Frobenius-UOV: A Very Efficient Multivariate Public Key Signature Scheme

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

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

Abstract. We present Frobenius-UOV, 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: post-quantum signatures · multivariate cryptography · UOV · Frobenius forms · quadratic forms

1 Introduction

Modern digital infrastructures rely on digital signatures to authenticate software updates, protect online transactions, and secure communications. In the post-quantum setting, *multivariate* cryptography is an attractive direction: verification amounts to evaluating polynomial equations over a finite field, while security is linked to the apparent difficulty of solving the corresponding systems.

A prominent multivariate design is Unbalanced Oil and Vinegar (UOV) [31]. In UOV, the signer holds a trapdoor that enables efficient inversion of a structured quadratic system (to produce a signature), whereas verification consists of evaluating a public quadratic map.

Contributions. We introduce Frobenius-UOV (F-UOV), a UOV-type signature scheme whose public key admits a substantially more compact representation. The main idea is to restrict the public quadratic equations to a structured subclass built from *Frobenius forms*, i.e., sparse sums of terms of the shape $x^{p^a} y^{p^b}$ over an extension field. This preserves the standard UOV hash-and-sign workflow while reducing the number of coefficients needed to describe the public system.

Specifically, we provide:

- a complete specification of F-UOV (key generation, signing, and verification), including the derivation of the public system from a compact secret description;
- a security analysis that identifies the main attack avenues (direct algebraic attacks and key recovery) and isolates the computational problems induced by Frobenius sparsity;
- concrete parameter sets and a proof-of-concept implementation, reporting public-key and signature sizes as well as timings.

Technical overview. Shrinking multivariate public keys has been a central theme in recent proposals [22], including MAYO [3], QR-UOV [16], and SNOVA [21]. 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 [20,11]. More generally, both the base-field characteristic and the choice between homogeneous and affine public maps affect design and cryptanalysis [32,9].

Within the UOV family, F-UOV offers a favorable efficiency–security tradeoff:

- The secret oil space has dimension 1 (generated by a single vector), which limits the direct applicability of the "one-vector" viewpoint of Pébereau [23].
- With overwhelming probability, a public key determines a unique secret key; equivalently, the defining systems for the secret are typically overdetermined.
- Public keys are designed to be computationally indistinguishable from uniform data.
- 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 $n \approx \ell$):

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

  We do not know whether these problems admit polynomial-time algorithms; consequently, we do not have rigorous complexity bounds. We conjecture hardness and provide supporting estimates.
- The scheme depends on parameters o (oil variables), v (vinegar variables), p (characteristic/field size), e (extension degree), and ℓ (Frobenius sparsity limit). 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): $o \cdot (o+v)^2 \ell e \log_2(p)$ bits,
  - Secret key (compressed): $O(\lambda)$ bits,
  - Secret key (uncompressed): $o \cdot (o+v)^2 \ell e \log_2(p)$ bits,
  - Security: $e \log_2(p)/2$ bits (our assessment).

  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 [21] 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): $o^3 l^2 \log_2(p)$,
  - Secret key (uncompressed): $\big((v+2o+1)\, vo\, l^2 + 4l^4\big) \log_2(p)$,
  - Security: $O\big((vo^2 l^4)^{1/3}\big)$ (empirical[2] fit to the parameters in [21]).

  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 $[3, 4]$ [2]; ratios near 2.5 are considered aggressive, and ratios above 4 are often viewed as inefficient [23]. 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: $31^{30} \approx 2^{150}$), and we do not expect direct algebraic attacks to be effective. We therefore expect an adversary to descend to the base field $\mathbb{F}_p$, leading to systems with too many variables to be tractable; see Sec. 4.1.
- Public-key compression can be viewed as follows. Over $\mathbb{F}_p$, representing e generic quadratic equations in e variables (e.g., $\sum \alpha_{kij} x_i x_j$) costs $e \cdot \frac{e(e+1)}{2}$ 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). F-UOV reduces the number of monomials to $\sum_{i=1}^{\ell} \alpha_i X^{p^{a_i}} X^{p^{b_i}}$ with $\ell \ll \frac{e(e+1)}{2}$; we show that this sparsification does not compromise security.

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

[1] For smooth asymptotics, take, e.g., $p \approx e/2$; 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.

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., [27].

