Key Recovery Attacks on UOV Using p^ℓ-truncated Polynomial Rings

Hiroki Furue[1] and Yasuhiko Ikematsu[2]

[1] NTT Social Informatics Laboratories, Tokyo, Japan
hiroki.furue@ntt.com
[2] Kyushu University, Fukuoka, Japan
ikematsu@imi.kyushu-u.ac.jp

Abstract. The unbalanced oil and vinegar signature scheme (UOV) was proposed by Kipnis et al. in 1999 as a multivariate-based scheme. UOV is regarded as one of the most promising candidates for post-quantum cryptography owing to its short signatures and fast performance. Recently, Ran proposed a new key recovery attack on UOV over a field of even characteristic, reducing the security of its proposed parameters. Furthermore, Jin et al. generalized Ran's attack to schemes over a field of arbitrary characteristic by exploiting the structure of the symmetric algebra.

In this work, we propose a new framework for recovering the secret subspace of UOV over a finite field $\mathbb{F}_{p^e}$ by generalizing these preceding results. First, we show that a key recovery against UOV can be successfully performed using the XL algorithm by exploiting the structure of the p-truncated polynomial ring $R^{(p)} = \mathbb{F}_{p^e}[x_1, \ldots, x_n]/\langle x_1^p, \ldots, x_n^p \rangle$. This result simplifies the description of the attacks proposed by Jin et al. by formulating them in terms of the polynomial ring, independent of the structure of the symmetric algebra. Second, we generalize this result to the polynomial rings of more general forms, namely, the p^ℓ-truncated polynomial rings $R^{(p^\ell)}$ for any $1 \leq \ell \leq e$. This result is due to our description in terms of the polynomial ring and can relax the constraints on the solving degree of the XL algorithm using $R^{(p^\ell)}$ by taking a larger ℓ. Finally, we consider performing the reconciliation and intersection attacks using the p^ℓ-truncated polynomial rings against UOV. In particular, we consider the intersection attack using this framework, which has not been addressed in previous analyses.

Based on our complexity estimation, we confirm that the optimal complexity of the reconciliation attack using the proposed framework is consistent with that of the symmetric-algebra attack by Jin et al. We further show that the intersection attack using the proposed framework outperforms the reconciliation attack against the proposed parameters of UOV and reduces the security of multiple parameters compared to their claimed security levels.

Keywords: post-quantum cryptography, multivariate-based cryptography, unbalanced oil and vinegar (UOV), reconciliation attack, intersection attack

1 Introduction

Currently used public-key cryptosystems such as RSA and ECC can be broken in polynomial time using Shor's quantum algorithm [21]. This fact has motivated research on post-quantum cryptography (PQC), which aims to remain secure against attacks by quantum computers. Several classes of PQC candidates have been proposed, including lattice-based, code-based, and isogeny-based cryptography. Multivariate-based cryptography is one such candidate for PQC and is constructed based on the hardness of solving the multivariate quadratic (MQ) problem. This problem is to find a solution to a system of multivariate quadratic polynomial equations over a finite field. It is proven to be NP-hard [8] and is widely believed to be secure against attacks by quantum computers.

The unbalanced oil and vinegar signature scheme (UOV) was proposed by Kipnis et al. at EUROCRYPT 1999 [11]. UOV is a well-established multivariate signature scheme owing to its short signatures and fast performance. It is also highly regarded from a security perspective, as it has withstood various types of attacks for over 20 years. The security of UOV can be reduced to the hardness of two problems: solving the MQ problem and distinguishing a UOV public map from a random MQ map. In particular, key recovery attacks on UOV aim to recover a solution space given as a vector space from a given UOV public key map in order to distinguish it from a random MQ map.

The U.S. National Institute of Standards and Technology (NIST) has been running a PQC standardization project since 2016 [15]. Through this process, NIST has selected two key-encapsulation mechanisms and three digital signature algorithms for standardization [17, 18]. In addition, NIST announced the start of a new selection process for additional digital signature schemes in 2022 [16] to ensure algorithmic diversity. More recently, in May 2026, NIST selected 9 signature schemes as third-round candidates in this additional selection process [19]. Among these candidates, UOV and its three variants, MAYO [2], QR-UOV [7], and SNOVA [22], were selected as multivariate-based signature schemes. These facts indicate that UOV plays an important role in cryptographic standardization, making its security analysis an important task.

Recently, novel key recovery attacks on UOV that effectively exploit the algebraic structure of the public key maps have been proposed. In this direction, Ran [20] proposed a key recovery attack on UOV over a field of characteristic $p = 2$. This attack regards UOV public maps as elements of the exterior algebra and recovers the secret key by exploiting this structure. As a result, the attack is more efficient than existing key recovery attacks for some of the proposed parameters of UOV and MAYO. In addition, Jin et al. [10] proposed new key recovery attacks on UOV by generalizing Ran's attack to arbitrary characteristic p. They achieve this generalization by replacing the exterior algebra used in Ran's attack with the reduced symmetric algebra $\overline{\mathrm{Sym}}(\mathbb{F}_q^n)$. Notably, when $p = 2$, the reduced symmetric algebra coincides with the exterior algebra. This approach can be interpreted as performing the reconciliation attack, which is one of the most standard key recovery attacks on UOV, using $\overline{\mathrm{Sym}}(\mathbb{F}_q^n)$. Jin et al.'s attack transforms the UOV public key map into elements of the reduced symmetric

algebra $\overline{\mathrm{Sym}}(\mathbb{F}_q^n)$ and applies the XL algorithm [5] to recover the secret key on $\overline{\mathrm{Sym}}(\mathbb{F}_q^n)$. Consequently, this attack extends the applicability of Ran's original attack, in particular enabling its application to schemes defined over a field of odd characteristic, such as QR-UOV.

Our contributions. In this work, we propose a new framework for performing key recovery attacks on UOV. This framework is constructed as a generalization of the line of attacks initiated by prior works by Ran [20] and Jin et al. [10]. Using this framework, we evaluate the security of the proposed parameters of UOV and its variants.

We first provide a theoretical analysis of a multivariate polynomial system in N variables $\mathbf{x} = (x_1, \ldots, x_N)^\top$ over $\mathbb{F}_q$, where $q = p^e$ and p is a prime number, whose solution space contains a vector space $\mathcal{A} \subset \mathbb{F}_q^N$. This polynomial system corresponds to the public key map of UOV. In our analysis, we show that the system can be solved on the p-truncated polynomial ring $R^{(p)} = \mathbb{F}_q[x_1, \ldots, x_N]/\langle x_1^p, \ldots, x_N^p \rangle$. This result simplifies the description of the attack by Jin et al. by using only polynomial ring techniques, without relying on the reduced symmetric algebra $\overline{\mathrm{Sym}}(\mathbb{F}_q^n)$. It is known that when constructing the Macaulay matrix from the given system in $\mathbb{F}_q[\mathbf{x}]$, its right kernel has vectors derived from the solution space $\mathcal{A}$. In this work, we show that some of these vectors also lie in the right kernel of the Macaulay matrix constructed over the p-truncated polynomial ring $R^{(p)}$. Furthermore, we show that the system can be solved in a similar manner using the p^ℓ-truncated polynomial ring $R^{(p^\ell)} = \mathbb{F}_q[x_1, \ldots, x_N]/\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell} \rangle$ for $\ell \in [e]$. This generalization is achieved through our polynomial ring–based description. From an efficiency point of view, our complexity analysis indicates that solving the system using the p^ℓ-truncated polynomial ring is most efficient when $\ell = 1$. On the other hand, larger values of ℓ make the attack applicable to a larger range of parameters by increasing the admissible range of the degree d of the Macaulay matrix.

We then apply the proposed framework based on the p^ℓ-truncated polynomial ring to recover the secret subspace of UOV. Specifically, we consider the application to two types of key recovery attacks, the reconciliation attack [6] and the intersection attack [1]. The symmetric-algebra attack by Jin et al. only considers applying the reconciliation attack, but we also take into account the intersection attack within the proposed framework. In addition, we explain how to recover a solution vector from a kernel vector of the Macaulay matrix constructed over the p^ℓ-truncated polynomial ring. We estimate the complexity of both the reconciliation and intersection attacks using the proposed framework.

We finally evaluate the security of the proposed parameter sets of UOV and its variants against the key recovery attacks using the proposed framework. For the parameters of UOV, we confirm that the intersection attack is more efficient than the reconciliation attack, which corresponds to the original attack considered in [10], when applied using the proposed p^ℓ-truncated framework. In particular, for the parameter sets `uov-Ip`, `uov-III`, and `uov-V` with $q = 256$, the

intersection attack using the proposed framework reduces the security by 15–50 bits compared to their claimed security levels.

Organization. The rest of this paper is organized as follows. Section 2 describes the construction of UOV and existing attacks on the scheme. Section 3 provides a theoretical analysis of solving polynomial systems whose solution space contains a vector space, which corresponds to key recovery on UOV, over the truncated polynomial rings with specific forms. Section 4 describes the proposed framework based on the p^ℓ-truncated polynomial ring and applies it to key recovery attacks on UOV. Section 5 evaluates the security of the proposed parameters of UOV and its variants against the proposed framework. Finally, Section 6 summarizes the main results of this paper.

2 Preliminaries

This section provides a description of the unbalanced oil and vinegar signature scheme (UOV) [11] and reviews known attacks on the scheme. Throughout this paper, let $\mathbb{F}_q$ denote the finite field with q elements. For a polynomial ring $\mathbb{F}_q[\mathbf{x}]$, we denote by $\mathrm{Mon}_d(\mathbf{x})$ the set of degree d monomials in $\mathbf{x}$. We also denote by $[n]$ the set of positive integers $\{1, \ldots, n\}$ for $n \in \mathbb{N}$.

2.1 UOV

This subsection describes the structure of the unbalanced oil and vinegar signature scheme (UOV). Let v, o, and m be positive integers, and set $n = v+o$. Here, n and m denote the numbers of variables and equations, respectively. For variables $\mathbf{x} = (x_1, \ldots, x_n)^\top$ over $\mathbb{F}_q$, the first v variables $x_1, \ldots, x_v$ are called *vinegar variables*, and the remaining o variables $x_{v+1}, \ldots, x_n$ are called *oil variables*.

We first describe the key generation procedure of UOV. We design $\mathcal{F} = (f_1, \ldots, f_m) : \mathbb{F}_q^n \to \mathbb{F}_q^m$, called a *central map*, such that each f_k with $k \in [m]$ is a quadratic polynomial of the form

$$f_k(x_1, \ldots, x_n) = \sum_{i=1}^{v} \sum_{j=i}^{n} \alpha_{i,j}^{(k)} x_i x_j \tag{1}$$

where $\alpha_{i,j}^{(k)} \in \mathbb{F}_q$. Next, we randomly choose an invertible linear map $\mathcal{S} : \mathbb{F}_q^n \to \mathbb{F}_q^n$ to hide the structure of $\mathcal{F}$. The public key $\mathcal{P}$ is then provided as a quadratic map computed by

$$\mathcal{P} = \mathcal{F} \circ \mathcal{S} : \mathbb{F}_q^n \to \mathbb{F}_q^m,$$

and the secret key consists of $\mathcal{F}$ and $\mathcal{S}$.

We next describe the inversion of the central map $\mathcal{F}$. Given a vector $\mathbf{y} \in \mathbb{F}_q^m$, we aim to find $\mathbf{x} \in \mathbb{F}_q^n$ such that $\mathcal{F}(\mathbf{x}) = \mathbf{y}$. To this end, we first choose random values $a_1, \ldots, a_v \in \mathbb{F}_q$ for the vinegar variables. Then, by substituting these values into the central map, the equation $\mathcal{F}(a_1, \ldots, a_v, x_{v+1}, \ldots, x_n) = \mathbf{y}$

becomes a linear system in the oil variables $x_{v+1}, \ldots, x_n$, consisting of o variables and m equations, due to the construction of the central map (1). In UOV, the parameters are typically chosen such that $o = m$, and hence this linear system admits a solution with high probability. If no solution exists, new random values $a'_1, \ldots, a'_v$ are selected, and the above procedure is repeated.

