X24 Down: Cryptanalysis of Hankel-based Multivariate Signatures

Alexandre Camelin[1], Thai Hung Le[1,2], Brice Minaud[1], Phong Q. Nguyen[1], and Florian Tousnakhoff[1]

[1] École normale supérieure, PSL University, CNRS, Inria, France
{firstname.name}@ens.psl.eu
[2] LTCI, Telecom Paris, Institut Polytechnique de Paris, France
thai.le@telecom-paris.fr

Abstract. The X24 multivariate signature scheme was introduced by Di Muzio, Feussner, and Semaev at PQCrypto 2026. It offers remarkably short signatures, together with a new design approach for multivariate signatures that departs from the typical UOV and HFE frameworks.
In this work, we present an efficient cryptanalysis of X24. Our attack recovers the secret key from the public key in time $\mathcal{O}(q \cdot \mathsf{poly}(n))$, where n is the number of field elements in the signature, and q is the order of the finite field. An implementation of the attack recovers the secret key in a few minutes on the full X24 parameters.
The attack makes essential use of the exterior algebra, and shows a different way of using that algebra for multivariate cryptanalysis, compared to the wedge attack introduced by Ran at Eurocrypt 2026. Another notable feature of the attack is that it eventually reduces the cryptanalysis of X24 to the cryptanalysis of a McEliece variant using Generalized Reed-Solomon codes, drawing an unexpected connection between multivariate and code-based cryptanalysis.

1 Introduction

The transition towards post-quantum security is an important target of current cryptographic research. The intent is to prepare for the event where efficient quantum computers become available. While the probability and timeline of that event remain a matter of debate, the potential impact of adversaries equipped with powerful enough quantum machines is well-established [19,33]. Such adversaries would be able to break most public-key cryptography deployed today, potentially crippling communication networks, banking systems, and indeed much of the digital infrastructure modern life relies on. Post-quantum cryptography aims to lay the groundwork for the post-quantum transition, by building primitives and protocols that remain secure in the face of such adversaries.

This effort is pursued in conjunction with standardization bodies including the National Institute of Standards and Technology (NIST) as well as industrial and governmental stakeholders, with the aim to bring secure post-quantum solutions to market within a relatively short time frame (in regard to the considerable amount of time-consuming work required to achieve that goal, from

construction to cryptanalysis, optimization, vetting, standardization, implementation, industrial deployment, etc). The National Security Agency (NSA) has set a timeline for the transition to post-quantum cryptography by 2035 [23]. Recently, Google and Cloudflare have both accelerated their own internal timeline for post-quantum security to 2029, reflecting growing industrial concern [17,39]. This decision was due in part to new results achieving factorization and discrete logarithm computation over elliptic curves with quantum circuits much smaller than was previously known to be possible [1].

Multivariate cryptography. The scheme targeted in this article is X24 [14]. It is a post-quantum digital signature algorithm, based on multivariate cryptography. In the general landscape of post-quantum cryptography, digital signatures are a central tool, and multivariate signatures in particular are attractive in view of their short signature sizes. Among candidates selected by NIST at the outcome of their first call for post-quantum primitives, the lattice-based scheme FALCON achieves the shortest signatures, at around 650 bytes [15]. For comparison, multivariate signatures such as UOV, currently a second-round candidate in the ongoing second NIST call for additional post-quantum signatures, achieve signatures in the order of 100 bytes [6].

The main drawback of multivariate signatures is large public keys. Currently, all four multivariate signatures that are part of the second round of the NIST call for additional digital signatures are variants of UOV (UOV, QR-UOV, MAYO and SNOVA [5,6,16,38]). The main goal of these constructions is to decrease the size of the UOV public key by introducing additional structure. To our knowledge, few modern designs for multivariate signatures depart from the UOV framework, with exceptions including the HFE-based scheme from [10].

The X24 scheme. The X24 multivariate signature scheme was presented by Di Muzio, Feussner, and Semaev at PQCrypto 2026 [14]. The scheme is notable in at least two regards: first, it achieves remarkably short signatures. Second, it departs from the current main paradigms to build multivariate signatures: it is not an UOV variant, nor even an HFE variant [10,29]. Instead, X24 introduces an elegant new construction based around the convolution product, where the signature algorithm amounts to polynomial factorization. Because of its new structure, as well as its promising signature and key sizes, the question of analyzing the security of X24 is natural.

1.1 Contributions

Our main contribution is a theoretically and practically efficient cryptanalysis of X24. Our attack recovers an equivalent secret key using only the public key, in time $\mathcal{O}(q \cdot \mathsf{poly}(n))$, where n denotes the number of field elements in the signature, and $q > n$ denotes the finite field order. (Here, we talk about recovering *an equivalent secret key*, rather than *the secret key*, because the public key of X24 does not uniquely define the secret key; but the effect is the same: recovering

an equivalent secret key enables arbitrary forgeries, using the original signature algorithm.) Let us highlight a few features of the attack.

- The starting point is to compute the bi-orthogonal of the public key matrices with respect to a certain bilinear map. The bilinear map arises from the exterior algebra (Section 3.2), which plays an essential role. The exterior algebra has seen a recent surge of interest for multivariate cryptanalysis, mainly following the pioneering wedge attack introduced by Ran at the rump session of Eurocrypt 2025 [30], and subsequently published at Eurocrypt 2026 [31]. Our usage is quite different. Among other properties, our use of the exterior algebra works over fields of any characteristic, whereas the wedge attack currently requires even characteristic[3] [31,8,37]. This is interesting insofar as our result provides evidence that the exterior algebra is a powerful tool for the cryptanalysis of multivariate schemes beyond the wedge attack and its variants.
- The expanded public key obtained via the bi-orthogonal can be described exactly as the vector space of matrices $E_S = \{S^\mathsf{T} H S : H \text{ is Hankel}\}$, where S is the secret key of X24, and Hankel matrices are matrices with constant anti-diagonals (cf. Section 3.4). We characterize matrices of rank one in E_S, and prove that E_S must contain at least $n+2$ (pairwise linearly independent) such matrices whenever $q > n$. We use the kernel algorithm for MinRank computation [18] in order to extract those rank-one matrices from E_S. Proving unconditional bounds on the performance of a MinRank algorithm in a situation where there are many rank-one matrices would usually be tricky, especially considering known analyses of the kernel algorithm are heuristic. Thanks to specific properties of Hankel matrices, in our setting it is possible to formally prove that the MinRank step terminates in expected time $\mathcal{O}(q^2 n^\omega \ln(\frac{q}{q-n})) = \mathcal{O}(q \cdot \mathsf{poly}(n))$.[4]
- Finally, it turns out that the set of rank-one matrices produced by the MinRank step can be directly transformed into a public key for a different signature scheme: a McEliece variant based on Generalized Reed-Solomon codes [34]. The secret key of the resulting McEliece variant instance is essentially the same as the secret key of the target X24 instance, so that breaking the McEliece variant implies breaking X24. The McEliece variant in question is known to be insecure: its secret key can be recovered in polynomial time using an algorithm by Sidelnikov and Shestakov [34]. This yields a complete break of the original X24 instance. To our knowledge, this is the first time that the public key of a given scheme (X24, which is a multivariate scheme) essentially reduces to the public key of a scheme belonging to a different branch of cryptography (a McEliece variant, which is a code-based scheme). This is certainly surprising, although some similarities between multivariate and code-based cryptography are discussed in the related work section.

[3] Variants of the wedge attack over fields of odd characteristic do exist, but they do not use the exterior algebra, and are less efficient [21,31].

[4] It may appear that the first expression is only $\tilde{\mathcal{O}}(q^2 \cdot \mathsf{poly}(n))$, but it is in fact $\mathcal{O}(q \cdot \mathsf{poly}(n))$ (cf. Proposition 3).

While the attack combines several different techniques, they are all quite efficient: the total complexity is bounded by $\mathcal{O}(qn^{\omega+1} + n^{3\omega})$ field operations, where q is the field order, n is the number of field elements in the signature, and $2 \leq \omega \leq 3$ is the matrix multiplication exponent, assuming $q = (1 + \Omega(1))n$.

In order to evaluate practical performance, we have implemented the full attack. The authors of X24 provide example parameters aiming for 128 bits of post-quantum security[5]. Our implementation is able to recover an equivalent secret key of the X24 scheme with those parameters, using only the public key, within a few minutes on a laptop computer. (In fact, the attack does not need the full public key: any 4 matrices among the 40 matrices defining the public key suffice.) The implementation is available here:

`https://gitlab.di.ens.fr/ftousnakhoff/x24-cryptanalysis`

Remark. The complexity of our attack has a linear dependency on q due to the kernel attack in the MinRank step. In theory, increasing the parameters of X24 to have an exponential q would thwart the attack as it stands. However, we do not expect that this countermeasure would be viable in terms of performance or security (see Appendix B).

Disclaimer. This is a preliminary version of our work: the description of the attack is complete, but the proofs of some propositions are not yet included; neither is a deeper analysis of the case $q \leq n$. They will be added in an upcoming full version.

1.2 Related work

Exterior algebra. The exterior algebra has made an appearance in cryptography constructively [36], and for the cryptanalysis of multivariate signatures following the UOV framework [21,8,37]. The latter series of works are extensions of Ran's wedge attack [31]. As noted earlier, our usage is different: in the wedge attack, the exterior algebra appears because in characteristic 2, the polar forms of quadratic forms (which make up the public key of UOV and its variants) are alternating. In our case, the exterior algebra arises mainly from Lemma 1, and from the bilinear map ϕ in Section 4.1. This suggests that the exterior algebra has a role to play in the cryptanalysis of multivariate schemes beyond the wedge attack.

Bi-orthogonal. Beside the exterior algebra, the other key technique in the initial step of our attack is to consider the bi-orthogonal of the public key matrices with respect to the bilinear map ϕ. The public key of X24 can be regarded as truncated, in the same vein as the truncation techniques that have been used notably in the context of multivariate HFE-based schemes [28,10]. Taking the bi-orthogonal reveals a "completion" of the original public key, in the sense that

[5] The target security level seems to be somewhat overstated due to generic attacks, see Appendix A.

it removes this truncation. We are not aware of a technique similar to the bi-orthogonal in multivariate cryptanalysis before. However, to some extent, this is reminiscent of lattice-based cryptanalysis using orthogonal lattices introduced in [24] (see the survey [25]), where the completion of an integer lattice $\bar{L}$ is defined, and is exactly the bi-orthogonal. There, the bilinear map is the Euclidean inner product: a typical attack builds a lattice L, computes a reduced basis of its orthogonal lattice, removes a few elements, and computes a reduced basis of the new orthogonal lattice. When the map ϕ is a non-degenerate symmetric bilinear functional over a vector space V, we also remark that the notion of bi-orthogonal with respect to ϕ matches the notion of bidual with respect to ϕ. In all of these settings, the bi-orthogonal or bidual of a set x can be interpreted as a completion or closure of x (for a notion of completion or closure that depends on each setting).

MinRank problem. The MinRank problem plays an important role in both multivariate and code-based cryptography [3,2]. For algorithms based on the minors or support-minors modeling, the complexity analysis typically relies on a regularity or genericity assumption [4]. Here, we use the kernel attack by Goubin and Courtois [18], because specific properties of Hankel matrices in our setting allow for a provable bound on the running time without requiring additional assumptions. We note that other MinRank algorithms would likely also perform well, but leave their analysis in our setting as an open problem.

Relationship with code-based cryptography. A striking aspect of our attack is that it eventually reduces the cryptanalysis of a multivariate signature scheme (X24) to the cryptanalysis of a code-based signature scheme (a McEliece variant using Generalized Reed-Solomon codes [34]). In general, multivariate and code-based schemes are quite different branches of cryptography, and we do not want to make too much of a connection between the two areas. Nevertheless, the fact that such a reduction could occur may be explained to some degree by limited similarities between multivariate and code-based signatures.

In both cases, the traditional paradigm to build hash-and-sign signatures is similar: build a trapdoor function by obfuscating a trapdoored quadratic system (resp. a linear code with a known efficient decoding algorithm). In both cases, in addition to relying on a general assumption such as the hardness of average-case MQ (resp. average-case hardness of decoding a random code), such signature schemes rely on a more ad-hoc structured assumption that the trapdoor function is hard to distinguish from a random quadratic system (resp. a random code). Furthermore, in the case of algebraic codes such as Reed-Solomon, Gabidulin or Goppa codes, polynomial systems over a finite field play a central role, as they do in multivariate cryptosystems.

Bridging the two areas was explored in an earlier work by Smith-Tone and Tone [35], on the constructive side (rather than the cryptanalytic side in this work). The recent syzygy-based McEliece distinguisher by Randriambololona [32] is another example of a work that targets code-based cryptography, using tools

that could also apply to multivariate cryptography (for example, one could try to use Betti numbers to distinguish polynomial systems arising from the public keys of multivariate cryptosystems from random polynomial systems, although we are not yet aware of a successful attempt to do so).

Overall, the use of McEliece-style structured assumptions outlined above results in schemes that bring code-based cryptography a little closer with multivariate cryptography, even if the two areas are quite different in general. This contrasts with other approaches that try to bring code-based cryptography closer to lattice-based cryptography, and try to avoid relying on McEliece-style indistinguishability assumptions, such as [13,12].

2 Technical overview

A Hankel matrix is a matrix whose entries along anti-diagonals are constant:

$$\begin{bmatrix} h_1 & h_2 & h_3 & \dots & h_{n-1} & h_n \\ h_2 & h_3 & h_4 & \dots & h_n & h_{n+1} \\ h_3 & h_4 & h_5 & \dots & h_{n+1} & h_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ h_{n-1} & h_n & h_{n+1} & \dots & h_{2n-3} & h_{2n-2} \\ h_n & h_{n+1} & h_{n+2} & \dots & h_{2n-2} & h_{2n-1} \end{bmatrix}$$

The vector space of n-dimensional Hankel matrices is denoted by $\mathcal{H}_n$. Observe that $\mathcal{H}_n$ has dimension $2n-1$.

An instance of the X24 signature scheme is defined by the choice of a finite field order q, and two integers m and n such that $n \leq m \leq 2n-1$. The recommended parameters are $q = 256$, $n = 24$, $m = 40$.

- The secret key of X24 consists of an invertible matrix $S \in \mathrm{GL}_n(\mathbb{K})$, and m Hankel matrices $H_1, \dots, H_m$, sampled uniformly at random (conditioned on being linearly independent).
- The public key of X24 is the m-tuple of symmetric matrices:

$$A_1 = S^\mathsf{T} H_1 S, \dots, A_m = S^\mathsf{T} H_m S.$$

The main attack in this work recovers S (up to equivalence), on input the public key $A_1, \dots, A_m$. An important quantity in the attack is the space $E_S = S^\mathsf{T} \mathcal{H}_n S$, of dimension $2n-1$. The linear span of public key matrices $\mathrm{span}(A_1, \dots, A_m)$ is distributed as a uniformly random subspace of E_S of dimension m. The attack proceeds in three steps.

1. **Step 1** (Section 4) recovers E_S from $A_1, \dots, A_m$. This step runs in time $\mathcal{O}(n^{3\omega})$.
2. **Step 2** (Section 5) extracts $n+2$ independent rank-one matrices from E_S, assuming $q > n$. This step runs in time $\mathcal{O}(qn^{\omega+1})$, provided $q = (1+\Omega(1))n$.

3. **Step 3** (Section 6) computes S given the $n+2$ rank-one matrices produced by Step 2. This step runs in time $\mathcal{O}(n^\omega)$.

Let us consider each step in turn.

2.1 Step 1: expanding the public key into E_S

Exterior algebra. For the purpose of this overview, we will be content with an informal outline of the exterior algebra. A formal treatment can be found in Section 3.

In the polynomial algebra $\mathbb{K}[X_1, \ldots, X_n]$, unknowns are assumed to commute, in the sense that for any $1 \leq i, j \leq n$, we have $X_i X_j = X_j X_i$. If we discard the assumption that unknowns commute, and suppose that there is no special relationship between $X_i X_j$ and $X_j X_i$, then we obtain the tensor algebra over $\mathbb{K}^n$. If we assume that unknowns anticommute, that is $X_i X_j = -X_j X_i$, and also that $X_i^2 = 0$ for all i, then we obtain the exterior algebra. In each case, multiplication is written differently: in the tensor algebra, the product is denoted $\otimes$; in the exterior algebra, it is denoted $\wedge$. In all cases, we write X_i^2 for the product of X_i with itself.

Example 1. We can use distributivity to compute $(X_1 + X_2)^2$ in each algebra. In the polynomial algebra, this yields the well-known identity:

$$(X_1 + X_2) \cdot (X_1 + X_2) = X_1^2 + 2X_1X_2 + X_2^2.$$

In the tensor algebra, the identity becomes:

$$(X_1 + X_2) \otimes (X_1 + X_2) = X_1^2 + X_1 \otimes X_2 + X_2 \otimes X_1 + X_2^2.$$

In the exterior algebra, something special happens:

$$(X_1 + X_2) \wedge (X_1 + X_2) = X_1^2 + X_1 \wedge X_2 + X_2 \wedge X_1 + X_2^2 = 0.$$

From now on, we will think of $X_1, \ldots, X_n$ not as unknowns, but as the canonical basis of $\mathbb{K}^n$. The behavior of computations we have just seen is unaffected, this is just a conceptual change. Observe that the space $\{\sum_{i=1}^n \lambda_i X_i : (\lambda_1, \ldots, \lambda_n) \in \mathbb{K}^n\}$ of linear combinations of $X_1, \ldots, X_n$ is now equal to $\mathbb{K}^n$. Thus $\mathbb{K}^n$ appears as a subspace of the exterior algebra, corresponding to homogeneous polynomials of degree 1. In general, we write $\bigwedge^k(\mathbb{K}^n)$ for the homogeneous polynomials of degre k within the exterior algebra, with $\bigwedge^1(\mathbb{K}^n) = \mathbb{K}^n$ (as just discussed), and $\bigwedge^0(\mathbb{K}^n) = \mathbb{K}$. The exterior algebra has a whole is denoted by $\bigwedge(\mathbb{K}^n)$.

