






Exploiting SNOVA’s Structure in the Wedge Product Attack

Hung Le^{1,2} , Maxime Bros³ , Jacob Lichtinger⁴ , Brice Minaud², Ray Perlner⁴ ,
Daniel Smith-Tone^{4,5} , and Cristian Valenzuela⁵

¹ LTCI, Telecom Paris, Institut Polytechnique de Paris, France
`thai.le@telecom-paris.fr`

² École normale supérieure, PSL University, CNRS, Inria, France
`{hung.le,brice.minaud}@ens.fr`

³ Associate, National Institute of Standards and Technology (NIST), Maryland, USA
`maxime.bros@nist.gov`

⁴ National Institute of Standards and Technology (NIST), Maryland, USA
`{jacob.lichtinger,ray.perlner,daniel.smith}@nist.gov`

⁵ University of Louisville, Louisville KY, USA
`{dcsmit11,cristian.valenzuela}@louisville.edu`

Abstract. Post-quantum cryptography (PQC) aims to develop cryptographic schemes secure against quantum adversaries. One promising class of digital signature schemes is based on multivariate quadratic equations, with Unbalanced Oil and Vinegar (UOV) being a leading example. UOV has been extensively studied since its introduction in 1999 (Kipnis, Patarin, Goubin, Eurocrypt 1999), and it has remained secure. It offers very small signatures but suffers from very large public keys; to remediate this, some schemes, such as MAYO, QR-UOV, SNOVA, add a structure to reduce the size of the public key. These four multivariate schemes are candidates that made it to the Second Round of NIST PQC Additional Call for Post-Quantum Signature schemes. In this work, we revisit a new algebraic attack proposed recently by Lars Ran at Eurocrypt 2025 Rump Session by showing how to exploit the block-ring structure of SNOVA to reduce the cost of the attack. Our improved attack, which relies on a conjecture (work in progress to confirm it experimentally), improves significantly on the previous one for almost all SNOVA parameters; for instance bringing the security of SNOVA-I $((v, o, \ell) = (24, 5, 4))$ down to 94 bits of security when the previous attack was at 160 bits. A consequence of our attack is that all parameters of SNOVA updated for Round 2 of NIST Standardization are now broken.

Keywords: Post-Quantum Cryptography · UOV · SNOVA · Algebraic Attack.

1 Introduction

In 2016, the National Institute of Standards and Technology (NIST) began the process for standardizing quantum-resistant public-key cryptography. After selecting some algorithms to standardize, NIST opened an additional standardization process [11] to increase the variety of digital signature standards available. A significant number of submissions belong to the family of multivariate public key schemes, with many belonging to the subfamily of Unbalanced Oil and Vinegar (UOV) schemes. A specific secondary goal mentioned in the NIST call for proposals is to select schemes with short signatures.

Most current post-quantum schemes have larger signatures compared to those currently in use such as ECDSA. Among leading approaches, only two families of constructions offer signatures in the order of 100 bytes or less: isogeny-based schemes, and multivariate schemes. Isogeny-based schemes are very attractive, offering both short signatures and short public keys. On the other hand, isogeny-based schemes are still relatively new. As such, their security is not yet fully understood, as illustrated by the attacks that broke SIDH and SIKE, see [16,7].

Unbalanced Oil and Vinegar. By comparison, Unbalanced Oil and Vinegar (UOV) is time-tested: despite many cryptanalysis efforts since its introduction in 1999, its security has stood firm. UOV is a hash-and-sign multivariate signature scheme defined over a field \mathbb{F}_q with $n = v + o$ variables and $m = o$ quadratic equations. Its trapdoor structure splits the variables into *vinegar* (v) and *oil* (o) parts; each public polynomial is quadratic in all variables except that oil-oil terms are absent. Knowing the secret affine transformation

T and the central map F lets the signer invert the public map $P = F \circ T$, while outsiders face an NP-hard MQ system.

A significant drawback to the UOV scheme is its large public key size, which typically ranges in the dozens of kilobytes for compressed key variants. To address this issue, structured variants of UOV have been introduced, such as MAYO, QR-UOV, or SNOVA [4,10,19]. Similar to structured lattices or low-density codes, this additional structure allows for more efficient and compact schemes; but careful analysis is required to ensure that security is not weakened in the process. A notable previous attempt to reduce the public key size of UOV by adding extra structure was Rainbow [9]; however, the extra structure was able to be exploited to break the scheme [2] during the original NIST PQC Standardization process.

