

Attacking UOV-based Signatures over divided power algebras

Peigen Li¹ and Siyong Tao²

¹ Beijing Institute of Mathematical Sciences and Applications
lpg22@bimsa.cn

² Department of Mathematical Sciences, Tsinghua University
taosy22@mails.tsinghua.edu.cn

Abstract. Merz and Ran derive additional equations from divided powers of the public polar forms in characteristic two. We extend this construction to finite fields of arbitrary characteristic and apply it to the security analysis of QR-UOV.

Keywords: UOV problem · divided power algebras

1 Introduction

The Unbalanced Oil and Vinegar (UOV) signature scheme [8] hides an oil space on which its public quadratic map vanishes. A basis of this space serves as a signing trapdoor: it reduces the search for a preimage of a target to solving linear systems [1]. Recovering the oil space from the public map is therefore a central objective of structural cryptanalysis. Wedge attacks approach this problem through linear algebra in suitable graded algebras.

Ran [10] introduced the wedge attack using the exterior algebra in characteristic two. Subsequent work developed formulations in reduced symmetric algebras [7] and truncated polynomial rings [4].

These approaches share the ideal inclusion $I \subseteq J$, where I is generated by the public quadratic relations and J by linear forms defining the oil space. In a suitable graded algebra, multiplication by monomials assembles the public relations of a fixed degree into a Macaulay matrix. Its right kernel contains a subspace determined by the oil space, but may be larger than this subspace. Adding further public relations that remain in J can reduce the kernel while preserving the information needed for recovery.

Merz and Ran [9] extend the exterior algebra attack, previously applicable to only a few parameter sets, in two directions. First, in characteristic two, they derive additional equations from divided powers of the public polar forms. Second, they extend the linear algebra computations from a single exterior power to arbitrary Schur modules.

Divided powers provide analogues of $f^k/k!$ when division by $k!$ is unavailable. In the attack, their role is to produce additional homogeneous relations from the public quadratics. The key algebraic requirement is that these relations remain

in the hidden ideal J . Their usefulness also depends on the degree at which the kernel reveals the oil space and on the size and sparsity of the resulting linear systems. This connects the divided power construction to questions about graded quotient dimensions and the cost of recovery.

Our Contribution Building on their construction in [9, §3], we develop divided powers in p -truncated polynomial rings over finite fields of arbitrary characteristic, obtaining additional public relations that remain in the hidden ideal J . Merz and Ran also implicitly use a divided power structure (over a field of characteristic odd) on a multigraded exterior algebra in [9, §5.3]. Their construction supports a Schur–Macaulay attack, while ours focuses on enlarging the public ideal and analyzing recovery criteria and costs in the truncated ring.

This framework directly handles systems over any fields, including QR-UOV. Under Conjecture 3.7, it also enables recovery for the two examples in Section 4 with more than twice as many vinegar variables as equations ($v > 2m$). The classical wedge attack does not apply to these instances, and our comparisons predict lower costs than for the generalized wedge attack of Ikematsu and Furue [6] for these selected parameters.

As discussed in [9, Remark 3.8], the divided power approach underlying our method and the method of Ikematsu and Furue follow distinct strategies and are effective in different parameter regimes. We therefore do not claim that our method is generally superior to theirs: the improvements reported here are limited to the selected parameter sets. Our work should be viewed as complementary to both the construction in [9, §3] and the wedge attack in [6].

The implementation of the experiments, complexity estimators and toy examples can be found in the repository below:

<https://github.com/Yulan24/divided-power-attack>

Outline §2 recalls UOV and the wedge attack. §3 develops the recovery criterion, the divided power construction, and the degree and cost estimates. §4 and §5 are reserved for applications and future directions.

2 Preliminaries

2.1 Unbalanced Oil and Vinegar (UOV)

The UOV signature scheme was introduced by Kipnis, Patarin, and Goubin in 1999 [8]. We follow Beullens’ description of UOV in terms of its hidden oil space [1]. Let $n = v + o$ with $v > o$. The public key is a homogeneous quadratic map $\mathcal{P} = (p_1, \dots, p_m) : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on a hidden o -dimensional subspace $\mathcal{O}$, called the oil space. In the standard signature scheme, $m = o$, and a basis of $\mathcal{O}$ serves as the signing trapdoor. Define the bilinear polar map by

$$\mathcal{P}'(\mathbf{x}, \mathbf{y}) := \mathcal{P}(\mathbf{x} + \mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y}) = (p'_1(\mathbf{x}, \mathbf{y}), \dots, p'_m(\mathbf{x}, \mathbf{y})).$$

Since $\mathcal{O}$ is a linear subspace, we have

$$\mathcal{P}(\mathbf{o}) = 0, \quad \mathcal{P}'(\mathbf{o}, \mathbf{o}') = 0 \quad \text{for all } \mathbf{o}, \mathbf{o}' \in \mathcal{O}. \quad (2.1)$$

To sign a target $\mathbf{t} \in \mathbb{F}_q^m$ obtained by hashing the message, choose a random $\mathbf{v} \in \mathbb{F}_q^n$ and seek $\mathbf{o} \in \mathcal{O}$ such that $\mathcal{P}(\mathbf{v} + \mathbf{o}) = \mathbf{t}$. By the definition of the polar form and Eq. (2.1),

$$\mathcal{P}(\mathbf{v} + \mathbf{o}) = \underbrace{\mathcal{P}(\mathbf{v})}_{\text{fixed by the choice of } \mathbf{v}} + \underbrace{\mathcal{P}(\mathbf{o})}_{=0} + \underbrace{\mathcal{P}'(\mathbf{v}, \mathbf{o})}_{\text{linear function of } \mathbf{o}} = \mathbf{t}. \quad (2.2)$$

Equivalently,

$$\mathcal{P}'(\mathbf{v}, \mathbf{o}) = \mathbf{t} - \mathcal{P}(\mathbf{v}). \quad (2.3)$$

For fixed $\mathbf{v}$, this is a linear system in the o coordinates of $\mathbf{o}$ with respect to a secret basis of $\mathcal{O}$. There are $m = o$ equations, so a nonsingular system gives a unique solution; if no solution exists, choose a new $\mathbf{v}$ and retry. The resulting signature vector is $\mathbf{s} = \mathbf{v} + \mathbf{o}$, and verification checks $\mathcal{P}(\mathbf{s}) = \mathbf{t}$. Thus recovering the oil space provides sufficient information to sign. For our cryptanalytic setting, we allow arbitrary m and consider the following recovery problem.

