Extending the Schur–Macaulay Framework for UOV Attacks via Truncated Polynomial 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. Building on the ideal-inclusion framework of Ding, Li, and Tao [6], we reformulate and extend the Schur–Macaulay approach of Merz and Ran [15] using multigraded p-truncated polynomial algebras over arbitrary finite fields. The construction admits a reconciliation interpretation involving multiple oil vectors. Combining divided powers with transfer relation spaces, we derive quotient-dimension criteria for recovering graded components of the hidden oil-space ideal. Hilbert-series conjectures and structured rank predictions underpin the complexity estimates. Applications to NIST third-round submissions yield reductions of up to 61 bits in estimated attack costs for selected additional MAYO parameter sets and improvements over the classical wedge attack for selected QR-UOV parameters. Certain alternating parameters provided by authors in MAYO fail to meet the NIST security level III and V.

Keywords: UOV problem · divided powers · truncated polynomial rings

1 Introduction

The Unbalanced Oil and Vinegar (UOV) signature scheme [13] 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 [3]. 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.

The security of UOV-based constructions has been studied from a variety of cryptanalytic perspectives. Classical key recovery attacks include the Kipnis–Shamir (KS) attack [14], which exploits linear-algebraic properties of the hidden oil space, as well as subsequent reconciliation [7], intersection [3], and rectangular MinRank attacks [3] that recover partial information about the secret structure or reduce the complexity of the associated MQ systems. These developments have significantly influenced the selection of secure parameters and the design of more sophisticated UOV-based schemes.

Recently, Ran [17] introduced the wedge attack using the exterior algebra. Subsequent work developed formulations in reduced symmetric algebras [12] and truncated polynomial rings [8]. Wedge attacks were also independently applied to SNOVA in [5] and [18].

Ding, Li, and Tao [6] reformulated the recovery mechanism of wedge attack in terms of Macaulay right kernels and the ideal inclusion $I \subseteq J$, where I is generated by the public quadratic relations and J by linear forms defining the hidden oil space. Their formulation applies over fields of arbitrary characteristic. A central recovery criterion is the existence of a degree $d > 0$ such that $I_d = J_d$ and the degree-d quotient is nonzero. Under these conditions, a normalized basis of the Macaulay right kernel reveals the linear forms defining the oil space. This viewpoint provides a common algebraic description of the attacks in [17, 12, 8, 5, 18]: they realize the inclusion $I \subseteq J$ in different graded algebras.

The equality $I_d = J_d$, however, may fail at every relevant degree when $v > 2m$, where v is the number of vinegar variables and m is the number of public equations. Ikematsu and Furue [10] address this regime by analyzing the structure of Macaulay kernels beyond the case in which the standard recovery criterion holds.

Merz and Ran [15] extended the exterior algebra attack in two directions. First, in characteristic two, they derived additional equations from divided powers of the public polar forms. Second, they extended the linear algebra computations from a single exterior power to arbitrary Schur modules. The first of these directions was extended to arbitrary characteristic in [**?**].

Contributions. Our contributions are threefold.

1. Building on the ideal-inclusion framework of Ding, Li, and Tao [6], we interpret the Schur–Macaulay approach of Merz and Ran [15] as a reconciliation attack involving multiple oil vectors.
2. We develop a multigraded construction over arbitrary finite fields using p-truncated polynomial algebras, divided powers, and transfer relation spaces. For two or more variable sets, we establish quotient-dimension criteria for recovering graded components of the hidden oil-space ideal and formulate Hilbert-series conjectures and structured rank predictions.
3. Under these predictions, we obtain lower estimated attack complexities for selected NIST third-round QR-UOV and additional MAYO parameter sets.

Code: The magma implementation of the experiments, complexity estimators can be found in the repository below

https://github.com/Yulan24/Extended-Schur-Macaulay

Outline § 2 reviews UOV, wedge attacks, and divided powers on truncated polynomial rings. § 3 develops the bigraded construction and its rank analysis, and § 4 extends it to multiple sets of variables. § 5 explains the relation between the work developed in this paper and the exterior-algebra approach of Merz and Ran [15]. § 6 is devoted to experiments, cost analysis, and applications. § 7 concludes with directions for future work.

2 Preliminaries

2.1 Unbalanced Oil and Vinegar (UOV)

The UOV signature scheme was introduced by Kipnis, Patarin, and Goubin in 1999 [13]. We follow Beullens' description of UOV in terms of its hidden oil space [3]. 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

$$\mathcal{P}'(\mathbf{x}, \mathbf{y}) := \mathcal{P}(\mathbf{x}+\mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y}) = (q_1(\mathbf{x}, \mathbf{y}), \ldots, q_m(\mathbf{x}, \mathbf{y})).$$

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 $v \leq \frac{m(o-1)}{2}$.

The UOV problem underlies signature schemes such as QR-UOV, MAYO, and SNOVA, which have advanced alongside UOV to the third round of NIST's standardization process for additional digital signatures [16], highlighting the cryptographic significance of this problem.

2.2 Wedge attack

Ran [17] introduced the wedge attack on UOV exterior algebras. 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 $\mathcal{V} = L_1 \wedge \cdots \wedge L_v$ 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 $\mathcal{V}$ up to a nonzero scalar, and hence the hidden oil space.

Jin et al. [12] 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. [18] studied the wedge map through a lifting to an extension field.

Ding, Li, and Tao [6] gave a common description of these recovery mechanisms of wedge attack through Macaulay right kernels and the ideal inclusion $I \subseteq J$, valid in arbitrary characteristic. Ikematsu and Furue [10] studied Macaulay kernels beyond the standard recovery criterion, addressing the regime $v > 2m$. Merz and Ran [15] strengthened the exterior algebra attack by adding divided-power relations and extending the linear algebra computations to arbitrary Schur modules. Using the language developed by Ding, Li, and Tao [6] and drawing on insights from Merz et al. [15], we extend the latter's work to truncated polynomial algebras over any finite field.

2.3 Divided powers on p-truncated polynomial rings

Divided powers generalize the operations $f^r/r!$ to positive characteristic. We recall their defining identities and the properties needed for our construction.

Definition 2.3. *[2, Definition 3.1] 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 $k_1, k_2 > 0, f \in I$, $\gamma_{k_1}(\gamma_{k_2}(f)) = \frac{(k_1 k_2)!}{k_1!(k_2)!^{k_1}}\gamma_{k_1 k_2}(f)$.*

q is a power of a prime number p. We consider the p-truncated polynomial algebra $R := \mathbb{F}_q[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p \rangle$. We usually do not distinguish the polynomial $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ from its coset $\overline{f}$ in R. We will use the following vector notation:

- $\mathcal{A} := \{\alpha = (\alpha_1, \ldots, \alpha_n) \in \mathbb{N}^n \setminus \{(0, \ldots, 0)\} : 0 \leq \alpha_i < p\}$.
- For each n-tuple $\alpha \in \mathbb{N}^n$, $|\alpha| := \sum_{i=1}^n \alpha_i$, $\mathbf{x}^\alpha := \prod_{i=1}^n x_i^{\alpha_i}$.
- For polynomials $f_1, \ldots, f_m \in \mathbb{F}_q[x_1, \ldots, x_n]$, $\langle f_1, \ldots, f_m \rangle_R$ denotes the ideal in R generated by $\overline{f_i}$.

Lemma 2.4. *There is a divided power structure on $\mathfrak{m} := \langle x_1, \ldots, x_n \rangle_R \subseteq R$. Any $f \in \mathfrak{m}$ can be uniquely written as $f = \sum_{\alpha \in \mathcal{A}} c_\alpha \mathbf{x}^\alpha$ with $c_\alpha \in \mathbb{F}_q$. We set*

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

If $b_\alpha = 0$, then we require $b_\alpha! = 1$. $\{\gamma_k\}_{k \geq 0}$ is a divided power structure on $\mathfrak{m}$.

For the proof of this lemma see Appendix A. In particular, $\gamma_k(f) = \frac{f^k}{k!}$ when $0 \leq k < p$. For future reference, we collect in the following remarks some easy computations of $\gamma_k(f)$.

Remark 2.5. If the p-adic expansion of k is $\sum_{i=0}^s k_i p^i$ with $k_i \in \{0, \ldots, p-1\}$, then repeatedly using Lucas's Theorem and the fourth condition in Definition 2.3, we obtain

$$\gamma_k(f) = \prod_{i=0}^s \frac{(\gamma_{p^i}(f))^{k_i}}{k_i!}.$$

This implies that $\langle \gamma_k(f) : k \geq 1 \rangle = \langle \gamma_{p^s}(f) : s \geq 0 \rangle$.

Remark 2.6. For any $f = \sum_{i=1}^m u_i g_i$ with $u_i, g_i \in \mathfrak{m}$ and for $k > 0$,

$$\gamma_k(f) = \sum_{\alpha \in \mathbb{N}^m, |\alpha| = k} \prod_{i=1}^m \gamma_{\alpha_i}(u_i g_i) = \sum_{\alpha \in \mathbb{N}^m, |\alpha| = k} \prod_{i=1}^m u_i^{\alpha_i} \gamma_{\alpha_i}(g_i) \in \langle u_i : 1 \leq i \leq m \rangle.$$

The second equality can be proven by induction on m. It holds for $m = 2$, which is exactly the fifth condition of Definition 2.3.

3 A bigraded truncated-polynomial construction

Let $k = \mathbb{F}_q$, where q is a power of the prime p. Consider a UOV-like system consisting of m quadratic polynomials $p_1, \ldots, p_m \in k[x_1, \ldots, x_n]$, each vanishing on an o-dimensional subspace $\mathcal{O}$. The reconciliation attack [7] begins by finding a nonzero oil element $\mathbf{u} \in \mathcal{O}$ satisfying

$$p_i(\mathbf{u}) = 0, \quad 1 \leq i \leq m.$$

Once such an element $\mathbf{u}$ has been found, we seek a second oil element $\mathbf{v}$, linearly independent of $\mathbf{u}$, by solving the system

$$\begin{cases} p_i(\mathbf{v}) = 0, & 1 \le i \le m, \\ q_i(\mathbf{u}, \mathbf{v}) = 0, & 1 \le i \le m, \end{cases}$$

where

$$q_i(\mathbf{x}, \mathbf{y}) = p_i(\mathbf{x} + \mathbf{y}) - p_i(\mathbf{x}) - p_i(\mathbf{y}), \quad 1 \le i \le m.$$

Repeating this procedure, with polar constraints imposed against all previously found oil elements, we seek to recover a basis of $\mathcal{O}$. Alternatively, we can search for multiple oil elements simultaneously.

Instead of considering the ideal $\langle p_1, \cdots, p_m \rangle$, one can consider the ideal

$$\langle p_i(\mathbf{x}), p_i(\mathbf{y}), q_i(\mathbf{x}, \mathbf{y}) : 1 \le i \le m \rangle \subseteqq k[\mathbf{x}, \mathbf{y}].$$

In the following, we implement this idea using a bigraded truncated-polynomial algebra. This yields a truncated-polynomial formulation of the construction of Merz and Ran [15, §4] over a field of positive characteristic. We first describe the hidden oil space using a bigraded quotient, then strengthen the public relations using divided powers and transfer maps. Finally, we analyze the resulting quotient dimensions to obtain a criterion for recovering a component of the hidden ideal.

