On the Security of the Cubic UOV Variant CML

Manit Gupta

DuPont Manual High School

Abstract. CML is a cubic multivariate signature scheme proposed over $\mathbb{F}_q$ with $q = 2^8$. Earlier cubic Oil–Vinegar variants such as CUOV and CSSv were broken by differential and high-rank attacks. CML adds an ℓ-invertible cycle system and a secret position-hiding map to hide the cubic positions. We show that the cubic homogeneous parts of the public key span a space of dimension at most two, so a publicly derived output combination has degree at most two. After restriction to the oil-direction space, the polar matrices of the publicly derived quadratic components span rank-two matrices associated with the edges of a path. This structure yields a procedure for recovering an equivalent key. Our implementation recovered an equivalent key in under 5 minutes for each tested parameter set, including the set targeting 256-bit security.

1 Introduction

Multivariate public-key cryptography (MPKC) builds public maps from low-degree polynomial systems over finite fields. Such maps are fast to evaluate in the forward direction, but inverting them—as required to forge a signature or recover a private key—is intended to require solving hard systems of multivariate polynomial equations, making them attractive for post-quantum cryptography. The Oil–Vinegar family is one of the central examples. Its original balanced form was broken by Kipnis and Shamir [13], while the Unbalanced Oil–Vinegar variant (UOV) has more vinegar than oil variables and remains one of the most studied multivariate signature paradigms [12,5,21,1]. In UOV, the secret central map is easy to invert after fixing vinegar variables, and affine transformations hide this structure in the public map.

A persistent practical limitation of UOV is that its public keys are relatively large [6]. Cubic variants were introduced to improve efficiency while preserving a secret Oil–Vinegar-type structure that lets the signer invert the central map efficiently. Nie, Liu, Xiong, and Lu proposed CUOV, a cubic unbalanced Oil–Vinegar signature scheme using a stepwise iteration method to obtain shorter signatures and faster signing [17]. Hashimoto later showed that CUOV contains an exploitable differential structure that leads to key recovery [10]. Duong, Petzoldt, Wang, and Takagi revisited CUOV and proposed modified variants, including CSSv and SVSv, to avoid the differential weakness [7]. Hashimoto subsequently showed that these variants were also insecure [10]. Shim, Koo, and Park later presented closely related key-recovery attacks against CSSv and SVSv [20].

The CML proposal [22] follows this history. It introduces a cubic multivariate signature scheme with an ℓ-invertible cycle system, aiming to keep the advantages of cubic MPKC—fast signing, short signatures, and a public key size comparable to that of CUOV—while hiding the positions of the cubic components against the known CUOV/CSSv attacks.

Contributions. We model the relevant CML layer as a masked cycle: in secret coordinates the cycle layer contains products $z_i z_{i+1}$ hidden by input mixing, position hiding, and vinegar masks. After removing the two cubic edges adjacent to the seed vertex $j^\star$, the remaining vertices form a hidden Shamir-type chain, so an equivalent key can be recovered as in [2,3]. We prove a cubic-layer collapse, extract the quadratic polar space, recover the oil-direction space by a public block split, and recover chain ends by public candidate searches followed by rank-one verification. We then estimate the cost for the proposed parameters.

Roadmap. After giving the linear-algebra notation in Section 2, we recall the CML construction and the relevant prior attacks in Section 3. In Section 4, we introduce polar forms. Sections 5 and 6 present the cubic collapse and the exposed quadratic chain, respectively. Section 7 states the public equivalent-key recovery algorithm and proves its correctness. We analyze the concrete cost in Section 8 and conclude in Section 9. The appendix collects the probability estimates and their proofs.

2 Linear-algebra preliminaries

Throughout, $\mathbb{F}_q$ is a finite field with q elements and characteristic 2. Vectors are columns, and matrices act on the left. For any vector or matrix A, $A^\top$ denotes its transpose. For $a, b \in \mathbb{Z}$ with $a \le b$, let $[\![a, b]\!]$ denote the discrete interval from a to b, namely $[\![a, b]\!] = \{i \in \mathbb{Z} : a \le i \le b\}$. For a subspace $U \subseteq \mathbb{F}_q^n$, write

$$\operatorname{ann}(U) = \{\boldsymbol{v} \in \mathbb{F}_q^n : \boldsymbol{v}^\top \boldsymbol{u} = 0 \text{ for all } \boldsymbol{u} \in U\},$$

for its orthogonal complement. A *direction* in an $\mathbb{F}_q$-vector space is a one-dimensional subspace. A nonzero vector belonging to a direction is called a *representative* of that direction. Thus, $\boldsymbol{a}$ represents $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{a}\}$, and two nonzero vectors represent the same direction exactly when they are nonzero scalar multiples of one another. For a subspace $U \subseteq X$, a *direction modulo* U is a one-dimensional subspace of the quotient X/U. A vector $\boldsymbol{a} \in X \setminus U$ represents the quotient direction $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{a} + U\} \subseteq X/U$. We use the standard notions of span, rank, kernel, image, and column space, writing $\operatorname{col} A$ for the column space of a matrix A.

A square matrix $A \in \mathbb{F}_q^{d \times d}$ is *alternating* if $A^\top = A$ and its diagonal entries are zero. Given square matrices P and Q of the same size, their *matrix pencil* consists of the matrices $P - \lambda Q$ for $\lambda \in \mathbb{F}_q$, together with Q.

Linear forms are represented by their coefficient vectors, written with bold $\mathbf{L}$ symbols. For a polynomial f, write $\deg f$ for its total degree and $f_{[d]}$ for

its degree-d homogeneous part. Throughout the paper, we order the degree-d monomials lexicographically with $x_1 > x_2 > \cdots > x_n$, and write $\mathrm{coeff}_d(f)$ for the row vector of coefficients of $f_{[d]}$ in that monomial basis.

Given a family of matrices $(M_i)_{i\in[\![1,k]\!]}$, we write

$$\mathrm{span}_{\mathbb{F}_q}\{M_i : i \in [\![1,k]\!]\} = \left\{ \sum_{i=1}^{k} a_i M_i : a_i \in \mathbb{F}_q \right\}.$$

For a linear subspace $\mathcal{M} \subseteq \mathbb{F}_q^{n\times n}$, its *common radical* is

$$\mathrm{rad}(\mathcal{M}) = \{\boldsymbol{u} \in \mathbb{F}_q^n : M\boldsymbol{u} = \boldsymbol{0} \text{ for every } M \in \mathcal{M}\} = \bigcap_{M\in\mathcal{M}} \ker M.$$

For a subspace $\mathcal{K} \subseteq \mathbb{F}_q^n$, let $U_{\mathcal{K}}$ be any matrix whose columns form a basis of $\mathcal{K}$. The restriction of M to $\mathcal{K}$ is the matrix $U_{\mathcal{K}}^\top M U_{\mathcal{K}}$, which represents the restriction to $\mathcal{K}\times\mathcal{K}$ of the bilinear form with matrix M. Its rank is independent of the chosen basis of $\mathcal{K}$. For a polynomial f, its restriction to $\mathcal{K}$ is the polynomial

$$f|_{\mathcal{K}}(\boldsymbol{y}) = f(U_{\mathcal{K}}\boldsymbol{y}), \qquad \boldsymbol{y} \in \mathbb{F}_q^{\dim \mathcal{K}},$$

where the columns of $U_{\mathcal{K}}$ form a basis of $\mathcal{K}$. In particular, if $\lambda(\boldsymbol{x}) = \mathbf{L}_\lambda^\top \boldsymbol{x}$ is a homogeneous linear form, then its *restricted linear form* on $\mathcal{K}$ is

$$\lambda|_{\mathcal{K}}(\boldsymbol{y}) = \mathbf{L}_\lambda^\top U_{\mathcal{K}}\boldsymbol{y} = (U_{\mathcal{K}}^\top \mathbf{L}_\lambda)^\top \boldsymbol{y}.$$

Thus, in the coordinates determined by $U_{\mathcal{K}}$, the coefficient column vector of $\lambda|_{\mathcal{K}}$ is $U_{\mathcal{K}}^\top \mathbf{L}_\lambda$.

3 The CML construction

We follow the CML proposal [22]. CML is a cubic multivariate signature scheme whose public key is a polynomial map

$$\mathcal{P} : \mathbb{F}_q^n \longrightarrow \mathbb{F}_q^o.$$

Here, $n = o + v$, with the first o variables called oil variables and the last v called vinegar variables. Table 1 summarizes the composition used throughout this paper. Although CML describes $\mathcal{S}$, $\mathcal{H}$, and $\mathcal{T}$ as affine, their translations can be normalized to zero: absorb the translation of $\mathcal{H}$ into the constants of $\widehat{\mathcal{F}}$, the translation of $\mathcal{S}$ into translated central coordinate polynomials and vinegar-only masks, and the translation of $\mathcal{T}$ into the constant terms of the masks. As a result, we take $\mathcal{S}(\boldsymbol{x}) = S\boldsymbol{x}$, $\mathcal{H}(\boldsymbol{y}) = H\boldsymbol{y}$, and $\mathcal{T}(\boldsymbol{y}) = T\boldsymbol{y}$ to be invertible linear maps.

Table 1. The CML composition used in this paper. The public map is $\mathcal{P} = \mathcal{T} \circ \overline{\mathcal{F}} \circ ((\mathcal{H} \circ \widehat{\mathcal{F}}) \times \mathrm{id}_v) \circ \mathcal{S}$: the input is first transformed by the secret input-hiding map $\mathcal{S}$, then mapped by the seed and position-hiding layers, combined with the unchanged vinegar coordinates, passed through the ℓ-invertible cycle layer, and finally mixed by $\mathcal{T}$.

Layer	Map	Domain $\to$ codomain	Secret?
Input hiding	$\mathcal{S}$	$\mathbb{F}_q^n \to \mathbb{F}_q^n$	Yes
Seed map	$\widehat{\mathcal{F}}$	$\mathbb{F}_q^n \to \mathbb{F}_q^o$	Yes
Position hiding	$\mathcal{H}$	$\mathbb{F}_q^o \to \mathbb{F}_q^o$	Yes
Vinegar carry-through	id_v	$\mathbb{F}_q^v \to \mathbb{F}_q^v$	No
Cycle layer	$\overline{\mathcal{F}}$	$\mathbb{F}_q^o \times \mathbb{F}_q^v \to \mathbb{F}_q^o$	Yes
Output mixing	$\mathcal{T}$	$\mathbb{F}_q^o \to \mathbb{F}_q^o$	Yes
Public key	$\mathcal{P}$	$\mathbb{F}_q^n \to \mathbb{F}_q^o$	Public

3.1 CML trapdoor map

The central trapdoor is built from two layers. The first layer is a map $\widehat{\mathcal{F}} : \mathbb{F}_q^n \to \mathbb{F}_q^o$ with one quadratic coordinate polynomial and $o-1$ affine coordinate polynomials:

$$\widehat{f}^{(1)} = \sum_{1 \le i \le j \le n} a_{ij}^{(1)} x_i x_j + \sum_{i=1}^{n} b_i^{(1)} x_i + c^{(1)},$$
$$\widehat{f}^{(k)} = \sum_{i=1}^{n} b_i^{(k)} x_i + c^{(k)}, \qquad k \in [\![2, o]\!].$$

CML then applies the invertible linear map $\mathcal{H}(\boldsymbol{y}) = H\boldsymbol{y}$ to move the quadratic coordinate to a secret position. As specified in [22, Section 4.2], there is a secret index $j^\star \in [\![1, o]\!]$ for which H satisfies

$$H_{i1} = 0 \quad (i \in [\![1, o]\!],\ i \neq j^\star), \qquad H_{j^\star 1} \neq 0.$$

Equivalently, the $j^\star$-th row of H contains $\widehat{f}^{(1)}$ with a nonzero coefficient, and all other rows use only $\widehat{f}^{(2)}, \ldots, \widehat{f}^{(o)}$. Write

$$\boldsymbol{z} = (z_1, \ldots, z_o)^\top = \mathcal{H} \circ \widehat{\mathcal{F}},$$

where $z_i : \mathbb{F}_q^n \to \mathbb{F}_q$ is the i-th coordinate polynomial. Then $z_{j^\star}$ is the unique quadratic coordinate polynomial, while every z_i with $i \neq j^\star$ is affine. Although $\mathcal{H}$ is linear, these z_i may not be linear because the coordinate polynomials $\widehat{f}^{(2)}, \ldots, \widehat{f}^{(o)}$ may have nonzero constant terms.