The SNOVA scheme. SNOVA is a round-2 candidate in the ongoing NIST standardization process for additional post-quantum signatures, noted for its significant public key size reduction relative to UOV. For example, the smallest public key size achieved by the submitted parameters of SNOVA is 1016 bytes (with a signature size of 248 bytes). SNOVA can be viewed as an augmentation of UOV with additional structure in two different ways: it embeds a MAYO-like “whipping” structure; and it uses a specific block-ring structure as used in QR-UOV. Earlier analysis by Ward Beullens showed that in the initial design document, the whipping structure was vulnerable to attacks [3]. Independently, Cabarcas, Li, Verbel and Villanueva-Polanco showed that the block-ring structure could be exploited to speed up polynomial system solving for relevant SNOVA systems [6].

Our Contributions. In this work, we continue this line of research analyzing how the SNOVA structure interacts with various cryptanalysis techniques. Building on the wedge attack against generic UOV systems by Lars Ran [15], we identify and exploit the block-ring structure that is intrinsic to SNOVA’s *mini-UOV* map (using the terminology of [1]). This insight lowers the *dimension* of the exterior space that must be solved for, from $\binom{\ell(v+o')}{\ell}$ down to $\binom{v+o'}{v}^\ell$. Our improved attack, which relies on a conjecture (work in progress to confirm it experimentally), improves significantly on the previous one for almost all SNOVA parameters; for instance bringing the security of SNOVA-I $((v, o, \ell) = (24, 5, 4))$ to 94 bits of security when the previous attack was at 160 bits. In addition, for one particular set of SNOVA parameters, namely SNOVA-V with $(v, o, \ell) = (60, 10, 4)$, the previous attack [15] did not apply (because there does not exist a value for the variable o' fulfilling the required condition, given by Equation (4), for the attack to work); whereas in our approach we are able to satisfy the required constraints and the attack brings the security below the security level. We also prove the success condition for the generic attack and provide a conjectured condition for the attack applied to SNOVA. Overall, a consequence of our attack is that all parameters of SNOVA updated for Round 2 of NIST Standardization are now broken.

1.1 Notation

- \mathbb{F}_q is the finite field with $q = 2^k$ elements with $k \geq 1$; \mathbb{F}_{q^ℓ} is its degree- ℓ extension.
- For a given field \mathbb{F}_q , the set of matrices with a rows, b columns, and entries in \mathbb{F}_q , is denoted $\mathbb{F}_q^{a \times b}$.
- The general linear group of non singular matrices in $\mathbb{F}_q^{n \times n}$ is denoted $\text{GL}_n(\mathbb{F}_q)$.
- Matrices and vectors are written in boldface font: \mathbf{M} and \mathbf{v} , respectively. By extension, we identify a point \mathbf{p} to its position vector, thus it is also written in boldface font.
- The vector space \mathbb{F}_q^n where $n = v + o$ splits into the vinegar space V and the oil space O , of dimension v and o , respectively; so $\mathbb{F}_q^n = V \oplus O$.

2 Wedge product attack improvement

Let $q = 2^k$ and \mathbb{F}_q be the finite field of characteristic 2. The public key of UOV consists of quadratic maps $p = (p_1, \dots, p_m) : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ where $p_k(\mathbf{x}) = \mathbf{x} \mathbf{P}_k \mathbf{x}^\top$, $\mathbf{P}_k \in \mathbb{F}_q^{n \times n}$ with $1 \leq k \leq m$. The private key of a UOV instance is a secret subspace O on which the public key vanishes. An important map that can be associated to any quadratic map is its polar form, sometimes referred to as the discrete differential.

Definition 1 (Polar form (Discrete Differential)). Let $p = (p_1, \dots, p_m) : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ where each quadratic form $p_k(\mathbf{x})$ is defined by $p_k(\mathbf{x}) = \mathbf{x}^\top \mathbf{P}_k \mathbf{x}$ with $\mathbf{P}_k \in \mathbb{F}_q^{n \times n}$ for $1 \leq k \leq m$.

The polar form Q of p is defined as

$$\begin{aligned} \mathbb{F}_q^n \times \mathbb{F}_q^n &\rightarrow \mathbb{F}_q^m \\ Q(\mathbf{x}, \mathbf{y}) &\mapsto p(\mathbf{x} + \mathbf{y}) - p(\mathbf{x}) - p(\mathbf{y}) \end{aligned} \quad (1)$$

Note that the polar form is a symmetric bilinear map.

