-type algorithm [30] for sparse Gaussian elimination, this amounts to $\tau T^2 \geq 2^{80}$ in both cases.

Figure 3 offers some experimental evidence in support of this argument. It plots the running time of MAGMA’s F_4 algorithm to recover the plaintext from the ciphertext and the public key. The graph on the left starts out with $q = 2$, $n = 35$ and $a = 1$; from there on out, the parameter a increases. The graph on the right lets n vary from 15 to 38 with $q = 2$, and keeps a constant at 10 for the basic trapdoor EFC_p^- (blue circles) and at 8 for the Frobenius tail equivalent $EFC_{pt^2}^-$ (red crosses).

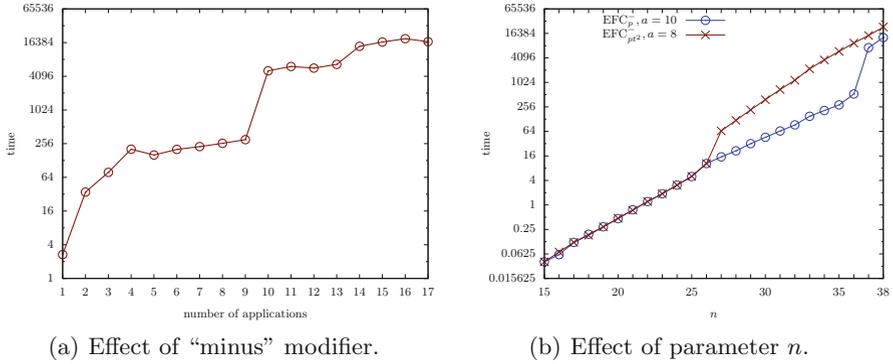


Fig. 3. Running time of algebraic attack for various parameters (Color figure online).

The graphs indicate two things. First, the minus modifier enhances security with (nearly) every application, occasionally lifting the system into the next degree of regularity. Second, the Frobenius tail modifier enhances security, even compensating for the rank drop associated with going from $a = 10$ to $a = 8$.

4.2 Complexity

The basic trapdoor, as well as all the modified variants, feature only quadratic terms. Therefore, the transformations T and S should be linear and not affine, and consequently also the public key will consist of only quadratic terms.

The public key consists of $2n - a$ polynomials of degree 2 in n variables. Thus the number of coefficients from \mathbb{F}_q in the public key is $(2n - a) \times \frac{n(n-1)}{2} = n^3 - (a + 1)n^2 + an = O(n^3)$ because $a \ll n$. However, we note that there is a considerable amount of redundancy in the public key which we expect can be exploited to produce smaller keys.

The private key consists of two linear transformations S and T , along with a degree- n irreducible polynomial $\psi(z)$, and matrices A and B . This amounts to $n^2 + (2n)^2 + 2(n^2) + n = 7n^2 + n = O(n^2)$ coefficients in \mathbb{F}_q .

The most computationally intensive part of the key generation algorithm is the symbolic matrix-vector multiplication — once in $\varphi(A\mathbf{x})\mathbf{x}$ and once in $\varphi(B\mathbf{x})\mathbf{x}$. Both procedures require n^2 polynomial-multiplications, each of which

consists of n multiplications in \mathbb{F}_q . Since the other steps in the key generation algorithm are less complex, the asymptotic time complexity of this entire algorithm is $O(n^3)$. For the Frobenius tail modifier, this complexity is worse because the additional extension field products $\varphi(A\mathbf{x})(Q\mathbf{A}\mathbf{x})$ and $\varphi(B\mathbf{x})(Q\mathbf{B}\mathbf{x})$ (where Q is the matrix associated with the Frobenius map $x \mapsto x^2$) have dense right-side multiplicands. Consequently, the cost of polynomial multiplication rises to n^2 multiplications and the total time complexity of the key generation to $O(n^4)$.

Encryption consists of evaluating $2n - a$ quadratic polynomials in n variables. This comes down to two time steps with unlimited parallelism. Without parallelism, however, each of the $(2n - a) \times (n(n - 1) + 2n)$ base field operations must be executed sequentially and the time complexity is therefore $O(n^3)$.