The phenomenon $(X_1 + X_2)^2 = 0$ we just saw at the end of Example 1 is not an accident. The exterior algebra obeys two defining rules:

1. $\forall \boldsymbol{v} \in \mathbb{K}^n$, $\quad \boldsymbol{v} \wedge \boldsymbol{v} = 0$,
2. $\forall \boldsymbol{v}, \boldsymbol{w} \in \mathbb{K}^n$, $\quad \boldsymbol{v} \wedge \boldsymbol{w} = -\boldsymbol{w} \wedge \boldsymbol{v}$.

(Technically, the second rule is implied by the first one; we write both rules explicitly because of their usefulness in computations.)

Bracket notation. Now let us define a bracket notation on vectors in $\mathbb{K}^n$:

$$\forall \boldsymbol{v}, \boldsymbol{w} \in \mathbb{K}^n, \quad [\boldsymbol{v}, \boldsymbol{w}] = \boldsymbol{v}\boldsymbol{w}^\mathsf{T} - \boldsymbol{w}\boldsymbol{v}^\mathsf{T} \in \mathcal{M}_n(\mathbb{K}).$$

We claim that $[\boldsymbol{v}, \boldsymbol{w}]$ can be identified with $\boldsymbol{v} \wedge \boldsymbol{w} \in \bigwedge^2(\mathbb{K}^n)$. A proper treatment of this point will be given in Section 3. For now, as a sanity check, we point out that $[\boldsymbol{v}, \boldsymbol{w}]$ abides by the defining rules of the exterior algebra mentioned earlier:

1. $\forall \boldsymbol{v} \in \mathbb{K}^n, \qquad [\boldsymbol{v}, \boldsymbol{v}] = \boldsymbol{v}\boldsymbol{v}^\mathsf{T} - \boldsymbol{v}\boldsymbol{v}^\mathsf{T} = 0,$
2. $\forall \boldsymbol{v}, \boldsymbol{w} \in \mathbb{K}^n, \qquad [\boldsymbol{v}, \boldsymbol{w}] = \boldsymbol{v}\boldsymbol{w}^\mathsf{T} - \boldsymbol{w}\boldsymbol{v}^\mathsf{T} = -(\boldsymbol{w}\boldsymbol{v}^\mathsf{T} - \boldsymbol{v}\boldsymbol{w}^\mathsf{T}) = -[\boldsymbol{w}, \boldsymbol{v}].$

Due to the identification, we write $[\boldsymbol{v}, \boldsymbol{w}] = \boldsymbol{v} \wedge \boldsymbol{w}$.

Expanding the bracket notation to matrices, define for any two matrices $X, Y \in \mathbb{K}^{n \times k}$ (with n rows and k columns):

$$[X, Y] = XY^\mathsf{T} - YX^\mathsf{T}.$$

The bracket over matrices has the following useful property. If $X, Y \in \mathcal{M}_n(\mathbb{K})$ and $S \in \mathrm{GL}_n(\mathbb{K})$, then:

$$[S^\mathsf{T} X S, S^\mathsf{T} Y S^{-\mathsf{T}}] = S^\mathsf{T}[X, Y]S. \tag{1}$$

In addition, letting $\boldsymbol{x_1}, \ldots, \boldsymbol{x_k}$ (resp. $\boldsymbol{y_1}, \ldots, \boldsymbol{y_k}$) denote the columns of X (resp. Y):

$$[X, Y] = \sum_{i=1}^{k} [\boldsymbol{x}_i, \boldsymbol{y}_i] = \sum_{i=1}^{k} \boldsymbol{x}_i \wedge \boldsymbol{y}_i. \tag{2}$$

A proof of this equality is given in Lemma 1. We do not dwell on it here, in order to keep the focus on the overall attack outline. In the remainder, we will work with $\mathbb{K} = \mathbb{F}_q$.

Expanding the public key into E_S. All the pieces are now in place to explain the first step of the attack. Define:

$$\begin{aligned} \phi : \left(\bigwedge^2 \mathbb{K}^n\right)^n \times \left(\bigwedge^1 \mathbb{K}^n\right)^n &\to \bigwedge^3 \mathbb{K}^n \\ \left((x_i)_{i=1}^n, (y_i)_{i=1}^n\right) &\mapsto \sum_{i=1}^{n} x_i \wedge y_i. \end{aligned}$$

We start by explaining how ϕ is used to compute E_S. We will then explain where ϕ is coming from. (We proceed in that order because knowing what we intend to do with ϕ helps understand where we are going in the second part.)

The right input of ϕ lies in $(\bigwedge^1 \mathbb{K}^n)^n = (\mathbb{K}^n)^n$, which we can identify with $\mathcal{M}_n(\mathbb{K})$ by mapping $(\boldsymbol{c_1}, \ldots, \boldsymbol{c_n}) \in (\mathbb{K}^n)^n$ to the matrix in $\mathcal{M}_n(\mathbb{K})$ with columns $\boldsymbol{c_1}, \ldots, \boldsymbol{c_n}$.

Given a subset $E \subseteq \mathcal{M}_n(\mathbb{K})$, define the *left orthogonal* of E (with respect to ϕ) as:

$$E^{\perp} = \left\{ x \in \left(\bigwedge^2 (\mathbb{K}^n) \right)^n : \forall M \in E, \phi(x, M) = 0 \right\}.$$

Similarly for $F \subseteq (\bigwedge^2(\mathbb{K}^n))^n$, define the *right orthogonal* of F (with respect to ϕ) as:

$$F^{\perp} = \left\{ M \in \mathcal{M}_n(\mathbb{K}) : \forall x \in F, \phi(x, M) = 0 \right\}.$$

(We use the same notation for the left and right orthogonal; since the left and right inputs of ϕ lie in different spaces there is no ambiguity.) Finally, define the *bi-orthogonal* $E^{\perp\perp}$ of $E \subseteq \mathcal{M}_n(\mathbb{K})$ as the right orthogonal of the left orthogonal of E.

Cutting to the chase, the way we reconstruct the secret space $E_S = S^{\mathsf{T}} \mathcal{H}_n S$ from the public key matrices $A_1, \ldots, A_m$ is:

$$E_S = \{A_1, \ldots, A_m\}^{\perp\perp}. \tag{3}$$

First, let us point out that this equation allows us to compute a basis of E_S in time $\mathcal{O}(n^{3\omega})$. Indeed, because ϕ is bilinear, computing $\{A_1, \ldots, A_m\}^{\perp}$ amounts to computing the kernel of the linear map:

$$x \mapsto (\phi(x, A_i))_{i=1}^m.$$

Because the input x in that linear system lies in $(\bigwedge^2(\mathbb{K}^n))^n$, which is a space of dimension $n\binom{n}{2} = \mathcal{O}(n^3)$, the complexity of solving the system is $\mathcal{O}(n^{3\omega})$. Computing the bi-orthogonal amounts to computing the right orthogonal of the previous kernel. This is cheaper than the previous computation, because the input now lies in $\mathcal{M}_n(\mathbb{K})$, which is of dimension $\mathcal{O}(n^2)$ instead of $\mathcal{O}(n^3)$.

Let us now sketch why Equation (3) holds. What really happens is:

$$\{A_1, \ldots, A_m\}^{\perp} = E_S^{\perp}. \tag{4}$$

The reasoning behind that statement will be explained shortly, but we note that it is a conjecture: we do not formally prove it. Instead, we provide several supporting arguments. (Since our full attack implementation recovers the secret key of X24 with its real parameters, there is no question that the attack works; the point of the analysis is to provide insight.)

1. First, $\{A_1, \ldots, A_m\} \subseteq E_S$, which implies $E_S^{\perp} \subseteq \{A_1, \ldots, A_m\}^{\perp}$. In Section 4, we prove that the dimension of $E_S^{\perp}$ is $2n-2$, with an exact description of the elements it contains (Theorem 3). Hence an equivalent reformulation of the conjecture from Equation (4) is: we know that $\{A_1, \ldots, A_m\}^{\perp}$ *must* contain those elements; we conjecture that it does not contain any other.
2. A heuristic argument why $\{A_1, \ldots, A_m\}^{\perp}$ is not expected to contain any other element "by accident" is that the linear system that arises from computing the left orthogonal of $A_1, \ldots, A_m$ is overdetermined. Indeed, the kernel lies within $(\bigwedge^2 \mathbb{K})^n$, which has dimension $n\binom{n}{2}$ as a $\mathbb{K}$-vector space (this

quantity can be thought of as the number of scalar unknowns in the linear system); whereas the image of ϕ lies in $\bigwedge^3 \mathbb{K}$, which has dimension $\binom{n}{3}$, and the left orthogonal of $\{A_1, \ldots, A_m\}$ must annihilate all A_i's, which amounts to $m\binom{n}{3}$ equations (discounting possible linear dependencies among equations). As soon as $m \geq 4$, we have $m\binom{n}{3} \geq n\binom{n}{2}$ (except for very low values of n, which are not relevant for security reasons), and we have more linear equations than unknowns. For reference, the real parameters of X24 have $m = 40$, so the real system is vastly overdetermined.

3. We have performed extensive experiments (Section 7). The conjecture has held in every experiment, as long as $m \geq 4$ (as predicted), and for every choice of q we have tried, including both even and odd characteristic.

Additional supporting analysis and experiments are provided in Sections 4 and 7. From the conjecture, Equation (3) (which was the equation allowing us to compute E_S) follows by taking the right orthogonal on both sides, and using the fact that $E_S^{\perp\perp} = E_S$ (proved in Corollary 1).

Before moving on to the next step of the attack, let us provide some insight regarding where ϕ and the related conjecture are coming from. Let Z be the shift matrix:

$$Z = \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 \\ 1 & 0 & 0 & \ldots & 0 \\ 0 & 1 & 0 & \ldots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \ldots & 0 & 1 & 0 \end{bmatrix}.$$

Let $H \in \mathcal{H}_n$ be an arbitrary Hankel matrix. Let $\boldsymbol{h}$ denote the first column vector of $-ZH$, and let $\boldsymbol{e_1}$ denote the first canonical basis vector of $\mathbb{K}^n$. Compute:

$$\begin{aligned} & [H, Z] \\ &= HZ^\mathsf{T} - ZH^\mathsf{T} \\ &= \begin{bmatrix} h_1 & h_2 & h_3 & \ldots & h_{n-1} & h_n \\ h_2 & h_3 & h_4 & \ldots & h_n & h_{n+1} \\ h_3 & h_4 & h_5 & \ldots & h_{n+1} & h_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ h_{n-1} & h_n & h_{n+1} & \ldots & h_{2n-3} & h_{2n-2} \\ h_n & h_{n+1} & h_{n+2} & \ldots & h_{2n-2} & h_{2n-1} \end{bmatrix} Z^\mathsf{T} - Z \begin{bmatrix} h_1 & h_2 & h_3 & \ldots & h_{n-1} & h_n \\ h_2 & h_3 & h_4 & \ldots & h_n & h_{n+1} \\ h_3 & h_4 & h_5 & \ldots & h_{n+1} & h_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ h_{n-1} & h_n & h_{n+1} & \ldots & h_{2n-3} & h_{2n-2} \\ h_n & h_{n+1} & h_{n+2} & \ldots & h_{2n-2} & h_{2n-1} \end{bmatrix} \\ &= \begin{bmatrix} 0 & h_1 & h_2 & h_3 & \ldots & h_{n-1} \\ 0 & h_2 & h_3 & h_4 & \ldots & h_n \\ 0 & h_3 & h_4 & h_5 & \ldots & h_{n+1} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & h_{n-1} & h_n & h_{n+1} & \ldots & h_{2n-3} \\ 0 & h_n & h_{n+1} & h_{n+2} & \ldots & h_{2n-2} \end{bmatrix} - \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 & 0 \\ h_1 & h_2 & h_3 & \ldots & h_{n-1} & h_n \\ h_2 & h_3 & h_4 & \ldots & h_n & h_{n+1} \\ h_3 & h_4 & h_5 & \ldots & h_{n+1} & h_{n+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ h_{n-1} & h_n & h_{n+1} & \ldots & h_{2n-3} & h_{2n-2} \end{bmatrix} \end{aligned}$$

$$= \begin{bmatrix} 0 & h_1 & h_2 & \dots & h_{n-2} & h_{n-1} \\ -h_1 & 0 & 0 & \dots & 0 & 0 \\ -h_2 & 0 & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -h_{n-2} & 0 & 0 & \dots & 0 & 0 \\ -h_{n-1} & 0 & 0 & \dots & 0 & 0 \end{bmatrix}$$
$$= \boldsymbol{h}\boldsymbol{e_1}^\mathsf{T} - \boldsymbol{e_1}\boldsymbol{h}^\mathsf{T}$$
$$= [\boldsymbol{h}, \boldsymbol{e_1}].$$

Recall that $A_1, \dots, A_m$ are the public key matrices of X24, with $A_i = S^\mathsf{T} H_i S$. Applying the above result to H_i, we have:

$$[H_i, Z] = [\boldsymbol{h_i}, \boldsymbol{e_1}],$$

where $\boldsymbol{h_i}$ is the first column vector of $-ZH_i$. Applying Equation (1), we deduce:

$$\begin{aligned}[S^\mathsf{T} H_i S, S^\mathsf{T} Z S^{-\mathsf{T}}] &= S^\mathsf{T}[\boldsymbol{h_i}, \boldsymbol{e_1}]S \\ &= [S^\mathsf{T}\boldsymbol{h_i}, S^\mathsf{T}\boldsymbol{e_1}],\end{aligned}$$

whence, letting $M = S^\mathsf{T} Z S^{-\mathsf{T}}$, $\boldsymbol{h'_i} = S^\mathsf{T}\boldsymbol{h_i}$, $\boldsymbol{s} = S^\mathsf{T}\boldsymbol{e_1}$:

$$[A_i, M] = [\boldsymbol{h'_i}, \boldsymbol{s}] = \boldsymbol{h'_i} \wedge \boldsymbol{s}.$$

Let $\boldsymbol{a_{i,1}}, \dots, \boldsymbol{a_{i,n}}$ (resp. $\boldsymbol{m_1}, \dots, \boldsymbol{m_n}$) denote the column vectors of A_i (resp. M). We can apply Equation (2) to the previous equality and obtain:

$$\sum_{k=1}^{n} \boldsymbol{a_{i,k}} \wedge \boldsymbol{m_k} = \boldsymbol{h'_i} \wedge \boldsymbol{s}.$$

Now that everything is neatly expressed within the exterior algebra, we can start to take advantage of its properties. Applying the map $x \mapsto x \wedge \boldsymbol{s}$ on both sides of the equation yields:

$$\sum_{k=1}^{n} \boldsymbol{a_{i,k}} \wedge \boldsymbol{m_k} \wedge \boldsymbol{s} = \boldsymbol{h'} \wedge \boldsymbol{s} \wedge \boldsymbol{s} = 0.$$

That equation translates exactly to:

$$\phi\big((\boldsymbol{m_k} \wedge \boldsymbol{s})_{k=1}^{n}, A_i\big) = 0.$$

The above reasoning is what originally led us to consider ϕ. An important point here is that the left input $(\boldsymbol{m_k} \wedge \boldsymbol{s})_{k=1}^{n}$ depends *only* on S. Indeed $\boldsymbol{s} = S^\mathsf{T}\boldsymbol{e_1}$, and the $\boldsymbol{m_k}$'s are the columns of $M = S^\mathsf{T} Z S^{-\mathsf{T}}$. In particular, $(\boldsymbol{m_k} \wedge \boldsymbol{s})_{k=1}^{n}$ must be in the left orthogonal of any subset of E_S.

We have thus found one element, $(\boldsymbol{m_k} \wedge \boldsymbol{s})_{k=1}^{n}$, in the left orthogonal of $\{A_1, \ldots, A_m\}$. Initially, that element arose from computing $[H, Z]$, then doing a change of basis by S (i.e. applying Equation (1)). In Section 4, we show that additional elements arise from considering $[H, Z^2], \ldots, [H, Z^{n-1}]$ (stopping at $n-1$ because $Z^n = 0$), although the computation is more involved. All those elements have the same key property that they depend only on S. This yields $n-1$ elements in $\{A_1, \ldots, A_m\}^\perp$.

Finally, $\{A_1, \ldots, A_m\}^\perp$ contains an additional n "trivial" elements that lie in the left orthogonal of any set of symmetric matrices (recall that the A_i's are symmetric); see Section 4 for more information. That brings us to $n-1+n-1 = 2n-2$ (linearly independent) elements in the left orthogonal of $\{A_1, \ldots, A_m\}$, because the element arising from considering $[H, Z^{n-1}]$ is also one of the "trivial" elements (see Section 4.3). This is indeed the dimension of $E_S^\perp$: we prove that $E_S^\perp$ is equal to the linear span of those elements (Theorem 3). The earlier conjecture amounts to stating: the same is true for any subset of E_S spanning a subspace of high enough dimension. (Specifically, we conjecture that dimension 4 suffices, for analytical and experimental reasons discussed in Section 7.2.)

Remark. Computing $\dim(\{A_1, \ldots, A_m\}^\perp)$ constitutes a polynomial-time distinguisher between the public key of X24 and a random m-tuple of symmetric matrices: in the random case, the dimension is n, while in the real case, the dimension is $2n-2$.

