Attacking UOV-based Signatures over divided power algebras

Peigen Li[1] and Siyong Tao[2]

[1] Beijing Institute of Mathematical Sciences and Applications
lpg22@bimsa.cn
[2] Department of Mathematical Sciences, Tsinghua University
taosy22@mails.tsinghua.edu.cn

Abstract. We introduce a family of divided power wedge attacks on UOV-type systems over finite fields of arbitrary characteristic. The construction extends the characteristic-two exterior-algebra construction of Merz and Ran to the p-truncated polynomial ring, where p is the field characteristic. An attack level controls which divided-power relations are adjoined, balancing a lower solving degree against increased average Macaulay row weight. These relations remain in the ideal defining the hidden oil space and can remove extraneous kernel elements, including those obstructing the classical attack when $v \geq 2m$, where v and m denote the numbers of vinegar variables and public equations. We establish an oil-space recovery criterion and use conjectural Hilbert-series predictions to estimate solving degrees and attack costs. A multigraded extension accommodates lifted SNOVA systems. We evaluate the framework on UOV, MAYO, QR-UOV, and SNOVA. Relative to the comparison baselines reported in this paper, our conditional estimates yield reductions of 2 bits for three UOV parameter sets, 3–14 bits for selected additional MAYO parameter sets, and 2–24 bits for three selected additional SNOVA parameter sets. No improvement is obtained for the QR-UOV parameter sets studied.

Keywords: UOV problem · wedge attack · divided power algebras

1 Introduction

The Unbalanced Oil and Vinegar (UOV) signature scheme [12] 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 [2]. 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 [15] introduced the wedge attack using the exterior algebra in characteristic two. Subsequent work developed formulations in reduced symmetric algebras [11] and truncated polynomial rings [8].

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 [14] 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 In this work, we propose a family of new attacks indexed by $\mathbb{N}$ for UOV-like systems over finite fields $\mathbb{F}_q$ of arbitrary characteristic. Our construction combines the p-truncated polynomial algebra of Furue and Ikematsu [8] with the divided power method of Merz and Ran [14]. We extend the latter from characteristic two to arbitrary characteristic by working directly with the public quadratic polynomials in

$$R = \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p \rangle, \quad p = \mathrm{char}(\mathbb{F}_q).$$

This formulation does not require the alternating polar forms, which are central to the exterior-algebra construction in characteristic two.

Given public quadratic polynomials $\{f_i\}_{i=1}^m$, the level-r divided power wedge attack is based on the ideal (see Eq. (3.3))

$$I_{r,\mathrm{dp}} := \langle \gamma_{p^e}(f_i) : 0 \leq e \leq r, 1 \leq i \leq m \rangle \subseteq R$$

We have the ascending chain of ideals

$$I = I_{0,\mathrm{dp}} \subseteq I_{1,\mathrm{dp}} \subseteq \cdots \subseteq J,$$

where J is the ideal generated by the hidden linear forms defining the oil space.

Level-0 divided power wedge attack recovers the p-key recovery attack of [8], which seeks an integer d such that $I_d = J_d \neq 0$. The existence of such d allows us to recover the oil space from the right nullspace of the Macaulay matrix M_d of I. Generally, the level-r divided power wedge attack aims to find the integer d such that $(I_{r,\mathrm{dp}})_d = J_d \neq 0$. Each additional level adds m generators and lowers the solving degree d. Therefore we can control the trade-off between matrix size

and the average number of nonzero entries per row of the Macaulay matrix by the level r.

The additional relations also address the kernel obstruction that limits the classical recovery criterion to $v < 2m$. Ikematsu and Furue [10] recover oil-space information from kernels containing additional public components, while our approach enlarges the public ideal to remove such components in a suitable degree. We establish a recovery criterion and use Conjecture 3.6, supported by experiments on small instances, to predict suitable degrees. For some parameter sets with $v \geq 2m$, this approach yields lower estimated costs than the extended wedge attack of [10] (Table 2).

We apply the framework to UOV, MAYO, QR-UOV, and SNOVA, using a multigraded formulation for the lifted SNOVA systems. Conditional on Conjecture 3.6 about Hilbert series and its multigraded analogue, the estimates reduce the stated comparison baselines by 2 bits for `uov-Ip`, `uov-III`, and `uov-V` (Table 3), by 3–14 bits for four additional MAYO parameter sets (Table 4), and by 2, 18, and 24 bits for the three additional SNOVA parameter sets in Table 6. For the QR-UOV parameter sets in Table 5, increasing the level does not lower the estimated cost, and the resulting estimates do not improve on the listed comparison baselines.

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

https://github.com/Yulan24/divided_power_wedge_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 [12]. We follow Beullens' description of UOV in terms of its hidden oil space [2]. Let $n = v + o$ with $v > o$. The public key is a homogeneous quadratic map $\mathcal{P} = (p_1, \ldots, p_m) : \mathbb{F}_q^n \to \mathbb{F}_q^m$ 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

$$\mathcal{P}'(\mathbf{x}, \mathbf{y}) := \mathcal{P}(\mathbf{x}+\mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y}) = (p_1'(\mathbf{x}, \mathbf{y}), \ldots, p_m'(\mathbf{x}, \mathbf{y})).$$

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

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

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),

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

Equivalently,

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

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 $(p_1, \cdots, p_m) : \mathbb{F}_q^n \to \mathbb{F}_q^m$ 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 [15] 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 $L_1, \ldots, L_v$ 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 $\Omega = L_1 \wedge \cdots \wedge L_v$ lies in the kernel of

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

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

Jin et al. [11] developed a reduced symmetric-algebra formulation that also applies in odd characteristic. Furue and Ikematsu [8] reformulated these attacks using p-truncated polynomial rings and extended them to p^ℓ-truncated rings, where $p = \mathrm{char}(\mathbb{F}_q)$, including their use in reconciliation and intersection attacks. Wedge methods were also adapted to SNOVA by exploiting its additional algebraic structure: Bros et al. [5] used its block-ring structure, while Tseng et al. [16] studied the wedge map through a lifting to an extension field.

3 Extending to the divided power algebra

Throughout this section, we let R be the p-truncated polynomial algebra:

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

We have the natural map $\mathbb{F}_q[x_1, \ldots, x_n] \to R$, $f \mapsto \overline{f} := f + \langle x_1^p, \ldots, x_n^p \rangle$. Usually, we do not distinguish f from its coset $\overline{f}$ in R. R is an $\mathbb{N}$-graded algebra. For an integer $d \geq 0$, we denote by $\mathcal{A}_d$ the set of monomials in R of degree d:

$$\mathcal{A}_d := \{\mathbf{x}^\alpha = x_1^{\alpha_1} \cdots x_n^{\alpha_n} : 0 \leq \alpha_i \leq p-1, \alpha_1 + \cdots + \alpha_n = d\}.$$

R_d is the $\mathbb{F}_q$-vector space spanned by $\mathcal{A}_d$. $R_d = 0$ when $d > N := n(p-1)$. Let $\omega := (p-1, \ldots, p-1)$. $R_N = \mathbb{F}_q\{\mathbf{x}^\omega\}$ is of dimension 1. $R = \oplus_{d=0}^N R_d$. Set

$$F_{n,p}(t) := \left(\sum_{i=0}^{p-1} t^i \right)^n \in \mathbb{Z}[t]$$

$\dim R_d$ is exactly the coefficient of t^d in the polynomial $F_{n,p}(t)$. We denote this dimension by $T(n, p, d)$.