Decryption consists of the following steps for q^a different guesses, which may be executed in parallel if the resources are available: (1) inversion of T , which requires $(2n)^2$ operations; (2) computation of $\varphi(\mathbf{d}_1)$ and $\varphi(\mathbf{d}_2)$, which requires n vectorized additions for a total of n^2 operations; (3) two matrix multiplications of n^3 operations each, followed by a matrix subtraction; (4) a Gaussian elimination of some $2n^3/3$ operations; (5) inversion of S requiring some n^2 operations; and finally (6) pruning, which has an almost constant expected running time. Thus, decryption has an expected running time of $O(q^a n^3)$. While this expression does involve an exponential factor, the exponent is rather small — on the order of $a \approx \log n$, so that decryption is still practically speaking a polynomial-time algorithm.

Figure 4 emphasizes this exponential behavior by logarithmically plotting the decryption time as a function of a . Even a moderate increase in the number of dropped parameters can make decryption impractically slow.

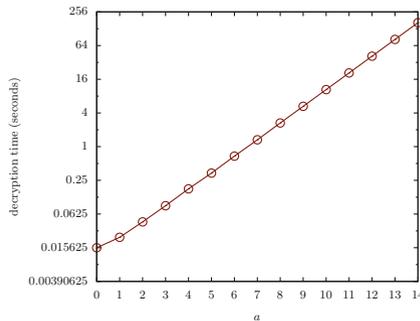


Fig. 4. Decryption time as a function of a for $n = 83$ and $q = 2$.

4.3 Speed

Table 1 shows some timing results obtained from a straightforward C++ implementation on a 64-bit 3.3 GHz Intel CPU. Despite the scheme's obvious capacity

Table 1. Implementation results — timings of key generation, encryption and decryption algorithms along with public key, secret key and ciphertext size.

Construction	Sec. key	Pub. key	Ctxt	Key gen	Enc	Dec
$\text{EFC}_p^-, q = 2, n = 83, a = 10$	48.3 KB	509 KB	20 B	2.45 s	0.004 s	9.074 s
$\text{EFC}_{pt^2}^-, q = 2, n = 83, a = 8$	48.3 KB	523 KB	20 B	3.982 s	0.004 s	2.481 s
$\text{EFC}_p^-, q = 3, n = 59, a = 6$	48.8 KB	375 KB	28 B	2.938 s	0.004 s	12.359 s

for parallelism, it is not exploited beyond bit packing and vectorized addition (byte-wise xor) for \mathbb{F}_2 . The only other optimization that was used was the compiler’s optimization flag. For $q = 3$, the sizes are computed by representing elements of \mathbb{F}_3 by two bits.

5 Conclusion

Extension Field Cancellation (EFC) is a new construction for central trapdoors in MQ cryptosystems which exploits the commutativity of the extension field in order to cancel the complexity of the extension field polynomials. After cancellation, the plaintext can be obtained by solving a linear system. We anticipate several known attacks and use the projection and minus modifiers to inoculate EFC against these attacks.

We estimate parameters associated with 80 bits of security from the running time of an algebraic attack and offer some experimental validation of its complexity. Our implementation confirms the correctness of our schemes as well as their practical efficiency. Encryption can be done in only a few milliseconds, on par with other post-quantum cryptosystems such as NTRU [16] and McEliece [20]. However, due to the missing information from the minus modifier, decryption takes several seconds.

This minus modifier is an obvious candidate for improvement. While it is necessary for security, any significant number of dropped polynomials constitutes an onerous cost on the decryption function because its running time is exponential in this number. In fact, the minus modifier is ideally suited for MQ *signature* schemes, but ill-suited for MQ *encryption* schemes. The reason is that for signatures, any assignment to the missing variables will do; in contrast, the decryption algorithm must iterate over all possible assignments in order to find the correct plaintext. Any alternative modifier that has the same effect on security but obviates the need for exhaustive search can drastically accelerate decryption.

Another question is to determine to which extent the public keys can be shrunk. While it is difficult to shrink the secret keys without throwing away entropy, the public keys contain a large amount of redundancy. Even a relatively moderate reduction in the public key size can make the cryptosystem a feasible option for applications where the public key size is critical and currently too large.

Acknowledgments. The authors would like to thank the anonymous reviewers for their helpful comments. This work was supported in part by the Research Council KU Leuven: C16/15/058. In addition, this work was supported by the Flemish Government, FWO WET G.0213.11N and by the European Commission through the ICT programme under contract FP7-ICT-2011-284833 PUFFIN, FP7-ICT-2013-10-SEP-210076296 PRACTICE, through the Horizon 2020 research and innovation programme under grant agreement No H2020-ICT-2014-644371 WITDOM and H2020-ICT-2014-645622 PQCRYPTO; as well as by grant USDC (NIST) 60NAN15D059 from the Nation Institute of Standards of Technology. Alan Szepieniec is funded by a research grant of the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen).