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 $f(x,y) = xy$, write $f(x,\cdot) : y \mapsto xy$ and $f(\cdot,y) : x \mapsto xy$. If $x = \sum_{i=1}^{e} x_i\xi^{i-1}$ and $y = \sum_{i=1}^{e} y_i\xi^{i-1}$, then $xy = \sum_{i,j=1}^{e} x_iy_j\,\xi^{i-1}\xi^{j-1}$. For $i,j \in \{1,\dots,e\}$, 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

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

Equivalently,

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

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 $k \in \{1,\dots,e\}$, 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 $(1,\xi,\dots,\xi^{e-1})$, 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 $Y \in \mathbb{F}_q^m$ we have

$$\Phi(XM\,{}^tY) = \big(\Phi_{\otimes n}(X)M_1\,{}^t\Phi_{\otimes m}(Y), \dots, \Phi_{\otimes n}(X)M_e\,{}^t\Phi_{\otimes m}(Y)\big).$$

In particular, right-multiplication by M induces an $\mathbb{F}_p$-linear map $\mathbb{F}_p^{en} \to \mathbb{F}_p^{em}$, i.e., there exists $\widetilde{M} \in \mathbb{F}_p^{en\times em}$ such that

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

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., $f(x,y) = xM\,{}^ty$, 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 [31] 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, while remaining indistinguishable from a generic quadratic system for an adversary.

All constructions are over an extension field $\mathbb{F}_q$ with $q = p^e$, where p is prime. Variables are partitioned into o oil and v vinegar variables, with $n = o + v$. As in classical UOV, the number of equations is set to $m = o$ (we keep the notation m when the equation count is to be emphasized). Taking e prime excludes proper intermediate subfields between $\mathbb{F}_p$ and $\mathbb{F}_q$, thereby restricting potential algebraic relations between public and secret data to these two natural field levels.

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 [26,18]. 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. 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. 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

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

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,

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

with the induced matrix relation

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

UOV trapdoor constraint. The secret system is required to satisfy the UOV condition, i.e., it contains no oil–oil monomials. 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 [24], one may assume (up to equivalence) 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.

Two interpretations of the embedded constraints. Equation (1) can be exploited both constructively (by the defender) and analytically (by the attacker).

Defender side. During key generation, one can sample L (equivalently R) uniformly at random, select a large fraction of the public coefficients α_{rkij}, and solve (1) for the remaining coefficients so that the oil–oil terms cancel by design. Since (1) is linear in the unknown coefficients α_{rkij}, this completion step reduces to solving linear systems over $\mathbb{F}_q$. Concretely, there are $o^2m\ell$ scalar constraints and $n^2m\ell$ public coefficients. These decompose into $m\ell$ independent linear systems (one for each pair (r,k)), each consisting of o^2 equations in n^2 unknowns. Thus $(n^2-o^2)m\ell$ coefficients may be chosen freely, and the remaining $o^2m\ell$ coefficients are derived.

This perspective also motivates a compression method: publish a seed that deterministically expands to the freely chosen coefficients, and publish explicitly only the derived part (the $o^2m\ell$ completed coefficients). The explicit component requires $o^2m\ell\log_2(q) = o^2m\ell\, e\log_2(p)$ bits, which incentivizes small o (and hence small m), potentially compensated by a larger extension degree e.

Attacker side. Conversely, since (1) holds for the true secret L and contains only public coefficients α_{rkij}, it provides an attacker with additional algebraic relations that may leak information about L. Using the block form (2), a direct approach yields a system of $o^2m\ell$ equations in the ov unknown entries of R. A basic balancing heuristic is $o^2m\ell \approx ov$, equivalently $om\ell \approx v$; with $m = o$ this becomes $o^2\ell \approx v$. In the classical UOV regime, taking $v = O(o^2)$ enables asymptotic attacks, so the designer objective of minimizing o conflicts with this heuristic. The parameter sets considered here target concrete (non-asymptotic) security; this trade-off is revisited in Sec. 4.

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

This subsection explains how to *flatten* F-UOV 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 $\bar{\Phi}(.A_{rk}.) = (A_{rk1}, \ldots, A_{rke})$ — 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 $\mathbb{F}_q$-F-UOV 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:

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

However, these equations already lie in the span of the original ones, since $\Phi(\xi.)\bar{\Phi}(X^{[a]}) = \bar{\Phi}(X^{[a]})\Phi_{\otimes n}(.\xi^{[-a]})$, and $\langle\Phi(.\xi A.)\rangle = \langle\Phi(.A.)\rangle$. Therefore, in order to convert and solve the key recovery of F-UOV over $\mathbb{F}_p$, it is sufficient to assume that a solution satisfies:

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

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

4.2 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 [6], 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

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

(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 [6], the ideal (over $\mathbb{F}_q[x_{11}, \ldots, x_{ne}]$) generated by the quadratic forms $P_{\otimes n}(\mathcal{F}^a_{\otimes n}\bar{\Phi}(.A.)\,{}^t\mathcal{F}^b_{\otimes n})\,{}^tP_{\otimes n}$ 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 [6].