By using this inversion procedure, a signature is generated as follows: Given a message $\mathbf{m} \in \mathbb{F}_q^m$ to be signed, we find a vector $\mathbf{m}_1 \in \mathbb{F}_q^n$ satisfying the equation $\mathcal{F}(\mathbf{m}_1) = \mathbf{m}$. The signature for the message $\mathbf{m}$ is then computed as $\mathbf{s} = \mathcal{S}^{-1}(\mathbf{m}_1) \in \mathbb{F}_q^n$. The verification step is performed by checking whether $\mathcal{P}(\mathbf{s}) = \mathbf{m}$.

Finally, we introduce matrix representations of the public and secret keys of UOV. For each polynomial p_k with $k \in [m]$ in the public key $\mathcal{P} = (p_1, \ldots, p_m)$, there exists an n-by-n matrix P_k such that $p_k(\mathbf{x}) = \mathbf{x}^\top \cdot P_k \cdot \mathbf{x}$. Similarly, for each polynomial f_k with $k \in [m]$ in the central map $\mathcal{F}$, we associate an n-by-n matrix F_k. These matrices P_k and F_k are taken to be symmetric matrices when q is odd, and upper triangular matrices when q is even. In the even characteristic case, we define $P'_k = P_k + P_k^\top$ and $F'_k = F_k + F_k^\top$, which represent the polar forms of the corresponding quadratic forms. For convenience, in the odd characteristic case, we set $P'_k = P_k$ and $F'_k = F_k$. From equation (1), each matrix F'_k can be taken in the following form

$$\begin{pmatrix} *_{v\times v} & *_{v\times o} \\ *_{o\times v} & 0_{o\times o} \end{pmatrix}, \tag{2}$$

whose o-by-o lower-right submatrix is the zero matrix. Furthermore, from the relation $\mathcal{P} = \mathcal{F} \circ \mathcal{S}$, each matrix P'_k satisfies

$$P'_k = S^\top \cdot F'_k \cdot S, \tag{3}$$

where $S \in \mathbb{F}_q^{n\times n}$ is defined by $\mathcal{S}(\mathbf{x}) = S \cdot \mathbf{x}$.

2.2 Known Attacks on UOV

This paper deals with key recovery attacks on UOV, which are formulated as follows. Given a public key $\mathcal{P}$, the goal is to find a corresponding central map $\mathcal{F}$ and a linear map $\mathcal{S}$ such that $\mathcal{P} = \mathcal{F} \circ \mathcal{S}$. More specifically, key recovery attacks aim to recover the subspace $S^{-1}\mathcal{O}$ of $\mathbb{F}_q^n$, where $\mathcal{O}$ denotes the oil subspace defined by

$$\mathcal{O} := \left\{ (0, \ldots, 0, a_1, \ldots, a_o)^\top \;\middle|\; a_i \in \mathbb{F}_q \right\}.$$

Various key recovery attacks on UOV have been proposed in the literature, including the Kipnis–Shamir [12], reconciliation [6], intersection [1], and rectangular MinRank [1] attacks. In this paper, we focus on two of these key recovery attacks, the reconciliation and intersection attacks, and describe their details in what follows.

Reconciliation attack. The reconciliation attack treats $\mathbf{x} \in S^{-1}\mathcal{O}$ as an unknown. From the construction of the public key map, the subspace $S^{-1}\mathcal{O}$ vanishes on $\mathcal{P}$; that is, $\mathbf{x}$ can be obtained by solving the following quadratic system

$$\mathbf{x}^\top P_k \mathbf{x} = 0 \quad (k \in [m]).$$

The dimension of $S^{-1}\mathcal{O}$ is o, and thus, by imposing affine constraints, we can reduce it to a system of m equations in $n - o = v$ variables.

In the case where $v > m$, which is generally satisfied in UOV to prevent the Kipnis-Shamir attack [12], the system of equations $\mathbf{x}^\top P_k \mathbf{x} = 0$ admits a large number of solutions. Therefore, in order to determine the subspace uniquely, it is necessary to simultaneously find multiple elements $\mathbf{x}_1, \ldots, \mathbf{x}_\ell$ in $S^{-1}\mathcal{O}$ by solving the following system:

$$\begin{cases} \mathbf{x}_i^\top P_k \mathbf{x}_i = 0 & (k \in [m], i \in [\ell]), \\ \mathbf{x}_i^\top P'_k \mathbf{x}_j = 0 & (k \in [m], 1 \le i < j \le \ell). \end{cases}$$

Intersection attack. Let W_1 and W_2 be two invertible matrices randomly chosen from the vector space spanned by the symmetric public key matrices $P'_1, \ldots, P'_m$. We consider multiplying W_1 and W_2 by a vector $\mathbf{y} \in S^{-1}\mathcal{O}$. Then, from (2) and (3), the vectors $W_i\mathbf{y}$ for $i = 1, 2$ belong to $S^\top\mathcal{V}$, where $\mathcal{V}$ denotes the vinegar space defined as the orthogonal complement of the oil space $\mathcal{O}$. This implies that $W_i S^{-1}\mathcal{O} \subseteq S^\top\mathcal{V}$. In the case where $v < 2o$, the subspaces $W_1 S^{-1}\mathcal{O}$ and $W_2 S^{-1}\mathcal{O}$ always intersect, since the dimensions of $\mathcal{V}$ and $\mathcal{O}$ are v and o, respectively. The intersection attack exploits this fact by regarding the intersection space as the solution space. Concretely, the attack solves the following system of equations for $\mathbf{x} \in \mathbb{F}_q^n$:

$$\begin{cases} (W_1^{-1}\mathbf{x})^\top P_k (W_1^{-1}\mathbf{x}) = 0, \\ (W_1^{-1}\mathbf{x})^\top P'_k (W_2^{-1}\mathbf{x}) = 0, \\ (W_2^{-1}\mathbf{x})^\top P_k (W_2^{-1}\mathbf{x}) = 0, \end{cases} \quad (k \in [m]). \tag{4}$$

These equations hold because, if $\mathbf{x} \in (W_1 S^{-1}\mathcal{O}) \cap (W_2 S^{-1}\mathcal{O})$, then $W_i^{-1}\mathbf{x} \in S^{-1}\mathcal{O}$ for $i = 1, 2$.

The intersection attack can also be performed in the case where $v \ge 2o$. In this setting, the subspaces $W_1 S^{-1}\mathcal{O}$ and $W_2 S^{-1}\mathcal{O}$ intersect with probability $q^{-(v-2o+1)}$ [1]. Therefore, by solving the system (4), we can recover a vector in $S^{-1}\mathcal{O}$ with this probability.

Furthermore, the attack can be generalized by choosing more than two matrices from the vector space spanned by $P'_1, \ldots, P'_m$. However, in this paper, we restrict our attention to the case of two matrices W_1 and W_2, since the attack using the proposed framework is most efficient in this setting against the proposed parameters of UOV and its variants. For details of the general case, we refer the reader to [1].

3 Macaulay Matrix using Truncated Polynomial Ring

Let $R := \mathbb{F}_q[\mathbf{x}]$ in N variables $\mathbf{x} = (x_1, \ldots, x_N)^\top$ over $\mathbb{F}_q$, where $q = p^e$ and p is a prime number, and $f_1(\mathbf{x}), \ldots, f_M(\mathbf{x}) \in R$ be M homogeneous quadratic polynomials. We suppose that the solution space of the system

$$f_1(\mathbf{x}) = \cdots = f_M(\mathbf{x}) = 0 \tag{5}$$

contains a vector space $\mathcal{A} \subset \mathbb{F}_q^N$ over $\mathbb{F}_q$ with $\dim_{\mathbb{F}_q} \mathcal{A} = \alpha$. In this section, we revisit the XL algorithm [5] that solves such a system, and study the structure of the Macaulay matrix constructed using the p^ℓ-truncated polynomial ring $R^{(p^\ell)} := R \Big/ \left\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell} \right\rangle$ with $\ell \in [e]$.

3.1 XL Algorithm for Linear Solution Space

This subsection first recalls how the XL algorithm [5] constructs and uses the Macaulay matrix for the polynomial system (5). We then analyze the structure of the right kernel space of the Macaulay matrix in the case where $\alpha > 1$.

We recall the construction of the Macaulay matrix for a chosen degree $d \in \mathbb{N}$. We multiply each polynomial f_i by all monomials of degree $d-2$ in $\mathbf{x}$ to generate the following set of degree-d polynomials:

$$G = \{af_i \mid a \in \mathrm{Mon}_{d-2}(\mathbf{x}), i \in [M]\}.$$

We then construct a coefficient matrix $\mathcal{M}_d$ for the set G, which is called the Macaulay matrix. The rows of $\mathcal{M}_d$ correspond to the polynomials in G, the columns correspond to the monomials in $\mathrm{Mon}_d(\mathbf{x})$, and each entry corresponding to $g \in G$ and $a \in \mathrm{Mon}_d(\mathbf{x})$ is given by the coefficient of a in g.

We then show that we can construct a right kernel vector of the Macaulay matrix $\mathcal{M}_d$ from a solution of the system (5). We denote by $\mathbf{v}_d(\mathbf{x})$ the vector of monomials in $\mathrm{Mon}_d(\mathbf{x})$ corresponding to the columns of $\mathcal{M}_d$, defined as

$$\mathbf{v}_d(\mathbf{x}) = (x_1^d, x_1^{d-1}x_2, \ldots, x_N^d)^\top.$$

Although the order of monomials is not essential, we use the lexicographical order for convenience. For a solution $\mathbf{a} \in \mathbb{F}_q^n$ of the system (5), i.e., $\mathbf{a} \in \mathcal{A}$, we evaluate the vector $\mathbf{v}_d(\mathbf{a})$. This vector turns out to be a right kernel vector of $\mathcal{M}_d$, since for any row vector $\mathbf{b}^\top$ of $\mathcal{M}_d$ corresponding to a polynomial $g = af_i \in G$, we have

$$\mathbf{b}^\top \cdot \mathbf{v}_d(\mathbf{a}) = g(\mathbf{a}) = af_i(\mathbf{a}) = 0.$$

In the case where the dimension of the solution space is expected to be 1, the XL algorithm finds a solution using the Macaulay matrix $\mathcal{M}_d$ as follows. Let I be the homogeneous ideal generated by $f_1, \ldots, f_M$ in R, and let I_d denote the subspace of degree-d homogeneous polynomials in I. If I_d has codimension 1 in R_d, then the vector $\mathbf{v}_d(\mathbf{a})$ corresponding to the solution $\mathbf{a}$ can be recovered by computing the right kernel of the Macaulay matrix $\mathcal{M}_d$. Once the kernel vector

$\mathbf{v}_d(\mathbf{a})$ is obtained, it is straightforward to recover the solution $\mathbf{a}$. In the case of $n > m$, the XL algorithm is generally applied after reducing the dimension of the solution space to 1 by fixing some variables in $\mathbf{x}$.