An alternative way of defining the polar form is to symmetrize the matrices \mathbf{P}_k . For the specific case of characteristic 2:

$$Q(\mathbf{x}, \mathbf{y}) := (Q_1, Q_2, \dots, Q_m),$$

where each Q_k is defined as

$$Q_k(\mathbf{x}, \mathbf{y}) := \mathbf{x}(\mathbf{P}_k + \mathbf{P}_k^\top)\mathbf{y}^\top. \quad (2)$$

Note that for simplicity of exposition the above definition addresses only the homogeneous case, relevant for any candidate scheme in the UOV subfamily. For the more general definition, applicable to any quadratic map, a constant correction term of $p(\mathbf{0})$ must be added to Equation (1) to achieve bilinearity.

An important observation, specific to characteristic 2, is that the polar form, see Definition 1, is an alternating bilinear form. In Eurocrypt 2025 Rump Session, Lars Ran leveraged this fact and revealed that these public quadratics can be manipulated with exterior algebra to effect a new attack, see [15].

Since the hidden linear transformation that mixes the coordinates of the secret map is the same in every Q_k , there exists a non-zero wedge product

$$V := \mathbf{v}_1 \wedge \dots \wedge \mathbf{v}_v \in \bigwedge^v \mathbb{F}_q^n$$

such that, first, its wedge product with each Q_k vanishes, and second, its orthogonal complement is the oil subspace. (Here, by the orthogonal complement of a vector space V , we mean the space of vectors whose dot product with all vectors in V is zero⁶.)

This property motivates us to build the linear map

$$\Phi : \bigwedge^v \mathbb{F}_q^n \longrightarrow \left(\bigwedge^{v+2} \mathbb{F}_q^n \right)^m, \quad V \xrightarrow{\Phi} (V \wedge Q_1, \dots, V \wedge Q_m). \quad (3)$$

If $\ker(\Phi)$ is one-dimensional, the unique kernel vector mapping the monomial 1 to 1 recovers V , from which an oil-vinegar decomposition follows via Gaussian elimination.

The author also points out in [15] that the algorithm succeeds for those (v, o, m) -UOV instances that satisfy the non-positivity criterion

$$\sum_{i=0}^{\lfloor \frac{o}{2} \rfloor} (-1)^i \binom{m+i-1}{i} \binom{v+o}{v+2i} \leq 0. \quad (4)$$

We prove condition given by Equation (4) in Section 2.2. Whenever it holds, one only needs to find the unique kernel vector of a sparse matrix of size $\binom{v+o}{v} \times \binom{v+o}{v}$ whose row density is $\binom{v+2}{2}$ [14]. Then, let E be the number of linearly independent equations and U be the number of variables, the complexity of the attack is given by $\mathcal{O}\left(\binom{v+2}{2}EU\right)$ operations in \mathbb{F}_q according to Wiedemann algorithm for sparse matrix. This wedge-product view therefore transforms the algebraic cryptanalysis of UOV into a single structured linear problem that can be efficiently solved by modern sparse-matrix solvers.

Applying this attack to SNOVA, the variables in Lars's modeling correspond to the maximal minors of a generic $\ell v \times \ell(v+o)$ matrix, \mathbf{V} , representing the vinegar space (or more properly the dual space of the oil space). This matrix is not generic, however, but is known to have a blockwise structure where each $\ell \times \ell$ block is given by an element of $\mathbb{F}_q[\mathbf{S}]$. This extra structure from the SNOVA construction changes the analysis of the resulting system.

⁶ Although this naming convention is common, note that the orthogonal complement of a subspace may non-trivially intersect the subspace.

First, note that this structure induces linear dependencies among the maximal minors, effectively reducing the number of variables in the system. In particular, for any $\mathbf{\Gamma} \in \mathbb{F}_q[\mathbf{S}]$, we have the relation:

$$\mathbf{\Gamma}^{\otimes v} \cdot \mathbf{V} = \mathbf{V} \cdot \mathbf{\Gamma}^{\otimes(v+o)}.$$

This identity can be used to express the minors of \mathbf{V} given by $(\mathbf{V} \cdot \mathbf{e}_{i_1}) \wedge (\mathbf{V} \cdot \mathbf{e}_{i_2}) \wedge \cdots \wedge (\mathbf{V} \cdot \mathbf{e}_{i_{v\ell}})$ as a linear combination of other minors of \mathbf{V} , specifically:

$$(\mathbf{V} \cdot \mathbf{e}_{i_1}) \wedge (\mathbf{V} \cdot \mathbf{e}_{i_2}) \wedge \cdots \wedge (\mathbf{V} \cdot \mathbf{e}_{i_{v\ell}}) = \det(\mathbf{\Gamma})^{-v} \cdot (\mathbf{V} \cdot \mathbf{\Gamma}^{\otimes(v+o)} \mathbf{e}_{i_1}) \wedge (\mathbf{V} \cdot \mathbf{\Gamma}^{\otimes(v+o)} \mathbf{e}_{i_2}) \wedge \cdots \wedge (\mathbf{V} \cdot \mathbf{\Gamma}^{\otimes(v+o)} \mathbf{e}_{i_{v\ell}}).$$

