UFOs: A Very Efficient Multivariate Public Key Signature Scheme

Gilles Macario-Rat[1][0000−0002−6156−8295]

Orange Research, Châtillon, FRANCE
gilles.macariorat@orange.com

Abstract. We present UFOs, a multivariate public-key signature scheme in the Unbalanced Oil and Vinegar (UOV) family. The scheme replaces generic quadratic polynomials with a structured subclass based on Frobenius-type quadratic forms, yielding a compressed public-key representation while retaining the efficient UOV signing procedure. We describe the key-generation, signing, and verification algorithms, and we detail the derivation of the public system from a compact secret description. We discuss security in the standard multivariate setting, including direct algebraic attacks and key-recovery approaches, and we formalize the underlying computational problems induced by the proposed structure. Finally, we report implementation results quantifying the costs of key generation, signing, and verification, as well as the resulting public-key and signature sizes.

Keywords: UOV · post-quantum signatures · multivariate cryptography · Frobenius forms · quadratic forms

1 Introduction

Modern digital infrastructures rely on digital signatures; post-quantum designs based on multivariate polynomial systems remain attractive due to fast verification and conservative hardness assumptions. A canonical representative is Unbalanced Oil and Vinegar (UOV) [34], where signing inverts a structured MQ map using an oil–vinegar trapdoor.

Contributions. We introduce UFOs (**U**nbalanced **F**robenius **O**il and Vinegar with **S**chedule), a UOV-type signature scheme whose public key is specified by a small number of *Frobenius blocks* over $\mathbb{F}_q$, i.e., sparse sums of monomials $x^{p^a}y^{p^b}$. This reduces public-key size while preserving the standard UOV signing procedure (vinegar sampling followed by an $\mathbb{F}_p$-linear solve in the oil variables).

This version updates [24] in light of the exponent-only attack of [2] by formalizing schedule-selection criteria and instantiating hardened exponent schedules.

We provide (i) a complete specification, (ii) a concrete security analysis against the main known attack families, and (iii) parameter sets with proof-of-concept performance measurements.

Technical overview. The scheme works over an odd-characteristic base field and an extension $\mathbb{F}_q = \mathbb{F}_{p^e}$, and uses homogeneous public maps. After descent to $\mathbb{F}_p$, verification remains evaluation of an MQ system, while signing reduces to solving an $oe \times me$ linear system over $\mathbb{F}_p$ (square when $m = o$). The main design parameter beyond (o, v, m, p, e) is the exponent schedule $\{(a_k, b_k)\}_{k=1}^{\ell} \subset (\mathbb{Z}/e\mathbb{Z})^2$, which governs the resistance to direct low-dimensional algebraic models.

2 Preliminaries

2.1 Notation

Throughout the paper, Roman letters (e.g., a, b, c, p, q) denote integers, Greek letters (e.g., $\alpha, \xi, \Lambda, \Omega$) denote finite-field elements, and capital letters denote vectors or matrices. We reserve x, y for scalar variables and X, Y for vectors of variables.

For a positive integer n, we write 1_n for the $n \times n$ identity matrix. For a square matrix M of size ℓ, we set

$$M_{\otimes n} := 1_n \otimes M,$$

the Kronecker product of 1_n and M; it is a block-diagonal matrix of size $n\ell$ with n copies of M on the diagonal.

For a row vector X, ${}^tX$ denotes its transpose (a column vector), and for a matrix M, ${}^tM$ denotes its transpose.

2.2 Finite fields

We work with a prime field $\mathbb{F}_p$, where p is an odd prime, and an extension field $\mathbb{F}_q$ with $q = p^e$, where $e > 2$ is prime. The primeness of e ensures that $\mathbb{F}_q/\mathbb{F}_p$ has no intermediate extensions.

We choose p close to a power of 2 so that sequences of elements of $\mathbb{F}_p$ can be efficiently packed into bit arrays.

An element of $\mathbb{F}_q$ is represented via $\mathbb{F}_q = \mathbb{F}_p[\xi]/(i(\xi))$, where $i(\xi)$ is an irreducible polynomial of degree e. The coefficient vector in the basis $(1, \xi, \ldots, \xi^{e-1})$ identifies $\mathbb{F}_q$ with the vector space $\mathbb{F}_p^e$ through the isomorphism $\Phi : \mathbb{F}_q \to \mathbb{F}_p^e$ defined by

$$\Phi\Big(\sum_{i=1}^{e} x_i \xi^{i-1}\Big) = (x_1, \ldots, x_e).$$

When the context is clear, we use bold letters for elements of $\mathbb{F}_p^e$, e.g., $\mathbf{x} = (x_1, \ldots, x_e)$. For efficient extension-field arithmetic, see, e.g., [30].

We extend Φ componentwise to $\mathbb{F}_q^n$ by

$$\Phi_{\otimes n}(X) = (\Phi(X_1), \ldots, \Phi(X_n)) \in \mathbb{F}_p^{ne},$$

and write $\bar{\Phi}$ for $\Phi_{\otimes n}$ when there is no ambiguity.

2.3 Linearity, bilinearity, and quadraticity

We briefly recall standard notions used in multivariate cryptography.

Linear and quadratic forms in n variables over a field $\mathbb{F}$ are multivariate polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ of total degree 1 and 2, respectively. Their homogeneous parts can be written as $\sum_{i=1}^{n} \alpha_i x_i$ and $\sum_{i,j=1}^{n} \alpha_{ij} x_i x_j$, with coefficients in $\mathbb{F}$.

Linearity is, by default, with respect to all variables, but it is often useful to consider linearity in only a subset. For example, xy is quadratic in (x, y), yet linear in x and linear in y (bilinear).

A (linear or quadratic) (n, m)-map is a collection of m forms in n variables. These notions interact naturally with field extensions.

For instance, consider the $(2, 1)$-$\mathbb{F}_q$ quadratic (equivalently, bilinear) map $f(x, y) = xy$ (multiplication in $\mathbb{F}_q$). It can be represented as a $(2e, e)$-$\mathbb{F}_p$ quadratic map via

$$\mathbf{f}(\mathbf{x}, \mathbf{y}) = \Phi\big(f(\bar{\Phi}^{-1}(\mathbf{x}, \mathbf{y}))\big).$$

We also extend Φ to linear and quadratic maps and use the usual composition rules. For the multiplication map $(x, y) \mapsto xy$, write $x\cdot : y \mapsto xy$ and $\cdot y : x \mapsto xy$. If $x = \sum_{i=1}^{e} x_i \xi^{i-1}$ and $y = \sum_{i=1}^{e} y_i \xi^{i-1}$, then $xy = \sum_{i,j=1}^{e} x_i y_j \, \xi^{i-1} \xi^{j-1}$. For $i, j \in \{1, \ldots, e\}$, write $\xi^{i-1}\xi^{j-1} = \sum_{k=1}^{e} \xi_{ijk} \, \xi^{k-1}$. Let $\mathbb{M}_k \in \mathbb{F}_p^{e \times e}$ be the matrix with entry (i, j) equal to ξ_{ijk}. Then

$$\Phi(xy) = (\Phi(x)\mathbb{M}_1 \, {}^t\Phi(y), \ldots, \Phi(x)\mathbb{M}_e \, {}^t\Phi(y)).$$

Equivalently,

$$\begin{aligned}
\Phi(xy) &= \Phi(x)\Phi(\cdot 1\cdot) \, {}^t\Phi(y) = \Phi(x, \cdot) \, {}^t\Phi(y) = \Phi(x)\Phi(\cdot y), \\
\Phi(x\cdot) &= \Phi(x)\Phi(\cdot 1\cdot) = (\Phi(x)\mathbb{M}_1, \ldots, \Phi(x)\mathbb{M}_e), \\
\Phi(\cdot y) &= \Phi(\cdot 1\cdot) \, {}^t\Phi(y) = (\mathbb{M}_1 \, {}^t\Phi(y), \ldots, \mathbb{M}_e \, {}^t\Phi(y)).
\end{aligned}$$

Linear, bilinear, and quadratic maps can be represented by matrices. The isomorphism Φ also extends to matrices over $\mathbb{F}_q$. Let $M = (\mu_{ij}) \in \mathbb{F}_q^{n \times m}$. For each $k \in \{1, \ldots, e\}$, define a matrix $M_k \in \mathbb{F}_p^{en \times em}$ as an $n \times m$ block matrix with $e \times e$ blocks, where the (i, j) block is the $\mathbb{F}_p$-matrix representing the map $z \mapsto \mu_{ij} z$ on $\mathbb{F}_q$ in the basis $(1, \xi, \ldots, \xi^{e-1})$, composed with projection onto the k-th coordinate (equivalently, the block is $\Phi(\cdot \mu_{ij}) \, \mathbb{M}_k$).

With this notation, for $X \in \mathbb{F}_q^n$ and $Y \in \mathbb{F}_q^m$, the left-right multiplication by M can be defined as $\Phi(\cdot M \cdot)$ where :

$$\begin{aligned}
\Phi(XM \, {}^tY) &= \Phi_{\otimes n}(X)\Phi(\cdot M\cdot) \, {}^t\Phi_{\otimes m}(Y) \\
&= \Phi_{\otimes n}(X)(M_1, \ldots, M_e) \, {}^t\Phi_{\otimes m}(Y) \\
&= \big(\Phi_{\otimes n}(X)M_1 \, {}^t\Phi_{\otimes m}(Y), \ldots, \Phi_{\otimes n}(X)M_e \, {}^t\Phi_{\otimes m}(Y)\big)
\end{aligned}$$

In the same manner, right-multiplication by M induces an $\mathbb{F}_p$-linear map $\mathbb{F}_p^{en} \to \mathbb{F}_p^{em}$, i.e., there exists $\Phi(\cdot M) \in \mathbb{F}_p^{en \times em}$ such that

$$\Phi_{\otimes m}(XM) = \Phi_{\otimes n}(X)\Phi(\cdot M).$$

For a map f, we define its differential by $\Delta_f(x,y) = f(x+y) - f(x) - f(y) + f(0)$. If f is quadratic, then Δ_f is bilinear (it is the polar form of f).

For a bilinear map f represented by a matrix M, i.e., $f(x,y) = xM\,{}^t y$, we define its symmetric and alternating parts by $f^+(x,y) = (f(x,y)+f(y,x))/2$ and $f^-(x,y) = (f(x,y) - f(y,x))/2$. These correspond to the symmetric and skew-symmetric matrices $M^+ = (M + {}^tM)/2$ and $M^- = (M - {}^tM)/2$, respectively. We assume $\mathrm{char}(\mathbb{F}) \neq 2$ here.

2.4 Frobenius automorphisms

A central tool for field extensions is the Frobenius automorphism. For $\mathbb{F}_q/\mathbb{F}_p$ with $q = p^e$, define $\phi : \mathbb{F}_q \to \mathbb{F}_q$ by $\phi(x) = x^p$. The map ϕ is $\mathbb{F}_p$-linear. Writing $x = \sum_{i=1}^{e} x_i \xi^{i-1}$, we have $\phi(x) = \sum_{i=1}^{e} x_i\, \phi(\xi^{i-1})$. Therefore, the matrix $\mathcal{F} \in \mathbb{F}_p^{e\times e}$ whose i-th column is

$$\Phi\big(\phi(\xi^{i-1})\big) = \Phi(\xi^{(i-1)p}), \quad i = 1, \ldots, e,$$

satisfies $\Phi\big(\Phi^{-1}(\mathbf{x})^p\big) = \mathbf{x}\mathcal{F}$ for all $\mathbf{x} \in \mathbb{F}_p^e$.

The map ϕ is not $\mathbb{F}_q$-linear. It generates the cyclic group of field automorphisms of $\mathbb{F}_q$ that fix $\mathbb{F}_p$, i.e., the Galois group $\mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p)$.

For a vector X (resp. a matrix M), we write $X^{[i]}$ (resp. $M^{[i]}$) for the result of raising each entry to the power p^i. Indices i are always taken modulo e.

2.5 UOV scheme and its trapdoor

We now give a short overview of the UOV scheme; we refer to the specification document [34] for full details.

UOV is a trapdoor construction for multivariate quadratic (MQ) systems. It starts from a "central" quadratic map $F = (F_1, \ldots, F_m) : \mathbb{F}^n \to \mathbb{F}^m$ in variables $(x_1, \ldots, x_n)$, with $n = v + m$ split into v "vinegar" variables and m "oil" variables. The defining property is that F contains no oil–oil quadratic terms: once the vinegar variables are fixed, each F_r becomes an affine function of the oil variables.

The public key is the disguised system $\mathbf{Pub} = F \circ S$, where $S : \mathbb{F}^n \to \mathbb{F}^n$ is a secret invertible affine transformation. Given a target $h \in \mathbb{F}^m$, inverting $\mathbf{Pub}(x) = h$ without the secret is assumed hard.

To use the trapdoor, one first chooses random values for the vinegar variables, then solves the resulting linear system in the oil variables to obtain u such that $F(u) = h$. Finally, one outputs $x = S^{-1}(u)$, which satisfies $\mathbf{Pub}(x) = h$ by construction.

3 Design Rationale

Main design idea The objective is to preserve the UOV trapdoor paradigm while enabling a compact public-key representation. Inspired by recent structured MQ signature proposals (e.g., SNOVA, QR-UOV, MAYO), the public map is

constrained to a family that can be generated from a small amount of randomness. Note that in an information-theoretic sense such keys are *not* distributed like fully random MQ instances: key generation produces (typically overdetermined) systems with a guaranteed hidden solution (the oil space). The relevant question is whether this structural bias yields an *efficient* distinguisher or can be turned into practical key-recovery/forgery attacks. However, if such a distinguisher could be constructed, it could be transformed into a procedure that distinguishes between univariate polynomials that possess roots and those that do not. This appears highly implausible, given that (for several natural input models) the general problem of determining the existence of a root is NP-hard [5].

Shrinking multivariate public keys has been a central theme in recent proposals [25], including MAYO [4], QR-UOV [18], and SNOVA [23]. We work with homogeneous equations over an odd-characteristic base field. This avoids characteristic-two subtleties, where the Frobenius map is $\mathbb{F}_2$-linear and quadratic/bilinear representations behave differently [22,13]. More generally, both the base-field characteristic and the choice between homogeneous and affine public maps affect design and cryptanalysis [35,11].

Within the UOV family, UFOs offers a favorable efficiency–security trade-off:

- The secret oil space has dimension 1 (generated by a single vector), which limits the direct applicability of the "one-vector" viewpoint of Pébereau [26].
- A public key typically (with overwhelming probability) determines a unique secret key up to the standard equivalences; equivalently, the defining systems for the secret are typically overdetermined (see our toy experiments in Appendix C).
- The central security tasks (key recovery and forgery) reduce to two closely related problem families over an extension field:
  - Solve the univariate equation

    $$\sum_{i=1}^{\ell} \alpha_i x^{p^{a_i}+p^{b_i}} = \gamma,$$

    where ℓ is small (e.g., $2 \le \ell \le 5$) and the exponents may be arbitrary.
  - Solve the related multivariate system of ℓ equations in n variables (typically $n \approx \ell$) and the exponents may be arbitrary:

    $$\sum_{j=1}^{n} \alpha_{ij} x_j^{p^{a_i}+p^{b_i}} = 0, \quad i = 1, \ldots, \ell.$$

  Up to our knowledge, and for well chosen exponents, these problems do not admit polynomial-time algorithms; consequently, we conjecture hardness and provide supporting estimates under standard heuristics (e.g., random-equations models, dense-MQ behavior after descent, and generic linear-algebra costs).
- The scheme depends on parameters o (oil variables), v (vinegar variables), p (characteristic/field size), e (extension degree), and ℓ (Frobenius sparsity limit). Additionally, the scheme requires ol pairs of Frobenius exponents, integers interpreted modulo e. Typical sizes are:

- Signature: $(o+v)e\log_2(p)$ bits,
- Public key (compressed): $o^3\ell e\log_2(p)$ bits,
- Public key (uncompressed): $(o+v)^2o\ell e\log_2(p)$ bits,
- Secret key (compressed): $O(\lambda)$ bits,
- Secret key (uncompressed): $(o+v)^2o\ell e\log_2(p)$ bits,
- Security: $e\log_2(p)/2$ bits (heuristic baseline). Rationale: the tightest generic baseline attacks we model scale like $p^{e/2}=\sqrt{q}$ field operations, hence correspond to about $\frac{1}{2}\log_2(q)=\frac{1}{2}e\log_2(p)$ bits. *Concrete* security for each parameter set is then taken as the minimum of the modeled attack costs reported in Tab. 2.