Definition 2.1 (UOV problem). *Let $v > o$ and m be positive integers, and set $n = v + o$. Given a homogeneous quadratic map $(p_1, \dots, p_m) : \mathbb{F}_q^n \rightarrow \mathbb{F}_q^m$ that vanishes on a subspace $\mathcal{O} \subset \mathbb{F}_q^n$ of dimension o , called the oil space, recover $\mathcal{O}$.*

Remark 2.2. We assume that the oil space $\mathcal{O}$ is unique under the assumption that $v \leq \frac{m(o-1)}{2}$. In the SNOVA setting, the uniqueness condition may differ due to its additional algebraic structure.

2.2 Wedge attack

Ran [10] introduced the wedge attack on UOV over fields of characteristic two. Its starting point is that the polar forms $q_i := p'_i$ are alternating and vanish on $\mathcal{O} \times \mathcal{O}$ by Eq. (2.1). Let $E = (\mathbb{F}_q^n)^*$ and let $L_1, \dots, L_v$ be a basis of the annihilator of $\mathcal{O}$ in E . Viewing the q_i as elements of $\bigwedge^2 E$, the nonzero decomposable vector $\Omega = L_1 \wedge \dots \wedge L_v$ lies in the kernel of

$$\mathcal{M} : \bigwedge^v E \longrightarrow \left(\bigwedge^{v+2} E \right)^m, \quad \omega \longmapsto (q_1 \wedge \omega, \dots, q_m \wedge \omega).$$

If this kernel is one-dimensional, linear algebra recovers Ω up to a nonzero scalar, and hence the hidden oil space.

Jin et al. [7] developed a reduced symmetric-algebra formulation that also applies in odd characteristic. Furue and Ikematsu [4] reformulated these attacks using p -truncated polynomial rings and extended them to p^ℓ -truncated rings, where $p = \text{char}(\mathbb{F}_q)$, including their use in reconciliation and intersection attacks. Wedge methods were also adapted to SNOVA by exploiting its additional algebraic structure: Bros et al. [2] used its block-ring structure, while Tseng et al. [13] studied the wedge map through a lifting to an extension field.

3 Extending to the Divided Power Algebra

A common idea behind the wedge-attack formulations in [10, 7, 4] is to study the ideal inclusion

$$I := \langle p_1, \dots, p_m \rangle_A \subseteq J := \langle L_1, \dots, L_v \rangle_A \quad (3.1)$$

in different graded algebras A , such as exterior algebras and truncated polynomial rings; see also [3]. Here the L_i are linearly independent linear forms whose common kernel is the oil space. In these settings, the attack seeks a degree $d > 0$ for which $I_d = J_d$ and $(A/J)_d \neq 0$, allowing the linear generators of J to be recovered from I_d . This suggests enlarging I with further publicly computable homogeneous relations that remain in J .

Merz and Ran [9, §3] construct such additional relations using divided powers in characteristic two. Their paper contains a broader range of results; the present article focuses specifically on extending the construction in [9, §3] to finite fields of arbitrary characteristic.

Throughout this section, let $p = \text{char}(\mathbb{F}_q)$ and work in the graded ring

$$R := \mathbb{F}_q[x_1, \dots, x_n] / \langle x_1^p, \dots, x_n^p \rangle, \quad \mathfrak{m} := \langle x_1, \dots, x_n \rangle_R.$$

We use the same notation for a polynomial and its class in R . From now on, I and J denote the corresponding ideals in R . Set $o = n - v > 0$ and $N = n(p - 1)$, the largest nonzero homogeneous degree of R .

The central construction is an enlarged public ideal $\gamma(I)$ satisfying

$$I \subseteq \gamma(I) \subseteq J. \quad (3.2)$$

We first establish a recovery method for homogeneous ideals $K \subseteq J$ under a dimension condition. We then construct $\gamma(I)$ using divided powers and apply this method to it. Finally, a conjectural Hilbert-series estimate predicts the degrees at which this condition may hold, and a bound on the matrix row supports gives an estimate of the sparse linear-algebra cost.

3.1 Oil-Space Recovery from Macaulay Kernels

To make clear what the additional equations must achieve, we first give a recovery criterion for a general known homogeneous ideal

$$K := \langle f_1, \dots, f_s \rangle_R \subseteq J := \langle L_1, \dots, L_v \rangle_R,$$

where $\deg f_i = d_i \geq 1$. This includes both the original ideal I and its enlargement $\gamma(I)$ constructed below. We write $R_d = 0$ for $d < 0$ or $d > N$. For an $\mathbb{F}_q$ -vector space V , let $V^\vee = \text{Hom}_{\mathbb{F}_q}(V, \mathbb{F}_q)$.

Multiplying the generators by monomials collects all known relations of a fixed degree. For $0 \leq d \leq N$, this gives the Macaulay map

$$\mathcal{M}_d : \bigoplus_{i=1}^s R_{d-d_i} \longrightarrow R_d, \quad (u_i)_{i=1}^s \longmapsto \sum_{i=1}^s u_i f_i.$$

Its image is K_d . The dual map $\mathcal{M}_d^\vee : R_d^\vee \rightarrow \bigoplus_{i=1}^s R_{d-d_i}^\vee$ is represented by the Macaulay matrix M_d : its columns are indexed by the monomial basis of R_d , its rows by pairs (i, x^β) with $x^\beta \in R_{d-d_i}$, and its entries are

$$(M_d)_{(i, x^\beta), x^\alpha} = [x^\alpha](x^\beta f_i).$$

Thus $F \in R_d^\vee$ belongs to $\ker \mathcal{M}_d^\vee$ if and only if $M_d(F(x^\alpha))_{x^\alpha \in R_d} = 0$.

