One vector to rule them all: Key recovery from one vector in UOV schemes

Pierre Pébereau

Sorbonne Université, LIP6, CNRS
Thales SIX
pierre.pebereau@lip6.fr

Keywords: Cryptanalysis · Unbalanced Oil and Vinegar · Multivariate Cryptography

Abstract. Unbalanced Oil and Vinegar is a multivariate signature scheme that was introduced in 1999. Most multivariate candidates for signature schemes at the NIST standardization competition are either based on UOV or closely related to it. The philosophy of the scheme is that the signer has to solve only a linear system to sign a message, while producing a forgery should be as hard as solving a random quadratic system. To achieve this, the signer uses the UOV trapdoor, which is a secret subspace, the "oil subspace".
We show how to recover an equivalent secret key from the knowledge of a single vector in the oil subspace in any characteristic. From this vector, we obtain linear equations that contain enough information to dismiss the public quadratic equations and retrieve the secret subspace with linear algebra for practical parametrizations of UOV, in at most 13 seconds for modern instantiations of UOV. This proves that the security of the UOV scheme lies in the complexity of finding exactly one vector in the oil space.
We show how to extend this result to schemes related to UOV, such as MAYO and VOX.

1 Introduction

In order to replace number-theoretic problems used in cryptography that are threatened by quantum computing, such as factorization or the discrete logarithm problem, several families of problems have been extensively studied. One of them is related to multivariate polynomial system solving, and is referred to as "multivariate cryptography". The underlying problem, Polynomial System Solving "PoSSo", is NP-hard, and this gives confidence in the hardness of this problem for quantum computers. The history of this field shows that one must be very careful in order to construct secure schemes with these tools, as in the past many cryptosystems that seemed secure turned out to be broken by a wide variety of methods. Lately, the attacks on the Rainbow signature scheme have motivated a return to the roots, in particular to the parent scheme "Unbalanced

Oil and Vinegar" of Kipnis, Patarin and Goubin [15]. Recently, NIST has issued a call for alternate post-quantum signatures not relying on standard lattice assumptions. Among the submitted schemes, 10 out of 40 are based on multivariate polynomial systems, and 7 of them are closely related to UOV ([3], [7], [14], [21], [10], [5], [8]). The main features of these schemes are the short signatures and quick signing time, which are competitive with classical cryptography, and suggest applications with constrained bandwidth, such as embedded systems.

1.1 Related work in cryptanalysis

Many contributions to the cryptanalysis of UOV stem from the study of Rainbow [11], a more structured scheme built upon the foundations of UOV. In particular, the reconciliation attack [12] targeted Rainbow but is easily applied to UOV. This attack finds vectors in the secret subspace of the UOV trapdoor by exploiting their relationship with one another. More recently, Beullens introduced the intersection attack [2] which improves the first step of the reconciliation attack (finding the first (two) vector(s)). Beullens describes this reconciliation process in more detail in [3]. In that paper, he mentions that once enough vectors of $\mathcal{O}$ are found, one can dismiss the quadratic equations and solve a linear system. Using his bound, this process requires finding α vectors in $\mathcal{O}$ before being able to conclude, where $\alpha = 2$ for modern UOV instantiations, and more generally α is the ceiling of the ratio between the number of variables and the number of equations. Another key recovery attack against UOV is the Kipnis-Shamir attack, which targets invariant subspaces of some linear functions related to the public key. This attack is the one that motivated the "unbalanced" property of UOV. The state of the art for forgery attacks against UOV (direct attacks) consists of exploiting the underdeterminedness of the system to eliminate equations with the Thomae-Wolf algorithm [20]. The attacker then has to solve a system in $m-1$ variables and equations for modern UOV parameters.

1.2 Previous work in side-channel attacks

In the context of side-channel attacks, more precisely fault-injection attacks, Aulbach, Campos, Krämer, Samardjiska and Stöttinger recently published a paper with a similar result [1]. Their result can be stated in the same manner, namely that one vector yields a fast key recovery, which is expected to run in polynomial time but the complexity is not given. There is a fundamental difference in the reasoning and in the complexity achieved however, as they follow the intuition of Beullens' reconciliation attack as described in [3]: he observes that one needs only two vectors of the secret subspace to conclude because they induce an overdetermined linear system whose solution space is exactly $\mathcal{O}$. They use an adapted Kipnis-Shamir attack to obtain a second vector from the first one to conclude with this observation. In our case, we focus on the geometric point of view instead of the algebraic one. We show that a single vector is enough to characterize $\mathcal{O}$, without using the reconciliation modelling. Therefore we skip directly from one vector to the full key, without using a second vector as a stepping stone. We obtain very efficient algorithms both in theory and in practice for practical parameters of UOV, while their attack suffers from the cost of the reconciliation attack. The largest instance they attack using their tools (targeting NIST security level 5) takes a total of 12 hours including the Kipnis-Shamir and the reconciliation step, while our attack takes only 13 seconds on the same instance.

1.3 Contribution

In this paper, we prove that the difficulty of retrieving the UOV secret key is not only dominated by the complexity of finding the first vector in $\mathcal{O}$, but that, in fact, the problem becomes polynomial given a single vector in the secret subspace. Therefore, retrieving the UOV secret key is not harder than finding a single vector in the secret subspace. In addition, we show how this yields a polynomial-time answer to the question "$\boldsymbol{x} \in \mathcal{O}$?" without the secret key, which may be of independent interest, and is not possible with the tools introduced by [1] without going through the entire attack. We stated our result in a form as general as possible, enabling us to apply them to schemes based on or close to UOV. In particular, we also analyze the impact of our attack on MAYO and VOX.

We provide an implementation of our algorithms in SageMath [19] on a GitHub repository.

2 Preliminaries

2.1 Notations

Let q be a power of a prime and let $\mathbb{F}_q$ denote the finite field with q elements. If $q = p^m$ for p prime, we write p the characteristic of $\mathbb{F}_q$. Vectors are assumed to be column vectors and are denoted by bold letters: $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{o}, \ldots$. Matrices are denoted by capital letters, and transposition is written A^T. The right kernel of a matrix A is denoted by $\ker(A)$ unless mentionned otherwise: $\boldsymbol{x} \in \ker(A) \iff A\boldsymbol{x} = 0$. Given a field $\mathbb{F}$ and an integer n, we denote $\mathbb{F}[x_1, \ldots, x_n]$ or $\mathbb{F}[\boldsymbol{x}]$ the polynomial ring of $\mathbb{F}$ in the n indeterminates $x_1, \ldots, x_n$. The restriction of a function f to a set E is denoted $f_{|E}$. The cost of multiplying two square matrices of dimension n is $O(n^\omega)$, with $2 \leq \omega < 3$.