We then investigate the right kernel space of the Macaulay matrix $\mathcal{M}_d$ when $\alpha > 1$. Let $\mathbf{a}_1, \ldots, \mathbf{a}_\alpha$ be a basis of the solution space $\mathcal{A}$. Any solution in $\mathcal{A}$ can then be written as $\sum_{i=1}^{\alpha} s_i \mathbf{a}_i$ for some $s_1, \ldots, s_\alpha \in \mathbb{F}_q$. For such a solution, the vector $\mathbf{v}_d\left(\sum_{i=1}^{\alpha} s_i \mathbf{a}_i\right)$ is also a right kernel vector of the Macaulay matrix $\mathcal{M}_d$. Let $\mathbf{t} = (t_1, \ldots, t_\alpha)^\top$ be a set of α variables. Since $\mathbf{v}_d\left(\sum_{i=1}^{\alpha} t_i \mathbf{a}_i\right)$ is a vector whose entries are homogeneous polynomials of degree d in $\mathbf{t} = (t_1, \ldots, t_\alpha)^\top$, it can be expanded with respect to the monomials in $\mathrm{Mon}_d(\mathbf{t})$ as

$$\mathbf{v}_d\left(\sum_{i=1}^{\alpha} t_i \mathbf{a}_i\right) = t_1^d \mathbf{b}_{t_1^d} + t_1^{d-1} t_2 \mathbf{b}_{t_1^{d-1} t_2} + \cdots + t_\alpha^d \mathbf{b}_{t_\alpha^d} = \sum_{\lambda \in \mathrm{Mon}_d(\mathbf{t})} \lambda \cdot \mathbf{b}_\lambda,$$

where each vector $\mathbf{b}_\lambda$ lies in $\mathbb{F}_q^{\binom{N+d-1}{d}}$, since $\mathbf{v}_d(\mathbf{x})$ is a $\binom{N+d-1}{d}$-dimensional vector. Then we have the following lemma.

Lemma 1. *For each $\lambda \in \mathrm{Mon}_d(\mathbf{t})$, the vector $\mathbf{b}_\lambda$ is a right kernel vector of $\mathcal{M}_d$. Moreover, the set of vectors $\mathbf{b}_\lambda$ is linearly independent.*

Proof. First, we show that $\mathbf{b}_\lambda$ is a right kernel vector of $\mathcal{M}_d$. It is clear that $\sum_{i=1}^{\alpha} s_i \mathbf{a}_i$ is a solution to the system (5) even for $s_1, \ldots, s_\alpha \in \bar{\mathbb{F}}_q$, where $\bar{\mathbb{F}}_q$ is an algebraic closure of $\mathbb{F}_q$. Thus, $\mathbf{v}_d\left(\sum_{i=1}^{\alpha} s_i \mathbf{a}_i\right) = \sum_{\lambda \in \mathrm{Mon}_d(\mathbf{t})} \lambda(\mathbf{s}) \cdot \mathbf{b}_\lambda$ is a right-kernel of $\mathcal{M}_d$ for $\mathbf{s} = (s_1, \ldots, s_\alpha) \in \bar{\mathbb{F}}_q^\alpha$. Since the monomial functions

$$\bar{\mathbb{F}}_q^\alpha \ni \mathbf{s} \mapsto \lambda(\mathbf{s}) \in \bar{\mathbb{F}}_q \qquad (\lambda \in \mathrm{Mon}_d(\mathbf{t})),$$

are linearly independent over $\bar{\mathbb{F}}_q$, we may choose sufficiently many vectors

$$\mathbf{s}_1, \ldots, \mathbf{s}_u \in \bar{\mathbb{F}}_q^\alpha$$

such that the matrix

$$(\lambda(\mathbf{s}_j))_{j,\lambda}$$

has full rank. Then, using the relation $\mathbf{v}_d\left(\sum_{i=1}^{\alpha} s_{j,i} \mathbf{a}_i\right) = \sum_{\lambda \in \mathrm{Mon}_d(\mathbf{t})} \lambda(\mathbf{s}_j) \cdot \mathbf{b}_\lambda$ for any $\mathbf{s}_j = (s_{j,1}, \ldots, s_{j,\alpha})^\top$, each coefficient vector $\mathbf{b}_\lambda$ can be recovered as a linear combination of the vectors $\mathbf{v}_d\left(\sum_{i=1}^{\alpha} s_{j,i} \mathbf{a}_i\right)$. Since each of these vectors lies in the right kernel of $\mathcal{M}_d$, it follows that each $\mathbf{b}_\lambda$ also lies in the right kernel of $\mathcal{M}_d$.

Next, we show that the set of vectors $\mathbf{b}_\lambda$ is linearly independent. Let $\mathbf{a}_{\alpha+1}, \ldots, \mathbf{a}_N$ be a set of linearly independent vectors such that $\mathbf{a}_1, \ldots, \mathbf{a}_\alpha, \mathbf{a}_{\alpha+1}, \ldots, \mathbf{a}_N$ form a basis of $\mathbb{F}_q^N$. Then, the transformation

$$x_i \mapsto \sum_{j=1}^{N} a_{j,i} x_j \quad (i \in [N])$$

induces a ring automorphism of $R = \mathbb{F}_q[\mathbf{x}]$, where $\mathbf{a}_j = (a_{j,1}, \ldots, a_{j,N})^\top$. Moreover, the ring automorphism induces a linear automorphism of the subspace R_d. Then the matrix representation of this linear automorphism with respect to a basis of R_d contains the vectors $\mathbf{b}_\lambda$ as its columns. Therefore, the set of vectors $\mathbf{b}_\lambda$ is linearly independent. □

From the above lemma, the dimension of the vector space spanned by all the $\mathbf{b}_\lambda$ is given by

$$\#\mathrm{Mon}_d(\mathbf{t}) = \binom{\alpha + d - 1}{d}.$$

Therefore, if the subspace I_d of the ideal I generated by $f_1, \ldots, f_M$ has codimension $\binom{\alpha+d-1}{d}$ in R_d, then we can recover the vector space generated by $\mathbf{b}_\lambda$ by computing the right kernel of $\mathcal{M}_d$.

Example 1. Let $f_1, \ldots, f_5$ be homogeneous quadratic polynomials in four variables $\mathbf{x} = (x_1, x_2, x_3, x_4)^\top$ over $\mathbb{F}_4$, where $\mathbb{F}_4 = \mathbb{F}_2[\omega]/\langle \omega^2 + \omega + 1 \rangle$:

$$\begin{aligned}
f_1(\mathbf{x}) &= x_1x_2 + x_1x_3 + (\omega+1)x_1x_4 + (\omega+1)x_3^2 + \omega x_3x_4,\\
f_2(\mathbf{x}) &= x_1^2 + x_1x_2 + x_1x_3 + x_2^2 + x_2x_3 + (\omega+1)x_2x_4 + (\omega+1)x_3^2 + x_3x_4 + \omega x_4^2,\\
f_3(\mathbf{x}) &= \omega x_1^2 + \omega x_1x_2 + (\omega+1)x_1x_3 + \omega x_1x_4 + (\omega+1)x_2^2 + (\omega+1)x_2x_4 + x_3x_4,\\
f_4(\mathbf{x}) &= (\omega+1)x_1^2 + (\omega+1)x_1x_2 + \omega x_1x_4 + x_2x_3 + \omega x_2x_4 + (\omega+1)x_3x_4 + \omega x_4^2,\\
f_5(\mathbf{x}) &= x_1^2 + \omega x_1x_2 + x_1x_3 + (\omega+1)x_1x_4 + x_2^2 + x_2x_3 + x_2x_4 + \omega x_3^2 + \omega x_3x_4.
\end{aligned}$$

We aim to solve the system of equations $f_i(\mathbf{x}) = 0$ with $1 \le i \le 5$. The solution space $\mathcal{A}$ of this system has dimension $\dim \mathcal{A} = 2$, and a basis $\mathbf{a}_1$ and $\mathbf{a}_2$ is given by

$$\begin{aligned}
\mathbf{a}_1 &= (1, 0, 1, \omega)^\top,\\
\mathbf{a}_2 &= (0, 1, 0, 1)^\top.
\end{aligned}$$

To solve the system, we construct the Macaulay matrix of degree 2 as follows:

$$\mathcal{M}_2 = \begin{array}{c} \begin{matrix} x_1^2 & x_1x_2 & x_1x_3 & x_1x_4 & x_2^2 & x_2x_3 & x_2x_4 & x_3^2 & x_3x_4 & x_4^2 \end{matrix} \\ \begin{pmatrix}
0 & 1 & 1 & \omega+1 & 0 & 0 & 0 & \omega+1 & \omega & 0 \\
1 & 1 & 1 & 0 & 1 & 1 & \omega+1 & \omega+1 & 1 & \omega \\
\omega & \omega & \omega+1 & \omega & \omega+1 & 0 & \omega+1 & 0 & 1 & 0 \\
\omega+1 & \omega+1 & 0 & \omega & 0 & 1 & \omega & 0 & \omega+1 & \omega \\
1 & \omega & 1 & \omega+1 & 1 & 1 & 1 & \omega & \omega & 0
\end{pmatrix} \end{array}.$$

For two variables t_1 and t_2, the vector $t_1\mathbf{a}_1 + t_2\mathbf{a}_2$ is given by

$$(t_1, t_2, t_1, \omega t_1 + t_2)^\top.$$

By evaluating $\mathbf{v}_2(\mathbf{x})$ at $t_1\mathbf{a}_1 + t_2\mathbf{a}_2$, we obtain

$$(t_1^2, t_1t_2, t_1^2, \omega t_1^2 + t_1t_2, t_2^2, t_1t_2, \omega t_1t_2 + t_2^2, t_1^2, \omega t_1^2 + t_1t_2, (\omega+1)t_1^2 + t_2^2)^\top,$$

and the vectors corresponding to the coefficients of t_1^2, t_2^2, t_1t_2 are given by

$$\begin{aligned}\mathbf{b}_{t_1^2} &= (1,0,1,\omega,0,0,0,1,\omega,\omega+1)^\top,\\ \mathbf{b}_{t_2^2} &= (0,0,0,0,1,0,1,0,0,1)^\top,\\ \mathbf{b}_{t_1t_2} &= (0,1,0,1,0,1,\omega,0,1,0)^\top.\end{aligned}$$

We can confirm that any vector in the vector space spanned by the above three vectors lies in the right kernel of the Macaulay matrix $\mathcal{M}_2$.

3.2 Macaulay Matrix Using p-truncated Polynomial Ring

In this subsection, we study the Macaulay matrix of degree d constructed from the system (5) using the p-truncated polynomial ring $R^{(p)} = R\Big/\langle x_1^p, \ldots, x_N^p\rangle$ where p is the characteristic of the field $\mathbb{F}_q$.

For the discussion using $R^{(p)}$, we define the following notations:

Definition 1. *1. The* $\mathrm{Mon}_d^{(p)}(\mathbf{x})$ *(resp.* $\mathrm{Mon}_d^{(p)}(\mathbf{t})$*) denotes the set of the p-truncated monomials defined as monomials in* $\mathrm{Mon}_d(\mathbf{x})\setminus\langle x_1^p,\ldots,x_N^p\rangle_d$ *(resp.* $\mathrm{Mon}_d(\mathbf{t})\setminus\langle t_1^p,\ldots,t_\alpha^p\rangle_d$*).*

2. The p-truncated Macaulay matrix $\mathcal{M}_d^{(p)}$ *denotes the submatrix of* $\mathcal{M}_d$ *whose columns correspond to monomials in* $\mathrm{Mon}_d^{(p)}(\mathbf{x})$*.*

3. The p-truncated right kernel vector $\mathbf{b}_\lambda^{(p)}$ *denotes the subvector of* $\mathbf{b}_\lambda$ *consisting of the entries corresponding to monomials in* $\mathrm{Mon}_d^{(p)}(\mathbf{x})$*.*

We then provide the following theorem for the right kernel vector of $\mathcal{M}_d^{(p)}$ obtained from the solution space $\mathcal{A}$.

Theorem 1. *Let* $2 \le d \le \alpha(p-1)$ *be a positive integer. For any monomial* $\lambda \in \mathrm{Mon}_d^{(p)}(\mathbf{t})$*,* $\mathbf{b}_\lambda^{(p)}$ *is a right kernel vector of the p-truncated Macaulay matrix* $\mathcal{M}_d^{(p)}$ *constructed from the system* (5)*.*

Proof. In this proof, we only discuss the cases of $d = p$ and $d > p$, since $R_d \simeq R_d^{(p)}$ if $d < p$.