For polynomials $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$, we will use $\mathbb{V}(f_1, \ldots, f_m)$ to denote all the common zeros of $f_1, \ldots, f_m$ in $\mathbb{F}_q^n$.

3.1 Oil-space recovery from Macaulay kernels

Let $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$ be homogeneous polynomials such that $f_i \in \langle L_1, \ldots, L_v \rangle$ for all i, where the L_i are linearly independent linear polynomials. The solution space of the system $f_1(\mathbf{a}) = \cdots = f_m(\mathbf{a}) = 0$ contains the hidden oil space $\mathcal{O} = \mathbb{V}(L_1, \ldots, L_v)$ of dimension $o := n - v$. Without loss of generality, we assume that $L_i = x_i + \sum_{j=v+1}^n a_{i,j} x_j$ with unknown coefficients $a_{i,j} \in \mathbb{F}_q$. In this section, we first revisit the key recovery framework proposed in [7]. It gives us a criterion for recovering $\mathcal{O}$ when we can obtain enough information about it from the Macaulay matrix of $f_1, \ldots, f_m$.

Passing to the quotient ring R, we have the containment

$$K := \langle f_1, \ldots, f_m \rangle_R \subseteq J := \langle L_1, \ldots, L_v \rangle_R$$

Suppose that $f_i \in R_{d_i}$. Fix an integer $d \geq 0$, define the Macaulay map of degree d to be the linear map

$$\mathcal{M}_d : \bigoplus_{i=1}^m R_{d-d_i} \to R_d, \quad (u_1, \ldots, u_m) \mapsto \sum_{i=1}^m u_i f_i.$$

Here we set $R_d = 0$ if $d < 0$. It is easy to verify that $\operatorname{Im} \mathcal{M}_d = K_d := K \cap R_d$. Its dual map

$$\mathcal{M}_d^\vee : R_d^\vee \to \left(\bigoplus_{i=1}^m R_{d-d_i} \right)^\vee, \quad (F : R_d \to \mathbb{F}_q) \mapsto F \circ \mathcal{M}_d$$

is represented by the Macaulay matrix M_d. The columns of M_d are indexed by the set $\mathcal{A}_d$, and its rows are indexed by $\{(\mathbf{x}^\beta, f_i) : 1 \le i \le m, \mathbf{x}^\beta \in \mathcal{A}_{d-d_i}\}$. The $((\mathbf{x}^\beta, f_i), \mathbf{x}^\alpha)$-th entry of M_d is $[\mathbf{x}^\alpha](\mathbf{x}^\beta f_i)$, the coefficient of $\mathbf{x}^\alpha$ in $\mathbf{x}^\beta f_i$.

For any $0 \le d \le N$, there exists a perfect $\mathbb{F}_q$-bilinear pairing $\lambda : R_d \times R_{N-d} \to \mathbb{F}_q$ defined by

$$\lambda(f, g) := [\mathbf{x}^\omega](fg).$$

Under the pairing λ, the monomial bases of R_d and R_{N-d} are visibly dual, with $\mathbf{x}^{\omega-\alpha}$ dual to $\mathbf{x}^\alpha$ $(0 \le \alpha_i \le p-1)$. λ induces the isomorphism

$$R_{N-d} \cong R_d^\vee, \quad f \mapsto (\lambda_f : g \in R_d \to \lambda(f, g)).$$

This observation enables us to describe $\ker \mathcal{M}_d^\vee$ explicitly. Restricting the isomorphism $R_{N-d} \cong R_d^\vee$ gives a one-to-one correspondence:

$$K_{N-d}^\perp := \{f \in R_{N-d} : fK_d = 0\} \cong \ker \mathcal{M}_d^\vee. \tag{3.1}$$

Hence $\dim K_{N-d}^\perp = \dim \ker \mathcal{M}_d^\vee = \dim \mathrm{Coker} \mathcal{M}_d = \dim R_d/K_d$. Naturally, $J_{N-d}^\perp \subseteq K_{N-d}^\perp$.

Lemma 3.1. *Let $M = (L_1 \cdots L_v)^{p-1} \in R_{v(p-1)}$. Then $J_{N-d}^\perp = \langle M \rangle_{N-d}$.*

Proof. Set $L_i = x_i$ when $i > v$. From the form of $L_1, \ldots, L_v$, we see that $\{L_i\}_{i=1}^n$ is a basis of R_1. In particular,

$$\{\mathbf{L}^\alpha := L_1^{\alpha_1} \cdots L_n^{\alpha_n} : 0 \le \alpha_i \le p-1, \alpha_1 + \cdots + \alpha_n = N - d\}$$

form a basis of R_{N-d}. Now suppose that $f = \sum c_\alpha \mathbf{L}^\alpha \in R_{N-d}$ annihilates J_d. If $c_\alpha \neq 0$, then $\mathbf{L}^{\omega-\alpha} := L_1^{p-1-\alpha_1} \cdots L_n^{p-1-\alpha_n} \in R_d$, and

$$\mathbf{L}^{\omega-\alpha} f = c_\alpha \mathbf{L}^\omega = c_\alpha \mathbf{x}^\omega \neq 0,$$

so $\mathbf{L}^{\omega-\alpha} \notin J_d$, which implies that $\alpha_1 = \cdots = \alpha_v = p-1$. Therefore,

$$J_{N-d}^\perp \subseteq \mathrm{Span}\{M\mathbf{x}^\alpha : \mathbf{x}^\alpha = x_{v+1}^{\alpha_{v+1}} \cdots x_n^{\alpha_n} \in \mathcal{A}_{N-d-v(p-1)}\}.$$

$J_{N-d}^\perp \subseteq \langle M \rangle_{N-d}$. Conversely, since $L_i M = 0$ for all $1 \le i \le v$, $\langle M \rangle_{N-d} \subseteq J_{N-d}^\perp$ is obvious. □

Combining Eq. (3.1) with the lemma above, we get $\dim R_d/J_d = T(o, p, d)$. When $d > o(p-1)$, $J_d = R_d$ and $J_{N-d}^\perp = 0$.

Theorem 3.2. *For some integer $0 \le d \le o(p-1)$, if*

$$\dim R_d/K_d = \dim R_d/J_d \neq 0$$

then we can recover a nontrivial oil element from the right nullspace of M_d.

