Extracting Linearization Equations from Noisy Sources

Daniel Smith-Tone^{1,2}

¹Department of Mathematics, University of Louisville, Louisville, Kentucky, USA ²National Institute of Standards and Technology, Gaithersburg, Maryland, USA

daniel.smith@nist.gov

Abstract. This note was originally written under the name On the Security of HMFEv and was submitted to PQCrypto 2018. The author was informed by the referees of his oversight of an eprint work of the same name by Hashimoto, see eprint article /2017/689/, that completely breaks HMFEv, rendering the result on HMFEv obsolete. Still, the author feels that the technique used here is interesting and that, at least in principal, this method could contribute to future cryptanalysis. Thus, with a change of title indicating the direction in which this work is leading, we present the original work with all of its oversights intact and with minimal correction (only references fixed).

At PQCRYPTO 2017, a new multivariate digital signature based on Multi-HFE and utilizing the vinegar modifier was proposed. The vinegar modifier increases the Q-rank of the central map, preventing a direct application of the MinRank attack that defeated Multi-HFE. The authors were, therefore, confident enough to choose aggressive parameters for the Multi-HFE component of the central map (with vinegar variables fixed). Their analysis indicated that the security of the scheme depends on the sum of the number of variables k over the extension field and the number v of vinegar variables with the individual values being unimportant as long as they are not "too small." We analyze the consequences of this choice of parameters and derive some new attacks showing that the parameter v must be chosen with care.

Key words: Multivariate Cryptography, HMFEv, Q-rank

1 Introduction

Note: The attack presented on HMFEv is obsolete, due to the attack by Hashimoto in [1]. The sections relevant to linearization equation extraction are Sections 5 (an original attack on multi-HFE) and 6 (filtering out noise). What follows is the original introduction.

Contributions to this work by NIST, an agency of the US government, are not subject to US copyright.

We are currently engaged in a massive international project to secure information and communication from adversaries with access to large scale quantum computers. Since Shor's algorithm broke public key cryptography in this paradigm, see [2], we have come a long way developing the mathematics of postquantum cryptography. The science is now sufficiently advanced for us to make educated decisions in how to move forward.

The National Institute of Standards and Technology (NIST) has begun evaluating submissions for post-quantum standards with the primary task of securing the internet in the coming quantum age. NIST's call for proposals, see [3], outlines the requirements of these technologies and illustrates the criteria by which they are evaluated. The principal prerequisite of any submission is to achieve certain security levels against quantum adversaries.

Multivariate public key cryptography (MPKC) provides a platform for potentially achieving these security levels. Multivariate cryptosystems rely on two known NP-complete problems for their hardness. The first is the MQ-problem: the problem of solving systems of nonlinear multivariate equations over a field. The second is the morphism of polynomials (MP) problem: the problem of determining whether there is a morphism between two polynomial systems. While typically multivariate cryptosystems lack a complexity theoretic reduction to one of these problems, there is a small collection of cryptanalytic techniques that often can be addressed specifically to derive security results.

In particular, MPKC has produced a few digital signature schemes that have withstood the test of time. Variations on the ideas of HFEv- and UOV, see[4, 5] have been around since the late 1990s without suffering any devestating attacks. PFLASH, see [6], which appeared to many to be weak, has now survived a decade and has fairly strong security arguments, see [7, 8].

Of course, we should not forget the other face of MPKC signatures. Oil-and-Vinegar (OV), SFLASH and Square, see [9–11], to name a few, were soundly defeated in [12–14]. Yet sometimes out of the ashes rises a new and more powerful scheme. The idea for UOV came from the attack on OV, and PFLASH is the progeny of SFLASH.

In this manuscript, we analyze a possible such phoenix. Multi-HFE, first proposed in [15], was completely broken in [16] by a clever MinRank attack exploiting the extremely low Q-rank of the central map of multi-HFE. The idea was breathed new life recently at PQCRYPTO 2017 in [17] where the idea of using the vinegar modifier on multi-HFE as a patch for the low Q-rank was proposed.

This new scheme, named HMFEv, is purported to have its security dependent upon the sum of two values: k, the number of variables and equations over the extension field defining the multi-HFE central map; and, v, the number of vinegar variables added to the central map. The authors in [17] claim that the exact values of k and v are not important as long as the sum is large and neither are "too small." In particular, they claim that $k, v \ge 2$, along with a large sum, suffices for securing the scheme from algebraic attacks. We offer a more precise justification for these claims by developing an explicit attack to filter out the vinegar variables when too few are included in the construction. The attack is statistical, bootstrapping an original attack on multi-HFE to form a distinguisher that successfully discerns whether a map is random or of multi-HFE shape. The attack depends on a disparity in the distribution of cubic forms generated from HMFEv instances with differing numbers of vinegar variables added. As the number of vinegar variables is increased, the distance between the distributions is decreased so that the addition of sufficiently many vinegar variables renders the attack impotent, thus demonstrating the need for a large number of vinegar variables.

The paper is organized as follows. In the next section, we describe the multi-HFE and HMFEv constructions. The following section describes Q-rank, an essential notion for understanding modern multivariate cryptography. In section 4, we review the previous cryptanalyses of multi-HFE and HMFEv. The subsequent section contains an original cryptanalysis of multi-HFE. Then, in Section 6, we extend this method into an attack filtering out the vinegar variables from HMFEv. Finally, we conclude, noting the affect these results have on parameter selection for HMFEv.

2 HFE Variants

Multivariate public key schemes can be broadly categorized as either "small field" or "big field" schemes. Small field schemes rely on the structure of a single field for their construction whereas the big field schemes rely on the multplicative structure of a hidden extension field. Given an extension \mathbb{E} of $\mathbb{F} = GF(q)$ of degree l, one can see that any monomial in $\mathbb{E}[X]$ of the form $X^{q^a+q^b}$ is the product of two Frobenius automorphisms, that is the product of two \mathbb{F} -linear functions. Therefore, this monomial can be written as a vector of quadratic functions over \mathbb{F} ; hence, we call such a monomial \mathbb{F} -quadratic. Big field multivariate schemes are based on easily invertible \mathbb{F} -quadratic maps from \mathbb{E} to \mathbb{E} with the structure hidden by an isomorphism of polynomials.

Definition 1 Two vector-valued multivariate polynomials F and G are said to be isomorphic if there exist two affine maps T, U such that $G = T \circ F \circ U$.