2.2 Step 2: extracting rank-one elements of E_S

At the outcome of the previous step, we have computed (a basis of) $E_S = S^\mathsf{T} \mathcal{H}_n S$. Define a *Vandermonde* vector to be a vector of the form

$$(1, \alpha^1, \ldots, \alpha^{n-1}) \in \mathbb{K}^n, \quad \text{for } \alpha \in \mathbb{K}.$$

As a special case, it will be convenient to also call $\boldsymbol{e_n} = (0, \ldots, 0, 1)$ a Vandermonde vector. In total there are $q+1$ Vandermonde vectors in $\mathbb{K} = \mathbb{F}_q$: one for each choice of α, plus the special case $\boldsymbol{e_n}$.

It can be observed that for any Vandermonde vector $\boldsymbol{v} = (1, \alpha^1, \ldots, \alpha^{n-1})$, the entry at row i and column j of $\boldsymbol{v}\boldsymbol{v}^\mathsf{T}$ is α^{i+j}, in particular $\boldsymbol{v}\boldsymbol{v}^\mathsf{T}$ is a Hankel matrix of rank one. The same is true for $\boldsymbol{e_n}\boldsymbol{e_n}^\mathsf{T}$. It follows that $S^\mathsf{T}\boldsymbol{v}\boldsymbol{v}^\mathsf{T}S = (S^\mathsf{T}\boldsymbol{v})(S^\mathsf{T}\boldsymbol{v})^\mathsf{T}$ is a rank-one matrix in E_S, for every Vandermonde vector $\boldsymbol{v}$. In Section 5, we prove that *all* rank-one matrices in E_S are of this form, up to a scalar multiple. Since there exist $q+1$ Vandermonde vectors, the previous point implies that there exist exactly $q+1$ non-equivalent rank-one matrices in E_S (for $n > 1$), where we treat matrices as equivalent whenever they are scalar multiples of each other.

The second step of the attack is to compute $n+2$ such matrices, which assumes $q+1 \geq n+2$, i.e. $q > n$. (The real parameters of X24 are $q = 256$, $n = 24$; the case $q \leq n$ is briefly discussed in Appendix B.) Finding matrices of a given low rank in a given subspace of $\mathcal{M}_n(\mathbb{K})$ is called the MinRank problem (cf. Section 1.2). Here, we want to find rank-one matrices in E_S. We use a specific

MinRank algorithm due to Goubin and Courtois [18], because it lends itself well to formal analysis in the Hankel setting (and also because that algorithm is very fast for the real X24 parameters).

In short, the idea of this algorithm is to pick two random vectors $\boldsymbol{v}$ and $\boldsymbol{w}$, and solve for $M \in E_S$ in the linear system:

$$M\boldsymbol{v} = 0, \quad M\boldsymbol{w} = 0.$$

That system has $2n$ linear equations (when viewed over the base field), and the dimension of E_S is $2n-1$, so naively we would expect one non-trivial solution on average. Crucially, when a solution does exist, it is very likely to be a matrix of rank one. This is because a given rank-one matrix $M = \boldsymbol{a}\boldsymbol{b}^\mathsf{T}$ satisfies $M\boldsymbol{v} = 0$ and $M\boldsymbol{w} = 0$ as soon as $\boldsymbol{b}^\mathsf{T}\boldsymbol{v} = 0$ and $\boldsymbol{b}^\mathsf{T}\boldsymbol{w} = 0$, which holds with probability roughly q^{-2}. That probability decreases exponentially with the rank of M. It follows that the sampling algorithm over E_S induced by picking $\boldsymbol{v}$, $\boldsymbol{w}$ at random and solving the previous system is heavily concentrated in rank-one matrices.

We repeat the process until we have sampled $n+2$ non-equivalent rank-one matrices in E_S. The formal analysis of the running time is quite long (unlike the running time itself), so we omit it here. In the end, we are able to prove that the expected time to obtain the desired $n+2$ matrices is bounded by $\mathcal{O}(qn^{\omega+1})$ for large enough q (the precise bound is given in Proposition 3).

2.3 Step 3: recovering the secret key from rank-one matrices

At the outcome of the previous step, we know $n+2$ rank-one matrices in E_S, and we know that those matrices are of the form $\boldsymbol{w}\boldsymbol{w}^\mathsf{T}$ (up to a scalar multiple), with $\boldsymbol{w} = S^\mathsf{T}\boldsymbol{v}$ for some Vandermonde vector $\boldsymbol{v}$. Given such a matrix, we obtain a scalar multiple of $S^\mathsf{T}\boldsymbol{v}$ simply by choosing any arbitrary non-zero row (or column) of the matrix.

For simplicity, let us set aside the case $\boldsymbol{v} = \boldsymbol{e_n}$. From the $n+2$ matrices, we get $n+2$ vectors of the form $\boldsymbol{w_i} = \beta_i S^\mathsf{T}(1, \alpha_1^i, \ldots, \alpha_i^{n-1})$, for $1 \leq i \leq n+2$, where the α_i's and β_i's are scalars. Let $G \in \mathbb{K}^{n\times(n+2)}$ denote the matrix whose columns are $\boldsymbol{w_1}, \ldots, \boldsymbol{w_{n+2}}$. The punchline is that G is the generating matrix of a Generalized Reed-Solomon code.

A Generalized Reed-Solomon code is a linear subspace of $\mathbb{K}^\ell$ of the form:

$$\Big\{\big(\gamma_1 f(\xi_1), \ldots, \gamma_\ell f(\xi_\ell)\big) : f \in \mathbb{K}[X] \text{ of degree at most } d\Big\},$$

where the γ_i's and ξ_i's are scalar constants that define the code (together with its parameters $\mathbb{K}$, ℓ and d). In our setting, the code generated by G is for the parameters $\ell = n+2$, $d = n-1$, $\xi_i = \alpha_i$, and $\gamma_i = \beta_i$.

Indeed, if we view the k-th row of S^T as the vector of coefficients of a polynomial in $\mathbb{K}[X]$ of degree at most $n-1$, then the k-th entry of $S^\mathsf{T}(1, \alpha_1^i, \ldots, \alpha_i^{n-1})$ is the evaluation of that polynomial at α_i. This implies that each row of G is the evaluation of a fixed polynomial (whose coefficients are given by the corresponding row of S^T) at the points $\alpha_1, \ldots, \alpha_{n+2}$, with columns subsequently scaled by

β_i. Because the rows of S are linearly independent (S is invertible), so are the rows of G, which means that the rows of G are not only included in the the code defined in the previous paragraph, but form a basis of it.

The matrix G is exactly the public key of the natural McEliece variant that uses Generalized Reed-Solomon codes in place of binary Goppa codes, already considered in [34]. (In the McEliece cryptosystem, the public key is obtained by "obfuscating" the generating matrix of a linear code with a known efficient decoding algorithm, such as a Generalized Reed-Solomon code; this obfuscation is performed by left-multiplying by an invertible matrix, which is exactly the role played by S^T in our setting, and permuting the columns, which in the Reed-Solomon case is not useful because it just amounts to permuting the α_i's).

The attack by Sidelnikov and Shestakov from [34] recovers S (up to equivalence) from G, in time $\mathcal{O}(n^\omega)$. We give a description of the attack in Section 6 for completeness. Once S is known, the rest of the X24 secret key follows by computing $H_i = S^{-\mathsf{T}} A_i S^{-1}$. This concludes the cryptanalysis of X24.

2.4 Roadmap

Section 3 contains necessary preliminaries, including an introduction to the exterior algebra, the bracket notation $[X, Y]$, and a definition of Hankel matrices and X24. Section 4 presents the first stage of the attack, which uses a bilinear map ϕ arising from the exterior algebra to expand the X24 public key, revealing the vector space $E_S = S^\mathsf{T} \mathcal{H}_n S$, where S is the secret key of X24, and $\mathcal{H}_n$ is the vector space of Hankel matrices. Section 5 studies E_S, showing that rank-one matrices in E_S arise from Vandermonde vectors, and analyzes the performance of the kernel MinRank algorithm on E_S. Section 6 shows that the rank-one matrices extracted from E_S form the public key of a McEliece variant using Generalized Reed-Solomon codes, and presents the attack by Sidelnikov and Shestakov to recover the secret key [34], completing the attack on X24. Finally, Section 7 discusses experimental results.

3 Preliminaries

3.1 Notation and terminology

Let $\mathbb{K} = \mathbb{F}_q$ denote the field of cardinality q. The set of matrices with a rows and b columns is written $\mathbb{K}^{a \times b}$. We may also write $\mathcal{M}_n(\mathbb{K}) = \mathbb{K}^{n \times n}$. The subset of invertible matrices of $\mathcal{M}_n(\mathbb{K})$ is denoted by $\mathrm{GL}_n(\mathbb{K})$. We write M^T for the transpose of a matrix M, and $M^{-\mathsf{T}} = (M^{-1})^\mathsf{T}$ for the transpose of its inverse. A matrix $M \in \mathcal{M}_n(\mathbb{K})$ is *symmetric* if $M = M^\mathsf{T}$. It is *antisymmetric* if $M = -M^\mathsf{T}$. It is *alternating* if it is antisymmetric and its diagonal is zero[6].

[6] In characteristic other than two, a matrix is antisymmetric if and only if it is alternating, but we will sometimes work in characteristic two, which makes the distinction necessary.

Throughout this work, vectors are written in boldface. If $\boldsymbol{v} \in \mathbb{K}^n$ and $1 \leq i \leq n$, we write $\boldsymbol{v}[i]$ for the i-th entry of $\boldsymbol{v}$. If $M \in \mathbb{K}^{a \times b}$ and $1 \leq i \leq a$, $1 \leq j \leq b$, we write $M[i,j]$ for the entry of M at the intersection of row i and column j.

Vectors are assumed to be column vectors by default. In other words, we identify $\mathbb{K}^n$ with $\mathbb{K}^{n \times 1}$. In particular, if $\boldsymbol{x} \in \mathbb{K}^a$, $\boldsymbol{y} \in \mathbb{K}^b$, then $\boldsymbol{x}\boldsymbol{y}^\mathsf{T}$ is the matrix of rank at most one in $\mathbb{K}^{a \times b}$ whose entry at the intersection of the i-th row and j-th column is $\boldsymbol{x}\boldsymbol{y}^\mathsf{T}[i,j] = \boldsymbol{x}[i]\boldsymbol{y}[j]$.

The linear algebra exponent is denoted by $2 \leq \omega \leq 3$.

3.2 Exterior algebra

We already gave an intuitive presentation of the exterior algebra in the technical overview (Section 2). That presentation gives a correct description of the exterior algebra (up to isomorphism), but it is not a standard definition. In that description, the unknowns $X_1, \ldots, X_n$ (assimilated with the canonical basis of $\mathbb{K}^n$) may appear to play a special role. That is not the case. Any automorphism of $\mathbb{K}^n$, sending $X_1, \ldots, X_n$ to any other basis, extends to an algebra automorphism. This makes it apparent that the unknowns $X_1, \ldots, X_n$ play no special role, beside forming a basis of the underlying vector space.

Overall, the simple description of the exterior algebra from Section 2 amounts to choosing a basis $X_1, \ldots, X_n$ of $\mathbb{K}^n$, then describing the behavior of the algebra over that basis (which also generates the algebra). In contrast, standard definitions of the exterior algebra are independent of a choice of basis. We provide such a definition next. This will also serve to introduce important properties that will be needed later.

The following presentation of the exterior algebra is valid over any field. Because the X24 scheme is over $\mathbb{F}_{2^8}$, we will be especially careful to write statements that are also true in characteristic two. We will sometimes treat two algebraic structures as being equal when they are equal up to isomorphism.

Remark. Compared to presentations of the exterior algebra in some prior works [31,8], we are more explicit regarding its relationship with alternating tensors and matrices, as we make frequent use of that relationship.

Formal definition. Let V be a vector space over $\mathbb{K}$. Let $\mathsf{T}_k(V) = V \otimes \cdots \otimes V$ denote the k-fold tensor product of V with itself, with $\mathsf{T}_0(V) = \mathbb{K}$. An element of $\mathsf{T}_k(V)$ is called a *k-tensor*. The *tensor algebra* over V is the graded algebra $\mathsf{T}(V) = \bigoplus_{k \in \mathbb{N}} \mathsf{T}_k(V)$.

Definition 1 (Exterior algebra). *The* exterior algebra *over V is the quotient algebra:*

$$\bigwedge(V) = \frac{\mathsf{T}(V)}{\langle \boldsymbol{v} \otimes \boldsymbol{v} : \boldsymbol{v} \in V \rangle},$$

where $\langle \boldsymbol{v} \otimes \boldsymbol{v} : \boldsymbol{v} \in V \rangle$ denotes the ideal of $\mathsf{T}(V)$ generated by elements $\boldsymbol{v} \otimes \boldsymbol{v} \in \mathsf{T}_2(V)$. The product over $\bigwedge(V)$ inherited from $\otimes$ via the quotient structure is called the wedge product*, written $\wedge$.*

Intuitively, quotienting by the ideal generated by elements $\boldsymbol{v} \otimes \boldsymbol{v}$ amounts to forcing $\boldsymbol{v} \wedge \boldsymbol{v} = 0$ for all $\boldsymbol{v} \in V$ (since $\boldsymbol{v} \otimes \boldsymbol{v}$ lies in the ideal, hence it is in the equivalence class of zero). It should be immediately noted that this also forces anti-commutativity, in the sense that $\boldsymbol{v} \wedge \boldsymbol{w} = -\boldsymbol{w} \wedge \boldsymbol{v}$ for all $\boldsymbol{v}, \boldsymbol{w} \in V$. Indeed:

$$0 = (\boldsymbol{v} + \boldsymbol{w}) \wedge (\boldsymbol{v} + \boldsymbol{w}) = \boldsymbol{v} \wedge \boldsymbol{v} + \boldsymbol{v} \wedge \boldsymbol{w} + \boldsymbol{w} \wedge \boldsymbol{v} + \boldsymbol{w} \wedge \boldsymbol{w} = \boldsymbol{v} \wedge \boldsymbol{w} + \boldsymbol{w} \wedge \boldsymbol{v}.$$

If V has dimension n and we let $X_1, \ldots, X_n$ be a basis of V, we recover the simple description of the exterior algebra from Section 2.

Remark. In characteristic other than two, the ideals $\langle \boldsymbol{v} \otimes \boldsymbol{w} + \boldsymbol{w} \otimes \boldsymbol{v} : \boldsymbol{v}, \boldsymbol{w} \in V \rangle$ and $\langle \boldsymbol{v} \otimes \boldsymbol{v} : \boldsymbol{v} \in V \rangle$ are equal, so that the exterior algebra can be thought of as the "anti-commutative" free algebra over V, but in characteristic two there is a strict inclusion $\subset$ (for sufficiently large n).

Useful properties. The exterior algebra inherits the graded structure from $\mathsf{T}(V)$. Define

$$\bigwedge^k(V) = \frac{\mathsf{T}_k(V)}{\langle \boldsymbol{v} \otimes \boldsymbol{v} : \boldsymbol{v} \in V \rangle \cap \mathsf{T}_k(V)},$$

then:

$$\bigwedge(V) = \bigoplus_{k \in \mathbb{N}} \bigwedge^k(V).$$

Definition 2 (k-form, k-blade, decomposability). *An element of $\bigwedge^k(V)$ is called a* k-form *(or alternatively, a* homogeneous element of degree k*). A k-form is said to be* decomposable *if and only if it can be written in the form $\boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k$ for some $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k \in V$. A decomposable k-form is called a* k-blade*.*

Note that most k-forms are not decomposable (in the same way that most multivariate polynomials of degree k are not products of linear polynomials, or most tensors are not pure).

An important property of k-blades is that $\boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k = 0$ whenever there exists $i \neq j$ such that $\boldsymbol{v}_i = \boldsymbol{v}_j$. More generally, $\boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k = 0$ if and only if $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k$ are linearly dependent. Another notable property is that two non-zero k-blades $v = \boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k$ and $w = \boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k$ are colinear (in the sense that one is a scalar multiple of the other) if and only if:

$$\operatorname{span}(\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}) = \operatorname{span}(\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\}),$$

where $\operatorname{span}(x)$ denotes the linear span of x. Intuitively, a k-blade $\boldsymbol{v}_1 \wedge \cdots \wedge \boldsymbol{v}_k$ only carries the information of $\operatorname{span}(\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_k\})$, plus a scalar. Conversely, any k-dimensional subspace of V uniquely identifies a k-blade, up to a scalar.

Assume V has dimension n, and let $\boldsymbol{e_1}, \ldots, \boldsymbol{e_n}$ be a basis of V. A basis of $\bigwedge^k(V)$ as a $\mathbb{K}$-vector space is given by the k-blades:

$$\boldsymbol{e}_{i_1} \wedge \cdots \wedge \boldsymbol{e}_{i_k} \quad \text{for } 1 \leq i_1 < \cdots < i_k \leq n.$$

In particular, $\dim(\bigwedge^k(V)) = \binom{n}{k}$.

Identification with alternating tensors. Recall that a map f defined over V^k is *k-linear* if it is linear in each coordinate, whenever the other coordinates are fixed. It is a k-linear functional if, in addition, it maps into the base field $\mathbb{K}$. A tensor of the form $\boldsymbol{v}_1 \otimes \cdots \otimes \boldsymbol{v}_k$ can be naturally mapped to the k-linear functional:

$$\begin{aligned}\rho(\boldsymbol{v}_1 \otimes \cdots \otimes \boldsymbol{v}_k) : V^k &\to \mathbb{K} \\ (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) &\mapsto \prod_{i=1}^{k} (\boldsymbol{x}_i \mid \boldsymbol{v}_i),\end{aligned}$$

where $(\boldsymbol{x} \mid \boldsymbol{v})$ denotes the dot product between $\boldsymbol{x}$ and $\boldsymbol{v}$. Extending ρ by linearity to all of $\mathsf{T}_k(V)$ yields an isomorphism between $\mathsf{T}_k(V)$ and the vector space of k-linear functionals.