Proof. Since $K_d \subseteq J_d$, if $\dim R_d/K_d = \dim R_d/J_d$, then $K_d = J_d$ and $K_{N-d}^{\perp} = J_{N-d}^{\perp}$. Fix any monomial $\mathbf{x}^{\alpha} = x_{v+1}^{\alpha_{v+1}} \cdots x_n^{\alpha_n} \in \mathcal{A}_d$. Next, we find a linear functional $F : R_d \to \mathbb{F}_q$ satisfying

$$M_d \cdot (F(\mathbf{x}^{\beta}))_{\mathbf{x}^{\beta} \in \mathcal{A}_d} = 0 \text{ and } F(\mathbf{x}^{\beta}) = \begin{cases} 0 & \text{if } \mathbf{x}^{\beta} = x_{v+1}^{\beta_{v+1}} \cdots x_n^{\beta_n} \neq \mathbf{x}^{\alpha} \\ 1 & \text{if } \mathbf{x}^{\beta} = \mathbf{x}^{\alpha}. \end{cases}$$

F can be computed from the right nullspace of M_d. As F vanishes on $I_d = J_d$, we have $F = \lambda_f$ for some $f \in J_{N-d}^{\perp}$. By Lemma 3.1, $f = M x_{v+1}^{p-1-\alpha_{v+1}} \cdots x_n^{p-1-\alpha_n}$. Suppose that $\alpha_{v+1} \neq 0$. Then, for $1 \leq i \leq v$,

$$F(x_i x_{v+1}^{\alpha_{v+1}-1} \cdots x_n^{\alpha_n}) = (p-1)a_{i,v+1} = -a_{i,v+1}.$$

We obtain a nontrivial oil element $(-a_{1,v+1}, \ldots, -a_{v,v+1}, 1, \ldots, 0)^{\top} \in \mathcal{O}$. □

The existence of an integer $d \leq o(p-1)$ such that $\dim R_d/K_d = \dim R_d/J_d \neq 0$, i.e., $K_{N-d}^{\perp} = J_{N-d}^{\perp}$ requires $\sum_{i=1}^{m} d_i < v$. Otherwise,

$$K_{N-d}^{\perp} \supseteq \langle P, M \rangle_{N-d} \supsetneqq \langle M \rangle_{N-d} = J_{N-d}^{\perp}, \quad P = (f_1 \cdots f_m)^{p-1}.$$

Generally the second containment is strict, as $\deg P < \deg M$ and $M \notin \langle P \rangle$.

Once a first vector in $\mathcal{O}$ is found, finding subsequent vectors is much easier. The complexity of the attack is dominated by the complexity of finding the first vector in $\mathcal{O}$.

3.2 Divided power structure on the p-truncated polynomial algebra

A divided power structure is a way of introducing items with similar properties as expressions of the form $\frac{x^n}{n!}$ have, also when it is not possible to actually divide by $n!$.

Definition 3.3. *Let A be a commutative ring, $I \subseteq A$ an ideal. A divided power structure on I we mean a collection of maps $\gamma_k : I \to A$ for all integers $k \geq 0$ such that:*

1. *For all $f \in I$, $\gamma_0(f) = 1, \gamma_1(f) = f, \gamma_k(f) \in I$ if $k > 0$.*
2. *For $f, g \in I$, $\gamma_k(f+g) = \sum_{i=0}^{k} \gamma_i(f)\gamma_{k-i}(g)$,*
3. *For $u \in A, f \in I$, $\gamma_k(uf) = u^k \gamma_k(f)$.*
4. *For $f \in I$, $\gamma_{k_1}(f)\gamma_{k_2}(f) = \binom{k_1+k_2}{k_1}\gamma_{k_1+k_2}(f)$.*
5. *For $k_1, k_2 > 0, f \in I$, $\gamma_{k_1}(\gamma_{k_2}(f)) = \frac{(k_1 k_2)!}{k_1!(k_2)!^{k_1}}\gamma_{k_1 k_2}(f)$.*

The fourth property implies that $f^k = \gamma_1(f)^k = k!\gamma_k(f)$. So if $\operatorname{char} A = p$ and I admits a divided power structure, then $f^p = 0$ for all $f \in I$. Moreover, recall that Lucas' theorem states that for integers $n, m \geq 0$ and the prime number p,

$$\binom{n}{m} = \prod_{i \geq 0} \binom{n^{[i]}}{m^{[i]}} \bmod p,$$

where for any integer n, $n = \sum_{i\geq 0} n^{[i]}p^i$ with $n^{[i]} \in \{0, \ldots, p-1\}$ is the base-p expansion of n. By Lucas' theorem, we deduce that if $\{\gamma_k\}_{k\geq 0}$ is a divided power structure on I, then for any $f \in I$,

$$\gamma_k(f) = \prod_{i=0}^{e} \frac{\gamma_{p^i}(f)^{k^{[i]}}}{k^{[i]}!}.$$

For $f_1, \ldots, f_n \in I$, repeatedly using the second property, we obtain

$$\gamma_k(f_1 + \cdots + f_n) = \sum_{a_1+\cdots+a_n=k} \gamma_{a_i}(f_i). \tag{3.2}$$

The above identities are important for us, and it is natural to consider the divided power structure on $\mathfrak{m} := \langle x_1, \ldots, x_n\rangle_R \subseteq R = \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p\rangle$.

Lemma 3.4. *Any $f \in \mathfrak{m}$ can be uniquely written as $f = \sum_{0\leq\alpha_i\leq p-1} c_\alpha \boldsymbol{x}^\alpha$ with $c_\alpha \in \mathbb{F}_q$. For $k \geq 0$, we set*

$$\gamma_k(f) := \sum_{\sum_\alpha b_\alpha = k, 0\leq b_\alpha\leq p-1} \prod_\alpha \frac{(c_\alpha \boldsymbol{x}^\alpha)^{b_\alpha}}{b_\alpha!}.$$

$\{\gamma_k\}_{k\geq 0}$ is a divided power structure on $\mathfrak{m}$.

For the proof of this lemma see Appendix A.

3.3 Divided power ideals in UOV systems

Suppose that $\{p_i\}_{i=1}^m$ are polynomials in a public key of the UOV signature scheme with the parameters (n, v, m, q). Let $I \subseteq R = \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p\rangle$ be the UOV-like ideal generated by $p_1, \ldots, p_m$. $I \subseteq J := \langle L_1, \ldots, L_v\rangle$ where each L_i is a linear polynomial. We now study the divided powers of p_i and propose the divided power wedge attack against UOV systems. This attack is based on the oil-space recovery framework which we introduced in Section 3.1

Lemma 3.5. *For an arbitrary integer $r \geq 0$, we have*

$$I \subseteq I_{r,\mathrm{dp}} := \langle \gamma_{p^e}(p_i) : 0 \leq e \leq r,\ 1 \leq i \leq m\rangle \subseteq J. \tag{3.3}$$

Proof. Since $\gamma_{p^0}(p_i) = p_i$, we have $I \subseteq I_{r,\mathrm{dp}}$. Now we show the second containment. Each p_i can be written as a finite sum $p_i = \sum_{j=1}^v u_{i,j}L_j$ with $u_{i,j} \in \mathfrak{m}$. Then by Eq. 3.2 and the fourth property of the divided powers,

$$\gamma_{p^e}(p_i) = \sum_{a_1+\cdots+a_v=p^e} \prod_{j=1}^{v} \gamma_{a_j}(u_{i,j}L_j) = \sum_{a_1+\cdots+a_v=p^e} \prod_{j=1}^{v} L_j^{a_j}\gamma_{a_j}(u_{i,j})$$

As $p^e > 0$, we have $\gamma_{p^e}(p_i) \in J$ for all e and i. Hence $I_{r,\mathrm{dp}} \subseteq J$ □

.

In the classical wedge attack (we will introduce it in Section 3.5), we search for the degree d such that:

$$I_d = J_d \neq 0, \quad 1 \leq d \leq o(p-1).$$

When $p = 2$ and o is small, we cannot find such d. Now the r-level divided power wedge attack searches for the degree d satisfying

$$(I_{r,\mathrm{dp}})_d = J_d \neq 0, \quad 1 \leq d \leq o(p-1).$$

Once we find the solving degree, we can apply Theorem 3.2 in the case $K = I_{r,\mathrm{dp}}$ to recover oil elements in $\mathbb{V}(L_1, \ldots, L_v)$. To ensure the existence of the solving degree, we need the condition that

$$\sum_{i=1}^{m} \sum_{e=0}^{r} \deg \gamma_{p^e}(p_i)^{p-1} \geq \sum_{i=1}^{v} \deg L_i^{p-1},$$

i.e., $2m(1 + p + \cdots + p^r) \geq v$. And this holds for almost all the parameters of UOV-like schemes if we choose a suitable $r > 0$.