We first observe the structure of the vectors $\mathbf{b}_\lambda$, which are right kernel vectors of the Macaulay matrix $\mathcal{M}_d$, in the case where $d = p$. When we substitute $\sum_{j=1}^{\alpha} t_j\mathbf{a}_j$ for $\mathbf{x}$ and evaluate the monomial $x_i^p \in \mathrm{Mon}_p(\mathbf{x})$, it is computed as

$$\left(\sum_{j=1}^{\alpha} t_j a_{j,i}\right)^p = \sum_{j=1}^{\alpha} t_j^p a_{j,i}^p,$$

where $a_{j,i}$ denotes the i-th component of the vector $\mathbf{a}_j$. In this expression, the coefficients of all monomials in $\mathrm{Mon}_p^{(p)}(\mathbf{t})$ are equal to zero. This implies that, for the vector $\mathbf{b}_\lambda$ with $\lambda \in \mathrm{Mon}_p^{(p)}(\mathbf{t})$, the entries at the positions corresponding to $x_1^p, \ldots, x_N^p$ are identically zero. From this discussion, for any monomial

$\lambda \in \mathrm{Mon}_p^{(p)}(\mathbf{t})$, the corresponding vector $\mathbf{b}_\lambda^{(p)}$ lies in the right kernel of $\mathcal{M}_p^{(p)}$ as follows:

$$\mathcal{M}_p \mathbf{b}_\lambda = \left[\begin{array}{c|c} \star & \mathcal{M}_p^{(p)} \end{array}\right] \begin{pmatrix} 0 \\ \mathbf{b}_\lambda^{(p)} \end{pmatrix} = 0 \;\Rightarrow\; \mathcal{M}_p^{(p)}\, \mathbf{b}_\lambda^{(p)} = 0\ .$$

We further consider the case where $d > p$. When we substitute $\sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ into the monomials in $\langle x_1^p, \ldots, x_N^p \rangle_d$, the coefficients of all monomials in $\mathrm{Mon}_d^{(p)}(\mathbf{t})$ are equal to zero. Consequently, for the vector $\mathbf{b}_\lambda$ with $\lambda \in \mathrm{Mon}_d^{(p)}(\mathbf{t})$, the entries at the positions corresponding to $\langle x_1^p, \ldots, x_N^p \rangle_d$ are equal to zero. By the same argument as in the case $d = p$, for any monomial $\lambda \in \mathrm{Mon}_d^{(p)}(\mathbf{t})$, the corresponding vector $\mathbf{b}_\lambda^{(p)}$ lies in the right kernel of $\mathcal{M}_d^{(p)}$. □

Note that if $d > \alpha(p-1)$, $\mathcal{M}_d^{(p)}$ does not have the right kernel vector of the form $\mathbf{b}_\lambda^{(p)}$, since $\mathrm{Mon}_d^{(p)}(\mathbf{t}) = \emptyset$.

Example 2. We consider solving the system $f_i(\mathbf{x}) = 0$ with $1 \le i \le 5$ given in Example 1 using the structure of $\mathbb{F}_4[\mathbf{x}]/\langle x_1^2, \ldots, x_4^2 \rangle$. The p-truncated Macaulay matrix $\mathcal{M}_2^{(2)}$ of degree 2 is given by

$$\begin{array}{c} \begin{array}{cccccc} x_1x_2 & x_1x_3 & x_1x_4 & x_2x_3 & x_2x_4 & x_3x_4 \end{array} \\ \begin{pmatrix} 1 & 1 & \omega+1 & 0 & 0 & \omega \\ 1 & 1 & 0 & 1 & \omega+1 & 1 \\ \omega & \omega+1 & \omega & 0 & \omega+1 & 1 \\ \omega+1 & 0 & \omega & 1 & \omega & \omega+1 \\ \omega & 1 & \omega+1 & 1 & 1 & \omega \end{pmatrix} \end{array},$$

which is obtained by removing the columns corresponding to x_i^2 with $1 \le i \le 4$ from the matrix in Example 1. For $\mathbf{b}_{t_1t_2}$ given in Example 1, one can confirm that all entries corresponding to x_i^2 with $1 \le i \le 4$ are equal to zero, as shown below:

$$\begin{array}{c} \begin{array}{cccccccccc} x_1^2 & x_1x_2 & x_1x_3 & x_1x_4 & x_2^2 & x_2x_3 & x_2x_4 & x_3^2 & x_3x_4 & x_4^2 \end{array}^\top \\ \begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 1 & \omega & 0 & 1 & 0 \end{pmatrix} \end{array}.$$

Thus, when we obtain $\mathbf{b}_{t_1t_2}^{(p)}$ given as

$$\mathbf{b}_{t_1t_2}^{(p)} = (1, 0, 1, 1, \omega, 1)^\top,$$

by removing the elements corresponding to x_i^2 with $1 \le i \le 4$, this vector is a right kernel vector of the p-truncated Macaulay matrix $\mathcal{M}_2^{(2)}$.

3.3 Macaulay Matrix Using p^ℓ-truncated Polynomial Ring

In this subsection, we study the Macaulay matrix of degree d constructed from the system (5) using the p^ℓ-truncated polynomial ring $R^{(p^\ell)} = R\Big/\left\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}\right\rangle$ with $\ell \in [e]$ by generalizing the results in Subsection 3.2.

As in Subsection 3.2, we define the following notations:

Definition 2. *(i) The* $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{x})$ *(resp.* $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$*) denotes the set of the* p^ℓ*-truncated monomials defined as monomials in* $\mathrm{Mon}_d(\mathbf{x}) \setminus \left\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}\right\rangle_d$ *(resp.* $\mathrm{Mon}_d(\mathbf{t}) \setminus \left\langle t_1^{p^\ell}, \ldots, t_\alpha^{p^\ell}\right\rangle_d$*).*

(ii) The p^ℓ*-truncated Macaulay matrix* $\mathcal{M}_d^{(p^\ell)}$ *denotes the submatrix of* $\mathcal{M}_d$ *whose columns correspond to monomials in* $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{x})$*.*

(iii) The p^ℓ*-truncated right kernel vector* $\mathbf{b}_\lambda^{(p^\ell)}$ *denotes the subvector of* $\mathbf{b}_\lambda$ *consisting of the entries corresponding to monomials in* $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{x})$*.*

We then provide the following theorem for the right kernel vector of $\mathcal{M}_d^{(p^\ell)}$ obtained from the solution space $\mathcal{A}$.

Theorem 2. *Let* $2 \le d \le \alpha(p^\ell - 1)$ *be a positive integer. For any monomial* $\lambda \in \mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$*,* $\mathbf{b}_\lambda^{(p^\ell)}$ *is a right kernel vector of the* p^ℓ*-truncated Macaulay matrix* $\mathcal{M}_d^{(p^\ell)}$ *constructed from the system* (5)*.*

Proof. In the case where $d = p^\ell$, we consider substituting $\mathbf{x} = \sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ into the monomial $x_i^{p^\ell}$ for $i \in [N]$. Then, this substitution yields

$$\left(\sum_{j=1}^{\alpha} t_j a_{j,i}\right)^{p^\ell} = \sum_{j=1}^{\alpha} t_j^{p^\ell} a_{j,i}^{p^\ell},$$

since the characteristic of $\mathbb{F}_q$ is p. Similarly, we consider the case where $d > p^\ell$. When we substitute $\mathbf{x} = \sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ into the monomials in $\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}\rangle_d$ and regard the result as a polynomial in $\mathbf{t}$, the coefficients of all monomials in $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$ are equal to zero. This implies that, for the vector $\mathbf{b}_\lambda$ corresponding to any monomial $\lambda \in \mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$, the entries at the positions corresponding to $R^{(p^\ell)}$ are identically zero. Then, by the same argument as in the proof of Theorem 1, for any monomial $\lambda \in \mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$, the corresponding vector $\mathbf{b}_\lambda^{(p^\ell)}$ lies in the right kernel of $\mathcal{M}_d^{(p^\ell)}$. □

As in the case of the p-truncated polynomial ring, there exist such kernel vectors $\mathbf{b}_\lambda^{(p^\ell)}$ of $\mathcal{M}_d^{(p^\ell)}$ only under the condition of $d \le \alpha(p^\ell - 1)$.

Remark 1 (Case of the k-truncated polynomial ring for general k). In this remark, we consider solving the system (5) using the Macaulay matrix constructed using the structure of the k-truncated polynomial ring $R^{(k)} = R/\langle x_1^k, \ldots, x_N^k \rangle$ for a general integer k.

For the case where $d = k$, when we substitute $\mathbf{x} = \sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ into the monomial x_i^k for $i \in [N]$ and regard the result as a polynomial in $\mathbf{t}$, the coefficients of the monomials in

$$B = \left\{ \prod_{j=1}^{\alpha} t_j^{i_j} \;\middle|\; \sum_{j=1}^{\alpha} i_j = k,\ p \mid \frac{k!}{i_1! \cdots i_\alpha!} \right\},$$

are equal to zero. This indicates that, for any monomial $\lambda \in B$, the corresponding vector $\mathbf{b}_\lambda^{(k)}$ is a right kernel vector of $\mathcal{M}_d^{(k)}$ when $d = k$.

However, this structure is not preserved for general degrees $d > k$. Indeed, when we substitute $\mathbf{x} = \sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ into the monomials in $\langle x_1^k, \ldots, x_N^k \rangle_d$, the coefficients of monomials in $\langle B \rangle_d$ are not necessarily equal to zero, unlike the case of the p^ℓ-truncated polynomial ring. For example, in the case where $p = 2$, $k = 6$, and $\alpha = 3$, we list, for each degree $d \geq k$, the monomials λ in $\mathbf{t} = (t_1, t_2, t_3)^\top$ such that the corresponding vector $\mathbf{b}_\lambda^{(k)}$ is a right kernel vector of $\mathcal{M}_d^{(k)}$ as follows:

$d = 6$	$t_1^5t_2, t_1^5t_3, t_1t_2^5, t_1t_3^5, t_2^5t_3, t_2t_3^5, t_1^4t_2t_3, t_1t_2^4t_3, t_1t_2t_3^4, t_1^3t_2^3$ $t_1^3t_3^3, t_2^3t_3^3, t_1^3t_2^2t_3, t_1^3t_2t_3^2, t_1^2t_2^3t_3, t_1^2t_2t_3^3, t_1t_2^3t_3^2, t_1t_2^2t_3^3, t_1^2t_2^2t_3^2$
$d = 7$	$t_1^5t_2t_3, t_1t_2^5t_3, t_1t_2t_3^5, t_1^3t_2^2t_3^2, t_1^2t_2^3t_3^2, t_1^2t_2^2t_3^3, t_1^3t_2^3t_3, t_1^3t_2t_3^3, t_1t_2^3t_3^3$
$d = 8$	$t_1^3t_2^3t_3^2, t_1^3t_2^2t_3^3, t_1^2t_2^3t_3^3$
$d = 9$	$t_1^3t_2^3t_3^3$
$d = 10$	-

Indeed, this example shows that, for larger degrees d, we cannot obtain a kernel vector derived from the solutions over the k-truncated polynomial ring for a general integer k.

The discussion above indicates that, for a general integer k, we can exploit the structure of the k-truncated polynomial ring only for a limited range of degrees d. Thus, in the rest of this paper, we focus on the p^ℓ-truncated polynomial ring, in which kernel vectors derived from the solutions can be obtained for any degree d satisfying $d \leq \alpha(p^\ell - 1)$.

Remark 2 (Relation to the symmetric-algebra attack [10]). This remark discusses a relation between the method based on the p-reduced symmetric algebra proposed by Jin et al. [10] and our proposed framework.

