Wedges, oil, and vinegar
An analysis of UOV in the exterior algebra

Lars Ran

Radboud Universiteit Nijmegen, Netherlands
lars.ran@ru.nl

Abstract. The Unbalanced Oil and Vinegar (UOV) construction has been a central framework in multivariate cryptography since its appearance in 1999. In fact, four schemes in the second round of NISTs additional call for signatures are UOV-based. For efficiency considerations, most of these schemes are defined over a field of characteristic 2. This has as a side effect that the polar forms of the UOV public maps are not only symmetric, but also alternating.
In this work, we propose a new key-recovery attack on UOV over fields of characteristic 2 that exploits this alternating property. We interpret the polar forms of the UOV public map as elements of the exterior algebra. Moreover, we show that these forms exhibit a structure that is dependent on the secret oil space. Using the Plücker embedding, we also express the dual of the secret oil space as an element of the exterior algebra. Utilizing the structure of the public maps, we can formulate relations on this secret element in this algebra. Finally, we demonstrate that the secret oil space can be recovered using sparse linear algebra techniques.
This new attack has a lower time complexity than previous methods and reduces the security of `uov-Ip`, `uov-III`, and `uov-V`, by 4, 11, and 20 bits respectively. In addition, the attack is applicable to MAYO_2 and reduces its security by 28 bits.

1 Introduction

Quantum computers could potentially break the currently used public-key cryptography such as RSA and ECDSA due to Shor's polynomial-time quantum algorithm [29] against the hidden subgroup problem. In order to be one step ahead, different institutes, such as NIST [22], began an effort to standardize new cryptographic schemes that should be quantum resistant. In the meantime, many schemes have been proposed, many have been broken, and a few remain standing. This led NIST to standardize three schemes last year [25,24,26] based on lattice-based cryptography and hash-based cryptography. However, lattice-based and hash-based hardness assumptions are not the only viable hardness assumptions, and to diversify assumptions, NIST opened an additional call for digital signature schemes [23].

The Multivariate Quadratic (MQ) problem is among these potential hard problems. This problem assumes that it is hard to find a solution to a multivariate quadratic system of equations over a finite field. There are two main types of

trapdoor-based schemes based on MQ, Unbalanced Oil and Vinegar (UOV) and Hidden Field Equations (HFE). Since the attack by Tao, Petzoldt, and Ding [30], most variations of HFE have been broken. The UOV-construction, proposed in 1999 by Kipnis, Patarin, and Goubin [19], on the other hand, still receives a lot of attention. In fact, 7 out of the 40 schemes proposed in the NIST additional call for digital signatures are based on UOV. Of these, 4 got into the second round out of the 14 second round schemes in total. It is safe to say that the community is still interested in UOV.

Initially, UOV was proposed as just Oil and Vinegar by Patarin [27]. However, just a year later, this was broken by Kipnis and Shamir [20], leading to the currently used UOV [8]. Since then, different variations of UOV have been proposed that aim to reduce the large public-key size of such systems. Rainbow [11] was one such scheme that was promising for a long time until it was famously broken by Beullens [6]. Currently, the new variations MAYO [7], QR-UOV [1], and SNOVA [31] comprise the three additional schemes that are in the second round of the NIST competition.

Meanwhile, the cryptanalytic toolkit for UOV has also expanded. For forging signatures, there have been significant improvements in underdetermined MQ solvers [16]. And for key recovery, the state-of-the-art is the reconciliation [12] and intersection [5] attacks, both relying heavily on polar forms.

Our contribution. In this work, we propose a new attack against UOV that gives a different mathematical perspective on the private and public keys and how they relate. In particular, we establish a new relation between the polar forms and the dual of the oil space by considering them as elements of the exterior algebra. Whereas polar forms q are commonly exploited by using their relation $q(o, o') = 0$ for oil vectors o, o', we take a different approach. We rewrite the polar form in terms of the dual of the oil space and use these to establish linear constraints on the Plücker embedding of this dual in the exterior algebra. Next, we provide a simple way to retrieve the oil space, the private key, as a kernel vector of an explicit, sparse but large, linear map. We show that the correctness of this attack is based on the well-known heuristic assumption on the existence of semi-regular sequences over $\mathbb{F}_2$. With this attack, we provide new security estimates for UOV and MAYO. We compute the new estimates for the second round parameter sets that are affected by the attack and provide them in Table 1. Furthermore, we generalize this attack to fields of odd characteristic by leveraging the full power of the Plücker embedding and considering symmetric 2-forms of the exterior algebra. This generalization is independent of the given odd characteristic p. Finally, we provide an implementation of the attack in `MAGMA` [9], exhibiting correctness, and providing timing results. The implementation, together with the complexity estimators, can be found in the repository below

`https://github.com/LarsMath/wedge-attack`

Subsequent work. After a first draft of this paper appeared, there have been two works that build upon these results. The first [18] is a generalization of

Set	Previous best	This work
uov-Ip	145	**141**
uov-Is	**143**	176
uov-III	218	**207**
uov-V	278	**258**
MAYO_2	141	**113**

Table 1. The $\log_2$ complexities of parameter sets affected by the attack

the attack presented in this paper to fields of odd characteristic. This is most effective when the characteristic is low and the authors exhibit parameters for which it achieves a lower complexity than the reconciliation attack. Note that this generalization is very different from the one presented in this paper. Where [18] considers symmetric d-forms on $\mathbb{F}_q^n$, here we consider symmetric 2-forms on the exterior algebra. Furthermore, its cost is dependent on the size of the base-field, whereas ours is independent of it.

The second work is an adaptation of the attack against SNOVA [21]. Because SNOVA consists of UOV systems with added structure, our analysis and attack do not immediately carry over to SNOVA systems. In [21], the authors analyze the attack with the added structure and show that SNOVA does not meet the security levels for some of the parameter sets.

2 Preliminaries

Here, we briefly recap the UOV problem together with some observations that will be useful to us. Next, we will provide a basic treatment of the exterior algebra to provide the necessary definitions and tools to understand the attack.

Notation We use small bold letters, $\mathbf{v}, \mathbf{x}$ for vectors and capital bold letters $\mathbf{V}, \mathbf{P}$ for matrices. By $\mathbf{e}_1, \ldots, \mathbf{e}_n$, we denote the standard basis vectors of $\mathbb{F}_q^n$. We denote the maximal minor of $\mathbf{V}$ indexed by J, i.e. the determinant of the submatrix consisting of columns indexed by J, by $|\mathbf{V}_{*J}|$. We denote $[n] = \{1, \ldots, n\}$ and for $J \subset [n]$ we denote $J^c = [n] \setminus J$ if n is clear from the context.

2.1 UOV

As the name suggests, most UOV-based schemes rely on the hardness assumption of unbalanced oil and vinegar.

Problem 1. Let $v > o$, and m be positive integers and set $n = v + o$. Given a homogeneous quadratic map $(p_1, \ldots, p_m) : \mathbb{F}_q^n \to \mathbb{F}_q^m$ that vanishes on a subspace $O \subset \mathbb{F}_q^n$ of dimension o, find O.

We denote such a problem by (v, o, m)-UOV. However, often the (v, o, m) are left implicit. The subspace O is called the oil space and, by extension, $O^\perp$ the

vinegar space. When the problem was first proposed, only the $v = o$ case was considered, and it was simply called oil and vinegar. Since this was broken in [20], only the unbalanced case $v > o$ is still viable.

Polar Forms. A common tool in the study of UOV is to consider its polar forms. The polar form of a homogeneous quadratic map p is the bilinear map q given by

$$q(\mathbf{x}, \mathbf{y}) = p(\mathbf{x} + \mathbf{y}) - p(\mathbf{x}) - p(\mathbf{y}) \qquad \forall\ \mathbf{x}, \mathbf{y} \in \mathbb{F}_q^n.$$

The reason why these forms are interesting in the study of UOV comes from the fact that for such forms and any oil vectors $\mathbf{o}, \mathbf{o}'$ we obtain $q(\mathbf{o}, \mathbf{o}') = 0$. Let $\mathbf{o}_1, \ldots, \mathbf{o}_o \in O$ be a basis of the oil space, then we get the following implications:

$$p(\mathbf{o}) = 0 \quad \forall\ \mathbf{o} \in O \iff \begin{matrix} q(\mathbf{o}_i, \mathbf{o}_j) = 0 \\ p(\mathbf{o}_i) = 0 \end{matrix} \quad \forall\ 1 \leq i < j \leq o$$

In other words, we can reduce from a property that should hold on every vector of the oil space to a property that only has to hold on a basis of the oil space.

A single oil vector is enough. In practice, it turns out that it is enough to find a single oil vector, i.e. a single vector in O. Then by [2,28] there exists a polynomial-time algorithm to retrieve the entire oil space. This is due to the polar forms. Given an oil vector $\mathbf{o}$, we obtain m linear constraints on all other oil vectors $\mathbf{o}'$ given by $q_i(\mathbf{o}, \mathbf{o}') = 0$.

However, note that not all vanishing vectors are oil vectors. There may be many solutions to $(p_1, \ldots, p_m)$ that do not belong to the oil space. In fact, the average zero-set of such systems is of size q^{n-m} while the oil space is only of size q^o. Therefore, if $n - m > o$, the zero-set is, with high probability, orders of magnitude larger than the oil space. In other words, not every solution to $(p_1, \ldots, p_m)$ results in a polynomial-time key-recovery attack that solves the UOV problem.