3.1 Truncation and the oil quotient

Let $\mathcal{O} \subseteq k^n$ be the hidden oil space of dimension o. Choose independent linear polynomials $L_1, \ldots, L_v$, with $v = n - o$, such that

$$\mathcal{O} = \{\mathbf{u} \in k^n : L_1(\mathbf{u}) = \cdots = L_v(\mathbf{u}) = 0\}.$$

We assume that each public quadratic has a decomposition $p_i = \sum_{j=1}^{v} u_{i,j} L_j$, where the $u_{i,j}$ are linear polynomials. For $1 \le i \le m$, the public equations in the two sets of variables and their polar equations therefore take the form

$$\begin{aligned} p_i(\mathbf{x}) &= \sum_{j=1}^{v} u_{i,j}(\mathbf{x}) L_j(\mathbf{x}), \\ p_i(\mathbf{y}) &= \sum_{j=1}^{v} u_{i,j}(\mathbf{y}) L_j(\mathbf{y}), \\ q_i(\mathbf{x}, \mathbf{y}) &= \sum_{j=1}^{v} \big(u_{i,j}(\mathbf{x}) L_j(\mathbf{y}) + u_{i,j}(\mathbf{y}) L_j(\mathbf{x}) \big). \end{aligned}$$

All these polynomials vanish on $\mathcal{O} \oplus \mathcal{O} \subseteq k^{2n}$, which is defined by the linear equations $L_j(\mathbf{x}) = L_j(\mathbf{y}) = 0$ for $1 \le j \le v$.

We work in the p-truncated polynomial algebra

$$R := k[x_1, \dots, x_n, y_1, \dots, y_n]/\langle x_1^p, \dots, x_n^p, y_1^p, \dots, y_n^p \rangle.$$

Give R the bigrading determined by $\deg x_i = (1,0)$ and $\deg y_i = (0,1)$. For $a, b \in \mathbb{N}$, its bihomogeneous component is

$$R_{a,b} = \operatorname{span}_k\{\mathbf{x}^\alpha \mathbf{y}^\beta : |\alpha| = a,\ |\beta| = b,\ 0 \le \alpha_i, \beta_i < p\}.$$

Components with a negative degree are understood to be zero. Define

$$T(d,p,r) := (1 + t + \cdots + t^{p-1})^d [t^r],$$

with value zero outside $0 \le r \le d(p-1)$. Then $\dim_k R_{a,b} = T(n,p,a)T(n,p,b)$. An element of $R_{a,b}$ is called bihomogeneous of bidegree (a,b). An ideal $I \subseteq R$ generated by bihomogeneous elements is bigraded, and it decomposes as

$$I = \bigoplus_{a,b \in \mathbb{N}} I_{a,b}, \qquad I_{a,b} := I \cap R_{a,b}.$$

Using the same notation for polynomials and their images in R, define

$$\mathcal{I} := \langle p_i(\mathbf{x}), p_i(\mathbf{y}), q_i(\mathbf{x}, \mathbf{y}) : 1 \le i \le m \rangle_R,$$
$$\mathcal{J} := \langle L_j(\mathbf{x}), L_j(\mathbf{y}) : 1 \le j \le v \rangle_R.$$

Both ideals are bigraded, and the decompositions above give

$$\mathcal{I} \subseteq \mathcal{J}.$$

The quotient by $\mathcal{J}$ can be described explicitly in oil coordinates. Choose a basis of $\mathcal{O}$ and let $U \in k^{n \times o}$ be the matrix whose columns are the basis vectors. Introduce $\mathbf{s} = (s_1, \dots, s_o)$ and $\mathbf{t} = (t_1, \dots, t_o)$, and put

$$S := k[s_1, \dots, s_o, t_1, \dots, t_o]/\langle s_1^p, \dots, s_o^p, t_1^p, \dots, t_o^p \rangle,$$

where $\deg s_j = (1,0)$ and $\deg t_j = (0,1)$.

Proposition 3.1. *The substitution map*

$$\rho_R : R \longrightarrow S, \qquad F(\mathbf{x}, \mathbf{y}) \longmapsto F(U\mathbf{s}, U\mathbf{t})$$

is surjective, preserves bidegrees, and has kernel $\mathcal{J}$. Thus

$$R/\mathcal{J} \cong S. \tag{3.1}$$

In particular,

$$\dim_k (R/\mathcal{J})_{a,b} = \dim_k S_{a,b} = T(o,p,a)T(o,p,b).$$

The largest degree in either set of oil variables is $d_o := o(p-1)$, and $T(o,p,d_o) = 1$.

Proof. It's easy to check or see the proof of [6, Theorem 4.1]. □

The inclusion $\mathcal{I} \subseteq \mathcal{J}$ is the starting point for our construction. Following the divided-power approach of [**?**], we next enlarge $\mathcal{I}$ and compare the two ideals in each bidegree after adding transfer relations. These relations correspond to the maps θ_ℓ in [15].

3.2 Divided powers and transfer relations in positive characteristic

Consider the maximal ideal $\mathfrak{m} := \langle x_1, \ldots, x_n, y_1, \ldots, y_n \rangle$ of R. Let $\{\gamma_k : \mathfrak{m} \to R\}_{k\in\mathbb{N}}$ be the divided power structure on $\mathfrak{m}$ (see Lemma 2.4). Fix $r > 0$, define

$$\mathcal{I}_{r,\mathrm{dp}} := \left\langle \begin{matrix} \gamma_{p^s}(p_i(\mathbf{x})),\ \gamma_{p^s}(q_i(\mathbf{x},\mathbf{y})),\ \gamma_{p^s}(p_i(\mathbf{y})) : \\ 1 \le i \le m, \quad 0 \le s \le r. \end{matrix} \right\rangle_R .$$

Its generators have bidegrees $(2p^s, 0)$, (p^s, p^s), and $(0, 2p^s)$. Remark 2.5 shows that all positive divided powers of the public quadratics generate this same ideal. By Remark 2.6, we have

$$\mathcal{I} \subseteq \mathcal{I}_{r,\mathrm{dp}} \subseteq \mathcal{J}.$$

Let t be an indeterminate. $\{\gamma_k\}_{k\ge 0}$ can be naturally extended to $\mathfrak{m}[t] \subseteq R[t]$ satisfying $\gamma_i(ut^j) = \gamma_i(u)t^{ij}$ for $u \in \mathfrak{m}$, which makes $R[t]$ a divided power ring.

We define an algebra endomorphism $\sigma : R[t] \to R[t]$ by sending $f(\mathbf{x}, \mathbf{y})$ to $f(\mathbf{x} + t\mathbf{y}, \mathbf{y})$. For each $\ell \ge 0$, we have the $\mathbb{F}_q$-linear map

$$D_\ell : R \to R, f \mapsto \sigma(f)[t^\ell].$$

Especially, $\{D_\ell\}_{\ell\ge 0}$ is a Hasse-Schmidt derivation on R, which means that the sequence of these linear maps satisfying

$$D_0 = \mathrm{id}_R, \quad D_\ell(fg) = \sum_{i=0}^{\ell} D_i(f)D_{\ell-i}(g) \text{ for all } f, g \in R.$$

They also have the property that

$$fD_\ell(g) = \sum_{i=0}^{\ell} (-1)^{\ell-i} D_i(D_{\ell-i}(f)g) \text{ for all } f, g \in R. \tag{3.2}$$

It is easy to verify that D_ℓ is a graded map of degree $(-\ell, \ell)$, i.e., $D_\ell(R_{a+\ell,b-\ell}) \subseteq R_{a,b}$. The degree-$(a, b)$ transfer relation space is defined to be

$$\Xi_{a,b} = \sum_{\ell\ge 0} D_\ell(R_{a-\ell,b+\ell}) \subseteq R_{a,b}.$$

Remark 3.2. 1. The operators are independent of the public quadratics and of the hidden oil space.
2. When $\mathbb{F}_q$ has characteristic 2, D_ℓ coincides with θ_ℓ and the quotient $\mathcal{R}_{a,b}/\Xi_{a,b}$ coincides with the Schur module $L_{(a,b)}\,\mathbb{F}_q^n$ defined in [15, §4]. See also §5.

We next describe how transfers interact with divided powers and show that the relations in the first variable set suffice modulo transfers.

Lemma 3.3. *The operators γ_i and σ are commutative. For $f \in \mathfrak{m}$ and all $i \ge 0$,*

$$\sigma(\gamma_i(f)) = \gamma_i(\sigma(f)).$$

Proof. If $f = x^\alpha y^\beta$ is a monomial and $\alpha \neq \mathbf{0}$, $\gamma_i(\sigma(f)) = \prod_{j=1}^{n}(x_i + ty_i)^{j\alpha_i}\gamma(y^\beta)$ $= \sigma(x^{i\alpha}\gamma_i(y^\beta)) = \sigma(\gamma_i(f))$. Generally, write $f = \sum_{j=1}^{l} u_i$ with u_j are monomials.

$$\begin{aligned}\sigma(\gamma_i(f)) &= \sum_{a_1+\cdots+a_l=i} \sigma\left(\prod_{j=1}^{l}\gamma_{a_j}(u_j)\right) = \sum_{a_1+\cdots+a_l=i}\prod_{j=1}^{l}\sigma(\gamma_{a_j}(u_j)) \\ &= \sum_{a_1+\cdots+a_l=i}\prod_{j=1}^{l}\gamma_{a_j}(\sigma(u_j)) = \gamma_i(\sigma(f)).\end{aligned}$$

Thus we prove this result. □

Proposition 3.4. *For every $f \in R_{2,0}$ and $a, b \geq 0$, write $g = D_1(f)$ and $h = D_2(f)$, we have:*

$$\langle \gamma_k(f) : k > 0\rangle_{a,b} + \Xi_{a,b} = \langle \gamma_k(f), \gamma_k(g), \gamma_k(h) : k > 0\rangle_{a,b} + \Xi_{a,b}.$$

Proof. From Lemma 3.3, for any $k, i \geq 0$,

$$\begin{aligned}D_k(\gamma_i(f)) &= \sigma(\gamma_i(f))[t^k] = \gamma_i(\sigma(t))[t^k] = \gamma_i(f + gt + ht^2)[t^k] \\ &= \left(\sum_{i_1+i_2+i_3=k}\gamma_{i_1}(f)\gamma_{i_2}(gt)\gamma_{i_3}(ht^2)\right)[t^k] \\ &= \left(\sum_{i_1+i_2+i_3=i}\gamma_{i_1}(f)\gamma_{i_2}(f)\gamma_{i_3}(f)t^{i_2+2i_3}\right)[t^k] \\ &= \sum_{\substack{i_1+i_2+i_3=i \\ i_2+2i_3=k}}\gamma_{i_1}(f)\gamma_{i_2}(f)\gamma_{i_3}(f).\end{aligned} \tag{3.3}$$

By Eq. (3.2), for $u \in R_{a-k,b-k}$ and $v \in R_{a,b-2k}$, we have

$$\begin{aligned}uD_k(\gamma_k(f)) - (-1)^k D_k(u)\gamma_k(f) \in \Xi_{a,b}, \\ vD_{2k}(\gamma_k(f)) - D_{2k}(v)\gamma_k(f) \in \Xi_{a,b},\end{aligned}$$

as $D_{k-i}(u)\gamma_k(f) \in R_{a+i,b-i}$ and $D_{2k-j}(v)\gamma_k(f) \in R_{a+j,b-j}$. Using Eq. (3.3),

$$\begin{aligned}u\gamma_k(g) = uD_k(\gamma_k(f)) - u\sum_{i=1}^{\lfloor \frac{k}{2}\rfloor}\gamma_i(f)\gamma_{k-2i}(g)\gamma_i(h) \in \langle \gamma_i(f) : i > 0\rangle_{a,b} + \Xi_{a,b}, \\ v\gamma_k(h) = vD_{2k}(\gamma_k(f)) \in \langle \gamma_i(f) : i > 0\rangle_{a,b} + \Xi_{a,b}.\end{aligned}$$

Thus the two subspaces are equal. □

Set $f_i = p_i(\mathbf{x})$, $g_i = q_i(\mathbf{x}, \mathbf{y})$, and $h_i = p_i(\mathbf{y})$. Then

$$\sigma_t(f_i) = f_i + tg_i + t^2h_i, \qquad D_1(f_i) = g_i, \qquad D_2(f_i) = h_i.$$

By Proposition 3.4, set $\mathcal{K} := \langle \gamma_{p^j}(p_i(\mathbf{x}) : 1 \leq i \leq m, 0 \leq j \leq r\rangle$, for each $a, b \geq 0$,

$$\mathcal{K}_{a,b} + \Xi_{a,b} = (\mathcal{I}_{r,\mathrm{dp}})_{a,b} + \Xi_{a,b}. \tag{3.4}$$

Remark 3.5. The Proposition 3.4 shows that, modulo $\Xi_{a,b}$, the truncated divided power ideal generated by f_i coincides with the ideal generated by the corresponding divided powers of the polynomials f_i, g_i, and h_i. This is the analogue of [15, Lemma 4.5] in the truncated ring.