To describe the information in this kernel, we identify functionals on R_d with polynomials of complementary degree. Let $\omega = (p-1, \dots, p-1)$. Multiplication induces the perfect pairing [6, Lemma 1(ii)]

$$\Phi_d : R_d \times R_{N-d} \longrightarrow \mathbb{F}_q, \quad (f, g) \longmapsto [x^\omega](fg). \quad (3.3)$$

This pairing induces the isomorphism

$$R_d \xrightarrow{\sim} \text{Hom}_{\mathbb{F}_q}(R_{N-d}, \mathbb{F}_q), \quad f \longmapsto (g \longmapsto [x^\omega](fg)). \quad (3.4)$$

Replacing d by $N-d$ also gives $R_{N-d} \cong R_d^\vee$, which we use below. For a homogeneous ideal $H \subseteq R$, define

$$H^\perp := \text{Ann}_R(H) = \{g \in R : gh = 0 \text{ for all } h \in H\}, \quad H_t^\perp := H^\perp \cap R_t.$$

Under the pairing in Eq. (3.3), we have

$$K_{N-d}^\perp = \{g \in R_{N-d} : gK_d = 0\} \cong (R_d/K_d)^\vee = \ker \mathcal{M}_d^\vee. \quad (3.5)$$

To justify the first equality, if $gK_d = 0$ and $d_i \leq d$, then gf_i pairs to zero with every element of R_{d-d_i} , so $gf_i = 0$ by the perfect pairing. If $d_i > d$, the same conclusion holds because $\deg(gf_i) > N$. Since $K \subseteq J$, the following diagram has an exact bottom row:

$$\begin{array}{ccccccc} & & J_{N-d}^\perp & & & & \\ & & \downarrow & & & & \\ 0 & \longrightarrow & K_{N-d}^\perp & \hookrightarrow & R_d^\vee & \xrightarrow{\mathcal{M}_d^\vee} & \bigoplus_{i=1}^s R_{d-d_i}^\vee \end{array} \quad (3.6)$$

The subspace $J_{N-d}^\perp$ depends only on the hidden linear forms. The next lemma describes it explicitly.

Lemma 3.1. *Let $M = (L_1 \cdots L_v)^{p-1}$. Then $J^\perp = \langle M \rangle_R$. In particular, $J_{N-d}^\perp = \langle M \rangle_{N-d} = 0$ if $d > o(p-1)$.*

Proof. Complete $L_1, \dots, L_v$ to a basis $y_1 = L_1, \dots, y_v = L_v, y_{v+1}, \dots, y_n$ of R_1 . An invertible linear change of variables preserves the ideal generated by the p th powers, so $R \cong \mathbb{F}_q[y_1, \dots, y_n]/\langle y_1^p, \dots, y_n^p \rangle$. In this monomial basis, $y_i g = 0$ holds precisely when every monomial appearing in g has exponent $p-1$ in y_i . Thus g annihilates J precisely when it is a multiple of $y_1^{p-1} \cdots y_v^{p-1} = M$. The degree assertion follows from $\deg M = v(p-1)$. $\square$

We can now compare the dimension of the computed kernel with that of this hidden subspace. Write

$$T(r, p, d) := [t^d](1 + t + \cdots + t^{p-1})^r, \quad T(r, p, d) = 0 \quad \text{if } d < 0.$$

Then $\dim R_d = T(n, p, d)$ and $\dim(R_d/J_d) = T(o, p, d)$. When the two quotient dimensions agree and are positive, the entire computed kernel is determined by the oil space, allowing recovery.

Theorem 3.2. *If, for some integer $d > 0$,*

$$\dim(R_d/K_d) = \dim(R_d/J_d) > 0, \quad (3.7)$$

then the oil space $\mathcal{O} = \{\mathbf{z} \in \mathbb{F}_q^n : L_i(\mathbf{z}) = 0 \text{ for all } i\}$ can be recovered from the right kernel of M_d .

Proof. The inclusion $K_d \subseteq J_d$ and the dimension equality give $K_d = J_d$, hence $K_{N-d}^\perp = J_{N-d}^\perp$. Moreover, $1 \leq d \leq o(p-1)$.

Choose o coordinate variables whose residue classes form a basis of R_1/J_1 . After permuting variables and changing the basis of J_1 , we may write

$$L_i = x_i + \sum_{j=v+1}^n c_{ij}x_j, \quad 1 \leq i \leq v. \quad (3.8)$$

This coordinate choice can be found from the kernel. Start with all variables and successively delete a variable whenever projection of $\ker M_d$ onto the degree- d monomial coordinates in the remaining variables still has rank $T(o, p, d)$. If their classes span r dimensions in R_1/J_1 , this projection has rank $T(r, p, d)$. Since $T(r, p, d) < T(o, p, d)$ for $r < o$, the test preserves exactly the sets whose classes span R_1/J_1 . The procedure therefore ends with o variables whose classes form a basis.

Set $\ell = o(p-1) - d$, and let $\mathcal{A}_\ell$ be the set of $\mathbf{a} = (a_{v+1}, \dots, a_n)$ with $0 \leq a_j < p$ and $\sum_{j=v+1}^n a_j = \ell$. Write $x^{\mathbf{a}} = \prod_{j=v+1}^n x_j^{a_j}$. By Lemma 3.1, the polynomials

$$g_{\mathbf{a}} = Mx^{\mathbf{a}}, \quad \mathbf{a} \in \mathcal{A}_\ell,$$

form a basis of $K_{N-d}^\perp$: the identities $ML_i = 0$ reduce every multiple of M to a polynomial multiple in the last o variables, and the normalization below shows linear independence. Eq. (3.8) gives the normalization

$$[x_1^{p-1} \cdots x_v^{p-1} x^{\mathbf{b}}]g_{\mathbf{a}} = \delta_{\mathbf{a}, \mathbf{b}}, \quad \mathbf{a}, \mathbf{b} \in \mathcal{A}_\ell. \quad (3.9)$$

Consequently, each $g_{\mathbf{a}}$ is the unique element of the computed kernel satisfying these coefficient constraints, and can be obtained by linear algebra without knowing the c_{ij} .