Once o, v, and ℓ are fixed, the main tuning knob is $e\log_2(p)$ (or simply e[1]). Doubling e doubles security and scales public-key and signature sizes proportionally.

– For comparison, SNOVA [23] uses parameters o (oil), v (vinegar), p (field size), and l (extension degree of a non-commutative ring), with representative sizes:
  - Signature: $(o+v)l^2\log_2(p)$,
  - Public key (compressed): $o^3l^2\log_2(p)$,
  - Secret key (uncompressed): $\big((v+2o+1)\,vo\,l^2+4l^4\big)\log_2(p)$,
  - Security: $O\big((vo^2l^4)^{1/3}\big)$ (empirical[2] fit to the parameters in [23]).

  Under this model, doubling l^2 doubles signature and public-key sizes, more than doubles secret-key size, and increases security only by a factor $2^{2/3}\approx 1.6$.
– Our choice $o=m=1$, $v=4$ yields $n/m=n/o=5$, which is atypical for UOV. Conservative UOV instantiations usually take this ratio in $[3,4]$ [3]; ratios near 2.5 are considered aggressive, and ratios above 4 are often viewed as inefficient [26]. Our results indicate that ratio 5 can remain practical in the present structured setting.
– The Frobenius structure yields compact, high-degree equations. In the extension-field representation, degrees of the form p^a+p^b are extremely large (typical value: $31^{30}\approx 2^{150}$). While high degree alone does not imply hardness, we are not aware of direct attacks in the $\mathbb{F}_q$ representation that exploit this structure more efficiently than descending to the base field $\mathbb{F}_p$; see Sec. 4.2.
– Public-key compression can be viewed as follows. Over $\mathbb{F}_p$, representing e generic quadratic equations in e variables (e.g., $\sum\alpha_{kij}x_ix_j$) costs $e\frac{e(e+1)}{2}$ coefficients in $\mathbb{F}_p$. Over $\mathbb{F}_p^e$, the corresponding representation is a single equation of the form $\sum\alpha_{ij}X^{p^i}X^{p^j}$, which costs $\frac{e(e+1)}{2}$ coefficients in $\mathbb{F}_{p^e}$ (and hence the same total information). UFOs reduces the number of monomials to $\sum_{i=1}^{\ell}\alpha_iX^{p^{a_i}}X^{p^{b_i}}$ with $\ell\ll\frac{e(e+1)}{2}$; we show that this sparsification does not compromise security.

[1] For smooth asymptotics, take, e.g., $p=O(e)$; for concrete parameters, choose p to balance efficiency and complexity.

[2] We searched among simple monomials in v, o, and l with small integral exponents for the best fit.

Frobenius forms. The field extension is exploited by restricting each public equation to a sum of *Frobenius forms*, generalizing Frobenius-twisted quadratic forms and Hermitian Frobenius forms; see [29,20]. For integers a, b (interpreted modulo e), an (a, b)-Frobenius form is a polynomial in $\mathbb{F}_q[x_1, \ldots, x_n]$ whose monomials have the shape $x_i^{p^a} x_j^{p^b}$. Each monomial has total degree $p^a + p^b$, so this class strictly contains ordinary quadratic forms as a special case ($a = b = 0$). Over $\mathbb{F}_p$-coordinates, these maps become quadratic in the usual MQ sense; see Sec. 4.2. The construction focuses on $a \neq b$, since $a = b$ yields additional symmetry that is not desirable here.

A crucial property is invariance under invertible $\mathbb{F}_q$-linear changes of variables: composing an (a, b)-Frobenius form with an invertible linear map over $\mathbb{F}_q$ produces another (a, b)-Frobenius form. (Indeed, $(XL)^{[a]} = X^{[a]}L^{[a]}$ entry-wise, so $(XL)^{[a]}A\,{}^t(XL)^{[b]} = X^{[a]}(L^{[a]}A\,{}^tL^{[b]})\,{}^tX^{[b]}$.) Consequently, Frobenius forms are compatible with the standard UOV masking technique.

Public equations and secret masking. Let ℓ denote an upper bound on the number of Frobenius blocks per public equation. The public key is given by a map

$$\begin{aligned} \mathbf{Pub} &= (\mathbf{Pub}_1, \ldots, \mathbf{Pub}_m) : \mathbb{F}_q^n \to \mathbb{F}_q^m \\ \mathbf{Pub}_r(x_1, \ldots, x_n) &= \sum_{1 \le k \le \ell,\ 1 \le i,j \le n} \alpha_{rkij}\, x_i^{p^{a_{rk}}} x_j^{p^{b_{rk}}} \\ &= \sum_{1 \le k \le \ell} X^{[a_{rk}]} A_{rk}\, {}^tX^{[b_{rk}]} \quad 1 \le r \le m, \end{aligned}$$

where $X = (x_1, \ldots, x_n)$, $A_{rk} = (\alpha_{rkij})_{1 \le i,j \le n} \in \mathbb{F}_q^{n \times n}$, and $a_{rk}, b_{rk} \in \mathbb{Z}/e\mathbb{Z}$.

As in UOV, the public and secret systems are related by an invertible $\mathbb{F}_q$-linear map $L \in \mathbb{F}_q^{n \times n}$ such that

$$F(X) = \mathbf{Pub}(XL).$$

Then each secret equation is again a sum of Frobenius forms,

$$F_r(x_1, \ldots, x_n) = \sum_{1 \le k \le \ell} X^{[a_{rk}]} F_{rk}\, {}^tX^{[b_{rk}]} \quad 1 \le r \le m,$$

with the induced matrix relation

$$F_{rk} = L^{[a_{rk}]} A_{rk}\, {}^tL^{[b_{rk}]} \quad 1 \le r \le m,\ 1 \le k \le \ell.$$

Proof : $\mathbf{Pub}_{rk}(XL) = (XL)^{[a_{rk}]} A_{rk}\, {}^t(XL)^{[b_{rk}]} = X^{[a_{rk}]}(L^{[a_{rk}]} A_{rk}\, {}^tL^{[b_{rk}]})\, {}^tX^{[b_{rk}]}$.

UOV trapdoor constraint. The secret system is required to satisfy the UOV condition, i.e., it contains no oil–oil monomials. With our $(o \mid v)$ split (oil variables first), the oil–oil coefficients are exactly the top-left $o \times o$ block of each matrix F_{rk}. Equivalently, for all $1 \le s, t \le o$ one imposes $(F_{rk})_{st} = 0$, which yields

$$\sum_{1 \le i,j \le n} \alpha_{rkij}\, \lambda_{si}^{p^{a_{rk}}} \lambda_{tj}^{p^{b_{rk}}} = 0 \quad 1 \le s, t \le o,\ 1 \le r \le m,\ 1 \le k \le \ell, \tag{1}$$

where $L = (\lambda_{ij})_{1 \le i,j \le n}$.

Using standard UOV arguments [27], one may assume (up to the usual UOV equivalences given by invertible linear changes of variables and a change of basis of the oil space) that L has the block form

$$L = \begin{pmatrix} 1_o & R_{o \times v} \\ 0_{v \times o} & 1_v \end{pmatrix}, \tag{2}$$

where R is random and the indices indicate block dimensions.

Exponent schedule rationale. A central design parameter in UFOs is the choice of the five[3] exponent pairs $(a_k, b_k) \in (\mathbb{Z}/e\mathbb{Z})^2$, $k = 1, \ldots, \ell$, that define the Frobenius blocks

$$\mathrm{Pub}_k(X) = X^{[a_k]} A_k \, {}^t X^{[b_k]}.$$

While the public-key format, the signing algorithm, and the trapdoor structure are unchanged for any such choice, the concrete resistance to direct algebraic forgery attacks depends strongly on the *relative position* of these exponents on the cyclic index set $\mathbb{Z}/e\mathbb{Z}$. This is because several direct attacks work by cancelling some public blocks and then exploiting a low-complexity representation of the remaining indices, either as a sparse univariate equation or as a low-bidegree bivariate system.

Why the exponents matter. After cancelling two blocks (see Hybrid Attack in Sec. 4.4, in the regime $\ell = 5$), one is left with three residual Frobenius terms. A first attack route is *univariate*: after a common Frobenius shift, the residual equation becomes a sparse polynomial in one unknown over $\mathbb{F}_{p^e}$. If the six residual endpoints

$$a_i, b_i, a_j, b_j, a_t, b_t$$

for the surviving blocks lie inside a short cyclic interval of $\mathbb{Z}/e\mathbb{Z}$, then the resulting univariate degree is comparatively small, which reduces the cost of root finding. Hence exponent schedules should avoid configurations for which, after cancelling two blocks, the remaining six indices are concentrated in a short arc of the cycle.

A second route is *bivariate*. Its basic idea is to rewrite the same residual three-block equation using two Frobenius coordinates

$$X = x^{[r]}, \qquad Y = x^{[s]},$$

for suitable anchors $r, s \in \mathbb{Z}/e\mathbb{Z}$. Each residual block may be represented in one of two orientations, according to whether a_k is assigned to the X-side and b_k to the Y-side, or vice versa. Combining the original residual equation with one Frobenius-shifted copy may produce two equations of low bidegrees in (X, Y). Eliminating one variable by a resultant then yields a univariate equation of degree governed by these bidegrees, and if this degree is too small, one obtains a practical direct forgery procedure. Therefore, one must also avoid exponent schedules for which some residual triplet admits a low-bidegree two-anchor representation.

[3] We consider here the most convenient choice of parameters where $\ell = 5$, $m = o = 1$.

A brief description of the bivariate attack. [2] Suppose two blocks have been canceled and three blocks remain. The attacker chooses:

- a nonzero Frobenius shift σ,
- two anchor indices r, s,
- and, for each of the three residual blocks, one orientation in the original equation and one orientation in the shifted equation.

This gives two representations of the residual relation as two bi-variate equations in variables $X = x^{[r]}$ and $Y = x^{[s]}$. If the corresponding bi-degrees are

$$(d_{x,1}, d_{y,1}) \qquad \text{and} \qquad (d_{x,2}, d_{y,2}),$$

then eliminating one variable leads to a resultant whose degree is controlled by

$$R = \max(d_{x,1} + d_{y,2},\ d_{x,2} + d_{y,1}),$$

possibly with a further saving coming from the $\mathbb{F}_p^*$-scaling symmetry: it exploits the invariance $(X, Y) \mapsto (\lambda X, \lambda^{-1} Y)$ for $\lambda \in \mathbb{F}_p^*$ by switching to invariants $Z = X^{p-1}$ and $T = XY$, reducing the concrete resultant degree. Thus small values of these degree parameters directly translate into cheaper forgery attacks. In particular, schedules with residual triplets admitting two "well-aligned" representations around two anchors must be avoided.

Design principle: evaluate residual triplets. This motivates a schedule-selection methodology based on the ten possible *residual triplets*, obtained by canceling any two out of the five blocks. For each residual triplet, we evaluate two complementary quantities.

First, we consider the *univariate span* L: this is the length of the smallest cyclic interval containing the six residual endpoints. A large value of L means that no common Frobenius shift can compress the triplet into a short univariate window.

Second, we consider the best *two-anchor bivariate model*. For every choice of shift σ, orientations, and anchors (r, s), we compute the two bidegrees induced by the residual indices and derive the corresponding resultant-based score. The attacker will of course choose the most favorable such representation, so the relevant defensive quantity is the *minimum* score over all these choices; good schedules maximize this minimum.

Geometric interpretation. These criteria admit a simple geometric interpretation on the cycle $\mathbb{Z}/e\mathbb{Z}$. Each block contributes an ordered pair of endpoints (a_k, b_k), and after two cancellations the attacker sees three such pairs, i.e., six marked points on the cycle.

The univariate criterion asks whether these six points fit inside one short arc. If they do, then one obtains a low-degree one-variable description.

The bivariate criterion asks whether the same six points can be organized around two centers (the anchors r and s), after possibly flipping the orientation

of each surviving pair and comparing the original triplet with a shifted copy. If the endpoints assigned to the X-side stay close to one anchor and those assigned to the Y-side stay close to the other, then the two equations have small bidegrees, which leads to a low-degree resultant attack. Thus, from a geometric viewpoint, a good exponent schedule is one for which *every* residual triplet is simultaneously:

- spread out enough to resist short-arc univariate compression, and
- misaligned enough to resist concentration around two anchors.

Practical guideline for choosing exponent pairs. In practice, this suggests the following rules. One should avoid schedules for which:

1. some two-block cancellation leaves six endpoints contained in a short cyclic interval;
2. some residual triplet can be partitioned, after orientation choices and a Frobenius shift, into two small neighborhoods around two anchors;
3. several differences $b_k - a_k \bmod e$ are repeated or nearly repeated up to sign, which tends to create aligned residual configurations.

Conversely, good schedules are obtained by maximizing, over all residual triplets, the worst-case univariate span and the worst-case bivariate score. This provides a concrete and reproducible exponent-selection strategy that is fully compatible with the original F-UOV format and algorithms, while explicitly targeting the direct algebraic attack surface created by the Frobenius index pattern.

Additional exponent-schedule diagnostics. Beyond the basic univariate-span and two-anchor bivariate tests, the exponent schedule can be screened through several further diagnostics of the type studied in [2]. These diagnostics are not security proofs, but they help identify schedules that create especially simple low-dimensional algebraic models after cancelling some Frobenius blocks.

A first family of diagnostics concerns *affine restrictions*. If one fixes a direction $U \in \mathbb{F}_q^n$ and evaluates the public map on an affine line $V + xU$, then after cancelling two blocks the remaining expression contains both quadratic Frobenius terms and linear Frobenius terms in x. In this setting, a useful question is whether the residual exponents can be covered by a small number of *anchors* on the cyclic index set $\mathbb{Z}/e\mathbb{Z}$, with small maximal distance to an anchor. Small covering radii indicate that the affine restriction may admit a compact few-variable representation. Hence one diagnostic is to measure, for the residual indices, the best two-, three-, or four-anchor covering radius.

A second family concerns *common zeros of several public blocks*. If one searches for a vector that annihilates three or four blocks simultaneously, then the resulting system can be represented as a bihomogeneous system in two projective blocks of variables, associated with two Frobenius layers. The size of such a model can be estimated by multihomogeneous Bézout numbers. Small values suggest that the corresponding common-zero system may be unusually accessible to direct algebraic solving. Conversely, large values indicate that the naive low-degree intersection model quickly becomes large.

A third diagnostic considers *multi-layer bilinear models using all five blocks.* Instead of representing only two Frobenius powers of the unknown, one introduces several formal layers

$$Z_t \approx X^{[\theta^t]}$$

for a suitable generator θ of the Frobenius cycle, and writes each public block as a bilinear relation between two layers whose indices differ by the block difference. This yields a graph-like model whose vertices are Frobenius layers and whose edges are induced by the five block differences. Useful diagnostics then include:

- the shortest consecutive window of layers supporting all five block differences;
- the size of smaller arbitrary connected layer subsets supporting enough bilinear relations to form a square or overdetermined system;
- the associated multihomogeneous Bézout-type counts.

These models do not by themselves produce an attack, since the formal layer variables must still satisfy the Frobenius-consistency relations linking them back to a single original unknown. Nevertheless, they are valuable indicators of whether the exponent schedule creates unexpectedly small relaxed bilinear systems.

A convenient summary of such behavior is obtained from the *difference pattern*

$$\delta_k = \langle b_k - a_k \rangle_e, \qquad k = 1, \ldots, \ell.$$

Following [2], one may define, for each $t \in \{3, 4, 5\}$, a consecutive-window diagnostic w_t by minimizing the total cyclic spread of every t-subset of differences after multiplication by a unit of $\mathbb{Z}/e\mathbb{Z}$. Intuitively, small values of w_t indicate that several block differences can be simultaneously aligned inside a short orbit window, which is precisely the type of configuration that favors compact multi-layer models. Likewise, the *endpoint span*, namely the length of the shortest cyclic interval containing all endpoints a_k, b_k, gives a coarse measure of global concentration of the schedule.

Taken together, these diagnostics suggest the following practical methodology for choosing exponent pairs. First, reject schedules with repeated or nearly repeated differences up to sign, or with obvious affine symmetries. Second, for every residual triplet obtained after cancelling two blocks, maximize the minimum univariate span and the minimum two-anchor bivariate score. Third, among schedules that pass these direct tests, prefer those with larger values of the orbit diagnostics (w_3, w_4, w_5), larger endpoint span, and larger covering radii or Bézout-type scores in the affine, common-zero, and multi-layer models. This does not prove resistance to all direct algebraic attacks, but it helps avoid exponent schedules that are visibly weak in the currently known families of low-dimensional representations.