Proposition 3.6. *Put $h_{a,b} := \dim R_{a,b}$, with $h_{a,b} = 0$ if either index is negative. Then*

$$\dim(R_{a,b}/\Xi_{a,b}) = \begin{cases} h_{a,b} - h_{a+1,b-1}, & a \geq b, \\ 0, & a < b. \end{cases}$$

In particular, $R_{a,b} = \Xi_{a,b}$ for $0 \leq a < b$, and

$$\sum_{a \geq b \geq 0} \dim(R_{a,b}/\Xi_{a,b})\, t^a s^b = \left[(1 + t + \cdots + t^{p-1})^n (1 + s + \cdots + s^{p-1})^n (1 - t^{-1}s)\right]_{a \geq b \geq 0},$$

where the brackets retain only monomials $t^a s^b$ with $a \geq b \geq 0$.

The proof is given in Appendix B.2.

Lemma 3.7. *At the oil top degree $d_o = o(p-1)$, the transfer relations satisfy*

$$\Xi_{d_o,b} \subseteq (\mathcal{J})_{d_o,b} \qquad (b \geq 0).$$

Proof. For $b = 0$, the transfer space is zero. Suppose $b \geq 1$ and fix $1 \leq \ell \leq b$. Write $\rho_{a,b} : R_{a,b} \longrightarrow S_{a,b}$ for the bihomogeneous components of the restriction map ρ_R. In the oil ring S, define

$$D_\ell^{\mathcal{O}}(G) := G(\mathbf{s} + \tau\mathbf{t}, \mathbf{t})[\tau^\ell],$$

where τ is an indeterminate. This is well defined because $(s_j + \tau t_j)^p = s_j^p + \tau^p t_j^p = 0$ in $S[\tau]$. The identity $U(\mathbf{s} + \tau\mathbf{t}) = U\mathbf{s} + \tau U\mathbf{t}$ shows that restriction commutes with the substitution defining transfer. Taking the coefficient of τ^ℓ gives the commutative diagram

$$\begin{array}{ccc} R_{a+\ell,b-\ell} & \xrightarrow{D_\ell} & R_{a,b} \\ \downarrow{\scriptstyle \rho_{a+\ell,b-\ell}} & & \downarrow{\scriptstyle \rho_{a,b}} \\ S_{a+\ell,b-\ell} = 0 & \xrightarrow[D_\ell^{\mathcal{O}}]{} & S_{a,b}. \end{array}$$

Indeed, every surviving monomial in S has s-degree at most $o(p-1) = d_o$, so the lower-left space is zero when $a = d_o$. Hence

$$\rho_{d_o,b} \circ D_\ell = D_\ell^{\mathcal{O}} \circ \rho_{d_o+\ell,b-\ell} = 0,$$

and therefore $\operatorname{im} D_\ell \subseteq \ker \rho_{d_o,b} = (\mathcal{J})_{d_o,b}$ by Proposition 3.1. Summing over $1 \leq \ell \leq b$ proves the inclusion. □

3.3 Rank analysis

The preceding results give

$$\mathcal{I}_{a,b} \subseteq \mathcal{I}_{a,b} + \Xi_{a,b} \subseteq (\mathcal{I}_{r,\mathrm{dp}})_{a,b} + \Xi_{a,b} = \mathcal{K}_{a,b} + \Xi_{a,b} \subseteq \mathcal{J}_{a,b} + \Xi_{a,b}. \tag{3.5}$$

By Lemma 3.7, the last term is $\mathcal{J}_{a,b}$ when $a = d_o$. We seek degrees a, b for which the last inclusion is an equality. At such degrees, the public relations determine information about $\mathcal{J}_{a,b}$ through $(\mathcal{I}_{\mathrm{dp}})_{a,b}$ and the transfer relations; see [6, Theorem 4.1].

Lemma 3.8. *Let $\mathcal{N}$ be a bihomogeneous ideal in R, and let $a, b \geq 0$. If*

$$\mathcal{N}_{a,b} + \Xi_{a,b} = R_{a,b},$$

then

$$\mathcal{N}_{a,b'} + \Xi_{a,b'} = R_{a,b'} \qquad \text{for every } b' \geq b.$$

Additionally, if $b = 0$, then $\mathcal{N}_{a',0} = R_{a',0}$ for every $a' \geq a$.

Proof. It suffices to prove the assertion for $b' = b+1$ and then iterate. Since the transfer substitution fixes each y_j, we have

$$D_h(y_j F) = y_j D_h(F) \qquad (F \in R,\ h \geq 0).$$

For $1 \leq h \leq b$ and $F \in R_{a+h,b-h}$, the product $y_j F$ belongs to $R_{a+h,b+1-h}$. Hence

$$y_j \Xi_{a,b} \subseteq \Xi_{a,b+1}.$$

Also, $y_j \mathcal{N}_{a,b} \subseteq \mathcal{N}_{a,b+1}$ because $\mathcal{N}$ is an ideal. Every truncated monomial of bidegree $(a, b+1)$ is y_j times a truncated monomial of bidegree (a, b) for some j, so

$$\begin{aligned} R_{a,b+1} &= \sum_{j=1}^{n} y_j R_{a,b} \\ &= \sum_{j=1}^{n} y_j (\mathcal{N}_{a,b} + \Xi_{a,b}) \\ &\subseteq \mathcal{N}_{a,b+1} + \Xi_{a,b+1} \subseteq R_{a,b+1}. \end{aligned}$$

Thus equality holds. The last statement follows from the fact that $\Xi_{c,0} = 0$ for any $c \geq 0$. □

The above lemma implies that following definition. In fact, Lemma 3.8 is an analogue of [15, Lemma 4.6].

Definition 3.9. *For a formal Laurent series $f = \sum_{(a,b)\in\mathbb{Z}^2} f_{ab} t^a s^b$ in t, s with integer coefficients, we write $[f]_+$ for the series supported on $\mathbb{Z}^2_{\geq 0}$ with coefficients*

$$[f]_+[t^a s^b] = \begin{cases} f_{ab}, & \text{if } f_{a'0} > 0 \text{ for all } 0 \leq a' \leq a \text{ and } f_{ab'} > 0 \text{ for all } 0 \leq b' \leq b, \\ 0, & \text{otherwise,} \end{cases}$$

.

Lemma 3.10. *For any $a, b \geq 0$, the map ρ_R in Proposition 3.1 induces an isomorphism*

$$R_{a,b}/(\mathcal{J}_{a,b} + \Xi_{a,b}) \cong S_{a,b}/\Xi^{\mathcal{O}}_{a,b},$$

where $\Xi^{\mathcal{O}}_{a,b} := \sum_{\ell=1}^{b} D^{\mathcal{O}}_{\ell}(S_{a+\ell,b-\ell})$ is the transfer relation space in S.

Proof. By Proposition 3.1, each restriction $\rho_{a,b} : R_{a,b} \to S_{a,b}$ is surjective with kernel $\mathcal{J}_{a,b}$. Restriction commutes with transfer:

$$\rho_{a,b} \circ D_{\ell} = D^{\mathcal{O}}_{\ell} \circ \rho_{a+\ell,b-\ell}.$$

Indeed, this follows by taking the coefficient of τ^{ℓ} in the identity $U(\mathbf{s} + \tau\mathbf{t}) = U\mathbf{s} + \tau U\mathbf{t}$. Since $\rho_{a+\ell,b-\ell}$ is surjective, we obtain

$$\rho_{a,b}(D_{\ell}(R_{a+\ell,b-\ell})) = D^{\mathcal{O}}_{\ell}(S_{a+\ell,b-\ell}).$$

Summing over $1 \leq \ell \leq b$ gives $\rho_{a,b}(\Xi_{a,b}) = \Xi^{\mathcal{O}}_{a,b}$, also when $b = 0$, when both transfer spaces are zero. Thus the composite map $R_{a,b} \to S_{a,b}/\Xi^{\mathcal{O}}_{a,b}$ is surjective and has kernel $\mathcal{J}_{a,b} + \Xi_{a,b}$: if $\rho_{a,b}(F) \in \Xi^{\mathcal{O}}_{a,b}$, choose $H \in \Xi_{a,b}$ with $\rho_{a,b}(H) = \rho_{a,b}(F)$, so that $F - H \in \mathcal{J}_{a,b}$. This implies the claimed isomorphism. □

Proposition 3.6 and Lemma 3.10 give

$$\begin{aligned}
\mathcal{H}_0(t,s) :&= \sum_{a \geq b \geq 0} \dim(R_{a,b}/(\mathcal{J}_{a,b} + \Xi_{a,b}))\, t^a s^b \\
&= \sum_{a \geq b \geq 0} \dim(S_{a,b}/\Xi^{\mathcal{O}}_{a,b})\, t^a s^b \\
&= \left[(1 + t + \cdots + t^{p-1})^o (1 + s + \cdots + s^{p-1})^o (1 - t^{-1}s)\right]_{a \geq b \geq 0},
\end{aligned} \tag{3.6}$$

Here $\Xi^{\mathcal{O}}_{a,b}$ in the second line denotes the transfer relation space in $S_{a,b}$.

Conjecture 3.11. Let $p_1, \ldots, p_m \in k[x_1, \ldots, x_n]_2$ be independent, uniformly random homogeneous quadratics, with no prescribed oil space. Retain the notation R, $\mathcal{I}_{r,\mathrm{dp}}$, and Ξ from above. With high probability,

$$\begin{aligned}
\mathcal{H}(t,s) &:= \sum_{a \geq b \geq 0} \dim\left(\frac{R_{a,b}}{(\mathcal{I}_{r,\mathrm{dp}})_{a,b} + \Xi_{a,b}}\right) t^a s^b \\
&= \sum_{a \geq b \geq 0} \dim\left(\frac{R_{a,b}}{(\mathcal{K})_{a,b} + \Xi_{a,b}}\right) t^a s^b \\
&= \left[\frac{F_{n,p}(t) F_{n,p}(s)(1 - t^{-1}s)}{F_{m,p^{r+1}}(t^2) F_{m,p^{r+1}}(ts) F_{m,p^{r+1}}(s^2)} \cdot \frac{t - s}{t}\right]_{+},
\end{aligned} \tag{3.7}$$

here $F_{n,p}(t) := (\sum_{i=0}^{p-1} t^i)^n$. The second equality follows from Eq. (3.4); only the final formula is conjectural.