Definition 3 (Alternating tensor). *A k-linear functional $f : V^k \to \mathbb{K}$ is said to be* alternating *if for all $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k \in V$, and for all $1 \leq i < j \leq n$,*

$$\boldsymbol{x}_i = \boldsymbol{x}_j \implies f(\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) = 0.$$

A tensor t in $\mathsf{T}_k(V)$ is said to be alternating *if $\rho(t)$ is alternating.*

If we write the tensor $t \in \mathsf{T}_k(V)$ along a basis of $\mathsf{T}_k(V)$ (as a $\mathbb{K}$-vector space):

$$t = \sum_{(i_1,\ldots,i_k) \in \{1,\ldots,n\}^k} t_{i_1,\ldots,i_k} \cdot \boldsymbol{e}_{i_1} \otimes \cdots \otimes \boldsymbol{e}_{i_k},$$

then t is alternating if and only if for all permutations σ of $\{1, \ldots, k\}$,

$$t_{i_{\sigma(1)},\ldots,i_{\sigma(k)}} = \epsilon(\sigma) \cdot t_{i_1,\ldots,i_k},$$

where $\epsilon(\sigma)$ denotes the sign of σ (in other words, the k-dimensional array defined by the $t_{i_1,\ldots,i_k}$'s is antisymmetric), and, in addition, $t_{i_1,\ldots,i_k} = 0$ whenever there exist $j \neq j'$ such that $i_j = i_{j'}$. (In characteristic other than two, the second condition is implied by the first, so that alternating tensors coincide with antisymmetric tensors.)

Let $\mathsf{A}_k(V)$ denote the subspace of alternating tensors within $\mathsf{T}_k(V)$. The point of introducing alternating tensors is that there is a natural mapping from $\bigwedge^k(V)$ to $\mathsf{A}_k(V)$, given by:

$$\tau(\boldsymbol{e}_{i_1} \wedge \cdots \wedge \boldsymbol{e}_{i_k}) = \sum_{\sigma \in S_k} \epsilon(\sigma) \cdot \boldsymbol{e}_{i_{\sigma(1)}} \otimes \cdots \otimes \boldsymbol{e}_{i_{\sigma(k)}},$$

where S_k denotes the set of permutations over $\{1, \ldots, k\}$. The map τ extends to all of $\bigwedge^k(V)$ by linearity, and yields a vector space isomorphism between $\bigwedge^k(V)$ and $\mathsf{A}_k(V)$. In the remainder of this article, we will often identify elements of $\mathsf{A}_k(V)$ with $\bigwedge^k(V)$, as is common in the literature [40]. For instance, we may write $\boldsymbol{x} \otimes \boldsymbol{y} - \boldsymbol{y} \otimes \boldsymbol{x} = \boldsymbol{x} \wedge \boldsymbol{y}$.

Remark. The map τ is more than a vector space isomorphism: it can be extended to an algebra isomorphism in a natural way, by equipping $\mathsf{A}_k(V)$ with the appropriate product. We will not directly use this product in the present paper.

Identification between $\bigwedge^2(V)$ and alternating matrices. We have seen in the previous section that k-tensors (resp. alternating k-tensors) can be naturally identified with k-linear functionals (resp. alternating k-linear functionals), which are themselves in bijection with k-dimensional arrays $\{1,\dots,n\}^k \to \mathbb{K}$ (resp. alternating arrays).

The public key of X24 is a tuple of bilinear forms. For that reason, we are especially interested in the case $k = 2$. In that case, the previous identifications can be summarized as follows (writing "$\leftrightarrow$" for each isomorphism):

$$\text{2-tensor (over } \mathbb{K}^n) \leftrightarrow \text{bilinear functional (over } \mathbb{K}^n) \leftrightarrow \text{matrix (in } \mathbb{K}^{n\times n}),$$

and similarly:

$$\text{2-form} \leftrightarrow \text{alternating bilinear functional} \leftrightarrow \text{alternating matrix}.$$

Looking more closely:

- The second arrow $\leftrightarrow$ corresponds to the familiar description of a bilinear functional as a matrix. Namely, given a bilinear functional $f : \mathbb{K}^n \times \mathbb{K}^n \to \mathbb{K}$, there exists a unique matrix $M \in \mathbb{K}^{n\times n}$ such that:
$$\forall \boldsymbol{x}, \boldsymbol{y} \in \mathbb{K}^n, \quad f(\boldsymbol{x}, \boldsymbol{y}) = \boldsymbol{x}^\mathsf{T} M \boldsymbol{y}.$$
- For $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{K}^n$, the composition of the two arrows $\leftrightarrow$ above sends a tensor $\boldsymbol{x} \otimes \boldsymbol{y}$ (resp. a form $\boldsymbol{x} \wedge \boldsymbol{y}$) to the matrix $\boldsymbol{x}\boldsymbol{y}^\mathsf{T}$ (resp. $\boldsymbol{x}\boldsymbol{y}^\mathsf{T} - \boldsymbol{y}\boldsymbol{x}^\mathsf{T}$). For that reason, we may also write such a matrix $\boldsymbol{x} \otimes \boldsymbol{y}$ (resp. $\boldsymbol{x} \wedge \boldsymbol{y}$).

Remark. There are multiple choices for the second isomorphism $\leftrightarrow$, depending on the choice of a basis for V. Whenever we work over $V = \mathbb{K}^n$, we implicitly use the canonical basis.

3.3 Bracket notation

Definition 4 (bracket notation). *Given two arbitrary matrices $X, Y \in \mathbb{K}^{n\times k}$, define the bracket $[X, Y]$ to be the alternating matrix:*

$$[X, Y] = XY^\mathsf{T} - YX^\mathsf{T}.$$

Remark. When X and Y are both symmetric matrices, $[X, Y]$ matches the commutator (and also the Lie bracket) of X and Y, which inspires the notation.

Lemma 1 (bracket lemma). *Let $X, Y \in \mathbb{K}^{n\times k}$. Denote by $\boldsymbol{x}_1, \dots, \boldsymbol{x}_k$ (resp. $\boldsymbol{y}_1, \dots, \boldsymbol{y}_k$) the columns of X (resp. Y). The following holds:*

$$[X, Y] = \sum_{i=1}^{k} \boldsymbol{x}_i \wedge \boldsymbol{y}_i.$$

Proof. The lemma follows directly from the observation:

$$XY^{\mathsf{T}} = \sum_{i=1}^{k} \boldsymbol{x}_i \boldsymbol{y}_i^{\mathsf{T}} = \sum_{i=1}^{k} \boldsymbol{x}_i \otimes \boldsymbol{y}_i.$$

The right equality holds by definition, via the earlier identification between $\mathsf{T}_2(\mathbb{K}^n)$ and $\mathcal{M}_n(\mathbb{K})$. The left equality holds because for all $1 \leq a, b \leq n$:

$$\Big(\sum_{i=1}^{k} \boldsymbol{x}_i \boldsymbol{y}_i^{\mathsf{T}}\Big)[a,b] = \sum_{i=1}^{k} (\boldsymbol{x}_i \boldsymbol{y}_i^{\mathsf{T}})[a,b] = \sum_{i=1}^{k} \boldsymbol{x}_i[a] \boldsymbol{y}_i[b] = \sum_{i=1}^{k} X[a,i] Y^{\mathsf{T}}[i,b],$$

which matches the standard definition of the matrix product XY^{T}. □

3.4 Hankel matrices

Hankel matrices were introduced in the technical overview (Section 2). Formally, they may be defined as follows.

Definition 5 (Hankel matrix). *A matrix $M \in \mathcal{M}_n(\mathbb{K})$ is a* Hankel matrix *if for all $1 \leq i, j, k, \ell \leq n$:*

$$i + j = k + \ell \implies M[i,j] = M[k,\ell].$$

The vector space of n-dimensional Hankel matrices is denoted by $\mathcal{H}_n(\mathbb{K})$, or just $\mathcal{H}_n$ when the ambient field is clear from the context.

3.5 Description of X24

An instance of the X24 signature scheme is defined by the choice of a finite field order q, and two integers m and n such that $n \leq m \leq 2n - 1$.

Secret key. The secret key is a uniformly random invertible matrix $S \in \mathrm{GL}_n(\mathbb{F}_q)$, and an m-tuple of Hankel matrices $(H_1, \ldots, H_m)$, sampled uniformly at random among m-tuples of linearly independent Hankel matrices.

Public key. The public key is the m-tuple of symmetric matrices:

$$S^{\mathsf{T}} H_1 S, \ldots, S^{\mathsf{T}} H_m S.$$

Our attack recovers the secret key (up to equivalence) from the public key, without interacting with the signature or verification algorithms. As a consequence, for our purpose, the structure of a signature is irrelevant. For the sake of completeness, we briefly outline it, and refer the reader to [14] for more details. The message is hashed into a vector $\boldsymbol{h} \in \mathbb{F}_q^m$. A valid signature is a vector $\boldsymbol{u} \in \mathbb{F}_q^n$ whose first non-zero coordinate is 1 and such that the overdetermined $m \times n$ linear system $\boldsymbol{u}^t A_i \boldsymbol{v} = \boldsymbol{h}[i]$, for $1 \leq i \leq m$ has a unique solution $\boldsymbol{v} \in \mathbb{F}_q^n$ which is larger than $\boldsymbol{u}$ in lexicographic order (for a given ordering over $\mathbb{F}_q$).

The goal of our attack is to recover the secret matrix S. Once S is known, the H_i's can be immediately computed from the public key. For that reason, throughout this article, we (slightly abusively) call S the secret key of X24.

The recommended parameters for X24 are $q = 256$, $n = 24$, $m = 40$. For those parameters, the signature size is 24B, with a 12KB public key and a 48KB private key. The targeted security level corresponds to 128 bits of post-quantum security (more precisely, NIST level-I security [26]).

Remark. Compared to the original article, we have renamed S_1 to S for notational convenience. We have also used the equivalent description of the scheme in terms of Hankel matrices, rather than its description in terms of convolution product. This is because our attack uses the Hankel viewpoint.

4 Subspace recovery using the exterior algebra

In this section, we show how to recover in polynomial time the secret subspace $E_S = S^\mathsf{T}\mathcal{H}_n S$ defined by the secret matrix S. In the process, we also obtain an efficient distinguisher between the X24 public key and a tuple of random symmetric matrices. Surprisingly, this stage of the attack does not require the full public key: for the recommended parameters of X24, 4 matrices out of m are sufficient in practice. All the results of this section are centered around a bilinear map ϕ (which can be viewed as a contraction).

4.1 A bilinear map distinguishing Hankel matrices

Define:

$$\begin{aligned}\phi : \left(\bigwedge^2 \mathbb{K}^n\right)^n \times \left(\bigwedge^1 \mathbb{K}^n\right)^n &\to \bigwedge^3 \mathbb{K}^n \\ \left((x_i)_{i=1}^n, (y_i)_{i=1}^n\right) &\mapsto \sum_{i=1}^n x_i \wedge y_i.\end{aligned}$$

Via the identification between $\bigwedge^1 \mathbb{K}^n$ and $\mathbb{K}^n$, and the identification between an $n \times n$ matrix and the n-tuple of its columns, the second input of ϕ can be viewed as a matrix. Thus, we view ϕ as a bilinear map:

$$\phi : E \times F \to \bigwedge^3 \mathbb{K}^n, \quad \text{with } E = \left(\bigwedge^2 \mathbb{K}^n\right)^n \text{ and } F = \mathcal{M}_n(\mathbb{K}),$$

where we recall that $F = \mathcal{M}_n(\mathbb{K})$ denotes the space of $n \times n$ matrices over $\mathbb{K}$.

For any $\boldsymbol{y} \in F$, define the subspace of E that is *left orthogonal* to $\boldsymbol{y}$ as:

$$\boldsymbol{y}^\perp = \{\boldsymbol{x} \in E \text{ s.t. } \phi(\boldsymbol{x}, \boldsymbol{y}) = 0\}.$$

By extension, for any subset $Y \subseteq F$, such as a subspace, define the left orthogonal of Y as: $Y^\perp = \bigcap_{\boldsymbol{y} \in Y} \boldsymbol{y}^\perp$.

Similarly, for any $\boldsymbol{x} \in E$ define the subspace of F that is *right orthogonal* to $\boldsymbol{x}$ as:

$$\boldsymbol{x}^\perp = \{\boldsymbol{y} \in F \text{ s.t. } \phi(\boldsymbol{x}, \boldsymbol{y}) = 0\}.$$

By extension, for any subset $X \subseteq E$, define the right orthogonal of X as: $X^\perp = \bigcap_{\boldsymbol{x} \in X} \boldsymbol{x}^\perp$.

Note that for any subspace $Y \subseteq F$, we have $Y \subseteq Y^{\perp\perp}$, but we do not necessarily have equality, unlike the special case where ϕ is the dot product over a finite-dimensional vector space and Y is a subspace: our attack is based on this gap, when Y is the subspace spanned by the public key.

4.2 Left orthogonal of symmetric matrices

We observe that all left orthogonals $V^\perp$ share a non-trivial subspace, when V only contains symmetric matrices. Indeed, if $M \in \mathcal{M}_n(\mathbb{K})$ is symmetric and I is the identity matrix, then $[I, M] = IM^\mathsf{T} - MI^\mathsf{T} = 0$ which can be rewritten by Lemma 1 as:

$$\sum_{i=1}^{n} \boldsymbol{e}_i \wedge \boldsymbol{m}_i = 0$$

where the $\boldsymbol{m}_i$'s are the columns of M. Thus for any $\boldsymbol{x} \in \mathbb{K}^n$,

$$0 = \boldsymbol{x} \wedge \sum_{i=1}^{n} \boldsymbol{e}_i \wedge \boldsymbol{m}_i = \sum_{i=1}^{n} (\boldsymbol{x} \wedge \boldsymbol{e}_i) \wedge \boldsymbol{m}_i$$

which is exactly of the form $\phi(f(\boldsymbol{x}), M)$ for some map f, namely:

Lemma 2. *Let $f : \mathbb{K}^n \to E$ be the linear map defined by*

$$f(\boldsymbol{x}) = (\boldsymbol{x} \wedge \boldsymbol{e_1}, \ldots, \boldsymbol{x} \wedge \boldsymbol{e_n}).$$

For any symmetric matrix $M \in \mathcal{M}_n(\mathbb{K})$ and any $\boldsymbol{x} \in \mathbb{K}^n$:

$$\phi(f(\boldsymbol{x}), M) = 0.$$

Furthermore, $\{f(\boldsymbol{e_1}), \ldots, f(\boldsymbol{e_n})\}$ is a basis of Im f.

Proof. We already proved the first statement. And observe that f is injective. Since $(\boldsymbol{e_1}, \ldots, \boldsymbol{e_n})$ is a basis of $\mathbb{K}^n$, it follows that $f(\boldsymbol{e_1}), \ldots, f(\boldsymbol{e_n})$ are linearly independent. □

Notice that $M^\perp$ is much bigger than its n-dimensional subspace Im f. Indeed, $M^\perp$ is the kernel of the linear map $\boldsymbol{x} \mapsto \phi(\boldsymbol{x}, M)$ whose input space has dimension $n^2(n-1)/2$ and whose output space has dimension $n(n-1)(n-2)/6$. Therefore $M^\perp$ has dimension at least $n(n-1)(n-(n-2)/3)/2 = n(n-1)(n+1)/3$.

However, we conjecture that for at least 4 random matrices, its left orthogonal is Im f, as 4 is the smallest integer such that the number $4n(n-1)(n-2)/6$ of equations defining the left orthogonal becomes larger than the dimension $n^2(n-1)/2$ of the input space

Conjecture 1. Let $M_1, \ldots, M_4 \in \mathcal{M}_n(\mathbb{K})$ be symmetric matrices chosen independently and uniformly at random. Then with overwhelming probability:

$$\text{span}(M_1, ..., M_4)^\perp = \text{Im } f,$$

where f is the rank-n injective linear map of Lemma 2.

Conjecture 1 is not strictly necessary for our overall attack to function: we state it to provide additional insight into the behavior of ϕ and left orthogonals. A weaker form of this conjecture, in which we replace 4 by the number m of public-key matrices and ask that the dimension of $\text{span}(M_1, ..., M_m)^\perp$ is only $< 2n-2$ (rather than being equal to n), is sufficient to warrant the correctness of a polynomial-time distinguisher. We have checked this conjecture experimentally and provided the arguments behind this conjecture in Section 7.2. As supporting evidence for the conjecture, we prove the following.