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; because the Frobenius action preserves multi-degree, the expanded system can be viewed as a multi-homogeneous system with $e\ell$ equations. Applying the multi-homogeneous XL strategy of [6] 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 $\mathbf{d}_{sol} \geq [1, \ldots, 1]$. At that first nontrivial multi-degree, we already have at least $\overline{\mathcal{M}}(\mathbf{d}_{sol}) \geq v^e$ relevant monomials[3]; therefore 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 $\mathbf{Pub}$ is a sum of ℓ Frobenius forms with *different* exponent differences $b_k - a_k$, hence the full equation mixes monomials of different Frobenius (multi-)degrees. After expanding to ground-field coordinates (as above), the resulting system over $\mathbb{F}_p[x_{11}, \ldots]$ is therefore *not* multi-homogeneous (there is no consistent multi-degree grading preserved by all terms). Consequently, the multi-homogeneous XL speedups exploited in [6] do not apply to the forgery system in any direct way.

[3] Intuitively, there are v "vinegar-type" variables in each of the e Frobenius layers, hence v^e products of one variable per layer.

4.3 Attacks based on finding a common isotropic subspace for many public matrices

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).

Heuristic counting for a common isotropic subspace. Fix t public matrices (more generally, t 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

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

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), this matrix can be written as $X = (1_s \; R_{s\times(n-s)})$, so W is parametrized by the $s(n-s)$ entries of R. The isotropy constraints become

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

i.e., $s^2 t$ equations in $s(n-s)$ unknowns over $\mathbb{F}_q$. Under the usual random-equations heuristic, the expected number of solutions is about $q^{s(n-s)-s^2 t}$. In particular, when $s(n-s) - s^2 t < 0$, we expect a solution to exist 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 $t = 1$, we recover the standard rule of thumb that a random nondegenerate quadratic form has maximal totally isotropic subspaces of dimension at most $\lfloor n/2 \rfloor$. Indeed, 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 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 at most dimension $t_1 := \lfloor n/2 \rfloor$. At step i, assume we have collected t_i independent vectors that are isotropic for i forms. Pick a new $(i+1)$st form and restrict it to the span of these t_i vectors. Over a space of dimension t_i, we can extract a totally isotropic subspace of at most dimension $t_{i+1} := \lfloor t_i/2 \rfloor$. Iterating while $t_i \ge 2$ (equivalently $2^i \le n$) suggests that one may hope to obtain a common isotropic vector for at most up to $\lfloor \log_2(n) \rfloor$ 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

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

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

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

For a random choice of U, the coefficients $U_i^{[a_{rk}]} A_{rk}\, {}^t U_j^{[b_{rk}]}$ behave like random field elements, so the restricted instance 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.

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 $\mathbb{F}_q$; we model this by an MQ instance with about te equations in about ve variables.

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. If one solves (6) as a univariate polynomial, the best known algorithms run in time linear in the degree ($\approx p^{e/2}$) [17,20,8,28,30], i.e.,

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

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, ?: not known.

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	∞	∞	∞	∞	∞	∞

exponential in e. If instead one rewrites (6) as a multivariate quadratic system, our experiments indicate 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 $e/2$. 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 [12]) and $\binom{e+e/2}{e}^2$.

4.4 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. 4.3), a natural strategy is to first find a vector U that cancels two ($t = 2$) Frobenius forms (which can be done in polynomial time), 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 $X^{p^{a_k}+p^{b_k}}$. No algorithm is currently known that solves such instances below the targeted security level for our parameters.

Solving $\mathbf{Pub}(X) = H$ in $\mathbb{F}_p$. After base-field expansion, we model this as an MQ instance with approximately oe equations in approximately ne variables over $\mathbb{F}_p$. 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 [12].

4.5 Direct key recovery

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 $\mathcal{O}$:

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

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 model its cost by the generic complexity of solving a univariate equation over $\mathbb{F}_{q^\ell}$, which we estimate by its degree and which we lower bound by $q^{e/2}$.

Solving $\mathbf{Pub}_{rk}(X) = 0$ in $\mathbb{F}_p$. As in Sec. 4.1, 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, this estimate is 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 [15] 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 F-UOV 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 [15, Sec. C] confirm this expectation.

4.6 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.

4.7 Minrank attacks

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

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

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). Given the construction of Λ, such dependencies could occur only in the unlikely event that the corresponding expanded basis Λ^* contains one or more rows with all coordinates in $\mathbb{F}_p$. Since Λ is sampled at random, this event has probability at most $(p/q)^v = 1/p^{v(e-1)}$, and we therefore 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 [29]. 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.