2.2 Quadratic forms

Let f be a quadratic form over a vector space $\mathbb{F}_q^n$.

A function $\mathcal{F} : \boldsymbol{x} \mapsto (f_1(\boldsymbol{x}), \ldots, f_m(\boldsymbol{x}))$ such that each f_i is a quadratic form is called a *quadratic map*. In fields of odd characteristic, a quadratic form f is characterized by its *polar form* $f^* := (\boldsymbol{x}, \boldsymbol{y}) \mapsto f(\boldsymbol{x} + \boldsymbol{y}) - f(\boldsymbol{x}) - f(\boldsymbol{y})$ which is a symmetric bilinear form. As such, it admits a symmetric matrix representation in $\mathbb{F}_q^{n \times n}$ that we identify with it, and with the original quadratic form. In other words, given a quadratic form f, there exists $M \in \mathbb{F}_q^{n\times n}$ such that for all $\boldsymbol{x} \in \mathbb{F}_q^n, f(\boldsymbol{x}) = f^*(\boldsymbol{x}, \boldsymbol{x}) = \boldsymbol{x}^T M\boldsymbol{x}$. In fields of even characteristic, there is no longer an equivalence with symmetric bilinear forms, as symmetric forms are also antisymmetric. Instead, we can represent quadratic forms using triangular matrices. Note that this is also true in fields of odd characteristic, but the set of triangular matrices is not stable by congruence, therefore changes of variables are more delicate in this setting. If q is odd, we say that f has *rank* r if the matrix associated to f has rank r. In particular, the rank is preserved by changes of variables in odd characteristic, which is not the case in even characteristic. A subspace $V \subset \mathbb{F}^n$ is *isotropic* for f if there exists $\boldsymbol{x} \in V$ such that $f(\boldsymbol{x}) = 0$, *totally isotropic* if for all $\boldsymbol{x} \in V, f(\boldsymbol{x}) = 0$, and *anisotropic* if for all $\boldsymbol{x} \in V \setminus \{0\}, f(\boldsymbol{x}) \neq 0$. For an introduction to quadratic forms, we refer the reader to [18].

We recall here a characterisation of totally isotropic subspaces that describes the secret key of UOV:

Lemma 1. *The subspace $\mathcal{O}$ is a totally isotropic subspace of a quadratic form f if and only if for all $(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{O}^2, f^*(\boldsymbol{x}, \boldsymbol{y}) = f^*(\boldsymbol{y}, \boldsymbol{x}) = 0$.*

Proof. Assume $\mathcal{O}$ is a totally isotropic subspace of the quadratic form f. Then for all $\boldsymbol{x}$ in $\mathcal{O}$, $f(\boldsymbol{x}) = 0$ by definition. Let $\boldsymbol{y} \in \mathcal{O}$. Then $f(\boldsymbol{y}) = 0$ and since $\mathcal{O}$ is a linear subspace, $\boldsymbol{x} + \boldsymbol{y} \in \mathcal{O}$, therefore $f(\boldsymbol{x} + \boldsymbol{y}) = 0$. Therefore,

$$f^*(\boldsymbol{x}, \boldsymbol{y}) = f(\boldsymbol{x} + \boldsymbol{y}) - f(\boldsymbol{x}) - f(\boldsymbol{y}) = f(\boldsymbol{x} + \boldsymbol{y}) = 0$$

Conversely, assume for all $(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{O}^2$, $f^*(\boldsymbol{x}, \boldsymbol{y}) = 0$. Notice in particular that $(\boldsymbol{x}, \boldsymbol{x}) \in \mathcal{O}^2$ therefore for all x in $\mathcal{O}$, $f(\boldsymbol{x}) = 0$. Therefore $\mathcal{O}$ is a totally isotropic subspace of f. □

Observe that the dimension of a totally isotropic subspace of a quadratic form of a certain rank is bounded:

Lemma 2. *Let f be a quadratic form of rank n defined over a field $\mathbb{K}$. Let $\mathcal{O}$ be a totally isotropic subspace of f. Then $\mathcal{O}$ has dimension less than or equal to $\lfloor \frac{n}{2} \rfloor$.*

Proof. By contradiction, assume that $\dim(\mathcal{O}) = r > \lfloor \frac{n}{2} \rfloor$. Let B be a basis of $\mathcal{O}$, let $\hat{B}$ be a completion of B into a basis of $\mathbb{K}^n$. Then the matrix representing f^* in basis $\hat{B}$ has a block of zeros of size $r \times r$ in the top left corner. Therefore its rank is less than n, which is a contradiction.

□

2.3 Cryptanalysis

Given a signature scheme instance $\Sigma = (\mathcal{S}, \mathcal{P})$ where $\mathcal{S}$ is the secret key and $\mathcal{P}$ is the public key, we define two goals of cryptanalysis:

– *Forgery*, which is achieved if an attacker can find a signature for *one* message in the message space of Σ.

– *Key recovery*, which is achieved if the attacker obtains an equivalent secret key $\mathcal{S}'$ enabling them to sign *any* message.

These notions can be refined to specify the tools and goals of the attacker, but this high-level description is enough for us.

2.4 Unbalanced Oil and Vinegar signatures

One of the oldest multivariate signature schemes was introduced in [17], and later generalised by [15], and remains standing after more than two decades. We formulate it in a more abstract manner than in the seminal paper, following the formalism of Beullens [2].

Definition 1 (Kipnis, Patarin, Goubin [15]). *A UOV instance is parametrized by the following parameters:*

– *m, the number of equations*
– *n, the number of variables*
– *q, the size of the finite field $\mathbb{F}_q$.*

The UOV public key is a set of m quadratic forms $G = (G_1, ..., G_m)$ of rank n over $\mathbb{F}_q$. The secret key is a totally isotropic subspace $\mathcal{O}$ of dimension m of the homogeneous component of degree two of each G_i.

This property is not generic for a family of quadratic forms, and the key generation will use a trick to introduce this structure. This trick was the original formulation of UOV in [15], and corresponds to a block of zeros of size m in the top left corner of the symmetric matrices representing the key in a secret basis. In particular, the secret key is a pair (A, F) where A is a linear change of variables (that characterizes $\mathcal{O}$) and F is a quadratic map where the variables $x_i, 1 \leq i \leq m$ appear linearly. We deduce the public key as $G = F \circ A$ by composing the secret quadratic map with the secret change of variables. Write $A^{-1} = [\boldsymbol{o}_1, ..., \boldsymbol{o}_m, \boldsymbol{v}_1, ..., \boldsymbol{v}_{n-m}]$ and observe that $\mathcal{O} = \mathsf{span}(\boldsymbol{o}_1, ..., \boldsymbol{o}_m)$. For $1 \leq i \leq m$, we call x_i an "oil" variable, and the remaining ones "vinegar" variables. We write $v = n - m$, the number of vinegar variables.