4 Concrete Security Analysis

We analyze the main attack strategies currently known for Frobenius-UOV and derive concrete cost estimates for the proposed parameter sets. The discussion

is kept as general as possible; nevertheless, we occasionally specialize to the practically relevant regime

- $o = m = 1$, and
- $p^{eo/2} > 2^{\lambda}$, where λ denotes the targeted security level.

4.1 Attack coverage and scope

We consider classical polynomial-time adversaries with full access to public keys and (for forgery) standard signing-oracle access as in EUF-CMA. We do not address side-channel leakage or fault attacks.

Covered attack families. The estimates in this section cover the following broad families: (i) direct algebraic solving of the public equations (over $\mathbb{F}_q$ and after descent to $\mathbb{F}_p$), (ii) key-recovery attacks exploiting public/secret relations (including multihomogeneous XL-style strategies and subspace-recovery viewpoints), (iii) MinRank-type approaches (over $\mathbb{F}_q$ and over $\mathbb{F}_p$ after descent), (iv) exterior-algebra (wedge) attacks based on bilinearization/differentials, and (v) generic collision/brute-force baselines.

Mapping to Tab. 2. The rows $MQ(e, 5e, p)$ and $MQ(2e, 4e, p)$ correspond to modeling direct forgery/key-recovery instances after base-field expansion, using Assumption 4.11. The terms $\log_2\left(\binom{e+e/2}{e}^2\right)$ and $\log_2(v^{2e})$ correspond to linear-algebra steps in (multi)homogeneous XL/Macaulay-style strategies under the dense-MQ/solving-degree heuristics used throughout this section. The MinRank and wedge rows correspond to the dedicated attacks discussed in their respective subsections.

4.2 Turning the $\mathbb{F}_q$-UFOs problem into a $\mathbb{F}_p$-problem.

This subsection explains how to *flatten* UFOs objects defined over the extension field $\mathbb{F}_q = \mathbb{F}_{p^e}$ into equivalent descriptions over the prime field $\mathbb{F}_p$. This translation is useful both conceptually and algorithmically: it lets us view the same objects from several angles and, in particular, reduces many steps to linear-algebraic manipulations over $\mathbb{F}_p$.

Since $\mathbb{F}_q$ is an extension of $\mathbb{F}_p$, any system of equations over $\mathbb{F}_q$ can be expanded into a (typically larger) system over $\mathbb{F}_p$. As seen in Sec. 2.3, the matrices $A_{rk} \in \mathbb{F}_q^{n\times n}$ defining Frobenius forms can be associated with a family of e block matrices in $\mathbb{F}_p^{ne\times ne}$, with $\bar{\Phi}(\cdot A_{rk}\cdot) = (A_{rk1}, \ldots, A_{rke})$ — where the (i, j) block of A_{rks} is $\Phi(.\alpha_{rkij})\mathbb{M}_s$ — such that

$$\Phi(XA_{rk}\,{}^tX) = \bar{\Phi}(X)(A_{rk1}, \ldots, A_{rke})\,{}^t\bar{\Phi}(X).$$

In turn, we obtain

$$\begin{aligned}\Phi(X^{[a_{rk}]}A_{rk}\,{}^tX^{[b_{rk}]}) &= (\bar{\Phi}(X)\mathcal{F}_{\otimes n}^{a_{rk}})(A_{rk1}, \ldots, A_{rke})\,{}^t(\bar{\Phi}(X)\mathcal{F}_{\otimes n}^{b_{rk}}) \\ &= \bar{\Phi}(X)(\mathcal{F}_{\otimes n}^{a_{rk}}A_{rk1}\,{}^t\mathcal{F}_{\otimes n}^{b_{rk}}, \ldots, \mathcal{F}_{\otimes n}^{a_{rk}}A_{rke}\,{}^t\mathcal{F}_{\otimes n}^{b_{rk}})\,{}^t\bar{\Phi}(X).\end{aligned}$$

For sake of simplicity, we note

$$\tilde{A}_{rks} = \mathcal{F}_{\otimes n}^{a_{rk}} A_{rks}\, {}^t\mathcal{F}_{\otimes n}^{b_{rk}} \quad 1 \le r \le m,\ 1 \le k \le \ell,\ 1 \le s \le e.$$

The o-dimensional oil subspace $\mathcal{O}$ of the $\mathbb{F}_q$-UFOs problem (with basis $\Lambda \subseteq \mathbb{F}_q^n$) can be viewed as a oe-dimensional oil subspace Ω of $\mathbb{F}_p^{ne}$ with basis $(\Omega_1, \ldots, \Omega_{oe})$. These vectors satisfy

$$\Omega_i \tilde{A}_{rks}\, {}^t\Omega_j = 0, \quad 1 \le i \le j \le oe,\ 1 \le r \le m,\ 1 \le k \le \ell,\ 1 \le s \le e. \tag{3}$$

Since the key-recovery equations are homogeneous, converting one equation from $\mathbb{F}_q$ to $\mathbb{F}_p^e$,

$$\bar{\Phi}(X)(\tilde{A}_{rk1}, \ldots, \tilde{A}_{rke})\, {}^t\bar{\Phi}(X) = 0$$

immediately yields e further equations by multiplying the original $\mathbb{F}_q$-equation by ξ^i and projecting to coordinates:

$$\Phi(\xi^i \cdot)\bar{\Phi}(X)(\tilde{A}_{rk1}, \ldots, \tilde{A}_{rke})\, {}^t\bar{\Phi}(X) = 0, \quad i = 0, \ldots, e-1.$$

However, these equations already lie in the $\mathbb{F}_p$-span of the original coordinate equations, since multiplication by ξ corresponds to an invertible $\mathbb{F}_p$-linear change of coordinates on $\mathbb{F}_p^e$ (and similarly after applying Frobenius), i.e., $\Phi(\xi\cdot)\bar{\Phi}(X^{[a]}) = \bar{\Phi}(X^{[a]})\Phi_{\otimes n}(\cdot\xi^{[-a]})$ and $\langle \Phi(\cdot\xi A\cdot)\rangle = \langle \Phi(\cdot A\cdot)\rangle$. Therefore, when solving over $\mathbb{F}_p$ we may restrict to candidate oil subspaces $\Omega \subseteq \mathbb{F}_p^{ne}$ that are compatible with the fixed $\mathbb{F}_p$-basis $(1, \xi, \ldots, \xi^{e-1})$ of $\mathbb{F}_q$; concretely, it is sufficient to assume that a solution satisfies

$$\Omega_{(i-1)e+j} = \Phi\big(\xi^j\, \Phi^{-1}(\Omega_{(i-1)e})\big), \quad 1 \le i \le o,\ 2 \le j \le e. \tag{4}$$

Alternately, we may look for not compatible solutions. This is equivalent to considered the "flatten" problem as a genuine UOV problem over $\mathbb{F}_p$. However in this case, a simple counting argument shows that the probability that such a solution exists is overwhelmingly negligible. Indeed, with similar arguments as in Sec. 4.4, Ω can be uniquely described over $\mathbb{F}_p$ by a system of $e\frac{eo(eo+1)}{2}$ equations in e^2ov variables. Since this system is massively overdetermined, it is not likely to have many spurious random solutions. Therefore its genuine solutions must be the ones defined by the descent.

We now review the main families of attacks relevant to multivariate signature schemes in this setting.

4.3 Attacks linked to the homogeneity of the Frobenius forms

For a given (a, b)-Frobenius form $f(X) = X^{[a]} A\, {}^tX^{[b]}$ and any scalar $\alpha \in \mathbb{F}_q$, one has $f(\alpha X) = \alpha^{p^a + p^b} f(X)$. Consequently, the ideal generated by f is stable under scalar multiplication.

Following [8], write $X = \sum_{i=1}^{e} x_i \xi^{i-1}$. Then $\xi X = \sum_{i=1}^{e} x'_i \xi^{i-1}$, where $(x'_1, \ldots, x'_e) = (x_1, \ldots, x_e)C_\xi$ and $C_\xi \in \mathbb{F}_p^{e\times e}$ is the companion matrix of the

irreducible polynomial $i(\xi)$. It follows that the ideal generated by the quadratic forms

$$\mathcal{F}^a_{\otimes n}\bar{\Phi}(\cdot A\cdot)\,{}^t\mathcal{F}^b_{\otimes n}$$

(i.e., the coordinate forms of f over $\mathbb{F}_p[x_{11},\ldots]$) is stable under the action of the matrix group generated by $(C_\xi)_{\otimes n}$. Let $D_\xi = \mathrm{Diag}(\xi, \xi^p, \ldots, \xi^{p^{e-1}})$ and let $P \in \mathbb{F}_q^{e\times e}$ be a change-of-basis matrix such that $D_\xi = PC_\xi P^{-1}$. By [8], the ideal (over $\mathbb{F}_q[x_{11},\ldots,x_{ne}]$) generated by the quadratic forms $P_{\otimes n}(\mathcal{F}^a_{\otimes n}\bar{\Phi}(\cdot A\cdot)\,{}^t\mathcal{F}^b_{\otimes n})\,{}^tP_{\otimes n}$ is stable under the action of the cyclic matrix group generated by $(D_\xi)_{\otimes n}$.

Specialization to $m = o = 1$. In this paragraph, we consider the two problems of key recovery and signature forgery in the light of [8].

Key recovery The key-recovery problem amounts to solving the public/secret relation (1) (equivalently (8)), which is a system of ℓ homogeneous equations in n unknowns over $\mathbb{F}_q$. Writing each $\mathbb{F}_q$-unknown in an $\mathbb{F}_p$-basis of size e yields ne ground-field variables.

Lemma 1 (Multi-grading after base-field expansion). *Let $X = (X_1,\ldots,X_n) \in \mathbb{F}_q^n$ and write each coordinate as $X_j = \sum_{s=0}^{e-1} x_{j,s}\,\xi^s$ with $x_{j,s} \in \mathbb{F}_p$. Consider the polynomial ring $\mathbb{F}_p[x_{j,s} \mid 1 \le j \le n,\ 0 \le s \le e-1]$ equipped with the $\mathbb{Z}^e_{\ge 0}$-grading defined by $\deg(x_{j,s}) = \mathbf{e}_s$. Then for any fixed Frobenius exponents (a,b) and any $A \in \mathbb{F}_q^{n\times n}$, each coordinate polynomial over $\mathbb{F}_p$ obtained by expanding and projecting $X^{[a]}A\,{}^tX^{[b]}$ is a homogeneous quadratic polynomial (total degree 2) that decomposes as a sum of multihomogeneous components of multidegree $\mathbf{e}_u + \mathbf{e}_v$. Moreover, applying Frobenius to X permutes the e degree coordinates, hence permutes these multidegrees.*

Proof. Since $x_{j,s} \in \mathbb{F}_p$, we have $x_{j,s}^{p^a} = x_{j,s}$. Writing $\xi^{p^a} = \sum_t c_t^{(a)}\,\xi^t$ for the fixed $\mathbb{F}_p$-linear action of Frobenius on the basis, each coordinate of $X^{[a]}$ is an $\mathbb{F}_p$-linear combination of the variables $\{x_{j,s}\}$, and similarly for $X^{[b]}$. Expanding the bilinear expression $X^{[a]}A\,{}^tX^{[b]}$ therefore yields an $\mathbb{F}_p$-linear combination of monomials of the form $x_{i,s}\,x_{j,t}$. Every such monomial has total degree 2 and multidegree $\deg(x_{i,s}x_{j,t}) = \mathbf{e}_s + \mathbf{e}_t$, so each projected coordinate polynomial is a homogeneous quadratic polynomial that splits as a sum of multihomogeneous pieces indexed by the pair (s,t). Finally, Frobenius acts on the basis indices s by a permutation of $\{0,\ldots,e-1\}$, so it permutes the degree coordinates and hence the multidegrees $\mathbf{e}_s + \mathbf{e}_t$.

Consequently, the expanded key-recovery system inherits a natural multi-graded structure with respect to the e coordinate layers corresponding to $(1,\xi,\ldots,\xi^{e-1})$. Applying the multihomogeneous XL strategy of [8] requires choosing a target multi-degree $\mathbf{d}_{sol}$; for our large e, we cannot reliably predict it, but their experiments suggest the conservative heuristic $\mathbf{d}_{sol} \ge [1,\ldots,1]$. At that first nontrivial multi-degree, we already have at least $\overline{\mathcal{M}}(\mathbf{d}_{sol}) \ge v^e$ relevant

monomials (choose one of the v vinegar coordinates independently in each of the e layers), hence the linear-algebra step in XL costs at least on the order of $\overline{\mathcal{M}}(\mathbf{d}_{sol})^2 \geq v^{2e}$ operations (Tab. 2).

Signature forgery In a direct forgery attack, the adversary fixes a target hash/output and tries to solve the public equation $\mathbf{Pub}(X) = H$ for a preimage X. In our setting, each coordinate of **Pub** is a sum of ℓ Frobenius forms with *different* exponent differences $b_k - a_k$, hence the full equation mixes monomials with incompatible Frobenius shifts. After expanding to ground-field coordinates (as above), the resulting system over $\mathbb{F}_p[x_{11}, \ldots]$ is therefore *not* multihomogeneous (there is no multigrading that is simultaneously preserved by all blocks). Consequently, the multihomogeneous XL speedups exploited in [8] do not apply to the forgery system in any direct way.

4.4 Attacks based on finding a common isotropic subspace for many public Frobenius forms

The structured form of the public key suggests attacks that attempt to recover (part of) the UOV trapdoor, namely a subspace that is simultaneously isotropic for several public matrices. This viewpoint covers direct forgeries (restricting the public map to a small subspace), key recovery up to equivalence (finding an alternative "oil" space), and hybrid attacks (first recover partial structure, then solve a simplified instance).

First let's have a look of the initial question of finding isotropic vectors of only one Frobenius form.

Isotropic plane for one Frobenius form. Let

$$Q(X) = X^{[a]} A\, {}^t X^{[b]}, \quad B(X,Y) = X^{[a]} A\, {}^t Y^{[b]}$$

be a Frobenius form over $\mathbb{F}_q$, and its associated bilinear form. To find an isotropic vector, one restricts Q to an affine line $x + ty$ and chooses y so that $B(x,y) = B(y,x) = 0$, which gives

$$Q(x + ty) = Q(x) + Q(y) t^{p^a + p^b}.$$

Hence isotropy reduces to solving a single equation of the form $t^{p^{b-a}+1} = \gamma^{p^{-a}}$ in $\mathbb{F}_q$; in the parameter regime of UFOs, this is essentially a quadratic residuosity problem. Starting from one isotropic vector u, a second isotropic vector orthogonal to u is then sought inside

$$V_u = \{z \in \mathbb{F}_q^n : \; B(u,z) = B(z,u) = 0\}.$$

Writing $w = r_1 + t r_2$ with $r_1, r_2 \in V_u$, the condition $Q(w) = 0$ yields a semilinear fractional equation in t, whose solutions provide vectors w such that

$$Q(u) = Q(w) = B(u,w) = B(w,u) = 0.$$

Thus one obtains a totally isotropic plane, which is the basic ingredient behind the two-block cancellation strategy used in the direct-forgery analysis.

Using the absence of totally isotropic 3*-spaces in dimension* 5. For a nondegenerate Frobenius form on F_q^5, a totally isotropic subspace has dimension at most $\lfloor 5/2 \rfloor = 2$. Therefore, the existence of a totally isotropic 3-dimensional subspace forces degeneracy of the underlying form. This simple dimension bound can be exploited in the common-isotropy search. Indeed, after constructing a totally isotropic plane $W = \langle u, w \rangle$ for one public block, one does not need to search for a larger isotropic space for the same block: in the nondegenerate case, no such extension exists. Consequently, the next step must be to restrict a second Frobenius form to W and solve for a common isotropic direction inside this plane, exactly as in the two-block cancellation method. Conversely, if one finds that the orthogonal restriction still contains an isotropic direction transverse to W, then this certifies that the first form is degenerate on F_q^5 (or at least on the relevant support), which may itself yield a shortcut for cryptanalysis by lowering the effective rank and reducing the search space. Thus, in dimension 5, the nonexistence of totally isotropic 3-spaces for nondegenerate forms justifies stopping the isotropic construction at dimension 2 and turning immediately to the next block.