Theorem 1. *If $\mathcal{S}_n$ denotes the subspace of all $n \times n$ symmetric matrices, then:*

$$\text{Im } f = \mathcal{S}_n^\perp \text{ and } (\text{Im } f)^\perp = \mathcal{S}_n.$$

Proof. For the first statement, it suffices to show that $\mathcal{S}_n^\perp \subseteq \text{Im } f$. Let $(v_1, \ldots, v_n) \in \mathcal{S}_n^\perp$. By considering the diagonal matrix with only one non-zero coefficient at the i-th position, we get that $v_i \wedge \boldsymbol{e}_i = 0$, which implies that $v_i = \boldsymbol{e}_i \wedge \boldsymbol{u}_i$ for some vector $\boldsymbol{u}_i$. By considering the symmetric matrix with only two non-zero coefficients at the (i,j)-th and (j,i)-th positions, we obtain for $i \neq j$:

$$(\boldsymbol{e}_i \wedge \boldsymbol{u}_i) \wedge \boldsymbol{e}_j + (\boldsymbol{e}_j \wedge \boldsymbol{u}_j) \wedge \boldsymbol{e}_i = 0,$$

thus $\boldsymbol{u}_i - \boldsymbol{u}_j \in \text{span}(\boldsymbol{e}_i, \boldsymbol{e}_j)$. If we write $\boldsymbol{u}_i = \sum_{j=1}^n u_{i,j} \boldsymbol{e}_j$, then for all $k \notin \{i, j\}$, $u_{i,k} = u_{j,k}$. This means that for a fixed k, all the $u_{i,k}$'s for $i \neq k$ are equal, say to some α_k. Letting $\boldsymbol{u} = (\alpha_1, \ldots, \alpha_n)^t$, we obtain that $\boldsymbol{u}_i - \boldsymbol{u} \in \text{span } \boldsymbol{e}_i$ for all i. Then $v_i = \boldsymbol{e}_i \wedge \boldsymbol{u}$, so $(v_1, \ldots, v_n) \in \text{Im } f$.

For the second statement, assume that $M \in (\text{Im } f)^\perp$, and denote its columns by $\boldsymbol{m}_1, \ldots, \boldsymbol{m}_n$. For any $\boldsymbol{x} \in \mathbb{K}^n$, we have:

$$0 = \phi(f(\boldsymbol{x}), M) = \boldsymbol{x} \wedge \sum_{i=1}^n \boldsymbol{e}_i \wedge \boldsymbol{m}_i$$

This implies that $\sum_{i=1}^n \boldsymbol{e}_i \wedge \boldsymbol{m}_i = 0$, which means that $[I, M] = 0$, and therefore M is symmetric. □

4.3 Left orthogonal of Hankel matrices

In this subsection, we show that Hankel matrices behave differently with respect to ϕ, compared to an arbitrary symmetric matrix, even if they are hidden by the secret non-singular matrix S.

Hankel matrices are known to have rank-2 displacements: there are pairs of matrices (M, N) in $\mathcal{M}_n(\mathbb{K})$ such that for any Hankel matrix $H \in \mathcal{H}_n(\mathbb{K})$, the

matrix $HM^\mathsf{T} - NH$ has rank ≤ 2. With suitable choices of (M, N), this allows fast linear algebra operations on Hankel matrices (see the book [27]). Here, we exploit a special rank-2 displacement (M, N) where $M = N$, which means that the matrix $HM^\mathsf{T} - NH$ not only has rank ≤ 2, but is also equal to $[H, M]$, which is alternating, and therefore can be rewritten as a 2-form.

Most displacements of Hankel matrices are related to the (downward) *shift matrix* Z:

$$Z = \begin{bmatrix} 0 & 0 & 0 & \dots & 0 \\ 1 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix}.$$

Multiplication by Z shifts the coefficients of a column vector by one position downward, setting the first coefficient to 0. The computation given in the technical overview (see also [20, Eq. (0.5)]) shows that if $H \in \mathcal{H}_n(\mathbb{K})$ is a Hankel matrix, then $HZ^\mathsf{T} - ZH$ is alternating of rank ≤ 2, namely

$$[H, Z] = HZ^\mathsf{T} - ZH = \boldsymbol{e_1}\boldsymbol{h}^\mathsf{T} - \boldsymbol{h}\boldsymbol{e_1}^\mathsf{T} = \boldsymbol{e_1} \wedge \boldsymbol{h} \tag{5}$$

where $\boldsymbol{h} = Z\boldsymbol{h}_1 = (0, h_1, h_2, \dots, h_{n-1})^\mathsf{T}$ with $\boldsymbol{h}_1 = (h_1, \dots, h_n)^\mathsf{T}$ being the first column of H.

This implies, thanks to (1), that for any $S \in \mathrm{GL}_n(\mathbb{K})$:

$$[S^\mathsf{T}HS, S^\mathsf{T}ZS^{-\mathsf{T}}] = S^\mathsf{T}[H, Z]S = S^\mathsf{T}(\boldsymbol{e_1} \wedge \boldsymbol{h})S = (S^\mathsf{T}\boldsymbol{e_1}) \wedge (S^\mathsf{T}\boldsymbol{h}). \tag{6}$$

We generalize the linear map f of Lemma 2 as follows: for any $M \in \mathcal{M}_n(\mathbb{K})$, let $f_M : \mathbb{K}^n \to E$ be the linear map defined by

$$f_M(\boldsymbol{x}) = (\boldsymbol{x} \wedge \boldsymbol{m}_1, \dots, \boldsymbol{x} \wedge \boldsymbol{m}_n),$$

where $\boldsymbol{m}_1, \dots, \boldsymbol{m}_n$ denote the columns of M. Thus, the linear map f of Lemma 2 corresponds to f_M with M being the identity matrix. Equation (5) gives rise to more left-orthogonal vectors for a Hankel matrix, which is the first sign that ϕ can exploit the structure of Hankel matrices:

Lemma 3. *For any Hankel matrix $H \in \mathcal{H}_n(\mathbb{K})$, we have: $\phi(f_Z(\boldsymbol{e_1}), H) = 0$. And we have for any $S \in GL_n(\mathbb{K})$:*

$$\phi(f_{S^\intercal ZS^{-\intercal}}(S^\mathsf{T}\boldsymbol{e_1}), S^\mathsf{T}HS) = 0.$$

Proof. The equation (5) implies that:

$$\boldsymbol{e_1} \wedge [H, Z] = 0$$

On the other hand, if we denote the columns of H and Z by respectively $\boldsymbol{h}_1, \dots, \boldsymbol{h}_n$ and $\boldsymbol{z}_1, \dots, \boldsymbol{z}_n$, we have by Lemma 1:

$$HZ^\mathsf{T} - ZH = \sum_{i=1}^{n} \boldsymbol{h}_i \wedge \boldsymbol{z}_i.$$

Therefore:

$$\begin{aligned}\boldsymbol{e_1} \wedge [H, Z] &= \boldsymbol{e_1} \wedge \left(\sum_{i=1}^{n} \boldsymbol{h}_i \wedge \boldsymbol{z}_i \right) = \sum_{i=1}^{n} \boldsymbol{e_1} \wedge (\boldsymbol{h}_i \wedge \boldsymbol{z}_i) \\ &= -\sum_{i=1}^{n} (\boldsymbol{e_1} \wedge \boldsymbol{z}_i) \wedge \boldsymbol{h}_i = -\phi(f_Z(\boldsymbol{e_1}), H).\end{aligned}$$

Therefore $\phi(f_Z(\boldsymbol{e_1}), H) = 0$. Similarly, (6) implies that:

$$\phi(f_{S^\intercal Z S^{-\intercal}}(S^\mathsf{T} \boldsymbol{e_1}), S^\mathsf{T} H S) = 0. \qquad \square$$

Lemma 3 provides a new orthogonal vector for a Hankel matrix, compared to Lemma 2 with an arbitrary symmetric matrix. Lemma 2 gave an n-dimensional subspace for $M^\perp$. By generalizing Lemma 3, we can give a new $n-2$-dimensional subspace for $H^\perp$, which is in direct sum with the previous n-dimensional subspace Im f.

Lemma 4. *Let $H \in \mathcal{H}_n(\mathbb{K})$ whose first column is $\boldsymbol{h}_1$ and an integer r such that $1 \leq r \leq n$. Then:*

$$[H, Z^r] = \sum_{s=1}^{r} \boldsymbol{e}_s \wedge Z^{r+1-s} \boldsymbol{h}_1 = \sum_{s=1}^{r} (Z^{s-1} \boldsymbol{e_1}) \wedge (Z^{r+1-s} \boldsymbol{h}_1). \tag{7}$$

Proof. The telescoping identity [20, Lemma 1.1] shows that:

$$\begin{aligned}[H, Z^r] &= \sum_{s=0}^{r-1} Z^s [H, Z] \ (Z^\mathsf{T})^{r-1-s} \\ &= \sum_{s=0}^{r-1} Z^s (\boldsymbol{e_1} (Z \boldsymbol{h}_1)^\mathsf{T} - Z \boldsymbol{h}_1 \boldsymbol{e_1}^\mathsf{T}) (Z^\mathsf{T})^{r-1-s} \\ &= \sum_{s=1}^{r} \Big[\boldsymbol{e}_s (Z^{r+1-s} \boldsymbol{h}_1)^\mathsf{T} - Z^s \boldsymbol{h}_1 \, \boldsymbol{e}_{r+1-s}^\mathsf{T} \Big].\end{aligned}$$

Now, we split this in two sums, then reverse the order of the second sum:

$$\begin{aligned}[H, Z^r] &= \sum_{s=1}^{r} \boldsymbol{e}_s (Z^{r+1-s} \boldsymbol{h}_1)^\mathsf{T} - \sum_{s=1}^{r} Z^s \boldsymbol{h}_1 \, \boldsymbol{e}_{r+1-s}^\mathsf{T} \\ &= \sum_{s=1}^{r} \boldsymbol{e}_s (Z^{r+1-s} \boldsymbol{h}_1)^\mathsf{T} - \sum_{s=1}^{r} Z^{r+1-s} \boldsymbol{h}_1 \, \boldsymbol{e}_s^\mathsf{T} \\ &= \sum_{s=1}^{r} \Big[\boldsymbol{e}_s (Z^{r+1-s} \boldsymbol{h}_1)^\mathsf{T} - Z^{r+1-s} \boldsymbol{h}_1 \, \boldsymbol{e}_s^\mathsf{T} \Big] \\ &= \sum_{s=1}^{r} \boldsymbol{e}_s \wedge Z^{r+1-s} \boldsymbol{h}_1. \qquad \square\end{aligned}$$

Equation (7) gives rise to the following generalization of Lemma 3.

Theorem 2. *For any Hankel matrix $H \in \mathcal{H}_n(\mathbb{K})$, and $r \in \{1, \dots, n\}$:*

$$\phi(f_{Z^r}(\boldsymbol{e_1}) + f_{Z^{r-1}}(\boldsymbol{e}_2) + \cdots + f_Z(\boldsymbol{e}_r), H) = 0.$$

Furthermore, for any $S \in GL_n(\mathbb{K})$:

$$\phi(f_{S^\intercal Z^r S^{-\intercal}}(S^\mathsf{T}\boldsymbol{e_1}) + \cdots + f_{S^\intercal Z S^{-\intercal}}(S^\mathsf{T}\boldsymbol{e}_r), S^\mathsf{T} H S) = 0.$$

Proof. Denote by $\boldsymbol{z}_i^{[s]}$ the i-th column of Z^s. We have:

$$f_{Z^r}(\boldsymbol{e_1}) + f_{Z^{r-1}}(\boldsymbol{e}_2) + \cdots + f_Z(\boldsymbol{e}_r) = \sum_{s=1}^{r} f_{Z^{r+1-s}}(\boldsymbol{e}_s)$$

The i-th entry of $f_{Z^{r+1-s}}(\boldsymbol{e}_s)$ is $\boldsymbol{e}_s \wedge \boldsymbol{z}_i^{[r+1-s]}$. Thus:

$$\begin{aligned}
\phi(f_{Z^r}(\boldsymbol{e_1}) + f_{Z^{r-1}}(\boldsymbol{e}_2) + \cdots + f_Z(\boldsymbol{e}_r), H) &= \sum_{i=1}^{n} \sum_{s=1}^{r} \boldsymbol{e}_s \wedge \boldsymbol{z}_i^{[r+1-s]} \wedge \boldsymbol{h}_i \\
&= \sum_{s=1}^{r} \boldsymbol{e}_s \wedge \sum_{i=1}^{n} \boldsymbol{z}_i^{[r+1-s]} \wedge \boldsymbol{h}_i \\
&= \sum_{s=1}^{r} \boldsymbol{e}_s \wedge [Z^{r+1-s}, H]
\end{aligned}$$

But (7) implies that:

$$[Z^{r+1-s}, H] = -[H, Z^{r+1-s}] = -\sum_{t=1}^{r+1-s} \boldsymbol{e}_t \wedge Z^{r+2-s-t}\boldsymbol{h}_1.$$

Hence:

$$\phi(f_{Z^r}(\boldsymbol{e_1}) + f_{Z^{r-1}}(\boldsymbol{e}_2) + \cdots + f_Z(\boldsymbol{e}_r), H) = -\sum_{s=1}^{r} \sum_{t=1}^{r+1-s} \boldsymbol{e}_s \wedge \boldsymbol{e}_t \wedge Z^{r+2-s-t}\boldsymbol{h}_1$$

In this double sum, the index (s, t) runs over half a square:

- If $s = t$, $\boldsymbol{e}_s \wedge \boldsymbol{e}_t = 0$
- If $s \neq t$, any term (s, t) can be paired with (t, s), and the sum of the corresponding terms is

$$\begin{aligned}
&\boldsymbol{e}_s \wedge \boldsymbol{e}_t \wedge Z^{r+2-s-t}\boldsymbol{h}_1 + \boldsymbol{e}_t \wedge \boldsymbol{e}_s \wedge Z^{r+2-s-t}\boldsymbol{h}_1 \\
&= (\boldsymbol{e}_s \wedge \boldsymbol{e}_t + \boldsymbol{e}_t \wedge \boldsymbol{e}_s) \wedge Z^{r+2-s-t}\boldsymbol{h}_1 \\
&= 0.
\end{aligned}$$

We have thus shown that the whole double sum is zero. The proof of the second equality is similar. □

For $1 \leq r \leq n-1$, define

$$\kappa_r = \sum_{s=1}^{r} f_{Z^{r+1-s}}(\boldsymbol{e}_s).$$

Theorem 2 shows that $H^{\perp}$ contains the $(n-1)$-dimensional subspace spanned by $\kappa_1, \ldots, \kappa_{n-1}$, in addition to the n-dimensional subspace Im f. However, it should be noted that κ_{n-1} can be expressed as

$$\begin{aligned} &(\boldsymbol{e_1} \wedge \boldsymbol{e_n}, 0, ..., 0) \\ &+ (\boldsymbol{e}_2 \wedge \boldsymbol{e}_{n-1}, \boldsymbol{e}_2 \wedge \boldsymbol{e}_n, 0, ..., 0) \\ &+ ... \\ &+ (\boldsymbol{e}_{n-1} \wedge \boldsymbol{e}_2, \boldsymbol{e}_{n-1} \wedge \boldsymbol{e}_3, ..., \boldsymbol{e}_{n-1} \wedge \boldsymbol{e_n}, 0) \end{aligned}$$

which also belongs to Im f. The next lemma shows that there is no other linear dependency.

Lemma 5. *$f(\boldsymbol{e_1}), \ldots, f(\boldsymbol{e_n}), \kappa_1, \ldots, \kappa_{n-2}$ are linearly independent.*

Proof. Suppose that, for some scalars $\alpha_r, \beta_k \in \mathbb{K}$,

$$\sum_{r=1}^{n} \alpha_r f(\boldsymbol{e}_r) + \sum_{k=1}^{n-2} \beta_k \, \kappa_k = 0. \tag{8}$$

The i-th entry of κ_r is $(\kappa_r)_i = \sum_{s=1}^{r} \boldsymbol{e}_s \wedge \boldsymbol{e}_{i+r+1-s}$. First, looking at all the first entry of $f(\boldsymbol{e}_r)$ and κ_r in (8), we have

$$-\sum_{r=2}^{n} \alpha_r \, \boldsymbol{e_1} \wedge \boldsymbol{e}_r + \sum_{k=1}^{n-2} \beta_k \, \boldsymbol{e_1} \wedge \boldsymbol{e}_{k+1} = 0,$$

thus $\alpha_{k+1} = \beta_k \; \forall \, 1 \leq k \leq n-2$ and $\alpha_n = 0$.

Next, we look at the second entry of $f(\boldsymbol{e}_r)$ and κ_r. That gives:

$$\alpha_1 \, \boldsymbol{e_1} \wedge \boldsymbol{e}_2 - \sum_{r=3}^{n} \alpha_r \, \boldsymbol{e}_2 \wedge \boldsymbol{e}_r + \sum_{k=1}^{n-2} \beta_k \big(\boldsymbol{e_1} \wedge \boldsymbol{e}_{k+2} + \boldsymbol{e}_2 \wedge \boldsymbol{e}_{k+1} \big) = 0.$$

Looking at the coefficients of $\boldsymbol{e_1} \wedge \boldsymbol{e}_{k+2}$ for $1 \leq k \leq n-2$ gives $\beta_1 = \beta_2 = \cdots = \beta_{n-2} = 0$. Finally, from the previous relations $\alpha_{k+1} = \beta_k$, we deduce that $\alpha_2 = \cdots = \alpha_{n-1} = 0$, and we already know that $\alpha_n = 0$. Thus $\alpha_1 = \cdots = \alpha_n = 0$ and $\beta_1 = \cdots = \beta_{n-2} = 0$. Therefore the family $f(\boldsymbol{e_1}), \ldots, f(\boldsymbol{e_n}), \kappa^{(1)}, \ldots, \kappa^{(n-2)}$ is linearly independent. □

We are now in a position to provide an exact characterization of the left orthogonal of $\mathcal{H}_n$.

Theorem 3. *We have:*

$$\mathcal{H}_n^{\perp} = \mathrm{Im}\ f \oplus \mathrm{span}(\kappa_1, \ldots, \kappa_{n-2}), \tag{9}$$

$$\mathcal{H}_n^{\perp\perp} = \mathcal{H}_n. \tag{10}$$

Proof. See the full version. □

The next lemma shows that the double orthogonal commutes with the conjugation $H \mapsto S^{\mathsf{T}} H S$, which reduces the case of a general $S \in \mathrm{GL}_n(\mathbb{F}_q)$ to the Theorem above.