The second layer is the ℓ-invertible cycle system of Ding, Wolf, and Yang [4]. For odd ℓ, the equations

$$u_i u_{i+1} = \gamma_i \quad (i \in [\![1, \ell]\!]), \quad u_{\ell+1} = u_1,$$

can be solved successively, provided the cycle values are nonzero [4]. CML applies it to the values of the secret coordinate polynomials $(z_i)_{i\in[\![1,o]\!]}$. For each $i \in [\![1, o]\!]$, the i-th component polynomial of its cycle map $\overline{\mathcal{F}} : \mathbb{F}_q^o \times \mathbb{F}_q^v \to \mathbb{F}_q^o$ is

$$\overline{\mathcal{F}}^{(i)} = z_i \cdot z_{i+1} + g_i(x_{o+1}, \ldots, x_n),$$

where indices are read modulo o. Since only $z_{j^\star}$ is quadratic, the components $\overline{\mathcal{F}}^{(j^\star-1)}$ and $\overline{\mathcal{F}}^{(j^\star)}$ are the only ones that can have nonzero cubic parts. Here $g_{j^\star-1}$ and $g_{j^\star}$ are the only polynomials g_i that can contain cubic terms, while the remaining g_i have degree at most two [22, Section 4.2].

Definition 3.1. *A* mask *is a polynomial*

$$g_i \in \mathbb{F}_q[x_{o+1}, \ldots, x_n]$$

depending only on the vinegar variables, added to the product $z_i \cdot z_{i+1}$ in the component polynomial $\overline{\mathcal{F}}^{(i)}$. A mask g_i is called cubic *if* $\deg g_i = 3$ *and* quadratic *if* $\deg g_i = 2$.

3.2 Key generation and signing

The private key consists of the maps $(\mathcal{T}, \overline{\mathcal{F}}, \mathcal{H}, \widehat{\mathcal{F}}, \mathcal{S})$. Here $\mathcal{S}(\boldsymbol{x}) = S\boldsymbol{x}$ is a random invertible linear map. The map $\mathcal{T}(\boldsymbol{y}) = T\boldsymbol{y}$ is also invertible and linear. The matrix T is a random invertible matrix subject to one restriction: the columns indexed by the two cubic-component positions of $\overline{\mathcal{F}}$, $j^\star - 1$ and $j^\star$, have support only in the first three output rows, meaning that for each $k > 3$,

$$T_{k,j^\star-1} = T_{k,j^\star} = 0.$$

All other columns are unrestricted. The public key is the composed map

$$\mathcal{P} = \mathcal{T} \circ \overline{\mathcal{F}} \circ \left((\mathcal{H} \circ \widehat{\mathcal{F}}) \times \mathrm{id}_v\right) \circ \mathcal{S} = (p^{(1)}, \ldots, p^{(o)}).$$

Consequently, $p^{(1)}, p^{(2)}, p^{(3)}$ are the only public component polynomials that can have degree three. The remaining $o - 3$ components have degree at most two.

Signing inverts the trapdoor by choosing vinegar values, solving the secret odd cycle, and then applying the inverse private maps. To verify a signature $\boldsymbol{s} \in \mathbb{F}_q^n$, check whether $\mathcal{P}(\boldsymbol{s}) = \boldsymbol{m}$, where $\boldsymbol{m} \in \mathbb{F}_q^o$ is the hashed message.

3.3 Attacks on cubic UOV-type schemes

In this section, we discuss the prior attacks most relevant to our attack.

Hashimoto's CUOV break. In the original CUOV central map [17], one coordinate polynomial z_1 is quadratic in the vinegar variables, while the rest are linear. The first two public component polynomials both contain the term $z_1 \cdot z_2$, since $y_1 = r_1 z_1 \cdot (1 + z_2) + g_1$ and $y_2 = r_2 z_1 \cdot z_2 + g_2$. Because this term is always tied to the same fixed pair (y_1, y_2), a linear combination of these two public component polynomials cancels it and recovers $z_1(x)$ up to a cubic vinegar term. Differentiating gives the rest of the information required for an equivalent key [10].

Hashimoto's and Shim's attacks on CSSv/SVSv. Hashimoto described SVSv as essentially a thin version of Tsujii's/Shamir's scheme and showed that CSSv is vulnerable to a high-rank key-recovery attack [10]. Shim, Koo, and Park later gave a more detailed security analysis of CSSv-type constructions, exploiting the fact that the cubic component polynomial can be located and removed, after which the remaining public component polynomials expose high-rank quadratic structure suitable for key recovery [20].

4 Polar forms

We use polar matrices to represent the bilinear polar forms of the quadratic polynomial components [1].

Definition 4.1. *For $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ in characteristic 2, its* polar matrix Pol_f *is defined by*

$$(\mathrm{Pol}_f)_{ij} = (\mathrm{Pol}_f)_{ji} = [x_i x_j](f) \quad (i \neq j), \qquad (\mathrm{Pol}_f)_{ii} = 0.$$

Remark 4.2. This is equivalent to the polar-form construction introduced by Beullens for UOV and Rainbow [1]: for a quadratic polynomial $f(\boldsymbol{x}) = \boldsymbol{x}^\top Q \boldsymbol{x} + \boldsymbol{v}^\top \boldsymbol{x} + c$, its symmetric bilinear polar form

$$f'(\boldsymbol{x}, \boldsymbol{y}) = f(\boldsymbol{x} + \boldsymbol{y}) - f(\boldsymbol{x}) - f(\boldsymbol{y}) + f(\boldsymbol{0})$$

satisfies $f'(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^\top (Q + Q^\top)\boldsymbol{y}$. In characteristic 2, the matrix $Q + Q^\top$ corresponds exactly to Pol_f. Specifically, the off-diagonal entries match because $(Q+Q^\top)_{ij} = [x_i x_j](f)$ for $i \neq j$, while the diagonal entries are zero since $2Q_{ii} = 0$. Consequently, Pol_f is the coefficient matrix of the bilinear polar form.

Lemmas 4.3–4.6 are the characteristic-two specializations of standard polar-form identities.

Lemma 4.3. *Let $f \in \mathbb{F}_q[x_1, \ldots, x_n]$ have degree at most two. For an affine map $\mathcal{S}(\boldsymbol{x}) = S\boldsymbol{x} + \boldsymbol{t}$ with linear part S,*

$$\mathrm{Pol}_{f \circ \mathcal{S}} = S^\top \mathrm{Pol}_f \, S.$$

Lemma 4.4. *For $r \geq 1$, scalars $(t_k)_{k=1}^r$, and polynomials $(f^{(k)})_{k=1}^r$,*

$$\mathrm{Pol}_{\sum_{k=1}^r t_k f^{(k)}} = \sum_{k=1}^r t_k \, \mathrm{Pol}_{f^{(k)}} \, .$$

Lemma 4.5. *Let $\mathcal{S}(\boldsymbol{x}) = S\boldsymbol{x}$ be linear, and let f_d be the degree-d homogeneous part of f. Then the degree-d homogeneous part of $f \circ \mathcal{S}$ is $f_d \circ S$.*

Proof. Substituting homogeneous linear forms into a monomial of degree k yields a homogeneous polynomial of degree k or the zero polynomial. Thus only the degree-d homogeneous part of f contributes to the degree-d homogeneous part of $f(S\boldsymbol{x})$, giving $f_d \circ S$. □

Lemma 4.6. *Let $\boldsymbol{x} \in \mathbb{F}_q^n$ be a column vector, let $\mathbf{L}_a, \mathbf{L}_b \in \mathbb{F}_q^n$ be column vectors, and let $\alpha, \beta \in \mathbb{F}_q$. If $a(\boldsymbol{x}) = \mathbf{L}_a^\top \boldsymbol{x} + \alpha$ and $b(\boldsymbol{x}) = \mathbf{L}_b^\top \boldsymbol{x} + \beta$, denote their pointwise product by $(a \cdot b)(\boldsymbol{x}) = a(\boldsymbol{x})b(\boldsymbol{x})$. Then*

$$\mathrm{Pol}_{a\cdot b} = \mathbf{L}_a \mathbf{L}_b^\top + \mathbf{L}_b \mathbf{L}_a^\top.$$

Consequently, for every $\boldsymbol{u} \in \mathbb{F}_q^n$,

$$\mathrm{Pol}_{a\cdot b}\, \boldsymbol{u} = (\mathbf{L}_b^\top \boldsymbol{u})\mathbf{L}_a + (\mathbf{L}_a^\top \boldsymbol{u})\mathbf{L}_b.$$

If $\mathbf{L}_a$ and $\mathbf{L}_b$ are linearly independent, then $\mathrm{Pol}_{a\cdot b}$ has rank 2 and its image is $\mathrm{span}_{\mathbb{F}_q}\{\mathbf{L}_a, \mathbf{L}_b\}$.

Proof. The degree-2 part of the pointwise product $a \cdot b$ is $(\mathbf{L}_a^\top \boldsymbol{x})(\mathbf{L}_b^\top \boldsymbol{x})$. Reading off the off-diagonal quadratic coefficients gives $\mathrm{Pol}_{a\cdot b} = \mathbf{L}_a \mathbf{L}_b^\top + \mathbf{L}_b \mathbf{L}_a^\top$. The formula for $\mathrm{Pol}_{a\cdot b}\, \boldsymbol{u}$, the image, and the rank are immediate. □

For $d \geq 1$ and $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{F}_q^d$, define

$$\mathsf{P}(\boldsymbol{a}, \boldsymbol{b}) = \boldsymbol{a}\boldsymbol{b}^\top + \boldsymbol{b}\boldsymbol{a}^\top \in \mathbb{F}_q^{d\times d}.$$

Lemma 4.7. *Let $d \geq 1$ and $1 \leq r \leq d$ be integers, and let $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_r \in \mathbb{F}_q^d$ be linearly independent. Then the matrices $\mathsf{P}(\boldsymbol{a}_i, \boldsymbol{a}_j)$, for $1 \leq i < j \leq r$, are linearly independent.*

Proof. Choose $L \in \mathrm{GL}_d(\mathbb{F}_q)$ such that $L\boldsymbol{a}_i = \boldsymbol{e}_i$ for $i \in [\![1, r]\!]$, where $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_d$ are the standard coordinate-basis vectors of $\mathbb{F}_q^d$. The invertible linear map $M \mapsto LML^\top$ sends $\mathsf{P}(\boldsymbol{a}_i, \boldsymbol{a}_j)$ to $\mathsf{P}(\boldsymbol{e}_i, \boldsymbol{e}_j)$. For each pair $i < j$, the latter matrix has nonzero entries only in positions (i, j) and (j, i). These pairs of positions are distinct, so the matrices are linearly independent. □

Definition 4.8. *The* oil subspace *is*

$$\mathcal{O} = \{\boldsymbol{w} \in \mathbb{F}_q^n : w_{o+1} = \cdots = w_n = 0\}, \quad \dim \mathcal{O} = o.$$

The vinegar subspace *is*

$$W = \{\boldsymbol{w} \in \mathbb{F}_q^n : w_1 = \cdots = w_o = 0\}, \quad \dim W = v.$$

Thus $\mathbb{F}_q^n = \mathcal{O} \oplus W$ as a direct sum of subspaces.

Lemma 4.9. *For every mask $g_i \in \mathbb{F}_q[x_j : j \in [\![o+1, n]\!]]$ and every $\boldsymbol{w} \in \mathcal{O}$, one has $\mathrm{Pol}_{g_i}\, \boldsymbol{w} = 0$.*

Proof. By definition, g_i depends only on the vinegar variables $x_{o+1}, \ldots, x_n$. For $\boldsymbol{w} \in \mathcal{O}$, the vinegar coordinates of $\boldsymbol{w}$ are zero, so $\mathrm{Pol}_{g_i}\, \boldsymbol{w} = 0$. □

Consequently, for every $\boldsymbol{w} \in \mathcal{O}$, we have

$$\mathrm{Pol}_{z_i \cdot z_{i+1} + g_i}\, \boldsymbol{w} = \mathrm{Pol}_{z_i \cdot z_{i+1}}\, \boldsymbol{w}.$$

5 Weakness 1. The cubic layer collapses

We begin with the degree-3 homogeneous parts of the public component polynomials. Only the first three public components can have nonzero cubic parts, and every cubic term comes from either $\overline{\mathcal{F}}^{(j^\star-1)}$ or $\overline{\mathcal{F}}^{(j^\star)}$. This gives a publicly derivable linear relation among the cubic parts.

Theorem 5.1. *The degree-3 homogeneous parts of the public key span a vector space of dimension at most 2. That is,*

$$\dim \operatorname{span}_{\mathbb{F}_q} \left\{ p^{(k)}_{[3]} : k \in [\![1, o]\!] \right\} \leq 2.$$

Proof. By the construction in Section 3, $z_{j^\star}$ is the only quadratic coordinate polynomial and all other z_i are affine. Hence a product $z_i \cdot z_{i+1}$ can have degree 3 only for $i \in \{j^\star - 1, j^\star\}$. The masks $g_{j^\star-1}$ and $g_{j^\star}$ are the only masks that can have nonzero cubic parts. Let $\Gamma_i = (z_i \cdot z_{i+1} + g_i)_{[3]}$ for $i \in \{j^\star - 1, j^\star\}$, and set $\Gamma_i := 0$ for every other index. By Lemma 4.5, the degree-3 homogeneous part of the i-th component polynomial of the map

$$\overline{\mathcal{F}} \circ \left((\mathcal{H} \circ \widehat{\mathcal{F}}) \times \mathrm{id}_v \right) \circ \mathcal{S}$$

is $\Gamma_i \circ \mathcal{S}$. Thus only the component polynomials at positions $j^\star - 1$ and $j^\star$ can have nonzero cubic homogeneous parts, namely $\Gamma_{j^\star-1} \circ \mathcal{S}$ and $\Gamma_{j^\star} \circ \mathcal{S}$.