Second, there are also linear dependencies among the equations themselves. In particular, observe that the equations arising from

$$(\mathbf{S}^a)^{\otimes n} [\mathbf{Q}_i] (\mathbf{S}^a)^{\otimes n}$$

can be obtained as linear combinations of the equations arising from $[\mathbf{Q}_i]$, as follows. The original equations from $[\mathbf{Q}_i]$ are of the form:

$$[\mathbf{Q}_i] \wedge \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_{v\ell} = 0.$$

By applying a linear transformation, we obtain expressions within the span of the component equations corresponding to:

$$(\mathbf{S}^a)^{\otimes n} [\mathbf{Q}_i] (\mathbf{S}^a)^{\otimes n} \wedge (\mathbf{v}_1 (\mathbf{S}^a)^{\otimes n}) \wedge \cdots \wedge (\mathbf{v}_{v\ell} (\mathbf{S}^a)^{\otimes n}) = 0.$$

Using the linear dependencies established earlier (with $\mathbf{\Gamma} = \mathbf{S}^a$), however, we conclude:

$$(\mathbf{S}^a)^{\otimes n} [\mathbf{Q}_i] (\mathbf{S}^a)^{\otimes n} \wedge \mathbf{v}_1 \wedge \cdots \wedge \mathbf{v}_{v\ell} = 0.$$

2.1 Number of variables (linearly independent maximal minors)

For a generic UOV system, the number of linearly independent maximal minors, which is also the number of variables in our attack, is $\binom{v+o}{v}$. There is extra structure built into SNOVA, however, due to its definition utilizing the extension field $\mathbb{F}_{q^\ell} \cong \mathbb{F}_q[\mathbf{S}]$. In other words, one might expect the number of linearly independent minors in SNOVA to be less than $\binom{\ell(v+o)}{\ell v}$.

We did some experiments in Magma⁷ with small parameters, see Table 1, and found the number of linearly independent minors of SNOVA to be lower than the total number of maximal minors. From these data, we hypothesized a formula for the number of linearly independent minors, see the third column in Table 1 which we then proved, yielding Proposition 1.

Table 1. Number of linearly independent minors for small SNOVA instances

Parameters (v, o, ℓ)	# l.i. minors	# from Prop. 1	Total # max minors
(2, 2, 2)	36	36	70
(2, 1, 3)	27	27	84
(2, 2, 3)	216	216	924
(3, 1, 2)	16	16	28
(2, 1, 4)	81	81	495
(2, 5, 2)	441	441	1001
(2, 10, 2)	4356	4356	10626

⁷ Certain commercial equipment, instruments, or materials (or suppliers, or software, ...) are identified in this paper to foster understanding. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

Proposition 1. *Let $\mathbf{S} \in \mathbb{F}_q^{\ell \times \ell}$ be a symmetric matrix such that $\mathbb{F}_q[\mathbf{S}]$ is a field, isomorphic to \mathbb{F}_{q^ℓ} . In particular, \mathbf{S} induces an isomorphism between \mathbb{F}_q^ℓ and \mathbb{F}_{q^ℓ} as \mathbb{F}_q -vector spaces. Given this isomorphism, a matrix $\mathbf{M} \in \mathbb{F}_{q^\ell}^{a \times b}$ induces a matrix $\phi(\mathbf{M}) \in \mathbb{F}_q^{\ell a \times \ell b}$. We are interested in the space $A = \phi(\mathbb{F}_q^{(v+o) \times v}) \subseteq \mathbb{F}_q^{\ell(v+o) \times \ell v}$. Our goal is to count the number of linearly independent maximal minors of matrices in A . More formally, letting $m = \binom{\ell(v+o)}{\ell v}$, we want to compute the dimension of $\text{span}\langle f(\mathbf{M}) : \mathbf{M} \in A \rangle$, where $f : A \rightarrow \mathbb{F}_q^m$ maps the matrix \mathbf{M} to the vector of its maximal minors, taken in a fixed arbitrary order. (The map f essentially amounts to taking the Plücker coordinates of the input matrix.)*

Out of all the m possible maximal minors of matrices in A , we show that the number of linearly independent minors is:

$$\binom{v+o}{v}^\ell.$$

Proof. We recall that a square matrix over a finite field with an irreducible characteristic polynomial is diagonalizable over the splitting field of the characteristic polynomial [13]. Thus, for the matrix $\mathbf{S} \in \mathbb{F}_q^{\ell \times \ell}$ defined in the specification [19], there exists a matrix $\mathbf{B} \in \text{GL}_\ell(\mathbb{F}_{q^\ell})$ such that $\mathbf{B}^{-1}\mathbf{S}\mathbf{B}$ is a diagonal matrix in $\mathbb{F}_{q^\ell}^{\ell \times \ell}$. In particular, all elements of $\mathbb{F}_q[\mathbf{S}]$ are simultaneously diagonalizable with this matrix \mathbf{B} .