The following diagram summarizes the above discussion in the case of multi-HFE. One thing to note is that a multivariate polynomial ring over the extension is used instead of an univariate polynomial ring.



2.1 Multi-HFE

The HMFEv digital signature scheme of [17] is based on the multi-HFE primitive originally specified in [15]. Recalling the construction of multi-HFE, we choose a finite field \mathbb{F} , a degree ℓ extension \mathbb{E} , and an integer k. Setting $n = k\ell$ as the number of variables, one constructs n polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ as follows. Select an \mathbb{F} -vector space isomorphism $\phi : \mathbb{F}^{\ell} \to \mathbb{E}$ and two affine isomorphisms $T, U : \mathbb{F}^{k\ell} \to \mathbb{F}^{k\ell}$. Select the quadratic map $f = (f_1, \ldots, f_k)$ where $f_i(X_1, \ldots, X_k) \in \mathbb{E}[X]$ is defined by

$$f_i(X) = \sum_{1 \le r, s \le k} \alpha_{i,r,s} X_r X_s + \sum_{1 \le r \le k} \beta_{i,r} X_r + \gamma_i,$$

for $1 \leq i \leq k$. One then composes these maps producing the public key

$$P(x_1,\ldots,x_n) = T \circ (\phi^{-1})^k \circ f \circ (\phi)^k \circ U,$$

where $(\phi)^k = \phi \times \phi \times \cdots \times \phi$ with k coordinates.

A signature is the preimage of a certificate and so verification is accomplished by evaluating the public key at the signature value. The signature is generated by inverting each of the maps. The inversion of f is accomplished by generating a univariate polynomial, for example with a Gröbner basis algorithm with an elimination ordering, solving for a single variable and then repeating.

2.2 HMFEv

The modification of multi-HFE producing HMFEv is to add v additional variables and augment the definitions of $\beta_{i,r}$ and γ_i . Specifically, we let $U : \mathbb{F}^{k\ell+v} \to \mathbb{F}^{k\ell+v}$ and $T : \mathbb{F}^{k\ell} \to \mathbb{F}^{k\ell}$ be affine isomorphisms and define the quadratic map $f = (f_1, \ldots, f_k)$ by

$$f_i(X) = \sum_{1 \le r, s \le k} \alpha_{i,r,s} X_r X_s + \sum_{1 \le r \le k} \beta_{i,r}(x_{n+1}, \dots, x_{n+v}) X_r + \gamma_i(x_{n+1}, \dots, x_{n+v}),$$

for $1 \leq i \leq k$ where $\beta_{i,r} : \mathbb{F}^v \to \mathbb{E}$ and $\gamma_i : \mathbb{F}^v \to \mathbb{E}$ are linear forms in the vinegar variables x_{n+1}, \ldots, x_{n+v} . The public key is given by

$$P(x_1, \dots, x_{n+v}) = T \circ (\phi^{-1})^k \circ f \circ [(\phi)^k \times id_v] \circ U.$$

A signature is the preimage of a certificate and so verification is accomplished by evaluating the public key at the signature value. Signature generation is accomplished by randomly selecting values for the vinegar variables, which collapses f into the central map of a multi-HFE scheme. Then one inverts P restricted to this choice as in multi-HFE.

3 Q-Rank

As with all of the schemes in the HFE lineage, Q-rank plays an important role in the cryptanalysis of multi-HFE. Adapting the definition to multivariate extension field maps we may write the following definition.

Definition 2 The Q-rank of any quadratic map $f(\overline{x})$ on $\mathbb{F}_q^{k\ell}$ with respect to the degree ℓ extension \mathbb{E} is the rank of the quadratic form $(\phi^{-1})^k \circ f \circ (\phi)^k$ in $\mathbb{K}[X_0, \ldots, X_{k\ell-1}]$ via the identification $X_{\ell(n-1)+i} = \phi(\pi_n(\overline{x}))^{q^i}$, where $1 \leq n \leq k, 0 \leq i < \ell$ and π_n is the projection on to nth group of ℓ coordinates of \overline{x} .

Note that in the case of multi-HFE, the total degree of the central map over \mathbb{E} is two. Therefore, the Q-rank is bounded by k in all instances.

It is also important to note that although Q-rank is not preserved by isomorphisms of polynomials, the min-Q-rank in the linear span of f is preserved by such isomorphisms. This quantity is what is relevant for cryptography, and this is the property that has led to the attacks on multi-HFE.

4 Previous Cryptanalysis of multi-HFE and HMFEv

Being derived from multi-HFE, which is well known to have been broken, we review the security analysis of HMFEv and the cryptanalyses of muli-HFE. We offer an original, but trivial, extension to the security analysis of [17] which fits well in this section.

4.1 Cryptanalyses of Multi-HFE

Multi-HFE has been cryptanalyzed in a couple of related ways. In [16], the low Q-rank property described in Section 3 is exploited. Specifically, one may construct the rank $k\ell$ representation $\Phi : \mathbb{E}^k \to \mathbb{A}$ defined by

$$\Phi(\alpha,\beta,\ldots,\gamma) = (\alpha,\alpha^q,\ldots,\alpha^{\ell-1},\beta,\beta^q,\ldots,\beta^{\ell-1},\ldots,\gamma,\gamma^q,\ldots,\gamma^{\ell-1}).$$

Since each component of the central map $f : \mathbb{E}^k \to \mathbb{E}^k$ is of total degree two, represented as a quadratic form over \mathbb{A} it can involve only k coordinates, that is, the coordinates of α , β ,..., γ above. Therefore, each coordinate of the central map has Q-rank at most k.

Since multi-HFE is typically presented with k = 2 or k = 3, which means that the Q-rank of the central map is at most 3, the scheme is quite vulnerable to a MinRank attack via minors modeling, which is exactly what was efficiently done in [16]. To perform the attack, the sum of the product of variables t_i and the matrix representations of the public quadratic forms is constructed. By the Q-rank property, this matrix has rank at most 3. So by collecting all of the 4×4 minors, one generates an ideal whose Gröbner basis can be computed over \mathbb{F} and whose variety is then computed over \mathbb{E} .

In [18], another attack for odd characteristic instances of multi-HFE is presented. The authors describe the attack as a diagonalization approach. The technique can be described as a differential invariant attack that takes advantage of the fact that the multi-HFE central map has total degree two over the extension.