3.4 Solving Degree and Linear-Algebra Cost

For graded ideals $I_{r,\mathrm{dp}} \subseteq J \subseteq R$ in Eq. (3.3), the Hilbert series of $R/I_{r,\mathrm{dp}}$ is

$$H_{R/I_{r,\mathrm{dp}}}(t) := \sum_{d=0}^{n(p-1)} (\dim R_d/(I_{r,\mathrm{dp}})_d)\, t^d \in \mathbb{Z}[t].$$

Similarly, write $H_{R/J}(t)$ for the Hilbert series of R/J. The proof of Lemma 3.1 shows that $H_{R/J}(t) = F_{o,p}(t) = (\sum_{i=0}^{p-1} t^i)^o$.

Conjecture 3.6. Let $I_{r,\mathrm{dp}} = \langle \gamma_{p^e}(f_i) : 1 \leq i \leq m, 0 \leq e \leq r \rangle \subseteq R$ be the ideal generated by the divided powers of generic UOV polynomials $\{p_i\}_{i=1}^m$ with the parameters $\tau = (n, v, m, q)$. If $o < v < \min\{\frac{(o-1)m}{2}, 2m(1 + p + \cdots + p^r)\}$, then

$$H_{R/I_{r,\mathrm{dp}}}(t) \equiv \max\{H_{n,m,p,r}(t), F_{o,p}(t)\} \bmod t^{o(p-1)+1},$$
$$H_{n,m,p,r}(t) := \frac{F_{n,p}(t)}{F_{m,p^{r+1}}(t^2)}.$$

Here for two formal power series $F = \sum a_i t^i$ and $G = \sum b_i t^i \in \mathbb{Z}[\![t]\!]$, $\max\{F, G\} := \sum_{i \leq i_0} a_i t^i + \sum_{i > i_0} b_i t^i$, where $i_0 \geq 0$ is the maximal integer such that $a_{i_0} > b_{i_0}$.

Assuming the validity of the above conjecture (we will verify it experimentally later), we estimate the solving degree for the level-r divided power wedge attack against UOV with the parameters $\tau = (n, v, m, q)$ as

$$D_{r,\mathrm{dp}}(\tau) := \min\{1 \leq d \leq o(p-1) : [t^d] H_{n,m,p,r}(t) \leq T(o, p, d)\}.$$

We simply write D_r for this generalized notion of the solving degree when there is no danger of confusion. We distinguish between solving degrees only when the parameters or the attack schemes under consideration differ.

Using Theorem 3.2, choose a monomial $\mathbf{x}^\alpha = x_{v+1}^{\alpha_{v+1}} \cdots x_n^{\alpha_n} \in \mathcal{A}_{D_r}$, we need to compute a nontrivial right kernel $X = (X_{\mathbf{x}^\beta})_{\mathbf{x}^\beta \in \mathcal{A}_{D_r}(R)}$ of the Macaulay matrix M_{D_r} satisfying

$$X_{\mathbf{x}^\beta} = \begin{cases} 0 & \mathbf{x}^\beta = x_{v+1}^{\beta_{v+1}} \cdots x_n^{\beta_n} \neq \mathbf{x}^\alpha, \\ 1 & \mathbf{x}^\beta = \mathbf{x}^\alpha. \end{cases}$$

Quite a lot of information of the oil space is available from X.

The linear system $M_{D_r} X = 0$ involves $T(n,p,D_r) - T(o,p,D_r) + 1$ variables. For $e \leq r$ and $\mathbf{x}^\beta \in \mathcal{A}_{D_r - 2p^e}$, the $\mathbf{x}^\beta \gamma_{p^e}(f_i)$-th row of M_{D_r} has at most

$$\#\{\mathbf{x}^\alpha \in \mathcal{A}_{2p^e} : [\mathbf{x}^\alpha]\gamma_{p^e}(f_i) \neq 0\} \leq \#\mathcal{A}_{2p^e} = T(n,p,2p^e)$$

nonzero entries. When $p = 2$, this upper bound can be sharpened:

$$\begin{aligned} &\#\{\text{non-zero entries in the } \mathbf{x}^\beta \gamma_{p^e}(f_i)\text{-th row of } M_{D_r}\} \\ \leq &\#\{\mathbf{x}^\alpha \in \mathcal{A}_{2^{e+1}} : \mathbf{x}^\beta \mathbf{x}^\alpha \neq 0\} = \binom{n - |\beta|}{2^{e+1}} = \binom{n - D_r + 2^{e+1}}{2^{e+1}}. \end{aligned}$$

Let ρ_{D_r} be the average number of nonzero entries in any row of M_{D_r}. Set

$$\rho_{D_r} := \begin{cases} \frac{\sum_{e=0}^{r} mT(n,p,D_r - 2p^e)T(n,p,2p^e)}{\sum_{e=0}^{r} mT(n,p,D_r - 2p^e)} & p \neq 2, \\ \frac{\sum_{e=0}^{r} m\binom{n}{D_r - 2^{e+1}}\binom{D_r + 2^{e+1}}{2^{e+1}}}{\sum_{e=0}^{r} m\binom{n}{D_r - 2^{e+1}}} & p = 2. \end{cases}$$

We apply the Wiedemann algorithm [18] to find such the right kernel X. The number of operations required can be estimated as

$$\mathrm{Com}_{r,\mathrm{dp}}(\tau) := 3\rho_{D_r}(T(n,p,D_r) - T(o,p,D_r) + 1)^2 A(q),$$

here $A(q) := 2\log_2(q)^2 + \log_2(q)$ is the number of operations required for one $\mathbb{F}_q$ multiplication. For simplicity, sometimes we omit the τ from the notation.

Remark 3.7. When we apply the divided power wedge attack at different levels to recover the oil space, the complexity may vary depending on the chosen level r. Specifically, $\{D_r\}_{r=0}^{\infty}$ is decreasing but $\{\rho_{D_r}\}$ is increasing in general. We need to balance the trade-off between D_r and ρ_{D_r} to optimize the overall complexity. Furthermore, when $2p^r \geq o(p-1)$, $D_r = D_{r+1}$, since the solving degrees cannot exceed $o(p-1)$ and $(I_{r,\mathrm{dp}})_d = (I_{r+1,\mathrm{dp}})_d$ for all $d \leq o(p-1)$.

Divided power wedge reconciliation attack

For a UOV-like system $\{p_i\}_{i=1}^m \subseteq \langle L_1, \ldots, L_v \rangle \subseteq \mathbb{F}_q[x_1, \ldots, x_n]$ with parameters (n, v, m, q), we choose an integer $0 \leq k \leq o-1$. By setting $x_{n-k+1} = \cdots = x_n = 0$, we obtain the restricted UOV-like system of parameters $(n-k, v, m, q)$:

$$\{\tilde{p}_i\}_{i=1}^m \subseteq \langle \tilde{L}_1, \ldots, \tilde{L}_v \rangle \subseteq \mathbb{F}_q[x_1, \ldots, x_{n-k}].$$

where for $f \in \mathbb{F}_q[x_1, \ldots, x_n]$, $\tilde{f} := f(x_1, \ldots, x_{n-k}, 0, \ldots, 0) \in \mathbb{F}_q[x_1, \ldots, x_{n-k}]$. Note that

$$\mathbb{V}(\tilde{p}_1, \ldots, \tilde{p}_m) \supseteq \tilde{\mathcal{O}} := \mathbb{V}(\tilde{L}_1, \ldots, \tilde{L}_v) = \{\mathbf{a} \in \mathbb{F}_q^{n-k} : (\mathbf{a}, \mathbf{0}_k) \in \mathcal{O}\}.$$

Once a nontrivial oil element in $\tilde{\mathcal{O}}$ is found, we immediately get a nontrivial oil element in $\mathcal{O}$.