We briefly recall the description of the method in [10]. They treat a given degree-d polynomial system as a system of p-reduced symmetric d-forms. Let $V = \mathbb{F}_q^N$, and consider the symmetric algebra $\mathrm{Sym}(V)$ defined by

$$\mathrm{Sym}(V) = \mathbb{F}_q \oplus V \oplus V^{\otimes 2} \oplus \cdots / \langle \mathbf{x} \otimes \mathbf{y} - \mathbf{y} \otimes \mathbf{x} \,|\, \mathbf{x}, \mathbf{y} \in V \rangle .$$

Then, the p-reduced symmetric algebra $\overline{\mathrm{Sym}}(V)$ is defined as

$$\overline{\mathrm{Sym}}(V) = \mathrm{Sym}(V) / \left\langle \mathbf{x}^{\otimes p} \,\middle|\, \mathbf{x} \in V \right\rangle.$$

Theorem 10 in [10] provides an isomorphism map from the space of symmetric d-forms to $\overline{\mathrm{Sym}}(V)_d$, constructed as follows. The map sends a symmetric d-form f to the element

$$\sum_{1 \le i_1 \le \cdots \le i_d \le N} \frac{f(\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d})}{m_1! \cdots m_k!} \mathbf{e}_{i_1} \odot \cdots \odot \mathbf{e}_{i_d}, \tag{6}$$

where $\{m_1, \ldots, m_k\}$ are the repetition counts of the distinct indices in $(i_1, \ldots, i_d)$. Note that if $m_j \ge p$ for some j, we regard the coefficient $f(\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d})/m_1! \cdots m_k!$ as 0, since $f(\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d}) = 0$ for such cases.

We then explain that the symmetric-algebra attack coincides with the proposed framework when $\ell = 1$. The p-reduced symmetric algebra $\overline{\mathrm{Sym}}(V)$ can be identified with the p-truncated polynomial ring $R^{(p)}$. Under this identification, the p-reduced symmetric monomial basis corresponds to the p-truncated monomial basis of $R^{(p)}$. Therefore, with respect to these corresponding bases, the Macaulay matrix generated in the symmetric-algebra attack coincides with the p-truncated Macaulay matrix $M_d^{(p)}$ used in the proposed framework with $\ell = 1$.

We next consider a generalization of the isomorphism (6) to the p^ℓ-reduced symmetric algebra for the case of $\ell \ge 2$. In this setting there exist indices $i_1, \ldots, i_d$ such that $f(\mathbf{e}_{i_1}, \ldots, \mathbf{e}_{i_d}) \neq 0$ while some repetition count satisfies $m_j \ge p$. This suggests that the method based on the symmetric algebra proposed by Jin et al. may not admit a straightforward generalization to the p^ℓ-reduced symmetric algebra.

In conclusion, the proposed framework with $\ell = 1$ is equivalent to the symmetric-algebra attack, while the generalization to the p^ℓ-truncated polynomial ring with $\ell \ge 2$ relies on our polynomial-ring-based description.

4 Proposed Framework and its Application to UOV

This section provides a framework for solving the system (5) by using the p^ℓ-truncated Macaulay matrix provided in Section 3. We also provide the estimation of the complexity of the proposed framework and consider applying it to recover the secret subspace of UOV.

4.1 Proposed Framework Using p^ℓ-truncated Polynomial Ring

We first provide the framework for the XL algorithm using the p^ℓ-truncated polynomial ring $R^{(p^\ell)} = R/\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell} \rangle$. We consider solving the system (5) in N variables and M equations, whose solution space contains a vector space $\mathcal{A}$ with $\dim \mathcal{A} = \alpha$. In the proposed framework, we choose the parameter ℓ, which

determines the truncation level of the polynomial ring, and the degree D of the Macaulay matrix with $2 \leq D \leq \alpha(p^\ell - 1)$. Given the system (5), the proposed framework proceeds as follows:

1. Construct the Macaulay matrix $\mathcal{M}_D^{(p^\ell)}$ from the given system.
2. Compute a right kernel vector of $\mathcal{M}_D^{(p^\ell)}$.
3. Recover a solution vector in $\mathcal{A}$ from the right kernel vector of $\mathcal{M}_D^{(p^\ell)}$.

From Theorem 2, we can obtain $\mathbf{b}_\lambda^{(p^\ell)}$ with $\lambda \in \mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$ by computing a right kernel space of $\mathcal{M}_D^{(p^\ell)}$ if D satisfies the following condition. The codimension of I_D, where I is the ideal generated by $f_1, \ldots, f_M$, is equal to $\#\mathrm{Mon}_D^{(p^\ell)}(\mathbf{t})$. In Subsection 4.2, we explain how to estimate such a suitable degree D for finding a solution.

We then explain how to recover an element of the original solution space $\mathcal{A}$ from the right kernel space of $\mathcal{M}_D^{(p^\ell)}$. We assume that a basis $\mathbf{a}_1, \ldots, \mathbf{a}_\alpha$ of $\mathcal{A}$ is of the form

$$\mathbf{a}_i = (\overbrace{0, \ldots, 0, \underset{i}{1}, 0, \ldots, 0}^{\alpha}, a_{i,\alpha+1}, \ldots, a_{i,N})^\top \quad (i \in [\alpha]). \tag{7}$$

Then, an element $\sum_{i=1}^{\alpha} t_i \mathbf{a}_i$ of $\mathcal{A}$ has the form

$$\left(t_1, \ldots, t_\alpha, \sum_{i=1}^{\alpha} t_i a_{i,\alpha+1}, \ldots, \sum_{i=1}^{\alpha} t_i a_{i,N} \right)^\top. \tag{8}$$

Let $\mathbf{b}_\lambda^{(p^\ell)}$ be the right kernel vector of $\mathcal{M}_D^{(p^\ell)}$ corresponding to a monomial $\lambda = \prod_{j=1}^{\alpha} t_j^{i_j} \in \mathrm{Mon}_D^{(p^\ell)}(\mathbf{t})$. For this vector $\mathbf{b}_\lambda^{(p^\ell)}$, the entry corresponding to the monomial $\prod_{j=1}^{\alpha} x_j^{i_j}$ equals 1, while the entries corresponding to other monomials only in $x_1, \ldots, x_\alpha$ equal 0, due to the form of (8). Thus, we can obtain such a vector $\mathbf{b}_\lambda^{(p^\ell)}$ by computing a right kernel vector of the submatrix of $\mathcal{M}_D^{(p^\ell)}$ obtained by removing the columns corresponding to monomials only in $x_1, \ldots, x_\alpha$, except for $\prod_{j=1}^{\alpha} x_j^{i_j}$. Then, from the kernel vector $\mathbf{b}_\lambda^{(p^\ell)}$ with $\lambda = \prod_{j=1}^{\alpha} t_j^{i_j}$, we can recover $\mathbf{a}_{j'}$ such that $i_{j'} \neq 0$ as follows. For $\alpha + 1 \leq i \leq N$, the i-th entry of $\mathbf{a}_{j'}$ can be obtained as the entry of $\mathbf{b}_\lambda$ at the position corresponding to the monomial $x_1^{i_1} \cdots x_{j'}^{i_{j'}-1} \cdots x_\alpha^{i_\alpha} \cdot x_i$, by the form of (8). Note that the first α entries of $\mathbf{a}_{j'}$ are determined by the form in (7).

We finally consider applying the optimization technique used in [10, 20]. Before solving the system using the proposed p^ℓ-truncated framework, we apply the following transformation. For an integer $0 \leq k < \alpha$, we transform the given system into a reduced system in $n-k$ variables whose solution space $\mathcal{A}'$ satisfies $\dim \mathcal{A}' = \alpha - k$. This transformation can be performed by simply substituting 0 for k variables. We then recover a secret vector for the transformed system with $n - k$ variables using the proposed framework. From the obtained solution, we can recover a solution to the original system by inserting zeros into the corresponding k positions.

4.2 Complexity Estimation

In this subsection, we estimate the complexity of the proposed p^ℓ-truncated framework. Recall that the number of monomials of degree d in the p^ℓ-truncated polynomial ring in N variables is given by

$$\sum_{i=0}^{\min\{N,\lfloor d/p^\ell \rfloor\}} (-1)^i \binom{N}{i} \binom{d - ip^\ell + N - 1}{N-1},$$

and we denote this number by $T(N, p^\ell, d)$.

We first estimate the solving degree D of the Macaulay matrix required to find a solution of the system (5). From the discussion in Section 3, the dimension of the right kernel space of $\mathcal{M}_d^{(p^\ell)}$ derived from the solutions of the system is equal to the number of monomials in $\mathrm{Mon}_d^{(p^\ell)}(\mathbf{t})$, that is, $T(\alpha, p^\ell, d)$. Thus, for the efficiency of the attacks, we choose the solving degree D as

$$D = \min\left\{ 2 \le d \le \alpha(p^\ell - 1) \,\middle|\, \dim\ker \mathcal{M}_d^{(p^\ell)} \le T(\alpha, p^\ell, d) \right\}.$$

By extending the results of Theorem 15 and Conjecture 16 in [10], we provide a conjecture to estimate $\dim\ker \mathcal{M}_d^{(p^\ell)}$. The rows of $\mathcal{M}_d^{(p^\ell)}$ span the degree-d component of the ideal generated by $f_1, \ldots, f_M$ in the p^ℓ-truncated polynomial ring. Hence, the dimension of the right kernel of $\mathcal{M}_d^{(p^\ell)}$ is equal to the dimension of the degree-d homogeneous component of the quotient algebra

$$\mathbb{F}_q[\mathbf{x}]/\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}, f_1, \ldots, f_M \rangle.$$

We estimate the dimension of this quotient algebra by using the Hilbert series, which is a generating function whose coefficient of z^d gives the dimension of the degree-d homogeneous component of $\mathbb{F}_q[\mathbf{x}]/\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}, f_1, \ldots, f_M \rangle$.

Conjecture 1. The given quadratic polynomials $f_1, \ldots, f_M$ of (5) are semi-regular on the p^ℓ-truncated polynomial ring $R^{(p^\ell)}$ with very high probability. Then, the Hilbert series of $\mathbb{F}_q[\mathbf{x}]/\langle x_1^{p^\ell}, \ldots, x_N^{p^\ell}, f_1, \ldots, f_M \rangle$ is given by

$$\left[\frac{(1 + z + \cdots + z^{p^\ell - 1})^N}{(1 + z^2 + \cdots + z^{2(p^\ell - 1)})^M} \right]. \tag{9}$$

This conjecture is heuristic, and we examine its validity experimentally in Subsection 4.4. Assuming the validity of the above conjecture, we can estimate the solving degree D as the minimum integer d satisfying

$$\left[\frac{(1 + z + \cdots + z^{p^\ell - 1})^N}{(1 + z^2 + \cdots + z^{2(p^\ell - 1)})^M} \right]_d \le T(\alpha, p^\ell, d),$$

where $[\cdot]_d$ denotes the coefficient of z^d in the given series.

We then estimate the number of operations over $\mathbb{F}_q$ required to solve the system (5) using the p^ℓ-truncated polynomial ring. In the proposed algorithm, the step of finding kernel vectors from $\mathcal{M}_D^{(p^\ell)}$ dominates the overall time complexity. As explained in Subsection 4.1, we work only with the submatrix of $\mathcal{M}_D^{(p^\ell)}$ obtained by removing all columns corresponding to monomials involving only the first α variables $x_1, \ldots, x_\alpha$, except for one column. Then, the number of columns of this submatrix is given by

$$T(N, p^\ell, D) - T(\alpha, p^\ell, D) + 1.$$

For the complexity estimation, we consider a square matrix obtained by removing redundant rows. We then apply the Wiedemann algorithm [23] to compute a kernel vector of the resulting matrix. As a result, the number of operations required for the attacks is estimated as

$$3\binom{N+1}{2}\left(T(N, p^\ell, D) - T(\alpha, p^\ell, D) + 1\right)^2,$$

where the factor $\binom{N+1}{2}$ is derived as an upper bound on the number of nonzero entries in each row.