4.2 Previous Security Analysis of HMFEv

In [17], a preliminary analysis of HMFEv is presented. The authors consider the two principal attacks that seem relevant to the new scheme: the minrank attack and the direct algebraic attack.

For the MinRank attack, they note that the as long as $v \leq \ell$ the vinegar variables can be modelled by another variable over the extension field, where it is easy to show that the Q-rank of the central map is bounded by k + v. Experiments support the claim that this bound is tight, so they conclude that the complexity of the MinRank attack on HMFEv is $\mathcal{O}(\ell^{(k+v+1)\omega})$.

We can verify this claim analytically for all v in the following manner. For simplicity we consider the odd characteristic case. The argument is similar for characteristic two.

Proposition 1 The min-Q-rank of an HMFEv public key with parameters ℓ , k and v is k + v.

Proof. Let $\phi : \mathbb{F}^{\ell} \to \mathbb{E}$ be a vector space isomorphism. Choose a representation $\psi : \mathbb{E}^k \to \mathbb{A}$ defined by $\psi(X_1, X_2, \dots, X_k) = (X_1, X_1^q, \dots, X_1^{q^{\ell-1}}, X_2, \dots, X_k^{q^{\ell}})$. We then construct the vector space isomorphism $\Phi : \mathbb{F}^{k\ell+v} \to \mathbb{A} \times \mathbb{F}^v$ defined by $\Phi = (\psi \times id_v) \circ (\phi^k \times id_v)$.

We may now express the coordinates of the central map over \mathbb{E} as quadratic forms over $\mathbb{A} \times \mathbb{F}^v$. We observe that, due to the degree bound of two in the multi-HFE component, each coordinate f_i of the central map f satisfies $f_i = Q_i \circ \Phi$, where Q_i is a quadratic form on $\mathbb{A} \times \mathbb{F}_v$ with the following shape

where the α_{irs} coefficients represent HFE monomials $\alpha_{irs}X_rX_s$, the β_{irs} coefficients represent mixing terms $\beta_{irs}X_rx_{n+s}$, and the γ_{irs} coefficients represent the

quadratics $\gamma_{irs} x_{n+r} x_{n+s}$. Thus the Q-rank of f, which is bounded by the Q-rank of f_i for $1 \le i \le k$ is bounded by k + v. It is easy to see that in probability it is exactly k + v.

Addressing the complexity of the algebraic attack on HMFEv, the authors assume the tightness of the bounds given in both [17, Theorem 3] and in [19, Theorem 3.1] to conclude that the degree of regularity of the HMFEv system is

$$d_{reg} \leq \begin{cases} (q-1) \left\lfloor \frac{k+v}{2} \right\rfloor + 2 \text{ if } q \text{ is even, and} \\ \frac{q-1}{2}(k+v) + 2 \text{ otherwise.} \end{cases}$$

Again, this claim is supported by experiments showing that the above bound is fairly tight. They noted specifically that the choice of k and v were irrelevant as long as they were not too small, indicating in [20] that $k, v \ge 2$ suffices. Using these estimates they conclude the complexity of the direct attack is $\mathcal{O}\left(\binom{n+d_{reg}}{d_{reg}}^{2}\binom{n}{2}\right)$.

5 A New Attack on Multi-HFE

Multi-HFE has a couple of successful cryptanalyses as mentioned in the previous section. In both [16] and [18], a full key recovery attack is developed. We now introduce a less sophisticated attack focusing on the choice of central map of HMFEv which produces a kth-root speed-up in preimage search. As we will soon see, this technique allows information to be filtered through the vinegar modifier when an insufficient number of vinegar variables is used.

We consider the characteristic two case with k = 3. Since any multi-HFE instance over the field \mathbb{F}_q where q is even can be equivalently defined over GF(2), we consider the case in which q = 2. Let \mathbb{E} be the degree ℓ extension over which the central multi-HFE map of HMFEv is given by

$$Y_{1} = X_{1}X_{2} + \alpha_{1,1}X_{1} + \alpha_{1,2}X_{2} + \alpha_{1,3}X_{3} + \alpha_{1,4}$$

$$Y_{2} = X_{2}X_{3} + \alpha_{2,1}X_{1} + \alpha_{2,2}X_{2} + \alpha_{2,3}X_{3} + \alpha_{2,4}$$

$$Y_{3} = X_{1}X_{3} + \alpha_{3,1}X_{1} + \alpha_{3,2}X_{2} + \alpha_{3,3}X_{3} + \alpha_{3,4}$$
(1)

Consider the graded ring $(\mathcal{A}_d) = \mathbb{E}[X_1, X_2, X_3]/I$, graded by total degree, where I is the ideal generated by the homogeneous quadratic components of the above three polynomials. It is clear that $\dim_{\mathbb{E}}(\mathcal{A}_d) = 3$ if d > 0 and is one if d = 0. Therefore the Hilbert Series of I is

$$HS_I(t) = \frac{1+2t}{1-t},$$

and the Hilbert regularity of I is one. Thus, the ideal I already has nontrivial syzygies on its generating set and there exists an \mathbb{E} -bilinear relation between the homogeneous polynomials \widehat{Y}_i and the variables X_i .

It is easy to see that the polynomials Y_i in Equation (1) inherit this relation. To be explicit, we compute

$$X_{3}Y_{1} + X_{1}Y_{2} = \alpha_{1,1}X_{1}X_{3} + \alpha_{1,2}X_{2}X_{3} + \alpha_{1,3}X_{3}^{2} + \alpha_{2,1}X_{1}^{2} + \alpha_{2,2}X_{1}X_{2} + \alpha_{2,3}X_{1}X_{3} + \alpha_{1,4}X_{3} + \alpha_{2,4}X_{1} = \alpha_{1,1}Y_{3} + \alpha_{1,2}Y_{2} + \alpha_{1,3}L(X_{3}) + \alpha_{2,1}L(X_{1}) + \alpha_{2,2}Y_{1} + \alpha_{2,3}Y_{3} + \alpha_{1,2}(\alpha_{2,1}X_{1} + \alpha_{2,2}X_{2} + \alpha_{2,3}X_{3}) + \alpha_{2,1}(\alpha_{1,1}X_{1} + \alpha_{1,2}X_{2} + \alpha_{1,3}X_{3}) + \alpha_{2,3}(\alpha_{3,1}X_{1} + \alpha_{3,2}X_{2} + \alpha_{3,3}X_{3}) + \alpha_{1,4}X_{3} + \alpha_{2,4}X_{1},$$

$$(2)$$

which is GF(2)-affine in both X and Y. Another linearly independent such relation can be derived from $X_3Y_1 + X_2Y_3$; however, the remaining relation $X_1Y_2 + X_2Y_3$ is the sum of the first two relations.