To sign a message $\mu \in \{0,1\}^*$ the signer solves the system: $F(\boldsymbol{x}) = \mathcal{H}(\mu) \in \mathbb{F}_q^m$ where $\mathcal{H}$ is a *cryptographic hash function.* This is a linear system in the oil variables, with m unknowns and equations after choosing random values for the vinegar variables. The verifier, given $\boldsymbol{y} = A^{-1}\boldsymbol{x}$ and μ, checks that $G(\boldsymbol{y}) = \mathcal{H}(\mu)$.

We introduce the *forgery variety* which is the set of signatures accepted for a given vector $\boldsymbol{z} \in \mathbb{F}_q^m$. In practice, we always sign $\boldsymbol{z} = \mathcal{H}(\mu)$.

Definition 2 (Forgery variety). *Let G be a UOV public key and $\boldsymbol{z} \in \mathbb{F}_q^m$. We define the* forgery variety *associated to $\boldsymbol{z}$ as the set of signatures of the vector $\boldsymbol{z}$:*

$$\mathbb{V}(\boldsymbol{z}) = \{\boldsymbol{x} \in \overline{\mathbb{F}_q}^n, 1 \leq i \leq m, G_i(\boldsymbol{x}) = z_i\}$$

This variety has dimension $n - m$, but we only care for solutions in $\mathbb{F}_q$.

Notice that $\mathcal{O} \subset \mathbb{V}(\mathbf{0})$. It is interesting to note that the distribution of UOV signatures is not uniform in this forgery variety.

We include as a reference the parameters chosen for UOV in recent submissions to the NIST competition.

	NIST SL	n	m	$\mathbb{F}_q$	$\lvert\mathsf{pk}\rvert$ (bytes)	$\lvert\mathsf{sk}\rvert$ (bytes)	$\lvert\mathsf{cpk}\rvert$ (bytes)	$\lvert\mathsf{sig}+\mathsf{salt}\rvert$ (bytes)
`ov-Ip`	1	112	44	$\mathbb{F}_{256}$	278 432	237 912	43 576	128
`ov-Is`	1	160	64	$\mathbb{F}_{16}$	412 160	348 720	66 576	96
`ov-III`	3	184	72	$\mathbb{F}_{256}$	1 225 440	1 044 336	189 232	200
`ov-V`	5	244	96	$\mathbb{F}_{256}$	2 869 440	2 436 720	446 992	260

Fig. 1: UOV parameters in [6]

3 Retrieving the UOV private key from one secret vector

In this section, we assume that $n \leq 3m$. This is the case for all recent instantiations of UOV, in particular the ones referred to in Figure 1. At the end of the section, we explain how we proceed for very unbalanced cases ($n > 3m$), and give some reasons why very unbalanced instances of UOV are unlikely to be used in practice.

We assume that we have acquired a single vector $\boldsymbol{x}$ in the secret subspace $\mathcal{O}$ either via a side-channel attack or from computations, and leverage this information to complete a key recovery attack in polynomial time. To summarize, the secret subspace is included in the kernel of each dual linear form $\boldsymbol{x}^T G_i$ by definition. The intersection of the m hyperplanes defined by these kernels is of dimension $n - m$, and still contains $\mathcal{O}$. Therefore, this intersection is a smaller subspace than the ambiant space $\mathbb{F}_q^n$ that still contains the secret subspace, and even small enough to entirely retrieve the secret subspace by considering the restriction of the public key quadratic forms to this subspace.

Before we start, we recall the Kipnis-Shamir attack that justifies why the cases $n \leq 2m$ are called "easy instances of UOV".

Lemma 3 (Kipnis-Shamir cryptanalysis of Oil and Vinegar [16], [15]). *Let G be a UOV public key with parameters n, m, q. Then the following holds:*

i. If $n = 2m$, there exists a probabilistic algorithm performing a key recovery attack against G in time $O(n^\omega)$.

ii. If $n > 2m$, there exists a probabilistic algorithm performing a key recovery attack against G in time $O(q^{n-2m} n^\omega)$.

iii. If $n < 2m$, there exists a deterministic algorithm performing a key recovery attack against G in time $O(n^\omega)$.

Proof. The first two cases are exactly the Kipnis-Shamir attack against OV [16] and the extension to the unbalanced case found in [15]. The last case comes from the observation that if $n < 2m$, then the existence of an m-dimensional totally isotropic subspace for a quadratic form implies by Lemma 2 that it is not full rank. Therefore we retrieve the subspace $\mathcal{O}$ by computing the kernels of the matrices representing the quadratic forms of the public key. Each kernel is a subspace included in $\mathcal{O}$ of positive dimension. Since generically a collection of m subspaces of positive dimension in $\mathcal{O}$ span a subspace of dimension at least m, we obtain a basis of $\mathcal{O}$ from the bases of these kernels. In the unlikely event where this process fails, we can repeat the procedure with the first vectors obtained. □

We note here that the Kipnis-Shamir attack described in the original paper [16] can be improved from a practical point of view. The essence of the attack resides in the fact that invariant subspaces of the public key matrices are eigenspaces of some closely related matrices in the set $\mathcal{T}$ defined as the closure of the $G_i^{-1}G_j$ under addition, mutliplication, and multiplication by an element of $\mathbb{F}_q$.

Therefore, the strategy of the attack is to find such eigenspaces by computing and factoring the characteristic polynomials of random elements of $\mathcal{T}$. The attack looks for irreducible factors of degree m which yield eigenspaces of dimension m, one of which is guaranteed to be the secret subspace $\mathcal{O}$. If the polynomial factors further than degree m, they consider that the attack has failed and move on to the next random element of $\mathcal{T}$.