Since $\ell < o(p-1)$, some $a_j < p-1$. For this j and each $1 \leq i \leq v$, expansion of M gives

$$\left[x_i^{p-2} \prod_{\substack{1 \leq k \leq v \\ k \neq i}} x_k^{p-1} x^{\mathbf{a} + \mathbf{e}_j} \right] g_{\mathbf{a}} = (p-1)c_{ij} = -c_{ij}, \quad (3.10)$$

where $\mathbf{e}_j$ is the unit vector corresponding to x_j among the last o coordinates. Indeed, the indicated monomial uses one $c_{ij}x_j$ term from L_i^{p-1} and only the x_k^{p-1} terms from all other factors. Thus

$$(-c_{1j}, \dots, -c_{vj}, \mathbf{e}_j) \in \mathbb{F}_q^v \oplus \mathbb{F}_q^o$$

is a nonzero oil vector. For every $j > v$, an exponent vector $\mathbf{a} \in \mathcal{A}_\ell$ with $a_j < p-1$ exists; repeating the construction yields a basis of $\mathcal{O}$. $\square$

Once the coordinate choice in the proof is fixed, a single normalized vector can be computed using a reduced Macaulay system. Put $r = T(n, p, d)$ and $h = T(o, p, d)$. Setting the $h-1$ zero coordinates in Eq. (3.9) leaves a homogeneous system with $\sum_{i=1}^s \dim R_{d-d_i}$ rows and $r-h+1$ unknowns. Its kernel is one-dimensional, and any nonzero solution has a nonzero remaining normalization coordinate, which can be scaled to 1. Equivalently, fixing that coordinate to 1 gives an affine system with $r-h$ unknowns. Finding the coordinate choice and computing further vectors are separate steps.

Remark 3.3. Since each f_i has positive degree, $f_i^p = 0$ in R . Hence $P := (f_1 \cdots f_s)^{p-1}$ satisfies $Pf_i = 0$ for every i , and therefore $P \in K^\perp$; this element may be zero. For UOV quadratic systems, Ikematsu and Furue [6, §4] study recovery under conditions of the form $K_t^\perp = \langle P, M \rangle_t$, together with additional hypotheses. Theorem 3.2 uses the stronger condition $K_{N-d}^\perp = \langle M \rangle_{N-d}$.

The remaining task is to construct further public relations in J that help attain the dimension required by the recovery criterion. We use divided powers for this purpose.

3.2 Divided Powers on p -Truncated Polynomial Rings

Divided powers satisfy the formal identities of $f^k/k!$ even when $k!$ is not invertible in the base field. We first recall the general definition [11, Definition 23.2.1], then construct the structure used in our attack.

Definition 3.4 (Divided power algebra). *Let k be a commutative ring with identity, let A be a commutative k -algebra with identity, and let $\mathfrak{a} \subseteq A$ be an ideal. A divided power structure on $\mathfrak{a}$ is a family of maps $\gamma_r : \mathfrak{a} \rightarrow \mathfrak{a}$, $r \geq 1$, satisfying the following identities for all $u, v \in \mathfrak{a}$, $b \in A$, and $r, s \geq 1$, with the convention $\gamma_0(u) = 1 \in A$:*

$$\begin{aligned} \gamma_1(u) &= u, \\ \gamma_r(u)\gamma_s(u) &= \binom{r+s}{r} \gamma_{r+s}(u), \\ \gamma_r(bu) &= b^r \gamma_r(u), \\ \gamma_r(u+v) &= \sum_{i=0}^r \gamma_i(u)\gamma_{r-i}(v), \\ \gamma_r(\gamma_s(u)) &= \frac{(rs)!}{r!(s!)^r} \gamma_{rs}(u). \end{aligned}$$

The binomial and factorial-ratio coefficients are integers, interpreted in A via $\mathbb{Z} \rightarrow A$; no division by a nonunit in A is involved. We call the triple $(A, \mathfrak{a}, \gamma)$ a divided power k -algebra.

The product identity implies $r! \gamma_r(u) = u^r$ for every $r \geq 1$. Consequently, if A is a $\mathbb{Q}$ -algebra, the unique divided power structure on any ideal is $\gamma_r(u) = u^r/r!$. In characteristic p , this formula still holds for $1 \leq r < p$. For $r = p$, the identity instead gives $u^p = 0$ for every $u \in \mathfrak{a}$, without forcing $\gamma_p(u)$ to vanish.

In our application, the algebra is R and the distinguished ideal is $\mathfrak{m}$. We now define the divided powers on this ideal.

Write $x^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ and set

$$\mathcal{A} := \{\alpha \in \{0, \dots, p-1\}^n : |\alpha| > 0\}.$$

Every $f \in \mathfrak{m}$ has a unique expression $f = \sum_{\alpha \in \mathcal{A}} c_\alpha x^\alpha$ with $c_\alpha \in \mathbb{F}_q$. Set $\gamma_0(f) = 1$ and, for $k \geq 1$, define

$$\gamma_k(f) := \sum_{\substack{\alpha \in \mathcal{A} \\ 0 \leq k_\alpha < p}} \prod_{k_\alpha = k} \frac{c_\alpha^{k_\alpha} x^{k_\alpha \alpha}}{k_\alpha!} \quad \text{in } R. \quad (3.11)$$