It is easy to extend this analysis to any $k \leq 5$, see Appendix A. For k > 5 it is still possible to recover relations between X and Y linear in X; however, they must in general be nonlinear in Y.

Passing these relations to GF(2) and generalizing to k, we obtain $(k-1)\ell$ linearly independent relations linear in both x and y. Composing with the affine transformations U and T^{-1} , we obtain $(k-1)\ell$ linearization equations (or higher order linearization equations if k > 5) on the multi-HFE instance. Thus we obtain a decryption oracle with runtime $\mathcal{O}(2^{\ell})$, performing a preimage search on a space of one kth the dimension of the signature space.

6 Distilling Vinegar

The existence of the linearization equations of the previous section function as a criterion for the image of a linear projection being orthogonal to the vinegar subspace. From this idea, we build a distinguisher acting on projected HMFEv keys of the form $P \circ \pi$, able to determine whether a subspace of the vinegar space is orthogonal to $Im(\pi)$. We then bootstrap this technique to a key recovery. For the simplicity of avoiding the higher dimensional tensors necessary in the analysis of the higher order linearization equations case, we restrict to the case that $k \leq 5$ guaranteeing $(k-1)\ell$ linearization equations, noting that the general theory works analogously.

6.1 Filtering Vinegar Variables

Regarding the linearization equations of Section 5 as cubic forms in X, we induce projections on these cubic forms by projecting to a subspace of the input space. If a projection π is orthogonal to the vinegar subspace, the rank of these cubic forms can be as much as $(k-1)\ell$ smaller than for random projections if rank (π) is sufficiently large. To make this clear, we review the theory of linear embeddings of homogeneous forms on a vector space.

Theorem 1 Let V be an n-dimensional \mathbb{F} -vector space. Let $\mathbb{F}[X]_d$ be the space of homogeneous polynomials of degree d on V and let $\pi : V \to V$ be a rank r projection. The rank of the linear map $T_{\pi} : \mathbb{F}[X]_d \to \mathbb{F}[X]_d$ defined by $T_{\pi}(f) = f \circ \pi$ is $\binom{r+d-1}{d}$ if $|\mathbb{F}| \geq d$.

Proof. We show that $T_{\pi}(f)$ is in the span of $\binom{r+d-1}{d}$ *d*-tensors. In particular, we may without loss of generality assume that π projects onto the first r standard basis vectors of V, so that $T_{\pi}(f)$ involves only r variables. Since there are exactly $\binom{r+d-1}{d}$ distinct degree d monomials in r variables, the rank of $T_{\pi}(f)$ is $\binom{r+d-1}{d}$.

We note that when $|\mathbb{F}| < d$ that some of the monomials in Theorem 1 are equivalent to smaller degree monomials, and thus such homogeneous forms of degree d are degenerate and the rank in this sense is smaller. A particular case to note is that of cubic forms over GF(2). The number of such monomials in this case with distinct values is $\binom{r}{3} + \binom{r}{2} + \binom{r}{1}$ due to the fact that $x_i^2 = x_i$. Thus, for GF(2), we obtain a rank bound of $\binom{r}{3} + \binom{r}{2} + \binom{r}{1} = \frac{r^3 + 5r}{6}$.

Corollary 1 Let V be an n dimensional \mathbb{F} -vector space. Let $g: V \to V$ be a quadratic map. For any rank r projection $\pi: V \to V$, the rank of the linear map $T_{g,\pi}: \mathbb{F}[X]_2 \to \mathbb{F}[X]_3$ defined by $T_{g,\pi}(f)(x) = \langle \pi(x)A, g(\pi(x)) \rangle$, where $f(x) = \langle xA, x \rangle$ is the inner product representation of f, is at most $\min\{\frac{r^3+5r}{6}, rn\}$ if $|\mathbb{F}| = 2$ and at most $\min\{\binom{r+2}{3}, rn\}$ otherwise.

Proof. First note that since π is composed with the matrix representing f, there are only actually rn degrees of freedom in choosing f. We therefore focus on establishing the bound when r is sufficiently low.

Clearly we may write $T_{g,\pi} = T_{\pi} \circ T_g$ where $T_g : \mathbb{F}[X]_2 \to \mathbb{F}[X]_3$ is defined by $T_g(f)(x) = \langle xA, g(x) \rangle$. Since the rank of T_{π} has the appropriate bound by Theorem 1, the only thing to show is that T_g is linear. That $T_g(ap) = aT_g(p)$ is obvious. Let $p, q \in \mathbb{F}[X]_2$ and let A_p and A_q be the matrix representations of pand q, respectively. First,

$$(p+q)(x) = p(x) + q(x) = \langle xA_p, x \rangle + \langle xA_q, x \rangle = \langle x(A_p + A_q), x \rangle.$$

Then we obtain

$$T_g(p+q) = \langle x(A_p + A_q), g(x) \rangle = \langle xA_p, g(x) \rangle + \langle xA_q, g(x) \rangle = T_g(p) + T_g(q).$$

The conclusion of Corollary 1 is weaker than that of Theorem 1 precisely because of the possible degeneracy of g. The failure of T_g to be full rank implies a linear relation among the coordinates of X and the polynomial g(X), thus T_g is of full rank generically in characteristic zero (since the hypersurface satisfying these linear relations has measure zero) and with high probability in finite fields. On the other hand, if g has linearization equations, the rank of T_g is diminished.