Lemma 6. *We have:*

$$\left(S^{\mathsf{T}} \mathcal{H}_n S\right)^{\perp\perp} = S^{\mathsf{T}} \mathcal{H}_n^{\perp\perp} S.$$

Proof. First, we define

$$\begin{aligned} \mu : \left(\bigwedge\nolimits^2 \mathbb{K}^n\right)^n &\longrightarrow \left(\bigwedge\nolimits^2 \mathbb{K}^n\right)^n \\ (\boldsymbol{x_1}, ..., \boldsymbol{x_n}) &\longmapsto (s(\boldsymbol{x_1}), ..., s(\boldsymbol{x_n}))\, S^{-\mathsf{T}} \end{aligned}$$

where $s(\boldsymbol{x_k}) = \sum_{i<j} \lambda_{ij} S^{\mathsf{T}} \boldsymbol{e_i} \wedge S^{\mathsf{T}} \boldsymbol{e_j}$ when we write $\boldsymbol{x_k}$ in the canonical basis of $\bigwedge^2 \mathbb{K}^n$. Then, we have $(S^{\mathsf{T}} \mathcal{H}_n S)^{\perp} = \mu(\mathcal{H}_n^{\perp})$. Indeed,

$$\begin{aligned} \phi\big(\mu(\boldsymbol{x}), S^{\mathsf{T}} \boldsymbol{y} S\big) &= \sum_{i=1}^{n} \mu(\boldsymbol{x})_i \wedge (S^{\mathsf{T}} \boldsymbol{y} S)_i \\ &= \sum_i \Big(\sum_k (S^{-1})_{ik}\, s(\boldsymbol{x}_k) \Big) \wedge \Big(\sum_l S_{li}\, S^{\mathsf{T}} \boldsymbol{y}_l \Big) \\ &= \sum_{k,l} \Big(\sum_i (S^{-1})_{ik} S_{li} \Big)\, s(\boldsymbol{x}_k) \wedge S^{\mathsf{T}} \boldsymbol{y}_l \\ &= \sum_{k,l} \Big(S^{-1} S \Big)\, s(\boldsymbol{x}_k) \wedge S^{\mathsf{T}} \boldsymbol{y}_l \\ &= \sum_{k=1}^{n} s(\boldsymbol{x}_k) \wedge S^{\mathsf{T}} \boldsymbol{y}_k = S^{-\mathsf{T}} \wedge S^{-\mathsf{T}} \wedge S^{-\mathsf{T}}\, \phi(\boldsymbol{x}, \boldsymbol{y}). \end{aligned} \tag{11}$$

Thus, if $y = \mathcal{H}_n$, then $x = \mathcal{H}_n^{\perp}$ and $\phi(\boldsymbol{x}, \boldsymbol{y}) = 0$. That means $\phi(\mu(\boldsymbol{x}), S^{\mathsf{T}} \boldsymbol{y} S) = 0$, in other words, $(S^{\mathsf{T}} \mathcal{H}_n S)^{\perp} = \mu(\mathcal{H}_n^{\perp})$. This implies $(S^{\mathsf{T}} \mathcal{H}_n S)^{\perp\perp} = \mu(\mathcal{H}_n^{\perp})^{\perp}$. However, from equation (11), we also see that $\boldsymbol{y} \in \mathcal{X}^{\perp}$ implies $S^{\mathsf{T}} \boldsymbol{y} S \in \mu(\mathcal{X})^{\perp}$. Therefore, $\mu(\mathcal{X})^{\perp} = S^{\mathsf{T}} \mathcal{X}^{\perp} S$. Hence, we conclude:

$$\left(S^{\mathsf{T}} \mathcal{H}_n S\right)^{\perp\perp} = \mu(\mathcal{H}_n^{\perp})^{\perp} = S^{\mathsf{T}} \mathcal{H}_n^{\perp\perp} S.$$ □

From Lemma 6 and Theorem 3, we have the following corollary:

Corollary 1. *We have:*

$$\left(S^{\mathsf{T}} \mathcal{H}_n S\right)^{\perp\perp} = S^{\mathsf{T}} \mathcal{H}_n S.$$

4.4 A polynomial-time distinguisher

The previous results already give us an efficient distinguisher, based on Conjecture 1 and Theorem 2.

Take any 4 matrices $A_{i_1}, \ldots, A_{i_4}$ from the X24 public key $(A_1, \ldots, A_m)$. We can compute the left orthogonal $\{A_{i_1}, \ldots, A_{i_4}\}^\perp$ in polynomial time, and we know that its dimension is at least $2n-2$. Indeed, $A_i \in E_S$. Therefore $E_S^\perp \subseteq A_i^\perp$, and by Theorem 2, the dimension of $E_S^\perp$ is at least $2n-2$.

On the other hand, if we are given 4 symmetric matrices $M_1, \ldots, M_4$ sampled uniformly at random, then Conjecture 1 says that the dimension of $\{M_1, \ldots, M_4\}^\perp$ is n.

Thus, by computing a left orthogonal subspace, we can distinguish between any 4 matrices of the public key and any 4 random symmetric matrices. Since computing a left orthogonal amounts to solving a linear system with

$$\dim\left(\left(\bigwedge^2 \mathbb{K}\right)^n\right) = n\binom{n}{2} = \mathcal{O}(n^3)$$

unknowns, the distinguisher runs in $\mathcal{O}(n^{3\omega})$ field operations.

4.5 Expanding the public key into E_S

Recall that $E_S = S^\mathsf{T}\mathcal{H}_n S$, and $A_1, \ldots, A_m$ is the X24 public key. We know that $E_S^\perp \subseteq \{A_1, \ldots, A_m\}^\perp$. We conjecture that these subspaces are actually equal.

Conjecture 2. Let $A_1, \ldots, A_m$ be sampled according to the X24 public key distribution. Then with overwhelming probability:

$$\{A_1, \ldots, A_m\}^\perp = E_S^\perp.$$

As discussed in the technical overview, Conjecture 2 is the only conjecture required for our overall key-recovery attack to proceed. A stronger form of the conjecture is that the equality holds as soon as $m \geq 4$, for any sufficiently large n. A corollary of this conjecture is:

$$(S^\mathsf{T}\mathcal{H}_n S)^{\perp\perp} = \{A_1, \ldots, A_m\}^{\perp\perp}. \tag{12}$$

Conjecture 2 and Corollary 1 translate to a polynomial-time algorithm to recover the secret subspace $E_S = S^\mathsf{T}\mathcal{H}_n S$.

– Compute a basis C of the left orthogonal $\{A_1, \ldots, A_m\}^\perp$. This amounts to computing the kernel of the linear map:

$$x \mapsto \Big(\phi(x, A_i)\Big)_{i=1}^m.$$

As discussed earlier, this costs $\mathcal{O}(n^{3\omega})$ field operations, because x lives in a space of dimension $\mathcal{O}(n^3)$ over the base field $\mathbb{K}$. If Conjecture 2 holds, C is also a basis of $(S^\mathsf{T}\mathcal{H}_n S)^\perp$. By Theorem 3 and Lemma 6, C must contain $2n-2$ elements $C_1, \ldots, C_{2n-2}$.

– Compute a basis B of the right orthogonal $B^{\perp}$. This amounts to computing the kernel of the linear map:

$$y \mapsto \Big(\phi(C_i, y)\Big)_{i=1}^{2n-2}.$$

Because y lives in a space of dimension n^2 over the based field $\mathbb{K}$, this costs $\mathcal{O}(n^{2\omega})$ field operations. By Theorem 3, B is a basis of $E_S = S^{\mathsf{T}}\mathcal{H}_n S$ (necessarily of cardinality $2n-1$).

The total number of field operations in this algorithm is $\mathcal{O}(n^{3\omega}) + \mathcal{O}(n^{2\omega}) = \mathcal{O}(n^{3\omega})$. (See Appendix C for a more detailed discussion of the complexity.)

Remark. The previous algorithm is very naturally interpreted as a "key expansion" algorithm. The set $\mathrm{span}\{A_1, \ldots, A_m\}$ is an m-dimensional subspace of E_S, which has dimension $2n-1$. If we take a basis $A_{m+1}, \ldots, A_{2n-1}$ of a complementary subspace of $\mathrm{span}\{A_1, \ldots, A_m\}$ within E_S, we can add it to the public key to form an expanded public key $\{A_1, \ldots, A_{2n-1}\}$. This results in a valid X24 public key for the same secret S as the original public key, with respect to a larger parameter $m = 2n-1$. (Note that $m = 2n-1$ is as large as possible, since $\dim(E_S) = 2n-1$.) In that sense, the algorithm we have just seen is a "key expansion" algorithm. For the purpose of our attack, however, any basis of E_S will do; it need not be an extension of the original public key.

5 Recovering rank-one matrices in E_S

The original X24 article [14, end of Sect. 5] noticed that selecting the parameter m too close to $2n-1$ could potentially make the scheme vulnerable to MinRank attacks. The previous section recovered a basis of the subspace E_S, which can allow us to expand the public key to a full public key of $2n-1$ matrices, as if $m = 2n-1$. In this section, we show that there are many rank-one matrices in E_S, and that a random rank-one matrix can provably be found in time linear in q (the finite field order) and polynomial in n, using the so-called kernel attack by Goubin and Courtois [18]. Our analysis of rank-one matrices in E_S shows that by repeating the kernel attack, we can obtain a set of $n+2$ pairwise linearly independent rank-one matrices in E_S, still in time linear in q and polynomial in n. This number $n+2$ is required to run the Sidelnikov-Shestakov attack (cf. Section 6).

Our analysis of the kernel attack is new.

1. The original kernel attack is heuristic. Thanks to the structure of Hankel matrices, the kernel attack can be made fully provable here.
2. The usual analysis of the kernel attack leads to a heuristic time bound *quadratic* in q and polynomial in n. We obtain a better bound that is only linear in q and still polynomial in n, due to the high number of rank-one matrices in the subspace E_S.

5.1 Preliminaries

If $\mathcal{S}$ is a finite set, denote by $x \xleftarrow{\$} \mathcal{S}$ the uniform sampling of an element x in $\mathcal{S}$. For two matrices M and M' in $\mathcal{M}_n(\mathbb{K})$, write $M \sim M'$ if and only if $M = cM'$ for some $c \in \mathbb{K}^*$. The relation $\sim$ is an equivalence relation. If $M \in \mathcal{M}_n(\mathbb{K}) \setminus \{0\}$, denote by $\mathsf{cl}(M)$ the class of M for $\sim$. If $M_1, M_2, \ldots, M_k \in \mathcal{M}_n(\mathbb{K}) \setminus \{0\}$ all have distinct equivalence classes for the relation $\sim$, they will be called *pairwise linearly independent*.

5.2 Rank-one matrices in E_S

In 1960, Daykin [11] showed that the number of rank-r matrices in $\mathcal{H}_n(\mathbb{F}_q)$ is exactly $(q^2-1)q^{2(r-1)}$ if $1 \leq r < n$, and $(q-1)q^{2(r-1)}$ if $r = n$. In particular, the number of rank-one matrices in $\mathcal{H}_n(\mathbb{F}_q)$ is exactly $q^2 - 1 = (q-1)(q+1)$.

It is well-known that if V is an $k \times n$ rectangular Vandermonde matrix $V = (\alpha_i^{j-1})$ where $\alpha_1, \ldots, \alpha_k \in \mathbb{F}_q$ and if D is $n \times n$ diagonal matrix, then $V^\mathsf{T} D V \in \mathcal{H}_n(\mathbb{F}_q)$ (which can easily be checked). Reciprocally, not all Hankel matrices are of this form, but all rank-one Hankel matrices are, provided that we add an extra Vandermonde vector "at infinity". More precisely, for any $\alpha \in \mathbb{F}_q$, let $\boldsymbol{v}_\alpha = (1, \alpha, \ldots, \alpha^{n-1})^\mathsf{T}$ and let $\boldsymbol{v}_\infty = \boldsymbol{e_n}$ be the last canonical vector. Let:

$$\mathcal{V} = \{c\boldsymbol{v}_\alpha \boldsymbol{v}_\alpha^\mathsf{T} \text{ where } c \in \mathbb{F}_q^*, \alpha \in \mathbb{F}_q \cup \infty\}.$$

Lemma 7. *Let $H \in \mathcal{H}_n(\mathbb{F}_q)$. Then H has rank 1 if and only if $H \in \mathcal{V}$.*

Proof. We already know the converse. Note that the number of distinct matrices in $\mathcal{V}$ in the reciprocal statement is exactly q^2-1, as there are $q+1$ choices for α and $q-1$ choices for c. This proves the direct statement, following the previously cited result by Daykin [11], which implies that there are exactly q^2-1 rank-one Hankel matrices. □

Corollary 2. *Let $M \in E_S$ for some $S \in \mathrm{GL}_n(\mathbb{F}_q)$. Then M has rank 1 if and only if:*

$$M \in S^\mathsf{T}\mathcal{V}S = \{c(S^\mathsf{T}\boldsymbol{v}_\alpha)(S^\mathsf{T}\boldsymbol{v}_\alpha)^\mathsf{T} \text{ where } c \in \mathbb{K}^*, \alpha \in \mathbb{F}_q \cup \infty\}.$$

The previous lemma is remarkable: it shows that rank-one Hankel matrices have a very specific algebraic structure, even if they are "hidden" by S. This is what will be exploited in Section 6 to recover the secret key.

Lemma 7 tells us that there are $\frac{q^2-1}{q-1} = q+1$ equivalence classes of rank-one Hankel matrices with respect to $\sim$. These equivalence classes exactly correspond to the elements of $\mathbb{F}_q \cup \{\infty\}$: two rank-one matrices are equivalent for $\sim$ if and only if they correspond to the same value of α.

5.3 The kernel attack

Let us recall the kernel attack [18] applied to our specific case (see Algorithm 1 for the pseudo-code). We know a basis $A_1, \ldots, A_{2n-1}$ of the space $E_S = S^\intercal \mathcal{H}_n(\mathbb{F}_q) S$. To retrieve a low-rank matrix in E_S, the kernel attack repeatedly samples two vectors $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{K}^n$ independently and uniformly at random, until the space of solutions to the linear system

$$(\Sigma) \begin{cases} \sum_{i=1}^{2n-1} x_i A_i \boldsymbol{u} = \boldsymbol{0} \\ \sum_{i=1}^{2n-1} x_i A_i \boldsymbol{v} = \boldsymbol{0} \end{cases}$$

with unknowns $x_1, \ldots, x_{2n-1}$ has dimension at least 1. This is a system of $2n$ equations over $2n-1$ unknowns, so whenever non-trivial solutions exist, we heuristically expect the solution space to have dimension 1, and the matrix $\sum_{i=1}^{2n-1} x_i A_i$ to have rank one.

Algorithm 1 KernelAttack

Input: $(A_1, \ldots, A_{2n-1})$ a basis of E_S
Output: $\mathsf{cl}(M)$ in $E_S/\sim$ such that $M \in S^\intercal \mathcal{V} S$

1: **while** true **do**
2: $\quad \boldsymbol{u}, \boldsymbol{v} \xleftarrow{\$} \mathbb{F}_q^n$
3: $\quad$ Find the linear space $F_{\boldsymbol{u},\boldsymbol{v}}$ of the solutions $(x_1, \ldots, x_{2n-1})$ to the system

$$(\Sigma) \begin{cases} \sum_{i=1}^{2n-1} x_i A_i \boldsymbol{u} = \boldsymbol{0} \\ \sum_{i=1}^{2n-1} x_i A_i \boldsymbol{v} = \boldsymbol{0} \end{cases}$$

4: $\quad$ **if** $\dim(F_{\boldsymbol{u},\boldsymbol{v}}) = 1$ **then**
5: $\quad\quad$ Retrieve $(x_1, \ldots, x_{2n-1}) \in F_{\boldsymbol{u},\boldsymbol{v}} \setminus \{0\}$ *//unique up to scalar multiplication*
6: $\quad\quad$ $M \leftarrow \sum_{i=1}^{2n-1} x_i A_i$ *//the rank of M is necessarily one*
7: $\quad\quad$ **return** $\mathsf{cl}(M)$
8: $\quad$ **end if**
9: **end while**

The standard analysis of the kernel attack is heuristic. For a fixed matrix H, the probability that $H\boldsymbol{u} = H\boldsymbol{v} = 0$ is exactly $1/q^{2\mathrm{rank}(H)}$. Unfortunately, in the general case, we do not know if the coefficients $(x_1, \ldots, x_{2n-1})$ correspond to our target H, and not even if the solution space has dimension 1. It could be that $\sum_{i=1}^{n-1} x_i A_i$ is a rank-two matrix, especially if there are many rank-two matrices in the search space. As a consequence, no provable result is known for the kernel attack in general.

Algorithm 2 RANK1RETRIEVAL

Input: $(A_1, \ldots, A_{2n-1})$ a basis of E_S
Output: $\{\mathsf{cl}(M_1), \ldots, \mathsf{cl}(M_{n+2})\} \subset S^\mathsf{T}\mathcal{V}S/\sim$

$\mathcal{S} \leftarrow \emptyset$ *// the solution set*
while $\mathrm{card}(\mathcal{S}) < n+2$ **do**
  $\mathsf{cl}(M) \leftarrow \textsc{KernelAttack}(A_1, A_2, \ldots, A_{2n-1})$
  if $\mathsf{cl}(M) \notin \mathcal{S}$ **then**
    $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathsf{cl}(M)\}$
  end if
end while
return $\mathcal{S}$

However, the structure of the Hankel matrices enables a provable version of the kernel attack. As stated in Proposition 1 (and proved in the upcoming full version of the paper), Algorithm 1 runs in expected time $\mathcal{O}(qn^\omega)$. Moreover, its output distribution is uniform over $S^\mathsf{T}\mathcal{V}S/\sim$. Another surprising fact resulting from the structure of the Hankel matrices is that as soon as the solution space $F_{\boldsymbol{u},\boldsymbol{v}}$ has dimension one, the matrix M recovered by Algorithm 1 is certain to have rank one.