Projecting down. For most of the proposed UOV parameters and most existing algorithms to solve them, o is too high to be practical. A common technique in UOV cryptanalysis is to first project the problem down. For this we pick a random injection $\mathbf{P} : \mathbb{F}_q^{n'} \to \mathbb{F}_q^n$ and pre-compose this with the UOV maps to obtain:

$$(p_1 \circ \mathbf{P}, \ldots, p_m \circ \mathbf{P}) : \mathbb{F}_q^{n'} \to \mathbb{F}_q^n \to \mathbb{F}_q^m.$$

This is again a new UOV instance with oil space $O' = \mathbf{P}^{-1}(O \cap \operatorname{Im} \mathbf{P})$. Note that since $\mathbf{P}$ is injective, its inverse is well-defined on its image. For random injections $\mathbf{P}$, this will have dimension $\dim(O') = \dim O - (n - n')$ with high probability. If we now find any reduced oil vector $\mathbf{o} \in O'$ we can construct $\mathbf{Po} \in O$. In other words, we can solve a (v, o, m)-UOV instance by solving a (v, o', m)-UOV instance with $o > o'$. There is a limit to this; taking $o' < \frac{2v}{m} - 1$, spurious "oil spaces" might appear by chance that cannot be translated to the original instance.

In general, projecting down loses information. However, it can be beneficial if one already has more information than the algorithm needs. This happens, for example, in the reconciliation attack and in the intersection attack.

2.2 The exterior algebra

The main intuition for the attack comes from properties of the exterior algebra. Here we will give a pragmatic overview of it. We refer the interested reader to [15] for a more thorough treatment of the exterior algebra.

The exterior algebra. The exterior algebra of the vector space $\mathbb{F}_q^n$ is the graded anti-commutative[1] associative algebra $\bigwedge \mathbb{F}_q^n$ with product $\wedge$ generated by the basis vectors $\mathbf{e}_1, \ldots, \mathbf{e}_n$ of $\mathbb{F}_q^n$ in degree 1. In particular, its k-th graded component is spanned by the k-fold wedge products of *distinct* basis vectors:

$$\bigwedge^k \mathbb{F}_q^n = \langle \mathbf{e}_{j_1} \wedge \cdots \wedge \mathbf{e}_{j_k} \mid 1 \leq j_1 < \ldots < j_k \leq n \rangle_{\mathbb{F}_q} \quad 1 \leq k \leq n.$$

Together with $\bigwedge^0 \mathbb{F}_q^n = \langle \mathbf{1} \rangle_{\mathbb{F}_q} \cong \mathbb{F}_q$, this leads to the graded decomposition

$$\bigwedge \mathbb{F}_q^n = \bigoplus_{k=0}^{n} \bigwedge^k \mathbb{F}_q^n.$$

By a simple counting argument, we find their dimensions:

$$\dim\left(\bigwedge^k \mathbb{F}_q^n\right) = \binom{n}{k} \quad \text{and} \quad \dim\left(\bigwedge \mathbb{F}_q^n\right) = 2^n.$$

The defining property of the wedge product is that all vectors of $\mathbb{F}_q^n$ square to zero. One can check that this is equivalent to the basis vectors squaring to zero

$$\mathbf{e}_i \wedge \mathbf{e}_i = \mathbf{0} \in \bigwedge^2 \mathbb{F}_q^n \qquad \forall\ 1 \leq i \leq n.$$

This leads to the following anti-commutative product structure on the basis vectors of $\mathbb{F}_q^n$:

$$\mathbf{e}_i \wedge \mathbf{e}_j \mapsto \begin{cases} -\mathbf{e}_j \wedge \mathbf{e}_i & \text{if } i \neq j \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

For subsets $J = \{j_1 < \ldots < j_k\} \subset [n]$, we write $\mathbf{e}_J = \mathbf{e}_{j_1} \wedge \ldots \wedge \mathbf{e}_{j_k}$ so that any element $\alpha \in \bigwedge^k \mathbb{F}_q^n$ can be written uniquely as

$$\alpha = \sum_{J \subset [n], |J| = k} \alpha_J \mathbf{e}_J$$

with all $\alpha_J \in \mathbb{F}_q$. In some cases, we will denote by $\alpha_{j_1,\ldots,j_k}$ the coefficient of $\mathbf{e}_{j_1} \wedge \ldots \wedge \mathbf{e}_{j_k}$ for arbitrary sequences so that the coefficients $\alpha_{j_1,\ldots,j_k}$ are alternating in their indices.

[1] For characteristic 2 it is therefore also graded commutative.

Remark 1. In general, it holds that for odd k, any element of $\bigwedge^k \mathbb{F}_q^n$ squares to zero. However, remarkably, in characteristic 2, all elements of the exterior algebra square to zero due to the freshman's dream.

Alternating 2-forms. We would like to consider alternating 2-forms as elements of the exterior algebra. First, let us consider another description of such forms using the $\mathbb{F}_q^n$ inner product $\langle \cdot, \cdot \rangle$:

$$\begin{aligned} q(\mathbf{x}, \mathbf{y}) &= \sum_{i,j} q(\mathbf{e}_i, \mathbf{e}_j) \langle \mathbf{e}_i, \mathbf{x} \rangle \langle \mathbf{e}_j, \mathbf{y} \rangle \\ &= \sum_{i<j} q(\mathbf{e}_i, \mathbf{e}_j)(\langle \mathbf{e}_i, \mathbf{x} \rangle \langle \mathbf{e}_j, \mathbf{y} \rangle - \langle \mathbf{e}_j, \mathbf{x} \rangle \langle \mathbf{e}_i, \mathbf{y} \rangle). \end{aligned}$$

This suggests identifying

$$q(\mathbf{x}, \mathbf{y}) \mapsto \sum_{i<j} q(\mathbf{e}_i, \mathbf{e}_j) \cdot \mathbf{e}_i \wedge \mathbf{e}_j \quad \in \bigwedge^2 \mathbb{F}_q^n. \tag{1}$$

Similarly, we can define an inverse map

$$\sum_i \mathbf{v}_i \wedge \mathbf{w}_i \mapsto ((\mathbf{x}, \mathbf{y}) \mapsto \sum_i \langle \mathbf{v}_i, \mathbf{x} \rangle \langle \mathbf{w}_i, \mathbf{y} \rangle - \langle \mathbf{w}_i, \mathbf{x} \rangle \langle \mathbf{v}_i, \mathbf{y} \rangle). \tag{2}$$

Both identifications are linear maps and we have an isomorphism of vector spaces.[2]

2.3 The Plücker embedding.

Let us now introduce the Plücker embedding. This is a map from the Grassmannian, the set of k-dimensional subspaces of $\mathbb{F}_q^n$, to the projectivization of $\bigwedge^k \mathbb{F}_q^n$. This identification will allow us to consider subspaces of $\mathbb{F}_q^n$ as elements of the exterior algebra.

Definition 1. *The Plücker embedding is given by*

$$\begin{aligned} \iota : \mathrm{Gr}_k\left(\mathbb{F}_q^n\right) &\to \mathbb{P}\left(\bigwedge^k \mathbb{F}_q^n\right) \\ \langle \mathbf{x}_1, \ldots, \mathbf{x}_k \rangle_{\mathbb{F}_q} &\mapsto \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k \end{aligned}$$

Initially, it is not immediately clear why this should be a well-defined map, let alone why it should be injective. However, the following lemma, which tells us that this embedding encodes the minors of $[\mathbf{x}_1^T, \ldots, \mathbf{x}_k^T]$, will help us to see why it is so. We leave their proofs to the appendix.

[2] In fact, the alternating k-forms are exactly the elements of $\bigwedge^k (\mathbb{F}_q^n)^\vee$ and this identification with $\bigwedge^k \mathbb{F}_q^n$ is just an extension of $(\mathbb{F}_q^n)^\vee \cong \mathbb{F}_q^n$.

Lemma 1. *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *and define the matrix*

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_k \end{bmatrix} = \begin{bmatrix} x_{11} & \cdots & x_{1n} \\ \vdots & \ddots & \vdots \\ x_{k1} & \cdots & x_{kn} \end{bmatrix}$$

then

$$\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \sum_{J \subset [n], |J|=k} |\mathbf{X}_{*J}| \cdot \mathbf{e}_J.$$

Proof. See Section A

Corollary 1. $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *is linearly dependent* $\iff \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \mathbf{0}$

Corollary 2. *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *and* $\mathbf{y}_1, \ldots, \mathbf{y}_k \in \mathbb{F}_q^n$ *be two sets of linear independent vectors. Then*

$$\langle \mathbf{x}_1, \ldots, \mathbf{x}_k \rangle_{\mathbb{F}_q} = \langle \mathbf{y}_1, \ldots, \mathbf{y}_k \rangle_{\mathbb{F}_q} \iff \exists c \in \mathbb{F}_q^* \text{ s.t. } \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = c \cdot \mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_k.$$

Proof. See Section A.

This concludes that the Plücker embedding is indeed well-defined and injective.

3 The wedge attack

In characteristic 2, we can view both the polar forms and the (dual of the) oil space as elements of the exterior algebra. The first is due to the fact that they are alternating 2-forms and the second due to the Plücker embedding. In this section, we show that the polar forms reside in a specific subspace of the exterior algebra and that there are relations between those two objects that we can use to attack UOV.