Remark 3.12. Definition 3.9 follows [15, Definition 4.7], and Conjecture 3.11 is inspired by [15, Conjecture 4.8].

We conduct experiments on UOV systems to verify the Conjecture 3.11. For different parameters (n, v, m, q), we verify the prediction for $r = 1$ and all degrees (d_1, d_2) such that $0 \le d_2 \le d_1 \le D_0 \le o(p-1)$. The test results are summarized in the table below.

Table 1: The parameters were tested for different characteristics

Paramaters (n, v, q, D_0)	Range of m	Number of annomalies
(6,4,3,4)	2,3,6,8,9,10	0
(6,4,3,4)	4,5,7	1
(7,5,7,3)	$2 \le m \le 8$	0
(8,5,5,3)	$3 \le m \le 8$	0

Remark 3.13. As explained in § 5, in characteristic 2, our method is equivalent to the exterior-algebra approach of Merz and Ran [15], and the corresponding conjectures have the same expressions. They have already provided experimental evidence supporting their conjecture in this case. We therefore focus our experiments on odd characteristic.

For the public quadratics with hidden oil space $\mathcal{O}$, consider $a = d_o$ and $0 \le b \le d_o$. By Lemma 3.7, $\Xi_{d_o,b} \subseteq \mathcal{J}_{d_o,b}$, so

$$(\mathcal{I}_{r,\mathrm{dp}})_{d_o,b} + \Xi_{d_o,b} = \mathcal{K}_{d_o,b} + \Xi_{d_o,b} \subseteq \mathcal{J}_{d_o,b}.$$

Proposition 3.1 and $T(o, p, d_o) = 1$ give

$$\mathcal{H}_0[t^{d_o} s^b] = \dim\left(\frac{R_{d_o,b}}{\mathcal{J}_{d_o,b}}\right) = T(o, p, b).$$

Consequently,

$$\dim\left(\frac{R_{d_o,b}}{(\mathcal{I}_{\mathrm{dp}})_{d_o,b} + \Xi_{d_o,b}}\right) \ge T(o, p, b),$$

with equality if and only if $\mathcal{K}_{d_o,b} + \Xi_{d_o,b} = \mathcal{J}_{d_o,b}$. Thus, verifying that the quotient dimension is $T(o, p, b)$ allows us to obtain this component of the hidden ideal from the public relations. Conjecture 3.11 provides a random-system benchmark for selecting b, while $\mathcal{H}_0[t^{d_o} s^b]$ gives the dimension that must be attained in the structured system.

4 A multigraded truncated-polynomial construction

Retaining the notation of § 3, we extend the construction from two oil elements to $\ell \ge 2$ elements, using one set of variables for each element and polar equations

for every pair. This corresponds to the Schur module attack of Merz and Ran [15, §5]. The basis matrix U is used only for analysis; the attack requires only the public quadratics and the parameters n, o, p, ℓ. We omit details that carry over directly from the bigraded case.

4.1 Truncation and the oil quotient

Put $\mathbf{x}_t = (x_{t1}, \dots, x_{tn})$ for $1 \le t \le \ell$ and define the multigraded truncated polynomial ring

$$\mathcal{R} := k[\mathbf{x}_1, \dots, \mathbf{x}_\ell]/\langle x_{tj}^p : 1 \le t \le \ell,\ 1 \le j \le n\rangle, \mathfrak{m} := \langle x_{tj} : 1 \le t \le \ell,\ 1 \le j \le n\rangle.$$

Give x_{tj} degree $\mathbf{e}_t$, the t-th standard basis vector of $\mathbb{Z}^\ell$. For $\lambda \in \mathbb{N}^\ell$,

$$\mathcal{R}_\lambda = \operatorname{span}_k \left\{ \prod_{t=1}^{\ell} \mathbf{x}_t^{\alpha_t} : |\alpha_t| = \lambda_t,\ 0 \le \alpha_{tj} < p \right\}.$$

Components with a negative degree are zero. For $1 \le i \le m$ and $1 \le j < k \le \ell$, write

$$f_i^{(t)} := p_i(\mathbf{x}_t), \qquad q_i^{(jk)} := p_i(\mathbf{x}_j + \mathbf{x}_k) - p_i(\mathbf{x}_j) - p_i(\mathbf{x}_k).$$

Let

$$\mathcal{I} := \langle f_i^{(t)}, q_i^{(jk)} : 1 \le i \le m,\ 1 \le t \le \ell,\ 1 \le j < k \le \ell\rangle_{\mathcal{R}}$$
$$\mathcal{J} := \langle L_j(\mathbf{x}_t) : 1 \le j \le v,\ 1 \le t \le \ell\rangle_{\mathcal{R}}.$$

Thus $\mathcal{I} \subseteq \mathcal{J}$. Introduce oil variables $\mathbf{s}_t = (s_{t1}, \dots, s_{to})$ and the ring

$$\mathcal{S} := k[\mathbf{s}_1, \dots, \mathbf{s}_\ell]/\langle s_{tj}^p : 1 \le t \le \ell,\ 1 \le j \le o\rangle.$$

The multigraded analogue of Proposition 3.1 is as follows.

Proposition 4.1. *The substitution map*

$$\rho : \mathcal{R} \longrightarrow \mathcal{S}, \qquad F \longmapsto F(U\mathbf{s}_1, \dots, U\mathbf{s}_\ell)$$

is surjective, preserves multidegrees, and has kernel $\mathcal{J}$. In particular,

$$\mathcal{R}/\mathcal{J} \cong \mathcal{S}, \qquad \dim \mathcal{S}_\lambda = \prod_{t=1}^{\ell} T(o, p, \lambda_t). \tag{4.1}$$

The top degree in each set of oil variables is $d_o = o(p-1)$.

4.2 Divided powers and transfer relations in positive characteristic

Similarly, one can define a divided power γ_h on $\mathfrak{m}\mathcal{R}$, now with all ℓn variables, on the maximal ideal $\mathfrak{m}_{\mathcal{R}} = \langle \mathbf{x}_1, \dots, \mathbf{x}_\ell\rangle$. Set

$$\mathcal{I}_{r,\mathrm{dp}} := \left\langle \begin{matrix} \gamma_{p^a}(f_i^{(t)}) :\ 1 \le i \le m,\ 1 \le t \le \ell,\ 0 \le a \le r, \\ \gamma_{p^a}(q_i^{(jk)}) : 1 \le i \le m,\ 1 \le j < k \le \ell,\ 0 \le a \le r \end{matrix} \right\rangle_{\mathcal{R}}. \tag{4.2}$$

Their multidegrees are $2p^a\mathbf{e}_t$ and $p^a(\mathbf{e}_j+\mathbf{e}_k)$, respectively. Remark 2.5 shows that all positive divided powers of the public quadratics generate this same ideal. By Remark 2.6, we have

$$\mathcal{I} \subseteq \mathcal{I}_{r,\mathrm{dp}} \subseteq \mathcal{J}. \tag{4.3}$$

Fix $1 \le i < j \le \ell$ and let τ be an indeterminate. Define a $k[\tau]$-algebra endomorphism

$$\sigma_\tau^{ij} : \mathcal{R}[\tau] \longrightarrow \mathcal{R}[\tau]$$

by its values on the variables:

$$\sigma_\tau^{ij}(x_{tk}) = \begin{cases} x_{ik} + \tau x_{jk}, & t = i, \\ x_{tk}, & t \neq i, \end{cases} \qquad 1 \le t \le \ell, \quad 1 \le k \le n.$$

Thus every coordinate in the i-th variable set is changed using the same parameter τ, while the j-th variable set and all other sets remain fixed. The parameter itself is also fixed: $\sigma_\tau^{ij}(\tau) = \tau$. For a polynomial class $F \in \mathcal{R}$, this means

$$\sigma_\tau^{ij}(F) = F(\mathbf{x}_1, \dots, \mathbf{x}_i + \tau\mathbf{x}_j, \dots, \mathbf{x}_\ell).$$

The substitution respects the truncation relations because $(x_{ik} + \tau x_{jk})^p = x_{ik}^p + \tau^p x_{jk}^p = 0$ in $\mathcal{R}[\tau]$, and the remaining relations are fixed. Put $D_h^{ij}(F) := \sigma_\tau^{ij}(F)[\tau^h]$. Then

$$D_h^{ij} : \mathcal{R}_{\lambda + h\mathbf{e}_i - h\mathbf{e}_j} \longrightarrow \mathcal{R}_\lambda. \tag{4.4}$$

Explicitly, on a truncated monomial,

$$D_h^{ij}\left(\prod_{t=1}^{\ell} \mathbf{x}_t^{\alpha_t}\right) = \sum_{\substack{\kappa \le \alpha_i \\ |\kappa| = h}} \binom{\alpha_i}{\kappa} \mathbf{x}_i^{\alpha_i - \kappa} \mathbf{x}_j^{\alpha_j + \kappa} \prod_{t \neq i,j} \mathbf{x}_t^{\alpha_t}. \tag{4.5}$$

Terms with any exponent at least p are zero. For each fixed pair (i, j), the operators form an iterative Hasse–Schmidt derivation, in particular,

$$D_h^{ij}(FG) = \sum_{a=0}^{h} D_a^{ij}(F) D_{h-a}^{ij}(G), \qquad D_a^{ij} D_b^{ij} = \binom{a+b}{a} D_{a+b}^{ij}.$$

Define

$$\Xi_\lambda := \sum_{1 \le i < j \le \ell} \sum_{h=1}^{\lambda_j} \operatorname{im} D_h^{ij} = \sum_{1 \le i < j \le \ell} \sum_{\substack{a \ge 0 \\ p^a \le \lambda_j}} \operatorname{im} D_{p^a}^{ij}, \tag{4.6}$$

where each image has target $\mathcal{R}_\lambda$. The second equality follows by expanding h in base p and applying the iteration identity for the fixed pair.

Remark 4.2. When $k = \mathbb{F}_q$ has characteristic 2, the quotient $\mathcal{R}_\lambda / \Xi_\lambda$ coincides with the Schur module $L_\lambda \mathbb{F}_q^n$ defined in [15, §5].

We have a similar result as [15, Lemma 5.1], which will be used later in the cost analysis.

Lemma 4.3. *For every $\lambda \in \mathbb{N}^\ell$, we have*

$$\Xi_\lambda = \sum_{i=1}^{\ell-1} \sum_{\substack{a\geq 0 \\ p^a \leq \lambda_{i+1}}} \operatorname{im} D_{p^a}^{i,i+1}$$

Proof. Write A_λ for the sum on the right-hand side. For each adjacent pair, the base-p iteration identity used in Eq. (4.6) gives

$$A_\lambda = \sum_{i=1}^{\ell-1} \sum_{h=1}^{\lambda_{i+1}} D_h^{i,i+1}\big(\mathcal{R}_{\lambda+h\mathbf{e}_i-h\mathbf{e}_{i+1}}\big).$$

Clearly $A_\lambda \subseteq \Xi_\lambda$. We prove the reverse inclusion by induction on $j-i$, simultaneously for all target multidegrees and all $h \geq 1$.

For $i < k < j$, the substitution operators satisfy

$$\sigma_{-u}^{ik}\sigma_v^{kj}\sigma_u^{ik} = \sigma_v^{kj}\sigma_{uv}^{ij}$$

on $\mathcal{R}[u,v]$, where the rightmost operator acts first. Indeed, both sides send $\mathbf{x}_i$ to $\mathbf{x}_i + uv\mathbf{x}_j$, send $\mathbf{x}_k$ to $\mathbf{x}_k + v\mathbf{x}_j$, and fix the other variable sets. Taking the coefficient of $u^h v^h$ gives

$$D_h^{ij} = \sum_{a=0}^{h} (-1)^a D_a^{ik} D_h^{kj} D_{h-a}^{ik}. \tag{4.7}$$

This identity involves no division and is valid in characteristic p.