Given the properties of Algorithm 1, retrieving $n+2$ pairwise linearly independent matrices of $S^\mathsf{T}\mathcal{V}S$ reduces to a simple coupon collector's problem (see Algorithm 2). The analysis of this classical problem is provided in Proposition 2.

5.4 Analysis of the attack

In this section, we provide some statements to analyze Algorithm 2, eventually showing that if $q > n$, then we can retrieve $n+2$ pairwise linearly independent matrices in expected time

$$\mathcal{O}\left(q^2 n^\omega \ln\left(\frac{q}{q-n}\right)\right) = \mathcal{O}(q \cdot \mathsf{poly}(n)).$$

Proposition 1. *Let $A_1, A_2, \ldots, A_{2n-1}$ be a basis of E_S.*
Then $\textsc{KernelAttack}(A_1, A_2, \ldots, A_{2n-1})$ *terminates in expected time $\mathcal{O}(qn^\omega)$ and returns a sample drawn from the uniform distribution on $S^\mathsf{T}\mathcal{V}S/\sim$.*

Given the previous proposition, analyzing Algorithm 2 corresponds to solving a simple coupon collector's problem.

Proposition 2 (Coupon collector's problem, [22, section 3.6]). *Assume that Alice draws coupons independently and uniformly at random with replacement from a pool of M coupons. Then the expected number of draws until Alice has drawn j pairwise distinct coupons is:*

$$M \sum_{i=0}^{j-1} \frac{1}{M-i}.$$

From Proposition 1 and Proposition 2, one deduces the following.

Proposition 3. *Assume that we are given a basis $A_1, \ldots, A_{2n-1}$ of E_S. Assume $n < q$. Then there exists a probabilistic algorithm that returns $n+2$ pairwise linearly independent rank-one matrices in E_S in expected time*

$$\mathcal{O}\left(q^2 n^\omega \ln\left(\frac{q}{q-n}\right)\right).$$

In particular, this complexity is $\mathcal{O}(q \cdot \mathsf{poly}(n))$, and in the case where $q = (1 + \Omega(1))n$, it can be simplified to

$$\mathcal{O}\left(qn^{\omega+1}\right).$$

6 Application of the Sidelnikov-Shestakov attack

6.1 The Sidelnikov-Shestakov attack

For completeness, we sketch here the Sidelnikov-Shestakov attack on the McEliece variant with Generalized Reed-Solomon codes (for more details, please refer to [34]). In this subsection, and this subsection only, n denotes the length of the code, and k is its dimension[7].

For pairwise distinct $\alpha_1, \ldots, \alpha_n \in \mathbb{K} \cup \{\infty\}$, and $z_1, \ldots, z_n \in \mathbb{K}^*$, define:

$$G_{\boldsymbol{\alpha},\boldsymbol{z}} = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ \alpha_1 & \alpha_2 & \cdots & \alpha_n \\ \alpha_1^2 & \alpha_2^2 & \cdots & \alpha_n^2 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^{k-1} & \alpha_2^{k-1} & \cdots & \alpha_n^{k-1} \end{pmatrix} \cdot \begin{pmatrix} z_1 & 0 & 0 & \cdots & 0 \\ 0 & z_2 & 0 & \cdots & 0 \\ 0 & 0 & z_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & z_n \end{pmatrix}$$

where as a special case, the column corresponding to α_i on the left-hand side is replaced by $\boldsymbol{e_n}$ whenever $\alpha_i = \infty$.

In the setting of Sidelnikov-Shestakov, the attacker is given $G = SG_{\boldsymbol{\alpha},\boldsymbol{z}}$, for unknown $\boldsymbol{\alpha}$, $\boldsymbol{z}$, $S \in \mathrm{GL}_n(\mathbb{K})$. The attacker's goal is to recover all unknowns, up to equivalence. That is, any solution $\boldsymbol{\alpha'}$, $\boldsymbol{z'}$, S' such that $S'G_{\boldsymbol{\alpha'},\boldsymbol{z'}} = SG_{\boldsymbol{\alpha},\boldsymbol{z}}$ is considered an equivalent solution, and accepted as valid.

Indeed, the solution is not unique. If $a, b, c, d \in \mathbb{K}$ are such that $ad - bc \neq 0$ then the birational transformation

$$f : x \mapsto \frac{ax+b}{cx+d}$$

maps any vector $\alpha_1, \ldots, \alpha_n$ that is part of a valid solution to another vector $f(\alpha_1), \ldots, f(\alpha_n)$ that is also part of a valid solution (paired with different column

[7] When we later apply the attack to our setting, n will become $n+2$, and k will become n, but in this subsection, it is more natural to treat the Sidelnikov-Shestakov attack for general k and n. It also simplifies the notation.

multipliers $z'_1, \ldots, z'_n$). Here, f acts on the point at infinity ∞ in the natural way, by sending it to ∞ if $c = 0$, to 0 if $a = 0$, and to a/c otherwise, noting that $a = c = 0$ is impossible.

Let us now describe how the Sidelnikov-Shestakov attack recovers a valid solution $\alpha_1, \ldots, \alpha_n$. Choose arbitrarily a constant $\gamma \in \mathbb{K} \backslash \{0, 1\}$. For any pairwise distinct $\alpha_1, \alpha_2, \alpha_{k+1}$, it is possible to find a birational transformation f such that $f(\alpha_1) = 0, f(\alpha_2) = 1, f(\alpha_{k+1}) = \gamma$. Hence, there exists a valid solution such that $\alpha_1 = 0, \alpha_2 = 1$, and $\alpha_{k+1} = \gamma$. We will look for a solution of that form. Therefore, we assume from now on that $\alpha_1 = 0, \alpha_2 = 1$, $\alpha_{k+1} = \gamma$.

Given B, one can find S' such that the generating matrix is in canonical form:

$$R = S'G = \left(\text{Id} \mid C \right)$$

where Id denotes the $k \times k$ identity matrix. Indeed, the leftmost $k \times k$ block of G is a Vandermonde matrix (for pairwise distinct α_i's), multiplied on the left by $S \in \mathrm{GL}_n\mathbb{K}$ and on the right by a diagonal matrix of full rank, which implies that it is invertible.

Each row of R contains the evaluations at the points $\alpha_1, \ldots, \alpha_n$ of an unknown polynomial of degree $k - 1$. Since the i-th row is zero at exactly the k first coordinates except for the i-th one, the corresponding polynomial P_i can only be of the form:

$$P_i = c_i \prod_{\substack{j=1 \\ j \neq i}}^{k} (X - \alpha_j) \,.$$

For $j \geq k + 1$, it holds that:

$$R[i, j] = z_j P_i(\alpha_j)$$

Therefore, for all $j \geq k + 1$, we can compute the ratio:

$$\frac{R_{2,j}}{R_{1,j}} = \frac{c_2}{c_1} \frac{\alpha_j - \alpha_1}{\alpha_j - \alpha_2} = \frac{c_2}{c_1} \frac{\alpha_j}{\alpha_j - 1},$$

which allows us in turn to compute:

$$\omega_j = \frac{R_{2,j}}{R_{1,j}} \frac{R_{1,k+1}}{R_{2,k+1}} = \frac{\alpha_j}{\alpha_j - 1} \frac{\alpha_{k+1} - 1}{\alpha_{k+1}} = \frac{\alpha_j}{\alpha_j - 1} \frac{\gamma - 1}{\gamma}.$$

With this cross-ratio we can recover α_j for all $j \geq k + 2$ as:

$$\alpha_j = \frac{\omega_j}{\omega_j - \frac{\gamma - 1}{\gamma}}.$$

Finally, to recover α_j for $3 \leq j \leq k$, we compute:

$$\frac{R_{j,k+1}}{R_{1,k+1}} \frac{R_{1,k+2}}{R_{j,k+2}} = \frac{\alpha_{k+1} - \alpha_1}{\alpha_{k+1} - \alpha_j} \frac{\alpha_{k+2} - \alpha_j}{\alpha_{k+2} - \alpha_1} \,.$$

Since we know all variables in this cross-ratio except α_j, we can recover it.

We now wish to recover the z_i's. First, remark that we have both $R_{i,i} = 1$ and $R_{i,i} = z_i c_i \prod_{\substack{j=1 \\ j \neq i}} (\alpha_i - \alpha_j)$ we can deduce that:

$$c_i = \frac{1}{z_i \prod_{\substack{j=1 \\ j \neq i}}^{k} (\alpha_i - \alpha_j)}.$$

Therefore:

$$R_{i,j} = \frac{z_j}{z_i} \prod_{\substack{l=1 \\ l \neq i}}^{k} \frac{\alpha_j - \alpha_l}{\alpha_i - \alpha_l},$$

hence for $i \geq k+1$:

$$z_i = z_1 \cdot R_{1,i} \prod_{j=2}^{k} \frac{\alpha_1 - \alpha_j}{\alpha_i - \alpha_j}$$

and for $2 \leq i \leq k$

$$z_i = \frac{z_{k+1}}{R_{i,k+1}} \prod_{\substack{j=1 \\ j \neq i}}^{k} \frac{\alpha_{k+1} - \alpha_j}{\alpha_i - \alpha_j}.$$

Since multiplying all the column multipliers z_i by the same constant does not affect the row-span of the resulting generating matrix, we are free to set $z_1 = 1$, which allows us to recover all the other z_i's.

We have now recovered valid solutions for $\boldsymbol{\alpha}$ and $\boldsymbol{z}$ (they must be valid since they were computed using equations that necessarily hold). This allows us to recover the desired S with an additional $O(k^\omega)$ field operations. If our only goal is to recover S, then we are free to truncate G to its first $k+2$ columns before running the attack. In that case, the total complexity to recover S is $O(k^\omega)$ field operations.

6.2 Application to X24 key recovery

By combining the results of Section 4 and Proposition 3, we know that if $q > n$, then we can recover $n+2$ rank-one matrices from E_S which correspond to distinct Vandermonde vectors. More precisely, by Lemma 7, we know $n+2$ matrices of the form $S^\mathsf{T} \boldsymbol{v}_{\boldsymbol{\alpha_i}} \boldsymbol{v}_{\boldsymbol{\alpha_i}}{}^\mathsf{T} S$ for $1 \leq i \leq n+2$. From there, we can easily recover the values of $n+2$ vectors of the form $cS^\mathsf{T} \boldsymbol{v}_{\boldsymbol{\alpha}}$ with $c \in \mathbb{K}^*$. This means that we exactly have access to a matrix of the form $S^\mathsf{T} V C$ where $S \in GL_n(\mathbb{K})$, $C = \mathrm{Diag}(c_1, \ldots, c_{n+2})$ and $V = [\boldsymbol{v}_{\boldsymbol{\alpha_1}}, \boldsymbol{v}_{\boldsymbol{\alpha_2}}, \ldots, \boldsymbol{v}_{\boldsymbol{\alpha_{n+2}}}]$ is the matrix that has $\boldsymbol{v}_{\boldsymbol{\alpha_i}}$ as columns. This is exactly the setting of the Sidelnikov-Shestakov attack, which allows one to retrieve the matrix S in time $O(n^\omega)$. This reveals the secret key corresponding to the original X24 public key $\{A_i \mid 1 \leq i \leq m\}$ (up to equivalence), and concludes the attack.

7 Experiments

7.1 Implementation of the attack

We implemented the complete key recovery attack using the Sagemath open-source mathematical software. For the parameters proposed for X24 ($q = 256$, $n = 24$), our implementation finds an equivalent secret key from a X24 public key in a few minutes on a standard laptop computer. A link to the implementation is given in Section 1.1.

Remark on sparsity. Beside time complexity, one of the main bottlenecks of our implementation is its memory complexity. Because the linear spaces we consider are quite high-dimensional, the matrices representing the linear applications that compute the left orthogonal and bi-orthogonal of the public key are rather large. However, these matrices are very sparse. Our current implementation does not take advantage of that fact (considering how fast the attack is for the real X24 parameters, it did not seem necessary). With larger parameters, using sparse linear algebra could save a large amount of memory.

7.2 Additional evidence to support Conjectures 1 and 2

Notation. In this section, we will systematically denote by $M, M_1, M_2, \dots$ matrices that are sampled from the uniform distribution over symmetric matrices in $\mathcal{M}_n(\mathbb{K})$, and by $A, A_1, A_2, \dots$ matrices that are sampled from the uniform distribution over E_S.

Grassmann's formula. Recall that if X and Y are two subspaces of a vector space Z that are in a general position, then we expect:

$$\dim(X \cap Y) = \max(\dim(X) + \dim(Y) - \dim(Z), 0).$$

Supporting evidence for the conjectures. We have performed experiments to confirm Conjectures 1 and 2 for many values of m, n, and q. Both conjectures always hold as soon as $m \geq 4$, regardless of the order q of the finite field, including both odd and even characteristic. In this section, we propose a deeper look into the behavior of these conjectures, both to help support them, and to provide some additional insight into the attack.

Fix a sequence of uniformly random full-rank symmetric matrices. Compute the left orthogonal of the first k matrices, as a function of k. This results in a sequence of nested subspaces: each subspace contains the next subspace in the sequence. Conjecture 1 can be reformulated as: this sequence remains stationary as soon as $k \geq 4$. This is enough to imply the conjecture, because for large enough k, we know that the left orthogonal is Im f (Theorem 1), as claimed by Conjecture 1. Similarly, if we compute the left orthogonal of the first k matrices in an X24 public key of maximal length (i.e. $m = 2n - 1$), the stronger form of Conjecture 2 can also be reformulated as: this sequence is stationary as soon as $k \geq 4$. Indeed, similar to the random case, Theorem 3 implies that for large

enough k (namely $k = 2n - 1$), the left orthogonal has the form predicted by Conjecture 2.

In the rest of this section, we analyze and perform experiments regarding what happens near the "critical point" $k = 4$, in both the random and real cases. Regarding conjecture 1, we already know that if $M_1, \ldots, M_4 \in \mathcal{M}_n(\mathbb{K})$ are symmetric matrices, then $\text{Im } f \subset \{M_1, \ldots, M_4\}^\perp$. We recall from section 4 that $\text{Im } f$ has dimension n. Similarly, regarding conjecture 2, we already know that

$$(S^\mathsf{T}\mathcal{H}_n S)^\perp \subset \{A_1, A_2, A_3, A_4\}^\perp$$

where A_1, A_2, A_3, A_4 are four matrices coming from an X24 public key. Because of these inclusions, studying the dimensions of these spaces suffices to gain a better understanding of what is happening.

Specifically, on the one hand, we computed the dimension of $\text{span}(M_1, \ldots, M_k)^\perp$ for $k \in \{2, 3, 4\}$ in the case where $M_1, \ldots, M_k$ are drawn independently and uniformly among the symmetric matrices of $\mathcal{M}_n(\mathbb{K})$. On the other hand, we have computed the dimension of $\text{span}(A_1, \ldots, A_k)^\perp$ for $k \in \{2, 3, 4\}$ in the case where $A_1, \ldots, A_k$ are matrices coming from an X24 public key. Throughout all our experiments, we found that these dimensions depend only on k.

This observation can be explained by the following two heuristic assumptions.

- *(Assumption 1).* We assume that $M_i^\perp$ can be understood as $M_i^\perp = \text{Im } f \oplus U_i$ where U_i is a generic subspace of E in direct sum with $\text{Im } f$. Similarly, we assume that $M_i^\perp = (S^\mathsf{T}\mathcal{H}S)^\perp \oplus V_i$ where V_i is a generic subspace of E in direct sum with $(S^\mathsf{T}\mathcal{H}S)^\perp$.
- *(Assumption 2).* We assume that $x \in E \mapsto \phi(x, M) \in \bigwedge^3 \mathbb{K}^n$ and $x \in E \mapsto \phi(x, A) \in \bigwedge^3 \mathbb{K}^n$ are surjective for a uniformly random symmetric matrix $M \in \mathcal{M}_n(\mathbb{K})$ and a random matrix $A \in S^\mathsf{T}\mathcal{H}_n S$, which is verified in practice and could be expected since $\dim(E) = \frac{n^2(n-1)}{2}$ and $\dim(\bigwedge^3 \mathbb{K}^n) = \frac{n(n-1)(n-2)}{6}$.

The second assumption was checked by applying the rank-nullity theorem. Specifically, we checked that for some uniformly random symmetric matrices M in $\mathcal{M}_n(\mathbb{K})$ and some uniformly random symmetric matrices A in E_S:

$$\dim(M^\perp) = \dim(A^\perp) = \dim(E) - \dim(\textstyle\bigwedge^3 \mathbb{K}^n) = n\binom{n}{2} - \binom{n}{3} = 2\binom{n+1}{3}.$$

Under assumptions 1 and 2, we expect to observe the following:

$$\begin{aligned}
\dim(\{M_1, M_2\}^\perp) &= \dim(\text{Im } f) + \max(\dim(U_1) + \dim(U_2) - \dim(E/\text{Im } f), 0) \\
&= \max(\dim(M_1^\perp) + \dim(M_2^\perp) - \dim(E), \dim(\text{Im } f)) \\
&= \max\left(4\binom{n+1}{3} - n\binom{n}{2}, n\right).
\end{aligned}$$

A similar computation shows:

$$\dim(\{M_1, M_2, M_3\}^\perp) = \max\left(6\binom{n+1}{3} - 2n\binom{n}{2}, n\right),$$

and by induction:

$$\dim(\{M_1, M_2, \ldots, M_k\}^\perp) = \max\left(2k\binom{n+1}{3} - (k-1)n\binom{n}{2}, n\right).$$

In an analogous way, one could show that

$$\dim(\{A_1, A_2, \ldots, A_k\}^\perp) = \max\left(2k\binom{n+1}{3} - (k-1)n\binom{n}{2}, 2n-2\right).$$

The validity of these last two formulas has been extensively tested and confirmed for $k \in \{2, 3, 4\}$ and $n \in \{10, 11, 12, 13, 14, 15\}$.
One can notice that for $n \geq 8$, we have

$$2k\binom{n+1}{3} - (k-1)n\binom{n}{2} > 2n-2 \quad \text{if } k \leq 3$$

and conversely

$$2k\binom{n+1}{3} - (k-1)n\binom{n}{2} \leq n \quad \text{if } k \geq 4.$$

We believe that this provides a satisfactory (if heuristic) explanation as to why just 4 (out of $m = 40$) matrices of the X24 public key suffice, both to distinguish it from a tuple of random symmetric matrices, and to perform key expansion in the first step of our attack.