3.1 The structure of UOV polar forms

Let us call $(q_1, \ldots, q_m)$ the polar forms associated with the UOV map $(p_1, \ldots, p_m)$,

$$q_i(\mathbf{x}, \mathbf{y}) = p_i(\mathbf{x} + \mathbf{y}) - p_i(\mathbf{x}) - p_i(\mathbf{y}).$$

In characteristic two, the polar forms are alternating

$$q_i(\mathbf{x}, \mathbf{x}) = p_i(\mathbf{x} + \mathbf{x}) - p_i(\mathbf{x}) - p_i(\mathbf{x}) = 0.$$

Therefore, we can consider them as elements of $\bigwedge^2 \mathbb{F}_q^n$. However, since they are coming from a UOV map, we can say a bit more. For any $\mathbf{o}, \mathbf{o}' \in O$, we obtain $q_i(\mathbf{o}, \mathbf{o}') = 0$. This gives us the following structure.

Lemma 2. *Let $q \in \bigwedge^2 \mathbb{F}_q^n$ be an alternating form such that $q(\mathbf{o}, \mathbf{o}') = 0$ for all $\mathbf{o}, \mathbf{o}'$ in a subspace $O \subset \mathbb{F}_q^n$ and let $\mathbf{v}_1, \ldots, \mathbf{v}_v$ be a basis for $O^\perp$. Then there exist $\mathbf{x}_i \in \mathbb{F}_q^n$ such that*

$$q = \sum_i \mathbf{v}_i \wedge \mathbf{x}_i.$$

Proof. Let $\mathbf{o}_1, \ldots \mathbf{o}_o$ be a basis for O and pick vectors $\mathbf{w}_1, \ldots, \mathbf{w}_o \in \mathbb{F}_q^n$ such that

$$\langle \mathbf{w}_i, \mathbf{o}_j \rangle = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j. \end{cases}$$

By construction, $(\mathbf{v}_1, \ldots, \mathbf{v}_v, \mathbf{w}_1, \ldots, \mathbf{w}_o)$ is a basis of $\mathbb{F}_q^n$. Therefore, there exist $\alpha_{ij}, \beta_{ij}, \gamma_{ij} \in \mathbb{F}_q$ such that:

$$q = \sum_{i<j} \alpha_{ij} \mathbf{v}_i \wedge \mathbf{v}_j + \sum_{i,j} \beta_{ij} \mathbf{v}_i \wedge \mathbf{w}_j + \sum_{i<j} \gamma_{ij} \mathbf{w}_i \wedge \mathbf{w}_j.$$

Using the identifications eq. (1) and eq. (2), this implies

$$\begin{aligned} 0 &= q(\mathbf{o}_a, \mathbf{o}_b) \\ &= \sum_{i<j} \gamma_{ij} (\langle \mathbf{w}_i, \mathbf{o}_a \rangle \langle \mathbf{w}_j, \mathbf{o}_b \rangle - \langle \mathbf{w}_j, \mathbf{o}_a \rangle \langle \mathbf{w}_i, \mathbf{o}_b \rangle) \\ &= \gamma_{ab} \end{aligned}$$

which proves the statement. □

Remark 2. This tells us that the alternating forms associated with a UOV public map lie in the subspace generated by the elements $\mathbf{v} \wedge \mathbf{x}$ with $\mathbf{v} \in O^\perp$ and $\mathbf{x} \in \mathbb{F}_q^n$, which is a subspace of dimension $\binom{n}{2} - \binom{o}{2}$. This is exactly the number of off-diagonal degrees of freedom when constructing the central map.

3.2 A special annihilator of the polar forms

We are now ready to use the structure that is present in the UOV maps. Recall that the wedge product acts multi-linearly and vanishes on linear dependent vectors as seen in Corollary 1. Denote by $\mathbf{v}_1, \ldots, \mathbf{v}_v$ a basis of $O^\perp$. Then, for a UOV polar form $q = \sum_i \mathbf{v}_i \wedge \mathbf{x}_i$, we observe

$$\begin{aligned} \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_v \wedge q &= \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_v \wedge \left(\sum_i \mathbf{v}_i \wedge \mathbf{x}_i \right) \\ &= \sum_i \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_v \wedge \mathbf{v}_i \wedge \mathbf{x}_i \\ &= \mathbf{0} \in \bigwedge^{v+2} \mathbb{F}_q^n. \end{aligned}$$

In practice we do not know the oil space nor its dual, let alone any of their bases. However, with the observation above we find a valuable property of the unknown oil space.

Proposition 1. *Let $q_1, \ldots, q_m$ be UOV polar forms, so that $q_i = \sum_j \mathbf{v}_j \wedge \mathbf{x}_{ij}$ for a basis $\mathbf{v}_1, \ldots, \mathbf{v}_v$ of $O^\perp$. Then, the wedge product $\mathcal{V} = \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_v \in \bigwedge^v \mathbb{F}_q^n$ lies in the kernel of the following linear map:*

$$\begin{aligned} \mathcal{M} : \bigwedge^v \mathbb{F}_q^n &\to \left(\bigwedge^{v+2} \mathbb{F}_q^n\right)^m \\ \alpha &\mapsto (\alpha \wedge q_1, \ldots, \alpha \wedge q_m). \end{aligned} \tag{3}$$

Note that, by Corollary 2, $\mathcal{V} \in \bigwedge^v \mathbb{F}_q^n$ is independent, up to scaling, of the choice of basis of $O^\perp$.

Proof. Apply the above observation to each entry. □

3.3 Retrieving the oil space from $\mathcal{V}$

With this property at hand, we obtain linear restrictions on the coefficients of $\mathcal{V}$ that come from the polar forms of the UOV public maps. Before we investigate whether we can now efficiently find $\mathcal{V}$, let us first show how to retrieve the oil space from $\mathcal{V}$.

We assume, without loss of generality, that the oil space trivially intersects the subspace $\langle \mathbf{e}_j : o < j \leq n \rangle$. We can then write the basis of O as the rows of

$$\begin{bmatrix} \mathbf{I}_o & \mathbf{O} \end{bmatrix} = \begin{bmatrix} & x_{1(o+1)} & \cdots & x_{1n} \\ \mathbf{I}_o & \vdots & \ddots & \vdots \\ & x_{o(o+1)} & \cdots & x_{on} \end{bmatrix}.$$

A basis of $O^\perp$ is then given by the rows of $\mathbf{V} = \begin{bmatrix} \mathbf{O}^T & \mathbf{I}_v \end{bmatrix}$. Since $\mathcal{V}$ is the wedge product of the rows of $\mathbf{V}$ up to scaling, it can be written in terms of its minors

$$\mathcal{V} = \mathcal{V}_{o+1,\ldots,n} \sum_{J \subset [n], |J| = v} |\mathbf{V}_{*J}| \cdot \mathbf{e}_J.$$

Now, we can verify that for $1 \leq i \leq o < j \leq n$ and $J = \{i\} \cup \{o+1, \ldots, n\}/\{j\}$, we obtain $|\mathbf{V}_{*J}| = x_{ij}$. In other words, we can simply read off a basis of O from the coordinates of $\mathcal{V}$ in the standard basis.

3.4 Undesirable kernel elements

Since we were able to construct the kernel vector $\mathcal{V}$ of $\mathcal{M}$, we might wonder whether there might be additional structural kernel elements. If such elements exist, then this prevents us from finding $\mathcal{V}$ as the unique kernel element. In what follows, we present two situations in which further structural kernel elements do arise. As will become clear in section 4.3, these are the only two exceptions that we encounter in our experiments.

When $v = 2m$, we can construct

$$q_1 \wedge \cdots \wedge q_m \in \bigwedge^v \mathbb{F}_q^n$$

as an element in the kernel of $\mathcal{M}$. By extension, if $v \geq 2m$, then, for every $\alpha \in \bigwedge^{v-2m} \mathbb{F}_q^n$, the element

$$\alpha \wedge q_1 \wedge \cdots \wedge q_m$$

lies in the kernel of $\mathcal{M}$ as well.

The other exception has a more UOV-like nature. For an oil space $\langle \mathbf{o}_1, \ldots, \mathbf{o}_o \rangle$ we need $p_i(\mathbf{o}_j) = 0$ and $q_i(\mathbf{o}_j, \mathbf{o}_{j'}) = 0$ for all $i, j \neq j'$. For small enough o, there might be spaces that act as an oil space by chance. Additionally, in $\mathcal{M}$, we only consider the polar forms of the UOV map, and there is nothing to enforce $p_i(\mathbf{o}_j) = 0$. We can approximate how many spaces O' we expect to be a solution to the polar forms $q_1, \ldots, q_m$. By counting the degrees of freedom we find that there are roughly

$$|\mathbb{F}_q|^{vo-\binom{o}{2}m}$$

spurious solutions. All of these lead to an undesirable kernel vector $\mathbf{v}'_1 \wedge \cdots \wedge \mathbf{v}'_v$ where $\mathbf{v}'_1, \ldots, \mathbf{v}'_v$ is a basis for the dual of such a spurious solution $O'^{\perp}$. To avoid both of these cases, we require $v < \min\{\frac{(o-1)}{2}, 2\} \cdot m$.