Theorem 2 Let P be a public key of HMFE with parameters q = 2, k, ℓ and v. Let π be the rank $r \leq n = k\ell$ projection orthogonal to the vinegar subspace. Then

$$Rank(T_{P,\pi}) \le \min\left\{\frac{r^3 + 5r}{6}, rn - (k-1)\ell\right\}.$$

Proof. We may calculate the rank of $T_{P,\pi}$ directly by specifying a quadratic form $f \in \mathbb{F}[X]_2$ by its matrix representation $A = (a_{ij})$ and directly computing the nullity of $T_{P,\pi}$. Recall that due to π we may consider A to have an $r \times n$ block of possibly nonzero values. Notice that by the previous section, there are $(k-1)\ell$ linearization equations; therefore, whenever $rn - \frac{r^3 + 5r}{6} < (k-1)\ell$, we know that the kernel of $T_{P,\pi}$ is at least $\left((k-1)\ell + \frac{r^3 + 5r}{6} - rn\right)$ -dimensional. Thus the rank of $T_{P,\pi}$ is at most $\frac{r^3 + 5r}{6}$ when this quantity is less than $rn - (k-1)\ell$, and is at most $rn - (k-1)\ell$ otherwise.

Restricting the codomain of $T_{P,\pi}$ to the image of T_{π} , we can see that if the rank of $T_{P,\pi}$ is less than min $\left\{\frac{r^3+5r}{6}, rn-(k-1)\ell\right\}$, then there is a nontrivial cokernel, which is to say that there are additional linearization equations. For random functions one should expect this event to occur with low probability. Unsurprisingly, there is a distinction in the behavior of a multi-HFE primitive and a random function under the vinegar modification in this respect. Furthermore, there is a noticable relationship between the rank of $T_{P,\pi}$ and the dimension of the intersection of the dual of the cokernel of π and the vinegar subspace, V_{vin} , as illustrated in Table 1.

[r	1	2	3	4	5	6	7	8	9	10	11	12
	α	1	3	7	14	25	41	63	92	108	120	132	144
d	0				13.99	24.95	40.95	62.80	91.64	108	120	132	144
	1			7	13.96	24.97	40.89	62.82	91.5	108	120	132	144
	2		2.99	7	13.97	24.91	40.78	62.52	91.02	108	120	132	144
	3	1	2.99	6.98	13.96	24.86	40.51	62.3	89.81	107.44	119.83	131.87	143.88
	4	1	2.99	6.9	13.85	24.63	40.32	61.2	84.91	99.76	111.99	124	136
ľ	β	1	3	7	14	25	41	63	88	100	112	124	136

Table 1. Average rank of $T_{P,\pi}$ over 100 trial runs where the rank of π is r and dim(coKer(π)^{*} $\cap V_{vin}$) = d for n = 12 and v = 4. For comparison, we also include the values $\alpha = \min\{\frac{r^3+5r}{6}, rn\}$ and $\beta = \min\{\frac{r^3+5r}{6}, rn - (k-1)\ell\}$.

Notice, in particular, that the data in columns r = 7 and 8 of Table 1 exhibit a larger range of values among rows d = 0 through d = 3 than the remaining columns. The reason is that for these values of r, the image of T_{π} is approaching rn-dimensional, and since P has, in this case, $8 = (k-1)\ell$ linearization equations, having fewer possible values for the vinegar variables reduces the degrees of freedom on $T_{P,\pi}$. This transition appears to grow sharper as $\dim(coKer(\pi)^* \cap V_{vin})$ approaches $\dim(V_{vin})$. Apparently, the symmetries in the multi-HFE structure skew the distribution of $\operatorname{Rank}(T_{P,\pi})$ away from that of $\operatorname{Rank}(T_{f,\pi})$ for a random function f and the degrees of freedom of the values of the vinegar variables in the image of π is a metric for passing between these two extremes. The disparity is more apparent with fewer vinegar variables, see Table 2.

Remark 1 The variance of the data increases near $r = \sqrt{6n-5}$; however, the variance is still small, is even smaller with larger q and is extremely small with d = 0 in all cases.

	r	1	2	3	4	5	6	7	8	9	10	11	12
	α	1	3	7	14	25	41	63	92	108	120	132	144
d	0		2.99	6.95	13.92	24.92	40.86	62.49	91.07	108	120	132	144
	1	1	3	6.97	3.9	24.72	40.63	62.25	89.7	107.21	119.55	131.57	143.67
	2	0.98	2.97	6.95	13.84	24.73	40.04	6.36	84.95	99.72	112	124	136
	β	1	3	7	14	25	41	63	88	100	112	124	136

Table 2. Average rank of $T_{P,\pi}$ over 100 trial runs where the rank of π is r and dim(coKer(π)* $\cap V_{vin}$) = d for n = 12 and v = 2. For comparison, we also include the values $\alpha = \min\{\frac{r^3+5r}{6}, rn\}$ and $\beta = \min\{\frac{r^3+5r}{6}, rn - (k-1)\ell\}$.

6.2 Key Recovery

The method for turning this statistical anomaly into a key recovery is as follows. First one randomly generates a large number of rank $r \approx \sqrt{6n-5}$ projections π and selects a cutoff rank R for $T_{P,\pi}$. For any π , if the rank of $T_{P,\pi}$ is bounded by R, π is placed in a database. Next one chooses a number s of $\operatorname{coKer}(\pi)^*$ to intersect in the hopes of filtering out a vector \overline{v} in V_{vin} . Success is measurable immediately by estimating the distribution of $\operatorname{Rank}(T_{P',\pi})$ near $r = \sqrt{6n-5}$, where P' is P composed with the projection onto the orthogonal complement of \overline{v} . Finally, one repeats this process, which becomes easier as v diminishes.

For this method to be effective, one must fine tune R and s to minimize the number of low rank $T_{P,\pi}$ required to be computed. To estimate the optimal s we note that the probability of a vector lying in the intersection of s subspaces of dimension r in a vector space of dimension n + v is roughly $q^{sr-(s-1)(n+v)}$. Similarly, under the assumption that each of the k subspaces has a d-dimensional intersection with a fixed v-dimensional subspace, the probability that a vector in this subspace lies in the intersection is roughly $q^{sd-(s-1)v}$. When these probabilities are equal, i.e. when $s \approx \frac{n}{n+d-r}$, essentially all of the vectors in the intersection should lie in the fixed v-dimensional subspace. Specifically, when n is sufficiently large ($n \geq 19$ for d = 1 and $r \approx \sqrt{6n-5}$, for example), s = 2 suffices.