Since T is linear, taking degree-3 homogeneous parts commutes with applying T. Therefore, for every $k \in [\![1, o]\!]$,

$$p^{(k)}_{[3]} = \sum_{i=1}^{o} T_{k,i}(\Gamma_i \circ \mathcal{S}).$$

For $k > 3$, all summands on the right are zero, so only $p^{(1)}, p^{(2)}$, and $p^{(3)}$ can have nonzero cubic parts. Every cubic homogeneous part of a public component lies in the span of $\Gamma_{j^\star-1} \circ \mathcal{S}$ and $\Gamma_{j^\star} \circ \mathcal{S}$, which has dimension at most two. □

Algorithm 1 (Cubic collapse). *Input:* the public key $\mathcal{P} = (p^{(1)}, \ldots, p^{(o)})$. *Output:* a matrix $\widehat{\Delta} \in \mathrm{GL}_o(\mathbb{F}_q)$ and a map $\widetilde{\mathcal{P}} = \widehat{\Delta}\mathcal{P}$ whose first component and final $o - 3$ components, namely components $1, 4, \ldots, o$, have degree at most 2.

1. Extract the coefficient vectors of the degree-3 homogeneous parts of $p^{(1)}, p^{(2)}$, and $p^{(3)}$ using the lexicographic monomial ordering in Section 2.
2. Form the matrix B whose three rows are the coefficient vectors of these cubic homogeneous parts.
3. Compute a nonzero vector $\boldsymbol{c} = (c_1, c_2, c_3)^\top \in \ker B^\top$ and set

$$Q = \sum_{k=1}^{3} c_k p^{(k)}.$$

4. Extend the nonzero column vector $\boldsymbol{c}$ to a basis $\{\boldsymbol{c}, \boldsymbol{d}_2, \boldsymbol{d}_3\}$ of $\mathbb{F}_q^3$. Set

$$\Delta = \begin{bmatrix} \boldsymbol{c}^\top \\ \boldsymbol{d}_2^\top \\ \boldsymbol{d}_3^\top \end{bmatrix} \in \mathrm{GL}_3(\mathbb{F}_q).$$

Define

$$\begin{bmatrix} \widetilde{p}^{(1)} \\ \widetilde{p}^{(2)} \\ \widetilde{p}^{(3)} \end{bmatrix} = \Delta \begin{bmatrix} p^{(1)} \\ p^{(2)} \\ p^{(3)} \end{bmatrix}.$$

Then $\widetilde{p}^{(1)} = Q$. Return

$$\widetilde{\mathcal{P}} = (\widetilde{p}^{(1)}, \widetilde{p}^{(2)}, \widetilde{p}^{(3)}, p^{(4)}, \ldots, p^{(o)}) = \widehat{\Delta}\mathcal{P}, \qquad \widehat{\Delta} = \mathrm{diag}(\Delta, I_{o-3}).$$

Theorem 5.1 guarantees that Step 3 of Algorithm 1 succeeds. To solve $\mathcal{P}(\boldsymbol{x}) = \boldsymbol{m}$, set $\widetilde{\boldsymbol{m}} = \widehat{\Delta}\boldsymbol{m}$. Then $\mathcal{P}(\boldsymbol{x}) = \boldsymbol{m}$ if and only if $\widetilde{\mathcal{P}}(\boldsymbol{x}) = \widetilde{\boldsymbol{m}}$. Indeed, the former equality implies the latter after multiplication by $\widehat{\Delta}$. Conversely, the latter equality implies the former after multiplication by $\widehat{\Delta}^{-1}$, which exists because $\widehat{\Delta} \in \mathrm{GL}_o(\mathbb{F}_q)$.

6 Weakness 2. The low-rank quadratic chain

The cubic collapse exposes a second weakness: canceling the two cycle-layer components containing $z_{j^\star}$ leaves a path on the nonseed coordinates. We first state the publicly checkable conditions under which the derived quadratic polynomials span the polar matrices associated with all of its edges. Then we define this path and the polar matrices associated with its edges. After that, we show that the public quadratic polar space is the span of these edge matrices, transformed by the secret input map. Finally, we show how a witness at a chain end produces a rank-one matrix whose column space reveals the linear-form direction of the neighboring vertex.

Publicly checkable conditions. The remainder of the attack is conditional on the publicly checkable conditions

$$\operatorname{rank} B = 2, \qquad \dim \operatorname{span}_{\mathbb{F}_q}\{\mathrm{Pol}_Q, \mathrm{Pol}_{p^{(4)}}, \ldots, \mathrm{Pol}_{p^{(o)}}\} = o - 2, \tag{1}$$

where B is the matrix in Algorithm 1. More precisely, the rows of B lie in

$$\operatorname{span}_{\mathbb{F}_q}\left\{\mathrm{coeff}_3\big((z_{j^\star-1} \cdot z_{j^\star} + g_{j^\star-1}) \circ \mathcal{S}\big), \mathrm{coeff}_3\big((z_{j^\star} \cdot z_{j^\star+1} + g_{j^\star}) \circ \mathcal{S}\big)\right\}.$$

The construction therefore guarantees only $\operatorname{rank} B \leq 2$. The two coefficient vectors generating the span above may be zero or linearly dependent. When equality holds, write the first row of $\widehat{\Delta}T$ as $(a_1, \ldots, a_o)$. Then

$$Q = \sum_{i=1}^{o} a_i (f_i \circ \mathcal{S}), \qquad a_{j^\star-1} = a_{j^\star} = 0,$$

so neither $f_{j^\star-1} \circ \mathcal{S}$ nor $f_{j^\star} \circ \mathcal{S}$ contributes to Q. The second condition ensures that the $o-2$ matrices $\mathrm{Pol}_Q, \mathrm{Pol}_{p^{(4)}}, \ldots, \mathrm{Pol}_{p^{(o)}}$ are linearly independent. As shown below, each of these matrices is a linear combination of the $o-2$ polar matrices associated with the chain edges. Equality of the two dimensions implies

$$\mathrm{span}_{\mathbb{F}_q}\{\mathrm{Pol}_Q, \mathrm{Pol}_{p^{(4)}}, \ldots, \mathrm{Pol}_{p^{(o)}}\} = \mathrm{span}_{\mathbb{F}_q}\{\mathrm{Pol}_{f_i \circ \mathcal{S}} : i \in [\![1, o]\!] \setminus \{j^\star - 1, j^\star\}\}.$$

Thus the public matrices recover the full span of the polar matrices of the surviving central components, not a proper subspace of it.

6.1 CML's quadratic layer

For a graph $\mathcal{G}$, let $V(\mathcal{G})$ and $E(\mathcal{G})$ denote its vertex and edge sets.

Definition 6.1. *The* cycle graph *is the undirected graph G with vertex set $[\![1, o]\!]$ and edge set $E = \{\{i, i+1\} : i \in [\![1, o]\!]\}$, where indices are taken modulo o. Vertex i is labeled by the coordinate polynomial z_i. The edge $\{i, i+1\}$ is labeled by the component polynomial $\overline{\mathcal{F}}^{(i)}$. Vertex $j^\star$ is called the* seed vertex.

Definition 6.2. *The* chain graph*, denoted by J, is the subgraph of G obtained by deleting the seed vertex $j^\star$ and the two cycle edges incident to it. Equivalently, J has vertex set $[\![1, o]\!] \setminus \{j^\star\}$. Writing E_J for the edge set of J, we have*

$$E_J = \{\{i, i+1\} \in E : i \notin \{j^\star - 1, j^\star\}\}.$$

For each edge $\{i, i+1\} \in E_J$, the polynomial $\overline{\mathcal{F}}^{(i)}$ labels that edge and has $z_i \cdot z_{i+1}$ as its non-mask term. The subgraph J is a path whose edges are called chain edges *and whose degree-1 vertices are called* chain ends*. Its chain ends are $j^\star - 1$ and $j^\star + 1$, labeled by $z_{j^\star-1}$ and $z_{j^\star+1}$, respectively. All other vertices are* interior vertices*. A chain edge incident to a chain end is called a* chain-end edge *and all other chain edges are called* interior edges*.*

Choose either orientation of the path J, and write $(i_1, \ldots, i_{o-1})$ for its vertices in path order. Along each chain edge, the product term $z_i \cdot z_{i+1}$ is the product of the two adjacent affine coordinate polynomials. Hence these product terms are

$$z_{i_1} \cdot z_{i_2},\ z_{i_2} \cdot z_{i_3}, \ldots, z_{i_{o-2}} \cdot z_{i_{o-1}}.$$

Shamir's birational permutation construction uses the same pattern of products between coordinate polynomials [19]. For $j \in [\![1, o-2]\!]$, define the j-th edge in this path order by $e_j = \{i_j, i_{j+1}\}$, and let $\kappa_j \in [\![1, o]\!]$ be the unique index such that $e_j = \{\kappa_j, \kappa_j + 1\}$, with cycle indices read modulo o. Thus i_j labels a vertex in the path ordering, whereas κ_j labels the component $\overline{\mathcal{F}}^{(\kappa_j)}$ attached to the edge e_j.

Remark 6.3. The indices i_j and κ_j are not identical in general. For example, let $o = 3$ and $j^\star = 3$, and orient the path as $(i_1, i_2) = (2, 1)$. Then $e_1 = \{2, 1\} = \{1, 2\}$, so $i_1 = 2$ but $\kappa_1 = 1$.

Definition 6.4. *For each $i \in [\![1, o]\!]$, the* secret linear coefficient vector $\mathbf{L}_i \in \mathbb{F}_q^n$ *is the coefficient vector of the degree-*1 *homogeneous part of z_i. For every $i \neq j^\star$, there is a unique $\alpha_i \in \mathbb{F}_q$ such that*

$$z_i(\boldsymbol{x}) = \mathbf{L}_i^\top \boldsymbol{x} + \alpha_i.$$

The coefficient vector of the degree-1 part of $z_i \circ \mathcal{S}$ is $S^\top \mathbf{L}_i$. The recovery procedure below targets the subspaces $\mathrm{span}_{\mathbb{F}_q}\{S^\top \mathbf{L}_i\}$ for $i \neq j^\star$. For each $i \in [\![1, o]\!]$, let f_i be the i-th component of the map $\overline{\mathcal{F}} \circ ((\mathcal{H} \circ \widehat{\mathcal{F}}) \times \mathrm{id}_v)$. Explicitly,

$$f_i(\boldsymbol{x}) = z_i(\boldsymbol{x}) \cdot z_{i+1}(\boldsymbol{x}) + g_i(x_{o+1}, \ldots, x_n).$$

Definition 6.5. *For a chain edge $e = \{i, i{+}1\} \in E_J$, the* edge polar *is the polar matrix of f_i, defined as*

$$C_e = \mathrm{Pol}_{f_i} = \mathbf{L}_i \mathbf{L}_{i+1}^\top + \mathbf{L}_{i+1} \mathbf{L}_i^\top + \mathrm{Pol}_{g_i}\,.$$

The bare edge polar *associated with e is the polar of its non-mask product:*

$$C_e^{\mathrm{bare}} = \mathrm{Pol}_{z_i \cdot z_{i+1}} = \mathbf{L}_i \mathbf{L}_{i+1}^\top + \mathbf{L}_{i+1} \mathbf{L}_i^\top.$$

6.2 The publicly derived quadratic polar space

Set $m = o - 2$ and collect the publicly derived quadratic polynomials as

$$(q_1, \ldots, q_m) = (Q, p^{(4)}, \ldots, p^{(o)}).$$

Under (1), the matrices Pol_{q_k}, for $k \in [\![1, m]\!]$, are linearly independent. Recall the path ordering defined after Definition 6.2. Since $m = o - 2$, its edges are $e_1, \ldots, e_m$. From the definition of $\mathcal{P}$,

$$\mathcal{P}(\boldsymbol{x}) = T \begin{bmatrix} f_1(S\boldsymbol{x}) \\ \vdots \\ f_o(S\boldsymbol{x}) \end{bmatrix} = T \begin{bmatrix} (f_1 \circ \mathcal{S})(\boldsymbol{x}) \\ \vdots \\ (f_o \circ \mathcal{S})(\boldsymbol{x}) \end{bmatrix}.$$

Let $\widehat{\Delta}$ be the matrix constructed in Algorithm 1. By the preceding identity,

$$\widehat{\Delta}\mathcal{P}(\boldsymbol{x}) = \widehat{\Delta} T \begin{bmatrix} (f_1 \circ \mathcal{S})(\boldsymbol{x}) \\ \vdots \\ (f_o \circ \mathcal{S})(\boldsymbol{x}) \end{bmatrix}.$$