Only factorials $k_\alpha!$ with $k_\alpha < p$ occur, so the formula is well defined for every k , including $k \geq p$. In particular, $\gamma_k(0) = 0$ for $k \geq 1$, while $\gamma_0(0) = 1$, and $\gamma_k(f) = f^k/k!$ for $0 \leq k < p$. For each nonconstant monomial x^α ,

$$\gamma_k(x^\alpha) = \begin{cases} x^{k\alpha}/k!, & 1 \leq k < p, \\ 0, & k \geq p. \end{cases}$$

Proposition 3.5. *The maps in Eq. (3.11) make $(R, \mathfrak{m}, \gamma)$ a divided power $\mathbb{F}_q$ -algebra in the sense of Definition 3.4. This structure is compatible with the grading: $\gamma_k(R_d) \subseteq R_{kd}$ for $k, d \geq 1$.*

Proof. Start with the divided power polynomial algebra $A = \mathbb{F}_q\langle X_1, \dots, X_n \rangle$ of [12]. It has $\mathbb{F}_q$ -basis $X_1^{[a_1]} \cdots X_n^{[a_n]}$ with $a_i \geq 0$, where $X_i^{[0]} = 1$ and $X_i = X_i^{[1]}$. Multiplication is determined by

$$X_i^{[r]} X_i^{[s]} = \binom{r+s}{r} X_i^{[r+s]}.$$

Its ideal spanned by the positive-degree basis monomials carries canonical divided powers with $\gamma_r(X_i) = X_i^{[r]}$. The ideal $\mathfrak{a} = \langle X_i^{[r]} : 1 \leq i \leq n, r \geq p \rangle_A$ is stable under divided powers: the divided powers of each generator are integer multiples of $X_i^{[kr]}$ with $k \geq 1$. Thus the divided power structure descends to $A/\mathfrak{a} \cong R$, where $X_i^{[r]}$ corresponds to $x_i^r/r!$ for $0 \leq r < p$. The induced maps are exactly those in Eq. (3.11), so they satisfy the axioms in Definition 3.4. If f is homogeneous of degree d , every surviving summand in Eq. (3.11) has degree kd , proving the grading compatibility. $\square$

Iterating the addition identity gives, for $f_1, \dots, f_l \in \mathfrak{m}$,

$$\gamma_k\left(\sum_{i=1}^l f_i\right) = \sum_{\substack{a_1, \dots, a_l \geq 0 \\ a_1 + \dots + a_l = k}} \prod_{i=1}^l \gamma_{a_i}(f_i).$$

For the attack, the values on the public quadratics must be compatible with the unknown oil coordinates. Although the structure on $\mathfrak{m}$ depends on the chosen coordinates, its restriction to $\mathfrak{m}^2$ does not. Indeed, for $u, v \in \mathfrak{m}$, any divided power structure satisfies $\gamma_k(uv) = (uv)^k/k!$ for $k < p$ and $\gamma_k(uv) = u^k \gamma_k(v) = 0$ for $k \geq p$, since $u^p = 0$. The addition formula then determines the maps on $\mathfrak{m}^2$ uniquely. In particular, the divided powers of a homogeneous quadratic are independent of linear coordinates.

Next, we reduce the number of divided powers that need to be computed. For $k \geq 1$, write $k = \sum_{e=0}^{\ell} k_e p^e$ with $0 \leq k_e < p$. The product identity and Lucas's theorem give

$$\gamma_k(f) = \prod_{e=0}^{\ell} \frac{\gamma_{p^e}(f)^{k_e}}{k_e!}. \quad (3.12)$$

Consequently,

$$\langle \gamma_k(f) : k \geq 1 \rangle_R = \langle \gamma_{p^e}(f) : e \geq 0 \rangle_R.$$

Thus only indices of the form p^e are needed to generate the ideal of divided power relations. For $1 \leq k < p$, the value $\gamma_k(f) = f^k/k!$ already belongs to $\langle f \rangle_R$, so potentially new generators come from the indices p^e with $e \geq 1$.

Finally, we need to know that the new relations remain in the hidden ideal. If $f = \sum_{i=1}^s u_i g_i$ with $u_i \in \mathfrak{m}$ and $g_i \in R$, then, for $k \geq 1$,

$$\gamma_k(f) = \sum_{\substack{(a_1, \dots, a_s) \in \mathbb{Z}_{\geq 0}^s \\ a_1 + \dots + a_s = k}} \prod_{i: a_i > 0} g_i^{a_i} \gamma_{a_i}(u_i) \in \langle g_1, \dots, g_s \rangle_R. \quad (3.13)$$

Thus $f \in \mathfrak{m}J$ implies $\gamma_k(f) \in J$ for every ideal $J \subseteq R$ and every $k \geq 1$.

3.3 Divided Power Ideals in UOV Systems

We now apply these operations to the public ideal I . Since each public quadratic p_i belongs to the ideal generated by the linear forms L_j , it admits a representation

$$p_i = \sum_{j=1}^v u_{ij} L_j, \quad u_{ij} \in R_1.$$

In particular, $p_i \in \mathfrak{m}J$. The preceding containment therefore applies to every divided power of every public quadratic.

Lemma 3.6. *With the notation above, we have*

$$I \subseteq \gamma(I) := \langle \gamma_{p^e}(p_i) : e \geq 0, 1 \leq i \leq m \rangle_R \subseteq J. \quad (3.14)$$

Proof. Since $\gamma_{p^0}(p_i) = \gamma_1(p_i) = p_i$, we have $I \subseteq \gamma(I)$. Moreover, $p_i \in \mathfrak{m}J$, so Eq. (3.13) gives $\gamma_{p^e}(p_i) \in J$ for every $e \geq 0$ and $1 \leq i \leq m$. Thus $\gamma(I) \subseteq J$. $\square$

The addition, scalar, and composition identities show that $\gamma(I)$ is the smallest ideal containing I and stable under all γ_k with $k \geq 1$. Equivalently,

$$\gamma(I) = \langle \gamma_k(p_i) : k \geq 1, 1 \leq i \leq m \rangle_R = \langle \gamma_k(f) : k \geq 1, f \in I \rangle_R.$$

In particular, it is independent of the chosen generators of I . It is homogeneous, since $\gamma_k(p_i) \in R_{2k}$. Only finitely many generators need to be computed: those with $2p^e > N$ vanish in R .