Under the assumption that an R can be found such that the fraction of π satisfying dim $(coKer(\pi)^* \cap V_{vin}) = d$ among all π for which the rank of $T_{P,\pi}$ is bounded by R is significant, the recovery of a vector $\overline{v} \in V_{vin}$ requires approximately $q^{d(r-s)+v(s-1)}$ calculations of the rank of $T_{P,\pi}$ plus some additional linear algebra steps. Thus the complexity of recovering the vinegar subspace is $\mathcal{O}\left((rn)^{\omega}q^{d(r-s)+v(s-1)}\right)$. Since $r \approx \sqrt{6n-5}$, this attack is subexponential in n for any fixed v. For the special case of v = 2, addressing the claim in [17], we have d = 1 and s = 2 for sufficiently large n and this formula simplifies to $\mathcal{O}\left((rn)^{\omega}q^{r+1}\right) = \mathcal{O}\left(n^{3\omega/2}q^{\sqrt{6n-5}+1}\right)$.

Experiments show that this method is effective in practice on small scale schemes. In the case of q = 2, l = 4, k = 3 and v = 2, only four rank 9 projections π satisfying rank $(T_{P,\pi}) \leq R = 107$ were needed to find a vector in V_{vin} . On average one in 384 projections π satisfied this property, supporting the above complexity estimate.

It is interesting to note that the cutoff phenomenon in the rank of $T_{P,\pi}$, though present, is not extremely sharp, at least when $r \approx \sqrt{6n-5}$; therefore, for small values of v a value of R can be found to make the attack effective. Still, for v > 4, we estimate that the cutoff is sharp enough that one would require d > 1 to find a suitable value of R; however, this would render the attack worse than brute force for essentially any parameters, due to the large q of HMFEv.

7 Other Techniques

In [21], new statistical methods for attacking HFEv- are advanced. In principle, the attacks are applicable to HMFEv as well; however the complexity analysis is not the same. Similar to the attack presented in Section 6, these techniques incorporate projections and the calculation of invariants to bootstrap a distinguisher to a key recovery attack, though the approaches are different.

The first approach combines projection with the MinRank method. There are two possible variants of this technique: Project-then-MinRank and MinRankthen-Project. The Project-then-MinRank strategy works by noting that there is a distinction in the Q-rank of the central map under a projection reducing the dimension of the vinegar subspace versus a projection that is full rank when restricted to the vinegar subspace. Both attacks, however, still require the Min-Rank step to be executed. In the context of HMFEv, as long as k + v is large, even a random reduction in the rank of two or three, which occurs with very low probability, will not reduce the complexity sufficiently to risk the integrity of the scheme. Due to the relatively large size of q in the proposed parameters of HMFEv, the MinRank-then-Project method seems to be the more efficient. In this case the complexity of the MinRank-then-Project approach is dominated by the cost of the MinRank step, which has complexity $\mathcal{O}\left(\ell^{(k+v+1)\omega}\right)$, as predicted in Section 4.2.

The second approach attempts to leverage the low Q-rank of HFE by way of the degree of regularity. Specifically, the attack proceeds by projecting an HFEvscheme to a subspace of a dimension at which random quadratic systems have degree of regularity d but random systems of smaller dimension have degree of regularity d-1. Because of the low Q-rank property of HFE, there is a higher probability that projected systems from an HFE instance will have a lower degree of regularity. The idea is that if the projection eliminates a vinegar variable it may be detectable in the degree of regularity of the projected scheme.

The complexity of this attack in application to HMFEv is also dependent on the sum k + v. The projection must have a large enough corank to have a sufficiently high probability of achieving a reduction in the degree of regularity and when a vector in the vinegar subspace is found, there is no additional information about the basis of the vinegar subspace revealed. Thus the attack must be repeated with the slight advantage of one additional equation specifying a one-dimensional subspace of the subsequent kernel.

The complexity of the entire attack is approximately $\mathcal{O}\left(q^{n-t}\binom{n+v+d}{d}^2\binom{n+v}{2}\right)$, where t is the co-dimension of projection optimal for distinguishing and d is the degree of regularity at which distinguishing occurs. The quantity t is a decreasing function of k + v. This fact along with the size of q make this new technique infeasible for the parameters suggested in [17].

8 Experiments

We ran a series of experiments with Magma¹, see [22], on a 3.2 GHz Intel[®] XeonTM CPU, testing the first step of the attack, recovering a vector in the vinegar subspace, for a variety of values of ℓ with v = 2 or v = 3 and k = 3 or k = 4. Tables 3 and 4 summarize some of our results in the v = 2 and v = 3 cases, respectively. The data support our complexity estimate of $\mathcal{O}((rn)^{\omega}q^{d(r-s)+v(s-1)})$.

	v = 2							
l	3	4	5	6	3	4		
k	3	3	3	3	4	4		
$\lfloor \sqrt{6n-5} \rfloor + 1$	8	9	10	11	9	10		
r	8	9	10	11	9	9		
Т	10	10	5	5	5	3		
avg. time	145.2s	1155.4s	6131.3s	$28910.4 \mathrm{s}$	$13070.7 \mathrm{s}$	31570.6s		
min. time	73.1s	827.7s	4583.4s	14859.6s	2997.7s	22180.0s		

Table 3. Average time (in s) for T instances of the vinegar recovery attack of Section 6 on HMFEv($q = 2, k, \ell, v = 2$) for various values of ℓ and k.

There is an interesting artifact in the data that should be pointed out explicitly. For almost all of the parameters tested, the value of n is too small for our estimate of s to achieve an upper bound of 2. That is, the number of $\operatorname{coKer}(\pi)^*$

¹ Any mention of commercial products does not indicate endorsement by NIST

one must intersect in order to essentially guarantee that a nontrivial intersection will reveal a vector in V_{vin} is, for most of these low values of n, greater than 2. Therefore, for these small tests, it is likely that we will see nonempty intersections constructed from two projections revealing a vector not contained in V_{vin} . Still, we found it more efficient for these small scale experiments to choose a value of s = 2. In nearly every experiment, we found nonempty intersections which did not intersect the V_{vin} . In fact, aside from a couple of lucky instances far from the mean, there was exactly one set of parameters for which the nonempty intersections always were in V_{vin} . That test was for $k = \ell = 4$, for which our estimate for s is

$$s = \frac{n}{n+d-r} = \frac{16}{16+1-9} = 2$$

Thus, the tests behave exactly as predicted in the analysis of Section 6.

	v = 3				
l	3	4	5		
k	3	3	3		
$\lfloor \sqrt{6n-5} \rfloor + 1$	8	9	10		
r	7	8	9		
Т	10	10	3		
avg. time	623.6s	1443.2s	47921.1s		
min. time	440.7s	930.0	28803.2s		

Table 4. Average time (in s) for T instances of the vinegar recovery attack of Section 6 on HMFEv($q = 2, k = 3, \ell, v = 3$) for various values of ℓ .