Set $r_1 = 1$ and $r_k = k + 2$ for all $k \in [\![2, m]\!]$. By the definition of the q_k, the r_k-th component of $\widehat{\Delta}\mathcal{P}$ is q_k. Therefore, for every $k \in [\![1, m]\!]$,

$$q_k(\boldsymbol{x}) = \sum_{i=1}^{o} \left(\widehat{\Delta} T\right)_{r_k, i} (f_i \circ \mathcal{S})(\boldsymbol{x}). \tag{2}$$

Define

$$\Gamma_- := (f_{j^\star-1} \circ \mathcal{S})_{[3]}, \qquad \Gamma_+ := (f_{j^\star} \circ \mathcal{S})_{[3]}.$$

Every row of the publicly derived cubic coefficient matrix B is a linear combination of $\mathrm{coeff}_3(\Gamma_-)$ and $\mathrm{coeff}_3(\Gamma_+)$. Hence

$$\operatorname{rank} B \le \dim \operatorname{span}_{\mathbb{F}_q}\{\mathrm{coeff}_3(\Gamma_-), \mathrm{coeff}_3(\Gamma_+)\} \le 2.$$

Under (1), $\operatorname{rank} B = 2$, so Γ_- and Γ_+ are linearly independent. The polynomials $f_{j^\star-1} \circ \mathcal{S}$ and $f_{j^\star} \circ \mathcal{S}$, whose cubic homogeneous parts are respectively Γ_- and Γ_+, are the only component polynomials of

$$\overline{\mathcal{F}} \circ \left((\mathcal{H} \circ \widehat{\mathcal{F}}) \times \mathrm{id}_v\right) \circ \mathcal{S}$$

with nonzero cubic homogeneous parts. For each $k \in [\![1, m]\!]$, taking cubic homogeneous parts in (2) therefore gives

$$0 = (q_k)_{[3]} = (\widehat{\Delta T})_{r_k, j^\star-1}\Gamma_- + (\widehat{\Delta T})_{r_k, j^\star}\Gamma_+,$$

where the first equality holds because every q_k has degree at most two. Since Γ_- and Γ_+ are linearly independent, we have $(\widehat{\Delta T})_{r_k, j^\star-1} = (\widehat{\Delta T})_{r_k, j^\star} = 0$. Thus the r_k-th row of $\widehat{\Delta T}$ has zero entries in columns $j^\star - 1$ and $j^\star$. Its remaining m columns, indexed by $\kappa_1, \ldots, \kappa_m$, correspond exactly to the chain edges $e_1, \ldots, e_m$.

Define $\theta_{kj} = (\widehat{\Delta T})_{r_k, \kappa_j}$ and $\Theta = (\theta_{kj}) \in \mathbb{F}_q^{m \times m}$. Equivalently, Θ is the $m \times m$ submatrix of $\widehat{\Delta T}$ obtained by selecting rows $r_1, \ldots, r_m$ and columns $\kappa_1, \ldots, \kappa_m$. Therefore, for all $k \in [\![1, m]\!]$,

$$q_k(\boldsymbol{x}) = \sum_{j=1}^{m} \theta_{kj}(f_{\kappa_j} \circ \mathcal{S})(\boldsymbol{x}).$$

Since $\widehat{\Delta T}$ is invertible, its selected rows are linearly independent. Every selected row has zero entries in the two omitted columns, so any linear dependence among the rows after deleting those columns would give the same dependence among the original selected rows. Hence the rows of Θ are linearly independent, and Θ is invertible. Lemmas 4.3 and 4.4 give

$$\mathrm{Pol}_{q_k} = S^\top \left(\sum_{j=1}^{m} \theta_{kj} C_{e_j}\right) S, \tag{3}$$

for all $k \in [\![1, m]\!]$. Define the *quadratic polar space*

$$V := \operatorname{span}_{\mathbb{F}_q}\{\mathrm{Pol}_{q_k} : k \in [\![1, m]\!]\}.$$

By (3) and the invertibility of Θ,

$$V = S^\top \left(\operatorname{span}_{\mathbb{F}_q}\{C_e : e \in E_J\}\right) S. \tag{4}$$

6.3 Column-space mechanism

For $\boldsymbol{u} \in \mathbb{F}_q^n$, define

$$A(\boldsymbol{u}) = \left[\mathrm{Pol}_{q_1}\,\boldsymbol{u}\ \mathrm{Pol}_{q_2}\,\boldsymbol{u} \cdots\ \mathrm{Pol}_{q_m}\,\boldsymbol{u}\right] \in \mathbb{F}_q^{n\times m},$$

so $\operatorname{col} A(\boldsymbol{u}) = \operatorname{span}_{\mathbb{F}_q}\{\mathrm{Pol}_{q_k}\,\boldsymbol{u} : k \in [\![1, m]\!]\}$.

Chain-end witnesses. Let $J' \subseteq J$ be a subpath with at least one edge, and let π be a chain end of J'. A vector $\boldsymbol{u} \in \mathbb{F}_q^n$ is a *chain-end witness at π for J'* if $\boldsymbol{w} = S\boldsymbol{u} \in \mathcal{O}$ satisfies

$$\mathbf{L}_\pi^\top \boldsymbol{w} = 1, \qquad \mathbf{L}_j^\top \boldsymbol{w} = 0 \quad \text{for every } j \in V(J') \setminus \{\pi\}.$$

For each chain end π of J, we now prove the existence of a chain-end witness $\boldsymbol{u}_\pi$ satisfying $\operatorname{rank} A(\boldsymbol{u}_\pi) = 1$.

Reading $A(\boldsymbol{u})$ one column at a time. Fix $\boldsymbol{u} \in \mathbb{F}_q^n$ and set $\boldsymbol{w} = S\boldsymbol{u}$. Using (3) and linearity,

$$\mathrm{Pol}_{q_k}\,\boldsymbol{u} = S^\top \left(\sum_{j=1}^{m} \theta_{kj} C_{e_j}\right) S\boldsymbol{u} = S^\top \sum_{j=1}^{m} \theta_{kj} C_{e_j} \boldsymbol{w}. \tag{5}$$

Thus the k-th column of $A(\boldsymbol{u})$ is a linear combination of the vectors $S^\top C_e \boldsymbol{w}$ for $e \in E_J$. If $\boldsymbol{w} \in \mathcal{O}$, then Lemmas 4.6 and 4.9 give, for $e = \{i, i+1\} \in E_J$,

$$C_e \boldsymbol{w} = (\mathbf{L}_{i+1}^\top \boldsymbol{w})\mathbf{L}_i + (\mathbf{L}_i^\top \boldsymbol{w})\mathbf{L}_{i+1}. \tag{6}$$

Consequently, $C_e \boldsymbol{w} \in \operatorname{span}_{\mathbb{F}_q}\{\mathbf{L}_i, \mathbf{L}_{i+1}\}$ and can be nonzero only if at least one of $\mathbf{L}_i^\top \boldsymbol{w}$ and $\mathbf{L}_{i+1}^\top \boldsymbol{w}$ is nonzero.

Isolating a single chain end. Assume that the linear forms $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$, for $i \neq j^\star$, are linearly independent on $\mathcal{O}$. Then the map $\mathcal{O} \to \mathbb{F}_q^{o-1}$ given by $\boldsymbol{w} \mapsto (\mathbf{L}_i^\top \boldsymbol{w})_{i \neq j^\star}$ has rank $o-1$. Since the codomain also has dimension $o-1$, this map is surjective. Therefore, for a chain end π, there exists $\boldsymbol{w}_\pi \in \mathcal{O}$ such that

$$\mathbf{L}_\pi^\top \boldsymbol{w}_\pi = 1, \qquad \mathbf{L}_j^\top \boldsymbol{w}_\pi = 0 \quad \text{for every chain vertex } j \neq \pi.$$

Then $\boldsymbol{u}_\pi := S^{-1}\boldsymbol{w}_\pi$ is a chain-end witness at π for J. With this choice, consider an edge $e = \{a, b\} \in E_J$.

- If e is not incident to π, then $\mathbf{L}_a^\top \boldsymbol{w}_\pi = \mathbf{L}_b^\top \boldsymbol{w}_\pi = 0$, so (6) gives $C_e \boldsymbol{w}_\pi = \mathbf{0}$.
- If e is incident to π, then since π is a chain end there is exactly one such edge. Write it as $e_j = \{\pi, \pi'\}$. Then (6) gives

$$C_{e_j} \boldsymbol{w}_\pi = (\mathbf{L}_{\pi'}^\top \boldsymbol{w}_\pi)\mathbf{L}_\pi + (\mathbf{L}_\pi^\top \boldsymbol{w}_\pi)\mathbf{L}_{\pi'} = \mathbf{L}_{\pi'}.$$

Thus every nonincident edge acts trivially on $\boldsymbol{w}_\pi$, whereas the unique incident edge maps it to $\mathbf{L}_{\pi'}$.

The rank-one consequence. For the chain-end witness $\boldsymbol{u}_\pi$, Equation (5) gives $\mathrm{Pol}_{q_k}\,\boldsymbol{u}_\pi = \theta_{kj}(S^\top \mathbf{L}_{\pi'})$ for every $k \in [\![1, m]\!]$, and therefore

$$A(\boldsymbol{u}_\pi) = \left(S^\top \mathbf{L}_{\pi'}\right)\left[\theta_{1j}\ \theta_{2j}\ \cdots\ \theta_{mj}\right]. \tag{7}$$

The independence assumption gives $\mathbf{L}_{\pi'} \neq \mathbf{0}$. Since S and Θ are invertible, both factors in (7) are nonzero, so $A(\boldsymbol{u}_\pi)$ has rank one, and

$$\operatorname{col} A(\boldsymbol{u}_\pi) = \operatorname{span}_{\mathbb{F}_q}\{S^\top \mathbf{L}_{\pi'}\}.$$

Thus any nonzero column gives the direction of the linear part of $z_{\pi'} \circ \mathcal{S}$, where π' is the neighbor of the chain end π.

7 Equivalent-key recovery

Apply Algorithm 1, and then, as in Section 6.2, form the quadratic polynomials

$$(q_1, \ldots, q_m) = (Q, p^{(4)}, \ldots, p^{(o)}),$$

where $m = o - 2$, together with the quadratic polar space V. The remainder of the attack is organized into Algorithms 2–4. Let $K = S^{-1}\mathcal{O}$ be the *oil-direction* space. The recovery is split into three stages: recover K, recover the chain graph, and reconstruct $z_{j^\star} \circ \mathcal{S}$ and an equivalent key.

7.1 Recovery procedure

The first stage recovers a candidate oil-direction space and the associated geometric data needed by the later stages. This algorithm was inspired by [13].

Algorithm 2 (Oil-direction-space recovery). *Input:* $o \geq 5$, a public key $\mathcal{P}$, and the publicly derived quadratic polar space V.
Output: a tuple $(\mathcal{V}', K', \mathcal{O}', \mathcal{D}')$, where K' is a candidate oil-direction space.

1. Set

$$\mathcal{M} = \operatorname{span}_{\mathbb{F}_q}\{\mathrm{Pol}_{p^{(4)}}, \ldots, \mathrm{Pol}_{p^{(o)}}\}.$$

Choose $N_1, \ldots, N_\ell$ independently and uniformly at random from $\mathcal{M}$, where $\ell \geq 2$, and set

$$\delta = \min_{\lambda \in \mathbb{F}_q} \dim \ker(N_1 - \lambda N_2).$$

Construct

$$\mathcal{D}_{\mathrm{sing}} = \sum_{\substack{1 \leq a < b \leq \ell,\ \lambda \in \mathbb{F}_q \\ \dim \ker(N_a - \lambda N_b) = \delta}} \ker(N_a - \lambda N_b).$$

The subscript sing abbreviates "singular", since $\mathcal{D}_{\mathrm{sing}}$ is constructed from kernels of singular matrices $N_a - \lambda N_b$.

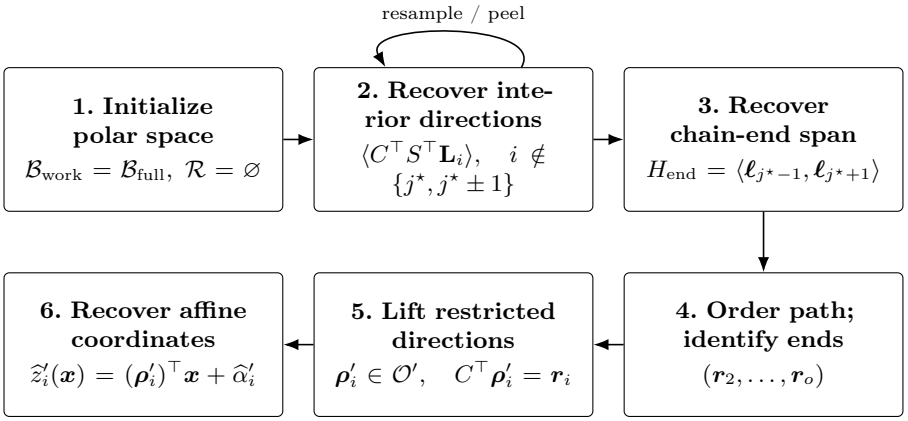


Fig. 1. Overview of Algorithm 3. Here $\boldsymbol{\ell}_i = C^\top S^\top \mathbf{L}_i$ denotes a restricted coefficient vector.