Heuristic counting for a common isotropic subspace. Fix t public Frobenius blocks : $A_1, \ldots, A_t \in \mathbb{F}_q^{n \times n}$ with associated exponents (a_k, b_k). We look for an s-dimensional subspace $W \subset \mathbb{F}_q^n$ such that

$$\forall X, Y \in W, \forall k \in \{1, \ldots, t\} \quad X^{[a_k]} A_k \, {}^t Y^{[b_k]} = 0.$$

Working assumption (random-isotropy model). For estimation purposes, we treat the constraints induced by distinct blocks as behaving like *independent random* quadratic constraints on the parameters describing W (after fixing a normal form for a basis of W). Concretely, we assume that for a uniformly random W and for "generic" public blocks A_k, the scalars $X_i^{[a_k]} A_k \, {}^t X_j^{[b_k]}$ behave close to uniform in $\mathbb{F}_q$ as (i, j, k) vary, except for the algebraic dependencies forced by symmetry and by the choice of basis.

Assuming such a W exists, we represent a basis of W by a matrix in reduced row echelon form. With overwhelming probability (for uniformly random W), the first s pivot columns are distinct, so this matrix can be written as $X = (1_s \; R_{s \times (n-s)})$, and W is parametrized by the $s(n-s)$ entries of R. The isotropy constraints become

$$X_i^{[a_k]} A_k \, {}^t X_j^{[b_k]} = 0, \quad 1 \le i, j \le s, \; 1 \le k \le t,$$

namely $s^2 t$ scalar equations in $s(n-s)$ unknowns over $\mathbb{F}_q$. Under the random-isotropy model above, the expected number of solutions is therefore about

$$q^{\,s(n-s) - s^2 t}.$$

In particular, when $s(n-s) - s^2 t < 0$, we expect a solution to exist only with probability about $q^{\,s(n-s)-s^2 t}$.

This suggests that searching for common isotropic subspaces is only plausible in the regime $s(t+1) \le n$ (equivalently $s(n-s) \ge s^2 t$). For instance, when $s = 1$,

one should not expect a nontrivial common isotropic vector for many more than $n-1$ random matrices.

When $t = 1$ and the induced quadratic form is nondegenerate, a standard dimension argument shows that a totally isotropic subspace has dimension at most $\lfloor n/2 \rfloor$: if a form had an isotropic subspace of dimension $> \lfloor n/2 \rfloor$, then in a basis extending that subspace the associated matrix would contain a large zero block, contradicting full rank.

A simple iterative heuristic. A simple method to obtain at least one vector that is simultaneously isotropic for several forms is as follows. Start with one form, and extract a totally isotropic subspace of dimension $t_1 := \lfloor n/2 \rfloor$ (or smaller, depending on degeneracy). At step i, assume we have collected t_i independent vectors spanning a subspace that is isotropic for the first i forms. Pick a new $(i+1)$st form and restrict it to this t_i-dimensional space. If the restriction behaves like a random nondegenerate quadratic form on t_i variables, we expect its maximal totally isotropic subspaces to have dimension about $\lfloor t_i/2 \rfloor$, hence we expect to obtain $t_{i+1} := \lfloor t_i/2 \rfloor$. Iterating while $t_i \geq 2$ (equivalently $2^i \leq n$) suggests that one may hope to obtain a common isotropic vector for at most on the order of $\lfloor \log_2(n) \rfloor$ forms.

A hybrid "partial key + restricted inversion" strategy. A generic approach is to pick an m-dimensional subspace $U \subset \mathbb{F}_q^n$ with basis $(U_1, \ldots, U_m)$ and to search for a preimage within that subspace. Writing a candidate as $u = \sum_{i=1}^{m} x_i U_i$ with unknowns $X = (x_1, \ldots, x_m)$ yields the restricted system

$$\sum_{k=1}^{\ell} \mathbf{Pub}_{rk}(X \cdot U) = h_r, \quad 1 \leq r \leq m,$$

where $X \cdot U$ denotes the corresponding linear combination of the basis vectors. Expanding one block gives

$$\mathbf{Pub}_{rk}(X \cdot U) = \sum_{i,j=1}^{m} x_i^{[a_{rk}]} x_j^{[b_{rk}]} \big(U_i^{[a_{rk}]} A_{rk} \, {}^t U_j^{[b_{rk}]} \big).$$

Working assumption (random-coefficient model for restricted instances). For estimation purposes, we model the coefficients $U_i^{[a_{rk}]} A_{rk} \, {}^t U_j^{[b_{rk}]} \in \mathbb{F}_q$ (for $1 \leq i, j \leq m$) as behaving close to independent uniform field elements when U is chosen uniformly at random, except for the algebraic relations forced by symmetry/alternation (when applicable) and by the shared choice of U across blocks. Under this random-coefficient model, the restricted system resembles an unstructured system in m variables.

The goal is therefore to choose U so that many coefficients vanish, i.e., so that the basis vectors are isotropic for as many public blocks as possible. Equivalently, we search for an "incomplete" trapdoor and then attempt a cheaper inversion on the remaining blocks.

Table 1. Trade-off between the common-isotropy step and the remaining polynomial step as a function of t (case $m = o = 1$) and $\ell = 5$. Legend: $+$: polynomial, ∞: over security level, $-$: depends on the exponents.

t	0	1	2	3	4	5
$CI(1, t, p, e, n)$	$+$	$+$	$+$	∞	∞	∞
$\ell - t$	5	4	3	2	1	0
$CP(\ell - t, p, e)$	∞	∞	∞	$-$	$+$	$+$
total	∞	∞	∞	∞	∞	∞

Specialization to $m = o = 1$. To make this trade-off explicit, consider $m = o = 1$. The attacker first tries to find a nonzero $U \in \mathbb{F}_q^n$ such that

$$\mathbf{Pub}_k(U) = U^{[a_k]} A_k \, {}^t U^{[b_k]} = 0, \quad k = 1, \ldots, t. \tag{5}$$

Given such a U, one then solves

$$\sum_{k=t+1}^{\ell} \mathbf{Pub}_k(xU) = \sum_{k=t+1}^{\ell} x^{[a_k]} x^{[b_k]} \mathbf{Pub}_k(U) = h. \tag{6}$$

Equation (5) is a common-isotropy problem for t matrices, whereas (6) is a univariate sparse polynomial equation with $\ell - t$ monomials of degrees $p^{a_k} + p^{b_k}$ and (heuristically) random coefficients.

Let $CI(s, t, p, e, n)$ denote[4] the cost of solving (5), and $CP(\ell - t, p, e)$ the cost of solving (6). The attacker would choose t to minimize $CI(1, t, p, e, n) + CP(\ell - t, p, e)$. Our heuristic counting above suggests that $CI(s, t, p, e, n)$ becomes exponential (in e) as soon as $s(t+1) > n$. In particular, for $s = 1$, beyond a small number $\lfloor \log_2(n) \rfloor$ (see the iterative method above) of matrices, the best known approach is to solve a genuine MQ system over $\mathbb{F}_p$ obtained by expanding $\mathbb{F}_q$. For estimation purposes, we then apply Assumption 4.11 and model the expanded instance as a dense-looking $MQ(te, ve, p)$ system (Tab. 2).

By contrast, $CP(1, p, e)$ is polynomial-time (essentially extracting a root of a monomial), while the cost grows rapidly—and in practice exponentially—with the number of monomials.

State of the art for sparse univariate instances. Equation (6) defines a sparse polynomial over $\mathbb{F}_q$ whose exponents have the special shape $p^{a_k} + p^{b_k}$. This is *not* a linearized (additive) polynomial, so the usual reductions to $\mathbb{F}_p$-linear algebra for sums of p-powers do not apply. Moreover, it is typically not a binomial/trinomial in a form where classical gcd tricks immediately reduce it to root extraction (e.g., when exponents share large common factors with $q - 1$). In the absence of such special structure, the generic approaches are: (i) treat it as an ordinary univariate polynomial and factor/find roots in time quasi-linear in the degree, and (ii) apply Weil descent to $\mathbb{F}_p$ and attack the resulting MQ system.

[4] C stands for complexity, I for isotropy, and P for polynomial.

Assumption (sparse Frobenius-sum root finding). For the concrete parameter sets we target, and in particular for well chosen exponents (a_k, b_k), we assume that for uniformly random coefficients $\alpha_1, \ldots, \alpha_L \in \mathbb{F}_q$ and $\gamma \in \mathbb{F}_q$ (with $L \in \{2, 3, 4, 5\}$), it is computationally infeasible for a PPT adversary to find $x \in \mathbb{F}_q$ satisfying

$$\sum_{i=1}^{L} \alpha_i x^{p^{a_i}+p^{b_i}} = \gamma,$$

in time below the targeted security level, beyond the generic approaches (i) and (ii) above.

If one solves (6) as a univariate polynomial, the best known algorithms run in time essentially linear in the degree (here on the order of $p^{e/2}$ in our parameter regime) [19,22,10,31,33], i.e., exponential in e.

If instead one rewrites (6) via descent as a multivariate quadratic system over $\mathbb{F}_p$, small-scale experiments suggest that for $\ell - t \geq 3$ (and possibly already for $\ell - t \geq 2$) it behaves like a dense quadratic system in e variables with solving degree about $e/2$. Concretely, on toy instances with small p and e (chosen so that Gröbner-basis computations are feasible), we measured the first degree d at which the degree-d Macaulay matrix reaches full rank / produces a univariate eliminant or a contradiction, and observed values clustering around $d \approx e/2$. This yields the crude estimate $O\left(\binom{e+e/2}{e}^2\right)$, again exponential in e.

For the parameter sets we target, this suggests choosing t small (so that CI remains manageable) while avoiding too small $\ell - t$; for instance, in the low-dimensional regime $n \leq 7$, we focus on $t = 2$, hence $\ell \geq 5$ (see Tab. 1).

Complexities are summarized in Tab. 2, where we report $MQ(te, ve, p)$ (from CryptographicEstimators [14]) and $\binom{e+e/2}{e}^2$.

4.5 Direct attacks

Specialization to $m = o = 1$ and $\ell = 5$.

Solving $\mathbf{Pub}(X) = H$ in $\mathbb{F}_q$. In this parameter regime (cf. Sec. 4.4), a natural strategy is to first find a vector U that cancels two ($t = 2$) Frobenius forms (which can be done in polynomial time by solving the corresponding common-isotropy conditions; proof sketch: descend the two Frobenius forms to $\mathbb{F}_p$ as quadratic forms in ne variables (Sec. 4.2); find an isotropic $\mathbb{F}_p$-subspace for the first form by standard quadratic-form linear algebra, restrict the second form to that subspace, and extract an isotropic vector; finally map it back to $\mathbb{F}_q^n$ via the fixed basis), and then solve the restricted equation $\mathbf{Pub}(X \cdot U) = H$.

The latter is a univariate sparse polynomial equation with $\ell - 2 = 3$ monomials of the form $X^{p^{a_k}+p^{b_k}}$. As discussed in the sparse-univariate paragraph above, these exponents do not yield a linearized (additive) polynomial, so there is no immediate reduction to $\mathbb{F}_p$-linear algebra; absent further exploitable coincidences in the exponents, the natural attack routes are generic univariate root finding

(cost dominated by the degree) or Weil descent followed by MQ solving. Under these standard models, this step remains above our targeted security level for the chosen parameters.

Solving $\mathbf{Pub}(\boldsymbol{X}) = \boldsymbol{H}$ in $\mathbb{F}_{\boldsymbol{p}}$. After base-field expansion, we obtain an MQ instance with approximately oe equations in approximately ne variables over $\mathbb{F}_p$. For estimation purposes, we apply Assumption 4.11 and treat it as a generic dense MQ system. This attack is also the most constraining one in our parameter selection (i.e., closest to the security threshold). Complexities are summarized in Tab. 2, where we report $MQ(e, 5e, p)$ using CryptographicEstimators [14].

4.6 Beyond univariate and bivariate direct models

The direct attack of [2] shows that, after cancelling two public blocks and restricting the public map to a public line $\mathbb{F}_q U$, the residual three-term equation can admit a low-degree bivariate representation, from which a forgery is obtained by elimination. More generally, once the residual equation is written as

$$\sum_{k \in I} c_k \, x^{[a_k]} x^{[b_k]} = h,$$

it is natural to introduce several Frobenius layers of the same unknown and to treat them as independent variables. In other words, one may replace the single scalar x by a family

$$Z_t \approx x^{[\theta_t]},$$

and rewrite each residual block as a bilinear relation between two selected layers. This strictly generalizes the univariate model of and the bivariate resultant model of [2].

This viewpoint leads to a family of multihomogeneous relaxed systems on products of projective spaces, indexed by the number and placement of the chosen Frobenius layers. As observed in [2], such models may have much smaller individual degree components than the original univariate equation: for example, once two suitable layers are separated, a term $x^{[a_k]} x^{[b_k]}$ becomes bilinear. This suggests that 3-variable, 4-variable, or even larger layer-based models may reveal compact algebraic representations that are invisible in the univariate picture. In particular, the orbit-style diagnostics and arbitrary-layer witnesses of [2] indicate that the first published exponent schedule admits small relaxed bilinear systems even beyond the explicit bivariate forgery route, and that exponent-only modifications should therefore be assessed against a broader class of multivariable representations.

At the same time, these multivariable systems are intrinsically relaxed: they do not automatically enforce the compatibility conditions requiring all variables to come from a single Frobenius orbit, namely

$$[Z_t] = [\theta_t(Z_0)].$$

Accordingly, a small multivariable representation does not by itself yield a forgery algorithm. Rather, it should currently be interpreted as a structural warning sign or as a cryptanalytic diagnostic. The present state of the analysis is therefore the following: the first published schedule is broken by an explicit low-degree bivariate elimination attack, whereas the modified schedules eliminate this specific direct model and significantly enlarge the corresponding enumerated diagnostics; however, unrestricted multivariable direct attacks remain open.

4.7 Direct key recovery

Specialization to $m = o = 1$ and $\ell = 5$. We next consider direct key-recovery strategies, i.e., attacks that attempt to recover (part of) the secret oil subspace from the public key. Let $\mathcal{O} \subset \mathbb{F}_q^n$ denote the secret oil subspace (of dimension o). By construction, every public block vanishes on $\mathcal{O}$:

$$\forall\, Y \in \mathcal{O},\ \forall\, r \in \{1, \ldots, m\},\ \forall\, k \in \{1, \ldots, \ell\}, \qquad \mathbf{Pub}_{rk}(Y) = 0.$$

In the regime $o = m = 1$, the (nonzero) secret vector $\Lambda \in \mathbb{F}_q^n$ spanning $\mathcal{O}$ is therefore a common isotropic vector for the ℓ Frobenius forms $\mathbf{Pub}_{1k}$ (up to scaling).

Solving $\mathbf{Pub}_{rk}(X) = 0$ in $\mathbb{F}_q$. For $o = m = 1$, a direct approach consists in finding a nonzero solution $X \in \mathbb{F}_q^n$ to

$$\mathbf{Pub}_k(X) = 0, \qquad k = 1, \ldots, \ell. \tag{7}$$

Since solutions are defined up to scaling, one may normalize one coordinate (when possible) and view (7) as a polynomial system of ℓ equations in v variables over $\mathbb{F}_q$. Since $\ell > v$, we use a conservative generic-model heuristic: we treat the overdetermined system as behaving like a single hard elimination equation of large degree (akin to a univariate equation over a comparable extension), and we lower bound the cost by $q^{e/2}$.

Solving $\mathbf{Pub}_{rk}(X) = 0$ in $\mathbb{F}_p$. As in Sec. 4.2, one may expand (7) to the base field. In the case $o = m = 1$, this yields an MQ system over $\mathbb{F}_p$ with about ℓe equations in about $(n-1)e$ variables. We therefore estimate this attack by

$$MQ(\ell e, ve, p).$$

Using the same convention as in the other rows of Tab. 2, for our parameter sets this estimate is numerically dominated by $MQ(2e, ve, p)$ and is therefore omitted from the concrete comparisons.