Let \mathbf{U} be the diagonal block matrix with n copies of the matrix \mathbf{B} , namely,

$$\mathbf{U}_{\ell(v+o)} = \mathbf{I}_n \otimes \mathbf{B} = \begin{pmatrix} \mathbf{B} & & \\ & \ddots & \\ & & \mathbf{B} \end{pmatrix} \in (\mathbb{F}_{16^\ell})^{\ell n \times \ell n}.$$

The block components of the matrix \mathbf{V} which is an element of $(\mathbb{F}_q[\mathbf{S}])^{\ell(v+o) \times \ell(v+o)}$ are diagonalized with $\mathbf{U}_{\ell v}$ and $\mathbf{U}_{\ell(v+o)}$ as follows:

$$\mathbf{V}_1 = \mathbf{U}_{\ell v} \mathbf{V} \mathbf{U}_{\ell(v+o)}^\top = \begin{pmatrix} \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array} & \cdots & \begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array} \\ \vdots & \ddots & \vdots \\ \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array} & \cdots & \begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array} \\ \vdots & \ddots & \vdots \\ \begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array} & \cdots & \begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array} \end{pmatrix}$$

Let $\Delta_I(\mathbf{V})$ denote the maximal minors of \mathbf{V} indexed by column subsets $I \subset \{1, \dots, \ell(v+o)\}$ of size ℓv , and define the vector of minors

$$\Delta(\mathbf{V}) := (\Delta_I(\mathbf{V})) \in \mathbb{F}_q^N, \quad \text{with } N = \binom{\ell(v+o)}{\ell v}.$$

Multiplying to the left by $\mathbf{U}_{\ell v}$ gives

$$\Delta_I(\mathbf{U}_{\ell v} \mathbf{V}) = \det(\mathbf{U}_{\ell v}) \cdot \Delta_I(\mathbf{V}) \quad \text{for all } I,$$

so all minors are rescaled by the same nonzero constant $\det(\mathbf{U}_{\ell v}) \in \mathbb{F}_q^\times$.

Right multiplication changes the columns: if $v_1, \dots, v_{\ell(v+o)}$ are the columns of \mathbf{V} , then the columns of \mathbf{V}_1

are $w_j = \sum_{k=1}^{\ell(v+o)} v_k \mathbf{U}_{n, kj}^\top$. This yields:

$$\Delta_I(\mathbf{V}_1) = \sum_{J, |J|=\ell v} \Delta_J(\mathbf{V}) \cdot \det(\mathbf{U}_{n, J, I}^\top)$$

This transformation is linear and invertible because \mathbf{U}_n^I is invertible, so the coefficient matrix $\mathbf{M} = (\det(\mathbf{U}_{n,J,I}^T))_{J,I}$ is also invertible. Thus,

$$(\Delta_I(\mathbf{V}))_I = \mathbf{M} \cdot (\Delta_I(A))_I.$$

Since \mathbf{M} is invertible, the linear span of the minors remains unchanged. Therefore, diagonalizing V over the extension field does *not* change the number of linearly independent maximal minors over the base field.

Because every $\ell \times \ell$ block in \mathbf{V}_1 is itself diagonal, the matrix \mathbf{V}_1 is block-diagonal when viewed in a different basis. In other words, there exist two permutation matrices \mathbf{P} and \mathbf{Q} such that

$$\mathbf{V}_2 = \mathbf{P}\mathbf{V}_1\mathbf{Q} = \text{diag}(\mathbf{B}_1, \mathbf{B}_2, \dots, \mathbf{B}_\ell) = \begin{pmatrix} \boxed{\mathbf{B}_1} & & \cdots & \\ & \boxed{\mathbf{B}_2} & \cdots & \\ & \vdots & \ddots & \vdots \\ & & & \boxed{\mathbf{B}_\ell} \end{pmatrix}$$

Since \mathbf{P} and \mathbf{Q} are permutation matrices, they only reorder the rows and columns of \mathbf{V}_1 and do not introduce any new linear dependencies among its maximal minors. As a result, the number of linearly independent maximal minors of \mathbf{V}_1 remains the same in \mathbf{V}_2 .