Acknowledgements This work was supported by the France 2030 ANR Project ANR-22-PECY-003 SecureCompute, ANR grant PQMC (ANR-25-CE39-7662), and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 885394).

References

1. Ryan Babbush, Adam Zalcman, Craig Gidney, Michael Broughton, Tanuj Khattar, Hartmut Neven, Thiago Bergamaschi, Justin Drake, and Dan Boneh. Securing elliptic curve cryptocurrencies against quantum vulnerabilities: Resource estimates and mitigations. Cryptology ePrint Archive, Paper 2026/625, 2026.
2. Magali Bardet, Pierre Briaud, Maxime Bros, Philippe Gaborit, and Jean-Pierre Tillich. Revisiting algebraic attacks on minrank and on the rank decoding problem. *Designs, Codes and Cryptography*, 91(11):3671–3707, 2023.
3. Magali Bardet, Maxime Bros, Daniel Cabarcas, Philippe Gaborit, Ray A. Perlner, Daniel Smith-Tone, Jean-Pierre Tillich, and Javier A. Verbel. Improvements of algebraic attacks for solving the rank decoding and minrank problems. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2020 - 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7-11, 2020, Proceedings, Part I*, Lecture Notes in Computer Science, pages 507–536. Springer, 2020.

4. Magali Bardet and Alban Gilard. Computation of the hilbert series for the support-minors modeling of the minrank problem. *arXiv preprint arXiv:2502.12721*, 2025.
5. Ward Beullens, Fabio Campos, Sofía Celi, Basil Hess, and Matthias J. Kannwischer. MAYO: Specification Document – Round 2. Technical Report Round 2, March 2025. Available at NIST CSRC.
6. Ward Beullens, Ming-Shing Chen, Jintai Ding, Boru Gong, Matthias J. Kannwischer, Jacques Patarin, Bo-Yuan Peng, Dieter Schmidt, Cheng-Jhih Shih, Chengdong Tao, and Bo-Yin Yang. UOV: Unbalanced oil and vinegar specification document - round 2. Submission to the nist post-quantum cryptography standardization process, National Institute of Standards and Technology, February 2025. Version 2.0.
7. Gilles Brassard, Peter Hoyer, and Alain Tapp. Quantum algorithm for the collision problem. *ACM SIGACT News (Cryptology Column)*, 1997.
8. Maxime Bros, Thai Hung Le, Jacob Lichtinger, Brice Minaud, Ray Perlner, Daniel Smith-Tone, and Cristian Valenzuela. Exploiting SNOVA's structure in the wedge product attack. Cryptology ePrint Archive, Paper 2026/237, 2026.
9. André Chailloux, María Naya-Plasencia, and André Schrottenloher. An efficient quantum collision search algorithm and implications on symmetric cryptography. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, Lecture Notes in Computer Science, pages 211–240. Springer, 2017.
10. Benoît Cogliati, Gilles Macariot-Rat, Jacques Patarin, and Pierre Varjabedian. State of the art of HFE variants: Is it possible to repair HFE with appropriate modifiers? In *International Conference on Post-Quantum Cryptography*, pages 144–167. Springer, 2024.
11. David E. Daykin. Distribution of bordered persymmetric matrices in a finite field. *Journal für die Reine und Angewandte Mathematik. [Crelle's Journal]*, 203:47–54, 1960.
12. Thomas Debris-Alazard, Philippe Gaborit, Romaric Neveu, and Olivier Ruatta. A minrank-based encryption scheme à la Alekhnovich-Regev. Cryptology ePrint Archive, Paper 2025/1879 (to appear at Eurocrypt 2026), 2025.
13. Thomas Debris-Alazard, Nicolas Sendrier, and Jean-Pierre Tillich. Wave: A new family of trapdoor one-way preimage sampleable functions based on codes. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part I*, Lecture Notes in Computer Science, pages 21–51. Springer, 2019.
14. Irene Di Muzio, Martin Feussner, and Igor Semaev. Multivariate signatures with polynomial factorization. In Magali Bardet and Ruben Niederhagen, editors, *Post-Quantum Cryptography*, Lecture Notes in Computer Science, pages 355–369, Cham, 2026. Springer Nature Switzerland.
15. Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Prest, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON: Fast-fourier lattice-based compact signatures over NTRU. Submission to the nist post-quantum cryptography standardization process, National Institute of Standards and Technology, October 2020. Version 1.2.
16. Hiroki Furue, Yasuhiko Ikematsu, Fumitaka Hoshino, Tsuyoshi Takagi, Haruhisa Kosuge, Kimihiro Yamakoshi, Rika Akiyama, Satoshi Nakamura, Shingo Orihara,

and Koha Kinjo. QR-UOV: Specification Document – Round 2. Technical Report Round 2, March 2025. Available at NIST CSRC.
17. Google Cloud. Post-quantum cryptography (PQC), 2026.
18. L. Goubin and N. T. Courtois. Cryptanalysis of the TTM cryptosystem. In T. Okamoto, editor, *Advances in Cryptology — ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, Berlin, Heidelberg, 2000. Springer.
19. Lov K Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996.
20. Georg Heinig and Karla Rost. *Algebraic Methods for Toeplitz-like Matrices and Operators*, volume 13 of *Operator Theory: Advances and Applications*. Birkhäuser, Basel, 1984.
21. Yi Jin, Yuansheng Pan, Xiaoou He, Boru Gong, and Jintai Ding. Security analysis on UOV families with odd characteristics: Using symmetric algebra. Cryptology ePrint Archive, Paper 2025/1137, 2025.
22. Rajeev Motwani and Prabhakar Raghavan. Randomized algorithms. *ACM Computing Surveys (CSUR)*, 28(1):33–37, 1996.
23. National Security Agency. Announcing the commercial national security algorithm suite 2.0. Cybersecurity advisory, National Security Agency, September 2022.
24. Phong Q. Nguyen and Jacques Stern. Merkle-Hellman revisited: A cryptanalysis of the Qu-Vanstone cryptosystem based on group factorizations. In Burton S. Kaliski Jr., editor, *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, Lecture Notes in Computer Science, pages 198–212. Springer, 1997.
25. Phong Q. Nguyen and Jacques Stern. The two faces of lattices in cryptology. In Joseph H. Silverman, editor, *Cryptography and Lattices, International Conference, CaLC 2001, Providence, RI, USA, March 29-30, 2001, Revised Papers*, Lecture Notes in Computer Science, pages 146–180. Springer, 2001.
26. NIST. Call for Additional Digital Signature Schemes for the Post-quantum Cryptography Standardization Process. NIST CSRC, 2022. https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/call-for-proposals-dig-sig-sept-2022.pdf.
27. Victor Y. Pan. *Structured Matrices and Polynomials: Unified Superfast Algorithms*. Birkhäuser, Boston, 2001.
28. Jacques Patarin, Nicolas Courtois, and Louis Goubin. Quartz, 128-bit long digital signatures: http://www. minrank. org/quartz. In *Cryptographers' Track at the RSA Conference*, pages 282–297. Springer, 2001.
29. Jacques Patarin and Pierre Varjabedian. Multivariate encryptions with LL' perturbations - is it possible to repair HFE in encryption? -. Cryptology ePrint Archive, Paper 2024/1999, 2024.
30. Lars Ran. Wedges, Oil, and Vinegar: A New Algorithm for UOV in Characteristic 2. Rump Session of Eurocrypt, 2025.
31. Lars Ran. Wedges, oil, and vinegar – an analysis of UOV in the exterior algebra. In *Proc. EUROCRYPT*, 2026. To appear.
32. Hugues Randriambololona. The syzygy distinguisher. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology - EUROCRYPT 2025 - 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4-8, 2025, Proceedings, Part VI*, Lecture Notes in Computer Science, pages 324–354. Springer, 2025.

33. Peter W Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.
34. Vladimir Michilovich Sidelnikov and Sergey O. Shestakov. On insecurity of cryptosystems based on generalized Reed-Solomon codes. *Discrete Mathematics and Applications*, 1992.
35. Daniel Smith-Tone and Cristina Tone. A nonlinear multivariate cryptosystem based on a random linear code. Cryptology ePrint Archive, Paper 2019/1355, 2019.
36. Gang Tang, Dung Hoang Duong, Antoine Joux, Thomas Plantard, Youming Qiao, and Willy Susilo. Practical post-quantum signature schemes from isomorphism problems of trilinear forms. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part III*, Lecture Notes in Computer Science, pages 582–612. Springer, 2022.
37. Po-En Tseng, Lih-Chung Wang, Peigen Li, and Yen-Liang Kuan. Investigating the wedge map on SNOVA. Cryptology ePrint Archive, Paper 2026/260, 2026.
38. Lih-Chung Wang, Chun-Yen Chou, Jintai Ding, Yen-Liang Kuan, Jan Adriaan Leegwater, Ming-Siou Li, Bo-Shu Tseng, Po-En Tseng, and Chia-Chun Wang. SNOVA: Specification Document – Round 2. Technical Report Round 2, March 2025. Available at NIST CSRC.
39. Bas Westerbaan. Cloudflare targets 2029 for full post-quantum security.
40. Takeo Yokonuma. *Tensor spaces and exterior algebra*. Number 108. American Mathematical Soc., 1992.

A Generic quantum forgery attacks

At the end ot their article, the authors of [14] provide an example parameter set, targeting NIST level-I security, viz. 128 bits of post-quantum security (the NIST call specifies the computation unit as one AES-128-computation-equivalent, so that the NIST-I security level translates to 2^{143} quantum gates). The signature size of X24 is $n = 24$ bytes. This is generically too small for the target security level.

Indeed, an attacker attempting to forge a message can sample a signature uniformly at random and run the verification algorithm to check whether the guessed signature is valid. Classically, this generic attack would cost $2^{8n} = 2^{192}$ attempts, each consisting of one run of the verification algorithm; but with Grover search, the cost comes down to 2^{96} evaluations of the verification circuit [19], which is below the target security level. This is assuming the worst-case scenario where there exists only one valid signature per message. In the case of X24, the central trapdoor function maps $2n = 48$ bytes to $m = 40$ bytes, so that each message has on average $2^{8(2n-m)} = 2^{64}$ valid signatures. This decreases the expected quantum cost of the forgery attack to 2^{64} evaluations of the verification circuit with Grover search.

A secondary concern is that $m = 40$ bytes at the output of the inner trapdoor function may also be generically too small. This is because, as with any signature following the hash-and-sign paradigm, the message must first be mapped to the output of the trapdoor function using a hash function h. If an adversary is able to find a collision $h(m_1) = h(m_2)$ for two messages m_1 and m_2, then she obtains a break in the standard EUF-CMA security model, by asking for a signature of m_1, then reusing the same signature to sign m_2. The quantum cost of a collision search over b bits is roughly $2^{b/3}$ [7] (compared to $2^{b/2}$ classically by the birthday paradox). This yields a quantum complexity of $2^{8m/3} \approx 2^{107}$, which is again below the security parameter. We note however that the cost of quantum collision depends on how quantum memory is modeled [9].

Because of these generic attacks, and because the security analysis in [14] seems to focus on classical attacks, we view the parameter set at the end of [14] as corresponding to 128 bits of classical security, moreso than 128 bits of post-quantum security (despite the stated NIST level-I security target). In any case, these generic attacks are exponential in nature, and could be easily thwarted by an increase in parameters, unlike our main attack.

B Protecting against the attack

The public key of X24 has size $\mathcal{O}(n^3 \log q)$ bits. Our attack has complexity $\mathcal{O}(q \cdot \mathsf{poly}(n))$, which means that it is technically not a polynomial-time attack: its running time is not polynomial in the size of its input, due to the linear dependency on q. (As noted in the experimental section, in practice, our attack recovers the secret key in a few minutes on the full parameters.)

In theory, modifying X24 so that the order q of the finite field becomes exponential in the security parameter would appear to defeat the attack. However, we do not expect that such a countermeasure is viable, for two reasons.

1. First, the only part of our attack that has a superpolynomial computation time of the form $q \cdot \mathsf{poly}(n)$ is the MinRank step. The linear dependency on q comes from using the kernel attack approach for MinRank. We believe this dependency could be removed by using other MinRank algorithms, such as minors modeling [3,2]. The reason we use the kernel attack is because, thanks to properties of Hankel matrices, we are able to bound the running time of the kernel attack unconditionally. With other MinRank approaches, a heuristic regularity assumption may be required. Still, we expect that those other approaches would work in practice, and remove the dependency on q. In short, it is likely that the linear dependency on q can be removed from the attack, at the cost of a heuristic regularity assumption.
2. Second, setting security aside for a moment, we do not see a way to pick a large enough q that would not also result in signature and key sizes that are worse than competing schemes such as UOV or MAYO, defeating the point of X24.

Another limitation of our attack is that the MinRank step requires $q > n$, in order to recover sufficiently many $(n+2)$ rank-one matrices for the next step of the attack. Thus, another apparent way to defeat the attack would be to set $q \leq n$, for instance $q = 2$. In the full version we will show that this case is not secure either.

C Complexity analysis

In the first step of our attack, when solving a linear system in k unknowns, we estimate the complexity to be $\mathcal{O}(k^\omega)$. This is true when the system has $\mathcal{O}(k)$ equations. If it has ℓ equations for some $\ell > k$, then a general upper bound is $\ell k^{\omega-1}$ (obtained by sampling $\mathcal{O}(k)$ random linear combinations of the equations, which costs $\mathcal{O}(\ell/k)$ multiplications of $k \times k$ matrices, and reduces to the first case).

In our setting, we can avoid this extra factor $\mathcal{O}(\ell/k)$ by simply taking fewer equations. For the left orthogonal computation, the stronger form of Conjecture 2 mentioned in Section 4.5 (and further analyzed in Section 7.2) states that we can use just 4 of the public key matrices. The rest of this section is devoted to explaining this point in more detail. This justifies why we use $\mathcal{O}(k^\omega)$ for our complexity estimates when computing the left orthogonal and bi-orthogonal.

Proposition 4 (Cost for computing the left kernel). *Given $B_1, \ldots, B_m \in \left(\bigwedge^1 \mathbb{K}^n\right)^n$ either sampled from the uniform distribution on symmetric matrices of $\mathcal{M}_n(\mathbb{K})$ or sampled uniformly from the set E_S. Assuming the stronger form of Conjecture 2 from Section 4.5, the linear space*

$$\{B_1, \ldots, B_m\}^\perp \subset \left(\bigwedge^2 \mathbb{K}^n\right)^n$$

can be computed in time $\mathcal{O}(n^{3\omega})$.

Proof. The remarks made in Section 4 imply that, assuming the stronger form of Conjecture 2,

$$\{B_1, B_2, \ldots, B_m\}^\perp = \{B_1, B_2, B_3, B_4\}^\perp .$$

Thus, computing $\{B_1, B_2, \ldots, B_m\}^\perp$ is equivalent to computing the kernel of the application:

$$\begin{aligned} \psi_1 : E &\to \left(\bigwedge^3 \mathbb{K}^n\right)^4 \\ x &\mapsto (\phi(x, B_1), \phi(x, B_2), \phi(x, B_3), \phi(x, B_4)). \end{aligned}$$

The domain of this application is a linear space of dimension $n\binom{n}{2} = \mathcal{O}(n^3)$ and its codomain is a linear space of dimension $4\binom{n}{3} = \mathcal{O}(n^3)$. So the kernel of this application can be computed in time $\mathcal{O}(n^{3\omega})$. □

Proposition 5 (Cost for computing the right kernel). *Let* $\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_{2n-2} \in \left(\bigwedge^2 \mathbb{K}^n\right)^n$. *Then the linear space*

$$\{\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_{2n-2}\}^\perp$$

can heuristically be computed in time $\mathcal{O}(n^{2\omega})$.

Proof. The computation of the linear space $\{\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_{2n-2}\}^\perp$ is exactly equivalent to the computation of the kernel of the linear application

$$\begin{aligned} \psi_2 : F &\to \left(\bigwedge^3 \mathbb{K}^n\right)^{2n-2} \\ (\boldsymbol{b_1}, \boldsymbol{b_2}, \ldots, \boldsymbol{b_{2n-2}}) &\mapsto (\phi(\boldsymbol{b}_1, y), \phi(\boldsymbol{b}_2, y), \ldots, \phi(\boldsymbol{b}_{2n-2}, y)). \end{aligned}$$

The domain of this application is a linear space of dimension n^2 and its codomain is a linear space of dimension $(2n-2)\binom{n}{3} = O(n^4)$. The resulting system of equations is largely overdetermined. We can heuristically extract $O(n^2)$ equations and still find the same solution space. Computing the kernel of a linear application between two spaces of dimension $\mathcal{O}(n^2)$ can be done in time $\mathcal{O}(n^{2\omega})$. □