References

1. NESSIE: New European Schemes for Signatures, Integrity, Encryption 05 November 2014 (2003). <https://www.cosic.esat.kuleuven.be/nessie/>
2. Bettale, L., Faugère, J., Perret, L.: Cryptanalysis of HFE, multi-HFE and variants for odd and even characteristic. *Des. Codes Crypt.* **69**(1), 1–52 (2013). <http://dx.doi.org/10.1007/s10623-012-9617-2>
3. Billet, O., Patarin, J., Seurin, Y.: Analysis of intermediate field systems. *IACR Cryptology ePrint Archive* 2009, p. 542 (2009). <http://eprint.iacr.org/2009/542>
4. Courtois, N.T.: The security of Hidden Field Equations (HFE). In: Naccache, D. (ed.) *CT-RSA 2001*. LNCS, vol. 2020, pp. 266–281. Springer, Heidelberg (2001). http://dx.doi.org/10.1007/3-540-45353-9_20
5. Ding, J., Hodges, T.J.: Inverting HFE systems is quasi-polynomial for all fields. In: Rogaway, P. (ed.) *CRYPTO 2011*. LNCS, vol. 6841, pp. 724–742. Springer, Heidelberg (2011). http://dx.doi.org/10.1007/978-3-642-22792-9_41
6. Ding, J., Kleinjung, T.: Degree of regularity for HFE. In: *IACR Cryptology ePrint Archive* 2011, p. 570 (2011). <http://eprint.iacr.org/2011/570>
7. Ding, J., Schmidt, D.: Rainbow, a new multivariable polynomial signature scheme. In: Ioannidis, J., Keromytis, A.D., Yung, M. (eds.) *ACNS 2005*. LNCS, vol. 3531, pp. 164–175. Springer, Heidelberg (2005). http://dx.doi.org/10.1007/11496137_12
8. Ding, J., Yang, B.-Y.: Degree of regularity for HFEv and HFEv. In: Gaborit, P. (ed.) *PQCrypto 2013*. LNCS, vol. 7932, pp. 52–66. Springer, Heidelberg (2013). http://dx.doi.org/10.1007/978-3-642-38616-9_4
9. Ding, J., Yang, B., Cheng, C., Chen, C.O., Dubois, V.: Breaking the symmetry: a way to resist the new differential attack. *IACR Cryptology ePrint Archive* 2007, p. 366 (2007). <http://eprint.iacr.org/2007/366>
10. Dubois, V., Fouque, P.-A., Shamir, A., Stern, J.: Practical cryptanalysis of SFLASH. In: Menezes, A. (ed.) *CRYPTO 2007*. LNCS, vol. 4622, pp. 1–12. Springer, Heidelberg (2007). http://dx.doi.org/10.1007/978-3-540-74143-5_1
11. Dubois, V., Fouque, P.-A., Stern, J.: Cryptanalysis of SFLASH with slightly modified parameters. In: Naor, M. (ed.) *EUROCRYPT 2007*. LNCS, vol. 4515, pp. 264–275. Springer, Heidelberg (2007). http://dx.doi.org/10.1007/978-3-540-72540-4_15
12. Faugère, J.C.: A new efficient algorithm for computing gröbner bases (f4). *J. Pure Appl. Algebra* **139**(1), 61–88 (1999)
13. Faugère, J.C.: A new efficient algorithm for computing gröbner bases without reduction to zero (f5). In: *Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation, ISSAC 2002*, pp. 75–83. ACM, New York (2002). <http://doi.acm.org/10.1145/780506.780516>