Fix $r \geq 0$ such that $v < 2m(1+\cdots+p^r)$. $\tilde{R} := \mathbb{F}_q[x_1, \ldots, x_{n-k}]/\langle x_1^p, \ldots, x_{n-k}^p \rangle$ is the polynomial algebra of $n-k$ variables. $\{\tilde{p}_i\}_{i=1}^m$ and their divided powers generate the ideal

$$(I_{\mathrm{rec}})_{r,\mathrm{dp}} := \langle \gamma_{p^e}(\tilde{p}_i) : 0 \leq e \leq r, 1 \leq i \leq m \rangle \subseteq \tilde{R}.$$

Under the assumption that

$$H_{\tilde{R}/(I_{\mathrm{rec}})_{r,\mathrm{dp}}}(t) \equiv \max\{H_{n-k,m,p,r}(t), F_{o-k,p}(t)\} \bmod t^{(o-k)(p-1)+1}, \tag{3.4}$$

the complexity of computing $\tilde{\mathcal{O}}$ is $\mathrm{Com}_{r,\mathrm{dp}}(\tau_k)$. Thus for the parameters τ, the improved estimate for the cost of the level-r divided power wedge reconciliation attack is

$$\mathrm{Com}^*_{r,\mathrm{dp}}(\tau) := \min\left\{ \mathrm{Com}_{r,\mathrm{dp}}(\tau_k) : 0 \leq k \leq o-1, v < \frac{(o-k-1)m}{2} \right\}.$$

We conduct experiments for small UOV systems with parameters (n, v, m, q). For some specific integers (k, r), we compute the Hilbert series of $\tilde{R}/(I_{\mathrm{rec}})_{r,\mathrm{dp}}$ as follows. For almost all parameters in our experiments, the experimental Hilbert series match the theoretical estimation Eq. (3.4).

Divided power wedge intersection attack

Write the public key $p_i(\mathbf{x})$ of a UOV-scheme with parameters (n, v, m, q) as $\mathbf{x}^\top M_i \mathbf{x}$, where M_i is an $n \times n$ symmetric matrix over $\mathbb{F}_q$ (or an upper triangular matrix if $2 \mid q$). Let L_1 and L_2 be randomly chosen invertible linear combinations of the M_i. If $\mathcal{O}$ is the oil space, then the $3m$ polynomials below

$$\begin{aligned} &p_{i,1}(\mathbf{x}) := p_i(L_1^{-1}\mathbf{x}), \quad p_{i,2}(\mathbf{x}) := p_i(L_2^{-1}\mathbf{x}), \\ &p_{i,3}(\mathbf{x}) := p_i(L_1^{-1}\mathbf{x} + L_2^{-1}\mathbf{x}) - p_i(L_1^{-1}\mathbf{x}) - p_i(L_2^{-1}\mathbf{x}). \end{aligned} \tag{3.5}$$

vanish on the subspace $L_1\mathcal{O} \cap L_2\mathcal{O}$. We only consider the case that $2o > v$. Then $L_1\mathcal{O} \cap L_2\mathcal{O}$ is nontrivial as

$$L_i\mathcal{O} \subseteq \mathcal{O}^\perp := \{\mathbf{a} \in \mathbb{F}_q^n : \mathbf{a}^\top \mathbf{b} = 0 \text{ for all } \mathbf{b} \in \mathcal{O}\}.$$

$\dim L_1\mathcal{O} \cap L_2\mathcal{O} \geq \dim L_1\mathcal{O} + \dim L_2\mathcal{O} - \dim \mathcal{O}^\perp = 2o - v$. So the complexity of classical intersection attack is dominated by the complexity of finding a $(2o-v)$-dimensional oil space of a UOV-like system with parameters $(n, 3v-2n, 3m-2,$

Table 1: The Hilbert series and the solving degrees for small UOV instances

Parameters (n, v, m, q, r, k)	Hilbert series (coefficients of $t^d : 0 \le d \le o(p-1)$)		Solving degrees $D_{r,\mathrm{dp}}(n-k, v, m, q)$
(9,6,3,16,1,0)	Theoretical	(1,9,33,57)	-
	Experimental	(1,9,33,57)	
(11,6,7,16,1,0)	Theoretical	(1,11,48,88,5,1)	4
	Experimental	(1,11,48,88,5,1)	
(9,5,4,16,2,0)	Theoretical	(1,9,32,48,1)	4
	Experimental	(1,9,32,48,1)	
(10,6,4,16,1,1)	Theoretical	(1,9,32,48)	-
	Experimental	(1,9,32,48)	
(10,7,3,9,1,0)	Theoretical	(1,10,52,180,453,852,1169)	-
	Experimental	(1,10,52,180,453,852,1169)	

q). Number $3m-2$ follows from the fact that the $3m$ polynomials in Eq. (3.5) are not linearly independent. They span a subspace of dimension $3m-2$.

According to the analysis of the previous section, we choose an parameter $0 \le k \le 2o-v-1$ and apply the level-r divided power wedge attack in the truncated algebra $\tilde{R}$ to $\mathrm{UOV}(n-k, 3v-2n, 3m-2, q)$. So the complexity of level-r divided power wedge intersection attack is given by

$$\mathrm{Com}_{r,\mathrm{dp},\mathrm{int}}(\tau) := \mathrm{Com}^*_{r,\mathrm{dp}}(n, 3v-2n, 3m-2, q).$$

Multi-divided power wedge attack

Now we consider the divided power wedge attack over the multi-graded truncated polynomial algebra

$$\mathcal{R} = \mathbb{F}_q[x_{i,j} : 1 \le i \le c, 1 \le j \le n_i]/\langle x^p_{i,j} : 1 \le i \le l, 1 \le j \le n_i\rangle.$$

For $1 \le i \le a$, there are m_i multi-homogeneous polynomials $\{f_{i,j}\}_{j=1}^{m_i}$ of multi-degree $\mathbf{d}_i \in \mathbb{N}^c$. $\mathbb{V}(f_{i,j} : 1 \le i \le c, 1 \le j \le m_i)$ contains the hidden oil subspace $\mathcal{O} := \mathcal{O}_1 \times \cdots \times \mathcal{O}_c$ with $\mathcal{O}_i \subseteq \mathbb{F}_q^{n_i}$ and $\dim \mathcal{O}_i = o_i$. For $\mathbf{r} = (r_1, \ldots, r_c)$, we conjecture that the multi-hilbert series of $\mathcal{R}/I_{\mathbf{r},\mathrm{dp}} = \mathcal{R}/\langle \gamma_{p^e}(f_{i,j}) : 0 \le e \le r_i\rangle$ satisfies

$$H_{R/I_{\mathbf{r},\mathrm{dp}}}(\mathbf{t}) \equiv \max\{\frac{\prod_{i=1}^c F_{n_i,p}(t_i)}{\prod_{i=1}^a F_{m_i,p}(\mathbf{t}^{\mathbf{d}_i})}, \prod_{i=1}^c F_{o_i,p^{r_i+1}}(t_i)\} \bmod \langle t_i^{o_i(p-1)} : 1 \le i \le c\rangle$$