Relation to the p^ℓ-truncated-ring XL framework. Furue and Ikematsu [17] propose a unified view of several UOV subspace-recovery attacks by modeling the key-recovery equations over p^ℓ-truncated polynomial rings and solving via XL. This framework is in principle relevant to our scheme once it is descended to $\mathbb{F}_p$, since UFOs induces a UOV-shaped instance with $N = ne$ base-field variables and an oil dimension $\alpha = oe = e$. However, for our concrete parameters ($n = 5$, $o = 1$, $v = 4$, $\ell = 5$, and $e \in \{61, 89, 127\}$), the resulting values $N = 5e$ are already in the hundreds, and the dominant Macaulay/Wiedemann linear-algebra step in their complexity model scales with (truncated) monomial counts at the solving degree; consequently, we do not expect this approach to yield a practical key recovery against our parameter sets or to outperform the other base-field attacks already accounted for in Tab. 2. Some tests with the code `complexity` at [17, Sec. C] confirm this expectation.

4.8 Kipnis Shamir attack

In $\mathbb{F}_q$. This attack is not applicable since there is no equivalent of the matrix representation of the public key. Indeed the public key is $\mathbb{F}_p$-quadratic, but not $\mathbb{F}_q$-quadratic.

Specialization to $p^{eo/2} > 2^\lambda$.

In $\mathbb{F}_p$. Taking into account our scheme, this attack, which aims at finding a vector in a kernel, has complexity lower bounded by $O(p^{e(v-o)})$. Our parameters make this attack inefficient as soon as $v - o \geq o/2$, since $p^{eo/2}$ is already above the security level.

4.9 Minrank attacks

In $\mathbb{F}_q$. Equations (8) can be rewritten as

$$\begin{aligned} \Lambda A_{rk}^{[-a_{rk}]}\, {}^t\Lambda^{[b_{rk}-a_{rk}]} &= 0_{o\times o}, \quad 1 \leq r \leq m,\ 1 \leq k \leq \ell, \\ \Lambda\, {}^tA_{rk}^{[-b_{rk}]}\, {}^t\Lambda^{[a_{rk}-b_{rk}]} &= 0_{o\times o}, \quad 1 \leq r \leq m,\ 1 \leq k \leq \ell. \end{aligned}$$

A Minrank-style approach would require exploitable linear dependencies between the subspaces spanned by (rows of) Λ, $\Lambda^{[b_{rk}-a_{rk}]}$, and $\Lambda^{[a_{rk}-b_{rk}]}$, for some (r, k).

Heuristic (uniform expanded rows). We treat the rows of the expanded basis Λ^* as behaving close to uniformly random elements of $\mathbb{F}_q^v$ (subject to the fixed structural constraints imposed by the key generation). Under this heuristic, the probability that a given row lies in the subspace $\mathbb{F}_p^v$ is $(p/q)^v$, and a union bound yields an overall probability at most $1/p^{v(e-1)}$. Accordingly, we do not expect an effective Minrank attack in the $\mathbb{F}_q$ model.

Specialization to $m = o = 1$.

In $\mathbb{F}_p$. Using (3), one can mount an (extended) Minrank attack over $\mathbb{F}_p$; see [32]. Using the trick in (4), we obtain a Minrank instance with $ve+1$ matrices of ℓe rows, $ve+1$ columns, and target rank ve. See Fig. 1.

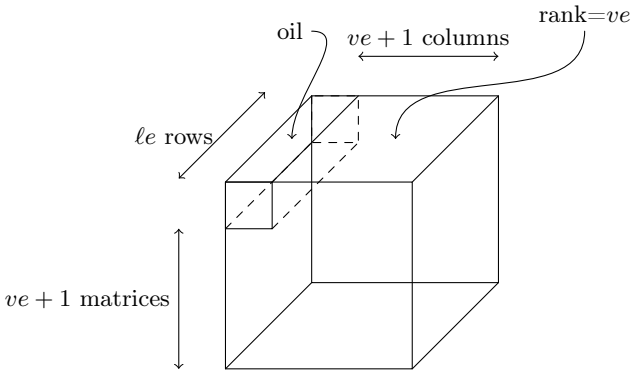

Fig. 1. Schematic 3D view of the secret matrix family. From above, one sees matrices with ℓe rows and ne columns; from the left, one sees the stack of ne matrices. The top-left corner corresponds to the secret oil space (all zeros). The schematic also indicates the number of matrices and their dimensions that are sufficient and necessary to mount the Minrank attack.

Complexity estimates are given in Tab. 2 (`MinRank` from [14]).

4.10 Wedge attack

This attack was introduced in [28]. It relies on exterior algebra and on the correspondence between wedge products, subspaces, and alternating forms. Concretely, it considers the linear map

$$M : \bigwedge^{v}(\mathbb{F}_q^n) \longrightarrow \Big(\bigwedge^{v+2}(\mathbb{F}_q^n)\Big)^m, \qquad \alpha \longmapsto (\alpha \wedge q_1, \ldots, \alpha \wedge q_m),$$

whose kernel is conjectured (and supported by heuristic arguments in [28]) to correlate with the orthogonal of the oil subspace, in the sense that nontrivial kernel elements may reveal wedge encodings of subspaces containing (or closely related to) $\mathcal{O}^{\perp}$. Here, $q_1, \ldots, q_m$ are derived from the alternating part of the bilinearization (equivalently, from differentials) of the public polynomials.

In our setting, the natural bilinearization $B(X,Y) = X^{[a]}A\,{}^tY^{[b]}$ is $\mathbb{F}_p$-bilinear but not $\mathbb{F}_q$-bilinear, so the symmetric/alternating decompositions live canonically over $\mathbb{F}_p$ rather than over $\mathbb{F}_q$. Consequently, an adversary would have to descend the instance and consider a $\mathbb{F}_p$-wedge attack. For our parameters, this appears infeasible, as it entails linear algebra in spaces of dimension on the order of $\binom{ne}{ve}$ (e.g., Gaussian elimination costing at least $\Omega\big(\binom{ne}{ve}^{\omega}\big)$ operations for some

$\omega \in [2, 3]$; for instance, when $n = 5$ and $o = 1$ one has

$$\log_2 \binom{ne}{ve} \approx ne\, H_2(v/n) = ne\, H_2(4/5) \approx 0.72\, ne$$

(Stirling), i.e., about 220, 321, and 458 bits for $e = 61, 89, 127$ respectively). See Tab. 2.

4.11 Collision attack

This is a generic attack: the adversary repeatedly evaluates the public map (and the salted hash) on trial inputs until the target value is matched. The cost is therefore dominated by public-key evaluations rather than by algebraic solving.

We estimate this work factor in the $\mathbb{F}_p$-expanded representation by considering the induced UOV-shaped instance with oe equations in ne variables over $\mathbb{F}_p$, and by using `UOV/collision` attack estimation from [14] to obtain a concrete baseline cost. The resulting estimates are reported in Tab. 2.

Table 2. Complexities of attacks. $MQ(m, n, p)$ is the problem of solving m random equations in n variables over $\mathbb{F}_p$.

Security Level	I	III	V
p, e	31, 61	61, 89	61, 127
o, v, ℓ	1, 4, 5		
NIST reference	143	207	272
$\log_2(p^{e/2})$	151	263	376
$MQ(e, 5e, p)$	154	236	327
$MQ(2e, 4e, p)$	309	545	631
$\log_2\left(\binom{e+e/2}{e}^2\right)$	159	236	340
$\log_2\left(\binom{ne}{ve}\right)$	216	316	453
$\log_2(v^{2e})$	244	356	508
MinRank	336	565	792
Collision	164	277	390

Methodology for concrete estimates. Whenever we write $MQ(m, n, p)$, we mean the estimated cost of solving a *random dense* $m \times n$ MQ instance over $\mathbb{F}_p$, as returned by CryptographicEstimators [14] under its default settings.

Assumption (dense MQ after descent). When we descend a structured $\mathbb{F}_q$-system to $\mathbb{F}_p$ (by expanding each $\mathbb{F}_q$-variable in a fixed $\mathbb{F}_p$-basis), we assume that the resulting $\mathbb{F}_p$-system behaves, from the viewpoint of the dominant linear-algebra step in Gröbner/XL-type solvers, comparably to a random dense MQ instance with the same (m, n, p). We therefore use $MQ(m, n, p)$ as a baseline estimate for this step.

Limitations of Assumption 4.11. This modeling choice may fail if the descended equations retain exploitable algebraic structure (e.g., unusually low rank, non-generic multihomogeneity, large symmetry groups, or other dependencies that provably reduce the solving degree or the linear-algebra dimension). We do not claim such structure is absent in full generality; rather, we use Assumption 4.11 as a conservative baseline and complement it with attack-specific structural discussions in Sec. 4.

5 Description of the scheme

5.1 Simplified description

We provide a compact, computation-oriented description of the objects and operations used by the scheme.

Public map. For $1 \le r \le m$ and $1 \le k \le \ell$, define the k-th Frobenius block of the r-th public equation by

$$\mathbf{Pub}_{rk}(X) = X^{[a_{rk}]} A_{rk}\, {}^{t}X^{[b_{rk}]},$$

where $A_{rk} \in \mathbb{F}_q^{n\times n}$ and $X \in \mathbb{F}_q^n$, $(a_{rk}, b_{rk}) \in (\mathbb{Z}/e\mathbb{Z})^2$. The public map is the sum of these blocks:

$$\mathbf{Pub} = (\mathbf{Pub}_1, \ldots, \mathbf{Pub}_m), \qquad \mathbf{Pub}_r(X) = \sum_{k=1}^{\ell} \mathbf{Pub}_{rk}(X).$$

Secret oil subspace. As in UOV, the trapdoor is an o-dimensional subspace (the "oil" space) that is totally isotropic with respect to every public block. Using the standard block form (2) for the secret change of variables, we represent a row basis of the oil space as

$$\Lambda = (1_o \; R_{o\times v}) \in \mathbb{F}_q^{o\times n}, \qquad \Lambda^* := R_{o\times v} \in \mathbb{F}_q^{o\times v},$$

so that the secret information is entirely captured by Λ^*.

The embedded UOV constraints (1) can be written as

$$\Lambda^{[a_{rk}]} A_{rk}\, {}^{t}\Lambda^{[b_{rk}]} = 0_{o\times o}, \quad 1 \le r \le m,\ 1 \le k \le \ell. \tag{8}$$

Equivalently, $\mathbf{Pub}_{rk}(X\Lambda) = 0$ holds for all $X \in \mathbb{F}_q^o$.

Block decomposition of the public matrices. To make the key-generation and signing formulas explicit, we decompose A_{rk} according to the $(o \mid v)$ partition of coordinates:

$$A_{rk} = \begin{pmatrix} A_{rk,o\times o} & A_{rk,o\times v} \\ A_{rk,v\times o} & A_{rk,v\times v} \end{pmatrix}.$$

We abbreviate

$$A^0_{rk} := A_{rk,o\times o}, \quad A^1_{rk} := A_{rk,v\times o}, \quad A^2_{rk} := A_{rk,o\times v}, \quad A^3_{rk} := A_{rk,v\times v},$$

and define the "known part"

$$A^*_{rk} := \begin{pmatrix} 0_{o\times o} & A^2_{rk} \\ A^1_{rk} & A^3_{rk} \end{pmatrix}.$$

Key-generation equation. Substituting $\Lambda = (1_o \ \Lambda^*)$ into (8) yields an explicit expression for the missing block A^0_{rk}:

$$A^0_{rk} = -\Lambda^{[a_{rk}]} A^*_{rk}\, {}^t\Lambda^{[b_{rk}]} = -\Big(\Lambda^{*[a_{rk}]} A^1_{rk} + A^2_{rk}\, {}^t\Lambda^{*[b_{rk}]} + \Lambda^{*[a_{rk}]} A^3_{rk}\, {}^t\Lambda^{*[b_{rk}]}\Big), \quad (9)$$

for all $1 \le r \le m$ and $1 \le k \le \ell$.

Key generation. Equation (9) gives a direct way to represent and generate both keys. Starting from a private seed, generate Λ^*. Starting from a public seed, generate the matrices A^*_{rk} for $r = 1, \ldots, m$ and $k = 1, \ldots, \ell$. Finally, compute the missing blocks A^0_{rk} via (9) and assemble the full matrices A_{rk}.

Trap door inversion. We describe how to solve the system $\mathbf{Pub}(X) = H$ over $\mathbb{F}_q^n$, i.e., m equations in n variables. The inversion procedure follows the standard UOV strategy and exploits the isotropy of the oil subspace; we adapt it here to the UFOs setting.

We look for a solution of the form $X\Lambda + V$, where $X \in \mathbb{F}_q^o$ is unknown, Λ is a basis matrix of the oil subspace, and $V \in \mathbb{F}_q^n$ is sampled uniformly at random. To reduce the computational cost, we restrict to vectors $V = (0_o \mid V^*)$, i.e., we fix the first o coordinates to 0 and sample the remaining v coordinates uniformly.

By isotropy of the oil subspace, we obtain

$$\begin{aligned} \mathbf{Pub}_{rk}(X\Lambda + V) = {} & X^{[a_{rk}]} \Lambda^{[a_{rk}]} A_{rk}\, {}^t V^{[b_{rk}]} + V^{[a_{rk}]} A_{rk}\, {}^t\Lambda^{[b_{rk}]}\, {}^t X^{[b_{rk}]} \\ & + \mathbf{Pub}_{rk}(V), \qquad r = 1, \ldots, m,\ k = 1, \ldots, \ell. \end{aligned} \quad (10)$$

As recalled in the preliminaries, the Frobenius map is $\mathbb{F}_p$-linear. Consequently, the non-constant part of $\mathbf{Pub}_r(X\Lambda + V)$ involves only $\mathbb{F}_p$-linear monomials in $X^{[a_{rk}]}$ and $X^{[b_{rk}]}$, and can be written as

$$\sum_{k=1}^{\ell} \Big(V^{*[b_{rk}]} La_{rk}\, {}^t X^{[a_{rk}]} + V^{*[a_{rk}]} Lb_{rk}\, {}^t X^{[b_{rk}]} \Big).$$

The matrices La_{rk} and Lb_{rk} can be precomputed as

$$La_{rk} = {}^t\big(\Lambda^{*[a_{rk}]} A^3_{rk} + A^2_{rk}\big), \quad Lb_{rk} = A^3_{rk}\, {}^t\big(\Lambda^{*[b_{rk}]} + A^1_{rk}\big). \quad (11)$$

The constant part is

$$\mathbf{Pub}_r(V) = \sum_{k=1}^{\ell} V^{*[a_{rk}]} A^3_{rk} \, {}^t V^{*[b_{rk}]}.$$

Indeed, since we restrict to $V = (0_o \mid V^*)$, only the $(v \times v)$ block A^3_{rk} contributes to $V^{[a_{rk}]} A_{rk} \, {}^t V^{[b_{rk}]}$.

The $\mathbb{F}_p$-linear system in X and the matrix M. Fix V. Define the $\mathbb{F}_q$-valued map

$$L_V(X) := \mathbf{Pub}(X\Lambda + V) - \mathbf{Pub}(V).$$

By the previous expansion, every coordinate of $L_V(X)$ is a sum of terms of the form $\alpha\, X^{[a]}$ and $\beta\, X^{[b]}$ with coefficients $\alpha, \beta \in \mathbb{F}_q$ depending (linearly) on V and the key; hence L_V is $\mathbb{F}_p$-linear in X. Therefore, after applying $\bar{\Phi}$, there exists a matrix $M(V) \in \mathbb{F}_p^{oe \times me}$ representing this linear map in the canonical $\mathbb{F}_p$-bases such that

$$\bar{\Phi}(L_V(X)) = \bar{\Phi}(X) M(V),$$
$$\text{i.e.,} \qquad \bar{\Phi}(\mathbf{Pub}(X\Lambda + V)) = \bar{\Phi}(X) M(V) + \bar{\Phi}(\mathbf{Pub}(V)).$$

Let $\Gamma = H - \mathbf{Pub}(V)$. Since $o = m$, the matrix $M(V)$ is square, and whenever it is invertible a solution is

$$X = \bar{\Phi}^{-1}(\bar{\Phi}(\Gamma) M(V)^{-1}).$$

Success probability (invertibility of $M(V)$). The entries of $M(V)$ are $\mathbb{F}_p$-linear forms in the $\mathbb{F}_p$-coordinates of V^*; in particular $\det(M(V))$ is a polynomial over $\mathbb{F}_p$ in those coordinates. Heuristically (as in standard UOV analyses), this determinant is not identically zero, and for uniformly random V^* the failure event $\det(M(V)) = 0$ then has probability at most $\deg(\det)/p$ (Schwartz–Zippel); empirically, the success rate is close to that of a random square matrix over $\mathbb{F}_p$.