2. If $\dim \mathcal{D}_{\text{sing}} = v + 1$, form

$$\mathcal{V}' = \operatorname{span}_{\mathbb{F}_q}\{M\boldsymbol{d} : M \in \mathcal{M},\ \boldsymbol{d} \in \mathcal{D}_{\text{sing}}\}, \qquad K' = \operatorname{ann}(\mathcal{V}'),$$
$$\mathcal{O}' = \operatorname{span}_{\mathbb{F}_q}\{M\boldsymbol{x} : M \in \mathcal{M},\ \boldsymbol{x} \in K'\}, \qquad \mathcal{D}' = \operatorname{ann}(\mathcal{O}').$$

Here $\mathcal{V}'$ is the candidate vinegar-form space, K' is the candidate oil-direction space, and $\mathcal{O}'$ is the candidate for $\operatorname{span}_{\mathbb{F}_q}\{S^\top \mathbf{L}_i : i \neq j^\star\}$. The space $\mathcal{D}'$ is the common kernel of the homogeneous linear parts of the nonseed coordinates. Return this tuple if

$$\dim \mathcal{V}' = v, \qquad \dim \mathcal{O}' = o - 1. \tag{8}$$

No separate dimension tests are needed for K' or $\mathcal{D}'$: since $K' = \operatorname{ann}(\mathcal{V}')$, $\mathcal{D}' = \operatorname{ann}(\mathcal{O}')$, and $n = o+v$, these tests imply $\dim K' = o$ and $\dim \mathcal{D}' = v+1$.

Given a candidate tuple from the first stage, the next algorithm recovers the chain graph J and orders its nonseed coordinates. The following algorithm is inspired by the attacks of [2,3]. For this algorithm, we follow the procedure in Figure 1.

Algorithm 3 (Chain-graph recovery).
Input: $o \geq 5$, the polynomials $q_1, \ldots, q_m$ from a run of Algorithm 1, and a candidate oil-space tuple returned by Algorithm 2.
Output: candidate affine polynomials $\widehat{z}'_2, \ldots, \widehat{z}'_o$, listed in path order. A successful output consists of polynomials individually proportional to the nonseed polynomials $z_i \circ \mathcal{S}$, up to reversal of the path order.

1. Choose $C \in \mathbb{F}_q^{n \times o}$ with $\operatorname{col} C = K'$, and set

$$\mathcal{B}_{\text{full}} := \operatorname{span}_{\mathbb{F}_q}\{C^\top \operatorname{Pol}_{q_k} C : k \in [\![1, m]\!]\}, \qquad \mathcal{B}_{\text{work}} := \mathcal{B}_{\text{full}}.$$

Thus $\mathcal{B}_{\text{full}}$ is the quadratic polar space V, restricted to K' and represented in the basis C. Initially no chain end has been removed, so $\mathcal{B}_{\text{work}} = \mathcal{B}_{\text{full}}$. Initialize $\mathcal{R} = \varnothing$ to store the distinct directions associated with the chain vertices after restriction to K'.

2. Repeat this step while $\dim \mathcal{B}_{\text{work}} \geq 2$. Call the computation performed at a fixed $\mathcal{B}_{\text{work}}$ a *peeling stage*. At the start of each peeling stage, choose a basis $B_1, \ldots, B_s \in \mathbb{F}_q^{o\times o}$ of $\mathcal{B}_{\text{work}}$, and, for $\boldsymbol{c} \in \mathbb{F}_q^o$, let

$$R_{\text{work}} = \bigcap_{i=1}^{s} \ker B_i, \qquad A_{\text{work}}(\boldsymbol{c}) = [B_1\boldsymbol{c} \mid \cdots \mid B_s\boldsymbol{c}].$$

When $\operatorname{rank} A_{\text{work}}(\boldsymbol{c}) = 1$, the quotient direction represented by $\boldsymbol{c}$ is a candidate for which $C\boldsymbol{c}$ may be a chain-end witness.
Sample $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_s)^\top$ and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_s)^\top$ independently and uniformly at random from $\mathbb{F}_q^s$, and form $P_\alpha = \sum_{i=1}^s \alpha_i B_i$ and $P_\beta = \sum_{i=1}^s \beta_i B_i$. For $\lambda \in \mathbb{F}_q$, let $U_\lambda = \ker(P_\alpha - \lambda P_\beta)$, and let $U_\infty = \ker P_\beta$. Since B_i annihilates R_{work} for all $i \in [\![1, s]\!]$, one has $R_{\text{work}} \subseteq U_\lambda$ for every $\lambda \in \mathbb{F}_q \cup \{\infty\}$, so the following quotients are well-defined. For each $\lambda \in \mathbb{F}_q \cup \{\infty\}$ with $\dim(U_\lambda/R_{\text{work}}) \in \{1, 2\}$, enumerate the directions of $U_\lambda/R_{\text{work}}$. Choose a representative $\boldsymbol{c}$ for each direction and accept it if

$$\mathcal{B}_{\boldsymbol{c}} := \{B \in \mathcal{B}_{\text{work}} : B\boldsymbol{c} = \boldsymbol{0}\}, \quad \operatorname{rank} A_{\text{work}}(\boldsymbol{c}) = 1, \quad \dim \operatorname{rad}(\mathcal{B}_{\boldsymbol{c}}) = \dim R_{\text{work}} + 1.$$

Call a candidate satisfying this condition *accepted*. Choose the first t for which $B_t\boldsymbol{c} \neq \boldsymbol{0}$, set $\boldsymbol{\varphi}_{\boldsymbol{c}} = B_t\boldsymbol{c}$, and add $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{\varphi}_{\boldsymbol{c}}\}$ to $\mathcal{R}$. The following iterative process is called *peeling*. If $\dim \mathcal{B}_{\text{work}} \geq 3$, replace $\mathcal{B}_{\text{work}}$ by $\mathcal{B}_{\boldsymbol{c}}$ and continue. Otherwise, if $\dim \mathcal{B}_{\text{work}} = 2$, then proceed to Step 3. This process recovers the $o - 3$ restricted directions associated with the interior vertices of J.

3. To recover the span of the directions associated with the chain ends, compute

$$\mathcal{C}_{K'} = \operatorname{span}_{\mathbb{F}_q}\{p^{(k)}_{[3]}(C\boldsymbol{y}) : k \in [\![1, 3]\!]\}.$$

Choose two linearly independent cubic forms $F_1, F_2 \in \mathcal{C}_{K'}$. If no such pair exists, return failure. Solve $sF_1 + tF_2 = 0$ for homogeneous linear forms $s, t \in \operatorname{span}_{\mathbb{F}_q}\{y_1, \ldots, y_o\}$. If no nonzero solution exists, return failure. For a nonzero solution (s, t), set $H_{\text{end}} = \operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{a}_s, \boldsymbol{a}_t\}$, where $\boldsymbol{a}_s, \boldsymbol{a}_t$ are the coefficient vectors of s, t. The space H_{end} is a candidate for $\operatorname{span}_{\mathbb{F}_q}\{C^\top S^\top \mathbf{L}_{j^\star-1}, C^\top S^\top \mathbf{L}_{j^\star+1}\}$.

4. For each direction $D \in \mathcal{R}$, choose a representative $\boldsymbol{r}_D \in D \setminus \{\boldsymbol{0}\}$. Construct the graph $G_{\text{rec}} = (\mathcal{R}, E_{\text{rec}})$, where

$$E_{\text{rec}} = \{\{D, D'\} : D \neq D',\ \mathsf{P}(\boldsymbol{r}_D, \boldsymbol{r}_{D'}) \in \mathcal{B}_{\text{full}}\}.$$

Require $|\mathcal{R}| = o - 3$ and G_{rec} to be a path; otherwise, return failure. Choose an orientation and label its vertices $D_3, \ldots, D_{o-1}$, setting $\boldsymbol{r}_i = \boldsymbol{r}_{D_i}$ for $i \in [\![3, o-1]\!]$. Choose a basis $\boldsymbol{h}_1, \boldsymbol{h}_2$ of H_{end}. For each $\boldsymbol{r} \in \{\boldsymbol{r}_3, \boldsymbol{r}_{o-1}\}$, compute

$$L_{\boldsymbol{r}} = \{(a, b) \in \mathbb{F}_q^2 : \mathsf{P}(a\boldsymbol{h}_1 + b\boldsymbol{h}_2, \boldsymbol{r}) \in \mathcal{B}_{\text{full}}\}.$$

Require $L_{\boldsymbol{r}_3}$ and $L_{\boldsymbol{r}_{o-1}}$ to be distinct directions in $\mathbb{F}_q^2$; otherwise, return failure. Choose nonzero $(a_2, b_2) \in L_{\boldsymbol{r}_3}$ and $(a_o, b_o) \in L_{\boldsymbol{r}_{o-1}}$, and set $\boldsymbol{r}_i = a_i\boldsymbol{h}_1 + b_i\boldsymbol{h}_2$ for $i \in \{2, o\}$. Together with $\boldsymbol{r}_3, \ldots, \boldsymbol{r}_{o-1}$, these two chain-end representatives form the completed path.

5. Choose $L \in \mathbb{F}_q^{(o-1)\times n}$ with $\mathrm{col}(L^\top) = \mathcal{O}'$. For each $i \in [\![2, o]\!]$, solve $(C^\top L^\top)\boldsymbol{\xi}_i = \boldsymbol{r}_i$ for $\boldsymbol{\xi}_i \in \mathbb{F}_q^{o-1}$. If any system is inconsistent, return failure. Otherwise, choose a solution $\boldsymbol{\xi}_i$ for each i and set $\boldsymbol{\rho}'_i = L^\top\boldsymbol{\xi}_i$.
6. To recover the affine constants, for each $i \in [\![2, o]\!]$, define $\rho'_i(\boldsymbol{x}) = (\boldsymbol{\rho}'_i)^\top\boldsymbol{x}$. Choose vectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_v$ forming a basis of $\mathcal{V}'$, define $\nu_a(\boldsymbol{x}) = \boldsymbol{v}_a^\top\boldsymbol{x}$ for $a \in [\![1, v]\!]$, and, for each k, solve the following identity for the unknowns $w_{k,i} \in \mathbb{F}_q$ for $i \in [\![2, o-1]\!]$, $\lambda_{k,i} \in \mathbb{F}_q$ for $i \in [\![2, o]\!]$, and $h_k \in \mathbb{F}_q[X_1, \ldots, X_v]_{\le 2}$:

$$q_k = \sum_{i=2}^{o-1} w_{k,i}\rho'_i\rho'_{i+1} + \sum_{i=2}^{o} \lambda_{k,i}\rho'_i + h_k(\nu_1, \ldots, \nu_v).$$

Then solve for the unknowns $\widehat{\alpha}'_2, \ldots, \widehat{\alpha}'_o \in \mathbb{F}_q$ from

$$\lambda_{k,i} = w_{k,i-1}\widehat{\alpha}'_{i-1} + w_{k,i}\widehat{\alpha}'_{i+1},$$

where $w_{k,1} = w_{k,o} = \widehat{\alpha}'_1 = \widehat{\alpha}'_{o+1} = 0$.
Return $\widehat{z}'_i = \rho'_i + \widehat{\alpha}'_i$ for all $i \in [\![2, o]\!]$. Reject the current candidate if any test in the algorithm fails.

Note 7.1. At each peeling stage of Algorithm 3, the search for $\boldsymbol{c}$ with rank $A_{\mathrm{work}}(\boldsymbol{c}) = 1$ can also be formulated as a MinRank problem. Alternatives to the randomized pencil search include linear algebra search [9], Kipnis–Shamir modeling [14], minors modeling [8], and support-minors modeling [15]. Candidates obtained by these methods must still pass the radical-dimension test in Step 2. In our experiments, these approaches were slower than the randomized pencil search.

After recovering $\widehat{z}'_2, \ldots, \widehat{z}'_o$, we construct $z_{j^\star} \circ \mathcal{S}$ and an equivalent key.

Algorithm 4 (Equivalent-key reconstruction). *Input:* $o \ge 5$, the public key $\mathcal{P}$, the space $\mathcal{V}'$, and the recovered affine polynomials $\widehat{z}'_2, \ldots, \widehat{z}'_o$.
Output: an equivalent key $(\mathcal{S}', \mathcal{F}', \mathcal{T}')$ satisfying $\mathcal{P} = \mathcal{T}' \circ \mathcal{F}' \circ \mathcal{S}'$.