4 The rank of $\mathcal{M}$

Let us take a closer look at the linear map $\mathcal{M}$ that we constructed in eq. (3). Our goal is to retrieve $\mathcal{V}$ from the kernel of $\mathcal{M}$. Typically, the kernel is too large to be enumerated exhaustively. For that reason, we are interested in parameters such that $\mathcal{V}$ is the unique kernel element up to scaling. Ideally, we would like the kernel of $\mathcal{M}$ to be of dimension 1.

4.1 A Macaulay matrix in disguise

It turns out that the matrix $\mathcal{M}$ is actually equivalent to a Macaulay matrix of degree o of the quadratic maps $p_1, \ldots, p_m$ considered as homogeneous elements of

$$\mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^2, \ldots, x_n^2 \rangle.$$

Let us briefly recall what such matrices look like again.

Definition 2. *Let $\mathcal{F} = (f_1, \ldots, f_m)$ be a system of homogeneous quadratic polynomials in $\mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^2, \ldots, x_n^2 \rangle$. The degree d Macaulay matrix of $\mathcal{F}$,*

$$\mathcal{M}_{\mathcal{F}}^d : \mathbb{F}_q^{\binom{n}{d}} \to \mathbb{F}_q^{m\binom{n}{d-2}},$$

is the matrix whose rows are indexed by pairs (t, f_i) where t is a monomial of degree $d-2$, its columns are indexed by monomials u of degree d, and its coefficients are given by

$$(\mathcal{M}_{\mathcal{F}}^d)_{(t,f_i),u} = \mathrm{coeff}(u, tf_i).$$

Here $\mathrm{coeff}(u, g)$ *is the monomial coefficient of u in the polynomial g.*

Proposition 2. *Let $p_1, \ldots, p_m$ be UOV maps, and $q_1, \ldots, q_m$ their polar forms. Consider the homogeneous quadratic system $\mathcal{F} = (\tilde{p}_1, \ldots, \tilde{p}_m)$, where $\tilde{p}_i$ is the image of p_i in $\mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^2, \ldots, x_n^2\rangle$. Furthermore, let $\mathcal{M}$ be the map as in eq.* (3)*, then*

$$\dim \ker \mathcal{M} = \dim \ker \mathcal{M}^o_{\mathcal{F}}.$$

Proof. To understand how $\mathcal{M}$ is equivalent to such a Macaulay matrix, we will instantiate it by picking a basis. Let $\mathcal{M}$ be the matrix that has its rows and columns indexed by pairs (J, k) and subsets I, respectively. Here $1 \leq k \leq m$, $I, J \subset [n]$, $|J| = v + 2$, and $|I| = v$. Its coefficients are given by

$$\mathcal{M}_{(J,k),I} = \operatorname{coeff}\left(\mathbf{e}_J, q_k \wedge \mathbf{e}_I\right)$$

where $\operatorname{coeff}(\mathbf{e}_J, \alpha)$ is the coefficient α_J. Next, note that $\tilde{p}_k$ and q_k have exactly the same coefficients

$$\tilde{p}_k = \sum_{i<j} \alpha_{ij}^{(k)} x_i x_j \qquad \text{and} \qquad q_k = \sum_{i<j} \alpha_{ij}^{(k)} \mathbf{e}_i \wedge \mathbf{e}_j.$$

Let us denote $x_J = \prod_{i \in J} x_i$. Then, for any k, I, and J as above, we have

$$\mathcal{M}_{(J,k),I} = \operatorname{coeff}\big(\mathbf{e}_J, q_k \wedge \mathbf{e}_I\big) = \begin{cases} \alpha_{J \setminus I} & \text{if } I \subset J, \\ 0 & \text{otherwise,} \end{cases}$$

$$(\mathcal{M}^o_{\mathcal{F}})_{(x_{J^c}, \tilde{p}_k), x_{I^c}} = \operatorname{coeff}\big(x_{I^c}, \tilde{p}_k \cdot x_{J^c}\big) = \begin{cases} \alpha_{J \setminus I} & \text{if } J^c \subset I^c, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, there is a bijection between the rows and columns of both matrices that keeps the entries intact. Therefore, they must have the same nullity. □

The rank of Macaulay matrices of quadratic homogeneous systems in

$$\mathbb{F}_2[x_1, \ldots, x_n]/\langle x_1^2, \ldots, x_n^2\rangle$$

has been studied before in [4]. In that work the authors defined the notion of *semi-regular sequences over* $\mathbb{F}_2$ and posed that the Hilbert series of such systems is given by

$$T_{m,n} = \left[\frac{(1+t)^n}{(1+t^2)^m}\right]_+.$$

This is the polynomial obtained by truncating $\frac{(1+t)^n}{(1+t^2)^m}$ at the first negative coefficient. It is conjectured that most[3] random sequences are semi-regular over $\mathbb{F}_2$. Note that, by extension of scalars, these results carry over to $\mathbb{F}_{2^\ell}$ if we keep the $\mathbb{F}_2$ field equations. Hence, we conjecture the following.

[3] In [17] it was shown that there are families of parameters n, m for which this does not hold. In fact, for these parameters, semi-regular systems do not exist. This is in line with the fact that we were able to construct additional kernel elements in Section 3.4. However, for most parameters, this conjecture is still widely believed to be true.

Conjecture 1. Let v, o, m, ℓ be positive integers such that $\min\{\frac{(o-1)}{2}, 2\} \cdot m > v > o$ and let $n = v+o$. Then with high probability a random quadratic homogeneous sequence $p_1, \ldots, p_m \in \mathbb{F}_{2^\ell}[x_1, \ldots, x_n]/\langle x_1^2, \ldots, x_n^2 \rangle$ is semi-regular over $\mathbb{F}_2$.

Assuming that this conjecture holds and that UOV systems are indistinguishable from random homogeneous quadratic systems, we can predict the following.

$$\dim \ker \mathcal{M} = T_{m,n}[t^o].$$

Here $T_{m,n}[t^o]$ is the coefficient of t^o in $T_{m,n}$.

4.2 Describing the linear dependencies in the exterior algebra

As the Hilbert series shows, the constraints placed by $\mathcal{M}$ on the coefficients of $\mathcal{V} \in \bigwedge^v \mathbb{F}_q^n$ are not all linearly independent. In this section, we describe what these linear dependencies look like in the exterior algebra to create a more intuitive perspective. In fact, they appear in the same way as seen in the Support Minors model in [3], and we follow a similar approach to count them. Let us look at the following example.

Example 1. Let α_J be indeterminates over $\mathbb{F}_q$ and parameterize

$$\mathcal{V}_\alpha = \sum_{J \subset [n], |J| = v} \alpha_J \cdot \mathbf{e}_J.$$

Our system consists of linear equations $\mathcal{V}_\alpha \wedge q_i = \mathbf{0}$. Then, the coefficients of $\mathcal{V}_\alpha \wedge q_i \wedge q_i$ can be written as linear combinations of these linear equations. However, since $q_i \wedge q_i = \mathbf{0}$, these linear combinations must be identically zero independent of $\mathcal{V}_\alpha$. Similar interference appears due to $\mathcal{V}_\alpha \wedge q_i \wedge q_j - \mathcal{V}_\alpha \wedge q_j \wedge q_i = \mathbf{0}$. Therefore, we find $\binom{m+1}{2}\binom{n}{v+4}$ linear dependencies among the linear equations. We can continue this line of argument and find linear dependencies among the linear dependencies, and so on.

This suggests that the dimension of the kernel is given by

$$\dim \ker \ \mathcal{M} = \max\left\{ \sum_{i=0}^{\lfloor \frac{o}{2} \rfloor} (-1)^i \binom{m+i-1}{i} \binom{v+o}{v+2i},\ 1 \right\}.$$

Indeed, if we compare this sum with the Hilbert series, we obtain

$$\begin{aligned} \frac{(1+t)^{v+o}}{(1+t^2)^m}[t^o] &= \left(\sum_i \binom{v+o}{i} t^i \right) \left(\sum_i \binom{-m}{i} t^{2i} \right) [t^o] \\ &= \sum_i (-1)^i \binom{m+i-1}{i} \binom{v+o}{v+2i}. \end{aligned}$$

Therefore, in order to successfully run the algorithm, we require the prediction to be equal to 1. For this, we introduce the following notion of admissible parameters, which we will use as a condition for the attack.

Definition 3. *We call the tuple (v, o, m) of positive integers* admissible *if*

$$\min\left\{\frac{(o-1)}{2}, 2\right\} \cdot m > v > o \quad \textit{and} \quad \sum_{i=0}^{\lfloor \frac{o}{2} \rfloor} (-1)^i \binom{m+i-1}{i} \binom{v+o}{v+2i} \leq 1.$$

4.3 Experimental results

In [4,17] the authors conducted experiments on random systems to verify the Hilbert series of such systems. We will repeat such experiments against UOV-like systems to test the Hilbert series for UOV-like parameters and to test whether such systems indeed behave similarly to random quadratic systems up to solving. For different pairs of v, o, we verified the prediction for all relevant m such that $\min\{\frac{(o-1)}{2}, 2\} \cdot m > v > o$. The test results are summarized in Table 2.

o	3	4	5	6	7	8
v	≤ 36	≤ 26	≤ 19	≤ 14	≤ 11	≤ 9

Table 2. The (v, o) pairs that were tested for $q = 2^8$ and *all* relevant values of m. No anomalies with respect to Conjecture 1 were discovered.