Choose $\mathbf{D} = (D_1, \ldots, D_c) \in \mathbb{N}^c$ satisfying $[\mathbf{t}^{\mathbf{D}}]H_{R/I_{\mathbf{r},\mathrm{dp}}}(\mathbf{D}) = \prod_{i=1}^c T(o_i, p, D_i)$, then by the same method of Section 3.1 we can recover the oil space by the right

nullspace of the Macaulay matrix $M_{\mathbf{D}}$. The complexity can be estimated as

$$\mathrm{Com}_{\mathbf{r},\mathrm{dp}} := 3\rho_{\mathbf{D}} \prod_{i=1}^{c} (T(n_i, p, D_i) - T(o_i, p, D_i) + 1) A(q).$$

3.5 Comparsion between classical wedge attack

The level-0 divided power wedge attack is exactly the classical wedge attack. It is efficient for the UOV signature scheme, but is generally not applicable to the MAYO signature scheme. The public key of MAYO consists of UOV polynomials $\{p_i\}_{i=1}^m$ for some parameters (n, v, m, q), where $v \approx m \gg o$ and $q = 16$. Since o is small, the solving degree cannot be less than o. Furue and Ikematsu [8] proposed the p^ℓ-key recovery attack over the p^ℓ-truncated (graded) polynomial algebra

$$S := \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^{p^\ell}, \ldots, x_n^{p^\ell} \rangle.$$

The socle degree of S is $N_\ell = n(p^\ell - 1)$. $S_{N_\ell} = \mathbb{F}_q\{(x_1 \cdots x_n)^{p^\ell - 1}\}$. When $d > N_\ell$, $S_d = 0$. S has a duality property similar to that of R. For any $0 \leq d \leq N_\ell$ and $f \in S_{N_\ell - d}$, we associate a linear functional $\sigma_f : S_d \to \mathbb{F}_q$ such that, for $g \in S_d$,

$$fg = \sigma_f(g) x_1^{p^\ell - 1} \cdots x_n^{p^\ell - 1}.$$

Under the assignment $f \mapsto \sigma_f$, the following isomorphism holds for any graded ideal $I \subseteq S$:

$$I_{N_\ell - d}^{p^\ell, \perp} := \{f \in S_{N_\ell - d} : fI_d = 0\} \cong \{F \in S_d^\vee : F(I_d) = 0\}.$$

Let $\{p_i\}_{i=1}^m$ and $\{L_i\}_{i=1}^v$ be the public and private keys of UOV(n, v, m, q), respectively. They generate the ideals I and J of S, respectively. We can find oil elements with an estimated cost of

$$\mathrm{Com}_{p^\ell,\mathrm{rev}} := 3\binom{n+1}{2} (T(n, p^\ell, D_{[\ell]}) - T(o, p^\ell, D_{[\ell]}))^2 A(q).$$

Here, $T(n, p^\ell, D_{[\ell]}) = \dim S_{D_{[\ell]}}$ is the coefficient of $t^{D_{[\ell]}}$ in $F_{n,p^\ell}(t) = (\sum_{i=0}^{p^\ell - 1} t^i)^n$. And $1 \leq D_{[\ell]} \leq o(p^\ell - 1)$ is the smallest positive integer such that

$$I_{D_{[\ell]}} = J_{D_{[\ell]}} \neq 0 \quad (\Longleftrightarrow I_{N_\ell - D_{[\ell]}}^{p^\ell, \perp} = J_{N_\ell - D_{[\ell]}}^{p^\ell, \perp} = \langle (L_1 \cdots L_v)^{p^\ell - 1} \rangle_{N_\ell - D_{[\ell]}}).$$

In this situation, any $f \in I_{N_\ell - D_{[\ell]}}^{p^\ell, \perp}$ contains sufficient information about $L_1, \ldots, L_v$.

Determining the correct solving degree when $v \geq 2m$ is an interesting problem. The previous analysis does not directly apply, because for all $d \leq o(p^\ell - 1)$

$$I_{N_\ell - d}^{p^\ell, \perp} \supseteq \langle (L_1 \cdots L_v)^{p^\ell - 1}, (p_1 \cdots p_m)^{p^\ell - 1} \rangle_{N_\ell - d} \supsetneq J_{N_\ell - d}^{p^\ell, \perp}.$$

Furue and Ikematsu [10] showed that, for any d satisfying the following conditions:

$$I_{N_\ell-d}^{p^\ell,\perp} = \langle (L_1 \cdots L_v)^{p^\ell-1}, (p_1 \cdots p_m)^{p^\ell-1} \rangle_{N_\ell-d},$$
$$\langle (L_1 \cdots L_v)^{p^\ell-1}, (p_1 \cdots p_m)^{p^\ell-1} \rangle_{N-d+1} \neq \langle (p_1 \cdots p_m)^{p^\ell-1} \rangle_{N_\ell-d+1},$$

where they hold only if $n(p^\ell-1) - \lceil \frac{n+2m}{2}(p^\ell-1) \rceil \leq d \leq o(p^\ell-1)$, oil elements can be found using the following steps:

Step 1 Compute a nonzero $f \in I_{N_\ell-d}^{p^\ell,\perp}$.
Step 2 Determine the nontrivial subspace $\mathcal{L} := \{L \in S_1 : Lf \in \langle p_1 \cdots p_m \rangle^{p^\ell-1}\}$.
Step 3 Since $L_i \in \mathcal{L}$, $\mathbb{V}(L : L \in \mathcal{L}) \subseteq \mathcal{O} = \mathbb{V}(L_1, \ldots, L_v)$.

The complexity of Step 2 is at most $3\binom{n+1}{2}(\dim R_d/I_d)^2 A(q)$, and the complexity of Step 3 is $O((T(n, p^\ell, (n-2m)(p^\ell-1)-d+1)+n)^3)$. For many parameter choices (n, v, m, q) with $v \geq 2m$, the solving degree d is very large, the resulting complexity is higher than that of our divided power wedge attack. See Table 2.

Table 2: Complexity of Level-1 divided power wedge attack for the case $v \geq 2m$

Parameters (n, v, m, q)	Level-1 divided power wedge attack $(k, D_{1,\mathrm{dp}}, \log_2 \mathrm{Com}_{1,\mathrm{dp}})$	Extended wedge attack [10] $(k, D, \log_2 \mathrm{Com})$
(70,40,16,11)	(17,129,359)	(16,140,394)
(70,55,20,11)	(0,147,443)	-
(90,50,20,7)	(24,94,344)	(23,102,385)

4 Applications

In this section, we apply the preceding framework to UOV, QR-UOV, MAYO, and SNOVA.

UOV

In Table 3, we calculate the complexity of the divided power wedge reconciliation attack and the divided power wedge intersection attack at levels 0,1, and 2 for the four parameters of UOV, see [4]. For `uov-Ip`, `uov-III`, `uov-V`, we get lower attack complexities than previously known results.