The effect on the Macaulay kernels follows directly from the ideal inclusions. Dualizing the corresponding quotient maps gives

$$(R_D/J_D)^\vee \hookrightarrow (R_D/\gamma(I)_D)^\vee \hookrightarrow (R_D/I_D)^\vee. \quad (3.15)$$

Thus adding divided power relations can reduce the kernel while retaining the subspace determined by J . In terms of dimensions,

$$T(o, p, D) \leq \dim(R_D/\gamma(I)_D) \leq \dim(R_D/I_D). \quad (3.16)$$

Applying Theorem 3.2 with $K = \gamma(I)$ shows that recovery is possible as soon as the left-hand bound is attained at a degree $d > 0$ with $T(o, p, d) > 0$. For a given instance, this condition can be checked by computing the dimension of the Macaulay right kernel.

3.4 Solving Degree and Linear-Algebra Cost

The recovery criterion is exact, but it does not determine in advance the degree at which the required dimension is reached. We use a Hilbert-series conjecture to predict this degree, then estimate the cost of a sparse solve at that degree.

For a homogeneous ideal $K \subseteq R$, write $K_D := K \cap R_D$ and

$$H_{R/K}(t) := \sum_{D \geq 0} \dim_{\mathbb{F}_q}(R_D/K_D)t^D.$$

The coefficient of t^D is the dimension of the degree- D Macaulay right kernel. The lower bound in Eq. (3.16) has a particularly simple Hilbert series. Since the L_j are linearly independent,

$$R/J \cong \mathbb{F}_q[y_1, \dots, y_o]/\langle y_1^p, \dots, y_o^p \rangle, \quad H_{R/J}(t) = (1 + t + \dots + t^{p-1})^o.$$

The following conjecture combines a semi-regular Hilbert-series estimate with this lower bound imposed by the oil space.

Conjecture 3.7 (Hilbert-series estimate). For generic UOV instances with the parameters above, we conjecture the degreewise estimate

$$[t^D]H_{R/\gamma(I)}(t) = \max\{[t^D]H_{\text{sr}}(t), T(o, p, D)\}, \quad (3.17)$$

for all $D \leq o(p-1)$, where

$$H_{\text{sr}}(t) := [(1+t+\cdots+t^{p-1})^n(1-t^2)^m]_+.$$

Here $[\sum_{d \geq 0} a_d t^d]_+$ retains the terms strictly before the first nonpositive coefficient and discards all subsequent terms.

Since the characteristic-two case has already been studied by Merz and Ran [9], we report here only some experimental results in odd characteristic. Table 1 compares the coefficients of $H_{R/\gamma(I)}(t)$ predicted by Conjecture 3.7 with those obtained experimentally for the three parameter sets listed. In all three cases, $o = n - v = 3$ and $p = 3$, so the displayed terms cover the full range $0 \leq D \leq o(p-1) = 6$ considered in the conjecture. The theoretical and experimental coefficients agree in each case, providing experimental support for the conjectured estimate.

Table 1: Comparison between the theoretical estimate and the experimental Hilbert series

parameter (n, v, m, p)	Hilbert Series	
(10,7,3,3)	theoretical	$1 + 10t + 52t^2 + 180t^3 + 453t^4 + 852t^5 + 1169t^6$
	experimental	$1 + 10t + 52t^2 + 180t^3 + 453t^4 + 852t^5 + 1169t^6$
(10,7,6,3)	theoretical	$1 + 10t + 49t^2 + 150t^3 + 300t^4 + 342t^5 + t^6$
	experimental	$1 + 10t + 49t^2 + 150t^3 + 300t^4 + 342t^5 + t^6$
(9,6,4,3)	theoretical	$1 + 9t + 41t^2 + 120t^3 + 240t^4 + 312t^5 + 164t^6$
	experimental	$1 + 9t + 41t^2 + 120t^3 + 240t^4 + 312t^5 + 164t^6$

Assuming Conjecture 3.7, a degree $2 \leq d \leq o(p-1)$ satisfying

$$[t^d]H_{\text{sr}}(t) \leq T(o, p, d) \quad (3.18)$$

gives $\dim(R_d/\gamma(I)_d) = T(o, p, d) > 0$, so Theorem 3.2 applies. The upper bound on d ensures that the target kernel is nonzero. The resulting prediction of a solving degree is conditional on the conjecture.

Suppose that we get a solving degree $1 \leq d \leq o(p-1)$ such that $[t^d]H_{R/\gamma(I)}(t) = [t^d]H_{R/J}(t)$. By Theorem 3.2, we only need compute a normalized right kernel vector of the degree- d Macaulay matrix of $\gamma(I)$, which is a right kernel vector with zeros at $T(o, d, p) - 1$ fixed positions. The linear system $M_d X = 0$ only involves $T(n, d, p) - T(o, d, p) + 1$ variables. M_d is sparse. The number of nonzero entries in the $x^\beta \gamma_{p^i}(f_j)$ -th row of M_d is the number of monomials in $\gamma_{p^i}(f_j)$, which is at most $\dim R_{2p^i}$. So the average number of nonzero entries per row of the Macaulay matrix, is at most

$$\rho_d := \frac{\sum_{0 \leq 2p^i \leq d} m \dim R_{d-2p^i} \dim R_{2p^i}}{\sum_{0 \leq 2p^i \leq d} \dim m R_{d-2p^i}}.$$

We apply the Wiedemann algorithm to compute one normalized kernel vector. It has the heuristic cost

$$3 \cdot \rho_d \cdot (T(n, p, d) - T(o, p, d) + 1)^2 \cdot (2 \log_2(q)^2 + \log_2(q)).$$