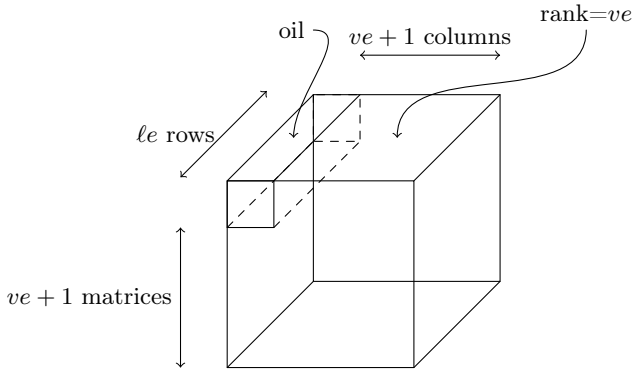

Fig. 1. Schematic 3D view of the secret matrix family. From above, one sees matrices with ℓe rows and ne columns; from the left, one sees the stack of ne matrices. The top-left corner corresponds to the secret oil space (all zeros). The schematic also indicates the number of matrices and their dimensions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

Complexity estimates are given in Tab. 2 (from CryptographicEstimators [12]).

4.8 Wedge attack

This attack was introduced in [25]. 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 expected to contain information about the orthogonal of the oil subspace. Here, $q_1, \ldots, q_m$ encode, in the exterior algebra, the alternating (antisymmetric) part of the public polynomials (equivalently, their differentials).

In our setting, the $\mathbb{F}_q$ viewpoint does not provide access to either the symmetric or the antisymmetric part of the underlying bilinear map, since $B(X, Y) = X^{[a]} A\, {}^tY^{[b]}$ is not $\mathbb{F}_q$-bilinear (only $\mathbb{F}_p$-bilinear), and thus neither are B^+ and B^-. Consequently, an adversary would have to 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 $\binom{ne}{ve}$ (Tab. 2).

4.9 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 CryptographicEstimators [12] 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

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 $X \in \mathbb{F}_q^n$. 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 F-UOV 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

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

The resulting linear system can be expressed via a matrix $M \in \mathbb{F}_p^{oe \times me}$ such that $\bar{\Phi}(\mathbf{Pub}(X\Lambda + V)) = \bar{\Phi}(X)M + \bar{\Phi}(\mathbf{Pub}(V))$. Let $\Gamma = H - \mathbf{Pub}(V)$. Since $o = m$, the matrix M is square, and whenever M is invertible a solution is $X = \bar{\Phi}^{-1}(\bar{\Phi}(\Gamma)M^{-1})$. Under the random choice of V, M is invertible with probability $1 - 1/p$.

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 $\Sigma \in \mathbb{F}_p^n$ 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 to 1 and, between Σ and $-\Sigma$, keep the lexicographically smaller representative.

5.2 Key generation

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

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 $\sigma \leftarrow (\Sigma, \tau)$.

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.

5.5 Proposed parameters

We propose the following parameter sets.

Throughout, we take $o = m = 1$, $\ell = 5$, $v = 4$, and hence $n = o + v = 5$. For NIST security levels I, III, and V we use $(p, e) = (31, 61)$, $(61, 89)$, and $(61, 127)$, respectively.

The pairs (a_k, b_k) are chosen so that the differences $b_k - a_k$ are pairwise distinct and as large as possible. This aims to mix the degrees and multi-degrees optimally (see Sec. 4.2). Concretely, $b_k - a_k = ([e/2] + 1 - k) \bmod (e)$ for $k = 1, \ldots, \ell$. Examples are given below. (In the table, a_k and b_k are ordered by decreasing $b_k - a_k$; in an implementation, any fixed order would do.)

Level	a_k	0	2	1	5	7
I	b_k	30	31	29	32	33
III	b_k	44	45	43	46	47
V	b_k	63	64	62	65	66

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 3. Proposed parameters. Sizes in bytes.

Level	o	v	ℓ	p	e	$\lvert\sigma\rvert$	$\lvert\mathbf{Pub}\rvert$	$\lvert\sigma\rvert + \lvert\mathbf{Pub}\rvert$	$\lvert\mathsf{Sec}\rvert$
I	1	4	5	31	61	223	207	430	32
III	1	4	5	61	89	366	350	716	32
V	1	4	5	61	127	509	493	1002	32

6 Implementation and Performances

We implemented a reference proof-of-concept implementation of the F-UOV scheme in the Magma language [5]. The source code is available at https://anonymous.4open.science/r/FUOV-BD39/ref%20F-uov.txt. The performance results reported in Tab. 4 were obtained using the online Magma Calculator [4]. 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.