1. For $i \in [\![2, o]\!]$, let $\boldsymbol{\rho}'_i$ be the coefficient vector of the homogeneous linear part of $\widehat{z}'_i$, and let $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_v$ be a basis of $\mathcal{V}'$. Complete $\boldsymbol{\rho}'_2, \ldots, \boldsymbol{\rho}'_o, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_v$ to a basis with a vector $\boldsymbol{\rho}'_1$, and set

$$R = [\boldsymbol{\rho}'_2 \ \cdots \ \boldsymbol{\rho}'_o \ \boldsymbol{\rho}'_1 \ \boldsymbol{v}_1 \ \cdots \ \boldsymbol{v}_v]^\top.$$

Set $\boldsymbol{u} = R\boldsymbol{x}$ and define $\widetilde{z}_i(\boldsymbol{u}) = \widehat{z}'_i(R^{-1}\boldsymbol{u})$.
2. To recover $z_{j^\star}$, find $\widetilde{Z}$ with $\deg\widetilde{Z} \le 2$ and $\boldsymbol{\lambda}, \boldsymbol{\mu} \in \mathbb{F}_q^o$ such that

$$\begin{aligned} &\sum_{k=1}^{o} \lambda_k p^{(k)}(R^{-1}\boldsymbol{u}) - \widetilde{z}_2\widetilde{Z}, \\ &\sum_{k=1}^{o} \mu_k p^{(k)}(R^{-1}\boldsymbol{u}) - \widetilde{Z}\widetilde{z}_o \end{aligned} \tag{9}$$

both lie in $\mathbb{F}_q[u_{o+1}, \dots, u_n]$. Let $\widetilde{Z}_0$ be a nonzero polynomial occurring in a solution to these conditions, and write $Z(\boldsymbol{x}) = \widetilde{Z}_0(R\boldsymbol{x})$.

3. To recover the equivalent key, set $S' = [\boldsymbol{\rho}'_1 \ \cdots \ \boldsymbol{\rho}'_o \ \boldsymbol{v}_1 \ \cdots \ \boldsymbol{v}_v]^\top$ and $\mathcal{S}'(\boldsymbol{x}) = S'\boldsymbol{x}$. In the coordinates $\boldsymbol{u} = \mathcal{S}'(\boldsymbol{x})$, set $z'_1 = Z \circ \mathcal{S}'^{-1}$ and $z'_i = \widehat{z}'_i \circ \mathcal{S}'^{-1}$ for $i \in [\![2, o]\!]$. For $i \in [\![1, o]\!]$, set $E_i = z'_i z'_{i+1}$, with subscripts modulo o.

4. For an unknown $T' \in \mathbb{F}_q^{o \times o}$, impose

$$\mathcal{P} \circ \mathcal{S}'^{-1} - T'(E_1, \dots, E_o)^\top \in (\mathbb{F}_q[u_{o+1}, \dots, u_n])^o .$$

Find an invertible solution T' of the resulting linear system and return failure if none exists. With this choice of T', set $\mathcal{T}'(\boldsymbol{y}) = T'\boldsymbol{y}$ and

$$(g'_1, \dots, g'_o)^\top = (T')^{-1}\big[\mathcal{P} \circ \mathcal{S}'^{-1} - T'(E_1, \dots, E_o)^\top\big] .$$

Then set $\mathcal{F}'(\boldsymbol{u}) = (E_1(\boldsymbol{u}) + g'_1(\boldsymbol{u}), \dots, E_o(\boldsymbol{u}) + g'_o(\boldsymbol{u}))^\top$, and return $(\mathcal{S}', \mathcal{F}', \mathcal{T}')$.

7.2 Correctness

Because the $o - 1$ nonseed forms $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$, $i \neq j^\star$, are assumed linearly independent on $\mathcal{O}$, their common kernel in $\mathcal{O}$ is one-dimensional. Fix a nonzero vector $\boldsymbol{w}_\star$ in this kernel. The probability that the following common-radical assumption fails is bounded in Lemma A.2.

Proposition 7.2. *Assume that*

$$\bigcap_{e \in E_J} \ker C_e = \operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{w}_\star\}. \tag{10}$$

Then $\operatorname{rad}(V) = \operatorname{span}_{\mathbb{F}_q}\{S^{-1}\boldsymbol{w}_\star\}$.

Proof. Equations (3) and (4), together with the invertibility of Θ, show that $\boldsymbol{u} \in \operatorname{rad}(V)$ if and only if $S\boldsymbol{u} \in \bigcap_{e \in E_J} \ker C_e$. The result follows from (10). □

Let $\mathcal{O}_0 = \operatorname{span}_{\mathbb{F}_q}\{S^\top \mathbf{L}_i : i \neq j^\star\}$, so that $\mathcal{D} = \operatorname{ann}(\mathcal{O}_0)$.

Theorem 7.3 (Correctness of Algorithm 2). *Assume that the nonseed forms are linearly independent on $\mathcal{O}$, that (10) holds, and that* $\operatorname{rad}(\mathcal{M}) = \operatorname{rad}(V)$. *If $\mathcal{D}_{\mathrm{sing}} = \mathcal{D}$, then Algorithm 2 recovers $(S^\top W, K, \mathcal{O}_0, \mathcal{D})$.*

Proof. The independence on $\mathcal{O}$ of the forms $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$, for $i \neq j^\star$, gives $\dim \mathcal{O}_0 = o - 1$ and $\dim \mathcal{D} = v + 1$. Their common kernel in $\mathcal{O}$ is $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{w}_\star\}$, so Proposition 7.2 gives $K \cap \mathcal{D} = \operatorname{rad}(V) = \operatorname{rad}(\mathcal{M})$.

For every $M \in \mathcal{M}$, its mask contribution annihilates K, whereas its bare-edge contribution annihilates $\mathcal{D}$. Consequently, $MK \subseteq \mathcal{O}_0$, $M\mathcal{D} \subseteq S^\top W$, and $\boldsymbol{x}^\top M \boldsymbol{d} = 0$ for all $\boldsymbol{x} \in K$ and $\boldsymbol{d} \in \mathcal{D}$. If $\mathcal{D}_{\mathrm{sing}} = \mathcal{D}$, Step 2 of Algorithm 2 gives $\mathcal{V}' \subseteq S^\top W$. The first dimension test in (8) therefore gives $\mathcal{V}' = S^\top W$, hence $K' = K$ and $\mathcal{O}' \subseteq \mathcal{O}_0$. Applying the second dimension test in (8) then gives $\mathcal{O}' = \mathcal{O}_0$, and taking annihilators gives $\mathcal{D}' = \mathcal{D}$. □

Lemma 7.4. *Under the conclusions of Theorem 7.3, let the columns of C and $L^\top$ be bases of K and $\mathcal{O}'$, respectively. Then $\mathcal{O}' \cap \mathcal{V}' = \{\mathbf{0}\}$, and $C^\top L^\top$ has full column rank.*

Proof. By Theorem 7.3, $(K', \mathcal{V}', \mathcal{O}') = (K, S^\top W, \mathrm{span}_{\mathbb{F}_q}\{S^\top \mathbf{L}_i : i \neq j^\star\})$. If $\boldsymbol{\rho} \in \mathcal{O}' \cap \mathcal{V}'$, write $\boldsymbol{\rho} = S^\top \sum_{i \neq j^\star} a_i \mathbf{L}_i$. Since $\boldsymbol{\rho} \in S^\top W$ and $S^\top$ is invertible, $\sum_{i \neq j^\star} a_i \mathbf{L}_i \in W$. Its associated linear form therefore vanishes on $\mathcal{O}$. The independence of the nonseed forms on $\mathcal{O}$ forces $a_i = 0$ for every $i \in [\![1, o]\!] \setminus \{j^\star\}$. Thus $\mathcal{O}' \cap \mathcal{V}' = \{\mathbf{0}\}$.

Since $K = \mathrm{ann}(\mathcal{V}')$ and the standard pairing is nondegenerate, $\mathrm{ann}(K) = \mathcal{V}'$. If $C^\top L^\top \boldsymbol{\xi} = \mathbf{0}$, then $L^\top \boldsymbol{\xi} \in \mathcal{O}' \cap \mathrm{ann}(K) = \mathcal{O}' \cap \mathcal{V}'$, so $\boldsymbol{\xi} = \mathbf{0}$. Therefore $C^\top L^\top$ has full column rank. □

With the oil-direction space recovered, we now establish the correctness of Algorithm 3.

Lemma 7.5 (Correctness of the peeling stage). *Under the hypotheses of Theorem 7.3, suppose Step 2 returns an accepted candidate at every peeling stage. Then the peeling process records the $o - 3$ directions associated with the interior vertices of J.*

Proof. We analyze each peeling stage in coordinates associated with the remaining secret subpath. For $i \neq j^\star$, set $\boldsymbol{d}_i = C^\top S^\top \mathbf{L}_i$. For an edge $e = \{u, v\}$ of J, restriction to K removes the mask and gives

$$D_e = C^\top S^\top C_e S C = \boldsymbol{d}_u \boldsymbol{d}_v^\top + \boldsymbol{d}_v \boldsymbol{d}_u^\top. \tag{11}$$

Hence

$$\mathcal{B}_{\mathrm{full}} = \mathrm{span}_{\mathbb{F}_q}\{D_e : e \in E_J\}. \tag{12}$$

At a peeling stage, let J_{work} denote the remaining secret subpath, with vertices $\nu_1, \ldots, \nu_r$ in order. Set $e_i = \{\nu_i, \nu_{i+1}\}$. Then

$$\mathcal{B}_{\mathrm{work}} = \mathrm{span}_{\mathbb{F}_q}\{D_{e_i} : i \in [\![1, r-1]\!]\}.$$

By Lemma 4.7, these $r - 1$ matrices are linearly independent, so $\dim \mathcal{B}_{\mathrm{work}} = r - 1$. Moreover, because $\boldsymbol{d}_{\nu_1}, \ldots, \boldsymbol{d}_{\nu_r}$ are linearly independent, (11) shows that $\boldsymbol{x} \in R_{\mathrm{work}}$ if and only if $\boldsymbol{d}_{\nu_i}^\top \boldsymbol{x} = \boldsymbol{d}_{\nu_{i+1}}^\top \boldsymbol{x} = 0$ for every $i \in [\![1, r-1]\!]$. Therefore,

$$R_{\mathrm{work}} = \bigcap_{i=1}^{r} \ker(\boldsymbol{d}_{\nu_i}^\top).$$

In particular,

$$\Psi : \mathbb{F}_q^o / R_{\mathrm{work}} \longrightarrow \mathbb{F}_q^r, \qquad \boldsymbol{x} + R_{\mathrm{work}} \longmapsto (\boldsymbol{d}_{\nu_1}^\top \boldsymbol{x}, \ldots, \boldsymbol{d}_{\nu_r}^\top \boldsymbol{x}) \tag{13}$$

is an isomorphism.

In these coordinates, we now analyze Step 2. For $r \geq 4$, it isolates a chain end, and the case $r = 3$ is treated separately below. For a candidate $\boldsymbol{c}$, write $x_i = \boldsymbol{d}_{\nu_i}^{\top}\boldsymbol{c}$. Equation (11) gives

$$D_{e_i}\boldsymbol{c} = x_{i+1}\boldsymbol{d}_{\nu_i} + x_i\boldsymbol{d}_{\nu_{i+1}}.$$

Since the B_t and the D_{e_i} are two bases of $\mathcal{B}_{\mathrm{work}}$,

$$\operatorname{rank} A_{\mathrm{work}}(\boldsymbol{c}) = \dim \operatorname{span}_{\mathbb{F}_q}\{D_{e_i}\boldsymbol{c} : i \in [\![1, r-1]\!]\}.$$

If an interior coordinate x_i were nonzero, the vectors $D_{e_{i-1}}\boldsymbol{c}$ and $D_{e_i}\boldsymbol{c}$ would be independent: the first has a nonzero $\boldsymbol{d}_{\nu_{i-1}}$-coefficient and the second a nonzero $\boldsymbol{d}_{\nu_{i+1}}$-coefficient. Rank one therefore forces all interior x_i to vanish. The only possible nonzero images are $D_{e_1}\boldsymbol{c} = x_1\boldsymbol{d}_{\nu_2}$ and $D_{e_{r-1}}\boldsymbol{c} = x_r\boldsymbol{d}_{\nu_{r-1}}$. For $r \geq 4$, the directions spanned by $\boldsymbol{d}_{\nu_2}$ and $\boldsymbol{d}_{\nu_{r-1}}$ are independent. Since $\operatorname{rank} A_{\mathrm{work}}(\boldsymbol{c}) = 1$, exactly one of x_1, x_r is nonzero. Thus every rank-one candidate is a scaled chain-end witness. Let π denote the corresponding chain end, let π' be its unique neighbor, and let $e_\pi = \{\pi, \pi'\}$. Then, for some $\gamma \in \mathbb{F}_q \setminus \{0\}$, $D_e\boldsymbol{c} = \boldsymbol{0}$ for $e \neq e_\pi$ and $D_{e_\pi}\boldsymbol{c} = \gamma\boldsymbol{d}_{\pi'}$. Writing $B_t = \sum_e \omega_{t,e}D_e$, we obtain the nonzero matrix

$$A_{\mathrm{work}}(\boldsymbol{c}) = \gamma\boldsymbol{d}_{\pi'}\begin{bmatrix}\omega_{1,e_\pi} & \cdots & \omega_{r-1,e_\pi}\end{bmatrix}. \tag{14}$$

Also $P_\alpha\boldsymbol{c}$ and $P_\beta\boldsymbol{c}$ are collinear, so $\boldsymbol{c} \in U_\lambda$ for some $\lambda \in \mathbb{F}_q \cup \{\infty\}$.