The case $j-i=1$ follows from the displayed description of A_λ. For $j-i>1$, take $k=i+1$ and $F \in \mathcal{R}_{\lambda+h\mathbf{e}_i-h\mathbf{e}_j}$. In Eq. (4.7), every term with $a \geq 1$ belongs to A_λ: its outermost operator is the adjacent transfer D_a^{ik}, and its input has degree $\lambda + a\mathbf{e}_i - a\mathbf{e}_k$. If $a > \lambda_k$, this input component is zero. The term with $a=0$ is $D_h^{kj}D_h^{ik}(F)$ and belongs to A_λ by the induction hypothesis for the shorter pair (k,j). Thus $D_h^{ij}(F) \in A_\lambda$, proving $\Xi_\lambda \subseteq A_\lambda$. □

Lemma 4.4. *For $\lambda = (d_o, \ldots, d_o, b)$ with $b \geq 0$,*

$$\Xi_\lambda \subseteq \mathcal{J}_\lambda. \tag{4.8}$$

Proof. This follows immediately from the commutative diagram

$$\begin{array}{ccc} \mathcal{R}_{\lambda+h\mathbf{e}_i-h\mathbf{e}_j} & \xrightarrow{D_h^{ij}} & \mathcal{R}_\lambda \\ \downarrow{\scriptstyle \rho_{\lambda+h\mathbf{e}_i-h\mathbf{e}_j}} & & \downarrow{\scriptstyle \rho_\lambda} \\ \mathcal{S}_{\lambda+h\mathbf{e}_i-h\mathbf{e}_j} = 0 & \xrightarrow[D_h^{ij,\mathcal{O}}]{} & \mathcal{S}_\lambda, \end{array}$$

which shows that each D_h^{ij} maps into the kernel of ρ_λ, which is precisely $\mathcal{J}_\lambda$. □

Remark 4.5. For more than two variable sets, $\lambda_1 = d_o$ alone is insufficient. For example, take $p = 3$, $o = 1$, $\ell = 3$, and $\lambda = (2, 1, 1)$. In the oil ring,

$$D_1^{23,\mathcal{O}}(s_{11}^2 s_{21}^2) = 2s_{11}^2 s_{21} s_{31} \neq 0.$$

Thus transfers between later variable sets can survive restriction to $\mathcal{O}$. The lemma uses a top degree in every variable set from which a forward transfer can originate.

As in the bigraded setting, we can reduce the study of $\mathcal{I}_{\mathrm{dp}}$ to the first variable set modulo transfers. The following lemma makes this reduction precise.

For any equation index i and pair $j < k$,

$$\sigma_\tau^{jk}(f_i^{(j)}) = f_i^{(j)} + \tau q_i^{(jk)} + \tau^2 f_i^{(k)}.$$

Put

$$\mathcal{K} := \langle \gamma_{p^a}(f_i^{(1)}) : 1 \leq i \leq m, 0 \leq a \leq r \rangle_{\mathcal{R}}.$$

Lemma 4.6. *For every $\lambda \in \mathbb{N}^\ell$,*

$$\mathcal{K}_\lambda + \Xi_\lambda = \mathcal{I}_{r,\mathrm{dp},\lambda} + \Xi_\lambda. \tag{4.9}$$

Proof. The proof is analogous to that of Proposition 3.4.

4.3 Rank analysis

As in § 3.3, we seek multidegrees λ for which the last inclusion

$$\mathcal{I}_\lambda \subseteq \mathcal{I}_\lambda + \Xi_\lambda \subseteq \mathcal{I}_{r,\mathrm{dp},\lambda} + \Xi_\lambda = \mathcal{K}_\lambda + \Xi_\lambda \subseteq \mathcal{J}_\lambda + \Xi_\lambda. \tag{4.10}$$

is an equality. At such multidegrees λ, the public relations and transfers determine information about $\mathcal{J}_\lambda$ from $\mathcal{I}_{\mathrm{dp},\lambda}$; see [6, Theorem 4.5]. We compare the dimensions of the corresponding quotients using multigraded Hilbert series. By Lemma 4.4, the last term is $\mathcal{J}_\lambda$ when $\lambda = (d_o, \ldots, d_o, b)$ with $0 \leq b \leq d_o$.

Write $\mathbf{t} = (t_1, \ldots, t_\ell)$, $\mathbf{t}^\lambda = \prod_{i=1}^{\ell} t_i^{\lambda_i}$, and recall that $F_{n,p}(t) = (\sum_{i=0}^{p-1} t^i)^n$. Let

$$\Lambda_\ell := \{\lambda \in \mathbb{N}^\ell : n(p-1) \geq \lambda_1 \geq \cdots \geq \lambda_\ell\}.$$

The notation $[F]_{\Lambda_\ell}$ means that only monomials with multidegrees in Λ_ℓ are retained. We also extend Definition 3.9 to ℓ variables as follows.

Definition 4.7. *For a formal Laurent series $F = \sum_{\mu \in \mathbb{Z}^\ell} c_\mu \mathbf{t}^\mu$ with integer coefficients, define $[F]_+$ to be supported on $\mathbb{N}^\ell$, with*

$$[F]_+[\mathbf{t}^\lambda] = \begin{cases} c_\lambda, & \text{if } c_{(\lambda_1,\ldots,\lambda_{i-1},j,0,\ldots,0)} > 0 \text{ for all } 1 \leq i \leq \ell \text{ and } 0 \leq j \leq \lambda_i, \\ 0, & \text{otherwise.} \end{cases}$$

For $\ell = 2$, this agrees with Definition 3.9.

The oil quotient has the following Hilbert series after imposing the transfer relations:

$$\begin{aligned}\mathcal{H}_0(\mathbf{t}) &:= \sum_{\lambda\in\Lambda_\ell} \dim\left(\frac{\mathcal{R}_\lambda}{\mathcal{J}_\lambda+\Xi_\lambda}\right)\mathbf{t}^\lambda \\ &= \sum_{\lambda\in\Lambda_\ell} \dim\left(\frac{\mathcal{S}_\lambda}{\Xi_\lambda^{\mathcal{O}}}\right)\mathbf{t}^\lambda \\ &= \left[\prod_{i=1}^{\ell} F_{o,p}(t_i)\prod_{1\le i<j\le\ell}(1-t_i^{-1}t_j)\right]_{\Lambda_\ell}.\end{aligned} \tag{4.11}$$

Here $\Xi_\lambda^{\mathcal{O}}$ is the transfer relation space in $\mathcal{S}_\lambda$, defined by the same substitutions as Ξ_λ. This is the multigraded counterpart of the formula for $\mathcal{H}_0(t,s)$ in § 3.3; we state the conclusion without repeating the argument.

For random quadratics without a prescribed oil space, the corresponding Hilbert series is predicted by the following generalization of Conjecture 3.11.

Conjecture 4.8. Let $p_1,\dots,p_m \in k[x_1,\dots,x_n]_2$ be independent, uniformly random homogeneous quadratics, with no prescribed oil space. Retain the notation $\mathcal{R}$, $\mathcal{I}_{\mathrm{dp}}$, $\mathcal{K}$, and Ξ from above. With high probability,

$$\begin{aligned}\mathcal{H}(\mathbf{t}) &:= \sum_{\lambda\in\Lambda_\ell} \dim\left(\frac{\mathcal{R}_\lambda}{\mathcal{I}_{\mathrm{dp},\lambda}+\Xi_\lambda}\right)\mathbf{t}^\lambda \\ &= \sum_{\lambda\in\Lambda_\ell} \dim\left(\frac{\mathcal{R}_\lambda}{\mathcal{K}_\lambda+\Xi_\lambda}\right)\mathbf{t}^\lambda \\ &= \left[\left[\prod_{i=1}^{\ell} F_{n,p}(t_i)\cdot\prod_{1\le i<j\le\ell}(1-t_i^{-1}t_j)\cdot\prod_{1\le i\le j\le\ell} F_{m,p}(t_it_j)\right]_+\right]_{\Lambda_\ell}.\end{aligned} \tag{4.12}$$

The second equality is already proved by Lemma 4.6; only the final formula is conjectural. For $\ell=2$, it specializes to Conjecture 3.11 on $a\ge b\ge 0$.

For the public quadratics with hidden oil space $\mathcal{O}$, put

$$h_\lambda := \dim\left(\frac{\mathcal{R}_\lambda}{\mathcal{I}_{\mathrm{dp},\lambda}+\Xi_\lambda}\right).$$

The inclusion (4.10) gives $h_\lambda \ge \mathcal{H}_0[\mathbf{t}^\lambda]$ for $\lambda\in\Lambda_\ell$. Equality of these dimensions holds if and only if

$$\mathcal{I}_{r,\mathrm{dp},\lambda}+\Xi_\lambda = \mathcal{J}_\lambda+\Xi_\lambda.$$

The random-system conjecture provides a benchmark for selecting multidegrees, while the oil quotient gives the dimension that must be attained in the structured system. In particular, for $\lambda=(d_o,\dots,d_o,b)$ with $0\le b\le d_o$, Lemma 4.4 and Proposition 4.1 give

$$\mathcal{H}_0[\mathbf{t}^\lambda] = T(o,p,b).$$

Thus, at such a multidegree, verifying $h_\lambda = T(o,p,b)$ establishes $\mathcal{K}_\lambda + \Xi_\lambda = \mathcal{J}_\lambda$, allowing us to obtain this component of the hidden ideal from the public relations.

5 Relation to exterior algebra

In this section, we compare the method developed in this paper with the approach of Merz and Ran [15]. The comparison depends on the characteristic of the underlying field: in characteristic two, our construction gives a polynomial-ring description of their method, whereas in odd characteristic the two constructions use different algebras.

Let $k = \mathbb{F}_q$ be a finite field and retain $V = k^n$ and $E = V^*$. Write $e_1, \dots, e_n$ for the coordinate basis of E, and set

$$W = \bigwedge(E \oplus E), \quad \bigwedge^{a,b} E = \bigwedge^a E \otimes \bigwedge^b E.$$

Merz and Ran [15] work with exterior algebras on V, whereas we use $E = V^*$ to regard the public relations as forms on V. This difference does not affect the construction: after choosing a basis of V and its dual basis, the resulting identification $V \cong V^*$ induces an isomorphism of the corresponding exterior algebras.

5.1 Characteristic two

Assume $\operatorname{char} k = 2$. In this case, exterior multiplication is commutative, and each e_i has square zero. Hence there is a graded algebra isomorphism

$$\Phi : R \xrightarrow{\sim} W, \qquad x_i \mapsto e_i \otimes 1, \quad y_i \mapsto 1 \otimes e_i.$$

To describe the public relations under this identification, write

$$p_i(\mathbf{x}) = \sum_j a_{i,j} x_j^2 + \sum_{j<k} c_{i,jk} x_j x_k.$$

The square terms vanish in R, while the polar form is $q_i(\mathbf{x}, \mathbf{y}) = \sum_{j<k} c_{i,jk}(x_j y_k + x_k y_j)$. Define

$$\omega_i = \sum_{j<k} c_{i,jk} e_j \wedge e_k \in \bigwedge^2 E, \qquad \beta_i = \sum_{j<k} c_{i,jk}(e_j \otimes e_k + e_k \otimes e_j) \in E \otimes E.$$

Here ω_i represents the alternating polar form as an exterior two-form, whereas β_i represents the same form as a bilinear relation between the two copies of V. Thus, for $i = 1, \dots, m$, we have

$$\Phi(p_i(\mathbf{x})) = \omega_i \otimes 1, \quad \Phi(p_i(\mathbf{y})) = 1 \otimes \omega_i, \quad \Phi(q_i(\mathbf{x}, \mathbf{y})) = \beta_i.$$

The corresponding multiplication maps are also identified by Φ.