Table 4. Timings using Online Magma Calculator [4].

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

7 Conclusions

We proposed Frobenius-UOV (F-UOV), 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 even diagonal,
- optimized, constant-time, leakage-resilient implementations,
- extensions such as threshold variants of F-UOV.

References

1. Arratia, R., Barbour, A.D., Tavaré, S.: Logarithmic Combinatorial Structures: A Probabilistic Approach. Cambridge University Press (2003)
2. 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)
3. 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)

4. Bosma, W., Cannon, J., Playoust, C.: Online magma calculator. https://magma.maths.usyd.edu.au/calc/
5. 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)
6. Cabarcas, D., Li, P., Verbel, J., Villanueva-Polanco, R.: Improved attacks for SNOVA by exploiting stability under a group action. In: Annual International Cryptology Conference. pp. 358–389. Springer (2025)
7. Cam, L.L.: An approximation theorem for the Poisson binomial distribution. Pacific Journal of Mathematics **10**(4), 1181–1197 (1960)
8. 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
9. Cox, D., Little, J., O'Shea, D.: Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra. Springer, 4 edn. (2015)
10. 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
11. 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)
12. Esser, A., Verbel, J.A., Zweydinger, F., Bellini, E.: SoK: CryptographicEstimators - a Software Library for Cryptographic Hardness Estimation. In: AsiaCCS. ACM (2024)
13. 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)
14. 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)
15. 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
16. 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)
17. von zur Gathen, J., Gerhard, J.: Modern Computer Algebra. Cambridge University Press, 3 edn. (2013)
18. 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/
19. Knopfmacher, J.: Abstract Analytic Number Theory. North-Holland (1975)
20. Lidl, R., Niederreiter, H.: Finite Fields, Encyclopedia of Mathematics and its Applications, vol. 20. Cambridge University Press, 2 edn. (1997)

21. 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
22. Nguyen, Q.D., Nguyen, M.H.: Towards compact UOV-based mq signatures: Rectangular and lifted whipping structures. Cryptology ePrint Archive (2026)
23. 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)
24. 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/
25. 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
26. 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
27. Rousseau, É.: Arithmétique efficace des extensions de corps finis. Theses, Institut Polytechnique de Paris (Jul 2021), https://theses.hal.science/tel-03299466
28. 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
29. 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
30. 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
31. 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
32. 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

A 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

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

A.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 F-UOV 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 [13,7,1,14,19]. This is also consistent with (12) being well-posed.

In Tab. 5, 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 $O(1/s!)$.

A.2 Some insights into the public equation

We also estimate how the complexity of solving a single equation of the form $X^{[a]} A \, {}^t X^{[b]} = h$ 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

[5] Of course, we are forced to choose small parameters to perform the experiment, but we expect the qualitative behavior to extend to other parameters.

[6] Computing these alternate keys becomes challenging for large values of p and e.

Table 5. 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

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 $U_1 = U^{[a_1]} A_1 {}^t U^{[b_1]}$ and $U_2 = U^{[a_2]} A_2 {}^t U^{[b_2]}$. 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 $1/p^{2e}$. Finding such a U algorithmically is expected to be costly, at least $O(p^{e/2})$.
- $U_1 = 0$, $U_2 \neq 0$ 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 $1/\delta$), 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 F-UOV 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.

[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 [10].

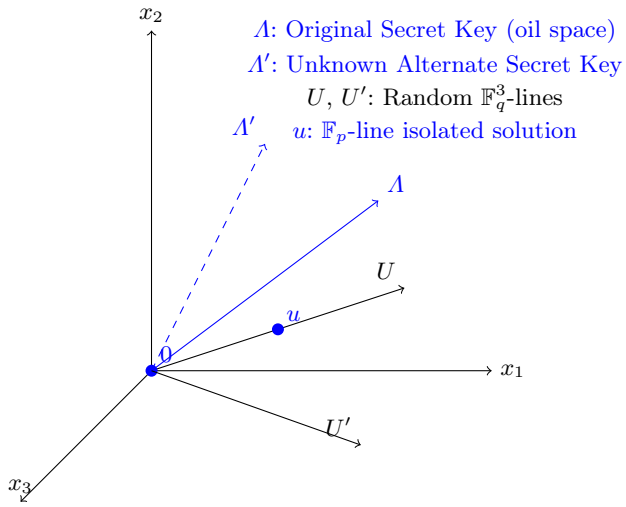

Fig. 2. Locus of $\mathbf{Pub}(x_1, x_2, x_3) = 0$ (in blue)