9 Conclusion

We have demonstrated that the security of HMFEv is not symmetrically dependent upon the values k, the number of multi-HFE variables over the extension field, and v, the number of vinegar variables. Due to the extremely low Q-rank structure of multi-HFE, symmetries in the multi-HFE map can percolate through the noise added by the vinegar variables and affect the rank of tensors derived from the public key when there are very few vinegar variables. Thus HMFEv really requires a large number of vinegar variables.

HMFEv adds to the interesting story of the interplay among Q-rank, the minus modifier, projection and the vinegar modifier. The new attack on the multi-HFE primitive shows that multi-HFE is actually a degenerate case of HFE, similar to C^* ; however, the vinegar modifier is still strong enough to secure the scheme. The lesson seems to be that vinegar is very good.

References

- Hashimoto, Y.: On the security of hmfev. IACR Cryptology ePrint Archive 2017 (2017) 689
- 2. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Sci. Stat. Comp. 26, 1484 (1997)
- 3. Group, C.T.: Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. NIST CSRC (2016) http://csrc.nist.gov/groups/ST/post-quantum-crypto/documents/call-for-proposals-final-dec-2016.pdf.
- Patarin, J.: Hidden Fields Equations (HFE) and Isomorphisms of Polynomials (IP): Two New Families of Asymmetric Algorithms. In: EUROCRYPT. (1996) 33–48
- Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. EUROCRYPT 1999. LNCS 1592 (1999) 206–222
- Ding, J., Dubois, V., Yang, B.Y., Chen, C.H.O., Cheng, C.M.: Could SFLASH be Repaired? In Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I., eds.: ICALP (2). Volume 5126 of Lecture Notes in Computer Science., Springer (2008) 691–701
- Chen, M.S., Yang, B.Y., Smith-Tone, D.: Pflash secure asymmetric signatures on smart cards. Lightweight Cryptography Workshop 2015 (2015) http://csrc.nist.gov/groups/ST/lwc-workshop2015/papers/session3-smithtone-paper.pdf.
- 8. Cartor, R., Smith-Tone, D.: An updated security analysis of PFLASH. [23] 241-254
- 9. Patarin, J.: The oil and vinegar algorithm for signatures. Presented at the Dagstuhl Workshop on Cryptography (1997)
- Patarin, J., Courtois, N., Goubin, L.: Flash, a fast multivariate signature algorithm. CT-RSA 2001, LNCS 2020 (2001) 297–307
- Clough, C., Baena, J., Ding, J., Yang, B.Y., Chen, M.S.: Square, a New Multivariate Encryption Scheme. In Fischlin, M., ed.: CT-RSA. Volume 5473 of Lecture Notes in Computer Science., Springer (2009) 252–264
- Shamir, A., Kipnis, A.: Cryptanalysis of the oil & vinegar signature scheme. CRYPTO 1998. LNCS 1462 (1998) 257–266
- Dubois, V., Fouque, P.A., Shamir, A., Stern, J.: Practical Cryptanalysis of SFLASH. In Menezes, A., ed.: CRYPTO. Volume 4622 of Lecture Notes in Computer Science., Springer (2007) 1–12
- Billet, O., Macario-Rat, G.: Cryptanalysis of the square cryptosystems. ASI-ACRYPT 2009, LNCS 5912 (2009) 451–486
- Chen, C.O., Chen, M., Ding, J., Werner, F., Yang, B.: Odd-char multivariate hidden field equations. IACR Cryptology ePrint Archive 2008 (2008) 543
- 16. Bettale, L., Faugère, J., Perret, L.: Cryptanalysis of HFE, multi-HFE and variants for odd and even characteristic. Des. Codes Cryptography **69** (2013) 1–52
- Petzoldt, A., Chen, M., Ding, J., Yang, B.: Hmfev an efficient multivariate signature scheme. [23] 205–223
- Hashimoto, Y.: Cryptanalysis of multi-hfe. IACR Cryptology ePrint Archive 2015 (2015) 1160
- Ding, J., Yang, B.Y.: Degree of regularity for hfev and hfev-. In Gaborit, P., ed.: PQCrypto. Volume 7932 of Lecture Notes in Computer Science., Springer (2013) 52–66

- 16 D. Smith-Tone
- Petzoldt, A., Chen, M.S., Ding, J., Yang, B.Y.: Hmfev an efficient multivariate signature scheme. Presentation - Post-Quantum Cryptography - 8th International Conference on Post-Quantum Cryptography, PQCrypto 2017, Utrecht, Netherlands, June 26-28, 2017 (2017) https://2017.pqcrypto.org/conference/slides/mqI/HMFEv.pdf.
- Ding, J., Perlner, R., Petzoldt, A., Smith-Tone, D.: Improved cryptanalysis of hfevvia projection. In Current Submission (2017)
- Bosma, W., Cannon, J., Playoust, C.: The Magma algebra system. I. The user language. J. Symbolic Comput. 24 (1997) 235–265 Computational algebra and number theory (London, 1993).
- Lange, T., Takagi, T., eds.: Post-Quantum Cryptography 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings. Volume 10346 of Lecture Notes in Computer Science., Springer (2017)

A Existence of Linearization Equations

Theorem 3 Let $3 \le k \le 5$ and for $i \in \{1, \ldots, k\}$ let

$$Y_{i} = X_{i}X_{i+1} + \alpha_{i,1}X_{1} + \dots + \alpha_{i,k}X_{k} + \alpha_{i,k+1},$$

be a multi-HFE central map over \mathbb{E} , where the indices are computed as least positive residues modulo k. Then there exists an \mathbb{E} -bilinear relation between Y_i and X_i .

Proof. Consider the variable X_i for $1 \le i \le k$. There are two nontrivial syzygies given by $LT(X_iY_{i+1}) - LT(X_{i+2}Y_i)$ and $LT(X_iY_{i-2}) - LT(X_{i-2}Y_{i-1})$ involving X_i . One can clearly see that summing over i we obtain exactly k syzygies since every one is counted exactly twice. The case of k = 3 is the exceptional case in which the span of these syzygies is less than k-dimensional; however, we have already seen that the result holds for k = 3.