Remark 3. As a generalization of the discussion in Subsection 3.4 of Ran [20], we claim that if

$$D \leq (N - 2M)(p^\ell - 1)$$

holds, then the complexity estimations are not expected to be accurate.

Let $y_1, \ldots, y_{N-\alpha}$ be linearly independent linear forms in $\mathbf{x}$ whose kernels contain the solution space $\mathcal{A}$ of $f_1, \ldots, f_M$, and $\langle y_1, \ldots, y_{N-\alpha}\rangle_d^{(p^\ell)}$ be the image of $\langle y_1, \ldots, y_{N-\alpha}\rangle_d$ in $R_d^{(p^\ell)}$. We note that in our analysis the complexity estimations are derived under the expectation that the equality

$$\langle y_1, \ldots, y_{N-\alpha}\rangle_D^{(p^\ell)} = \langle f_1, \ldots, f_M\rangle_D^{(p^\ell)} \tag{10}$$

holds for $d = D$. We explain that this equality may not hold under the above condition. We consider the following non-degenerate bilinear form:

$$B_D : R_D^{(p^\ell)} \times R_{N(p^\ell-1)-D}^{(p^\ell)} \ni (g, h) \mapsto \frac{gh}{y_1^{p^\ell-1}\cdots y_N^{p^\ell-1}} \in \mathbb{F}_q.$$

Here, $y_{N-\alpha+1}, \ldots, y_N$ are chosen so that $y_1, \ldots, y_N$ form a basis, and the quotient is the coefficient of $y_1^{p^\ell-1}\cdots y_N^{p^\ell-1}$ in gh. The annihilator of $\langle y_1, \ldots, y_{N-\alpha}\rangle_D^{(p^\ell)}$ with respect to B_D is given by

$$Y_D := \left\{ y_1^{p^\ell-1}\cdots y_{N-\alpha}^{p^\ell-1} \cdot u \;\middle|\; u \in \mathbb{F}_q[y_{N-\alpha+1}, \ldots, y_N]_{\alpha(p^\ell-1)-D}^{(p^\ell)} \right\}.$$

If $D \leq (N - 2M)(p^\ell - 1)$ holds, for any $r \in R_{(N-2M)(p^\ell-1)-D}^{(p^\ell)}$ the product $r \cdot f_1^{p^\ell-1}\cdots f_M^{p^\ell-1}$ is an element of the annihilator of $\langle f_1, \ldots, f_M\rangle_D^{(p^\ell)}$ but may

not belong to Y_D. In this case, the equality (10) does not hold, since the two annihilators differ. This makes the above complexity estimations invalid.

Therefore, in the case $D \leq (N-2M)(p^\ell - 1)$, the complexity estimations are not expected to be accurate. This case corresponds to the lifting approach of SNOVA (see Remark 4).

4.3 Application to Key Recovery Attacks on UOV

From the description in Subsection 2.2, the reconciliation and intersection attacks on UOV can be regarded as solving a polynomial system of the form (5) whose solution space contains a vector space. Thus, we consider applying these key recovery attacks on UOV using the proposed p^ℓ-truncated framework.

We provide complexity estimations for the reconciliation and intersection attacks using the proposed p^ℓ-truncated framework, applied to the UOV system with parameters (q, v, o, m), where $q = p^e$.

In the following estimations, we assume that the systems arising in the reconciliation and intersection attacks satisfy the semi-regularity assumption, that is, their behaviors follow the Hilbert series (9) in Conjecture 1.

Reconciliation attack. The reconciliation attack solves a system with $n = v + o$ variables and m equations, where the dimension of the solution space is o. We choose two parameters $1 \leq \ell \leq e$ and $0 \leq k < o$ for the attack. Under the semi-regularity assumption, we take the solving degree D as the minimum integer $2 \leq d \leq (o-k)(p^\ell - 1)$ satisfying

$$\left[\frac{(1 + z + \cdots + z^{p^\ell - 1})^{n-k}}{(1 + z^2 + \cdots + z^{2(p^\ell - 1)})^m}\right]_d \leq T(o-k, p^\ell, d). \tag{11}$$

Then, the complexity is estimated as

$$3\binom{n-k+1}{2}\left(T(n-k, p^\ell, D) - T(o-k, p^\ell, D) + 1\right)^2. \tag{12}$$

Intersection attack. We first deal with the case where $v < 2o$. In this case, the intersection attack solves a system with n variables and $3m - 2$ linearly independent equations, where the dimension of the solution space is $2o - v$. We choose two parameters $1 \leq \ell \leq e$ and $0 \leq k < 2o - v$ for the attack. Under the semi-regularity assumption, we take the solving degree D as the minimum integer $2 \leq d \leq (2o - v - k)(p^\ell - 1)$ satisfying

$$\left[\frac{(1 + z + \cdots + z^{p^\ell - 1})^{n-k}}{(1 + z^2 + \cdots + z^{2(p^\ell - 1)})^{3m-2}}\right]_d \leq T(2o - v - k, p^\ell, d). \tag{13}$$

Then, the complexity is estimated as

$$3\binom{n-k+1}{2}\left(T(n-k, p^\ell, D) - T(2o - v - k, p^\ell, D) + 1\right)^2. \tag{14}$$

Table 1: The complexity of the reconciliation attack using the p^ℓ-truncated polynomial ring with parameters ℓ and k, with the solving degree D shown in parentheses, on UOV with parameters $(q, v, o, m) = (16, 40, 20, 30)$

k	ℓ			
	1	2	3	4
0	100 (14)	178 (34)	262 (71)	356 (143)
1	96 (13)	174 (33)	256 (69)	348 (139)
2	95 (13)	170 (32)	250 (67)	340 (134)
3	95 (13)	166 (31)	245 (65)	331 (129)
4	94 (13)	162 (30)	239 (63)	324 (125)
5	90 (12)	161 (30)	231 (60)	315 (120)
6	89 (12)	156 (29)	226 (58)	307 (116)
7	88 (12)	152 (28)	220 (56)	299 (111)
8	87 (12)	148 (27)	214 (54)	291 (107)
9	83 (11)	144 (26)	208 (52)	282 (102)
10	-	140 (25)	202 (50)	274 (98)
11	-	139 (25)	198 (49)	266 (94)
12	-	134 (24)	192 (47)	258 (90)
13	-	-	186 (45)	249 (85)
14	-	-	-	241 (81)
15	-	-	-	-

On the other hand, in the case where $2o \leq v$, the system has a solution with probability $q^{-(v-2o+1)}$ [1]. To simplify the complexity estimation, we ignore cases where the intersection has dimension larger than 1, since such cases arise with smaller probability and do not improve the overall attack cost. Then, we estimate the solving degree D as the minimum integer $2 \leq d \leq p^\ell - 1$ satisfying

$$\left[\frac{(1 + z + \cdots + z^{p^\ell - 1})^n}{(1 + z^2 + \cdots + z^{2(p^\ell - 1)})^{3m-2}} \right]_d \leq T(1, p^\ell, d) = 1, \tag{15}$$

and the complexity is estimated as

$$3q^{v-2o+1} \binom{n+1}{2} T(n, p^\ell, D)^2. \tag{16}$$

Note that, in such cases where the dimension of the solution space equals 1, our proposed framework is equivalent to the standard XL algorithm. This is because we have to choose D such that $D \leq p^\ell - 1$ when $\alpha = 1$, and hence $R_D^{(p^\ell)}$ is isomorphic to R_D.

Finally, we discuss how the parameters ℓ and k affect the efficiency of the attacks using the proposed p^ℓ-truncated framework. We observe that the attacks are generally more efficient over smaller polynomial rings, especially in the case $\ell = 1$. Moreover, a larger k tends to make the attacks more efficient, since it decreases the number of variables to be solved. By contrast, increasing k reduces the gap between the solving degree D and its upper bound $(\alpha - k)(p^\ell - 1)$. Thus,

taking k too large may prevent the selection of a feasible degree D. From this discussion, it is reasonable to select the largest k for which a feasible D exists. For example, we compute the complexity of the reconciliation attack using the proposed framework against UOV with parameters $(q, v, o, m) = (16, 40, 20, 30)$ in Table 1. Each block shows the base-2 logarithm of the number of operations over $\mathbb{F}_q$, together with the solving degree D. Blank cells indicate that there is no feasible degree d in the range $2 \leq d \leq (o-k)(p^\ell - 1)$. Indeed, this table shows that larger k and smaller ℓ make the attacks more efficient, whereas larger ℓ enlarges the range of feasible values of k.

4.4 Validity of Semi-regularity Assumption

This subsection discusses the validity of Conjecture 1 used in our complexity estimation.

We examine the validity of the semi-regularity assumption in Conjecture 1 through experiments for specific parameters in Table 2. These experiments show that the theoretical predictions remain accurate in several cases.

However, we also observed discrepancies between the theoretical predictions and the experimental ranks for some parameter sets. For example, for the reconciliation attack with parameters $(p, e, v, o, m) = (5, 1, 3, 3, 3)$ and $(\ell, k, D) = (1, 0, 9)$, the theoretical and experimental ranks coincide up to degree 7, but differ at degree 8, as follows:

$\mathrm{rank}(\mathcal{M}_d)$	degree d							
	2	3	4	5	6	7	8	9
theoretical	3	18	60	150	316	576	921	1236
experimental	3	18	60	150	316	576	918	1236

Thus, while the estimation appears to capture the overall behavior well, these experiments indicate that it is not exact in all cases.

We emphasize that, to the best of our knowledge, the present estimation is the best available estimate for this setting, as it is based on a natural extension of the estimation method in [10]. In our experiments on random MQ maps, we did not observe any such discrepancy, suggesting that it may be caused by the specific structure of UOV maps. We regard developing more accurate rank estimates that take into account the specific structure of UOV maps as an important open problem for the security evaluation of UOV.

5 Complexity Evaluation on UOV and its Variants

In the NIST PQC standardization process for additional digital signatures, UOV and its variants, MAYO [2], QR-UOV [7], and SNOVA [22], have been selected

Table 2: Comparison between the theoretical estimate from (9) and the experimental rank of the Macaulay matrix $\mathcal{M}_d^{(p^\ell)}$ at degree d, constructed using the p^ℓ-truncated polynomial ring in the reconciliation and intersection attacks against UOV with parameters (q, v, o, m).

(i) reconciliation attack

(q,v,o,m)	(ℓ,k)	rank($\mathcal{M}_d$)	degree d					
			2	3	4	5	6	7
(16,10,7,7)	(1,0)	theoretical	7	119	924	4284	12369	
		experimental	7	119	924	4284	12369	
(16,10,7,7)	(1,1)	theoretical	7	112	812	3472	8007	
		experimental	7	112	812	3472	8007	
(16,8,6,6)	(1,0)	theoretical	6	84	525	1890	3002	
		experimental	6	84	525	1890	3002	
(16,8,6,6)	(1,1)	theoretical	6	78	447	1286		
		experimental	6	78	447	1286		
(16,5,5,5)	(2,2)	theoretical	5	40	170	520	1260	2466
		experimental	5	40	170	520	1260	2466
(16,5,5,5)	(2,3)	theoretical	5	35	130	350	727	
		experimental	5	35	130	350	727	

(ii) intersection attack

(q,v,o,m)	(ℓ,k)	rank($\mathcal{M}_d$)	degree d					
			2	3	4	5	6	7
(16,13,11,11)	(1,4)	theoretical	31	620	4840			
		experimental	31	620	4840			
(16,13,11,11)	(1,5)	theoretical	31	589	3875			
		experimental	31	589	3875			
(16,12,10,10)	(1,3)	theoretical	28	532	3871			
		experimental	28	532	3871			
(16,12,10,10)	(1,4)	theoretical	28	504	3059			
		experimental	28	504	3059			
(16,11,9,9)	(2,4)	theoretical	25	400	3100	15236		
		experimental	25	400	3100	15236		
(16,11,9,9)	(2,5)	theoretical	25	375	2700	11401		
		experimental	25	375	2700	11401		

Table 3: The complexity of the best known attacks in the specification document [3], the symmetric-algebra attack by Jin et al. [10], and the reconciliation and intersection attacks using our p^ℓ-truncated framework, on the parameters of UOV in [3]. The lowest complexity is highlighted in bold.

parameter	(q, n, m)	[3]	[10]	our reconciliation		our intersection	
				complexity	(ℓ, k, D)	complexity	(ℓ, k, D)
`uov-Ip`	$(256, 112, 44)$	145	140	141	$(1, 26, 18)$	**128**	$(1, 7, 13)$
`uov-Is`	$(16, 160, 64)$	**143**	175	176	$(1, 42, 22)$	159	$(1, 16, 16)$
`uov-III`	$(256, 184, 72)$	218	206	208	$(1, 46, 26)$	**182**	$(1, 14, 18)$
`uov-V`	$(256, 244, 96)$	278	257	258	$(1, 64, 32)$	**223**	$(1, 22, 22)$

as third-round candidates. This section evaluates the complexity of the key recovery attacks using the proposed p^ℓ-truncated framework against the proposed parameters of these schemes.

