D-James: Ultra Short Multivariate Signatures

Jacques Patarin[1,2] and Alexandre Roullet[1,2]

[1] Thales DIS France SAS, Meudon, France
[2] Université Paris–Saclay, UVSQ, CNRS, Laboratoire de mathématiques de Versailles, 78000, Versailles, France
jacques.patarin@thalesgroup.com, alexandreroullet3@gmail.com

Abstract. Multivariate signature schemes are among the few post-quantum candidates capable of providing very short signatures, but designing secure constructions has proven challenging. HFE-based schemes such as GeMSS were compromised by algebraic MinRank attacks. This motivates the HFE^-_{IP} framework, which combines IP and minus modifiers to address these attacks. We introduce James and D-James, the latter achieving signatures of only 156 bits at the 128-bit classical security level and 348 bits at the 256-bit classical security level, among the shortest signatures reported for practical post-quantum public-key signature schemes, with estimated signing and verification costs comparable to those of GeMSS. The main technical contribution is the introduction of Dragon terms, which decouple the number of public equations from the hash output length, allowing the signature size to be reduced independently of the security parameter. We characterize the algebraic structure introduced by Dragon terms and show that it does not enable known MinRank attacks when combined with the HFE^-_{IP} countermeasures. We also show experimentally that the known differential attack does not appear to extend to the minus variant. We further present parameter sets over both binary and small non-binary finite fields. For small values of $q > 2$, the public-key size decreases by a factor of about 3 to 5 for $q \in \{4, 5\}$, and by up to a factor of 10 for larger q, at the cost of slightly longer signatures and slower signing.

Keywords: Multivariate cryptography · Post-quantum cryptography · Ultra-short signatures

Note added (October 8th 2026)

After the first version of this paper was made public, we found that the MinRank-S estimate used in Section 6.2, namely $d' \geq n/2$ for $r > 1$, taken from [9], is not correct. We have identified an adaptation of the MinRank-S attack of [23] to the IP modifier whose target rank is

$$d' = d + 1 + r,$$

regardless of the minus modifier and of the Dragon terms. The subsequent recovery of an equivalent key is polynomial-time. For the parameter sets of this

paper this is far below the values reported in Tables 7 and 8. These results are currently being verified and refined, and a complete analysis will be published in a forthcoming work. In the meantime, the MinRank-S column of Tables 7 and 8 should be considered invalid, and the proposed parameter sets should not be used.

1 Introduction

In cryptographic systems, signature size is a critical design criterion alongside computational efficiency. Digital signatures are often transmitted and stored in large quantities, for example in distributed systems, software authentication infrastructures, and resource-constrained devices. In these settings, reducing signature size directly improves communication overhead, storage requirements, and overall scalability. This consideration becomes particularly important in the post-quantum setting, where many secure signature schemes introduce significant size increases compared to classical alternatives. Achieving compact signatures is therefore an important step toward the practical deployment of post-quantum cryptography [21].

Among post-quantum signature candidates, multivariate schemes are attractive because they address this problem while maintaining efficient signing and verification. However, several prominent proposals have been broken by increasingly powerful algebraic cryptanalysis.

A notable example is GeMSS, submitted to the NIST Post-Quantum Cryptography standardization process [8]. GeMSS was designed around the goal of achieving extremely compact signatures while retaining practical signing and verification performance. Based on the Hidden Field Equations (HFE) framework, GeMSS illustrated the potential of multivariate cryptography for communication-efficient applications. However, it was eventually broken by Tao *et al.* in [23]. This attack, known as MinRank-S, highlighted the difficulty of designing HFE-based signature schemes that simultaneously achieve compact signatures, practical efficiency, and resistance against known attacks.

To address these weaknesses, previous work [9] proposed the $\mathrm{HFE}^{-}_{\mathrm{IP}}$ framework. The construction combines an Internal Perturbation (IP) modifier, which masks the rank deficiency exploited by the MinRank-S attack, with minus modifiers that hide part of the public equations, thereby providing resistance against attacks such as the hybrid-Gröbner attack [4] and the Kipnis–Shamir attack [17]. However, several directions remained unexplored. In particular, alternatives to the Feistel–Patarin transformation, the use of finite fields larger than $\mathbb{F}_2$, and the resistance of this construction to differential attacks remained open questions.

In this work, we introduce James and Dragon-James (D-James), concrete designs based on the $\mathrm{HFE}^{-}_{\mathrm{IP}}$ construction that can be viewed as repairs of

GeMSS. We introduce new structural and parameter choices aimed at reducing signature size while maintaining practical signing and verification costs. Our use of Dragon terms differs from the approach of Patarin *et al.* [22], who also achieve ultra-short signatures using Dragon terms, but combine them with the LL' modifier rather than with the $\mathrm{HFE}^{-}_{\mathrm{IP}}$ construction.

Our contributions.

1. We incorporate into our scheme the Dragon terms of [20], *i.e.*, bilinear interaction terms between the HFE variables and the hash variables. Unlike the Feistel–Patarin transformations used in previous approaches, these terms allow the number of public equations to be reduced independently of the hash length, significantly decreasing the signature size. We analyze the algebraic properties introduced by Dragon terms and show that the complete $\mathrm{HFE}^{-}_{\mathrm{IP}}$ construction with Dragon terms retains the protection provided by the IP and minus modifiers.
2. We propose new parameter sets for $\mathrm{HFE}^{-}_{\mathrm{IP}}$ and study the impact of different choices of finite fields. In particular, we investigate small prime fields and extension fields with $q > 2$, showing that non-binary fields can significantly reduce public-key sizes, at the cost of some performance.
3. We provide an analysis of differential attacks against $\mathrm{HFE}^{-}_{\mathrm{IP}}$-based constructions. While previous work focused primarily on MinRank resistance, we investigate the differential properties of the construction and evaluate the effect of minus modifiers against these attacks. Our analysis suggests that a small number of minus modifiers makes the currently known differential attacks ineffective for the considered parameter sets.

2 Preliminaries

2.1 Univariate–Multivariate Correspondence

The central idea behind hidden field cryptography is to exploit the distinction between the extension field representation of $\mathbb{F}_{q^n}$ and its vector space representation over $\mathbb{F}_q$. Algebraically simple operations over $\mathbb{F}_{q^n}$ become highly structured multivariate polynomial systems when expressed in coordinates over $\mathbb{F}_q$. HFE exploits this phenomenon by hiding a low-degree univariate polynomial over $\mathbb{F}_{q^n}$ behind secret linear transformations, thereby yielding a public multivariate quadratic system over $\mathbb{F}_q$.