14. Faugère, J.-C., Joux, A.: Algebraic cryptanalysis of Hidden Field Equation (HFE) cryptosystems using Gröbner bases. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 44–60. Springer, Heidelberg (2003). http://dx.doi.org/10.1007/978-3-540-45146-4_3
15. Gaborit, P. (ed.): PQCrypto 2013. LNCS, vol. 7932. Springer, Heidelberg (2013). <http://dx.doi.org/10.1007/978-3-642-38616-9>
16. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: a ring-based public key cryptosystem. In: Buhler, J.P. (ed.) ANTS 1998. LNCS, vol. 1423, pp. 267–288. Springer, Heidelberg (1998). <http://dx.doi.org/10.1007/BFb0054868>
17. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592. Springer, Heidelberg (1999). http://dx.doi.org/10.1007/3-540-48910-X_15
18. Kipnis, A., Shamir, A.: Cryptanalysis of the HFE public key cryptosystem by relinearization. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 19–30. Springer, Heidelberg (1999). http://dx.doi.org/10.1007/3-540-48405-1_2
19. Matsumoto, T., Imai, H.: Public quadratic polynomial-tuples for efficient signature-verification and message-encryption. In: Günther, C.G. (ed.) EUROCRYPT 1988. LNCS, vol. 330, pp. 419–453. Springer, Heidelberg (1988). http://dx.doi.org/10.1007/3-540-45961-8_39
20. McEliece, R.J.: A public-key cryptosystem based on algebraic coding theory. DSN Prog. Rep. **42**(44), 114–116 (1978)
21. Patarin, J.: Cryptanalysis of the Matsumoto and Imai public key scheme of Eurocrypt 1988. In: Coppersmith, D. (ed.) CRYPTO 1995. LNCS, vol. 963, pp. 248–261. Springer, Heidelberg (1995). http://dx.doi.org/10.1007/3-540-44750-4_20
22. Patarin, J.: Hidden fields equations (HFE) and isomorphisms of polynomials (IP): two new families of asymmetric algorithms. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, pp. 33–48. Springer, Heidelberg (1996). http://dx.doi.org/10.1007/3-540-68339-9_4
23. Patarin, J., Goubin, L., Courtois, N.T.: C_{-+}^* and HM: variations around two schemes of T. Matsumoto and H. Imai. In: Ohta, K., Pei, D. (eds.) ASIACRYPT 1998. LNCS, vol. 1514, pp. 35–50. Springer, Heidelberg (1998). http://dx.doi.org/10.1007/3-540-49649-1_4
24. Porras, J., Baena, J., Ding, J.: ZHFE, a new multivariate public key encryption scheme. In: Mosca, M. (ed.) PQCrypto 2014. LNCS, vol. 8772, pp. 229–245. Springer, Heidelberg (2014). http://dx.doi.org/10.1007/978-3-319-11659-4_14
25. Shamir, A.: Efficient signature schemes based on birational permutations. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 1–12. Springer, Heidelberg (1994). http://dx.doi.org/10.1007/3-540-48329-2_1
26. Smith-Tone, D.: Properties of the discrete differential with cryptographic applications. In: Sendrier, N. (ed.) PQCrypto 2010. LNCS, vol. 6061, pp. 1–12. Springer, Heidelberg (2010). http://dx.doi.org/10.1007/978-3-642-12929-2_1
27. Tao, C., Diene, A., Tang, S., Ding, J.: Simple matrix scheme for encryption. In: Gaborit, P. (ed.) PQCrypto 2013. LNCS, vol. 7932, pp. 231–242. Springer, Heidelberg (2013). http://dx.doi.org/10.1007/978-3-642-38616-9_16
28. Thomae, E.: About the security of multivariate quadratic public key schemes. Ph.D. thesis, Ruhr-Universität Bochum (2013)
29. Thomae, E., Wolf, C.: Cryptanalysis of enhanced TTS, STS and all its variants, or: why cross-terms are important. In: Mitrokovtsa, A., Vaudenay, S. (eds.) AFRICACRYPT 2012. LNCS, vol. 7374, pp. 188–202. Springer, Heidelberg (2012). http://dx.doi.org/10.1007/978-3-642-31410-0_12

30. Wiedemann, D.H.: Solving sparse linear equations over finite fields. *IEEE Trans. Inf. Theor.* **32**(1), 54–62 (1986)
31. Wolf, C., Braeken, A., Preneel, B.: On the security of stepwise triangular systems. *Des. Codes Crypt.* **40**(3), 285–302 (2006). <http://dx.doi.org/10.1007/s10623-006-0015-5>
32. Wolf, C., Preneel, B.: Taxonomy of public key schemes based on the problem of multivariate quadratic equations. In: *IACR Cryptology ePrint Archive 2005*, p. 77 (2005). <http://eprint.iacr.org/2005/077>