Remark 3.8. The classical wedge attack over the p -truncated polynomial algebra (see [4]) search for the minimal integer $1 \leq d \leq o(p-1)$ such that $[t]^d H_{R/I}(t) = [t]^d H_{R/J}(t) > 0$. We can determine the oil space by finding a normalized right kernel vector of the degree- d Macaulay matrix of I . Each row of this matrix has at most $\binom{n+1}{2}$ non-zero entries. The complexity can be estimated as

$$3 \cdot \binom{n+1}{2} \cdot (T(n, p, d) - T(o, p, d) + 1)^2 \cdot (2 \log_2(q)^2 + \log_2(q)).$$

The classical wedge attack requires the validity of the conjection that

$$[t^d] H_{R/I}(t) = \max \left\{ [t^d] \left[\frac{(\sum_{i=0}^{p-1} t^i)^n}{(\sum_{i=0}^{p-1} t^i)^m} \right]_+, T(o, p, d) \right\}. \quad (3.19)$$

However, this conjection is false when $v > 2m$ (this will not influence the validity of Conjection 3.7). Because in this case, the equality $I_d = J_d \subsetneq R_d$ is impossible for all $1 \leq d \leq o(p-1)$, see also Remark 3.3 and [6].

4 Applications

In this section, we apply the preceding framework to UOV and QR-UOV, and consider additional UOV parameter sets with $v \geq 2m$.

4.1 Scope and Relation to Prior Work

Before presenting the estimates, we clarify the scope of these applications and their relation to the work of Merz and Ran [9].

1. The UOV parameter sets in Table 2 are all defined over fields of characteristic two, and the estimates in that table can already be obtained using the method of Merz and Ran [9, §3]. Another major contribution of their work is the use of Schur modules to analyze the security of MAYO, QR-UOV, and SNOVA. We include Table 2 for comparison and do not claim these estimates as a new contribution. For UOV, our claim is that the method developed here would also apply to parameter sets in odd characteristic, should such sets be proposed.
2. The characteristic-two divided power construction in [9, §3] does not directly apply to the odd-characteristic QR-UOV parameter sets in Table 3. We use divided powers in p -truncated polynomial rings to enlarge the public ideal and estimate the cost of recovery for these parameters. Merz and Ran also implicitly use a divided power structure on a multigraded exterior algebra in [9, §5.3] to support their Schur–Macaulay attack. Our estimates therefore complement their Schur-module analysis of QR-UOV.

3. Merz and Ran [9] have already analyzed the recommended SNOVA parameter sets in characteristic two. Our method can also be applied to SNOVA parameter sets in odd characteristic. However, since SNOVA's main recommended parameter sets are confined to characteristic two, we do not pursue a separate analysis of its odd-characteristic parameter sets here.

4.2 UOV

We begin with UOV as a comparison case, before turning to QR-UOV as an application of our framework in odd characteristic.

The number of oil variables of a UOV instance with parameters (n, v, m, q) is m . In Table 2, we give complexity estimates for the reconciliation and intersection variants of both the classical wedge attack and the extended wedge attack using the divided power structure.

The solving degree D of the extended wedge attack is smaller than the solving degree in the classical wedge attack. But the density ρ_D of the Macaulay matrix in the extended wedge attack is higher than the density in the classical wedge attack. For the two parameter sets uov-Ip and uov-Is, the extended wedge attack achieves the same complexity as the classical intersection wedge attack.

Table 2: Complexity estimates given in the round 3 specifications of UOV and the number of fixed variables k , the solving degree D , and $\log_2$ density of the Macaulay matrix of the extended wedge attack.

parameter (n, v, m, p)	classical wedge attack (bit, $k, D, \log_2 N, \log_2 \rho_D$)		extended wedge attack (bit, $k, D, \log_2 N, \log_2 \rho_D$)	
	Reconciliation	Intersection	Reconciliation	Intersection
uov-Ip (119,74,45,256)	154 (25,20,67,12)	144 (1,15,62,13)	153 (27,18,63,20)	143 (2,14,59,18)
uov-Is (160,96,64,16)	176 (42,22,79,13)	158 (16,16,70,14)	181 (43,21,77,23)	158 (17,15,67,20)
uov-III (193,121,72,256)	227 (43,29,103,13)	213 (1,22,96,14)	232 (45,27,98,29)	216 (2,21,93,24)
uov-III (259,163,96,256)	296 (58,38,137,14)	272 (1,28,125,15)	298 (62,34,127,36)	280 (2,27,121,31)

4.3 QR-UOV

To estimate the complexity of the extended wedge attack against the parameters (n, v, m, q, ℓ) of QR-UOV, we identify it with a plain UOV over the extension

field $\mathbb{F}_{q^\ell}$ with v/ℓ vinegar variables, m/ℓ oil variables, and m equations. When $q = 127$, the solving degree D is less than q . In this case, there is no difference between the extended wedge attack and the classical wedge attack. Indeed,

$$(1 + \dots + t^{126})^n (1 - t^2)^m \equiv \frac{(1 + \dots + t^{126})^n}{(1 + t^2 + \dots + t^{252})^m} \equiv \frac{(1 - t^2)^m}{(1 - t)^{127}} \pmod{t^{127}}$$

When $q = 7$ or 31 , the extended wedge attack and the classical wedge attack have the same solving degree, See Table 3.

Table 3: The complexity of the best known attacks in the specification document [5], the complexity of the reconciliation wedge attack given in the round 3 specifications of QR-UOV and the number of fixed variables k , the solving degree D , and $\log_2$ density of the Macaulay matrix of the extended wedge attack.