Let q be a prime power, and let $\mathbb{F}_q$ denote the finite field with q elements. Let $\{1, \alpha, \ldots, \alpha^{n-1}\}$ be a polynomial basis of the extension field $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$. Every element of $\mathbb{F}_{q^n}$ can therefore be uniquely identified with a vector in $\mathbb{F}_q^n$. We use the following notation throughout the paper. Elements of $\mathbb{F}_{q^n}$ are denoted by uppercase letters, while elements of $\mathbb{F}_q$ are denoted by lowercase letters. Row vectors over $\mathbb{F}_q$ are written as $\underline{x} \in \mathbb{F}_q^n$.

For a vector $\underline{y} = (y_0, \ldots, y_{n'-1}) \in \mathbb{F}_q^{n'}$, let

$$Y = \sum_{i=0}^{n'-1} y_i \alpha^i,$$

and define its Frobenius vector by

$$\underline{Y} = (Y, Y^q, \ldots, Y^{q^{n-1}}) \in \mathbb{F}_{q^n}^n.$$

Since Frobenius powers are the basic building blocks of HFE monomials, it is convenient to represent all conjugates of an element simultaneously. Then

$$\underline{Y} = \underline{y} M_{n' \times n}, \qquad \text{where} \qquad M_{n' \times n} = \left(\alpha^{iq^j}\right)_{\substack{0 \le i < n' \\ 0 \le j < n}}.$$

In the particular case $n' = n$, we obtain the square (Moore) matrix

$$M_n = M_{n \times n},$$

which is invertible. Hence, for every $\underline{x} \in \mathbb{F}_q^n$,

$$\underline{X} = \underline{x} M_n, \qquad \underline{x} = \underline{X} M_n^{-1}.$$

In our construction, hash vectors may have a dimension $n' \neq n$; the corresponding matrix $M_{n' \times n}$ is then rectangular and hence not invertible. Since only a linear map from the hash space $\mathbb{F}_q^{n'}$ into $\mathbb{F}_{q^n}$ is required, invertibility is unnecessary in this case. Signature vectors, on the other hand, always have dimension n, so M_n provides the inverse transformation back to $\mathbb{F}_q^n$.

2.2 HFE Framework

A fundamental property of finite fields is that the Frobenius maps

$$X \longmapsto X^{q^i}$$

are $\mathbb{F}_q$-linear endomorphisms of $\mathbb{F}_{q^n}$. Since multiplication of two $\mathbb{F}_q$-linear forms yields a quadratic form, monomials of the form

$$X^{q^i + q^j} = X^{q^i} X^{q^j}$$

correspond, through the above vector space representation, to quadratic forms over $\mathbb{F}_q$. More generally, any polynomial whose monomial exponents are sums of two powers of q induces a quadratic multivariate map on $\mathbb{F}_q^n$.

The original HFE framework, introduced by Patarin in [20] as a repair of the Matsumoto–Imai scheme [18], exploits this correspondence. Its secret key is a univariate polynomial of the form

$$H(X) = \sum_{q^i + q^j \le D} \lambda_{i,j} X^{q^i + q^j},$$

where the coefficients belong to $\mathbb{F}_{q^n}$. Equivalently, we can represent H as a matrix

$$H = (\lambda_{i,j})_{0 \le i,j < n} \in \mathbb{F}_{q^n}^{n \times n}, \qquad \lambda_{i,j} = 0 \text{ whenever } q^i + q^j > D,$$

via

$$H(X) = \underline{X}\, H\, \underline{X}^t.$$

We will use these representations interchangeably. Throughout the paper, the rank of H refers to the rank of its coefficient matrix.
The degree bound of H is an integer of the form

$$D = q^d + 1,$$

for some integer d, which we will use repeatedly. Since $q^i + q^j \le q^d + 1$ implies $i, j \le d$, all nonzero coefficients of H lie in its upper-left $(d+1) \times (d+1)$ block, so that

$$\operatorname{rank}(H) \le d + 1,$$

independently of n. This low rank is precisely the structure exploited by the MinRank attacks recalled in Section 4. The choice of D is subject to two conflicting requirements. On the one hand, D must remain sufficiently small so that the legitimate user can efficiently solve

$$Y = H(X)$$

using polynomial factorization algorithms over finite fields. On the other hand, the polynomial must contain enough nonlinear structure to resist known cryptanalytic attacks. In particular, numerous cryptanalytic results, including [7], show that secure HFE instances should contain several nonlinear terms in order to avoid known structural attacks.

To conceal the special algebraic structure of H, two random secret invertible linear transformations

$$S, T : \mathbb{F}_q^n \to \mathbb{F}_q^n$$

are chosen. We denote the public variables by

$$\underline{a} = \underline{x}\, S^{-1}, \qquad \underline{b} = \underline{y}\, T.$$

The public key is then defined as the composition

$$P = T \circ H \circ S,$$

which is represented as a system of n quadratic polynomials in n variables over $\mathbb{F}_q$:

$$P : \begin{cases} b_0 &= p_0(a_0, \ldots, a_{n-1}) \\ &\vdots \\ b_{n-1} &= p_{n-1}(a_0, \ldots, a_{n-1}). \end{cases}$$

The linear transformations S and T are intended to mask the hidden-field structure of the central map, making the public key appear as a random multivariate quadratic (MQ) system. The security of HFE relies on the assumption that recovering an equivalent secret representation from the public MQ system is computationally difficult.

The signing is done as follows:

1. Compute $\text{Hash}(M||r) = (b_0, \ldots, b_{n-1})$, where M is the message and r is a random salt.
2. Compute $\underline{y} = \underline{b}\, T^{-1}$.
3. Interpret $\underline{y}$ in $\mathbb{F}_{q^n}$ via $\underline{Y} = \underline{y}\, M_n$.
4. Solve for X in $Y = H(X)$. If no preimage exists, choose a fresh random salt r and restart from Step 1.[3]
5. Interpret X in $\mathbb{F}_q^n$ via $\underline{x} = \underline{X}\, M_n^{-1}$.
6. Apply S^{-1} to $\underline{x}$. This gives $(a_0, \ldots, a_{n-1}) = \underline{x}\, S^{-1}$.
7. The signature is $(\underline{a}, r)$.