Actually, compared with other attacks against UOV, the classical wedge attack (level-0 divided power wedge attack) is efficient. Increasing the level of the divided power wedge attack from 0 to 1 reduces the solving degree ($D_1 < D_0$). Consequently, the Macaulay matrix M_{D_1} has fewer columns, while its density remains manageable. This explains why the level-1 divided power wedge attack has lower complexity than the attacks at the other levels.

Table 3: Complexity of r-level divided power wedge attack against UOV

Attack level r	Reconciliation $(k, D_r, \log_2 \rho_{D_r}, \log_2 \mathrm{Com}^*_{r,\mathrm{dp}})$	Intersection $(k, D_r, \log_2 \rho_{D_r}, \log_2 \mathrm{Com}^*_{r,\mathrm{dp}})$
`uov-Ip`: $n = 119, v = 74, m = 45, q = 256$ $\log_2$ complexity of the best known attack: 144		
$r = 0$	(25,20,11,153)	(1,15,12,144)
$r = 1$	(27,18,15,149)	(2,14,15,142)
$r = 2$	(27,18,19,153)	(2,14,17,143)
uov-Is: $n = 160, v = 96, m = 64, q = 16$ $\log_2$ complexity of the best known attack: 143		
$r = 0$	(42,22,12,175)	(16,16,13,158)
$r = 1$	(43,21,17,175)	(17,15,16,155)
$r = 2$	(43,21,22,180)	(17,15,19,158)
`uov-III`: $n = 193, v = 121, m = 72, q = 256$ $\log_2$ complexity of the best known attack: 213		
$r = 0$	(43,29,12,226)	(1,22,13,212)
$r = 1$	(45,27,18,222)	(2,21,18,211)
$r = 2$	(45,27,25,229)	(2,21,23,216)
`uov-V`: $n = 259, v = 163, m = 96, q = 256$ $\log_2$ complexity of the best known attack: 272		
$r = 0$	(58,38,13,295)	(1,28,14,271)
$r = 1$	(62,34,20,282)	(2,27,20,270)
$r = 2$	(62,34,28,291)	(2,27,27,277)

MAYO

A MAYO public key is defined by a UOV-like parameter set (n, v, m, q). The oil space has very small dimension, the divided-power attack does not apply to the main MAYO parameter sets [3]. This is because for these main parameters, no matter which the level r we select, $o(p-1) < C_{n,m,p,r}$ and

$$[t^d]H_{n,m,p,r}(t) \gg [t^d]F_{o,p}(t) \text{ for all } 1 \leq d \leq o(p-1).$$

We cannot determine the solving degree D_r.

However, we find that for some additional parameters of the Security Level III and V with the dimension of the oil space ≥ 30, our level-1 divided power wedge attack lowers the estimated bit security by 3-14 bits, see Table 4.

Table 4: Complexity of divided power wedge attack against MAYO

Parameters (n, v, m, q)	Direct attack	Wedge attack	Our attack $(k, \log_2 \mathrm{Com}^*_{1,\mathrm{dp}})$
Security level III			
(161,129,94,16)	217	214	(7,215)
(174,128,92,16)	217	219	(21,214)
Security level V			
(212,179,130,16)	292	-	(1,278)
(222,179,128,16)	292	288	(11,278)
(242,179,126,16)	291	287	(31,278)

QR-UOV

A QR-UOV public key with parameters (v, m, q, l) can be identified with a plain UOV-system with parameters $((v+m)/l, v/l, m, q^l)$. q is odd. For the main parameters with $q = 127$ in [9], the solving degree of the level-0 divided power wedge attack is less than q, so the complexity of the divided power wedge attack does not depend on the level r, as we mention in Remark 3.7. For the addition parameters with $q = 7$, compared to level-0 divided power wedge attack, level-1 divided power wedge does not lower the solving degree. But the density of the Macaulay matrix increase, which leads to the higher complexity, see Table 5.

SNOVA

It remains to consider SNOVA. By the lifting technique of [6, 13], a SNOVA instance with parameters (v, o, q, ℓ) can be regarded as a multi-UOV problem

Table 5: Complexity of divided power wedge attack against QR-UOV

Parameters (q, v, m, ℓ)	Best attack	Level-0 $(k, \log_2 \mathrm{Com}^*_{0,\mathrm{dp}})$	Level-1 $(k, \log_2 \mathrm{Com}^*_{1,\mathrm{dp}})$
(127,154,54,3)	150	(17,164)	(17,164)
(127,820,90,10)	211	(8,211)	(8,211)
(127,1040,110,10)	273	(10,276)	(10,276)

over the extension field $\mathbb{F}_{q^\ell}$. Especially, when $\ell = 2$, SNOVA$(v, o, q, 2)$ turns into a multi-UOV problem with o public polynomials of multi-degree $(1, 0)$ and $(0, 1)$, $2o$ public polynomials of multi-degree $(1, 1)$ (Here both of these polynomials lie in $\mathbb{F}_{q^\ell}[x_{i,j} : i = 1, 2, 1 \leq j \leq n]$). Table 6 summarizes the complexity of multi-divided power wedge attack on some SNOVA parameter sets proposed in the third round of NIST PQC standardization process for additional signatures [17].

Table 6: Complexity of divided power wedge attack against SNOVA

Parameters (q, v, m, ℓ)	Best attack	Multi-divided power wedge attack $(\mathbf{k}, \mathbf{r}, \mathbf{D}, \log_2 \mathrm{Com}_{\mathbf{r},\mathrm{dp}})$
(16,50,17,2)	158	((14,0),(0,1,1),(3,17),156)
(16,76,25,2)	223	((23,0),(0,1,1),(2,25),205)
(16,104,33,2)	287	((31,0),(0,0,1),(2,33),263)

5 Conclusions and Future Work

Conclusion Building on the characteristic-two divided-power construction of Merz and Ran [14, Section 3], we have developed a family of wedge attacks in p-truncated polynomial rings over finite fields of arbitrary characteristic. The level controls the additional public relations, allowing the solving degree and average Macaulay row weight to be balanced. We give an explicit recovery criterion and predict the relevant degrees through conjectural Hilbert-series formulas. A multigraded extension applies the same approach to lifted SNOVA systems. For selected parameter sets with $v \geq 2m$, our attack has lower estimated costs than the extended wedge attack [10].

Our applications cover UOV, MAYO, QR-UOV, and SNOVA. Relative to the comparison baselines reported in this paper, the estimates give reductions of

2 bits for three UOV parameter sets, 3–14 bits for selected additional MAYO parameter sets, and 2–24 bits for selected additional SNOVA parameter sets. They do not improve the costs for the QR-UOV parameters studied. These complexity estimates remain conditional on the conjectured Hilbert-series predictions, including the multigraded prediction used for SNOVA.

Future Work Our complexity estimates assume that the actual Hilbert series agrees with the prediction in Conjecture 3.6 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 [14] 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 [14, §3]. The authors take full responsibility for the content and any errors that may remain.