This decision to dismiss such factors is coherent with the result proven in [16, Theorem 9], but experiments show that even smaller factors of the characteristic polynomial still yield part of the private key. We observe the following facts:

Fact 1 *Let $t \in \mathcal{T}$, let χ be the characteristic polynomial of t. Then χ is always a square*

Proof. This comes from the shape of the secret key in the balanced Oil and Vinegar scheme:

$$F_1 = A^T P_1 A = \begin{pmatrix} 0 & D \\ D^T & B \end{pmatrix}$$

where D is invertible if and only if F_1 is and D, B are square blocks of size m. In this case,

$$F_1^{-1} = \begin{pmatrix} * & D^{-1} \\ (D^{-1})^T & 0 \end{pmatrix}$$

The same structure holds for all $F_i, 1 \leq i \leq m$.
Therefore, $P_1^{-1}P_2 = (AF_1^{-1}A^T)(A^{-1T}P_2A^{-1}) = AF_1^{-1}F_2A^{-1}$. From the above, $F_1^{-1}F_2$ has diagonal blocks that are transpose of each other. Therefore the characteristic polynomial of $F_1^{-1}F_2$ is square, and $P_1^{-1}P_2$ is similar to $F_1^{-1}F_2$, therefore its characteristic polynomial differs only by a constant factor. This applies immediatly to linear combinations of the P_i and therefore to all $t \in \mathcal{T}$. □

Fact 2 *If $\chi = r^2$, then $\mathcal{O} = \ker(r(M))$.*

We insist on the fact that χ is always a square, because this is the main difference with the work presented in [16]: In [16, section 4.2], the authors remark that once a factorisation of $\chi(x) = P_1(x)P_2(x)$, with P_1, P_2 irreducible has been obtained, then $\mathcal{O}$ is either $\ker(P_1(M))$ or $\ker(P_2(M))$.

Since P_1 and P_2 are irreducible and χ is a square, it holds that $P_1 = P_2$. This was not exploited by [16] but it is a significant improvement from a practical point of view, as one only needs to consider one element of $\mathcal{T}$ obtained easily from the public key and never has to draw again (the original attack had a chance of failure, these observations show that the attack actually succeeds on every attempt). Technically, this observation also improves the complexity result because this turns the attack into a deterministic algorithm instead of a probabilistic one.

Lemma 4. *Let $G = (G_1, \ldots, G_m)$ be a homogeneous quadratic map of rank n represented by m matrices. Let $\mathcal{O}$ be a common totally isotropic subspace of $G_1, \ldots, G_m$. Let $\boldsymbol{x} \in \mathcal{O} \setminus \{\boldsymbol{0}\}$ and let $J(\boldsymbol{x}) = (\boldsymbol{x}^T G_1, \ldots, \boldsymbol{x}^T G_m)$. Then $\mathcal{O} \subset \ker(J(\boldsymbol{x}))$, and $\ker(J(\boldsymbol{x}))$ is generically an $(n-m)$-dimensional linear subspace of $\mathbb{F}_q^n$.*

Remark 1. $2J(\boldsymbol{x})$ is the Jacobian of G if the characteristic is not 2, hence the notation J.

Proof. Let $\boldsymbol{x} \in \mathcal{O} \setminus \{\boldsymbol{0}\}$. By Lemma 1, for all $\boldsymbol{z} \in \mathcal{O}$ and for all $g \in G$, we have: $g(\boldsymbol{x}) = g(\boldsymbol{z}) = 0$ and $g^*(\boldsymbol{z}, \boldsymbol{x}) = 0$. In particular, this implies that the kernel of the linear form $g_{\boldsymbol{x}} = g^*(\boldsymbol{x}, .)$ contains $\mathcal{O}$. By hypothesis, all the quadratic forms are of rank n, therefore this linear form is non-zero. Since it is a non-zero linear form, its kernel is a hyperplane.

We have shown that for all $1 \le i \le m, \mathcal{O} \subset \ker(\boldsymbol{x}^T G_i)$. Therefore,

$$\mathcal{O} \subset \bigcap_{1 \le i \le m} \ker(\boldsymbol{x}^T G_i) = \ker(J(\boldsymbol{x}))$$

We assume that these hyperplanes are generic among hyperplanes that contain $\mathcal{O}$. The intersection of m hyperplanes in general position has dimension $n - m$, which yields the conclusion. The probability that these hyperplanes are not in general position can be bounded by a constant by the Schwartz-Zippel lemma. Therefore, if it is not the case, we try again with a new vector and must retry at most a constant number of times.

$\square$

This lemma is the key to our attack. We apply it to the formalism of UOV in the following theorem:

Theorem 1 (Key recovery from one vector). *Let $G = (G_1, \ldots, G_m)$ be a UOV public key, let $\mathcal{O}$ be the secret subspace of G, and let $\boldsymbol{x} \in \mathcal{O} \setminus \{\boldsymbol{0}\}$.*

There exists an algorithm taking as input $(G, \boldsymbol{x})$ that outputs a basis of $\mathcal{O}$ in polynomial time. More precisely, Algorithm 2a performs this task and has complexity $O(mn^\omega)$, where $2 \le \omega \le 3$ is the exponent of matrix multiplication.

Proof. Recall that $n \leq 3m$. To keep notations simple, denote $K(\boldsymbol{x}) := \ker(J(\boldsymbol{x}))$. By Lemma 4, it holds that $\mathcal{O} \subset K(\boldsymbol{x})$. Let $B \in \mathbb{F}_q^{n \times (n-m)}$ be a basis of $K(\boldsymbol{x})$, which we can compute in time $O(n^\omega)$. We restrict the public key to $K(\boldsymbol{x})$:

$$\forall 1 \leq i \leq m, \quad G_{i|K(\boldsymbol{x})} = B^T G_i B \tag{1}$$

Computing these restrictions requires two matrix multiplications per element of the public key, which takes time $O(mn^\omega)$ in total.
Define $G_{|K(\boldsymbol{x})} = (G_{1|K(\boldsymbol{x})}, \ldots, G_{m|K(\boldsymbol{x})})$. By definition, $G_{|K(\boldsymbol{x})}$ is a UOV instance for parameters $(n-m, m)$. By hypothesis, $n \leq 3m$ therefore $n - m \leq 2m$. By Lemma 3, such an instance is broken in time $O(n^\omega)$, yielding a basis of the subspace $\hat{\mathcal{O}}$, the secret subspace of $G_{|K(\boldsymbol{x})}$. More precisely, in practice $n < 3m$ therefore we use the kernel approach (Lemma 3, iii.) instead of the Kipnis-Shamir attack. Once we obtain $\hat{\mathcal{O}}$, we take it back to the initial space $\mathbb{F}_q^n$ using B:

Let $C \in \mathbb{F}_q^{(n-m) \times m}$ be a basis of $\hat{\mathcal{O}}$. Then, for all $g \in G$:

$$(B \cdot C)^T g (B \cdot C) = C^T (B^T g B) C = C^T g_{|K(\boldsymbol{x})} C = 0 \in \mathbb{F}_q^{m \times m}$$

This proves that $B \cdot C$ is a basis of $\mathcal{O}$ since it is a free family of maximal cardinality included in $\mathcal{O}$. This matrix product costs $O(n^\omega)$, which yields a total complexity $O(mn^\omega)$. $\square$

one_vector_to_key$(G, \boldsymbol{x})$

- $m = |G|$
- $K(\boldsymbol{x}) = [\boldsymbol{x}^T G_i$ **for** $1 \leq i \leq m]$
- $B = \ker(K(\boldsymbol{x}))$
- $\hat{G} = [B^T G_i B$ **for** $1 \leq i \leq m]$
- $C = []$
- **for** $1 \leq i \leq m$:
  - **for** $\boldsymbol{z} \in \ker(\hat{G}_i)$:
    - **if** $z \notin \mathsf{span}(C)$:
      - $C = C \cup \{z\}$
  - **if** $|C| = m$:
    - **break**
- **return** BC