The non-leading terms of X_iY_{i+1} are either linear in X, specifically, $\alpha_{ii}X_i^2$; linear in X and Y, i.e. of the form $X_iX_{i\pm 1} = Y_{(i\pm 1-1)/2} + \sum_{j\neq i\pm 1} \alpha_jX_j$; or quadratic in X, such as X_iX_s where $s \notin \{i-1, i, i+1\}$. There are exactly $\frac{k^2-3k}{2}$ quadratics X_iX_j with $j \notin \{i-1, i, i+1\}$. Each such term can be eliminated with linear combinations of $X_iY_{i+1} - X_{i+2}Y_i$ provided $\frac{k^2-3k}{2} \leq k$, which occurs if $k \leq 5$.

B Toy Example

We illustrate the attack on a very small scale example, showing the extraction of the vinegar subspace of the plaintext space. Specifically, we consider the case $HMFEv(q = 2, k = 3, \ell = 2, v = 2)$, and recover a transformation of the plaintext space and a multi-HFE instance that can be used to produce forgeries with the vinegar variables of the central map set to zero. We simplify the exposition by considering a homogeneous key.

B.1 The Public Key

We construct the degree ℓ extension $\mathbb{E} = \mathbb{F}_2(b)$ where $b^2 + b + 1 = 0$. We specify the canonical isomorphism $\phi : \mathbb{F}_2^2 \to \mathbb{E}$. We randomly fix the central map F, specifying its coordinates,

$$Y_{1} = X_{1}X_{2} + \phi(L_{11}(X_{V}))X_{1} + \phi(L_{12}(X_{V}))X_{2} + \phi(L_{13}(X_{V}))X_{3} + \phi(x_{7}^{2}, x_{7}x_{8})$$

$$Y_{2} = X_{2}X_{3} + \phi(L_{21}(X_{V}))X_{1} + \phi(L_{22}(X_{V}))X_{2} + \phi(L_{23}(X_{V}))X_{3} + \phi(x_{7}x_{8}, x_{7}^{2} + x_{8}^{2})$$

$$Y_{3} = X_{1}X_{3} + \phi(L_{31}(X_{V}))X_{1} + \phi(L_{32}(X_{V}))X_{2} + \phi(L_{33}(X_{V}))X_{3} + \phi(x_{8}^{2}, x_{7}x_{8})$$

where

$$\begin{split} &L_{11} = X_V \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}, \ &L_{12} = X_V \begin{bmatrix} 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \end{bmatrix}, \ &L_{13} = X_V \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 \end{bmatrix}, \\ &L_{21} = X_V \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \end{bmatrix}, \ &L_{22} = X_V \begin{bmatrix} 1 & 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \end{bmatrix}, \ &L_{23} = X_V \begin{bmatrix} 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 1 \end{bmatrix}, \\ &L_{31} = X_V \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 \end{bmatrix}, \ &L_{32} = X_V \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 \end{bmatrix}, \ &L_{33} = X_V \begin{bmatrix} 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \end{bmatrix}, \end{split}$$

and two invertible linear transformations T and U:

(U is chosen in this manner to make it easier for the reader to identify the vinegar subspace. Note that this choice (1) does not mix the multi-HFE and vinegar variables and (2) this fact is irrelevant for the attack because random projections are selected and typically all variables will be mixed.)

Composing $P = T \circ \phi^{-1} \circ F \circ (\phi \times id_V) \circ U$, we obtain the public key expressed here in polar form over \mathbb{F}_2 :

$\mathbf{P_0} =$	$\begin{bmatrix} 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 0 & 1 \end{bmatrix}, \mathbf{P}$	$\mathbf{P}_{1} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \end{bmatrix}, \mathbf{P}_{2}$	$\mathbf{a}_{2} = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix},$
$\mathbf{P_3} =$	$\begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}, \mathbf{P}$	$\mathbf{P}_{4} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 & 1 & 1 \end{bmatrix}, \mathbf{P}_{2}$	$5 = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 &$

B.2 Recovering the Vinegar Subspace

We choose random rank r = 5 projections π and compute the rank of

$$T_{P,\pi}(\overline{x}) = \pi(\overline{x}) \begin{bmatrix} t_{1,1} \cdots t_{1,6} \\ \vdots & \ddots & \vdots \\ t_{5,1} \cdots & t_{5,6} \end{bmatrix} \overline{y}^{\top},$$

where $P(\pi(\overline{x})) = \overline{y}$. Setting a rank bound of 19, we generate maps π satisfying rank $(T_{P,\pi}) \leq 19$. As we find solutions, π_i we compute $\operatorname{coKer}(\pi_i)^* \cap \operatorname{coKer}(\pi_i)^*$.

In this experiment, the first two projections π_1 and π_2 satisfying the rank bound are

	11110100		[10001000]
	$0\ 0\ 1\ 0\ 1\ 0\ 1\ 1$	and $\Pi_2 =$	01100100
$\Pi_1 =$	01011111		01111100
	0 0 0 1 0 0 1 1		11101000
	00000011		11000100

and the linear form $\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix}$ is in $\operatorname{coKer}(\pi_1)^* \cap \operatorname{coKer}(\pi_2)^*$. The reader recognizes that this linear form is identified with an element of V_{vin} .

We may now project the entire scheme onto the orthogonal complement of this vector and will have eliminated one vinegar variable from HMFEv. Equivalently, one may choose projections π from those containing this vector in coKer(π)^{*}. Because of the size of this example, we chose the latter option, though the former is more efficient, in general. Continuing in this manner, after six additional projections were collected, we obtained

$$\boldsymbol{\Pi}_8 = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix},$$

and both $[0\ 0\ 0\ 0\ 0\ 0\ 1]$ and $[0\ 0\ 0\ 0\ 0\ 1\ 0]$ lie in $\operatorname{coKer}(\pi_2)^* \cap \operatorname{coKer}(\pi_8)^*$. Thus the entire vinegar subspace has been found. At this point we project onto the orthogonal complement of this subspace and obtain the multi-HFE key:

At this point the scheme is readily broken via the methods of [16].