5 Hybridization

Given that there is such a sharp boundary in admissible parameters, a natural question that arises is whether we can relax this constraint by hybridizing. Since the algorithm consists of solving large linear systems, fixing these variables is unlikely to help. Instead, we guess zeros in the generator matrix of $O^{\perp}$.

By Gaussian elimination, we know that $O^{\perp}$ has a, likely unique, vector of the form $\mathbf{v} = (0, \ldots, 0, *, *, \ldots, *)$ that has its $v-1$ initial coefficients equal to zero. With probability q^{-g}, the initial $v-1+g$ coefficients of $\mathbf{v}$ are equal to zero for $0 \leq g \leq o$. Consider the following generator matrix of $O^{\perp}$:

$$\mathbf{V} = \begin{bmatrix} \multicolumn{2}{c}{\mathbf{v}} \\ \hline & * \cdots * \\ \mathbf{I}_{v-1} & \vdots \ddots \vdots \\ & * \cdots * \end{bmatrix}$$

If we compute the minors of this matrix, we notice that all the minors contained in the first $v-1+g$ columns are identically zero. So, with the guessing cost q^g we can reduce the number of columns of $\mathcal{M}$ by $\binom{v-1+g}{v}$. Of course, this will mean that some rows in $\mathcal{M}$ are identically zero too. So instead we pose the following nullity conjecture

$$\dim \ker \mathcal{M}_g = \max\left\{ \sum_{i=0}^{\lfloor \frac{o}{2} \rfloor} (-1)^i \binom{m+i-1}{i} \left(\binom{v+o}{v+2i} - \binom{v-1+g}{v+2i} \right), 1 \right\}.$$

We verified these rank predictions for 3692 tuples (v, o, m, g) with $2m > v > o \geq g$.

Note that this approach works for any permutation of the columns. So in order to iterate through guesses, we can simply choose a new set $J \subset [n]$ of size $v - 1 + g$, eliminate all columns indexed by a set $I \subset J$ and compute the kernel.

The maximum amount of zeros we can guess in the first vector is equal to o, and this corresponds exactly to a $(v-1, o, m)$-UOV system. So in total, to solve a (v, o, m)-UOV system, one can instead solve (v', o, m)-UOV systems with $1 \leq g < o$ zeros with a probability of $q^{-o(v-v')-g}$.

6 Complexity analysis

Assuming Conjecture 1, we can analyze the complexity of the attack. Furthermore, we will compute new security estimates for the affected parameter sets in the second round of NISTs additional call for signatures.

6.1 Pruning $\mathcal{M}$

An obstacle that arises in handling $\mathcal{M}$ is that it is extremely rectangular once $o \geq 4$. This problem is two-fold. First, when doing linear algebra, we do a lot of superfluous operations on linearly dependent rows. Secondly, sparse linear algebra techniques require square matrices. Nevertheless, if we randomly select as many rows as there are columns in $\mathcal{M}$, then, with high probability, the kernel of the induced square submatrix coincides with that of $\mathcal{M}$. In fact, this approach allowed us to verify some of the higher parameters in Table 2 and was successful in $> 99\%$ of the cases.

6.2 The complexity formula

We can now proceed to compute the bit complexity of the algorithm. As a cost model, we assume free access to the coefficients of $\mathcal{M}$. Furthermore, we use that the cost of finding a kernel vector of a $N \times N$ matrix over $\mathbb{F}_q$ with ρ non-zero coefficients per row using sparse linear algebra [10] has a cost

$$3 \cdot \rho \cdot N^2 \cdot (2\log_2(q)^2 + \log_2(q)).$$

The matrix on which we will perform the sparse linear algebra is the pruned $\mathcal{M}$ corresponding to (v, o', m) where o' is the lowest $o' < o$ such that (v, o, m) is admissible. This matrix has size $\binom{v+o'}{v} \times \binom{v+o'}{v}$ and $\binom{v+2}{2}$ non-zero coefficients per row resulting in a total cost of

$$3\binom{v+2}{2}\binom{v+o'}{v}^2 (2\log_2(q)^2 + \log_2(q)).$$

Memory cost. As with any Wiedemann XL algorithm, this algorithm is highly memory-intensive, as we have to store the entire sparse matrix $\mathcal{M}$. Since the Wiedemann algorithm is used as a subroutine, the algorithm inherits its limitations and admits the same optimizations. The memory complexity of the algorithm is

$$\mathcal{O}\left(\rho \cdot N \cdot (\log N + \log q)\right).$$

Alternatively, we can re-compute the entries of $\mathcal{M}$ each time we need them. This introduces extra overhead, but does not use field arithmetic. In that case, the memory cost is the internal state of the Wiedemann algorithm, which has memory complexity

$$\mathcal{O}\left(\log q \cdot N\right).$$

6.3 Results

We apply the cost estimates to the parameters in the UOV signature scheme and the MAYO signature scheme as in the second round of NISTs additional call for signatures. For MAYO_1, MAYO_3, and MAYO_5, the tuples (v, o, m) are not admissible and therefore the attack does not apply. For `uov-Ip`, `uov-III`, `uov-V`, and MAYO_2, we obtain a lower attack complexity than previously known results. Interestingly, fixing zeros in $O^{\perp}$ did not find admissible parameters for MAYO_1, MAYO_3 and MAYO_5 and did not improve any of the complexities. It seems that for these field sizes the cost of hybridization is too steep. The results are summarized in Table 3.

	Specifications[8,7]					**This attack**	
Set	v	o	m	q	**Complexity**	o'	**Complexity**
`uov-Ip`	68	44	44	256	145	18	**141**
`uov-Is`	96	64	64	16	**143**	22	176
`uov-III`	112	72	72	256	218	26	**207**
`uov-V`	148	96	96	256	278	32	**258**
MAYO_2	64	17	64	16	141	13	**113**

Table 3. The different parameter sets and previous complexity estimates as given in the round 2 specifications together with the lowest admissible o' and $\log_2$ complexity of the new attack.

Additionally, we present the complexity growth of UOV-like parameters ($v \approx \frac{3}{2}m$) and MAYO-like parameters ($v = m$) in Figure 1. In the same figure, we depict the minimum value of o' that is *admissible*.

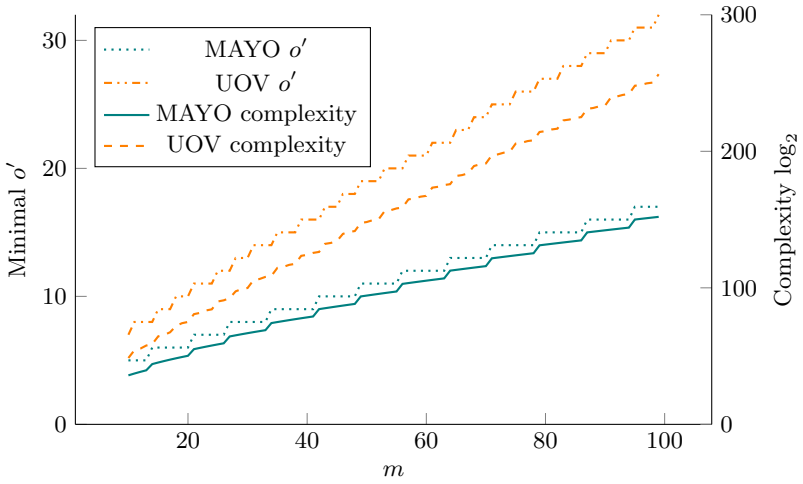


Fig. 1. Growth of attack complexity for UOV- and MAYO-like parameters.

6.4 Timings

We implemented the complete attack in `MAGMA` to test the attack in practice and obtain timings[4] for sets of small parameters. We chose to limit ourselves to parameters with $v = m$ (MAYO-like) and $v = \frac{3}{2} \cdot m$ (UOV-like). For the MAYO-like parameters, we assume that o is large enough for the attack to succeed. The results of those can be found in table 4 and table 5.

v	o	m	o'	Time (s)
9	6	6	6	4
10	6	6	6	13
11	7	7	6	35
12	8	8	6	88

Table 4. Timings for UOV-like parameters

v	o'	m	Time (s)
10	5	10	2
11	5	11	4
12	5	12	9
13	5	13	21
14	6	14	648
15	6	15	1365

Table 5. Timings for MAYO-like parameters

7 A generalization to odd characteristic

The main hurdle for this attack in fields of odd characteristic is that the polar forms are not alternating but only symmetric. Therefore, we cannot consider

[4] These experiments were carried out on a Dell PowerEdge R6515 with 1 AMD 7502P 2.5GHz 32-core processor and 64GB RAM using MAGMA version V2.28-8.

them as elements of the exterior algebra and hence the main idea fails. In [18], the authors overcome this by considering the polar forms and the vinegar space as elements of the p-reduced symmetric algebra. In this section, we will generalize the attack in a different fashion by considering the problem in a higher degree in the Plücker coordinates. Phrased differently, we will consider symmetric 2-forms on the exterior algebra and consider the polar forms as symmetric 2-forms of alternating 1-forms. The benefit of this is that the complexity does not depend on the odd characteristic p.

7.1 The structure of symmetric polar forms

With symmetric polar forms we can play the same game as before. Let us call $\mathrm{Sym}^2(W)$ the symmetric 2-forms of W which is the subspace of $W \otimes W$ generated by $\mathbf{x} \odot \mathbf{y} := \mathbf{x} \otimes \mathbf{y} + \mathbf{y} \otimes \mathbf{x}$ for all pairs $\mathbf{x}, \mathbf{y} \in W$.