We next show that this candidate passes Step 2. Replacing $\mathcal{B}_{\mathrm{work}}$ by $\mathcal{B}_{\boldsymbol{c}}$ then removes exactly one chain-end edge. For $r \geq 4$, a combination of the D_e annihilates $\boldsymbol{c}$ exactly when its D_{e_π}-coefficient vanishes. Therefore

$$\mathcal{B}_{\boldsymbol{c}} = \operatorname{span}_{\mathbb{F}_q}\{D_e : e \in E(J_{\mathrm{work}}),\ e \neq e_\pi\}, \qquad \operatorname{col} A_{\mathrm{work}}(\boldsymbol{c}) = \operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{d}_{\pi'}\}.$$

Because e_π is the only edge incident to π, deleting it removes the condition $\boldsymbol{d}_\pi^{\top}\boldsymbol{x} = 0$, while all other conditions remain, which implies $\dim \operatorname{rad}(\mathcal{B}_{\boldsymbol{c}}) = \dim R_{\mathrm{work}} + 1$. Thus an accepted candidate deletes a chain-end edge of J_{work} and records its neighboring direction.

When $r = 3$, rank one instead means $(x_1, x_2, x_3) = (x_1, 0, x_3) \neq \boldsymbol{0}$, and both vectors $D_{e_1}\boldsymbol{c}$ and $D_{e_2}\boldsymbol{c}$ lie in $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{d}_{\nu_2}\}$. Furthermore,

$$\mathcal{B}_{\boldsymbol{c}} = \operatorname{span}_{\mathbb{F}_q}\{x_3D_{e_1} - x_1D_{e_2}\}.$$

Modulo R_{work}, the generator $x_3D_{e_1} - x_1D_{e_2}$ is represented by the matrix

$$\begin{bmatrix} 0 & x_3 & 0 \\ x_3 & 0 & -x_1 \\ 0 & -x_1 & 0 \end{bmatrix}.$$

Since $(x_1, x_3) \neq (0, 0)$, this matrix has rank two. Hence $\dim \operatorname{rad}(\mathcal{B}_{\boldsymbol{c}}) = \dim R_{\mathrm{work}} + 1$, so the candidate is accepted. It records the direction $\operatorname{span}_{\mathbb{F}_q}\{\boldsymbol{d}_{\nu_2}\}$. Accordingly, the algorithm records this direction and stops. Therefore, Step 2 recovers all $o-3$ directions associated with the interior vertices of J.

□

Theorem 7.6 (Correctness of Algorithm 3). *Under the hypotheses of Lemma 7.5, suppose Step 3 finds two linearly independent restricted cubic forms. Then Algorithm 3 recovers the nonseed coordinates $z_i \circ \mathcal{S}$, up to nonzero rescaling of each coordinate and reversal of the path order.*

Proof. By Lemma 7.4, $C^\top L^\top$ has full column rank, so every consistent system in Step 5 has a unique solution. Lemma 7.5 shows that the peeling process records all $o-3$ directions associated with the interior vertices of J. For $i \neq j^\star$, set $\boldsymbol{d}_i = C^\top S^\top \mathbf{L}_i$.

We now analyze Step 3, which recovers the span of the two chain-end directions. Let $\sigma(\boldsymbol{y}) = (z_{j^\star} \circ \mathcal{S})_{[2]}(C\boldsymbol{y})$ and $\varphi_\pm(\boldsymbol{y}) = \mathbf{L}_{j^\star \pm 1}^\top SC\boldsymbol{y}$. On K, the cubic parts of the masks vanish and the only cubic product terms are $z_{j^\star-1}z_{j^\star}$ and $z_{j^\star}z_{j^\star+1}$. Since the corresponding columns of T are independent, $\mathcal{C}_{K'} = \{\sigma\eta : \eta \in \mathrm{span}_{\mathbb{F}_q}\{\varphi_-, \varphi_+\}\}$. For the independent forms chosen in Step 3, write $F_i = \sigma\eta_i$. Their independence forces $\sigma \neq 0$ and makes η_1, η_2 independent. Since the ring is a domain, $sF_1 + tF_2 = 0$ reduces to $s\eta_1 + t\eta_2 = 0$. Since $\gcd(\eta_1, \eta_2) = 1$, $(s,t) = a(\eta_2, \eta_1)$ for some $a \in \mathbb{F}_q^\times$. Hence

$$H_{\mathrm{end}} = \mathrm{span}_{\mathbb{F}_q}\{C^\top S^\top \mathbf{L}_{j^\star-1}, C^\top S^\top \mathbf{L}_{j^\star+1}\}.$$

We now justify the graph construction in Step 4. The membership test recovers adjacency among the interior directions and then separates the two chain ends. For distinct vertices a, b, Lemma 4.7 and (12) give $\mathsf{P}(\boldsymbol{r}_a, \boldsymbol{r}_b) \in \mathcal{B}_{\mathrm{full}}$ if and only if $\{a,b\} \in E_J$, for nonzero representatives $\boldsymbol{r}_a \in \mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{d}_a\}$ and $\boldsymbol{r}_b \in \mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{d}_b\}$. Thus G_{rec} is exactly the path obtained by deleting the two ends of J. If a, b are those ends and t is the interior-path end adjacent to a, then for $\boldsymbol{h} = \mu\boldsymbol{d}_a + \nu\boldsymbol{d}_b$, $\mathsf{P}(\boldsymbol{h}, \boldsymbol{d}_t) \in \mathcal{B}_{\mathrm{full}}$ if and only if $\nu = 0$. The spaces $L_{\boldsymbol{r}_3}$ and $L_{\boldsymbol{r}_{o-1}}$ therefore separate the two end directions. After a possible reversal, Step 4 obtains a bijection $\tau : [\![2, o]\!] \to [\![1, o]\!] \setminus \{j^\star\}$ and nonzero scalars γ_i with $\boldsymbol{r}_i = \gamma_i C^\top S^\top \mathbf{L}_{\tau(i)}$.

After the path has been ordered, Step 5 reconstructs the corresponding coefficient vectors in $\mathcal{O}'$. Since $S^\top \mathbf{L}_{\tau(i)} \in \mathcal{O}'$, each $\boldsymbol{r}_i$ lies in $C^\top \mathcal{O}'$. Hence the corresponding linear system is consistent. Full column rank makes its solution unique, and $\boldsymbol{\rho}'_i = \gamma_i S^\top \mathbf{L}_{\tau(i)}$ for $i \in [\![2, o]\!]$.

It remains to verify Step 6, which restores the affine constants. Set $\beta_i = \gamma_i \alpha_{\tau(i)}$, with $\beta_1 = \beta_{o+1} = 0$. We can write

$$\begin{aligned} w_{k,i}(\rho'_i + \beta_i)(\rho'_{i+1} + \beta_{i+1}) = {} & w_{k,i}\rho'_i\rho'_{i+1} + w_{k,i}\beta_{i+1}\rho'_i \\ & + w_{k,i}\beta_i\rho'_{i+1} + w_{k,i}\beta_i\beta_{i+1}. \end{aligned}$$

Thus the coefficient of ρ'_i is

$$\lambda_{k,i} = w_{k,i-1}\beta_{i-1} + w_{k,i}\beta_{i+1}.$$

The constant terms and vinegar masks are absorbed into h_k, so $(\beta_2, \ldots, \beta_o)$ is among the solutions retained in Step 6. For this solution, $\widehat{z}'_i = \rho'_i + \beta_i = \gamma_i(z_{\tau(i)} \circ \mathcal{S})$ for $i \in [\![2, o]\!]$, which proves the claim. □

Theorem 7.6 recovers the $o-1$ nonseed coordinates up to the scalings and path ordering allowed by an equivalent key. It remains to recover the seed coordinate.

Lemma 7.7. *Under the conclusions of Theorem 7.6, the matrix R constructed in Algorithm 4 is invertible. Let*

$$Z_\star(\boldsymbol{u}) = (z_{j^\star} \circ \mathcal{S})(R^{-1}\boldsymbol{u}).$$

Then $Z_\star$ is a nonzero polynomial of degree at most two satisfying the conditions in (9).

Proof. The vectors $\boldsymbol{\rho}'_2, \dots, \boldsymbol{\rho}'_o$ form a basis of $\mathcal{O}'$, and Lemma 7.4 gives $\mathcal{O}' \cap \mathcal{V}' = \{\mathbf{0}\}$. Together with the chosen bases of $\mathcal{V}'$ and the completing vector $\boldsymbol{\rho}'_1$, this shows that the rows used to construct R are linearly independent.

With the orientation fixed in the proof of Theorem 7.6, $\widetilde{z}'_2 = \gamma_2(z_{j^\star-1} \circ \mathcal{S})$ and $\widetilde{z}'_o = \gamma_o(z_{j^\star+1} \circ \mathcal{S})$. Since T is invertible, linear combinations of the public polynomials recover $f_{j^\star-1} \circ \mathcal{S}$ and $f_{j^\star} \circ \mathcal{S}$. Writing $\widetilde{z}_i = \widetilde{z}'_i \circ R^{-1}$, we have

$$\gamma_2(f_{j^\star-1} \circ \mathcal{S})(R^{-1}\boldsymbol{u}) - \widetilde{z}_2 Z_\star = \gamma_2(g_{j^\star-1} \circ \mathcal{S})(R^{-1}\boldsymbol{u}),$$
$$\gamma_o(f_{j^\star} \circ \mathcal{S})(R^{-1}\boldsymbol{u}) - Z_\star \widetilde{z}_o = \gamma_o(g_{j^\star} \circ \mathcal{S})(R^{-1}\boldsymbol{u}).$$

The right-hand sides are vinegar-only, so $Z_\star$ satisfies (9). Its quadratic homogeneous part is nonzero, so it is a nonzero admissible choice for $\widetilde{Z}$. □

The recovered coordinates now provide the data needed to assemble an equivalent key.

Theorem 7.8 (Correctness of Algorithm 4). *If Algorithm 4 returns a triple, then the returned maps satisfy*

$$\mathcal{P} = \mathcal{T}' \circ \mathcal{F}' \circ \mathcal{S}'.$$

Proof. Since the algorithm returns only after finding an invertible solution T' to the system in Step 4, every component of $\mathcal{P} \circ \mathcal{S}'^{-1} - T'(E_1, \dots, E_o)^\top$ is vinegar-only. Thus

$$(g'_1, \dots, g'_o)^\top = (T')^{-1}\big[\mathcal{P} \circ \mathcal{S}'^{-1} - T'(E_1, \dots, E_o)^\top\big]$$

is also vinegar-only, and

$$\mathcal{P} = \mathcal{T}' \circ (E_1 + g'_1, \dots, E_o + g'_o)^\top \circ \mathcal{S}'.$$

Therefore the returned triple is a valid equivalent key. □

Appendix A bounds the failure probabilities of the oil-space assumptions.

8 Experiments

This section empirically assesses CML's security. We test parameter sets (q, o, v) targeting 128-, 192-, and 256-bit security [16]. We adjust the number of oil variables o because the CML proposal [22] requires o to be odd, whereas its parameter table lists even values. For the 128-bit target, we increase the tabulated value of o to the next odd integer, giving the tuple $(2^8, 45, 22)$ while leaving the field size q and the number of vinegar variables v unchanged. The 192- and 256-bit labels are calculated from [18, Eq. 5.17], which states that for ℓ bits of security, we need

$$\left\lceil \frac{\ell - 12.1}{2.65} \right\rceil$$

oil variables.

The results in Table 2 were obtained using a C++ implementation compiled with Apple Clang version 21.0.0 (clang-2100.1.1.101). Wall-clock running times were measured on a MacBook Pro equipped with an Apple M1 processor and running macOS 26.2. The attack recovered equivalent keys in all trials in less than 5 minutes for the tested parameters.

For context, NIST Security Level V targets security comparable to AES-256 [16]. A 2010 estimate placed an exhaustive AES-256 search on a supercomputer at $4.17 \cdot 10^{56}$ years [11]. The measured key-recovery times therefore indicate that the tested CML parameter sets do not meet their claimed security targets.