With respect to these coordinates, the divided power operations on $\mathfrak{m}$ in R correspond to the divided power operations in the exterior construction. In other words,

$$\Phi(\gamma_{2^r}(f)) = (\Phi(f))^{(2^r)} \qquad (f \in R_2).$$

Similarly, the transfer maps D_ℓ in our setting correspond to θ_ℓ in the exterior construction:

$$\Phi(D_\ell(f)) = \theta_\ell(\Phi(f)).$$

For the precise definitions of $\omega^{(2^r)}$ and θ_ℓ, see [15, §4]. Consequently, the public relations, their divided powers, and the transfer relations agree under Φ. In particular, Φ induces an isomorphism

$$\overline{\Phi}\colon \mathcal{R}_{ab}/\Xi_{ab} \xrightarrow{\sim} L_{(a,b)}\,\mathbb{F}_q^n = \bigwedge^{a,b} \mathbb{F}_q^n \,/\Theta^{a,b},$$

where the right hand side is the Schur module as defined in [15, §4]. Moreover, the same identification extends naturally to several sets of variables, relating our construction in § 4 to the construction of Merz and Ran [15, §5].

Under this identification, the Schur–Macaulay map

$$M_{a,b} : \bigoplus_{\substack{1\le i\le m\\ 1\le j\le \log_2 n}} \bigwedge^{a-2^j,b} E \;\oplus\; \bigoplus_{0\le j\le \log_2 b} \bigwedge^{a+2^j,b-2^j} E \longrightarrow \bigwedge^{a,b} E$$

corresponds to the transpose of the Macaulay matrix whose row space is $(\mathcal{I}_{\mathrm{dp}})_{a,b}+\Xi_{a,b}$. Here $I_{\mathrm{dp}} := \cup_{r\ge 0} I_{r,\mathrm{dp}}$ Thus,

$$\dim\ker M_{a,b}^T = \dim\big(R_{a,b}/((\mathcal{I}_{\mathrm{dp}})_{a,b} + \Xi_{a,b})\big). \tag{5.1}$$

In general, we have the inclusion

$$(\mathcal{I}_{\mathrm{dp}})_{a,b} + \Xi_{a,b} \subseteq \mathcal{J}_{a,b} + \Xi_{a,b}.$$

Let $\mathcal{H}_0(s,t)$ be the Hilbert series of $R/(\mathcal{J}+\Xi)$ and $\mathcal{H}$ be the Hilbert series defined in Conjecture 3.11. For structured UOV systems, we use the following heuristic rank prediction:

$$\dim\ker M_{a,b}^T = \max\{\mathcal{H}[t^a s^b], \mathcal{H}_0[t^a s^b]\}. \tag{5.2}$$

When $a = o$, Lemma 3.7 gives $\Xi_{o,b} \subseteq \mathcal{J}_{o,b}$, and hence

$$\mathcal{I}_{o,b} + \Xi_{o,b} \subseteq \mathcal{J}_{o,b}.$$

This gives a lower bound on the kernel dimension. Since the hidden quotient satisfies

$$\dim \bigwedge^{o,b} \mathcal{O} = \dim\big(R_{o,b}/\mathcal{J}_{o,b}\big),$$

we have $\mathcal{H}_0[t^o s^b] = \dim \bigwedge^{o,b} \mathcal{O}$, so the predicted kernel dimension becomes

$$\dim\ker M_{o,b}^T = \max\{\mathcal{H}[t^o s^b], \dim \bigwedge^{o,b} \mathcal{O}\}. \tag{5.3}$$

Remark 5.1. Merz and Ran [15, §4.4] propose a prediction for $\dim\ker M_{a,b}^T$ of the form $\max\{\mathcal{H}[t^a s^b], \dim \bigwedge^{a,b} \mathcal{O}\}$. The preceding argument justifies the hidden-space term $\dim \bigwedge^{a,b} \mathcal{O}$ when $a = o$: by Lemma 3.7, $\Xi_{o,b} \subseteq \mathcal{J}_{o,b}$, so the transfer relations do not further reduce the hidden quotient. For general a, the corresponding term is instead $\mathcal{H}_0[t^a s^b]$, as in Eq. (5.2). This distinction does not affect their conclusions, since their analysis assumes $a = o$.

5.2 Odd characteristic

Assume $\operatorname{char} k = p > 2$. The preceding identification no longer holds: exterior generators have square zero and satisfy $e_i e_j = -e_j e_i$, whereas the p-truncated polynomial algebra is commutative and allows nonzero squares.

In this case, the approach of Merz and Ran [15, §5.3] is based on the exterior-algebra method of Ran [17, §7]. Since the polar form q_i is symmetric and generally not alternating on V, it does not give an exterior two-form ω_i as in characteristic two. Nevertheless, it still defines a tensor β_i between the two exterior factors.

Using the coefficients of p_i introduced above, the symmetric matrix B_i representing q_i has entries

$$(B_i)_{jj} = 2a_{i,j}, \qquad (B_i)_{jk} = (B_i)_{kj} = c_{i,jk} \quad (j < k).$$

Over any field, the tensor

$$\beta_i := \sum_{j,k} (B_i)_{jk} e_j \otimes e_k \in \bigwedge^{1,1} E = E \otimes E \tag{5.4}$$

corresponds in W to the alternating form on $V \oplus V$ given by

$$\widetilde{\beta}_i\big((u_1, u_2), (w_1, w_2)\big) = q_i(u_1, w_2) - q_i(w_1, u_2).$$

Since q_i vanishes on $\mathcal{O} \times \mathcal{O}$, this gives

$$\langle \beta_1, \dots, \beta_m \rangle_W \subseteq \langle L_i \otimes 1, 1 \otimes L_i : 1 \le i \le v \rangle_W.$$

This inclusion is also noted in [6, Remark 4.8]. Analogous transfer maps and divided power operations can also be defined in the odd-characteristic exterior setting, although their explicit forms and interactions with the bilinear relations differ from those in characteristic two.

Our method works directly in the commutative p-truncated polynomial algebra. In particular, the square terms of the public quadratics survive, and their polarizations contribute diagonal terms to the bilinear relations. The graded dimensions and the degrees used for recovery also change: the top degree of the truncated algebra on o oil variables is $o(p-1)$, compared with o in the exterior algebra. Thus the two methods extend the characteristic-two construction in different ways. Their relation spaces and matrix sizes need not agree.

Despite these differences, both methods can be described in the algebraic language developed in this paper. The ideal inclusion is their common foundation, see also [6, §4.2]. Divided powers and transfer maps add further relations that preserve the hidden structure, with the aim of lowering the solving degree and hence the computational cost. Thus the construction of Merz and Ran [15, §5.3] and our construction share the same underlying principle, while working in different algebras: their method uses exterior algebras, whereas ours uses p-truncated polynomial algebras.

Remark 5.2. In odd characteristic, the relative efficiency of the exterior-algebra and p-truncated polynomial-algebra methods remains to be determined. The exterior construction starts from the m bilinear tensors β_i, whereas our construction includes both the pure quadratic relations $f_i^{(s)}$ and the mixed relations $q_i^{(jk)}$. These additional relations may lower the degree required for recovery, although the transfer maps introduce dependencies among them. On the other hand, at the same multidegree λ, the ambient monomial basis in the exterior construction has $\prod_{s=1}^{\ell} \binom{n}{\lambda_s}$ elements, compared with $\prod_{s=1}^{\ell} T(n,p,\lambda_s)$ in our construction. The latter is at least as large, since every squarefree monomial is also a p-truncated monomial. Thus which method is more efficient will be an interesting question for further investigation.

6 Results and Applications

The ideal inclusions established above yield the following criterion for recovering a graded component of the hidden ideal from public relations.

Theorem 6.1. *Let $\lambda \in \mathbb{N}^{\ell}$ have at least one positive entry, where $\ell = 2$ denotes the bigraded construction. Put $W_\lambda = \mathcal{I}_{r,\mathrm{dp},\lambda} + \Xi_\lambda$. Suppose that*

$$W_\lambda = \mathcal{J}_\lambda + \Xi_\lambda, \qquad W_\lambda \neq \mathcal{R}_\lambda.$$

Then one can recover $\mathcal{J}$ from the kernel of the Macaulay matrix M_λ representing W_λ in $\mathcal{R}_\lambda$.

Proof. See [6, Theorem 4.5].

Remark 6.2. [15, §6] formulates their recovery condition as $\dim \operatorname{coker} M_\lambda \leq 1$. Their matrix convention differs from ours by transposition, so the recovery condition in [15, §6] becomes $\dim \ker M_\lambda \leq 1$ in our setting. Indeed,

$$\dim \ker M_\lambda = \dim(\mathcal{R}_\lambda / W_\lambda) = \dim \operatorname{coker}(M_\lambda^{\top}).$$

Theorem 6.1 allows recovery when this dimension is greater than one, provided that $W_\lambda = \mathcal{J}_\lambda + \Xi_\lambda$ and $W_\lambda \neq \mathcal{R}_\lambda$.

6.1 Cost analysis

We construct the Macaulay matrix M_λ with row space $\mathcal{I}_{r,\mathrm{dp},\lambda} + \Xi_\lambda + \mathcal{L}_\lambda$. By Lemmas 4.3 and 4.6, it suffices to include monomial multiples of the first-slot generators $\gamma_{p^a}(f_i^{(1)})$ and the images of source monomials under the adjacent transfer maps $D_{p^a}^{i,i+1}$. Put $\mu = (n_1, \dots, n_\ell)$, $\lambda = (\lambda_1, \dots, \lambda_\ell)$,

$$T(n,p,d) := (1 + t + \cdots + t^{p-1})^n [t^d], \qquad H_{\mu,\lambda} := \prod_{s=1}^{\ell} T(n_s, p, \lambda_s),$$

where $T(n,p,d) = 0$ for $d < 0$ or $d > n(p-1)$. Thus M_λ has $N = H_{\mu,\lambda}$ columns.

Table 2: Row distribution of M_λ in characteristic p. Here density means the number of nonzero entries per row of the Macaulay matrix.

Type	Bounds	Number of rows	Density
$\gamma_{p^j}(f_i^{(1)})$	$0 \le j \le r$	$mH_{\mu,\lambda-2p^j\mathbf{e}_1}$	$\le T(n_1,p,2p^j)$
$D_{p^j}^{i,i+1}$	$1 \le i < \ell$ $0 \le j \le \log_p(\lambda_{i+1})$	$H_{\mu,\lambda+p^j\mathbf{e}_i-p^j\mathbf{e}_{i+1}}$	$\le \binom{\lambda_i+p^j}{p^j}$

Therefore, the complexity of finding the oil space is

$$3 \cdot \rho \cdot H_{\mu,\lambda}^2 \cdot (2(\log_2 q)^2 + \log_2 q),$$

where ρ is the average row density of M_λ:

$$\rho := \frac{1}{P} \sum_{u=1}^{P} \#\{v : 1 \le v \le H_{\mu,\lambda},\ (M_\lambda)_{u,v} \ne 0\} = \frac{\mathrm{nnz}(M_\lambda)}{P}.$$

Here P is the number of rows of M_λ, and $\mathrm{nnz}(M_\lambda)$ denotes its total number of nonzero entries. For the two row families in Table 2, retaining zero and dependent rows, we have

$$P = m \sum_{0 \le j \le r} H_{\mu,\lambda-2p^j\mathbf{e}_1} + \sum_{i=1}^{\ell-1} \sum_{\substack{j \ge 0 \\ p^j \le \lambda_{i+1}}} H_{\mu,\lambda+p^j\mathbf{e}_i-p^j\mathbf{e}_{i+1}}.$$

Taking the weighted average of the density bounds gives

$$\rho \le \overline{\rho} := \frac{1}{P}\left(m \sum_{0 \le j \le r} H_{\mu,\lambda-2p^j\mathbf{e}_1}\, T(n_1,p,2p^j) + \sum_{i=1}^{\ell-1} \sum_{\substack{j \ge 0 \\ p^j \le \lambda_{i+1}}} H_{\mu,\lambda+p^j\mathbf{e}_i-p^j\mathbf{e}_{i+1}} \binom{\lambda_i+p^j}{p^j} \right). \tag{6.1}$$

Thus $\overline{\rho}$ provides an explicit upper bound for the average row density; truncation and vanishing coefficients may make the actual density smaller.