In the following tables, we evaluate the complexity of the reconciliation and intersection attacks using the proposed p^ℓ-truncated framework. In the tables, the attack complexity is expressed as $\log_2$ of the number of required gates. This value is computed by multiplying $2(\log_2 q)^2 + \log_2 q$ by the number of operations over $\mathbb{F}_q$ obtained using (12), (14), and (16). Recall that the proposed parameters are designed to meet the security levels I, III, and V, where a classical adversary is expected to require more than 2^{143}, 2^{207}, and 2^{272} classical gates to break the scheme, respectively [16]. The tables also show the optimal parameters ℓ and k that minimize the complexity, as well as the degree D derived from the parameters using the conditions in (11), (13), and (15). We also list the best complexity among the known attacks in the specification documents and the complexity of the symmetric-algebra attack [10]. The lowest complexity among these attacks is highlighted in bold in the table for each parameter set.

UOV. In Table 3, we evaluate the complexity of the reconciliation and intersection attacks using the proposed p^ℓ-truncated framework for the four parameters of UOV, `uov-Ip`, `uov-Is`, `uov-III`, and `uov-V` [3]. Recall that for these parameters, the number of oil variables is set to be m. From the results in Table 3, we observe that the complexities of the reconciliation attack are consistent with those of the symmetric-algebra attack [10]. Although the complexity values are slightly different between these two cases, this difference appears to come from the complexity estimation rather than from an algorithmic difference. Moreover, the table shows that, under the proposed framework, the intersection attack is more efficient than the reconciliation attack for all the proposed parameters. In particular, for `uov-Ip`, `uov-III`, and `uov-V`, the complexity of the intersection attack using the proposed framework is lower than the claimed security level by 10–50 bits.

MAYO. In Table 4, we evaluate the complexity of the attacks using the proposed framework against the MAYO parameters presented in [2]. For this scheme,

Table 4: The complexity of the best known attacks in the specification document [2], the symmetric-algebra attack by Jin et al. [10], and the reconciliation and intersection attacks using our p^ℓ-truncated framework, on the parameters of MAYO in [2]. The lowest complexity is highlighted in bold.

parameter	(q, n, m, o)	[2]	[10]	our reconciliation		our intersection	
				complexity	(ℓ, k, D)	complexity	(ℓ, k, D)
MAYO_1	$(16, 86, 78, 8)$	**156**	-	238	$(3, 2, 41)$	351	$(4, 0, 9)$
MAYO_2	$(16, 81, 64, 17)$	141	**112**	113	$(1, 4, 13)$	227	$(4, 0, 10)$
MAYO_3	$(16, 118, 108, 10)$	**229**	-	314	$(3, 2, 53)$	478	$(4, 0, 11)$
MAYO_5	$(16, 154, 142, 12)$	**298**	-	397	$(3, 2, 66)$	629	$(4, 0, 14)$

four parameter sets are provided, namely MAYO_1 and MAYO_2 for security level I, MAYO_3 for security level III, and MAYO_5 for security level V. The key recovery attacks on MAYO are equivalent to those on plain UOV, although the parameter settings differ. From our complexity evaluation, we confirm that the reconciliation attack reduces the security of MAYO_2, and this estimate is consistent with the symmetric-algebra attack [10]. Our generalization in the parameter ℓ makes the reconciliation attack applicable to the other parameters from the symmetric-algebra attack, although it does not reduce their security. For the parameter sets MAYO_1, MAYO_3, and MAYO_5, the complexity of the reconciliation attack using the proposed framework is higher than that of the original attack in [2]. This is because the attack in [2] exploits fixing a subset of variables (i.e., the hybrid approach) to improve efficiency. We leave it as future work to efficiently combine the proposed framework with the hybrid approach. The table also shows that the intersection attack using the proposed framework does not outperform the reconciliation attack, unlike the case of UOV. This is because the dimension of the solution space in the attack equals 1, and thus we cannot find a feasible degree D with a small ℓ.

QR-UOV. In Table 5, we evaluate the complexity of the attacks using the proposed framework against the three main parameters of QR-UOV, corresponding to security levels I, III, and V, presented in [7]. For QR-UOV, we apply the key recovery attacks by regarding its public key map as that of plain UOV over the extension field $\mathbb{F}_{q^l}$ with v/l vinegar variables, m/l oil variables, and m equations. From our complexity evaluation, we confirm that the complexity of the attacks using the proposed framework is consistent with that of the original attacks presented in [7]. This is because $p = 127$ is larger than the solving degree D, and thus $R_D \simeq R_D^{(p)}$.

Remark 4 (On the application of the proposed framework to SNOVA). There are two main approaches to applying key recovery attacks to SNOVA: (i) The first regards the SNOVA public key as that of plain UOV over $\mathbb{F}_{q^l}$ with v vinegar variables, o oil variables, and o equations [14], which is called the lifting approach. (ii) The second regards the SNOVA public key as that of plain UOV over $\mathbb{F}_q$

Table 5: The complexity of the best known attacks in the specification document [7], the symmetric-algebra attack by Jin et al. [10], and the reconciliation and intersection attacks using our p^ℓ-truncated framework, on the parameters of QR-UOV in [7]. The lowest complexity is highlighted in bold.

SL	(q, v, m, l)	[7]	[10]	our reconciliation complexity	our reconciliation (ℓ, k, D)	our intersection complexity	our intersection (ℓ, k, D)
I	$(127, 156, 54, 3)$	**150**	164	165	$(1, 17, 28)$	454	$(1, 0, 9)$
III	$(127, 228, 78, 3)$	**211**	231	231	$(1, 25, 40)$	652	$(1, 0, 12)$
V	$(127, 306, 105, 3)$	**277**	291	291	$(1, 34, 49)$	851	$(1, 0, 15)$

with vl vinegar variables, ol oil variables, and ol^2 equations [9, 13]. The instance in method (i) is given as a randomly chosen UOV public key map over $\mathbb{F}_{q^l}$. However, since this UOV instance satisfies the inequality in Remark 3, we refrain from presenting a complexity table for method (i). By contrast, for method (ii), the ol^2 equations are not sampled independently, since they are constructed by multiplying certain matrices by the original o public key matrices. This indicates that the semi-regularity assumption is unlikely to hold for this system, and thus it is difficult to precisely evaluate the complexity of method (ii) using the proposed framework. For the sake of reference, in Appendix C, we evaluate the complexity of the proposed framework for method (ii) under the semi-regularity assumption.

6 Conclusion

In this work, we propose a new framework for recovering the secret subspace of UOV, that is, for solving polynomial systems whose solution space contains a vector space. This framework is developed over truncated polynomial rings, which are obtained by setting higher-degree terms to zero. We apply the proposed framework to evaluate the security of the proposed parameters of UOV and its variants.

We first theoretically show that the secret subspace of UOV can be recovered using the XL algorithm by exploiting the structure of the p-truncated polynomial ring $\mathbb{F}_q[\mathbf{x}]/\langle x_1^p, \ldots, x_n^p \rangle$. Our discussion proceeds by focusing on the structure of right kernel vectors of the Macaulay matrix for polynomial systems whose solution space contains a vector space. Furthermore, this result can be naturally extended to the case of the p^ℓ-truncated polynomial ring. These results can be viewed as a generalization of the attacks on UOV based on the symmetric algebra proposed by Jin et al. For simplicity, this paper focuses on quadratic polynomial systems, but the framework can be naturally extended to systems of arbitrary degree. For the cryptanalysis of UOV, we consider applying the reconciliation attack using the proposed framework, following the approach of Jin et al. We further consider applying the intersection attack, which solves the quadratic system whose solution space contains a vector space, obtained by transforming the UOV public key map.

Using the proposed framework, we evaluate the security of the proposed parameters of UOV and its variants selected as third-round candidates in the NIST process for additional digital signatures. For UOV, we show that the intersection attack using the proposed framework reduces the security of several parameter sets. In particular, the attack reduces the security of `uov-V` by approximately 50 bits compared to its claimed security level.

Acknowledgements

The authors thank the anonymous reviewers for their valuable comments, especially those regarding the security evaluation of SNOVA.

This work was supported by JST CREST Grant Number JPMJCR2113, Japan, JST K Program Grant Number JPMJKP24U2, Japan, JSPS KAKENHI Grant Number JP25K00140 and JP24K02939, Japan, and Institute of Mathematics for Industry, Joint Usage/Research Center in Kyushu University (FY2025 Short-term Joint Research "Speeding up of symbolic computation and its application to solving industrial problems 3" (Reference No.2025a012)).

References

1. W. Beullens. Improved cryptanalysis of UOV and Rainbow. In *EUROCRYPT 2021*, volume 12696 of *LNCS*, pages 348–373. Springer, 2021.
2. W. Beullens, F. Campos, S. Celi, B. Hess, and M. Kannwischer. MAYO specification document - round 2. `https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/mayo-spec-round2-web.pdf`, 2025.
3. W. Beullens, M.-S. Chen, J. Ding, B. Gong, M. J. Kannwischer, J. Patarin, B.-Y. Peng, D. Schmidt, C.-J. Shih, C. Tao, and B.-Y. Yang. UOV specification document - round 2. `https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/uov-spec-round2-web.pdf`, 2025.
4. W. Bosma, J. Cannon, and C. Playoust. The Magma algebra system. I. The user language. *J. Symb. Comput.*, 24(3-4):235–265, 1997.
5. N. Courtois, A. Klimov, J. Patarin, and A. Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In *EUROCRYPT 2000*, pages 392–407. Springer, 2000.
6. J. Ding, B. Yang, C.-O. Chen, M. Chen, and C. Cheng. New differential-algebraic attacks and reparametrization of Rainbow. In *ACNS 2008*, volume 5037 of *LNCS*, pages 242–257. Springer, 2008.
7. H. Furue, Y. Ikematsu, F. Hoshino, T. Takagi, H. Kosuge, K. Yamakoshi, R. Akiyama, S. Nakamura, S. Orihara, and K. Kinjo. QR-UOV specification document - round 2. `https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/qruov-spec-round2-web.pdf`, 2025.
8. M.-R. Garey and D.-S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman, 1979.
9. Y. Ikematsu and R. Akiyama. Revisiting the security analysis of SNOVA. In *APKC 2024*, pages 54–61. ACM, 2024.