For verification, given a message M and a signature $(\underline{a}, r)$, the verifier accepts if and only if

$$P(\underline{a}) = \text{Hash}(M||r).$$

2.3 HFE Variants

The HFE scheme has been shown to be susceptible to super-polynomial algebraic attacks for certain parameter choices [14]. Several modifications have therefore been proposed to strengthen the scheme against known attacks. In this work, we focus on the two perturbations used in HFE^-_{IP}: the minus modifier and the IP modifier.

The minus modifier consists of hiding part of the public map. One chooses an integer $a < n$ and publishes only the first $n - a$ equations, so that the public key becomes

$$P_- = \Pi \circ T \circ H \circ S,$$

where $\Pi : \mathbb{F}_q^n \to \mathbb{F}_q^{n-a}$ denotes the projection onto the first $n - a$ coordinates. The intuition is that the missing equations reduce the amount of information available to an attacker, which makes algebraic reconstruction of the hidden structure more difficult and, in particular, complicates MinRank-based attacks [24].

To sign a hashed message of length $n - a$, one appends a uniformly random elements of $\mathbb{F}_q$ and signs the resulting vector of $\mathbb{F}_q^n$ using the standard HFE signing procedure. The appended random elements provide the randomness required during signing and therefore replace the random salt used in the original construction; the signature then consists of $\underline{a}$ only. This also applies to JAMES.

[3] The expected number of trials is about 1.582 when the characteristic is 2 and 3.07 otherwise. This was studied in [15].

The IP modifier introduced in [10] works as follows. Let $Z = (\ell_1, \ldots, \ell_r)$ denote a linear map of rank r, where $\ell_1, \ldots, \ell_r$ are linear forms, so that $Z(X) = (z_1, \ldots, z_r)$ with $z_i = \ell_i(x_0, \ldots, x_{n-1})$.

The central map is modified by adding random bilinear terms in the variables x and z, together with quadratic terms in the variables z. This can be described by considering an HFE polynomial of the form

$$H_{\mathrm{IP}}(X) = H(X) + M(X, Z(X)) + \overline{H}(Z(X)),$$

where $\overline{H}$ is a random quadratic polynomial and M is a random $\mathbb{F}_q$-bilinear map whose degree with respect to X is less than q^d. Note that, unlike the HFE$_{v-}$ scheme used in GeMSS, we do not add new variables here, since the z variables are linear combinations of the x variables.

To sign, we guess the values $z_1, \ldots, z_r$. Fixing the guessed values transforms the central map into a standard HFE polynomial, which can then be inverted using the secret key. This yields a candidate preimage $(x_0, \ldots, x_{n-1})$, which is accepted only if it is consistent with the guess, *i.e.*,

$$\ell_i(x_0, \ldots, x_{n-1}) = z_i, \qquad i = 1, \ldots, r.$$

Otherwise, we try another assignment. If no valid preimage is obtained after testing all q^r assignments, we append a random salt before hashing the message and repeat the process. This incurs an expected slowdown by a factor q^r compared to unmodified HFE. Hence, q^r must remain small in practice to keep the signing cost acceptable.

3 Ultra-Short Signatures and the Dragon Scheme

A fundamental limitation on signature size arises from the fact that the message digest $\underline{b}$ is the output of a cryptographic hash function. To provide λ-bit collision resistance against generic birthday attacks, the hash output should contain at least 2λ bits. Consider a public system of the form

$$P : \begin{cases} b_0 = p_0(a_0, \ldots, a_{n-1}) \\ \qquad \vdots \\ b_{m-1} = p_{m-1}(a_0, \ldots, a_{n-1}). \end{cases}$$

For a random quadratic system over a finite field, one expects a solution with non-negligible probability only when the number of variables is comparable to the number of equations. Hence one typically chooses $n \approx m$. It follows that 128-bit security requires a 256-bit hash output, yielding signatures of approximately 256 bits.

To overcome this limitation and obtain ultra-short signatures, Patarin proposed in [20] a modified construction in which the hash variables appear directly

in the central equations:

$$P : \begin{cases} c_0 = p_0(a_0, \ldots, a_{n-1}) & + g_0(a_0, \ldots, a_{n-1}, b_0, \ldots, b_{n'-1}) \\ & \vdots \\ c_{n-1} = p_{n-1}(a_0, \ldots, a_{n-1}) + g_{n-1}(a_0, \ldots, a_{n-1}, b_0, \ldots, b_{n'-1}), \end{cases}$$

where $\underline{b} = (b_0, \ldots, b_{n'-1})$ is the hash value, each polynomial g_i is bilinear in the variables a and b (and therefore has total degree two), and $\underline{c} \in \mathbb{F}_q^n$ is a *fixed nonzero constant vector* (for instance, the all-ones vector). We call the bilinear terms g_i *Dragon terms*. The vector $\underline{c}$ is fixed once and for all as part of the scheme's public parameters (it is not part of the public key). A nonzero right-hand side is necessary: the homogeneous system is always satisfied by $\underline{a} = \underline{0}$ regardless of $\underline{b}$, which would allow a trivial universal forgery, since $(\underline{0}, r)$ would verify for every message M and every salt r.
After fixing the hash value $\underline{b}$, every bilinear term $a_i b_j$ becomes a linear term in the variables a_i, so the system reduces to a quadratic system with n equations in n unknowns. Since the number of equations is independent of the hash length, one may choose $n' = 256$ (over $\mathbb{F}_2$) to obtain 128-bit security while keeping the signature size close to $n = 128$ bits (in practice, n is usually chosen slightly larger to take into account direct attacks on random quadratic systems such as [16]). This decouples the number of public equations from the hash output length.
This construction extends the HFE central map by introducing bilinear terms involving both X and Y:

$$H(X, Y) = \sum_{q^i + q^j \le D} \lambda_{i,j} X^{q^i + q^j} + \sum_{\substack{q^i \le D \\ j < n}} \mu_{i,j} X^{q^i} Y^{q^j}.$$

An equivalent formulation is obtained by introducing secret linear maps

$$L_0, \ldots, L_d,$$

defined over the hash variables. The central polynomial can then be written as

$$H(X, Y) = \sum_{q^i + q^j \le D} \lambda_{i,j} X^{q^i + q^j} + \sum_{q^i \le D} L_i(y_0, \ldots, y_{n'-1}) X^{q^i}.$$

The equivalence follows from the fact that every $\mathbb{F}_q$-linear map admits a representation as a linearized polynomial,

$$L_i(y_0, \ldots, y_{n'-1}) = \sum_j \mu_{i,j} Y^{q^j}, \qquad \text{with} \qquad Y = \sum_{k=0}^{n'-1} y_k \alpha^k.$$

After fixing the hash value, the quantities $L_i(y)$ become constants in $\mathbb{F}_{q^n}$, and the central polynomial reduces to an HFE polynomial of degree D. Hence, finding

a root of the central polynomial can be performed with essentially the same complexity as in the original HFE scheme.