(a) Key recovery from one vector

in_secret_subspace$(G, \boldsymbol{x})$

- $n = |\boldsymbol{x}|$
- $m = |G|$
- $K(\boldsymbol{x}) = [\boldsymbol{x}^T G_i$ **for** $1 \leq i \leq m]$
- $B = \ker(K(\boldsymbol{x}))$
- $\hat{G} = [B^T G_i B$ **for** $1 \leq i \leq m]$
- **for** $1 \leq i \leq m$:
  - **if** $\mathsf{rank}(\hat{G}_i) > 2(n - 2m)$:
    - **return false**
- **return true**

(b) Fast test of $\boldsymbol{x} \in \mathcal{O}$?

Fig. 2: Algorithms

Remark 2. The previous result is true regardless of the characteristic of the field. In characteristic two, to ensure we are considering full rank matrices, we use the symmetric bilinear form defined by $G_i + G_i^T$ instead of G_i, which shares the same properties as G_i when considered as a linear map (namely a large block of zeroes in some basis). This observation is credited to Coppersmith in [16, Remark after Lemma 4.].

Remark 3. Notice that Lemma 4 relies on a genericity assumption. This assumption describes the usual case encountered in practice, but the non-generic cases are not signficantly harder:

1. If $\boldsymbol{x}$ is not "generic in $\mathcal{O}$", $\boldsymbol{x}$ may be a singular point of the variety, in which case $\dim(K(\boldsymbol{x})) = n - m + 1$, which does not prevent the success of Algorithm 2a unless $n \geq 3m$.
2. If the G_i are not chosen uniformly at random as required by the KeyGen algorithm, the dimension of $K(\boldsymbol{x})$ may be arbitrarily large (up to n). For example if some equation is a linear combination of the others, the variety is no longer a complete intersection and $K(\boldsymbol{x})$ has a larger dimension. But in this case, the public key system becomes easier to solve and therefore the key does not reach the claimed security level.
3. The algorithm will succeed in polynomial time as long as $\dim(K(\boldsymbol{x})) \leq 2m$ (in case of equality, we use the Kipnis-Shamir attack).

Notice that in Algorithm 2a, we include a break statement because with overwhelming probability, a subset of the kernels are enough to retrieve the secret key. We also obtain the following result as a corollary of this theorem, which was the initial motivation for this work.

Corollary 1. *Given G a UOV public key and $\boldsymbol{x} \in \mathbb{F}_q^n$, there exists a polynomial-time algorithm deciding whether $\boldsymbol{x} \in \mathcal{O}$.*

Note that this question is interesting only if $\boldsymbol{x}$ is in the forgery variety of the vector $\mathbf{0} \in \mathbb{F}_q^m$, as any vector that does not vanish the public key has no chance of being part of the secret subspace.

Intuitively, to prove the corollary, it suffices to apply the algorithm of Theorem 1 and conclude from a success or a failure. We do not need to apply all of the algorithm, as we distinguish the case $\boldsymbol{x} \in \mathcal{O}$ using the rank of the restrictions of the public key to $J(\boldsymbol{x})$. More precisely, any vector in the forgery variety induces a restriction of the public key by Lemma 4 where all the matrices have kernel of dimension at least 1. These kernels coincide on a dimension 1 subspace corresponding to the span of the original vector written in the new basis. But, if the vector belongs to a common totally isotropic subspace of dimension at least two, these kernels are larger and their intersection has a dimension that matches that of the totally isotropic subspace. This is the distinguisher we use.

We use the following lemma to specialize the algorithm of Theorem 1 for this task.

Lemma 5. *Let G be a collection of quadratic forms, $\boldsymbol{x} \in \mathbb{V}(\mathbf{0}) \setminus \{\mathbf{0}\}$ and $J(\boldsymbol{x}) = (\boldsymbol{x}^T G_1, \ldots, \boldsymbol{x}^T G_m)$. Let B be a basis of $\ker(J(\boldsymbol{x}))$. Then for all $g \in G$, $B^T g B$ has rank at most $n-m-1$.*

Proof. Let $1 \leq i \leq m$. Let us show that $B^T G_i B \in \mathbb{F}_q^{n-m \times n-m}$ has rank at most $n-m-1$. For this, we exhibit a non-zero element of the kernel of G_i. Notice that $J(\boldsymbol{x})\boldsymbol{x} = (\boldsymbol{x}^T G_1 \boldsymbol{x}, \ldots, \boldsymbol{x}^T G_m \boldsymbol{x}) = 0$ therefore $\boldsymbol{x} \in \ker(J(\boldsymbol{x}))$. Consequently there exist $(\lambda_1, ..., \lambda_{n-m}) \in \mathbb{F}_q^{n-m}$ not all zero such that $\boldsymbol{x} = \sum_{i=1}^{n-m} \lambda_i B_i$. Let $\boldsymbol{x'} = (\lambda_1, \ldots, \lambda_{n-m})^T$. Then $\boldsymbol{x'}^T B^T G_i B = \boldsymbol{x}^T (G_i B) = (\boldsymbol{x}^T G_i) B$ and by definition B is a basis of $\cap_{j \leq m} \ker(\boldsymbol{x}^T G_j)$ which implies that $\mathsf{span}(B) \subset \ker(\boldsymbol{x}^T G_i)$ and therefore $\boldsymbol{x'} \in \ker(B^T G_i B)$ (this is a left kernel!), which yields the upper bound on the rank of G_i. ◻

Proof (of Corollary 1).

For all i, the rank of $B^T G_i B$ is upper bounded by Lemma 5 since $\mathcal{O} \subset \mathbb{V}(\mathbf{0})$. Assume that $\cap_{i=1}^m \ker(B^T G_i B)$ has dimension at least 2. We show that this implies that $\boldsymbol{x}$ belongs to a linear subspace included in $\mathbb{V}(\mathbf{0})$ of dimension at least 2.