Remark 6.3. In our cost analysis, we only consider the time complexity, the memory complexity is not taken into account.

Remark 6.4. Merz and Ran [15] consider unbalanced projection down. In our setting, however, we have not obtained an effective estimate of the Hilbert series using the same technique. We therefore restrict our analysis to balanced projection down.

6.2 Applications

Table 3 gives the complexity of different kinds of wedge attack on the Round three parameters of QR-UOV [9]. QR-UOV gives four parameters (q, v, m, l). We regard the public key as a plain UOV over $\mathbb{F}_{q^l}$ with v/l vinegar variables, $o = m/l$ oil variables, and m equations. We use the unblanced projection down technique (set $x_{n-k_1+1} \cdots = x_n = y_{n-k_2+1} = \cdots = y_n = 0$) to estimate the complexity of our attack against QR-UOV.

For the classical wedge attack, we randomly fix $0 \leq k \leq o-1$ variables and take the solving degree D as the minimum integer $2 \leq d \leq (o-k)(p-1)$ ($\operatorname{char} \mathbb{F}_q = p$) satisfying

$$\left[\frac{(\sum_{i=0}^{p-1} t^i)^n}{(\sum_{i=1}^{p-1} t^{2i})^m}\right]_+ [t^d] \leq T(o-k, p, d).$$

The the complexity of classical wedge attack is estimated as

$$\min_{0\leq k\leq o-1} 3\binom{n-k+1}{2}(T(n-k,D,p)-T(o-k,D,p)+1)^2(2(\log_2(q^l))^2+\log_2(q^l)).$$

For the level-r divided power wedge attack, see [**?**], the solving degree D is the minimun integer $2 \leq d \leq (o-k)(p-1)$ such that

$$\left[\frac{(\sum_{i=0}^{p-1} t^i)^n}{(\sum_{i=1}^{p^{r+1}-1} t^{2i})^m}\right]_+ [t^d] \leq T(o-k, p, d).$$

And the complexity is estimated via the formula

$$\min_{0\leq k\leq o-1} 3\rho_d(T(n-k,D,p) - T(o-k,D,p)+1)^2(2(\log_2(q^l))^2 + \log_2(q^l)),$$

where ρ_d is the density of the extended Macaulay matrix. From Table 3 we see that complexity of our level-1 schur wedge attack proposed in this paper is lower than the classical wedge attack on some parameters. We highlighted in bold fonts.

We also apply our new attack (set r=1) against MAYO [4] for some addition parameters which the level-1 divided power wedge attack can not apply on them, see Table 4.

Table 3: Complexity of different kinds of wedge attack against QR-UOV

Parameter (q, v, m, l)	[8]	[**?**]	This work $(k_1, k_2, D_1, D_2, \text{bit})$
Security level I			
(127,156,54,3)	164	164	(16,16,16,5,167)
(127,540,60,10)	150	150	(4,4,20,2,165)
(31,165,60,3)	150	150	(18,18,10,9,165)
Security level III			
(127,228,78,3)	231	231	(24,24,17,10,228)
(127,820,90,10)	211	211	(7,7,29,2,227)
(31,246,87,3)	228	228	(27,27,15,12,235)
Security level V			
(127,306,105,3)	290	290	(33,33,20,14,291)
(127,1040,110,10)	276	276	(9,9,29,7,292)

Table 4: Complexity of different kinds of attacks against MAYO

Parameter (n, o, m, q, k)	Best attack	This work $(k_1, k_2, D_1, D_2, \text{ bit})$
Security level III		
(113,13,100,16,8)	217	(0,3,13,10,215)
(112,14,98,16,7)	217	(0,7,17,10,198)
(118,16,96,16,6)	215	(0,11,16,5,196)
Security level V		
(152,16,136,16,9)	294	(0,3,16,13,281)
(151,17,134,16,8)	293	(0,6,17,11,269)
(152,19,132,16,7)	291	(0,14,19,5,230)
(168,22,132,16,6)	293	(0,17,22,5,254)

7 Conclusions and future work

We have developed a formulation of Schur–Macaulay attacks in multigraded p-truncated polynomial algebras, building on the ideal-inclusion framework of Ding, Li, and Tao [6]. This formulation interprets the approach of Merz and Ran [15] as a reconciliation attack involving multiple oil vectors. It agrees with their exterior-algebra construction in characteristic two and gives a distinct extension in odd characteristic. For each pair of variable sets, the transfer operators form an iterative Hasse–Schmidt derivation encoding relations between the corresponding systems.

Divided powers enlarge the public ideal while preserving its inclusion in the hidden oil-space ideal. Combining its graded components with transfer relation spaces yields quotient-dimension criteria for recovering the hidden ideal from Macaulay kernels. The recovery criterion also accommodates kernels of dimension greater than one when the required equality of relation spaces holds. These additional relations can permit recovery at lower degrees, although the resulting cost also depends on matrix dimensions and row density. Under the stated dimension predictions, our analysis gives lower estimated attack complexities for selected NIST third-round QR-UOV and additional MAYO parameter sets.

The Hilbert-series formulas for random systems and the rank predictions for structured UOV systems remain conjectural. Our experiments examine only the latter and are restricted to small instances. Proving the random-system conjectures and determining the range of validity of the structured predictions are therefore central directions for future work. A further question is how the truncated polynomial and exterior-algebra constructions compare in odd characteristic, where their relation spaces and matrix sizes differ. Following Merz and Ran [15, § 7], applications to MinRank and general multivariate quadratic systems also merit investigation.

Acknowledgements We thank Jintai Ding, Lars Ran and Simon-Philipp Merz for helpful discussions. On the other hand, Merz and Ran's work [15] was a major source of inspiration for this paper. Initially, we did not know how to estimate the Hilbert series of $R/\mathcal{I}$ in § 3. Reading their paper provided the framework we needed to formulate Hilbert-series estimates after adjoining divided powers and transfer relations. Once this framework was in place, much of the remaining development followed naturally. Most of the results presented here were thus inspired by the ideas of Merz and Ran, which we adapt to truncated polynomial rings.

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 to edit the text for clarity and grammar. The proof of Proposition 3.6 was generated using GPT-6, checked by the authors, and further validated by experimental checks. The main

framework was developed by the authors building on the work of Ding, Li, and Tao [6], and was also largely inspired by the work of Merz and Ran [15]. 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] Berthelot, P., Ogus, A.: Notes on crystalline cohomology. Princeton University Press, Princeton, NJ; University of Tokyo Press, Tokyo (1978)

[3] 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)

[4] 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

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

[7] Ding, J., Yang, B.Y., Chen, C.H.O., Chen, M.S., Cheng, C.M.: New differential-algebraic attacks and reparametrization of Rainbow. In: Applied Cryptography and Network Security: 6th International Conference, ACNS 2008, New York, NY, USA, June 3-6, 2008. Proceedings 6. pp. 242–257. Springer (2008)

[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] Jantzen, J.C.: Representations of Algebraic Groups, Mathematical Surveys and Monographs, vol. 107. American Mathematical Society, 2 edn. (2003)

[12] 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)

[13] 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)

[14] Kipnis, A., Shamir, A.: Cryptanalysis of the oil and vinegar signature scheme. In: Annual international cryptology conference. pp. 257–266. Springer (1998)

[15] 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
[16] National Institute of Standards and Technology: Nine candidates advance to the third round of the additional digital signatures for the PQC standardization process (2026), https://csrc.nist.gov/News/2026/nist-advances-9-candidates-to-the-3rd-round-of-pqc, announcement, 14 May 2026
[17] 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)
[18] Tseng, P.E., Wang, L.C., Li, P., Kuan, Y.L.: Investigating the wedge map on SNOVA. Cryptology ePrint Archive (2026)

A Divided power structure

Now we prove that the maps $\{\gamma_k\}_{k\geq 0}$ we given in Lemma 2.4 defines 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 $\{f_i\}_{i=1}^{r_0} \in \mathfrak{m}$,

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

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 $u = \sum_{i=1}^{r} u_1$ 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

$$\gamma_p(uf) = \sum_{\alpha\in\mathbb{N}^r,|\alpha|=p} \prod_{i=1}^{r} \gamma_{\alpha_i}(u_i f) = \gamma_p(u_1 f) + \sum_{\alpha\in\mathcal{A}_r(p)} \prod_{i=1}^{r} \gamma_{\alpha_i}(u_i f)$$

$$= u_1^p\gamma_p(f) + \sum_{\alpha\in\mathcal{A}_r(p)} f^p \prod_{i=1}^{r} \frac{u_i^{\alpha_i}}{\alpha_i!} = u_1^p\gamma_p(f) = u^p\gamma_p(f).$$

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

B Explanation of transfer relations

This appendix explains the role of transfer relations, gives their differential description, and proves the Hilbert-series formula for the transfer quotient.

B.1 Why transfers appear naturally

The difficulty with predicting the Hilbert series of $R/\mathcal{I}$, where

$$\mathcal{I} = \langle p_i(\mathbf{x}), q_i(\mathbf{x}, \mathbf{y}), p_i(\mathbf{y}) : 1 \le i \le m\rangle_R,$$

is that these generators are not independent generic forms of their respective bidegrees. They come from the same quadratics and satisfy identities such as

$$p_i(\mathbf{x} + t\mathbf{y}) = p_i(\mathbf{x}) + tq_i(\mathbf{x}, \mathbf{y}) + t^2 p_i(\mathbf{y}).$$

Thus the three families form a system stable under changes of the two variable sets, and their relations couple different bidegrees. A prediction based only on the numbers and bidegrees of the generators would treat the three families as independent generic forms and suggest

$$\left[\frac{h(t)^n h(s)^n}{h(t^2)^m h(ts)^m h(s^2)^m}\right]_+, \quad h(t) = 1 + t + \cdots + t^{p-1}.$$

See also [18, 6]. However, this expression does not account for the identities linking the three families, which arise from the same public quadratics. It therefore need not give the Hilbert series of $R/\mathcal{I}$. The challenge is to incorporate this structure into a useful Hilbert-series prediction.

The transfer maps encode precisely this interaction: D_ℓ moves degree from the $\mathbf{x}$ variables to the $\mathbf{y}$ variables while preserving total degree. Passing to $R_{a,b}/\Xi_{a,b}$ removes the images of these transfers from other bidegrees. As explained in the proof of Proposition 3.6, this is a weight component of the coinvariants for the substitutions $\mathbf{x} \mapsto \mathbf{x} + t\mathbf{y}$. For the ambient ring, its dimension has the explicit formula

$$\dim(R_{a,b}/\Xi_{a,b}) = h_{a,b} - h_{a+1,b-1} \qquad (a \ge b), \qquad h_{a,b} = \dim R_{a,b}.$$

This explains the factor $1 - t^{-1}s$ in the Hilbert-series prediction: it records the passage from the full bigraded space to the transfer quotient.