Combination with the IP modifier. In D-JAMES, the IP perturbation also acts on the Dragon part: in addition to $M(X, Z(X))$ and $\overline{H}(Z(X))$, we add random bilinear terms $z_k y_j$ between the perturbation variables and the hash variables,

$$N(Z(X), Y) = \sum_{k=1}^{r} \sum_{j=0}^{n'-1} \nu_{k,j}\, z_k\, y_j, \qquad \nu_{k,j} \in \mathbb{F}_{q^n}.$$

The central map of D-JAMES is therefore

$$H_{\mathrm{IP}}(X, Y) = \underbrace{H(X) + M(X, Z(X)) + \overline{H}(Z(X))}_{\mathrm{HFE_{IP}\ part}} + \underbrace{\sum_{q^i \le D} L_i(y_0, \ldots, y_{n'-1}) X^{q^i} + N(Z(X), Y)}_{\mathrm{Dragon\ part}}.$$

Since the z_k are linear in the x variables, the terms $z_k y_j$ contribute to the monomials $a_i b_j$ of the public key, and play for the Dragon part the role that $M(X, Z(X))$ plays for the HFE part. These terms come at no cost: during signing, y is known and z is guessed, so $N(Z(X), Y)$ is a constant. The public key size is unchanged.
As in HFE, the public key is obtained by composing the central map with two secret linear transformations masking the algebraic structure of the central map. There is, however, one major difference from vanilla HFE: here T acts linearly on the output equations, and not on the hash variables as in vanilla HFE. In particular, the hash value is used directly as $\underline{y} = \underline{b}$.

The signing is done as follows:

1. Compute $\underline{C} = \underline{c}\, T^{-1} M_n$.
2. Compute $\underline{y} = \underline{b} = \mathrm{Hash}(M||r) = (b_0, \ldots, b_{n'-1})$, where M is the message and r is a random salt.
3. Map $\underline{y} \in \mathbb{F}_q^{n'}$ to $Y \in \mathbb{F}_{q^n}$ via $\underline{Y} = \underline{y}\, M_{n' \times n}$.
4. Substitute the resulting value Y in $H(X, Y)$. This gives a univariate polynomial of degree D in X.
5. Find a root X of $H(X, Y) = C$. If no root exists, choose a fresh random salt r and restart from Step 2.
6. Interpret X in $\mathbb{F}_q^n$ via $\underline{x} = \underline{X}\, M_n^{-1}$.
7. Apply S^{-1} to $\underline{x}$. This gives $(a_0, \ldots, a_{n-1}) = \underline{x}\, S^{-1}$.
8. The signature is $(\underline{a}, r)$.[4]

[4] We can also omit the salt from the signature and let the verifier make an exhaustive search for the salt if the salt-space is small. This gives shorter signatures but longer verification times, which is what we preferred.

With the IP modifier, Steps 4 and 5 are repeated for each guess of $(z_1, \ldots, z_r)$, and a root is accepted only if it is consistent with the guess, as in $\mathrm{HFE}_{\mathrm{IP}}$.

For verification, given a message M and a signature $(\underline{a}, r)$, the verifier computes $\underline{b} = \mathrm{Hash}(M||r)$ and accepts if and only if

$$P(\underline{a}, \underline{b}) = \underline{c}.$$

4 MinRank Attacks on Dragon Terms

4.1 Original Attacks

Recovering T.
Recall that in [5], the authors showed that the public key can be expressed as

$$P = (P_0, \ldots, P_{n-1}) = (SM_nH^{*0}M_n^tS^t, \ldots, SM_nH^{*n-1}M_n^tS^t)\, M_n^{-1}T. \quad (1)$$

Here, H^{*k} denotes the matrix representation of the q^k-th Frobenius power applied to the internal polynomial, namely

$$\left(H^{*k}\right)_{i,j} = \left(H_{i-k,j-k}\right)^{q^k}.$$

Setting $U = T^{-1}M_n$ and $W = SM_n$, we obtain

$$(P_0, \ldots, P_{n-1})\, U = (WH^{*0}W^t, \ldots, WH^{*n-1}W^t). \quad (2)$$

Taking the first column of U yields

$$\sum_{i=0}^{n-1} u_{i,0}P_i = WHW^t,$$

which corresponds to a MinRank instance since H has the form

$$H = \begin{pmatrix} A & 0 \\ 0 & 0 \end{pmatrix},$$

with A of size $d+1$. Solving this instance yields the first column of U, from which the remaining columns can be easily recovered [5]. This attack is known as the MinRank-T attack and was introduced in [17].

Recovering S.
A second attack, known as MinRank-S, was introduced in [23]. Equation (2) can be rewritten as

$$(W^{-1}P_0W^{-t}, \ldots, W^{-1}P_{n-1}W^{-t}) = (H^{*0}, \ldots, H^{*n-1})\, U^{-1}. \quad (3)$$

[5] Because T^{-1} has coefficients in $\mathbb{F}_q$ and M_n is the Moore matrix, we have $u_{i,j} = u_{i,j-1}^q$.