Let $\boldsymbol{x'}, \boldsymbol{y'}$ be a basis of $\cap_{i=1}^m \ker(B^T G_i B)$. Then define $\boldsymbol{x''} := B\boldsymbol{x'}$ and $\boldsymbol{y''} := B\boldsymbol{y'}$. Observe that for all i, $\boldsymbol{x''}^T G_i \boldsymbol{y''} = \boldsymbol{x'}^T B^T G_i B \boldsymbol{y'} = 0 = \boldsymbol{y''}^T G_i \boldsymbol{x''}$. $\boldsymbol{x''}, \boldsymbol{y''}$ must be linearly independent since $\boldsymbol{x'}, \boldsymbol{y'}$ were and B is a free family by definition. By Lemma 1, this implies there is a dimension two totally isotropic subspace $\mathsf{span}(\boldsymbol{x}, \boldsymbol{y})$ shared by the G_i.

We obtain a more precise result if $\boldsymbol{x} \in \mathcal{O}$: the kernel of the G_i must be of a large dimension and included in $\mathcal{O}$. To prove this, assume that $B = B_1 \oplus B_2$ where B_1 is a basis of $\mathcal{O}$ which is possible since $\mathcal{O} \subset \mathsf{span}(B)$. In this case, for all i,

$$B^T G_i B = \begin{pmatrix} 0 & C_1^{(i)} \\ C_1^{(i)T} & C_2^{(i)} \end{pmatrix}$$

where $C_1^{(i)} \in \mathbb{F}_q^{m \times (n-2m)}, C_2^{(i)} \in \mathbb{F}_q^{(n-2m) \times (n-2m)}$ and C_2 is symmetric. Since the block of zeros have size m, such a matrix has rank at most $n-2m+n-2m = 2(n-2m)$.

The variety $\mathbb{V}(0)$ generically contains no linear subspaces of dimension m by the Debarre-Manivel bounds [9, Théorème 2.1]. Therefore, we expect $\mathcal{O}$ to be the only such subspace.

By Lemma 5, we distinguish a vector of $\mathcal{O}$ from a generic vector of $\mathbb{V}(\mathbf{0})$ if $2(n-2m) < n-m-1 \iff n < 3m-1$. If the parameters are such that $n = 3m$ or $3m-1$, we can apply the algorithm from Theorem 1 which succeeds only if $\boldsymbol{x} \in \mathcal{O}$, and is polynomial in any case. In practical instances of UOV, $n = \frac{5}{2}m$. Therefore the rank is at most $2(n-2m) = m$, which allows us to use this more specific result.

This yields Algorithm 2b.

◻

3.1 Very unbalanced instances of UOV

The key recovery attack described in the previous section only works if $n \leq 3m$ or if $n \leq 4m$ and q is even. We show here what happens in the $n > 3m$ case. The algorithm of Theorem 1 does not yield an easy UOV instance, but instead a smaller UOV instance that has a small rank deficiency.

Keeping the formalism of Theorem 1, let $\hat{G} = G_{|K(\boldsymbol{x})}$ using the basis B of $K(\boldsymbol{x})$. This restriction can be defined regardless of the ratio $\frac{n}{m}$, and always corresponds to a UOV instance in dimension $n - m$. Next, recall that $\boldsymbol{x} \in \mathsf{span}(B)$ and therefore we can define $\hat{\boldsymbol{x}} = (\lambda_1, \ldots, \lambda_{n-m})$ where $v = \sum_{i=1}^{n-m} \lambda_i B_i$. By construction, this vector $\hat{\boldsymbol{x}}$ is in the secret subspace of $\hat{G}$.

Notice that both instances are equivalent since a solution of either can be translated to the other with the restriction basis B, and the restricted one is in dimension $n - m$ instead of n.

$$(G, \boldsymbol{x}, \mathcal{O}) \overset{B}{\Longleftrightarrow} (\hat{G}, \hat{\boldsymbol{x}}, \hat{O})$$

Further, by Lemma 5, this new UOV instance is composed of quadratic forms that are not full rank, and in particular which share a kernel contained in $\mathcal{O}$. This information is redundant with the secret vector we had for the original instance, as this kernel corresponds to $\mathsf{span}(\hat{\boldsymbol{x}})$. We are tempted to use this new vector $\hat{\boldsymbol{x}}$ that belongs to $\hat{\mathcal{O}}$ to repeat the attack inductively, but this fails because this vector is in the kernel of each matrix of the public key, which means that the matrix $J(\boldsymbol{x})$ is the zero matrix. Therefore, we need to solve a new UOV instance (which has some more structure in the form of the kernel we observed in this paragraph) that is strictly weaker against key recovery attacks. For very unbalanced instances of UOV, we will need a constant number of vectors in the secret key to conclude, in a similar fashion as observed by Beullens in [3]. More precisely, each independent vector in $\mathcal{O}$ allows to reduce the search space by m dimensions. We can conclude with β vectors if $n - \beta m \leq 2m \iff \beta \geq \lceil \alpha - 2 \rceil$ since β is an integer. Naturally this yields $\beta = 1$ for practical instances of UOV, which is the result presented in Section 3.

It could thus seem to be a good idea to aim for very unbalanced parameters, but there are two reasons why these parameters are unlikely to be used in practice:

1. Random polynomial systems are easier to solve when they are heavily unbalanced. An argument that justifies this statement is the generic algorithm of Thomae and Wolf [20] (especially in characteristic two), and more generally the observation that any new variable is a degree of liberty that can be exploited for free.
2. UOV already has large keys. Linear increases in n yield quadratic increases in the key sizes.

This highlights an interesting tradeoff in the security of UOV: the larger the parameter $\alpha = \frac{n}{m}$, the stronger UOV is against key recovery attacks, and the weaker it is against forgery attacks. Reciprocally, the smaller α is, the weaker UOV is against key recovery attacks, and the stronger it is against forgery attacks.

3.2 Experimental results

The algorithms we obtain have polynomial complexities. We show that they are also fast in practice by providing an implementation in SageMath [19], using native linear algebra functions. We test them against the parameter sets of [6]. The strategy is as follows: the oracle providing a vector in $\mathcal{O}$ is obtained by a function that chooses a random element in $\mathsf{span}(\boldsymbol{o}_1, \dots, \boldsymbol{o}_m)$, which are the first m columns of A^{-1}.

The code can be found at :

`https://github.com/pi-r2/OneVector`

We test the attack against the parameters of [6], which are representative of the state-of-the-art instantiations of UOV. We include a key size with twice as many variables as the target for NIST level 5 to demonstrate that the attack scales beyond NIST parameters. The hardware used is a laptop with an Intel CPU i7-1165G7 running at 2.80GHz with 8GB of RAM. All experiments were ran on a single thread.

Parameters n,m,q	uov-Ip 112, 44, 256	uov-Is 160, 64, 16	uov-III 184, 72, 256	uov-V 244, 96, 256	"uov-X" $(480, 192, 256)$
Time	1.7s	4.4s	5.7s	13.3s	145.1s

Fig. 3: Key recovery from one vector with our attack in $\mathbb{F}_{256}$

To obtain a complete key recovery, one must first find a vector of the secret subspace $\mathcal{O}$. Then, the attacker uses the attack described in this paper to complete his basis of the secret subspace, in a matter of seconds on a laptop.