Lemma 3. *Let $q \in \mathrm{Sym}^2 \mathbb{F}_q^n$ be a symmetric form such that $q(\mathbf{o}, \mathbf{o}') = 0$ for all $\mathbf{o}, \mathbf{o}'$ in a subspace $O \subset \mathbb{F}_q^n$ and let $\mathbf{v}_1, \ldots, \mathbf{v}_v$ be a basis for $O^\perp$. Then there exists $\mathbf{x}_i \in \mathbb{F}_q^n$ such that*

$$q = \sum_i \mathbf{v}_i \odot \mathbf{x}_i.$$

Proof. The proof is is similar to the proof of lemma 2. □

Now let us define a multiplication structure on $\mathrm{Sym}^2(\bigwedge \mathbb{F}_q^n)$:

$$\dot{\wedge} : \mathrm{Sym}^2(\bigwedge \mathbb{F}_q^n) \times \mathrm{Sym}^2(\bigwedge \mathbb{F}_q^n) \to \mathrm{Sym}^2(\bigwedge \mathbb{F}_q^n)$$
$$(\alpha \odot \beta) \mathbin{\dot{\wedge}} (\gamma \odot \delta) = (\alpha \wedge \gamma) \odot (\beta \wedge \delta) + (\alpha \wedge \delta) \odot (\beta \wedge \gamma)$$

One can check that this is, in fact, bilinear. Now we observe that for symmetric polar forms q coming from UOV,

$$\begin{aligned}(\mathcal{V} \odot \mathcal{V}) \mathbin{\dot{\wedge}} q &= (\mathcal{V} \odot \mathcal{V}) \mathbin{\dot{\wedge}} \left(\sum_i \mathbf{v}_i \odot \mathbf{x}_i \right) \\ &= \sum_i (\mathcal{V} \odot \mathcal{V}) \mathbin{\dot{\wedge}} (\mathbf{v}_i \odot \mathbf{x}_i) \\ &= \sum_i 2(\mathcal{V} \wedge \mathbf{v}_i) \odot (\mathcal{V} \wedge \mathbf{x}_i) \\ &= \mathbf{0}\end{aligned}$$

We can now bundle these again into a linear map to obtain

Proposition 3. *Let $q_1, \ldots, q_m$ be UOV polar forms, i.e. $q_i = \sum_j \mathbf{v}_j \odot \mathbf{x}_{ij}$ for some basis $\mathbf{v}_1, \ldots, \mathbf{v}_v$ of $O^\perp$. Then, for the wedge product $\mathcal{V} = \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_v \in$*

$\bigwedge^v \mathbb{F}_q^n$, *its symmetric square* $\mathcal{V} \odot \mathcal{V}$ *lies in the kernel of the following linear map:*

$$\begin{aligned} \mathcal{M} : \mathrm{Sym}^2\left(\bigwedge^v \mathbb{F}_q^n\right) &\to \left(\mathrm{Sym}^2\left(\bigwedge^{v+1} \mathbb{F}_q^n\right)\right)^m \\ \alpha &\mapsto (\alpha \dot{\wedge} q_1, \ldots, \alpha \dot{\wedge} q_m). \end{aligned} \tag{4}$$

Proof. Apply the above observation to each entry. □

7.2 The Plücker relations

Note that the solution that we are looking for, $\mathcal{V} \odot \mathcal{V}$, has a special shape and lies in the subspace[5]

$$\mathcal{S}_k \mathbb{F}_q^n := \left\langle (\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k)^{\odot 2} : \mathbf{x}_1, \cdots, \mathbf{x}_k \in \mathbb{F}_q^n \right\rangle_{\mathbb{F}_q} \subset \mathrm{Sym}^2\left(\bigwedge^k \mathbb{F}_q^n\right).$$

It is a classical result that the coefficients of $\mathcal{X} \in \bigwedge^k \mathbb{F}_q^n$ satisfy additional quadratic relations whenever $\mathcal{X} = \mathbf{x}_1 \wedge \ldots \wedge \mathbf{x}_k$. These are called the Plücker relations. Specifically, this means that $\mathcal{S}_k \mathbb{F}_q^n$ lies in the dual of the space generated by the Plücker relations in degree 2. There are multiple equivalent ways to construct these relations, here we will present one.

Recall in Section 3.3 that we extracted oil vectors from the wedge, assuming that it was a simple wedge. We generalize this to construct a wide range of vectors

Lemma 4. *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *be linear independent vectors and write*

$$\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \sum_{J \subset [n], |J| = k} \alpha_J \mathbf{e}_J.$$

Furthermore, let $I = \{i_1, \ldots, i_{k-1}\} \subset [n]$*, then*

$$\sum_j \alpha_{j, i_1, \ldots, i_{k-1}} \mathbf{e}_j \in \langle \mathbf{x}_1, \ldots, \mathbf{x}_k \rangle_{\mathbb{F}_q}.$$

Proof. See Section A.

For every set $I \subset [n]$ with $|I| = k - 1$ we can thus construct a vector $\mathbf{x} \in \langle \mathbf{x}_1, \ldots, \mathbf{x}_k \rangle_{\mathbb{F}_q}$ from the coefficients of $\mathcal{X}$. By Corollary 1, we must have $\mathbf{x} \wedge \mathcal{X} = 0$. Written out, these relations have the form

$$\sum_\ell (-1)^\ell \mathcal{X}_{i_1 \cdots i_{k-1} j_\ell} \mathcal{X}_{j_1 \cdots j_{\ell-1} j_{\ell+1} \cdots j_{k+1}} = 0.$$

[5] The space $\mathcal{S}_k \mathbb{F}_q^n$ is a Weyl module and is written as $\mathbb{S}_{(2^k)} \mathbb{F}_q^n$ in the literature since it comes from the Schur functor with partition (2^k).

We construct the subspace of $\mathrm{Sym}^2\left(\bigwedge^k \mathbb{F}_q^n\right)$ generated by these relations

$$\mathcal{P} := \left\langle \sum_\ell (-1)^\ell \mathbf{e}_{i_1 \cdots i_{k-1} j_\ell} \mathbf{e}_{j_1 \cdots j_{\ell-1} j_{\ell+1} \cdots j_{k+1}} : \begin{smallmatrix} 1 \le i_1 < \ldots < i_{k-1} \le n \\ 1 \le j_1 < \ldots < j_{k+1} \le n \end{smallmatrix} \right\rangle_{\mathbb{F}_q}.$$

By construction, these vectors are now dual to elements in $\mathcal{S}_k \mathbb{F}_q^n$. In general, the generators of $\mathcal{P}$ are not linearly independent. However, for high enough characteristic, the dimension of the space they span is quite well-understood. From [13, Chapter 9, Lemma 1] we find that

$$\mathcal{P} = \left(\mathcal{S}_k \mathbb{F}_q^n\right)^\perp$$

and from [14, Theorem 6.3, Chapter 15.3] we find that

$$\dim \mathcal{S}_k \mathbb{F}_q^n = \frac{1}{k+1} \binom{n}{k} \binom{n+1}{k}.$$

7.3 Restricting $\mathcal{M}$

Since $\mathcal{V} \odot \mathcal{V} \in \mathcal{S}_v \mathbb{F}_q^n$ we can include the Plücker relations when constructing the linear map $\mathcal{M}$. However, additional linear dependencies with the original relations will appear from the relations

$$\begin{aligned}(\mathcal{V} \odot \mathcal{V}) \dot{\wedge} p_i &= \sum_i 2(\mathcal{V} \wedge \mathbf{v}_i) \odot (\mathcal{V} \wedge \mathbf{x}_i) \\ &= \sum_i \frac{1}{2} \left(\mathcal{V} \wedge (\mathbf{v}_i + \mathbf{x}_i)\right)^{\odot 2} - \frac{1}{2} \left(\mathcal{V} \wedge (\mathbf{v}_i - \mathbf{x}_i)\right)^{\odot 2}.\end{aligned}$$

If $\mathcal{V}$ satisfies the Plücker relations for v-wedges, then $\mathcal{V} \wedge (\mathbf{v}_i + \mathbf{x}_i)$ and $\mathcal{V} \wedge (\mathbf{v}_i - \mathbf{x}_i)$ must satisfy the Plücker relations for $(v+1)$-wedges. In other words, when we restrict $\mathcal{M}$ to $\mathcal{S}_k \mathbb{F}_q^n$ we obtain a map

$$\mathcal{M}|_{\mathcal{S}_v \mathbb{F}_q^n} : \mathcal{S}_v \mathbb{F}_q^n \to \left(\mathcal{S}_{v+1} \mathbb{F}_q^n\right)^m.$$

We can repeat the argument in Example 1 to predict the syzygies

$$(\mathcal{V} \odot \mathcal{V}) \dot{\wedge} p_i \dot{\wedge} p_j = (\mathcal{V} \odot \mathcal{V}) \dot{\wedge} p_j \dot{\wedge} p_i \text{ for } i \neq j.$$

This leads us to the following rank conjecture.