In the matrix \mathbf{V}_2 , each block \mathbf{B}_i contributes to the maximal minors. To form a non-zero maximal minor, we must select exactly v columns from the $v + o$ available columns in each \mathbf{B}_i ; selecting fewer or more would result in a minor with zero determinant due to the structure of \mathbf{V}_2 .

Moreover, because each \mathbf{B}_i occupies a distinct set of rows in \mathbf{V}_2 , the choices of columns in different blocks are independent of each other. Therefore, the total number of linearly independent maximal minors in \mathbf{V}_2 is:

$$\binom{v + o}{v}^\ell.$$

□

When applying the results from Proposition 1 to the parameters of SNOVA, it saves about 2.6 – 9.8 bits in variable count (see Table 2).

Table 2. Comparison of the number (given as a logarithm in base 2) of linearly independent (l.i.) minors vs. total number of minors across SNOVA instances.

Variants	v	o	q	ℓ	# l.i. minors	# minors
SNOVA-I	37	17	16	2	91	93
	25	8	16	3	71	76
	24	5	16	4	67	74
SNOVA-III	56	25	16	2	138	141
	49	11	16	2	115	120
	37	8	16	4	111	118
	24	5	16	5	85	95
SNOVA-V	75	33	16	4	185	188
	66	15	16	4	159	164
	60	10	16	4	154	159
	29	6	16	5	103	112

2.2 Number of linearly independent equations

Before going further, we need to recall the definition of a symmetric tensor, then we give a lemma which will be used to prove condition given by Equation (4).

Definition 2 (Tensor, Symmetric Tensor (as in [5])). Let m, b be integers such that $2 \leq b \leq m$. A tensor of dimension m and order b over \mathbb{F}_q is a set

$$(\mathbf{S}_{i_1, i_2, \dots, i_b})_{1 \leq i_1, i_2, \dots, i_b \leq m} \in \mathbb{F}_q^{m^b}.$$

Such a tensor is called a symmetric tensor of dimension m and order b over \mathbb{F}_q if for any permutation σ in the symmetric group \mathcal{S}_b , one has

$$\mathbf{S}_{i_1, i_2, \dots, i_b} = \mathbf{S}_{i_{\sigma(1)}, i_{\sigma(2)}, \dots, i_{\sigma(b)}}.$$

In other words, a symmetric tensor is a tensor which is invariant under the action of the symmetric group on its indices.

Lemma 1. Let $\mathbf{S}_{k_1 \dots k_i}$ be a symmetric tensor of dimension m and order i over \mathbb{F}_q . Let $\mathbf{V} = \mathbf{v}_1 \wedge \dots \wedge \mathbf{v}_v \in \Lambda^v(\mathbb{F}_q^n)$ be a v -form. For all $k_i \in \{1, \dots, m\}$, let each quadratic form $\mathbf{Q}_{k_i} \in \Lambda^2(\mathbb{F}_q^n)$ be an alternating form, then there are

$$\binom{m+i-1}{i} \binom{v+o}{v+2i}$$

relations of the form

$$\sum_{k_1, \dots, k_i} \mathbf{S}_{k_1 \dots k_i} (\mathbf{V} \wedge \mathbf{Q}_{k_1} \wedge \dots \wedge \mathbf{Q}_{k_i}) = \mathbf{0}.$$

Proof. We define $\mathbf{W}_{\mathbf{k}} = \mathbf{A}_{k_1 \dots k_i} := \mathbf{V} \wedge \mathbf{Q}_{k_1} \wedge \dots \wedge \mathbf{Q}_{k_i} \in \Lambda^{v+2i}(\mathbb{F}_q^n)$ for a tuple $\mathbf{k} = (k_1, \dots, k_i)$. Note that $\mathbf{W}_{\mathbf{k}}$ is totally antisymmetric in the k -indices:

$$\mathbf{W}_{\dots, k_t, k_{t+1}, \dots} = -\mathbf{W}_{\dots, k_{t+1}, k_t, \dots} \quad \text{and} \quad \mathbf{W}_{\dots, k, k, \dots} = \mathbf{0}.$$

Let $\mathbf{S}_{k_1 \dots k_i}$ is a symmetric tensor of order i over the index set $\{1, \dots, m\}$. The contraction of a symmetric tensor \mathbf{S} with an antisymmetric tensor \mathbf{A} is zero:

$$\sum_{k_1, \dots, k_i} \mathbf{S}_{k_1 \dots k_i} \mathbf{A}_{k_1 \dots k_i} = 0,$$

because if any index repeats, $\mathbf{A}_{k_1 \dots k_i} = 0$. If all are distinct, $\mathbf{S}_{k_1 \dots k_i} = \mathbf{S}_{k_{\sigma(1)} \dots k_{\sigma(i)}}$, whereas $\mathbf{A}_{k_1 \dots k_i} = \text{sign}(\sigma) \mathbf{A}_{k_{\sigma(1)} \dots k_{\sigma(i)}}$; thus, all terms cancel in pairs.