Remark 1. (Computation of M.) The theoretical expression for the (i, j) $e \times e$ block of M can be written as a sum of matrices of the form $\Phi(.\alpha)\mathcal{F}^a$. In practice, it is often faster to construct the matrix of the $\mathbb{F}_p$-linear map $x \mapsto \alpha x^{p^a}$ directly by evaluating it on the canonical basis $(\xi^0, \xi^1, \ldots, \xi^{e-1})$ and assembling the resulting image vectors.

Remark 2. (Homogeneity of **Pub**.) For every $\alpha \in \mathbb{F}_p$ and every $\Sigma \in \mathbb{F}_q^n$ we have $\mathbf{Pub}(\alpha\Sigma) = \alpha^2 \mathbf{Pub}(\Sigma)$; in particular, $\mathbf{Pub}(\Sigma) = \mathbf{Pub}(-\Sigma)$. To remove this sign ambiguity (and thus avoid trivial signature forgeries), one can impose a canonical normalization by fixing one $\mathbb{F}_p$-coordinate of h and Σ (or, more generally, any fixed $\mathbb{F}_p$-linear form such as the trace to $\mathbb{F}_p$). For instance, set the first $\mathbb{F}_p$-coordinate of h to 1 and, between Σ and $-\Sigma$, keep the lexicographically smaller representative.

5.2 Key generation

Algorithm 1: $\mathsf{KeyGen}(1^\lambda)$

Input: security parameter λ.
Output: secret key sk, public key pk.

1. Sample uniformly at random a 32-byte seed $\mathsf{SeedSec}$.
2. Expand $\mathsf{SeedSec}$ into ov elements of $\mathbb{F}_q$, and set these coefficients as Λ^*.
3. Derive from $\mathsf{SeedSec}$ a 16-byte seed $\mathsf{SeedPub}$.
4. Expand $\mathsf{SeedPub}$ into $N = m\ell(n^2 - o^2)$ elements of $\mathbb{F}_q$, and interpret them as the coefficients of the $m\ell$ matrices A^*_{rk}, each having a zero $o \times o$ block.
5. For all (r, k), compute the missing $o \times o$ block A^0_{rk} using (9), and form A_{rk}.
6. Output $\mathsf{sk} \leftarrow \mathsf{SeedSec}$.
7. Output $\mathsf{pk} \leftarrow \big(\mathsf{SeedPub}, \{A^0_{rk}\}_{r,k}\big)$.

5.3 Signature process

Algorithm 2: $\mathsf{Sign}(\mathsf{sk}, \mathcal{M})$

Input: secret key $\mathsf{sk} = \mathsf{SeedSec}$, message $\mathcal{M}$.
Output: signature $\sigma = (\Sigma, \tau)$.

1. Recompute Λ^* and $\mathsf{SeedPub}$ from $\mathsf{SeedSec}$.
2. Recompute the matrices A^*_{rk} from $\mathsf{SeedPub}$.
3. Compute the missing blocks A^0_{rk} using (9), and form A_{rk}.
4. Repeat:
   (a) Sample a salt τ uniformly at random and set $H := (h_1, \ldots, h_m) := \mathcal{H}(\mathcal{M} \mid \tau)$.
   (b) Sample $V := (0, \ldots, 0, V_{o+1}, \ldots, V_n)$ uniformly at random.
   (c) Build the $\mathbb{F}_p$-linear system corresponding to $\mathbf{Pub}(X\Lambda + V) = h$ (optionally using the precomputed matrices La_{rk} and Lb_{rk} from (11)).
   (d) Solve this linear system for X as in Sec. 5.1; denote a solution by X_0 (if no solution exists, restart the loop).
   (e) Set $\Sigma \leftarrow X_0\Lambda + V$.
5. Until $\mathbf{Pub}(\Sigma) = H$.
6. Output $\sigma \leftarrow (\Sigma, \tau)$.

Table 3. Exponent schedules used in this work (from the exponent-only hardening analysis of [2]).

Level	(p, e)	(a_1, b_1)	(a_2, b_2)	(a_3, b_3)	(a_4, b_4)	(a_5, b_5)
I	$(31, 61)$	$(0, 9)$	$(7, 54)$	$(13, 29)$	$(15, 37)$	$(35, 47)$
III	$(61, 89)$	$(0, 13)$	$(28, 43)$	$(40, 61)$	$(52, 84)$	$(53, 76)$
V	$(61, 127)$	$(0, 2)$	$(22, 83)$	$(29, 74)$	$(53, 79)$	$(94, 111)$

Table 4. Concrete parameters and sizes (bytes).

Level	p	e	$\lvert\sigma\rvert$	$\lvert\mathbf{Pub}\rvert$	$\lvert\sigma\rvert + \lvert\mathbf{Pub}\rvert$	$\lvert\mathsf{SeedSec}\rvert$	$\lvert\mathsf{SeedPub}\rvert$	$\lvert\tau\rvert$
I	31	61	223	207	430	32	16	32
III	61	89	366	350	716	32	16	32
V	61	127	509	493	1002	32	16	32

5.4 Verification process

Algorithm 3: $\mathsf{Verify}(\mathsf{pk}, \mathcal{M}, \sigma)$
Input: public key $\mathsf{pk} = (\mathsf{SeedPub}, \{A^0_{rk}\}_{r,k})$, message $\mathcal{M}$, signature $\sigma = (\Sigma, \tau)$.
Output: accept/reject.

1. From $\mathsf{SeedPub}$, recompute the matrices A^*_{rk} and complete them with the published blocks A^0_{rk} to form A_{rk}.
2. Compute $H := (h_1, \ldots, h_m) \leftarrow \mathcal{H}(\mathcal{M} \mid \tau)$.
3. For $r = 1, \ldots, m$, compute

$$\mathbf{Pub}_r(\Sigma) \leftarrow \sum_{k=1}^{\ell} \Sigma^{[a_{rk}]} A_{rk}\, {}^t\Sigma^{[b_{rk}]}.$$

4. Accept iff $h_r = \mathbf{Pub}_r(\Sigma)$ for all $r = 1, \ldots, m$; otherwise reject.

6 Parameters and implementation results

We instantiate UFOs with $o = m = 1$, $v = 4$, $n = 5$, and $\ell = 5$. Exponent schedules and concrete sizes/timings are summarized in Tables 3, 4 and 5.

7 Conclusions

UFOs compresses UOV public keys by restricting each public equation to a sparse sum of Frobenius blocks over $\mathbb{F}_{p^e}$, while retaining UOV-style signing via an $\mathbb{F}_p$-linear solve after fixing vinegar variables. The concrete parameters above yield compact keys and fast verification; improving the analysis of exponent

Table 5. Reference timings (Online Magma Calculator [6]).

Level	p	e	o	v	ℓ	Gen	Sign	Verif
I	31	61	1	4	5	1.7ms	7.1ms	1.0ms
III	61	89	1	4	5	2.9ms	13.1ms	1.7ms
V	61	127	1	4	5	3.2ms	20ms	2.1ms

schedules and understanding broader multivariable direct models remain key open problems.

Auxiliary Supporting Material

The material below is optional supporting material and is not required to read the main paper.

A Proposed parameters

Parameter selection. For the concrete UFOs instances considered in this work, we keep

$$o = m = 1, \qquad v = 4, \qquad n = 5, \qquad \ell = 5,$$

and select the field parameters (p, e) according to the targeted security level. Once these dimensions are fixed, the main remaining design parameter is the exponent schedule

$$\{(a_k, b_k)\}_{k=1}^{5} \subset (\mathbb{Z}/e\mathbb{Z})^2,$$

which determines the five Frobenius blocks of the public map. The original proposal chooses indices so that the differences $b_k - a_k$ are pairwise distinct and spread around $(e-1)/2$ [24]. However, the exponent-only analysis of Bariant shows that the first published schedules admit a low-degree bi-variate resultant representation after canceling two blocks, and proposes replacing only the ten exponents while keeping the dimensions, key format, signature format, signing procedure, and verification algorithm unchanged [2]. Following this methodology, we adopt the candidate schedules of Table 3, which are selected by screening all two-block cancellations against the direct univariate and bivariate models enumerated in [2]. These schedules should be understood as exponent-only hardened candidates: they improve the currently known direct-model diagnostics without changing the UFOs interface, but they do not constitute a proof of security against unrestricted multivariable attacks.

Sizes.

- Salt $|\tau|$: 256 bits (32 bytes),
- $|\mathsf{SeedSec}|$: 256 bits (32 bytes),
- $|\mathsf{SeedPub}|$: 128 bits (16 bytes),
- Signature size (in bytes): $|\sigma| = \lceil ne\lceil \log_2(p)\rceil/8\rceil + |\tau|/8$,
- Public key size (in bytes): $|\mathbf{Pub}| = \lceil mo^2 e\ell\lceil \log_2(p)\rceil/8\rceil + |\mathsf{SeedPub}|/8$,
- Secret key size (in bytes): $|\mathsf{Sec}| = |\mathsf{SeedSec}|/8$.

Table 6. Direct-model diagnostics for the first published and selected exponent schedules, reproduced from [2]. Here L is the shortest cyclic interval containing the six residual endpoints after canceling two blocks, R and $R-1+E$ come from the best two-anchor bivariate resultant model, (w_3, w_4, w_5) are the consecutive-window orbit diagnostics, and the endpoint span is the length of the shortest cyclic interval containing all ten endpoints.

Level	Schedule	L	R	$R-1+E$	(w_3, w_4, w_5)	Endpoint span
I	[24]	31	7	8	$(9, 16, 25)$	33
I	this article	34	49	61	$(18, 30, 53)$	47
III	[24]	45	7	8	$(9, 16, 25)$	47
III	this article	44	65	77	$(23, 36, 49)$	74
V	[24]	64	7	8	$(9, 16, 25)$	66
V	this article	61	84	112	$(30, 44, 80)$	103

Public output mixing of the Frobenius blocks. A natural hardening option is to post-process each public block $\mathbf{Pub}_k$ by a public $\mathbb{F}_p$-linear endomorphism of $\mathbb{F}_q$, derived from the same public seed.

More precisely, for each block index k, one may define

$$T_k(x) = \sum_{i=0}^{e-1} t_{k,i}\, x^{[i]}, \qquad t_{k,i} \in \mathbb{F}_q,$$

where the coefficients $t_{k,i}$ are deterministically expanded from $\mathsf{SeedPub}$. The public map would then use the mixed blocks

$$\widetilde{\mathbf{Pub}}_k = T_k \circ \mathbf{Pub}_k, \qquad \widetilde{\mathbf{Pub}} = \sum_{k=1}^{\ell} \widetilde{\mathbf{Pub}}_k.$$

Since each T_k is $\mathbb{F}_p$-linear, this transformation preserves the basic $\mathbb{F}_p$-quadratic structure of the scheme after descent: each coordinate of $T_k(\mathbf{Pub}_k(X))$ remains an $\mathbb{F}_p$-quadratic form in the base-field coordinates of X. In particular, the signing procedure is still compatible with the usual UOV inversion strategy. Indeed, once the vinegar variables are fixed, each block $\mathbf{Pub}_k(X\Lambda + V)$ contributes terms that are $\mathbb{F}_p$-linear in the oil parameter, and applying the $\mathbb{F}_p$-linear map T_k preserves this linearity. Hence the global inversion problem remains an $\mathbb{F}_p$-linear system of dimension e in the case $o = m = 1$, just as in the original construction. Verification is also unchanged in nature: one simply evaluates the mixed public map $\widetilde{\mathbf{Pub}}$ instead of $\mathbf{Pub}$.

The main interest of this modification is algebraic densification. In the original design, each block $\mathbf{Pub}_k(X) = X^{[a_k]} A_k\, {}^tX^{[b_k]}$ contributes a single Frobenius monomial pattern at the output level. After composition with T_k, the same block gives

$$T_k(\mathbf{Pub}_k(X)) = \sum_{i=0}^{e-1} t_{k,i}\, \mathbf{Pub}_k(X)^{[i]},$$

that is, a public $\mathbb{F}_p$-linear combination of all Frobenius shifts of the block. Thus, even though the hidden trapdoor relation is unchanged, the visible output of each block becomes substantially denser from the point of view of direct algebraic representations. In particular, attacks that rely on identifying a sparse low-dimensional model for individual blocks—for instance by isolating a few Frobenius layers, matching a low-degree bivariate rewriting, or exploiting the direct visibility of the original exponents—appear less natural after such mixing, since each block is no longer observed through a single Frobenius profile but through a full $\mathbb{F}_p$-linear orbit.

Heuristically, this should strengthen resistance against direct structured attacks, because it hides the output sparsity that underlies several elimination-based or layer-based representations. At the same time, the modification does not seem to affect the core functionality of the scheme: the system remains $\mathbb{F}_p$-linear on the oil variables during signing, and public verification still consists of evaluating a structured quadratic map over $\mathbb{F}_q$. Of course, this argument is only structural: the actual cryptanalytic effect of such public output mixing would require a dedicated analysis, since the resulting scheme still retains strong Frobenius symmetries and may introduce new linear dependencies at the output level.

B Implementation and Performances

We implemented a reference proof-of-concept implementation of the three UFOs functions (`generate`, `sign`, `verify`) in Magma language [7]. The source code is available at

https://anonymous.4open.science/r/FUOV-BD39/ref%20F-uov.txt.

The performance results reported in Tab. 5 were obtained using the online Magma Calculator [6]. Our benchmarks cover key generation, inversion of the public key, and verification. Message-and-salt hashing and coefficient expansion from seeds are not implemented; instead, these steps are emulated by sampling random values. A complete implementation should include the corresponding computational costs.

Cost estimates for key generation, signing, and verification. The computational cost of UFOs can be expressed directly in terms of the scheme parameters (p, e, o, v, m, ℓ), with $n = o + v$ and typically $m = o$. All costs below are counted in operations over $\mathbb{F}_{p^e}$, except when an $\mathbb{F}_p$-linear solve is explicitly mentioned.

Key generation. Key generation consists of:

1. sampling the secret oil-space description $\Lambda^* \in \mathbb{F}_{p^e}^{o \times v}$;
2. sampling the public "known part" A^*_{rk} for $1 \le r \le m$ and $1 \le k \le \ell$;
3. reconstructing the missing oil–oil blocks A^0_{rk} through

$$A^0_{rk} = -\Lambda^{*[a_{rk}]} A^1_{rk} - A^2_{rk}\, {}^t\Lambda^{*[b_{rk}]} - \Lambda^{*[a_{rk}]} A^3_{rk}\, {}^t\Lambda^{*[b_{rk}]},$$

or equivalently through the compact matrix formula given in the scheme description.

For each pair (r, k), the dominant term is the multiplication of an $o \times v$ matrix, a $v \times v$ matrix, and a $v \times o$ matrix, which costs on the order of

$$O(ov^2 + o^2v + o^2)$$

operations in $\mathbb{F}_{p^e}$, and simplifies to $O(ov^2)$ in the usual regime $o \leq v$. Hence the total key-generation cost is

$$O(m\ell\, ov^2)$$

operations in $\mathbb{F}_{p^e}$, up to lower-order sampling and Frobenius-application costs. In the proposed parameter sets, where $m = o = 1$, this reduces to

$$O(\ell v^2),$$

that is, essentially linear in the number of blocks and quadratic in the number of vinegar variables.

Signing. Signing samples a vinegar vector $V = (0, \ldots, 0, V^*)$ and solves, over $\mathbb{F}_p$, the linearized oil equation

$$\Phi(\Gamma) = \Phi(X)\, M(V), \qquad \Gamma = H(m\|\tau) - \mathrm{Pub}(V).$$

The main costs are:

1. evaluating $\mathrm{Pub}(V)$, which requires ℓm Frobenius blocks;
2. constructing the matrix $M(V)$ of size $(oe) \times (me)$ over $\mathbb{F}_p$;
3. solving the resulting linear system over $\mathbb{F}_p$.

The evaluation of $\mathrm{Pub}(V)$ costs about

$$O(m\ell\, v^2)$$

operations in $\mathbb{F}_{p^e}$, since only the vinegar–vinegar blocks contribute when $V = (0 \mid V^*)$. The construction of $M(V)$ is of the same asymptotic order, up to the cost of applying Frobenius maps and expanding to the fixed $\mathbb{F}_p$ basis. Finally, solving the square linear system over $\mathbb{F}_p$ costs

$$O((oe)^\omega)$$

operations in $\mathbb{F}_p$ with matrix-multiplication exponent ω, or $O((oe)^3)$ with classical Gaussian elimination. Expressed at the highest level, signing therefore consists of a moderate number of structured $\mathbb{F}_{p^e}$ operations together with one linear solve of dimension oe over $\mathbb{F}_p$. For the proposed instances, $o = 1$, so the solve dimension is exactly e.