SL	Parameters (q, v, m, ℓ)	[5]	reconciliation attack complexity, ($k, D, \log_2 N, \log_2 \rho_D$)	
			classical wedge attack	extended wedge attack
I	(127,156,54,3)	150	164 (17,28,72,11)	164 (17,28,72,11)
	(7,740,100,10)	-	159 (6,20,68,12)	166 (6,20,68,19)
	(31,165,60,3)	-	150 (19,23,65,11)	150 (19,23,65,11)
II	(7,1100,140,10)	-	230 (9,30,103,13)	246 (9,30,103,30)
	(31,890,100,10)	-	220 (8,32,98,13)	220 (8,32,98,13)
III	(7,1490,190,10)	-	299 (12,39,137,14)	321 (12,39,137,36)
	(31,324,114,3)	-	290 (36,47,134,13)	290 (36,47,134,13)
	(127,820,90,10)	211	207 (8,31,93,12)	207 (8,31,93,12)
V	(127,1040,110,10)	273	276 (10,43,125,13)	276 (10,43,125,13)

4.4 Parameter Sets with $v \geq 2m$

For some UOV-based systems with parameters (n, v, m, q) such that $v \geq 2m$ and q is odd. Neither the classical wedge attack [10, 7, 4] nor the method proposed in [9, §3] applies to these systems. The method proposed by Ikematsu and Furue [6] can deal with them. But the extended wedge attack over divided power algebras can obtain less complexity.

1. $(n, v, m, q) = (70, 40, 16, 11)$.

- The complexity of our method is 359. $k = 17$ and the solving degree is 129.
- The complexity of the method proposed by Ikematsu and Furue in [6] is 393, $k = 16$ and the solving degree is 140.
- 2. $(n, v, m, q) = (90, 50, 20, 7)$.
 - The complexity of our method is 383. $k = 24$ and the solving degree is 94.
 - The complexity of the method proposed by Ikematsu and Furue in [6] is 406, $k = 23$ and the solving degree is 102.

5 Conclusions and Future Work

Conclusion Building on the divided power construction of Merz and Ran [9, §3], we have developed an extension of the wedge attack in p -truncated polynomial rings over finite fields of arbitrary characteristic. This formulation directly handles UOV-type systems in odd characteristic, including QR-UOV, by adjoining additional public relations that remain in the hidden ideal J . Under Conjecture 3.7, it also enables recovery for the odd-characteristic instances with $v > 2m$ considered in Section 4, where the classical wedge attack [10, 7, 4] does not apply.

For the two examples with $v > 2m$, the comparisons in this paper predict lower costs for our method than for the generalized wedge attack of Ikematsu and Furue [6]. These examples illustrate that divided power relations can extend the applicability of the classical wedge attack and reduce the estimated cost for selected parameters. For the QR-UOV parameter sets studied here, the estimates do not improve on the classical wedge attack. Our complexity comparisons remain conditional on the conjectured Hilbert-series prediction.

Future Work Our complexity estimates assume that the actual Hilbert series agrees with the prediction in Conjecture 3.7 in the degrees relevant to recovery. We have not proved this conjecture; our evidence consists only of numerical experiments on small instances, in which no discrepancies were observed. The complexities reported here are therefore theoretical estimates conditional on this assumption. Further theoretical analysis of the validity and scope of this prediction is an important direction for future work.

The broader use of Schur modules in [9] suggests further connections beyond the construction studied here. A natural direction is to investigate how these ideas interact with divided powers in odd characteristic, and whether they can yield additional useful relations in J and more effective attacks on UOV-like schemes.

Acknowledgements We thank Lars Ran and Simon-Philipp Merz for helpful discussions, from which we have benefited greatly. This work was supported by the National Key R&D Program of China (No. 2023YFC3305501) and the National Natural Science Foundation of China (Grant No. 126011023).

AI disclosure In preparing this paper, we used GPT-6 Astra to edit the text for clarity and grammar. The idea of considering general divided power algebras was formulated by the authors, but was largely inspired by [9, §3]. The authors take full responsibility for the content and any errors that may remain.

Bibliography

- [1] Beullens, W.: Improved cryptanalysis of uov and rainbow. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 348–373. Springer (2021)
- [2] Bros, M., Le, T.H., Lichtinger, J., Minaud, B., Perlner, R., Smith-Tone, D., Valenzuela, C.: Exploiting SNOVA’s structure in the wedge product attack. Cryptology ePrint Archive (2026)
- [3] Ding, J., Li, P., Tao, S.: Revisiting the wedge attack on uov problem. Cryptology ePrint Archive (2026)
- [4] Furue, H., Ikematsu, Y.: Key recovery attacks on uov using p^l -truncated polynomial rings. Cryptology ePrint Archive (2026)
- [5] Furue, H., Ikematsu, Y., Hoshino, F., Takagi, T., Yasuda, K., Miyazawa, T., Saito, T., Nagai, A.: Qr-uov. Specification document of NIST PQC Standardization of Additional Digital Signature Scheme (2026)
- [6] Ikematsu, Y., Furue, H.: Extending the applicability of algebraic key recovery attacks on the uov signature scheme. Cryptology ePrint Archive (2026), <https://eprint.iacr.org/2026/1620>, paper 2026/1620
- [7] Jin, Y., Pan, Y., He, X., Gong, B., Ding, J.: Security analysis on uov families with odd characteristics: Using symmetric algebra. In: IACR International Conference on Public-Key Cryptography. pp. 428–454. Springer (2026)
- [8] Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: International Conference on the Theory and Applications of Cryptographic Techniques. pp. 206–222. Springer (1999)
- [9] Merz, S.P., Ran, L.: Attacking UOV-based signatures with Schur–Macaulay matrices. Cryptology ePrint Archive, Paper 2026/2181 (2026), <https://eprint.iacr.org/2026/2181>
- [10] Ran, L.: Wedges, oil, and vinegar: An analysis of uov in the exterior algebra. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 363–388. Springer (2026)
- [11] Stacks project authors, T.: The Stacks project. <https://stacks.math.columbia.edu/tag/07GK> (2026), section 23.2: Divided powers, Tag 07GK
- [12] Stacks project authors, T.: The Stacks project. <https://stacks.math.columbia.edu/tag/07H4> (2026), section 23.5: Divided power polynomial algebras, Tag 07H4
- [13] Tseng, P.E., Wang, L.C., Li, P., Kuan, Y.L.: Investigating the wedge map on SNOVA. Cryptology ePrint Archive (2026)