Table 2. Attack timing for 100 instances of the CML digital signature scheme for parameter sets targeting NIST security levels 1, 3, and 5.

Parameters (q, o, v)	Sec. Level	Least (s)	Average (s)	Median (s)	Most (s)
$(2^8, 45, 22)$	1	7.53	8.32	8.16	10.72
$(2^8, 69, 34)$	3	56.51	58.98	57.96	78.71
$(2^8, 93, 46)$	5	244.27	257.57	253.96	291.68

9 Conclusion

We showed that the cubic homogeneous parts of the CML public key span a space of dimension at most two. The attack cancels the cubic terms, recovers a path from the resulting rank-two matrices on the oil-direction space, and extends that path to an equivalent key. This demonstrates that hiding the locations of the cubic components is insufficient when their span has low dimension and the remaining quadratic layer retains a recoverable graph structure. Future cubic cryptosystems based on the UOV structure must account for both the dimension of the public cubic span and the structure exposed after its cancellation. Our experiments show that the tested CML parameter sets do not meet their targeted security levels.

Acknowledgments. The author thanks Professor Josh Holden of Rose-Hulman Institute of Technology for helpful discussions and guidance.

A Probability analysis

This section collects the probabilistic assumptions used by the attack.

Proposition 7.2 is used under the assumed linear independence of the restrictions $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$ on $\mathcal{O}$, for $i \neq j^\star$, and additionally assumes the common-radical equality in (10). The next two lemmas bound the two failure probabilities separately.

Lemma A.1. *Assume that the restrictions to $\mathcal{O}$ of the homogeneous linear forms $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$, for $i \neq j^\star$, are chosen independently and uniformly at random from the set of linear maps $\mathcal{O} \to \mathbb{F}_q$. Their probability of being linearly dependent is*

$$1 - \prod_{j=0}^{o-2} (1 - q^{j-o}) < \frac{1}{q(q-1)}.$$

Proof. After choosing a basis of $\mathcal{O}$, the coefficient vectors of the $o-1$ forms are independently and uniformly distributed over $\mathbb{F}_q^o$. Conditional on the first j vectors being independent, their span contains q^j vectors, so the next vector lies outside that span with probability $1 - q^{j-o}$. Multiplying over j, taking the complement, and applying the union bound gives

$$1 - \prod_{j=0}^{o-2} (1 - q^{j-o}) \leq \sum_{j=0}^{o-2} q^{j-o} = \frac{q^{-1} - q^{-o}}{q-1} < \frac{1}{q(q-1)}.$$

□

For a chain edge $e = \{i, i+1\}$, with indices read modulo o, define $G_e \in \mathbb{F}_q^{v \times v}$ by

$$(G_e)_{ab} = (\mathrm{Pol}_{g_i})_{o+a,o+b}$$

for $a, b \in [\![1, v]\!]$. Thus G_e is the lower-right vinegar–vinegar block of Pol_{g_i}. The distributional assumption on the matrices G_e below is used only to estimate the probability that the common-radical hypothesis of Proposition 7.2 fails.

Lemma A.2. *Assume that the restrictions to $\mathcal{O}$ of the linear forms $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$, for $i \neq j^\star$, are linearly independent, and that the matrices G_e are independent and uniform among the $v \times v$ alternating matrices over $\mathbb{F}_q$. Then*

$$\Pr\left(\bigcap_{e \in E_J} \ker C_e \neq \mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{w}_\star\} \right) \leq \frac{q^v - 1}{q-1} \, q^{-(o-2)(v-1)}.$$

Proof. Set

$$\mathcal{E}_{\mathrm{rad}} := \left\{ \bigcap_{e \in E_J} \ker G_e \neq \{\mathbf{0}\} \right\}.$$

We first show that failure of (10) implies $\mathcal{E}_{\mathrm{rad}}$. For $i \neq j^\star$, let $\boldsymbol{a}_i \in \mathbb{F}_q^o$ be the coefficient vector of the restriction of $\boldsymbol{w} \mapsto \mathbf{L}_i^\top \boldsymbol{w}$ to $\mathcal{O}$. Take $\boldsymbol{w} = (\boldsymbol{a}, \boldsymbol{b}) \in \bigcap_{e \in E_J} \ker C_e$ with respect to $\mathbb{F}_q^n = \mathcal{O} \oplus W$, where $\boldsymbol{a} \in \mathbb{F}_q^o$ and $\boldsymbol{b} \in \mathbb{F}_q^v$. For an edge $e = \{i, i+1\}$, the first o coordinates of $C_e \boldsymbol{w}$ are

$$(\mathbf{L}_{i+1}^\top \boldsymbol{w})\boldsymbol{a}_i + (\mathbf{L}_i^\top \boldsymbol{w})\boldsymbol{a}_{i+1},$$

because g_i is vinegar-only. The vectors $\boldsymbol{a}_i$ and $\boldsymbol{a}_{i+1}$ are independent, so $\mathbf{L}_{i+1}^\top \boldsymbol{w} = \mathbf{L}_i^\top \boldsymbol{w} = 0$. Because every vertex of J is incident to an edge, $\mathbf{L}_i^\top \boldsymbol{w} = 0$ for every $i \neq j^\star$, which gives $C_e^{\mathrm{bare}} \boldsymbol{w} = \mathbf{0}$ and $C_e \boldsymbol{w} = (\mathbf{0}, G_e \boldsymbol{b})$ for every $e \in E_J$, and hence $\boldsymbol{b} \in \bigcap_{e \in E_J} \ker G_e$.

If (10) fails, choose $\boldsymbol{w} \in \bigcap_{e \in E_J} \ker C_e$ such that $\boldsymbol{w} \notin \mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{w}_\star\}$. Then $\boldsymbol{b} \neq \mathbf{0}$, since otherwise $\boldsymbol{w} \in \mathcal{O}$ and the independence of the restricted forms would imply $\boldsymbol{w} \in \mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{w}_\star\}$. Hence

$$\boldsymbol{b} \in \left(\bigcap_{e \in E_J} \ker G_e \right) \setminus \{\mathbf{0}\},$$

so $\mathcal{E}_{\mathrm{rad}}$ occurs.

For fixed $\boldsymbol{b} \neq \mathbf{0}$, choose $P \in \mathrm{GL}_v(\mathbb{F}_q)$ such that $P\boldsymbol{e}_1 = \boldsymbol{b}$, and set $\widetilde{G}_e = P^\top G_e P$. Then $G_e \boldsymbol{b} = \mathbf{0}$ if and only if $\widetilde{G}_e \boldsymbol{e}_1 = \mathbf{0}$. The map $G \mapsto P^\top G P$ is a linear bijection on the space of alternating matrices, with inverse $H \mapsto P^{-\top} H P^{-1}$, so it preserves the uniform distribution on that space. The first column of the uniformly distributed alternating matrix $\widetilde{G}_e$ has zero diagonal entry and $v-1$ independent uniform off-diagonal entries. Hence $G_e \boldsymbol{b} = \mathbf{0}$ imposes $v-1$ independent linear conditions and holds with probability $q^{-(v-1)}$. Independence of the $o-2$ matrices gives probability $q^{-(o-2)(v-1)}$ for a fixed direction $\mathrm{span}_{\mathbb{F}_q}\{\boldsymbol{b}\}$. There are $(q^v - 1)/(q-1)$ directions in $\mathbb{F}_q^v$, and therefore

$$\begin{aligned} \Pr(\mathcal{E}_{\mathrm{rad}}) &\leq \sum_{\substack{D \subseteq \mathbb{F}_q^v \\ \dim D = 1}} \Pr\left(D \subseteq \bigcap_{e \in E_J} \ker G_e \right) \\ &= \frac{q^v - 1}{q - 1}\, q^{-(o-2)(v-1)}. \end{aligned}$$

Since failure of the common-radical equality implies $\mathcal{E}_{\mathrm{rad}}$, the estimate follows. □

References

1. Beullens, W.: Improved cryptanalysis of UOV and Rainbow. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12696, pp. 348–373. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_13.

2. Coppersmith, D., Stern, J., Vaudenay, S.: Attacks on the birational permutation signature schemes. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 435–443. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_37.
3. Coppersmith, D., Stern, J., Vaudenay, S.: The security of the birational permutation signature schemes. Journal of Cryptology 10(3), 207–221 (1997). https://doi.org/10.1007/s001459900028.
4. Ding, J., Wolf, C., Yang, B.-Y.: ℓ-Invertible cycles for multivariate quadratic (MQ) public key cryptography. In: Okamoto, T., Wang, X. (eds.) PKC 2007. LNCS, vol. 4450, pp. 266–281. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71677-8_18.
5. Ding, J., Yang, B.-Y., Chen, C.-H.O., Chen, M.-S., Cheng, C.-M.: New differential-algebraic attacks and reparametrization of Rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 242–257. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68914-0_15.
6. Ding, J., Zhang, Z., Deaton, J., Schmidt, K., Vishakha, FNU: New attacks on lifted unbalanced oil vinegar. In: Second PQC Standardization Conference. NIST (2019). https://csrc.nist.gov/CSRC/media/Events/Second-PQC-Standardization-Conference/documents/accepted-papers/ding-new-attacks-luov.pdf.
7. Duong, D.H., Petzoldt, A., Wang, Y., Takagi, T.: Revisiting the cubic UOV signature scheme. In: Hong, S., Park, J.H. (eds.) ICISC 2016. LNCS, vol. 10157, pp. 223–238. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-53177-9_12.
8. Faugère, J.-C., Levy-dit-Vehel, F., Perret, L.: Cryptanalysis of MinRank. In: Wagner, D.A. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 280–296. Springer (2008). https://doi.org/10.1007/978-3-540-85174-5_16.
9. Goubin, L., Courtois, N.T.: Cryptanalysis of the TTM cryptosystem. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 44–57. Springer (2000). https://doi.org/10.1007/3-540-44448-3_4.
10. Hashimoto, Y.: On the security of Cubic UOV and its variants. Cryptology ePrint Archive, Paper 2016/788 (2016). https://eprint.iacr.org/2016/788.
11. Kaminsky, A., Kurdziel, M., Radziszowski, S.: An overview of cryptanalysis research for the Advanced Encryption Standard. In: MILCOM 2010, pp. 1310–1316. IEEE (2010). https://doi.org/10.1109/MILCOM.2010.5680130. https://www.cs.rit.edu/~spr/gdn2010/aesark.pdf.
12. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 206–222. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_15.
13. Kipnis, A., Shamir, A.: Cryptanalysis of the oil and vinegar signature scheme. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 257–266. Springer, Heidelberg (1998). https://doi.org/10.1007/BFb0055733.
14. Kipnis, A., Shamir, A.: Cryptanalysis of the HFE public key cryptosystem by relinearization. In: CRYPTO 1999. LNCS, vol. 1666, p. 788. Springer (1999).
15. Bardet, M., Bros, M., Cabarcas, D., Gaborit, P., Perlner, R.A., Smith-Tone, D., Tillich, J., Verbel, J.A.: Improvements of algebraic attacks for solving the rank decoding and MinRank problems. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part I. LNCS, vol. 12491, pp. 507–536. Springer (2020). https://doi.org/10.1007/978-3-030-64837-4_17.
16. National Institute of Standards and Technology: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process (2016). https://csrc.nist.gov/CSRC/media/Projects/Post-Quantum-Cryptography/documents/call-for-proposals-final-dec-2016.pdf.

17. Nie, X., Liu, B., Xiong, H., Lu, G.: Cubic unbalance oil and vinegar signature scheme. In: Lin, D., Wang, X., Yung, M. (eds.) Inscrypt 2015. LNCS, vol. 9589, pp. 47–56. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-38898-4_3.
18. Petzoldt, A.: Selecting and reducing key sizes for multivariate cryptography. Ph.D. thesis, Technische Universität Darmstadt (2013). https://doi.org/10.26083/tuprints-00003523. https://tuprints.ulb.tu-darmstadt.de/3523/.
19. Shamir, A.: Efficient signature schemes based on birational permutations. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 1–12. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_1.
20. Shim, K.-A., Koo, N., Park, C.-M.: Security analysis of improved Cubic UOV signature schemes. In: Kim, H., Kim, D.-C. (eds.) ICISC 2017. LNCS, vol. 10779, pp. 310–324. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78556-1_17.
21. Szepieniec, A., Preneel, B.: Block-anti-circulant unbalanced oil and vinegar. In: Paterson, K.G., Stebila, D. (eds.) SAC 2019. LNCS, vol. 11959, pp. 574–588. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-38471-5_23.
22. Thomas, I., Zhang, N., Zhang, Z.: A cubic multivariate public key signature scheme with the ℓ-invertible cycle system. In: Proceedings of the 2026 5th International Conference on Cryptography, Network Security and Communication Technology (CNSCT 2026), pp. 7–11. ACM, New York (2026). https://doi.org/10.1145/3802927.3802929.