Verification. Verification recomputes the public matrices from the seed, evaluates

$$\mathrm{Pub}_r(\Sigma) = \sum_{k=1}^{\ell} \Sigma^{[a_{rk}]} A_{rk}\, {}^t\Sigma^{[b_{rk}]}, \qquad 1 \leq r \leq m,$$

and compares the result with the hash output. Each Frobenius block evaluation on a full vector of length n costs

$$O(n^2)$$

operations in $\mathbb{F}_{p^e}$, up to Frobenius applications. Thus verification costs

$$O(m\ell\, n^2)$$

operations in $\mathbb{F}_{p^e}$. In the proposed regime $m = 1$, this becomes simply

$$O(\ell n^2).$$

Unlike signing, verification does not require solving any linear system and is therefore especially lightweight.

Interpretation for the proposed parameter sets. With the concrete choices $m = o = 1$, $v = 4$, $n = 5$, and $\ell = 5$, the asymptotic expressions simplify considerably:

$$\mathsf{KeyGen} = O(\ell v^2), \qquad \mathsf{Verify} = O(\ell n^2),$$

while signing is dominated by:

- one evaluation of a very small structured public map,
- one construction of an $e \times e$ linear system over $\mathbb{F}_p$,
- one solve of that $e \times e$ system.

This explains the implementation profile observed in practice: verification is the fastest operation, key generation is also inexpensive, and signing is dominated by the base-field linear algebra step.

Qualitative comparison with other compact MQ signatures. From the viewpoint of high-level complexity, UFOs belongs to the same general family as compact structured MQ signatures such as MAYO, QR-UOV, and SNOVA: the public key is compressed by replacing a generic dense quadratic map with a structured subclass, while signing retains the basic UOV paradigm of fixing vinegar coordinates and solving a linear system in the oil variables. The main distinguishing feature of UFOs is that the compression is driven by *Frobenius sparsity over* $\mathbb{F}_{p^e}$. As a consequence:

- the public key size scales like

  $$O(o^3 \ell e \log p)$$

  bits in compressed form;
- the signature size scales like

  $$O((o+v)e \log p)$$

  bits;

– the online signing cost is governed by linear algebra in dimension oe over $\mathbb{F}_p$, rather than by solving a large dense MQ system.

Compared with other MQ candidates, this yields a different trade-off:

1. **Compared with generic UOV-like systems,** the public key is substantially smaller because only ℓ Frobenius blocks per equation are stored instead of a full quadratic form.
2. **Compared with ring- or module-structured variants such as QR-UOV or SNOVA,** the structure is concentrated in the Frobenius exponents rather than in a larger noncommutative or quotient-ring algebraic layer; this tends to keep the online algorithms conceptually simple, at the price of a more delicate exponent-schedule analysis.
3. **Compared with MAYO-like compact designs,** UFOs similarly aims at reducing key size while preserving efficient verification, but it does so through extension-field Frobenius structure rather than through the specific low-rank/oil-vinegar tensor organization used there.

Practical takeaway. In short, the computational profile of UFOs is:

– *very small secret key*,
– *compact public key*,
– *moderate signature size*,
– *very fast verification*,
– and *signing dominated by one structured linear solve over* $\mathbb{F}_p$.

This places UFOs in the same application niche as the most compact MQ signatures currently considered in the post-quantum literature, namely settings where public-key size and verification cost matter more than achieving the smallest possible signing latency.

C A toy example

We consider the scheme in the regime $o = m = 1$, $v = \ell = 2$, leaving p and e unfixed. This is a toy setting, and we make no security claim.

The purpose of this appendix is purely illustrative. Working in the smallest nontrivial regime makes the public map, the embedded isotropy constraints defining the secret oil line, and the geometry of the solution set fully explicit. In particular, we use this setting to heuristically gauge how often the embedded constraints admit a small number of *equivalent* secret keys, and to show that restricting $\mathbf{Pub}(X) = 0$ to random lines leads to sparse univariate Frobenius-type equations whose solutions may include a few $\mathbb{F}_q$-lines (including the trapdoor) together with many structured $\mathbb{F}_p$-components.

We fix the following toy-instance data:

– a secret oil line parametrized by $\Lambda = (1, \lambda_2, \lambda_3) \in \mathbb{F}_q^3$;

– integers $a_1, b_1, a_2, b_2 \in \{0, \ldots, e-1\}$ such that (w.l.o.g.) $a_1 > b_1$, $a_2 > b_2$, $a_1 > a_2$, $b_1 = 0$ or $b_2 = 0$, and $b_1 \neq b_2$;
– matrices $A_1, A_2 \in \mathbb{F}_q^{3\times 3}$.

We consider the public map

$$\mathbf{Pub}(X) = X^{[a_1]} A_1 \, {}^t X^{[b_1]} + X^{[a_2]} A_2 \, {}^t X^{[b_2]},$$

with embedded (isotropy) constraints

$$\Lambda^{[a_1]} A_1 \, {}^t \Lambda^{[b_1]} = 0, \qquad \Lambda^{[a_2]} A_2 \, {}^t \Lambda^{[b_2]} = 0. \tag{12}$$

C.1 Some insights into the embedded equation

We estimate how many *equivalent secret keys* (i.e., solutions (λ_2, λ_3) of (12)) a random toy instance may admit. We run the following experiment:

– Sample λ_2, λ_3, and all coefficients of A_1 and A_2 except α_{111} and α_{211} (their oil×oil coefficients).
– Solve (12) to recover the two missing coefficients α_{111} and α_{211}.
– With all coefficients now fixed, solve (12) again for (λ_2, λ_3) and count the number of solutions (at least one by construction).

The results suggest that the probability that a toy UFOs instance has s equivalent keys is well-approximated by $\frac{\exp(-1)}{(s-1)!}$ (Poisson law), matching the solution-count distribution of a generic random system conditioned on having at least one solution [15,9,1,16,21]. This is also consistent with (12) being well-posed.

In Tab. 7, we report observed counts among 100,000 random toy instances[5] for $p = 7$, $e = 7$ (with fixed $a_1 = 4$, $b_1 = 2$, $a_2 = 1$, $b_2 = 0$), together with the theoretical frequency $\frac{\exp(-1)}{(s-1)!}$. We estimate that a random toy instance has on average two[6] alternate secret keys that could be used to invert the public key. Moreover, the probability of having more than s such keys decreases rapidly with s, since it is $O(1/s!)$.

Table 7. Number of systems having a given number of equivalent keys for $p = 7$, $e = 7$, $a_1 = 4$, $b_1 = 2$, $a_2 = 1$, $b_2 = 0$ and 100,000 draws.

# of solutions	1	2	3	4	5	6	7	8
Observation	36575	37113	18348	6035	1545	332	45	7
Expectation	36787.9	36787.9	18393.9	6131.3	1532.8	306.6	51.1	7.3

[5] Of course, we are forced to choose small parameters to perform the experiment; the qualitative behavior may extend to other parameters, but we do not claim a proof.
[6] Computing these alternate keys becomes challenging for large values of p and e.

C.2 Some insights into the public equation

We also estimate how the complexity of solving a single equation of the form $X^{[a]}A\,{}^tX^{[b]} = h$ for $X = (x_1, x_2, x_3)$ depends on e. By fixing all variables but one, we obtain a univariate equation with monomials $x^{p^a+p^b}$, x^{p^a}, x^{p^b}, and a constant term. Its degree may be arbitrarily large; however, converting it into a quadratic system over $\mathbb{F}_p$ yields a structured instance. Since $a \neq b$, the resulting system can be viewed as Matsumoto–Imai-type (and such systems are known to be broken).[7]

We also estimate how the solutions of $\mathbf{Pub}(X) = 0$ are distributed. We first reduce to a univariate equation by choosing $U \in \mathbb{F}_q^3$ and looking for solutions on the line spanned by U, i.e., solving $\mathbf{Pub}(xU) = 0$. For convenience, let $U_1 = U^{[a_1]}A_1\,{}^tU^{[b_1]}$ and $U_2 = U^{[a_2]}A_2\,{}^tU^{[b_2]}$. The equation becomes $x^{p^{a_1}+p^{b_1}}U_1 + x^{p^{a_2}+p^{b_2}}U_2 = 0$. We distinguish the following cases.

- $U_1 = U_2 = 0$. Then the whole line spanned by U consists of solutions. The probability of finding such a U by chance is $1/p^{2e}$. Heuristically, finding such a U algorithmically should be costly (one expects at least generic square-root-type work in the relevant search space, i.e., about $O(p^{e/2})$).
- $U_1 = 0$, $U_2 \neq 0$ or $U_1 \neq 0$, $U_2 = 0$. Finding such U is easy (see above), but the only solution is the trivial one $x = 0$.
- $U_1 \neq 0$ and $U_2 \neq 0$. This is the most common case. The equation rewrites as $x^{p^{a_2}+p^{b_2}}\left(x^{p^{a_1}+p^{b_1}-p^{a_2}-p^{b_2}} + U_2/U_1\right) = 0$. The solutions are $x = 0$ and the solutions of $x^{p^{a_1}+p^{b_1}-p^{a_2}-p^{b_2}} = -U_2/U_1$. Let $\delta = \gcd(p^e - 1,\, p^{a_1} + p^{b_1} - p^{a_2} - p^{b_2})$. Then $x^{p^{a_1}+p^{b_1}-p^{a_2}-p^{b_2}} = -U_2/U_1$ has δ solutions if $(-U_2/U_1)^{\frac{p^e-1}{\delta}} = 1$ (which occurs with probability $1/\delta$ under the heuristic that $-U_2/U_1$ is uniformly distributed in $\mathbb{F}_q^*$ when $U_1 \neq 0$), and no solution otherwise. For the choice above ($p = e = 7$), one has $\delta = p - 1$. In summary, the solutions form an $\mathbb{F}_p$-line with probability $1/(p-1)$, and are reduced to $\{0\}$ otherwise.

In summary, the solutions of $\mathbf{Pub}(X) = 0$ in this toy UFOs regime form a set of s $\mathbb{F}_q$-lines (including the original trapdoor) and s' seemingly random $\mathbb{F}_p$-lines, where s is about 2 on average (with an exponentially decreasing tail) and $s' \approx p^{2e-1}$. All solutions are grouped into $\mathbb{F}_p$-lines due to $\mathbb{F}_p$-homogeneity. See Fig. 2.

D A Possible Public Output Mixing Layer

This appendix discusses a possible hardening variant of UFOs in which each public Frobenius block is post-processed by a public $\mathbb{F}_p$-linear endomorphism of $\mathbb{F}_q$. The purpose of this modification is not to alter the trapdoor paradigm, but rather to make the public algebraic representation of each block less sparse and less directly tied to a single visible exponent pair. In view of the direct algebraic

[7] We recall that (surprisingly, in odd characteristic) the equation $x^2 = y$ over the extension field, converted into a quadratic system over the base field, can have exponential solving complexity. See [12].

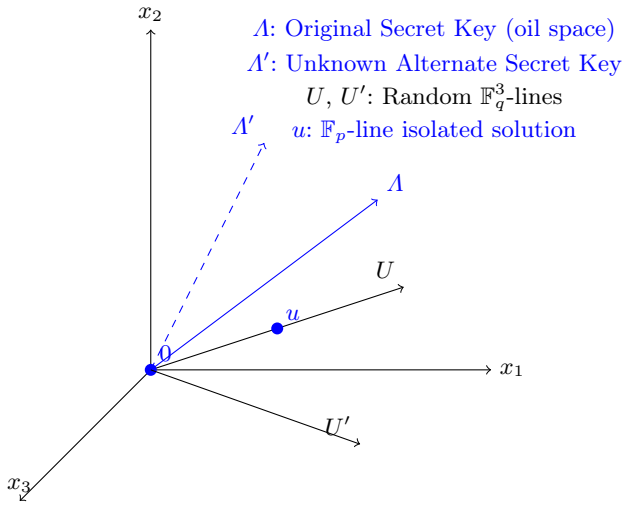

Fig. 2. Locus of $\mathbf{Pub}(x_1, x_2, x_3) = 0$ (in blue)

attacks discussed in the main text and in [2], this kind of public output mixing appears as a natural design option worth recording.

D.1 Definition of the output mixing layer

In UFOs the public map decomposes as

$$\mathbf{Pub}(X) = \sum_{k=1}^{\ell} \mathbf{Pub}_k(X), \qquad \mathbf{Pub}_k(X) = X^{[a_k]} A_k \, {}^t X^{[b_k]},$$

with $A_k \in \mathbb{F}_q^{n \times n}$. We modify the public key by post-composing each block with a public $\mathbb{F}_p$-linear endomorphism $T_k \in \mathrm{End}_{\mathbb{F}_p}(\mathbb{F}_q)$ of the form

$$T_k : \mathbb{F}_q \to \mathbb{F}_q, \qquad T_k(z) = \sum_{i=0}^{e-1} t_{k,i} \, z^{[i]},$$

where $t_{k,i} \in \mathbb{F}_q$ are deterministically derived from the public seed SeedPub. This yields

$$\widetilde{\mathbf{Pub}}(X) = \sum_{k=1}^{\ell} \widetilde{\mathbf{Pub}}_k(X), \qquad \widetilde{\mathbf{Pub}}_k(X) = T_k(\mathbf{Pub}_k(X)).$$

The maps T_k are fully public and introduce no additional secret material.

D.2 Compatibility with the original trapdoor structure

This transformation preserves the signing trapdoor: after fixing the vinegar variables, each block remains $\mathbb{F}_p$-linear in the oil coordinates, and post-composition with $\mathbb{F}_p$-linear T_k maintains $\mathbb{F}_p$-linearity.

More explicitly, in the original scheme one writes a signing candidate under the form

$$X\Lambda + V,$$

where Λ is a basis of the secret oil subspace and V is a random vinegar vector. By construction of the oil space, for each block one has

$$\mathbf{Pub}_k(X\Lambda + V) = X^{[a_k]}\Lambda^{[a_k]}A_k\,{}^tV^{[b_k]} + V^{[a_k]}A_k\,{}^t\Lambda^{[b_k]}\,{}^tX^{[b_k]} + \mathbf{Pub}_k(V),$$

because the oil–oil term vanishes identically. The nonconstant part is therefore $\mathbb{F}_p$-linear in the Frobenius transforms of the oil parameter X. Applying T_k gives

$$T_k(\mathbf{Pub}_k(X\Lambda + V)) = T_k\Big(X^{[a_k]}\Lambda^{[a_k]}A_k\,{}^tV^{[b_k]} + V^{[a_k]}A_k\,{}^t\Lambda^{[b_k]}\,{}^tX^{[b_k]}\Big) + T_k(\mathbf{Pub}_k(V)),$$

and this remains $\mathbb{F}_p$-linear in X, since T_k itself is $\mathbb{F}_p$-linear. Accordingly, after descent to $\mathbb{F}_p$, the signing step still reduces to solving an $\mathbb{F}_p$-linear system of the same dimension as in the original construction. In particular, for the concrete parameters $o = m = 1$, the inversion problem remains an $e \times e$ linear system over $\mathbb{F}_p$.

Verification is equally unaffected in nature. The verifier simply recomputes the public matrices from the public seed, derives the corresponding maps T_k, and checks whether

$$\widetilde{\mathbf{Pub}}(\Sigma) = H(m\|\tau).$$

Thus the interface of the scheme remains unchanged, and the extra work consists only in evaluating the chosen $\mathbb{F}_p$-linear output maps.

D.3 Preservation of the descended algebraic type

A second structural point is that the modified scheme remains quadratic over the base field. Each original block $\mathbf{Pub}_k$ is a Frobenius quadratic form over $\mathbb{F}_q$, and after expansion in an $\mathbb{F}_p$-basis it becomes a family of ordinary quadratic forms over $\mathbb{F}_p$. Since T_k is $\mathbb{F}_p$-linear, the composed block $T_k \circ \mathbf{Pub}_k$ is still an $\mathbb{F}_p$-quadratic map after descent.