Let

$$Q = U^{-t} \begin{pmatrix} h_0 \\ \vdots \\ h_{n-1} \end{pmatrix},$$

where h_i denotes the first row of H^{*i}. We then have

$$Q = U^{-t} \begin{pmatrix} H_1 \\ 0 \\ H_2 \end{pmatrix},$$

with $H_1 \in \mathbb{F}_{q^n}^{1\times n}$ and $H_2 \in \mathbb{F}_{q^n}^{d\times n}$, and a zero block of size $(n-d-1)\times n$. It follows that

$$\operatorname{rank}(Q) \leq d+1.$$

If we denote by $l = (l_0, \ldots, l_{n-1})$ the first row of W^{-1}, then stacking the first rows of the block equations in (3) yields

$$\begin{pmatrix} lP_0 \\ \vdots \\ lP_{n-1} \end{pmatrix} W^{-t} = Q.$$

For each $j \in \{0, \ldots, n-1\}$, let R_j be the matrix whose i-th row is the j-th row of P_i. Then

$$\begin{pmatrix} lP_0 \\ \vdots \\ lP_{n-1} \end{pmatrix} = \sum_{j=0}^{n-1} l_j R_j,$$

so that

$$\left(\sum_{j=0}^{n-1} l_j R_j \right) W^{-t} = Q,$$

or equivalently,

$$\sum_{j=0}^{n-1} l_j R_j = QW^t.$$

Hence, finding l[6] is equivalent to finding coefficients (l_j) such that the matrix

$$\sum_{j=0}^{n-1} l_j R_j$$

has rank at most $d+1$. This is precisely a MinRank instance over the linear space spanned by the matrices $R_0, \ldots, R_{n-1}$. Solving it gives l, from which we recover W^{-1}, and hence S.

[6] Here again we can recover W^{-1} from l by applying the Frobenius.

4.2 Extension to the Dragon Terms

We now consider a public key containing Dragon terms of the form

$$\begin{cases} \sum_{i,j} \alpha_{ij}^{(0)}\, a_i a_j \quad + \sum_{i,j} \beta_{ij}^{(0)}\, a_i b_j \quad = c_0 \\ \qquad\qquad\qquad\qquad \vdots \\ \sum_{i,j} \alpha_{ij}^{(n-1)}\, a_i a_j + \sum_{i,j} \beta_{ij}^{(n-1)}\, a_i b_j = c_{n-1}, \end{cases}$$

where the constant vector $\underline{c}$ plays no role in the quadratic forms considered below, with an internal polynomial

$$H(X,Y) = \sum_{q^i+q^j \le D} \lambda_{i,j} X^{q^i+q^j} + \sum_{\substack{q^i \le D \\ j<n}} \mu_{i,j} X^{q^i} Y^{q^j}.$$

The public key is then written as $\tilde{P} = (\tilde{P}_0, \ldots, \tilde{P}_{n-1})$, where each $\tilde{P}_k$ is an $(n+n') \times (n+n')$ matrix of the form[7]

$$\tilde{P}_k = \begin{pmatrix} \left(\alpha^{(k)}\right) & \left(\beta^{(k)}\right) \\ \left(\beta^{(k)}\right)^t & 0 \end{pmatrix}.$$

Recovering T.

Setting

$$U = T^{-1} M_n, \qquad \tilde{W} = \begin{pmatrix} W & 0 \\ 0 & M_{n' \times n} \end{pmatrix}, \qquad \tilde{H} = \begin{pmatrix} (\lambda_{ij}) & (\mu_{ij}) \\ 0 & 0 \end{pmatrix},$$

such that

$$H(X,Y) = (\underline{X}, \underline{Y})\, \tilde{H}\, (\underline{X}, \underline{Y})^t,$$

we obtain

$$(\tilde{P}_0, \ldots, \tilde{P}_{n-1})\, U = (\tilde{W}\tilde{H}^{*0}\tilde{W}^t, \ldots, \tilde{W}\tilde{H}^{*n-1}\tilde{W}^t). \tag{4}$$

Taking the first column of U in (4), we get

$$\sum_{i=0}^{n-1} u_{i,0} \tilde{P}_i = \tilde{W}\tilde{H}\tilde{W}^t.$$

Only the first $d+1$ rows of $\tilde{H}$ are nonzero, so that $\mathrm{rank}(\tilde{W}\tilde{H}\tilde{W}^t) \le \mathrm{rank}(\tilde{H}) \le d+1$. Therefore, we obtain a MinRank-T instance with the same target rank as in the original HFE case, but involving matrices of size $(n+n') \times (n+n')$.

[7] In characteristic 2, care must be taken since quadratic representations are not unique and the matrices are not necessarily symmetric.

Recovering S.
Similarly, for the MinRank-S attack, equation (4) can be rearranged into

$$(\tilde{H}^{*0}, \ldots, \tilde{H}^{*n-1})\, U^{-1} = (V\tilde{P}_0 V^t, \ldots, V\tilde{P}_{n-1} V^t)$$

with

$$V = (\tilde{W}^t \tilde{W})^{-1} \tilde{W}^t.$$

We define

$$Q = U^{-t} \begin{pmatrix} \tilde{h}_0 \\ \vdots \\ \tilde{h}_{n-1} \end{pmatrix},$$

where $\tilde{h}_k$ denotes the first row of $\tilde{H}^{*k}$. Then, Q is of the form

$$Q = U^{-t} \begin{pmatrix} \tilde{H}_1 \\ 0 \\ \tilde{H}_2 \end{pmatrix},$$

with $\tilde{H}_1 \in \mathbb{F}_{q^n}^{1\times 2n}$ and $\tilde{H}_2 \in \mathbb{F}_{q^n}^{d\times 2n}$. It follows that

$$\operatorname{rank}(Q) \leq d + 1.$$

Let l denote the first row of V. Then stacking the first rows of the block equation gives

$$\begin{pmatrix} l \cdot \tilde{P}_0 \\ \vdots \\ l \cdot \tilde{P}_{n-1} \end{pmatrix} V^t = Q.$$

Since $\tilde{P}_k$ is a linear combination of the matrices $\tilde{W}\tilde{H}^{*s}\tilde{W}^t$ and $V\tilde{W} = I$, the left-hand matrix above equals $Q\tilde{W}^t$, exactly as in the HFE case. Therefore,

$$\operatorname{rank} \begin{pmatrix} l \cdot \tilde{P}_0 \\ \vdots \\ l \cdot \tilde{P}_{n-1} \end{pmatrix} \leq d + 1.$$

This again yields a MinRank instance, with the same target rank but in dimension $n + n'$.