Bibliography

[1] Berthelot, P., Ogus, A.: F-isocrystals and de Rham cohomology. I. Invent. Math. **72**(2), 159–199 (1983). https://doi.org/10.1007/BF01389319, https://doi.org/10.1007/BF01389319

[2] 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)

[3] Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO specification document - round 3 (2026), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/mayo_specification.pdf

[4] Beullens, W., Chen, M.S., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B.Y., Schmidt, D., Shih, C.J., Tao, C., Yang, B.Y.: Uov specification document - round 3 (2026), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/UOV_Specification_R3.pdf

[5] 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)

[6] Cabarcas, D., Li, P., Verbel, J., Villanueva-Polanco, R.: Improved attacks for SNOVA by exploiting stability under a group action. In: Annual International Cryptology Conference. pp. 358–389. Springer (2025)

[7] Ding, J., Li, P., Tao, S.: Revisiting the wedge attack on uov problem. Cryptology ePrint Archive (2026)

[8] Furue, H., Ikematsu, Y.: Key recovery attacks on uov using p^l-truncated polynomial rings. Cryptology ePrint Archive (2026)

[9] Furue, H., Ikematsu, Y., Hoshino, F., Takagi, T., Sugai, R., Amagasa, H., Homma, N., Ueno, R., Sakata, K., Kosuge, H., Yamakoshi, K., Kinjo, K., Akiyama, R., Nakamura, S., Orihara, S., Kawahara, Y.: QR-UOV specification document - round 3 (2026), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/qruov_specifications.pdf

[10] 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

[11] 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)

[12] 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)

[13] Li, P., Ding, J.: Cryptanalysis of the SNOVA signature scheme. In: International Conference on Post-Quantum Cryptography. pp. 79–91. Springer (2024)

[14] 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
[15] 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)
[16] Tseng, P.E., Wang, L.C., Li, P., Kuan, Y.L.: Investigating the wedge map on SNOVA. Cryptology ePrint Archive (2026)
[17] Wang, L.C., Chou, C.Y., Ding, J., Guo, H., Kuan, Y.L., Leegwater, J.A., Li, M.S., Li, P., Tseng, B.S., Tseng, P.E., Wang, C.C.: SNOVA specification document - round 3 (2026), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/SNOVA_specifications.pdf
[18] Wiedemann, D.: Solving sparse linear equations over finite fields. IEEE Transactions on Information Theory **32**(1), 54–62 (1986). https://doi.org/10.1109/TIT.1986.1057137

A Divided power structures

Now we prove that the maps $\{\gamma_k\}_{k\geq 0}$ we gave in Lemma 3.4 define a divided power structure on $\mathfrak{m} = \langle x_1, \ldots, x_n\rangle_R \subseteq R = \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p\rangle$. By [1, A.1 Proposition], it suffices to prove that:

1. $p!\gamma_p(f) = f^p$ for all $f \in \mathfrak{m}$.
2. $\gamma_p(uf) = u^p\gamma_p(f)$ for all $f \in \mathfrak{m}$ and $u \in R$.
3. $\gamma_p(f+g) = \sum_{i=0}^{p} \gamma_i(f)\gamma_{p-i}(g)$ for all $f, g \in \mathfrak{m}$.

The first claim is obvious. Then we prove the third claim. Assume that $f = \sum_{i=1}^{r} k_i M_i$ and $g = \sum_{i=1}^{r} l_i M_i$, here $\{M_i\}_{i=1}^{r}$ are distinct monomials and $k_i, l_i \in \mathbb{F}_q$ may equal to 0. Set $\mathcal{A}_r(d) := \{\alpha \in \mathbb{N}^r : |\alpha| = d, 0 \leq \alpha_i \leq p-1\}$.

$$\begin{aligned}
\sum_{i=0}^{p} \gamma_i(f)\gamma_{p-i}(g) &= \sum_{i=0}^{p} \left(\sum_{\alpha\in\mathcal{A}_r(i)} \prod_{j=1}^{r} \frac{k_j^{\alpha_j}}{\alpha_j!} M_j^{\alpha_j} \right) \left(\sum_{\beta\in\mathcal{A}_r(p-i)} \prod_{j=1}^{r} \frac{l_j^{\beta_j}}{\beta_j!} M_j^{\beta_j} \right) \\
&= \sum_{i=0}^{p} \sum_{\alpha\in\mathcal{A}_r(i),\beta\in\mathcal{A}_r(p-i)} \prod_{j=1}^{r} \frac{k_j^{\alpha_j} l_j^{\beta_j}}{\alpha_j!\beta_j!} M_j^{\alpha_j+\beta_j} \\
&= \sum_{\beta\in\mathcal{A}_r(p)} \left(\sum_{\alpha\in\mathbb{N}^r,\alpha\leq\beta} \prod_{j=1}^{r} \frac{k_j^{\alpha_j} l_j^{\beta_j-\alpha_j}}{\alpha_j!(\beta_j-\alpha_j)!} \right) M_j^{\beta_j} \\
&= \sum_{\beta\in\mathcal{A}_r(p)} \prod_{j=1}^{r} \frac{(k_j+l_j)^{\beta_j}}{\beta_j!} M_j^{\beta_j} = \gamma_p(f+g).
\end{aligned}$$

Note that if $k < p, \gamma_k(f) = \frac{f^k}{k!}$. Thus for $k < p$, $r \geq 1$ and $f_1, \ldots, f_r \in \mathfrak{m}$,

$$\gamma_k\left(\sum_{i=1}^{r} f_i\right) = \sum_{\alpha\in\mathbb{N}^r,|\alpha|=k} \prod_{i=1}^{r} \gamma_{\alpha_i}(f_i).$$

We will show that this equation holds for $k = p$ by induction on r. $r = 2$ has already been proved. Now suppose that the equation holds for $r = r_0 - 1$. Then for $\{f_i\}_{i=1}^{r_0} \in \mathfrak{m}$,

$$\begin{aligned}
\gamma_p\left(\sum_{i=1}^{r_0} f_i\right) &= \sum_{i=0}^{p} \gamma_i\left(\sum_{j=1}^{r_0-1} f_j\right) \gamma_{p-i}(f_{r_0}) \\
&= \sum_{i=0}^{p} \sum_{\alpha\in\mathbb{N}^{r_0-1},|\alpha|=i} \left(\prod_{j=1}^{r_0-1} \gamma_{\alpha_j}(f_j) \right) \gamma_{p-i}(f_{r_0}) \\
&= \sum_{\alpha\in\mathbb{N}^{r_0},|\alpha|=p} \prod_{j=1}^{r_0} \gamma_{\alpha_j}(f_j).
\end{aligned}$$

For the second claim, if $u \in \mathbb{F}_q$ or u is a monomial, then we see that $\gamma_p(uf) = u^p\gamma_p(f)$ by the definition of γ_p. For general u, write $u = \sum_{i=1}^{r} u_i$ with $u_1 = u(\mathbf{0}) \in \mathbb{F}_q$ and u_i $(i > 1)$ are monomials. In the latter case, $\gamma_p(uf) = u^p\gamma_p(f) = 0$. So

$$\begin{aligned}\gamma_p(uf) &= \sum_{\alpha\in\mathbb{N}^r,|\alpha|=p}\prod_{i=1}^{r}\gamma_{\alpha_i}(u_if) = \gamma_p(u_1f) + \sum_{\alpha\in\mathcal{A}_r(p)}\prod_{i=1}^{r}\gamma_{\alpha_i}(u_if)\\ &= u_1^p\gamma_p(f) + \sum_{\alpha\in\mathcal{A}_r(p)} f^p\prod_{i=1}^{r}\frac{u_i^{\alpha_i}}{\alpha_i!} = u_1^p\gamma_p(f) = u^p\gamma_p(f).\end{aligned}$$

Here we use the fact that for any $f \in \mathbb{F}_q[x_1, \ldots, x_n]$, $f^p \in \langle x_1^p, \ldots, x_n^p\rangle$.