Equivalently, in extension-field notation one has

$$T_k(\mathbf{Pub}_k(X)) = \sum_{i=0}^{e-1} t_{k,i}\,\mathbf{Pub}_k(X)^{[i]}.$$

Using

$$\mathbf{Pub}_k(X)^{[i]} = X^{[a_k+i]}A_k^{[i]}\,{}^tX^{[b_k+i]},$$

this becomes

$$T_k(\mathbf{Pub}_k(X)) = \sum_{i=0}^{e-1} t_{k,i}\,X^{[a_k+i]}A_k^{[i]}\,{}^tX^{[b_k+i]}.$$

Hence each mixed block is itself a sum of Frobenius blocks, now ranging over all shifts of the original exponent pair. The construction therefore stays within the same broad algebraic universe: the public map is still built from Frobenius quadratic components, but each original visible block is replaced by a denser superposition of its entire Frobenius orbit.

D.4 Direct algebraic effect: densification of each public block

This densification is the main reason why such an output layer may be relevant from a cryptanalytic viewpoint. In the original design, a public block is exposed through a single exponent pair (a_k, b_k). Even though the full public map is a sum of several blocks, each component remains individually sparse and therefore potentially accessible to attacks that attempt to isolate a few blocks, cancel some of them, and exploit the low-dimensional structure of the residual equation.

After composition with T_k, the public contribution of the k-th block is no longer attached to a single visible pair, but to the full collection

$$(a_k + i, b_k + i), \qquad i = 0, \ldots, e-1,$$

weighted by the public coefficients $t_{k,i}$. From the attacker's point of view, the output of one block is therefore no longer a single sparse Frobenius profile, but an $\mathbb{F}_p$-linear combination of all its shifts. This tends to blur the original blockwise sparsity in at least three immediate ways.

First, a single block can no longer be read as a single bilinear relation between two visible Frobenius layers. Instead, it simultaneously contributes to many such layer interactions. This makes the direct identification of a residual equation with a very small number of exponent pairs less transparent.

Second, attacks based on restricting the public map to a public line and then trying to derive a very small univariate or bivariate model become less natural. Indeed, the direct resultant attack of [2] crucially exploits the fact that after cancelling two blocks, the remaining three terms still correspond to a very sparse visible exponent pattern admitting two compatible low-degree bivariate rewritings. If each residual block has first been mixed with all its Frobenius shifts, the attacker no longer sees one term per block in that sense, but rather a denser superposition whose elimination profile is much less obvious.

Third, the various orbit-style diagnostics discussed in [2] become harder to interpret directly at the level of visible blocks. Those diagnostics are driven by the placement of a small number of explicit endpoints or layer differences. Public output mixing does not destroy the underlying Frobenius symmetries, but it spreads each block over its full orbit at the output level, so that the naive visible endpoint schedule is no longer the relevant one for direct algebraic modeling.

D.5 Why the signing algorithm still works

For completeness, let us record the signing argument in a slightly more explicit form. In the original F-UOV inversion step, after fixing the vinegar part V, one

must solve

$$\mathbf{Pub}(X\Lambda + V) = H$$

for the oil parameter $X \in \mathbb{F}_q^o$. Since all oil–oil terms vanish, the map

$$X \longmapsto \mathbf{Pub}(X\Lambda + V) - \mathbf{Pub}(V)$$

is $\mathbb{F}_p$-linear. The modified inversion problem becomes

$$\widetilde{\mathbf{Pub}}(X\Lambda + V) = H,$$

that is,

$$\sum_{k=1}^{\ell} T_k(\mathbf{Pub}_k(X\Lambda + V)) = H.$$

Subtracting the constant term gives

$$\sum_{k=1}^{\ell} T_k(\mathbf{Pub}_k(X\Lambda + V) - \mathbf{Pub}_k(V)) = H - \widetilde{\mathbf{Pub}}(V).$$

Each summand is the image under an $\mathbb{F}_p$-linear map of an $\mathbb{F}_p$-linear form in X, and is therefore still $\mathbb{F}_p$-linear in X. Summing over k preserves this property. Hence the global system remains $\mathbb{F}_p$-linear in the unknown oil coordinates.

Accordingly, the only practical change in signing is that the matrix of the linear system must now incorporate the public endomorphisms T_k. This increases the constant factors in the computation, but does not alter the asymptotic structure of the trapdoor inversion procedure.

D.6 Possible implementation choices

Several implementation choices are possible for the maps T_k.

A first possibility is to sample a fully dense $\mathbb{F}_p$-linear map, i.e., to use all coefficients $t_{k,i}$. This gives maximal output spreading, but also increases the cost of evaluating the public map.

A second possibility is to restrict each T_k to a sparse linearized polynomial, for instance with only a few nonzero coefficients. This already destroys the one-pair visibility of a block while keeping evaluation costs lower.

A third possibility is to require T_k to be invertible as an $\mathbb{F}_p$-linear map. This is not necessary for correctness, since signing and verification only require public evaluation, but invertibility may be desirable if one wants to avoid collapsing too much information at the output level or introducing trivial rank defects. In practice, invertibility can be tested during parameter generation, or enforced by deriving T_k from a public seed until a suitable map is obtained.

D.7 Security interpretation

From the viewpoint of the two analyses considered in this work, the expected effect of public output mixing is favorable. The first part of this work emphasizes the role of sparse Frobenius representations in obtaining compact public keys and discusses direct algebraic attacks through line restrictions and common isotropic subspaces. The second part shows that for the published exponent schedule, a particularly small residual three-term equation admits a low-degree bivariate elimination attack. Both perspectives rely, in one form or another, on the fact that the public blocks remain individually sparse and directly visible.

The proposed output layer acts precisely on this point: it keeps the hidden oil space and the signing mechanism intact, but it makes each visible block much less elementary. Heuristically, this should increase the difficulty of:

- isolating a very small number of residual visible blocks after cancellation;
- recognizing a low-degree bivariate rewriting tied to a specific exponent schedule;
- constructing very small orbit-style or layer-based direct models from the public block decomposition alone.

In that sense, public output mixing appears to be a plausible hardening option against the currently known direct attacks.

D.8 Limitations and caveats

At the same time, this should not be overstated. The argument above is structural and heuristic, not a proof of increased security.

First, the mixed scheme still has strong Frobenius symmetries. The maps T_k are themselves linearized polynomials, so they do not remove the semilinear nature of the construction; they only make it less sparse at the visible output level.

Second, public output mixing may introduce new dependencies. For example, if some T_k have low rank or interact in a special way with the original blocks, one might obtain unexpected linear relations among coordinates of the public map. These would have to be excluded by parameter design or analyzed cryptanalytically.

Third, while the current resultant-style attack seems to rely on visible block sparsity, it is conceivable that more global multivariable or orbit-based attacks could be adapted to the mixed setting. In particular, once all Frobenius shifts of a block are made public simultaneously, new elimination strategies may become available at a different level of abstraction.

For these reasons, the present proposal should be viewed as a possible hardening direction rather than as a validated repair.

D.9 Summary

To summarize, composing each public block with a public $\mathbb{F}_p$-linear map

$$T_k(z) = \sum_{i=0}^{e-1} t_{k,i} z^{[i]}$$

derived from SeedPub appears to preserve all essential functional features of UFOs:

- the public system remains $\mathbb{F}_p$-quadratic after descent;
- the signing equation remains $\mathbb{F}_p$-linear in the oil variables once the vinegar variables are fixed;
- verification remains a straightforward evaluation of the public map.

Its main direct effect is to replace each visible Frobenius block by a public $\mathbb{F}_p$-linear combination of all its Frobenius shifts, thereby densifying the public output and obscuring the sparse exponent structure exploited by the currently known direct attacks. This makes public output mixing a plausible hardening option, although its precise cryptanalytic effect remains an open question.

E Conclusions

We proposed UFOs, a multivariate signature scheme that combines a UOV-style trapdoor with public equations restricted to sparse sums of Frobenius forms, leading to compact keys.

For our Level I parameter set, the public key consists of 5×61 field elements in $\mathbb{F}_{31}$ together with a short seed, totaling about 1600 bits (roughly 206 bytes). Signatures follow a hash-and-sign paradigm. Signing mainly requires arithmetic in $\mathbb{F}_{31}$ and inversion of a 61×61 linear system. Verification is performed over an extension of $\mathbb{F}_{31}$ and involves only a small number of multiplications. Each signature contains 5×61 field elements (about 1600 bits) plus a 256-bit salt, i.e., under 250 bytes.

At these parameters, public keys and signatures are competitive in size with widely deployed schemes and leading post-quantum candidates, while keeping operations simple.

From a security standpoint, the construction retains the UOV trapdoor structure and reduces the main attack avenues to either (i) recovering an equivalent oil subspace from the embedded isotropy constraints, or (ii) solving the resulting sparse Frobenius MQ systems over $\mathbb{F}_q$ (or, after expansion, quadratic systems over $\mathbb{F}_p$). We provide concrete cost estimates for these attacks in the corresponding models.

Several directions remain open, including:

- a tighter analysis of the embedded-constraint systems, particularly in the setting where the public "known part" (i.e. the matrices A^*) is fixed once and for all, possibly sparse,
- optimized, constant-time, leakage-resilient implementations,
- extensions such as threshold variants of UFOs.

References

1. Arratia, R., Barbour, A.D., Tavaré, S.: Logarithmic Combinatorial Structures: A Probabilistic Approach. Cambridge University Press (2003)
2. Bariant, A.: On exponent-only hardening of frobenius-UOV. Manuscript (2026), submitted to eprint iacr 2026
3. Beullens, W.: Improved cryptanalysis of UOV and Rainbow. In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021. pp. 348–373. Springer International Publishing, Cham (2021)
4. Beullens, W.: Mayo: Practical post-quantum signatures from oil-and-vinegar maps. In: AlTawy, R., Hülsing, A. (eds.) Selected Areas in Cryptography. pp. 355–376. Springer International Publishing, Cham (2022)
5. Bi, J., Cheng, Q., Rojas, J.M.: Sub-linear root detection, and new hardness results, for sparse polynomials over finite fields (2012), https://arxiv.org/abs/1204.1113
6. Bosma, W., Cannon, J., Playoust, C.: Online magma calculator. https://magma.maths.usyd.edu.au/calc/
7. Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system. I. The user language. J. Symbolic Comput. **24**(3-4), 235–265 (1997). https://doi.org/10.1006/jsco.1996.0125, http://dx.doi.org/10.1006/jsco.1996.0125, computational algebra and number theory (London, 1993)
8. Cabarcas, D., Li, P., Verbel, J., Villanueva-Polanco, R.: Improved attacks for SNOVA by exploiting stability under a group action. In: Annual International Cryptology Conference. pp. 358–389. Springer (2025)
9. Cam, L.L.: An approximation theorem for the Poisson binomial distribution. Pacific Journal of Mathematics **10**(4), 1181–1197 (1960)
10. Cantor, D.G., Zassenhaus, H.: A new algorithm for factoring polynomials over finite fields. Mathematics of Computation **36**(154), 587–592 (1981). https://doi.org/10.1090/S0025-5718-1981-0603967-0
11. Cox, D., Little, J., O'Shea, D.: Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra. Springer, 4 edn. (2015)
12. Ding, J., Clough, C., Araujo, R.: Inverting square systems algebraically is exponential. Finite Fields and Their Applications **26**, 32–48 (2014). https://doi.org/https://doi.org/10.1016/j.ffa.2013.10.004, https://www.sciencedirect.com/science/article/pii/S1071579713001093
13. Elman, R., Karpenko, N., Merkurjev, A.: The Algebraic and Geometric Theory of Quadratic Forms, American Mathematical Society Colloquium Publications, vol. 56. American Mathematical Society (2008)
14. Esser, A., Verbel, J.A., Zweydinger, F., Bellini, E.: SoK: CryptographicEstimators - a Software Library for Cryptographic Hardness Estimation. In: AsiaCCS. ACM (2024)
15. Flajolet, P., Odlyzko, A.M.: Random mapping statistics. In: Advances in Cryptology — EUROCRYPT '89. Lecture Notes in Computer Science, vol. 434, pp. 329–354. Springer (1990)
16. Flajolet, P., Soria, M.: Gaussian limiting distributions for the number of components in combinatorial structures. Journal of Combinatorial Theory, Series A **53**(2), 165–182 (1990)
17. Furue, H., Ikematsu, Y.: Key recovery attacks on UOV using p^l-truncated polynomial rings. Cryptology ePrint Archive, Paper 2026/298 (2026), https://eprint.iacr.org/2026/298

18. Furue, H., Ikematsu, Y., Kiyomura, Y., Takagi, T.: A new variant of unbalanced oil and vinegar using quotient ring: QR-UOV. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 187–217. Springer (2021)
19. von zur Gathen, J., Gerhard, J.: Modern Computer Algebra. Cambridge University Press, 3 edn. (2013)
20. Kadyrsizova, Z., Kenkel, J., Page, J., Singh, J., Smith, K., Vraciu, A., Witt, E.: Lower bounds on the F-pure threshold and extremal singularities. Transactions of the American Mathematical Society, Series B **9**(31), 977–1005 (Oct 2022). https://doi.org/10.1090/btran/106, https://www.ams.org/btran/2022-09-31/S2330-0000-2022-00106-0/
21. Knopfmacher, J.: Abstract Analytic Number Theory. North-Holland (1975)
22. Lidl, R., Niederreiter, H.: Finite Fields, Encyclopedia of Mathematics and its Applications, vol. 20. Cambridge University Press, 2 edn. (1997)
23. Lih-Chung, W., Chun-Yen, C., Jintai, D., Yen-Liang, K., Jan Adriaan, L., Ming-Siou, L., Bo-Shu, T., Po-En, T., Wang, C.C.: SNOVA proposal for NISTPQC: Additional digital signature schemes version 2.0 (2025), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/snova-spec-round2-web.pdf
24. Macario-Rat, G.: Frobenius-UOV: A very efficient multivariate public key signature scheme. Cryptology ePrint Archive, Paper 2026/968 (2026), https://eprint.iacr.org/2026/968
25. Nguyen, Q.D., Nguyen, M.H.: Towards compact UOV-based mq signatures: Rectangular and lifted whipping structures. Cryptology ePrint Archive (2026)
26. Pébereau, P.: One vector to rule them all: Key recovery from one vector in UOV schemes. In: International Conference on Post-Quantum Cryptography. pp. 92–108. Springer (2024)
27. Petzoldt, A.: Selecting and reducing key sizes for multivariate cryptography. Ph.D. thesis, Darmstadt University of Technology, Germany (2013), http://tuprints.ulb.tu-darmstadt.de/3523/
28. Ran, L.: Wedges, oil, and vinegar – an analysis of UOV in characteristic 2. Cryptology ePrint Archive, Paper 2025/1143 (2025), https://eprint.iacr.org/2025/1143
29. Riehm, C., Shrader-Frechette, M.: The equivalence of sesquilinear forms. Journal of Algebra **42**(2), 495–530 (1976). https://doi.org/https://doi.org/10.1016/0021-8693(76)90110-1, https://www.sciencedirect.com/science/article/pii/0021869376901101
30. Rousseau, É.: Arithmétique efficace des extensions de corps finis. Theses, Institut Polytechnique de Paris (Jul 2021), https://theses.hal.science/tel-03299466
31. Shoup, V.: A new polynomial factorization algorithm and its implementation. Journal of Symbolic Computation **20**(4), 363–397 (1995). https://doi.org/10.1006/jsco.1995.1055
32. Suzuki, T., Furue, H., Ito, T., Nakamura, S., Uchiyama, S.: An extended rectangular MinRank attack against UOV and its variants. Cryptology ePrint Archive, Paper 2025/739 (2025), https://eprint.iacr.org/2025/739
33. von zur Gathen, J., Panario, D.: Factoring polynomials over finite fields: A survey. Journal of Symbolic Computation **31**(1), 3–17 (2001). https://doi.org/https://doi.org/10.1006/jsco.1999.1002, https://www.sciencedirect.com/science/article/pii/S0747717199910027
34. Ward, B., Ming-Shing, C., Jintai, D., Boru, G., Matthias J., K., Jacques, P., Bo-Yuan, P., Dieter, S., Cheng-Jhih, S., Chengdong, T., Bo-Yin, Y.: UOV: Unbalanced oil and vinegar version 2.0 (2025), https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/uov-spec-round2-web.pdf

35. Wolf, C., Preneel, B.: Taxonomy of public key schemes based on the problem of multivariate quadratic equations. Cryptology ePrint Archive, Report 2005/077 (2005), https://eprint.iacr.org/2005/077