Parameters n,m,q	uov-Ip 112, 44, 256	uov-Is 160, 64, 16	uov-III 184, 72, 256	uov-V 244, 96, 256	"uov-X" $(480, 192, 256)$
Time	0.2s	0.5s	0.7s	1.5s	9.1s

Fig. 4: "$\boldsymbol{x} \in \mathcal{O}$?" with our algorithm in $\mathbb{F}_{256}$

4 Applications to UOV variants

4.1 MAYO

The MAYO signature scheme [3] was introduced by Beullens as a generalization of UOV in which we allow the subspace $\mathcal{O}$ to have a smaller dimension than

m. We switch to the notations of Beullens for clarity. The size of the secret subspace of a MAYO key is denoted o, m remains the number of quadratic forms in the public key, n remains the dimension of the vector space $\mathbb{F}_q^n$, and q is a small power of two. In the UOV formalism used so far, $m = o$. In MAYO, o is significantly smaller than m. This transformation makes the scheme much more compact, but increases signature size. Beullens introduces some additional structure in the form of a "Whipping" transformation that maps $\mathbb{F}_q^n \to \mathbb{F}_q^{ko}$, instead of UOV which maps $\mathbb{F}_q^n \to \mathbb{F}_q^m$. This is required to allow the signer to sign. We obtain the UOV scheme for $k = 1$. The following parameters of MAYO have been submitted to NIST, taken from [4].

Parameter set security level	MAYO_1 1	MAYO_2 1	MAYO_3 3	MAYO_5 5
n	66	78	99	133
m	64	64	96	128
o	8	18	10	12
k	9	4	11	12
q	16	16	16	16

If we try to attack the UOV map of MAYO, then we consider a collection of m quadratic maps P_i, the public key maps, that share an o dimensional totally isotropic subspace. The attack proceeds as follows: Given $\boldsymbol{x} \in \mathcal{O}$, we obtain m linear forms $P_i'(\boldsymbol{x}, .)$, therefore the intersection of their kernels generically defines $J(\boldsymbol{x})$ an $n - m$ dimensional subspace that still contains $\mathcal{O}$. In the context of MAYO, $n - m \leq o$. Therefore we recover $\mathcal{O}$ entirely from the kernels of the restriction of the public key to $J(\boldsymbol{x})$. Notice that this does not improve the reconciliation attack on MAYO, as this was already achieved by Beullens in [3] with an algebraic method. This shows that the work done in Section 3 is consistent with the state of the art when transferred to MAYO.

4.2 VOX

To have a result as general as possible, we apply our attack to $\text{UOV}^{\hat{+}}$ [13]. This corresponds to a more general formulation of VOX known as FOX, which is introduced in the same specification as VOX [8], based on [13]. Notably, it relies on less assumptions than VOX and still has competitive signature sizes with UOV, with a priori improved security. This signature scheme is a UOV-like signature scheme where t quadratic forms of the secret key are random. These random equations are called "vinegar forms" and the usual UOV quadratic forms are called 'oil forms' by analogy. This is the $\hat{+}$ perturbation. The private key is then composed with two changes of variables (S, T) where $S \in GL_o(\mathbb{F}_q), T \in GL_n(\mathbb{F}_q)$ In traditional UOV, $S = I_n$ and $T = A^{-1}$.

$$\mathcal{F} = S \circ \mathcal{P} \circ T$$

The tradeoff is that the signer now has to solve a small quadratic polynomial system with t equations to sign a message.

The transformation S adds "noise" to the equations: the oil quadratic forms are mixed with the vinegar quadratic forms. This implies that the public system does not have a high-dimensional totally isotropic subspace like the UOV one. More precisely, we have the following shape of S chosen in [8]. Here $S \circ \mathcal{P}'$ is a left product: $S \circ \mathcal{P}' = \mathcal{P}' \cdot S$.

$$S = \begin{pmatrix} I_t & S' \\ 0 & I_{o-t} \end{pmatrix}, \quad S' \in \mathbb{F}_q^{(o-t)\times t} \tag{2}$$

We have for $1 \leq i \leq o$, $f_i = \sum_{j=1}^{o} s_{i,j} p'_j$, and more precisely using (2):

$$\begin{cases} 1 \leq i \leq t : f_i = p'_i \\ t+1 \leq i \leq o : f_i = p'_i + \sum_{j=1}^{t} s'_{i,j} p'_j \end{cases} \tag{3}$$

The main takeaway is that S has $t(o-t)$ unknown coefficients. For vectors in $\mathcal{O}$, the contribution of the oil forms to these mixed equations is zero, therefore we can retrieve this linear change of variables with linear algebra from the evaluation of the public key on oil vectors. Each evaluation yields $o-t$ equations by expressing the last $o-t$ coefficients of $\mathcal{P}(\boldsymbol{x})$ as linear combinations of the first t coefficients. Therefore we need t vectors in the oil subspace to retrieve the change of variables S.

Once this is done, we can apply the tools introduced earlier to recover T from $\mathcal{P}' = S^{-1} \circ \mathcal{P}$ which is a UOV system with t random equations. If we are given $x \in \mathcal{O}$, we will observe that it only vanishes $m-t$ of the quadratic forms of $\mathcal{P}'$. Each of the remaining t vinegar forms have probability $\approx \frac{1}{q}$ to vanish coincidentally on this vector but the knowledge of S allows us to distinguish the oil forms. In any case, the algorithm $\boldsymbol{x} \in \mathcal{O}$? would enable one to distinguish oil forms from vinegar forms even if the equations were permuted.

Then, we are able to reduce the $\mathcal{P}'$ instance to a smaller subspace of dimension $n-(m-t)$, as we only consider $(m-t)$ linear forms instead of m. FOX with $S = I_n$, which is exactly what $\mathcal{P}'$ is, shares the weakness of UOV to the Kipnis-Shamir attacks (Lemma 3), therefore we complete the attack if $n - m + t \leq 2m \iff n + t \leq 3m$. The parameters of FOX from [8] are in figure 5.

Variant	Security Level	q	o	v	t	|sig|	|cpk|	|csk|
FOX-I	128	251	48	72	8	120 B	47,056 B	64 B
FOX-III	192	4093	68	106	8	261 B	211,156 B	64 B
FOX-V	256	65521	91	140	8	462 B	694,892 B	64 B

Fig. 5: FOX parameters in [8].