Because the \mathbb{F}_q -vector space of symmetric tensors of dimension m and order i over \mathbb{F}_q is isomorphic to the vector space of homogeneous polynomials of degree i in m variables over \mathbb{F}_q , it is of dimension $\binom{m+i-1}{i}$, hence the result. \square

Using Lemma 1, let i be an integer such that $1 \leq i < \frac{o}{2}$, we expect the number of linearly independent equations for a UOV system of m equations is

$$\sum_{i=1}^{\lfloor \frac{o}{2} \rfloor} (-1)^{i+1} \binom{m+i-1}{i} \binom{v+o}{v+2i}.$$

This means as long as this value is not smaller than the total number of monomials, i.e.,

$$\begin{aligned} & \sum_{i=1}^{\lfloor \frac{o}{2} \rfloor} (-1)^{i+1} \binom{m+i-1}{i} \binom{v+o}{v+2i} \geq \binom{v+o}{v} \\ \Leftrightarrow & \sum_{i=0}^{\lfloor \frac{o}{2} \rfloor} (-1)^i \binom{m+i-1}{i} \binom{v+o}{v+2i} \leq 0 \end{aligned}$$

we expect to find the unique kernel vector of a sparse matrix of size $\binom{v+o}{v}$ with density $\binom{v+2}{2}$.

We propose two hypotheses for the number of independent linear equations in the SNOVA system. These hypotheses differ only in the scaling ℓ of a binomial coefficient.

Hypothesis 1 (ℓ -Scaled Form):

$$E_d^{(1)} := \sum_{i=1}^{\lfloor \frac{o\ell}{2} \rfloor} \sum_{\substack{a_0, a_1, \dots, a_{\ell-1} \geq 0 \\ a_0 + a_1 + \dots + a_{\ell-1} = 2i}} (-1)^{i+1} \binom{m\ell + i - 1}{i} \prod_{j=0}^{\ell-1} \binom{v+o}{v+a_j}$$

Hypothesis 2 (Unscaled Form):

$$E_d^{(2)} := \sum_{i=1}^{\lfloor \frac{o\ell}{2} \rfloor} \sum_{\substack{a_0, a_1, \dots, a_{\ell-1} \geq 0 \\ a_0 + a_1 + \dots + a_{\ell-1} = 2i}} (-1)^{i+1} \binom{m+i-1}{i} \prod_{j=0}^{\ell-1} \binom{v+o}{v+a_j}$$

Therefore, let $U_d := \binom{v+o}{v}^\ell$ with d being the smallest positive integer such that $1 \leq d < \frac{o\ell}{2}$ and $U_d - 1 \leq E_d$, then the complexity of finding a unique kernel vector is in

$$\mathcal{O} \left(\min \left(E_d^{(i)} U_d^{\omega-1}, \tau E_d^{(i)} U_d \right) \right) \quad \text{where } i = 1 \text{ or } 2 \text{ reflects the hypotheses used for } E_d$$

operations in \mathbb{F}_q , where ω is the linear algebra constant and $\tau = \binom{v+2}{2}^d$ is the number non-zero entries in each row. The minimum in complexity comes from the use of the Strassen algorithms [18] or Wiedemann [8,12,17]. Note that we can delete (project away) some oil coordinates and keep the v vinegar coordinates together with the remaining $o' < o$ oil coordinates. Once an o' -dimensional subspace of the true oil space is known, the rest is recovered efficiently. However, there is a limit to projecting down as the attack can no longer distinguish the genuine oil subspace.

Table 3. Complexity (in bits) of our attack against SNOVA and comparison with the attack in [15].

Variants	v	o	o' (conj. 1)	o' (conj. 2)	ℓ	Our attack	Our attack	Attack in
						(conj. 1)	(conj. 2)	[15]
SNOVA-I ($q = 16$)	37	17	5	6	2	146	145	127
	25	8	3	5	3	103	104	123
	24	5	2	4	4	94	93	160
SNOVA-III ($q = 16$)	56	25	6	7	2	217	217	174
	49	11	4	7	3	164	164	249
	37	8	3	5	4	147	149	229
	24	5	2	4	5	109	110	172
SNOVA-V ($q = 16$)	75	33	6	9	2	288	288	225
	66	15	5	8	3	222	223	273
	60	10	4	7	4	204	204	-
	29	6	2	4	5	130	133	200

Overall, Table 3 summarizes our work by giving the complexity of our improved attack in comparison to the previous one by Lars Ran in [15]. The results are categorized by security levels of SNOVA under the names of SNOVA-I, SNOVA-III, and SNOVA-V. On all these parameters, our attack always worked with solving degree of $d = 1$.