Conjecture 2. Let $\tilde{\mathcal{M}}$ be the matrix consisting of the relations of $\mathcal{M}$ in *Equation* (4) and the Plücker relations $\mathcal{P}$, then with high probability

$$\dim \ker \tilde{\mathcal{M}} = \max \left\{ \sum_{i=0}^{o} (-1)^i \binom{m}{i} \frac{1}{v+i+1} \binom{v+o}{v+i} \binom{v+o+1}{v+i}, 1 \right\}$$

We tested this conjecture for 70 tuples (v, o, m) with $v > o$ in $\mathbb{F}_{31}$ of which 65 experiments verified this prediction. The tested tuples can be found in Table 6 and the instances in which the prediction was wrong in Table 7.

o	3									4			
v	4	5	6	7	8	9	10	11	12	5	6	7	8
m	$*$	$*$	$*$	$*$	$*$	$*$	≤ 8	≤ 4	≤ 3	$*$	$*$	≤ 2	≤ 1

Table 6. The (v, o, m) pairs that were tested for $q = 31$. A $*$ means all relevant values of m were tested.

v	o	m	Prediction	Experiment
4	3	2	126	127
4	3	3	1	6
5	4	2	792	835
6	4	2	4785	4786
6	4	3	1430	1466

Table 7. The (v, o, m)-tuples that were different from the prediction.

7.4 Complexity analysis

For the complexity analysis, we proceed similarly as in the case of an even characteristic. We assume that Conjecture 2 holds. It must be said that pruning $\tilde{\mathcal{M}}$ would likely lead to a drop in rank because the dependencies between the Plücker relations are quite sparse. Nevertheless, we optimistically assume that one can efficiently make a square submatrix without a rank drop. Furthermore, we also make the optimistic assumption that one can efficiently reduce the number of columns by choosing a basis of $\mathcal{S}_k \, \mathbb{F}_q^n$ in such a way that the number of nonzero coefficients per row does not increase. Then finally, with $(v+1)^2$ non-zero coefficients per row, we obtain the following cost in field operations

$$3 \cdot (v+1)^2 \cdot \left(\frac{1}{v+1} \binom{v+o'}{v} \binom{v+o'+1}{v} \right)^2 .$$

We computed the complexity for the parameter sets of (lifted-) QR-UOV and found that this attack does not reduce its security claim, even under these optimistic assumptions. The results are summarized in Table 8. Interestingly, the attack outperforms the reconciliation attack if UOV were chosen with odd parameters instead; see Table 9.

SL	q	v	o	m	o'	Complexity
I	127^3	52	18	54	13	182
I	31^3	55	20	60	13	186
III	127^3	76	26	78	17	249
III	31^3	82	29	87	18	266
V	127^3	102	35	105	22	328
V	31^3	108	38	114	23	344

Table 8. Complexity results for the attack against susceptible QR-UOV parameter sets. The complexity is in field operations.

q	v	m	Recon.	o'	Complexity
251	68	44	335	21	273
13	96	64	293	27	366
251	112	72	535	32	432
251	148	96	694	41	562

Table 9. Comparison of the new attack with the reconciliation attack against UOV in odd characteristic.

8 Conclusion and further research

With this new attack, we challenge the security claims of UOV and MAYO_2. For UOV the new complexities are just below the NIST thresholds, and increasing the parameters should bridge the gap. For MAYO only a single parameter set was susceptible, which is not a coincidence since this was also the outlier in the proposed sets. The small difference between $o = 17$ and $o' = 13$ for this set suggests that taking a smaller oil space and a larger whipping parameter would solve the problem.

Finally, considering this new structure opens up new research directions as well. In the even characteristic case of this attack, we only considered linear algebra on the exterior powers. However, the Plücker embedding provides us with quadratic relations with respect to the coefficients of $\mathcal{V}$ in the fields of characteristic 2 as well. The challenge here is to quantify the interaction between these quadratic relations and the linear relations used in the attack. Moreover, for both the even and odd cases, it would be interesting to investigate higher symmetric powers of the exterior algebra. This could provide a trade-off between lower o' but higher degree "monomials" to lower the total complexity. More importantly, this would provide a way around undesirable kernel elements and instances with a too small oil space. This, in turn, would make the attack applicable to a wider range of parameters.

Acknowledgments. The author thanks Simona Samardjiska for her involvement and advice in writing this paper. Furthermore, the author would like to thank Ray Perlner and Maxime Bros for pointing out an erroneous assumption on SNOVA systems in the earlier stages of this work. This research has been supported by the Dutch government through the NWO grant OCNW.M.21.193 (ALPaQCa).

References

1. Akiyama, R., Furue, H., Ikematsu, Y., Hoshino, F., Kinjo, K., Kosuge, H., Nakamura, S., Orihara, S., Takagi, T., Yamakoshi, K.: QR-UOV. Tech. rep., National Institute of Standards and Technology (2024), available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures`
2. 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 Transactions on Cryptographic Hardware and Embedded Systems **2023**(3), 221–245 (2023). `https://doi.org/10.46586/tches.v2023.i3.221-245`

3. Bardet, M., Bros, M., Cabarcas, D., Gaborit, P., Perlner, R.A., Smith-Tone, D., Tillich, J.P., Verbel, J.A.: Improvements of algebraic attacks for solving the rank decoding and MinRank problems. In: Moriai, S., Wang, H. (eds.) Advances in Cryptology – ASIACRYPT 2020, Part I. Lecture Notes in Computer Science, vol. 12491, pp. 507–536. Springer, Cham, Switzerland, Daejeon, South Korea (Dec 7–11, 2020). https://doi.org/10.1007/978-3-030-64837-4_17
4. Bardet, M., Faugère, J.C., Salvy, B.: Complexity of Gröbner basis computation for Semi-regular Overdetermined sequences over F_2 with solutions in F_2. Research Report RR-5049, INRIA (2003), https://inria.hal.science/inria-00071534
5. Beullens, W.: Improved cryptanalysis of UOV and Rainbow. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021, Part I. Lecture Notes in Computer Science, vol. 12696, pp. 348–373. Springer, Cham, Switzerland, Zagreb, Croatia (Oct 17–21, 2021). https://doi.org/10.1007/978-3-030-77870-5_13
6. Beullens, W.: Breaking rainbow takes a weekend on a laptop. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part II. Lecture Notes in Computer Science, vol. 13508, pp. 464–479. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15979-4_16
7. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
8. Beullens, W., Chen, M., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B., Schmidt, D., Shih, C., Tao, C., Yang, B.: UOV — Unbalanced Oil and Vinegar. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures
9. Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system. I. The user language. Journal of Symbolic Computation **24**(3-4), 235–265 (1997). https://doi.org/10.1006/jsco.1996.0125
10. Coppersmith, D.: Solving Homogeneous Linear Equations Over GF(2) via Block Wiedemann Algorithm. Mathematics of Computation **62**(205), 333–350 (1994). https://doi.org/10.2307/2153413
11. Ding, J., Chen, M.S., Petzoldt, A., Schmidt, D., Yang, B.Y., Kannwischer, M.J., Patarin, J.: Rainbow. Tech. rep., National Institute of Standards and Technology (2020), available at https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions
12. Ding, J., Yang, B.Y., Chen, C.H.O., Chen, M.S., Cheng, C.M.: New differential-algebraic attacks and reparametrization of Rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008: 6th International Conference on Applied Cryptography and Network Security. Lecture Notes in Computer Science, vol. 5037, pp. 242–257. Springer Berlin Heidelberg, Germany, New York, NY, USA (Jun 3–6, 2008). https://doi.org/10.1007/978-3-540-68914-0_15
13. Fulton, W.: Young Tableaux: With Applications to Representation Theory and Geometry. London Mathematical Society Student Texts, Cambridge University Press, Cambridge (1996). https://doi.org/10.1017/CBO9780511626241
14. Fulton, W., Harris, J.: Representation Theory, Graduate Texts in Mathematics, vol. 129. Springer, New York, NY (2004). https://doi.org/10.1007/978-1-4612-0979-9