We have $n = o+v$, where $o = m$ in our formalism. In all cases $n \leq 2.55o$, and in particular we respectively have $n+t = 122, 182, 239$ versus $3o = 144, 204, 273$ for security levels 1,3,5. Therefore our attacks apply to these parameter sets of FOX, but only with knowledge of S, which we obtain from t vectors of $\mathcal{O}$.

It is interesting to note that the signer has to solve a random system involving t quadratic equations, therefore the scheme does not allow much flexibility in the choice of t, as this task can only be done quickly for small values of t.

5 References

1. Aulbach, T., Campos, F., Krämer, J., Samardjiska, S., Stöttinger, M.: Separating oil and vinegar with a single trace side-channel assisted kipnis-shamir attack on UOV. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2023**(3), 221–245 (2023). https://doi.org/10.46586/tches.v2023.i3.221-245, https://doi.org/10.46586/tches.v2023.i3.221-245
2. Beullens, W.: Improved cryptanalysis of UOV and rainbow. In: Canteaut, A., Standaert, F. (eds.) Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12696, pp. 348–373. Springer (2021). https://doi.org/10.1007/978-3-030-77870-5_13, https://doi.org/10.1007/978-3-030-77870-5_13
3. Beullens, W.: MAYO: practical post-quantum signatures from oil-and-vinegar maps. In: AlTawy, R., Hülsing, A. (eds.) Selected Areas in Cryptography - 28th International Conference, SAC 2021, Virtual Event, September 29 - October 1, 2021, Revised Selected Papers. Lecture Notes in Computer Science, vol. 13203, pp. 355–376. Springer (2021). https://doi.org/10.1007/978-3-030-99277-4_17, https://doi.org/10.1007/978-3-030-99277-4_17
4. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: Mayo (2023), https://pqmayo.org, consulted 05/10/2023
5. Beullens, W., Chen, M.S., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B.Y., Schmidt, D., Shih, C.J., Tao, C., Yang, B.Y.: Uov (2023), uovsig.org, consulted 05/10/2023
6. Beullens, W., Chen, M., Hung, S., Kannwischer, M.J., Peng, B., Shih, C., Yang, B.: Oil and vinegar: Modern parameters and implementations. IACR Trans. Cryptogr. Hardw. Embed. Syst. **2023**(3), 321–365 (2023). https://doi.org/10.46586/tches.v2023.i3.321-365, https://doi.org/10.46586/tches.v2023.i3.321-365
7. Cogliati, B., Faugère, J.C., Fouque, P.A., Goubin, L., Larrieu, R., Macario-Rat, G., Minaud, B., Patarin, J.: Provable unbalanced oil and vinegar (2023), http://prov-sign.github.io, consulted 05/10/2023
8. Cogliati, B., Faugère, J.C., Fouque, P.A., Goubin, L., Larrieu, R., Macario-Rat, G., Minaud, B., Patarin, J.: Vox-sign (2023), http://vox-sign.com/files/vox_nist.pdf, consulted 05/10/2023
9. Debarre, O., Manivel, L.: Sur la variété des espaces linéaires contenus dans une intersection complète. Math. Annalen 312 (1998), 549-574 (1998). https://doi.org/https://doi.org/10.1007/s002080050235
10. Ding, J., Gong, B., Guo, H., He, X., Jin, Y., Pan, Y., Schmidt, D., Tao, C., Xie, D., Yang, B.Y., Zhao, Z.: Triangular unbalanced oil and vinegar (2023), tuovsig.org, consulted 05/10/2023
11. Ding, J., Schmidt, D.: Rainbow, a new multivariable polynomial signature scheme. In: Ioannidis, J., Keromytis, A.D., Yung, M. (eds.) Applied Cryptography and Network Security, Third International Conference, ACNS 2005, New York, NY,

USA, June 7-10, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3531, pp. 164–175 (2005). https://doi.org/10.1007/11496137_12, https://doi.org/10.1007/11496137_12
12. Ding, J., Yang, B., Chen, C.O., Chen, M., Cheng, C.: New differential-algebraic attacks and reparametrization of rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) Applied Cryptography and Network Security, 6th International Conference, ACNS 2008, New York, NY, USA, June 3-6, 2008. Proceedings. Lecture Notes in Computer Science, vol. 5037, pp. 242–257 (2008). https://doi.org/10.1007/978-3-540-68914-0_15, https://doi.org/10.1007/978-3-540-68914-0_15
13. Faugère, J.C., Macario-Rat, G., Patarin, J., Perret, L.: A new perturbation for multivariate public key schemes such as hfe and uov. Cryptology ePrint Archive, Paper 2022/203 (2022), https://eprint.iacr.org/2022/203
14. Furue, H., Ikematsu, Y., Kiyomura, Y., Takagi, T.: A new variant of unbalanced oil and vinegar using quotient ring: Qr-uov. In: Tibouchi, M., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2021. pp. 187–217. Springer International Publishing, Cham (2021)
15. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: Stern, J. (ed.) Advances in Cryptology - EUROCRYPT '99, International Conference on the Theory and Application of Cryptographic Techniques, Prague, Czech Republic, May 2-6, 1999, Proceeding. Lecture Notes in Computer Science, vol. 1592, pp. 206–222. Springer (1999). https://doi.org/10.1007/3-540-48910-X_15, https://doi.org/10.1007/3-540-48910-X_15
16. Kipnis, A., Shamir, A.: Cryptanalysis of the oil & vinegar signature scheme. In: Krawczyk, H. (ed.) Advances in Cryptology - CRYPTO '98, 18th Annual International Cryptology Conference, Santa Barbara, California, USA, August 23-27, 1998, Proceedings. Lecture Notes in Computer Science, vol. 1462, pp. 257–266. Springer (1998). https://doi.org/10.1007/BFb0055733, https://doi.org/10.1007/BFb0055733
17. Patarin, J.: The oil and vinegar signature scheme. In: Dagstuhl Workshop on Cryptography September, 1997 (1997)
18. Serre, J.P.: A course in arithmetic. Springer New York, NY (1978)
19. The Sage Developers: SageMath, the Sage Mathematics Software System (2022), https://www.sagemath.org, dOI 10.5281/zenodo.6259615
20. Thomae, E., Wolf, C.: Solving underdetermined systems of multivariate quadratic equations revisited. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) Public Key Cryptography - PKC 2012 - 15th International Conference on Practice and Theory in Public Key Cryptography, Darmstadt, Germany, May 21-23, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7293, pp. 156–171. Springer (2012). https://doi.org/10.1007/978-3-642-30057-8_10, https://doi.org/10.1007/978-3-642-30057-8_10
21. Wang, L.C., Chou, C.Y., Ding, J., Kuan, Y.L., Li, M.S., Tseng, B.S., Tseng, P.E., Wang, C.C.: Snova (2023), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/SNOVA-spec-web.pdf, consulted 02/01/2024