Adding Ξ does not by itself prove the proposed Hilbert series. The ambient transfer-quotient formula and the equality above are proved, but the dimension formula after imposing the public relations remains Conjecture 3.11. Transfers therefore supply a structural explanation for the form of the prediction, rather than a general guarantee of its validity. For multiple variable sets, the pairwise transfers play the same role and lead to the factors $\prod_{r<s}(1 - t_r^{-1}t_s)$ in Conjecture 4.8.

B.2 Hilbert series of the transfer quotient

Proof (of Proposition 3.6). All the spaces and maps involved commute with field extension, so we may assume that k is algebraically closed. Let $G = \mathrm{SL}_2(k)$ act simultaneously on each pair (x_i, y_i). This action descends to R, since the ideal generated by the pth powers is G-stable. Choose the diagonal torus so that x_i has weight 1 and y_i has weight -1, and let U_- be the unipotent subgroup acting by $x_i \mapsto x_i + \tau y_i$, $y_i \mapsto y_i$. For a fixed total degree d, put

$$V_d := \bigoplus_{a+b=d} R_{a,b}.$$

Its weight-$a - b$ space is $R_{a,b}$. The coefficient operators of U_- are precisely the D_ℓ. Consequently, the weight-$a - b$ component of its coinvariant space is

$$(V_d)_{U_-,\, a-b} = R_{a,b}/\Xi_{a,b}.$$

Here coinvariants mean the quotient by the span of $(u - 1)v$ for $u \in U_-$ and $v \in V_d$; over the infinite field k, this span equals the sum of the images of all D_ℓ, $\ell > 0$.

We use two standard facts about good filtrations: tensor products of modules with good filtrations again have good filtrations, and

$$\dim \mathrm{Hom}_G(\Delta(r), W) = [W : \nabla(r)]$$

for a module W with a good filtration, where $\Delta(r)$ and $\nabla(r)$ are the Weyl and induced modules of highest weight r; see [11, Part II, §4]. We first check that these facts apply here. Set $A = k[x, y]/(x^p, y^p)$ and let A_j be its degree-j component. For $0 \le j \le p - 1$, $A_j = \mathrm{Sym}^j(k^2)$ is simple and is both a Weyl and an induced module. For $p \le j \le 2p-2$, multiplication gives a perfect G-invariant pairing

$$A_j \times A_{2p-2-j} \longrightarrow A_{2p-2} = k\, x^{p-1}y^{p-1} \simeq k.$$

Thus $A_j \simeq A^*_{2p-2-j}$ is again a simple module of highest weight less than p. In particular, every A_j and its dual have good filtrations. The decomposition

$$V_d \simeq \bigoplus_{j_1+\cdots+j_n=d} A_{j_1} \otimes \cdots \otimes A_{j_n}$$

therefore shows that V_d^* has a good filtration.

Write $c_r = [V_d^* : \nabla(r)]$. The dual of $(V_d)_{U_-}$ is $(V_d^*)^{U_-}$. A U_--fixed vector of weight $-r$ is a lowest-weight vector and, by the universal property of the Weyl module, corresponds to a homomorphism $\Delta(r) \to V_d^*$. Such vectors have $r \geq 0$. Hence

$$\dim(V_d)_{U_-,r} = c_r \quad (r \geq 0), \qquad (V_d)_{U_-,r} = 0 \quad (r < 0).$$

Finally, the character of $\nabla(r)$ has weights $r, r-2, \ldots, -r$, each with multiplicity one. Taking the difference between weight multiplicities at r and $r+2$ in the good filtration gives, for $r = a - b \geq 0$,

$$c_r = h_{a,b} - h_{a+1,b-1}.$$

This proves the dimension formula and the vanishing for $a < b$. The series identity follows by taking coefficients in the Hilbert series of R and retaining the stated range of bidegrees. □

C Explanation of the Hilbert series

The formulas in Conjectures 3.11 and 4.8 combine the monomial count of the truncated ring, a transfer correction, and a formal count of the quadratic relations. We explain these ingredients below. The explanation motivates the formulas but does not prove the conjectures.

C.1 The ambient ring and the transfer correction

Put $h(t) = 1 + t + \cdots + t^{p-1}$. Each variable has exponents from 0 to $p-1$, so the multigraded Hilbert series of the ambient ring $\mathbb{F}_q[\mathbf{x}_1, \ldots, \mathbf{x}_\ell]/\langle x_{ij}^p : 1 \le i \le \ell,\ 1 \le j \le n\rangle$ is

$$A(\mathbf{t}) = \prod_{r=1}^{\ell} h(t_r)^n.$$

For two variable sets, Proposition 3.6 (proved in Appendix B.2) shows that quotienting by transfers replaces the coefficient $\dim R_{a,b}$ by $\dim R_{a,b} - \dim R_{a+1,b-1}$ on $a \ge b$. Thus it introduces the factor $1 - t^{-1}s$ and restriction to $a \ge b \ge 0$. In the multigraded formula, the analogous correction is

$$W(\mathbf{t}) = \prod_{r<s} (1 - t_r^{-1} t_s),$$

with restriction to Λ_ℓ. These factors encode transfers between variable sets; they do not represent additional polynomial generators of the public ideal.

C.2 The quadratic-relation factors

For each public quadratic, polarization supplies one relation of degree $2\mathbf{e}_r$ for each variable set and one of degree $\mathbf{e}_r + \mathbf{e}_s$ for each pair $r < s$. The formal correction associated with these m families is

$$Q(\mathbf{t}) = \prod_{r=1}^{\ell} (1 - t_r^2)^m \prod_{r<s} (1 - t_r t_s)^m.$$

To see the origin of such factors, let f be a homogeneous element of multidegree δ in a multigraded algebra B. The multiplication exact sequence is

$$0 \longrightarrow \mathrm{Ann}_B(f)[-\delta] \longrightarrow B[-\delta] \xrightarrow{\cdot f} B \longrightarrow B/(f) \longrightarrow 0.$$

Taking multigraded Hilbert series gives the exact identity

$$H_{B/(f)}(\mathbf{t}) = (1 - \mathbf{t}^\delta) H_B(\mathbf{t}) + \mathbf{t}^\delta H_{\mathrm{Ann}_B(f)}(\mathbf{t}).$$

If multiplication by f is injective, the annihilator term vanishes. For several homogeneous relations $f_1, \ldots, f_N$, write $\delta_i = \deg(f_i)$ and $I = \langle f_1, \ldots, f_N\rangle$. Introduce symbols u_i of multidegree δ_i and put $E = \bigoplus_{i=1}^N B[-\delta_i]u_i$. The Koszul complex $K_\bullet = K_\bullet(f_1, \ldots, f_N; B)$ is

$$0 \longrightarrow K_N \xrightarrow{d_N} K_{N-1} \longrightarrow \cdots \longrightarrow K_2 \xrightarrow{d_2} K_1 \xrightarrow{d_1} K_0 \longrightarrow 0,$$

where $K_0 = B$ and

$$K_k = \bigwedge_B^k E = \bigoplus_{1 \le i_1 < \cdots < i_k \le N} B[-(\delta_{i_1} + \cdots + \delta_{i_k})](u_{i_1} \wedge \cdots \wedge u_{i_k}).$$

Here the shift convention is $H_{B[-\eta]}(\mathbf{t}) = \mathbf{t}^{\eta} H_B(\mathbf{t})$. The degree-preserving differential is given by

$$d_k(b\, u_{i_1} \wedge \cdots \wedge u_{i_k}) = \sum_{a=1}^{k} (-1)^{a-1} b f_{i_a}\, u_{i_1} \wedge \cdots \wedge \widehat{u_{i_a}} \wedge \cdots \wedge u_{i_k},$$

where the hat means that the indicated factor is omitted. Commutativity of the f_i gives $d_{k-1} d_k = 0$. In particular,

$$d_1(b u_i) = b f_i, \qquad d_2(b u_i \wedge u_j) = b(f_i u_j - f_j u_i).$$

Thus $\mathrm{im}(d_1) = I$ and $H_0(K_\bullet) = B/I$. The kernel of d_1 consists of the syzygies $\sum_i b_i f_i = 0$, whereas the image of d_2 consists of the standard Koszul syzygies $f_i u_j - f_j u_i$ and their B-linear combinations. The higher differentials encode relations among these syzygies and, successively, relations among those relations.

The alternating Hilbert series of the terms of this complex is

$$\begin{aligned} \sum_{k=0}^{N} (-1)^k H_{K_k}(\mathbf{t}) &= H_B(\mathbf{t}) \sum_{S \subseteq \{1,\ldots,N\}} (-1)^{|S|} \mathbf{t}^{\sum_{i \in S} \delta_i} \\ &= H_B(\mathbf{t}) \prod_{i=1}^{N} (1 - \mathbf{t}^{\delta_i}). \end{aligned}$$

Consequently, the exact identity is

$$H_{B/I}(\mathbf{t}) = H_B(\mathbf{t}) \prod_{i=1}^{N} (1 - \mathbf{t}^{\delta_i}) + \sum_{k=1}^{N} (-1)^{k+1} H_{H_k(K_\bullet)}(\mathbf{t}).$$

If $f_1, \ldots, f_N$ form a regular sequence, the positive-degree Koszul homology vanishes, and the augmented complex ending in $B \longrightarrow B/I \longrightarrow 0$ is exact. In that case the product alone gives $H_{B/I}$. For arbitrary relations, $H_1(K_\bullet)$ measures syzygies not generated by the standard Koszul syzygies, and the higher homology provides further corrections. For $N = 1$, one has $H_1(K_\bullet) = \mathrm{Ann}_B(f_1)[-\delta_1]$, recovering the multiplication identity above.

In our finite truncated ring, positive-degree elements are nilpotent, so this regular-sequence argument cannot apply globally. Moreover, the polarized equations come from the same public quadratics. The product Q is therefore a formal dimension model, not the consequence of an established regular sequence. The Conjectures 3.11 and 4.8 assert that, after adjoining divided powers and passing to the transfer quotient, the surviving dimensions agree with the specified truncation of this model.

C.3 Divided powers and positive-part truncation

The higher divided powers are determined by the original quadratics and satisfy product identities; they are not independent generic equations. The conjectured expression does not assign a separate factor to each generator $\gamma_{p^e}(f)$; dependence alone, however, does not justify omitting such factors. Proposition 3.4 simplifies the divided-power relations modulo transfers, but they do not establish the proposed dimension count. Explaining fully why the divided-power closure yields this particular count remains part of the conjectural problem.

The formal product AWQ may have negative coefficients, whereas quotient dimensions are nonnegative. The operation $[\,\cdot\,]_+$ in Definitions 3.9 and 4.7 prescribes where the formal count is retained and where it is replaced by zero.

Combining the ingredients gives

$$\big[[A(\mathbf{t})W(\mathbf{t})Q(\mathbf{t})]_+\big]_{\Lambda_\ell}.$$

For $\ell = 2$, the expression inside the positive-part operation is

$$h(t)^n h(s)^n (1-t^{-1}s)(1-t^2)^m (1-ts)^m (1-s^2)^m,$$

which is the expression in Conjecture 3.11. These conjectures concern random systems without a prescribed oil space. The coefficientwise maximum with $\mathcal{H}_0$ used in our structured UOV experiments is a separate prediction, accounting for the quotient forced by the hidden oil space.