Effect of the modifiers. In D-James, the IP perturbation acts on both blocks of $\tilde{H}$: $M(X, Z(X)) + \overline{H}(Z(X))$ on (λ_{ij}) and $N(Z(X), Y)$ on (μ_{ij}). Since each z_k is a combination of all the X^{q^s}, the contribution of N fills all the rows of (μ_{ij}), exactly as $M(X, Z(X))$ does for (λ_{ij}). The MinRank-T and MinRank-S attacks therefore work in the same way as on $\mathrm{HFE}^-_{\mathrm{IP}}$, and the IP and minus modifiers increase the target rank in the same way on the Dragon block as on the HFE block.
The terms $z_k y_j$ are necessary: an attacker can restrict the instances to the blocks $\beta^{(k)}$, which M and $\overline{H}$ do not modify. Without N, this restricted MinRank-S

instance keeps the target rank $d+1$ for any r. On small instances ($q = 2$, $D = 5$, $n \in \{31, 47\}$, $r \leq 3$, with minus), the restricted matrix has rank 3 without N and full rank with N.

Discussion. The above analysis shows that the Dragon extension does not alter the fundamental MinRank structure underlying the attacks against HFE. In both the MinRank-T and MinRank-S attacks, the target rank remains bounded by the same value $d+1$ as in the original HFE setting; only the ambient dimension is increased by the additional Dragon variables. Thus, the Dragon terms do not provide a new low-rank structure with a smaller target rank that could be exploited by the same type of attack.

More importantly, the MinRank instances obtained above have the same structural origin as their HFE counterparts. Consequently, the structural countermeasures introduced against MinRank attacks on HFE apply to the Dragon extension as well. In particular, the minus and IP modifications disrupt the algebraic structures exploited by the MinRank-T and MinRank-S attacks, respectively, and the presence of the additional Dragon terms does not restore these structures, provided that the IP perturbation also includes the terms $z_k y_j$ described above. Hence, the Dragon extension does not introduce a new MinRank vulnerability beyond those already present in the underlying HFE construction.

From a complexity perspective, the target rank remains fixed while the ambient dimension increases from n to $n + n'$. Generic MinRank algorithms are therefore not expected to become more efficient merely from the addition of the Dragon terms; on the contrary, their cost typically increases with the ambient dimension for a fixed target rank. Thus, within the class of MinRank attacks considered here, there is no apparent advantage in exploiting the Dragon terms directly. The natural strategy is instead to disregard the additional Dragon coordinates and apply the corresponding MinRank attack to the underlying HFE instance. In this sense, the Dragon extension does not weaken the construction with respect to the MinRank attacks considered here, and its security against these attacks can be analyzed using the same structural considerations as for the underlying HFE construction.

5 Analysis of Differential Attacks on $\mathbf{HFE}^{-}_{\mathbf{IP}}$

The original attack. Dubois, Granboulan, and Stern showed in [12] that $\mathrm{HFE}_{\mathrm{IP}}$ is susceptible to differential attacks. In short, their attack recovers the secret kernel $\ker(S \circ Z)$ by studying the distribution of the kernel dimension of the differential of the public key. Once this structure is recovered, the system can be attacked as a standard instance of HFE.

We recall that, for a multivariate quadratic map F, its differential is the bilinear map

$$DF(\cdot,\cdot) : x, y \mapsto F(x+y) - F(x) - F(y) + F(0).$$

The attack relies on the observation that the differential of the public map carries information about the secret subspace

$$\mathcal{K} = \ker(S \circ Z).$$

More precisely, the distribution of the kernel dimension of the differential depends on whether the evaluation point belongs to this hidden subspace. This statistical bias is precisely what the attack exploits.
They prove that, for every x, the differential can be written as

$$DH_{IP}(\cdot, x) = L_{d,x} + l_{r,x},$$

with $L_{d,x}$ being an $\mathbb{F}_2$-linear polynomial of degree 2^d, and $l_{r,x}$ a linear map of rank r. Moreover,

$$L_{d,x}(x) = l_{r,x}(x),$$

and

$$\begin{aligned} x \in \mathcal{K} &\implies L_{d,x}(x) = l_{r,x}(x) = 0, \\ x \notin \mathcal{K} &\implies L_{d,x}(x) = l_{r,x}(x) \neq 0. \end{aligned}$$

Averaging over the possible choices of $L_{d,x}$, $l_{r,x}$, and x, they derive the conditional distributions

$$\begin{aligned} \pi^+(t) &= \Pr(\dim\ker(L_{d,x} + l_{r,x}) = t \mid L_{d,x}(x) = l_{r,x}(x) = 0)\,, \\ \pi^-(t) &= \Pr(\dim\ker(L_{d,x} + l_{r,x}) = t \mid L_{d,x}(x) = l_{r,x}(x) \neq 0)\,. \end{aligned}$$

These probabilities are obtained through combinatorial counting arguments. The attack relies on the fact that these two distributions are statistically distinguishable, even though the difference between them is small.

For a given x, they compute

$$t = \dim\ker DP_{IP}(\cdot, x).$$

If $\pi^+(t) \geq \pi^-(t)$, then x is deemed more likely to belong to $\mathcal{K}$ than not. Their distinguisher is therefore

$$\chi : \begin{cases} \chi(x) = 1 & \text{if } \pi^+(t) \geq \pi^-(t), \\ \chi(x) = 0 & \text{if } \pi^+(t) < \pi^-(t). \end{cases}$$

This can be viewed as a noisy membership test.
The distinguisher alone is too noisy to identify the kernel reliably. To amplify the signal, they exploit the linearity of $\mathcal{K}$. Indeed, if $x \in \mathcal{K}$, then for every y, either both y and $x + y$ belong to $\mathcal{K}$, or neither does. Consequently,

$$\chi(x + y) = \chi(y)$$

holds slightly more often than expected by chance.

For a large sample $y_1, \ldots, y_N$ of random binary vectors, they compute

$$\frac{1}{N} \sum_{i=1}^{N} \mathbf{1}_{\chi(x+y_i)=\chi(y_i)}.$$

If this empirical agreement exceeds the expected average, then x is declared to belong to $\mathcal{K}$. They choose N sufficiently large so that the probability of false positives becomes negligible.

Our experimental implementation. We investigated whether the differential attack of Dubois, Granboulan, and Stern can be extended to $\mathrm{HFE}^-_{\mathrm{IP}}$.
To this end, we first reimplemented the original attack and observed that our implementation reproduces the behavior reported in [12] for $\mathrm{HFE}_{\mathrm{IP}}$. We then applied the same methodology to $\mathrm{HFE}^-_{\mathrm{IP}}$ instances obtained by introducing minus components.
Even a single minus component prevented successful recovery of the hidden kernel in our experiments with the considered attack parameters (Table 1). A recovery is counted as successful if the algorithm outputs a basis of a subspace equal to the secret kernel $\mathcal{K}$.
Since the complexity of the original attack grows rapidly with the size of the parameters, we report complete key recovery experiments only for the smallest parameter set for which the attack remains computationally feasible. Larger parameter sets were also tested, and the distinguisher exhibited the same qualitative behavior, but we did not run the full recovery procedure due to its prohibitive cost.