10. Y. Jin, Y. Pan, X. He, B. Gong, and J. Ding. Security analysis on UOV families with odd characteristics: Using symmetric algebra. In *PKC 2026*, volume 16551 of *LNCS*, pages 428–454. Springer, 2026.
11. A. Kipnis, J. Patarin, and L. Goubin. Unbalanced oil and vinegar signature schemes. In *EUROCRYPT 1999*, volume 1592 of *LNCS*, pages 206–222. Springer, 1999.
12. A. Kipnis and A. Shamir. Cryptanalysis of the oil and vinegar signature scheme. In *CRYPTO 1998*, volume 1462 of *LNCS*, pages 257–266. Springer, 1998.
13. P. Li and J. Ding. Cryptanalysis of the SNOVA signature scheme. In *PQCrypto 2024*, volume 14772 of *LNCS*, pages 79–91. Springer, 2024.
14. S. Nakamura, Y. Tani, and H. Furue. Lifting approach against the SNOVA scheme. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.*, 108(10):1373–1381, 2025.
15. NIST. Post-quantum cryptography CSRC. `https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization`.
16. NIST. NIST: Call for additional digital signature schemes for the post-quantum cryptography standardization process. `https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf`, 2022.
17. NIST. Status report on the third round of the NIST post-quantum cryptography standardization process. NIST Internal Report 8413, 2022.
18. NIST. Status report on the fourth round of the NIST post-quantum cryptography standardization process. NIST Internal Report 8545, 2025.
19. NIST. Status report on the second round of the additional digital signature schemes for the NIST post-quantum cryptography standardization process. NIST Internal Report 8610, 2026.
20. L. Ran. Wedges, oil, and vinegar – an analysis of UOV in the exterior algebra. In *EUROCRYPT 2026*, volume 16544 of *LNCS*, pages 363–388. Springer, 2026.
21. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
22. L.-C. Wang, C.-Y. Chou, J. Ding, Y.-L. Kuan, J. A. Leegwater, M.-S. Li, B.-S. Tseng, P.-E. Tseng, and C.-C. Wang. SNOVA specification document - round 2. `https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/snova-spec-round2-web.pdf`, 2025.
23. D. H. Wiedemann. Solving sparse linear equations over finite fields. *IEEE Trans. Inf. Theor.*, 32(1):54–62, 1986.

A Alternative Parameters for UOV

In this appendix, we provide alternative parameters for UOV and compare their key and signature sizes with those of the original parameters.

We provide alternative parameters corresponding to `uov-Ip`, `uov-III`, and `uov-V`, whose security is reduced by the attacks using the proposed framework in Section 5. We propose three parameter sets with $q = 256$ satisfying the security levels I, III, and V, respectively, as follows:

Table 6: The complexity of the reconciliation and intersection attacks using our p^ℓ-truncated framework on the alternative parameters of UOV

parameter	(q, n, m)	our reconciliation complexity	(ℓ, k, D)	our intersection complexity	(ℓ, k, D)
`uov-Ip'`	$(256, 120, 44)$	160	$(1, 23, 21)$	224	$(2, 2, 29)$
`uov-III'`	$(256, 196, 72)$	239	$(1, 41, 31)$	331	$(2, 6, 42)$
`uov-V'`	$(256, 260, 96)$	297	$(1, 58, 38)$	273	$(1, 0, 28)$

Table 7: Comparison of the key/signature sizes (bytes) for UOV between the original and alternative parameters

(i) original parameters

parameter	(q, n, m)	$\lvert\mathsf{epk}\rvert$	$\lvert\mathsf{esk}\rvert$	$\lvert\mathsf{cpk}\rvert$	$\lvert\sigma\rvert$
`uov-Ip`	$(256, 112, 44)$	278,432	237,896	43,576	128
`uov-III`	$(256, 184, 72)$	1,225,440	1,044,320	189,232	200
`uov-V`	$(256, 244, 96)$	2,869,440	2,436,704	446,992	260

(ii) alternative parameters

parameter	(q, n, m)	$\lvert\mathsf{epk}\rvert$	$\lvert\mathsf{esk}\rvert$	$\lvert\mathsf{cpk}\rvert$	$\lvert\sigma\rvert$
`uov-Ip'`	$(256, 120, 44)$	319,440	279,256	43,576	136
`uov-III'`	$(256, 196, 72)$	1,390,032	1,209,776	189,232	212
`uov-V'`	$(256, 260, 96)$	3,257,280	2,826,080	446,992	276

$$\begin{aligned} \texttt{uov-Ip'}&: (q, n, m) = (256, 120, 44), \\ \texttt{uov-III'}&: (q, n, m) = (256, 196, 72), \\ \texttt{uov-V'}&: (q, n, m) = (256, 260, 96). \end{aligned}$$

For these parameters, we set n and m to be multiples of 4, as in the original parameters. To make these parameters secure, we increase the number n of variables while fixing the number m of equations; equivalently, we increase the number of vinegar variables $v = n - m$. Indeed, Table 6 confirms that the proposed parameters are secure against the attacks using the proposed framework.

In Table 7, we compare the key and signature sizes between the original and alternative parameters of UOV. In the table, $\lvert\mathsf{epk}\rvert$ and $\lvert\mathsf{esk}\rvert$ denote the public and secret key sizes of the classic version of UOV [3]. On the other hand, $\lvert\mathsf{cpk}\rvert$ denotes the public key size of the variant that compresses the keys using seed-expansion techniques. The table omits the compressed secret key size, since it does not depend on the parameters (q, n, m) of the scheme. We also report the signature size as $\lvert\sigma\rvert$. In the case without seed expansion, the alternative parameters increase $\lvert\mathsf{epk}\rvert$ and $\lvert\mathsf{esk}\rvert$ by approximately 10–20% compared to the original parameters. The signature size also increases by approximately 10 bytes for the alternative parameters. By contrast, the compressed public key size $\lvert\mathsf{cpk}\rvert$ is the same for the original and alternative parameters, since it depends only on q and m.

Table 8: The complexity of the symmetric-algebra attack [10], and the reconciliation and intersection attacks using our p^ℓ-truncated framework, on the parameters of UOV proposed by Jin et al. [10]

(n, m)	(p, e)	[10]	our reconciliation complexity	(ℓ, k, D)	our intersection complexity	(ℓ, k, D)
$(112, 44)$	$(3, 5)$	193	196	$(1, 29, 30)$	159	$(1, 11, 18)$
	$(5, 3)$	263	266	$(1, 31, 52)$	165	$(2, 19, 20)$
	$(7, 2)$	310	312	$(1, 32, 72)$	165	$(2, 19, 20)$
$(160, 64)$	$(3, 2)$	249	252	$(1, 45, 38)$	196	$(1, 21, 22)$
	$(5, 1)$	347	349	$(1, 47, 67)$	216	$(1, 25, 26)$
$(184, 72)$	$(3, 5)$	294	302	$(1, 49, 46)$	239	$(1, 18, 27)$
	$(5, 3)$	409	411	$(1, 52, 79)$	250	$(3, 31, 30)$
	$(7, 2)$	483	490	$(1, 53, 111)$	249	$(2, 31, 30)$
$(244, 96)$	$(3, 5)$	367	380	$(1, 67, 57)$	295	$(1, 27, 33)$
	$(5, 3)$	514	523	$(1, 71, 99)$	311	$(3, 43, 37)$
	$(7, 2)$	616	622	$(1, 73, 138)$	311	$(2, 43, 37)$

B Complexity of Proposed Framework for UOV parameters proposed by Jin et al.

In Table 8, we evaluate the complexity of the proposed framework against the UOV parameters considered by Jin et al. [10], which are obtained from the UOV parameters in [3] by changing only the values of p and e. As in the case of the UOV parameters in [3], the reconciliation attack using our framework gives complexity estimates comparable to those of the symmetric-algebra attack in [10]. In contrast, the intersection attack using our framework yields significantly lower complexities for all listed parameters.

C Complexity of Proposed Framework for SNOVA over $\mathbb{F}_q$ under the Semi-regularity Assumption

For the sake of reference, this appendix computes the complexity of the proposed framework for method (ii) applied to SNOVA, which treats the instance as the public key map of UOV with parameters (q, vl, ol, ol^2), under the semi-regularity assumption. (See Table 9.) As discussed in Remark 4, in method (ii), the ol^2 equations are correlated due to their construction. In this appendix, following the estimation approach in the specification document [22], we estimate the complexity of the proposed framework for method (ii) under the assumption that these ol^2 equations behave as a semi-regular system on the p^ℓ-truncated polynomial ring. Note that, for the intersection attack using method (ii), the number of linearly independent equations is $4ol^2 - 2l$ due to the non-symmetry of the public key matrices.

Table 9: The complexity of the reconciliation and intersection attacks using our p^ℓ-truncated framework via UOV with parameter (q, vl, ol, ol^2), on the parameters of SNOVA in [22]

SL	(q, v, o, l)	our reconciliation		our intersection	
		complexity	(ℓ, k, D)	complexity	(ℓ, k, D)
I	$(16, 37, 17, 2)$	129	$(1, 19, 15)$	153	$(4, 0, 12)$
	$(16, 25, 8, 3)$	124	$(1, 10, 14)$	221	$(4, 0, 10)$
	$(16, 24, 5, 4)$	162	$(1, 1, 19)$	349	$(4, 0, 11)$
III	$(16, 56, 25, 2)$	175	$(1, 30, 20)$	-	-
	$(16, 49, 11, 3)$	251	$(1, 2, 31)$	-	-
	$(16, 37, 8, 4)$	225	$(1, 6, 26)$	507	$(4, 0, 15)$
	$(16, 24, 5, 5)$	173	$(1, 6, 19)$	413	$(4, 0, 11)$
V	$(16, 75, 33, 2)$	226	$(1, 40, 26)$	-	-
	$(16, 66, 15, 3)$	325	$(1, 5, 40)$	-	-
	$(16, 60, 10, 4)$	695	$(2, 1, 116)$	-	-
	$(16, 29, 6, 5)$	201	$(1, 8, 22)$	495	$(4, 0, 13)$

D Complexity Estimation Code

This appendix provides a code to estimate the time complexity of the proposed framework in Magma [4].

Listing 1.1: Code to estimate the complexity of the proposed algorithm

```
// Compute the (log2) complexity and the optimal parameters
// p^e: the size of the finite fields
// N: the number of variables
// M: the number of equations
// a: the dimension of the linear solution space

function complexity(p,e,N,M,a)
    q:=p^e;
    com_min:=10^10;
    param_min:=[];
    for l in [1..e] do
        for k in [0..a-1] do
            nn:=N-k;
            aa:=a-k;
            R:=LazyPowerSeriesRing(RationalField(),1);
            G:=(&+[R.1^i: i in [0..p^l-1]])^nn*(&+[R.1^(2*i): i in
                [0..p^l-1]])^(-M);
            D:=0;
            for d in [2..aa*(p^l-1)] do
                if RationalField()! Coefficient(G,d) le &+[(-1)^i*
                    Binomial(aa,i)*Binomial(d-i*p^l+aa-1,aa-1): i in
                    [0..Min(aa,Floor(d/p^l))]] then
                    D:=d;
                    break;
                end if;
            end for;
            if D ne 0 then
```

```
25                  // r:=Log(2,q);
26                  // comp:=Log(2,(2*r^2+r)*3*Binomial(nn+1,2)
27                  com:=Log(2,3*Binomial(nn+1,2)
28                      *(
29                          &+[(-1)^i*Binomial(nn,i)*Binomial(D-i*p^l+nn
                                -1,nn-1): i in [0..Min(nn,Floor(D/p^l))]]
30                          - &+[(-1)^i*Binomial(aa,i)*Binomial(D-i*p^l+aa
                                -1,aa-1): i in [0..Min(aa,Floor(D/p^l))]]
31                          + 1
32                      )^2);
33                  if com lt com_min then
34                      com_min:=com;
35                      param_min:=[l,k,D];
36                  end if;
37              end if;
38          end for;
39      end for;
40      return com_min, param_min;
41  end function;
```