References

1. Alagic, G., Bros, M., Ciadoux, P., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y.K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Silberg, H., Smith-Tone, D., Waller, N.: Status report on the first round of the additional digital signature schemes for the nist post-quantum cryptography standardization process. NIST IR **8528** (2024)
2. Beullens, W.: Breaking Rainbow takes a weekend on a laptop. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology - CRYPTO 2022 - 42nd Annual International Cryptology Conference, CRYPTO 2022, Santa Barbara, CA, USA, August 15-18, 2022, Proceedings, Part II. Lecture Notes in Computer Science, vol. 13508, pp. 464–479. Springer (2022). https://doi.org/10.1007/978-3-031-15979-4_16, https://doi.org/10.1007/978-3-031-15979-4_16
3. Beullens, W.: Improved cryptanalysis of SNOVA. In: Advances in Cryptology - EUROCRYPT 2025. Lecture Notes in Computer Science, vol. 15606, pp. 277–293. Springer (2025)
4. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO: Specification document – round 2. Technical Report Round 2 (March 2025), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/mayo-spec-round2-web.pdf>, available at NIST CSRC
5. Bros, M.: Algebraic cryptanalysis and contributions to Post-Quantum Cryptography based on error-correcting codes in the rank-metric. Ph.D. thesis, Université de Limoges (2022)
6. Cabarcas, D., Li, P., Verbel, J.A., Villanueva-Polanco, R.: Improved attacks for SNOVA by Exploiting Stability under a Group action. In: CRYPTO (2025)
7. Castryck, W., Decru, T.: An efficient key recovery attack on SIDH. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 423–447. Springer (2023)
8. Coppersmith, D.: Solving homogeneous linear equations over $GF(2)$ via block Wiedemann algorithm. Mathematics of Computation **62**(205), 333–350 (1994)
9. Ding, J., Chen, M.S., Petzoldt, A., Schmidt, D., Yang, B.Y., Kannwischer, M., Patarin, J.: Rainbow: Proposal for NISTPQC: Digital signature algorithms version 3.0. Technical Report Round 3 (October 2020), <https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>, available at NIST CSRC
10. Furue, H., Ikematsu, Y., Hoshino, F., Takagi, T., Kosuge, H., Yamakoshi, K., Akiyama, R., Nakamura, S., Orihara, S., Kinjo, K.: QR-UOV. Technical Report Round 2 (March 2025), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/qruov-spec-round2-web.pdf>, available at NIST CSRC
11. Group, C.T.: Call for Additional Digital Signature Schemes for the Post-quantum Cryptography Standardization Process. NIST CSRC (2022), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf>
12. Kaltofen, E.: Analysis of Coppersmith's block Wiedemann algorithm for the parallel solution of sparse linear systems. Mathematics of Computation **64**(210), 777–806 (1995)
13. Nakamura, S., Tani, Y., Furue, H.: Lifting approach against the SNOVA scheme. IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences p. 2024EAP1124 (2025)
14. Ran, L.: Round 2 (additional signatures) official comment: UOV, MAYO, SNOVA. <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/NNv4C5PRIjs> (May 2025), posted to the `pqc-forum@list.nist.gov` mailing list
15. Ran, L.: Wedges, Oil, and Vinegar: A New Algorithm for UOV in Characteristic 2. Rump Session, Eurocrypt 2025 (2025), <https://larsmath.github.io/assets/pdf/rumpsessioneurocrypt25.pdf>, presented May 6, 2025
16. Robert, D.: Breaking SIDH in polynomial time. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 472–503. Springer (2023)
17. Villard, G.: A study of Coppersmith's block Wiedemann algorithm using matrix polynomials. Citeseer (1997)
18. Von Zur Gathen, J., Gerhard, J.: Modern computer algebra. Cambridge University Press (2003)
19. Wang, L.C., Chou, C.Y., Ding, J., Kuan, Y.L., Leegwater, J.A., Li, M.S., Tseng, B.S., Tseng, P.E., Wang, C.C.: SNOVA: Proposal for NISTPQC: Additional digital signature schemes version 2.0. Technical Report Round 2 (March 2025), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/snova-spec-round2-web.pdf>, available at NIST CSRC