Table 1: Experimental evaluation of the differential attack of Dubois, Granboulan, and Stern on $\mathrm{HFE}_{\mathrm{IP}}$ and $\mathrm{HFE}^-_{\mathrm{IP}}$. Each row corresponds to 100 independently generated secret keys.

n	D	r	Minus components (a)	Successful recoveries
32	3	1	0	100/100
32	3	1	1	0/100
32	3	1	2	0/100

Heuristic interpretation. As the number of minus components increases, the empirical distribution of kernel dimensions progressively approaches that of a random binary matrix. This suggests that the statistical information exploited by the original distinguisher becomes increasingly difficult to detect.

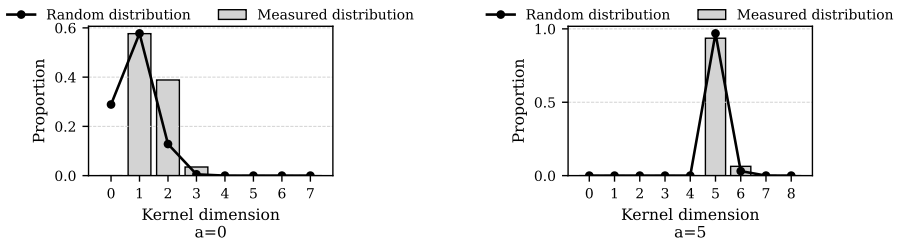

Fig. 1: Empirical distribution of $\dim\ker DP_{\mathrm{HFE}^-_{\mathrm{IP}}}(\cdot, x)$ for $n = 32$, $d = 1$, $r = 1$, $N_{\mathrm{samples}} = 10^4$, with $a = 0$ (left) and $a = 5$ (right), compared with the theoretical distribution of the kernel dimension of a random $(n - a) \times n$ binary matrix.

A plausible heuristic explanation is the following. The additional minus components appear to randomize the differential, whose rank empirically behaves like that of a random linear mapping of rank close to $n - a$, as shown in Figure 1.

We illustrate this behavior for $a = 5$. Similar behavior was observed for larger values of a, and for the recommended parameters the empirical and theoretical distributions are visually indistinguishable.

Under this heuristic, the kernel dimension is expected to follow approximately the same distribution as that of a random binary matrix. For reference, if $M \in \mathbb{F}_2^{m \times n}$, then

$$\Pr(\mathrm{rank}(M) = k) = \frac{1}{2^{mn}} \prod_{i=0}^{k-1} \frac{(2^n - 2^i)(2^m - 2^i)}{(2^k - 2^i)}.$$

In particular, for $M \in \mathbb{F}_2^{(n-a)\times n}$,

$$\Pr(\dim \ker M = a) = 1 - 2^{-a} + 2^{-n} + \mathcal{O}(2^{-2a}).$$

This would naturally explain why the original distinguisher appears ineffective in our experiments. If this heuristic is correct, then the conditional distributions π^+ and π^- are expected to become nearly identical, leaving the original distinguisher with little or no information about membership in $\mathcal{K}$.

Discussion. These observations do not exclude the existence of alternative distinguishers based on richer joint distributions of differential kernels. Whether such invariants can be exploited algorithmically remains an open question.

We emphasize that our goal is not to rule out the existence of differential attacks against $\mathrm{HFE}^-_{\mathrm{IP}}$. Rather, our experiments indicate that the statistical distinguisher introduced by Dubois, Granboulan, and Stern does not appear to extend directly once minus components are introduced: we did not observe the statistical bias exploited by the original attack in any of our experiments. A theoretical explanation of this phenomenon remains an interesting open problem.

It is worth noting that Dubois, Fouque, Shamir, and Stern described an attack on SFLASH in [11], which is a C^{*-} scheme. Their attack is algebraic and uses the fact that the hidden polynomial consists of a single monomial. It is unclear whether this attack can be adapted to an HFE^- scheme, let alone to $\mathrm{HFE}^-_{\mathrm{IP}}$.

6 Parameter Selection and Security Analysis

6.1 Parameter Selection

We select the parameters to meet the target security level against the quantitative attacks considered in this section, while also taking into account the experimental differential analysis. We balance security, public-key size, signature size, and signing performance. Table 2 summarizes the role of the different attacks in the parameter selection.

Table 2: Role of the different attacks in the parameter selection.

Attack	Constraint on parameters
Gröbner	Lower bound on m
MinRank-T	Lower bound on a (hence $n = m + a$)
MinRank-S	Lower bound on r
Differential	Accounted for by the choice of a
Birthday (MITM)	Lower bound on m for James
Hash collision	Lower bound on n' for D-James

We consider two families of parameter sets. The first (Tables 3 and 4) uses the Dragon terms and provides our main proposed instantiations. The second (Tables 5 and 6) omits the Dragon terms and serves as a conservative fallback in case future cryptanalysis weakens their security contribution.

For each parameter set, we report two signature sizes. The first, $|\mathtt{sig}_{\text{fast}}|$, corresponds to the standard signature size, *i.e.*, $n \log_2 q$ bits. We also consider a shortened variant, denoted $|\mathtt{sig}_{\text{short}}|$, obtained by truncating the signature following [21]. In this setting, the missing part of the signature is recovered during verification using a hybrid Gröbner basis attack, as proposed in [4]. The value of $|\mathtt{sig}_{\text{short}}|$ is selected so that signature verification completes in less than one minute using MAGMA V2.29-7 on a personal laptop.

Table 3: Recommended parameter sets for D-James at the 128-bit security level.