15. Greub, W.: Multilinear Algebra. Universitext, Springer, New York, NY (1978). https://doi.org/10.1007/978-1-4613-9425-9
16. Hashimoto, Y.: An improvement of algorithms to solve under-defined systems of multivariate quadratic equations. JSIAM Letters **15**, 53–56 (2023). https://doi.org/10.14495/jsiaml.15.53
17. Hodges, T.J., Molina, S.D., Schlather, J.: On the existence of homogeneous semi-regular sequences in F2[X1,...,Xn]/(X12,...,Xn2). Journal of Algebra **476**, 519–547 (2017). https://doi.org/https://doi.org/10.1016/j.jalgebra.2016.11.025
18. Jin, Y., Pan, Y., He, X., Gong, B., Ding, J.: Security analysis on UOV families with odd characteristics: Using symmetric algebra. Cryptology ePrint Archive, Report 2025/1137 (2025), https://eprint.iacr.org/2025/1137
19. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced Oil and Vinegar signature schemes. In: Stern, J. (ed.) Advances in Cryptology – EUROCRYPT'99. Lecture Notes in Computer Science, vol. 1592, pp. 206–222. Springer Berlin Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/3-540-48910-X_15
20. Kipnis, A., Shamir, A.: Cryptanalysis of the Oil & Vinegar signature scheme. In: Krawczyk, H. (ed.) Advances in Cryptology – CRYPTO'98. Lecture Notes in Computer Science, vol. 1462, pp. 257–266. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 23–27, 1998). https://doi.org/10.1007/BFb0055733
21. Le, H., Bros, M., Lichtinger, J., Minaud, B., Perlner, R., Smith-Tone, D., Valenzuela, C.: Exploiting SNOVA's Structure in the Wedge Product Attack. In: 6th PQC Standardization Conference (2025)
22. NIST (National Institute for Standards and Technology): Post-Quantum Cryptography Standardization (2017), https://csrc.nist.gov/Projects/Post-Quantum-Cryptography
23. NIST (National Institute for Standards and Technology): Post-Quantum Cryptography: Additional Digital Signature Schemes (2022), https://csrc.nist.gov/projects/pqc-dig-sig, published: NIST Official Website
24. NIST (National Institute for Standards and Technology): Module-Lattice-Based Digital Signature Standard. Tech. Rep. Federal Information Processing Standard (FIPS) 204, U.S. Department of Commerce (Aug 2024). https://doi.org/10.6028/NIST.FIPS.204
25. NIST (National Institute for Standards and Technology): Module-Lattice-Based Key-Encapsulation Mechanism Standard. Tech. Rep. Federal Information Processing Standard (FIPS) 203, U.S. Department of Commerce (Aug 2024). https://doi.org/10.6028/NIST.FIPS.203
26. NIST (National Institute for Standards and Technology): Stateless Hash-Based Digital Signature Standard. Tech. Rep. Federal Information Processing Standard (FIPS) 205, U.S. Department of Commerce (Aug 2024). https://doi.org/10.6028/NIST.FIPS.205
27. Patarin, J.: The oil and vinegar signature scheme. In: Presented at the Dagstuhl Workshop on Cryptography September 1997 (1997)
28. Pébereau, P.: One vector to rule them all: Key recovery from one vector in UOV schemes. In: Saarinen, M.J., Smith-Tone, D. (eds.) Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Part II. pp. 92–108. Springer, Cham, Switzerland, Oxford, UK (Jun 12–14, 2024). https://doi.org/10.1007/978-3-031-62746-0_5
29. Shor, P.W.: Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. SIAM Journal on Computing **26**(5), 1484–1509 (1997). https://doi.org/10.1137/S0097539795293172

30. Tao, C., Petzoldt, A., Ding, J.: Efficient key recovery for all HFE signature variants. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology – CRYPTO 2021, Part I. Lecture Notes in Computer Science, vol. 12825, pp. 70–93. Springer, Cham, Switzerland, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84242-0_4
31. Wang, L., Chou, C., Ding, J., Kuan, Y., Leegwater, J.A., Li, M., Tseng, B., Tseng, P., Wang, C.: SNOVA. Tech. rep., National Institute of Standards and Technology (2024), available at https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures

A The Plücker embedding

In this section, we will prove the results on the Plücker embedding that we left out in the text. Most results are standard and can be found in, for example, [14,15].

Lemma 5. *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *and define the matrix*

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_k \end{bmatrix} = \begin{bmatrix} x_{11} & \cdots & x_{1n} \\ \vdots & \ddots & \vdots \\ x_{k1} & \cdots & x_{kn} \end{bmatrix}$$

then

$$\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \sum_{J \subset [n], |J|=k} |\mathbf{X}_{*J}| \cdot \mathbf{e}_J.$$

Proof. We prove this by writing the vectors $\mathbf{x}_i$ in terms of basis vectors. Writing out the product, we see that most cross-terms vanish. Then we can group the sum by the basis vectors of $\bigwedge^k \mathbb{F}_q^n$ and the result follows.

$$\begin{aligned} \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k &= \left(\sum_j x_{1j} \mathbf{e}_j \right) \wedge \cdots \wedge \left(\sum_j x_{kj} \mathbf{e}_j \right) \\ &= \sum_{j_1, \cdots, j_k} x_{1j_1} \cdots x_{kj_k} \cdot \mathbf{e}_{j_1} \wedge \cdots \wedge \mathbf{e}_{j_k} \\ &= \sum_{j_1 < \cdots < j_k} \left(\sum_{\sigma \in S_k} \operatorname{sgn}(\sigma) x_{\sigma(1)j_1} \cdots x_{\sigma(k)j_k} \right) \mathbf{e}_{j_1} \wedge \cdots \wedge \mathbf{e}_{j_k} \\ &= \sum_{J \subset [n], |J|=k} |\mathbf{X}_{*J}| \cdot \mathbf{e}_J \end{aligned}$$

□

Corollary 3. $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *is linearly dependent* $\iff \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \mathbf{0}$. □

Corollary 4. *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{F}_q^n$ *and* $\mathbf{y}_1, \ldots, \mathbf{y}_k \in \mathbb{F}_q^n$ *be two linear independent sets of vectors. Then*

$$\langle \mathbf{x}_1, \ldots, \mathbf{x}_k \rangle_{\mathbb{F}_q} = \langle \mathbf{y}_1, \ldots, \mathbf{y}_k \rangle_{\mathbb{F}_q} \iff \exists c \in \mathbb{F}_q^* \text{ s.t. } \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = c \cdot \mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_k.$$

Proof. Consider

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_k \end{bmatrix} \quad \text{and} \quad \mathbf{Y} = \begin{bmatrix} \mathbf{y}_1 \\ \vdots \\ \mathbf{y}_k \end{bmatrix}.$$

Then $\langle \mathbf{x}_1, \dots, \mathbf{x}_k \rangle_{\mathbb{F}_q} = \langle \mathbf{y}_1, \dots, \mathbf{y}_k \rangle_{\mathbb{F}_q}$ if and only if there exists an invertible matrix $\mathbf{S}$ such that $\mathbf{X} = \mathbf{S}\mathbf{Y}$ which implies $|\mathbf{X}_{*J}| = |\mathbf{S}\mathbf{Y}_{*J}| = |\mathbf{S}| \cdot |\mathbf{Y}_{*J}|$. The first implication then follows from Lemma 1 and $c = |\mathbf{S}|$. Conversely, let $\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = c \cdot \mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_k$ and consider a vector $\mathbf{x} \in \langle \mathbf{x}_1, \dots, \mathbf{x}_k \rangle_{\mathbb{F}_q}$ such that $\mathbf{x} \notin \langle \mathbf{y}_1, \dots, \mathbf{y}_k \rangle_{\mathbb{F}_q}$. Then by Corollary 1 we obtain $\mathbf{x} \wedge \mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_k = c^{-1} \cdot \mathbf{x} \wedge \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \mathbf{0}$ which is then a contradiction by Corollary 1 again. □

Lemma 6. *Let $\mathbf{x}_1, \dots, \mathbf{x}_k \in \mathbb{F}_q^n$ be linear independent vectors and write*

$$\mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k = \sum_{J \subset [n], |J| = k} \alpha_J \mathbf{e}_J.$$

Furthermore, let $I = \{i_1, \dots, i_{k-1}\} \subset [n]$, then

$$\mathbf{v} = \sum_j \alpha_{j, i_1, \dots, i_{k-1}} \mathbf{e}_j \in \langle \mathbf{x}_1, \dots, \mathbf{x}_k \rangle_{\mathbb{F}_q}.$$

Proof. Denote $V = \langle \mathbf{x}_1, \dots, \mathbf{x}_k \rangle_{\mathbb{F}_q}$ and let H be the subspace of $\mathbb{F}_q^n$ defined by $H = \langle \mathbf{e}_i : i \notin I \rangle_{\mathbb{F}_q}$ and $h = \dim H \cap V$. Consider a basis

$$\langle \mathbf{y}_1, \dots, \mathbf{y}_k \rangle_{\mathbb{F}_q} = V \text{ such that } \langle \mathbf{y}_1, \dots, \mathbf{y}_h \rangle_{\mathbb{F}_q} = H \cap V$$

so that we can denote

$$\mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_h = \sum_{J' \subset I^c, |J'| = h} \beta_{J'} \mathbf{e}_{J'} \quad \text{and} \quad \mathbf{y}_{h+1} \wedge \cdots \wedge \mathbf{y}_k = \sum_{J'' \subset [n], |J''| = k-h} \gamma_{J''} \mathbf{e}_{J''}.$$

By Corollary 2, there is a $c \in \mathbb{F}_q^*$ such that

$$\begin{aligned} \sum_{J \subset [n], |J| = k} \alpha_J \mathbf{e}_J &= \mathbf{x}_1 \wedge \cdots \wedge \mathbf{x}_k \\ &= c \cdot \mathbf{y}_1 \wedge \cdots \wedge \mathbf{y}_k \\ &= c \cdot \left(\sum_{J' \subset I^c, |J'| = h} \beta_{J'} \mathbf{e}_{J'} \right) \wedge \left(\sum_{J'' \subset [n], |J''| = k-h} \gamma_{J''} \mathbf{e}_{J''} \right) \end{aligned}$$

If $h \geq 2$, then $\alpha_{j, i_1, \dots, i_{k-1}} = 0$ for any j by the wedge decomposition given above, and hence $\mathbf{v} = \mathbf{0}$. Instead, if $h = 1$, we find

$$\alpha_{j, i_1, \dots, i_{k-1}} = c \cdot \beta_{\{j\}} \cdot \gamma_I.$$

And thus $\mathbf{v} = c \cdot \gamma_I \cdot \mathbf{w}_1 \in \langle \mathbf{x}_1, \dots, \mathbf{x}_k \rangle_{\mathbb{F}_q}$. □