Parameters								$\lvert\mathtt{sig}_{\text{fast}}\rvert$	$\lvert\mathtt{sig}_{\text{short}}\rvert$	$\lvert\mathtt{PK}\rvert$
q	m	n	a	n'	r	D	Monomials	(bits)	(bits)	(Bytes)
2	162	189	27	256	2	5	X^3, X^5	189	156	1.35M
4	84	105	21	128	2	17	X^5, X^{17}	210	158	400k
5	73	94	21	111	2	6	X^2, X^6	219	163	317k
13	56	77	21	70	2	14	X^2, X^{14}	285	208	219k
23	53	74	21	57	2	24	X^2, X^{24}	335	245	211k

q: field size; m: number of public equations; n: extension field degree; a: number of removed equations (minus modifier, $a = n - m$); n': number of hash variables y; r: number of IP perturbation variables; D: degree of the internal HFE polynomial; Monomials: quadratic monomials appearing in H.

Table 4: Recommended parameter sets for D-James at the 256-bit security level (notation as in Table 3).

Parameters								$\lvert\mathtt{sig}_{\text{fast}}\rvert$	$\lvert\mathtt{sig}_{\text{short}}\rvert$	$\lvert\mathtt{PK}\rvert$
q	m	n	a	n'	r	D	Monomials	(bits)	(bits)	(Bytes)
2	324	390	66	512	2	5	X^3, X^5	390	348	11.1M
4	170	223	53	256	2	17	X^5, X^{17}	446	382	3.49M
5	153	206	53	221	2	6	X^2, X^6	479	409	2.97M
13	118	171	53	139	2	14	X^2, X^{14}	633	530	2.10M
23	110	163	53	114	2	24	X^2, X^{24}	738	616	1.99M

Table 5: Recommended parameter sets for James at the 128-bit security level (notation as in Table 3).

Parameters							$\lvert\mathtt{sig}_{\text{fast}}\rvert$	$\lvert\mathtt{sig}_{\text{short}}\rvert$	$\lvert\mathtt{PK}\rvert$
q	m	n	a	r	D	Monomials	(bits)	(bits)	(Bytes)
2	256	283	27	2	5	X^3, X^5	283	238	1.29M
4	128	149	21	2	17	X^5, X^{17}	298	242	358k
5	111	132	21	2	6	X^2, X^6	307	247	283k
13	70	91	21	2	14	X^2, X^{14}	337	245	136k
23	57	78	21	2	24	X^2, X^{24}	353	249	99.3k

Table 6: Recommended parameter sets for JAMES at the 256-bit security level (notation as in Table 3).

Parameters							$\|\mathtt{sig}_{\text{fast}}\|$	$\|\mathtt{sig}_{\text{short}}\|$	$\|\mathtt{PK}\|$
q	m	n	a	r	D	Monomials	(bits)	(bits)	(Bytes)
2	512	578	66	2	5	X^3, X^5	578	523	10.7M
4	256	309	53	2	17	X^5, X^{17}	618	542	3.07M
5	221	274	53	2	6	X^2, X^6	637	553	2.41M
13	139	192	53	2	14	X^2, X^{14}	711	596	1.19M
23	114	167	53	2	24	X^2, X^{24}	756	634	904k

All proposed parameter sets use $r = 2$. Nevertheless, it is conceivable that $r = 1$ already provides sufficient protection against the MinRank attack of [23]. As the signing cost increases by a factor of q^r, establishing the minimum secure value of r is an interesting avenue for future research and could substantially improve the performance of JAMES.

Working over larger fields such as $\mathbb{F}_4$ or $\mathbb{F}_5$ reduces the public-key size by a factor of about 3 to 5 at the cost of slightly longer signatures; larger values of q reduce it further, but at a significant cost in signing time (see Table 10).

We do not use the Feistel–Patarin transformation on JAMES in order to preserve good performance (see Section 7). Although it would produce shorter signatures, it would also slow down signing by up to a factor of four. Since the IP modifier already introduces a fourfold signing overhead, we chose not to apply this transformation. Nevertheless, the Feistel–Patarin transformation can be incorporated without affecting the security of the scheme, should shorter signatures be desired.

6.2 Security Analysis

We now evaluate the complexity of the main known attacks and verify that the resulting parameter sets meet their respective security targets.

Birthday. For JAMES, fixing X determines a unique value that can be compared against a hash value. Hence, a standard meet-in-the-middle (MITM) attack evaluates the public key on $q^{m/2}$ values of X and computes $q^{m/2}$ hash values, leading to an overall complexity of $q^{m/2}$.

On D-JAMES, fixing X turns the public-key system into a linear system

$$A_X Y = B_X$$

with n' variables and m equations. Assuming that A_X has full rank, the set of compatible hash values forms an affine subspace of dimension $n' - m$, and therefore contains $q^{n'-m}$ elements. Since the hash output consists of n' elements of $\mathbb{F}_q$,

a uniformly random hash value belongs to this affine subspace with probability

$$q^{-m}.$$

Consequently, for N_X values of X and N_H hash values, the expected number of compatible pairs (X, H) is

$$N_X N_H q^{-m}.$$

In particular, taking

$$N_X = N_H = q^{m/2}$$

gives a constant expected number of compatible pairs. This establishes a birthday-type *existence bound*, but not necessarily an attack of complexity $q^{m/2}$.
Indeed, unlike the standard MITM attack on JAMES, where each X is associated with a single value that can be inserted into a lookup table, here each X defines an affine subspace of hash values. Thus, a direct table lookup does not reveal whether a given hash value belongs to the corresponding subspace. In the absence of additional structure, a straightforward search must test the membership of each hash value in each affine subspace, resulting in

$$N_X N_H = q^m$$

membership tests. We therefore do not regard the $q^{m/2}$ birthday estimate as an attack complexity: it only describes the data size at which compatible pairs are expected to exist.
This search problem is related in spirit to several well-studied problems in coding theory and cryptanalysis, including syndrome decoding and generalized birthday algorithms [25], [19], [2]. However, these techniques do not immediately yield a $q^{m/2}$-time algorithm for our setting. In particular, the affine subspace associated with a given X depends on the full value of X, and we do not identify a decomposition into independently generated partial values that would enable a standard generalized-birthday or meet-in-the-middle approach. We therefore use the exhaustive matching strategy, with complexity q^m, as the generic attack considered in our security estimates. This is not a proof of optimality, but rather the best generic matching complexity currently known to us.
This raises the question of whether the linear structure in Y could be exploited to obtain an attack faster than the q^m exhaustive matching strategy considered above. If such an attack were found, one possible countermeasure would be to add quadratic terms in Y to the public-key equations. This modification would not affect the signing procedure, while preventing the direct linear-algebraic recovery of Y from X. The resulting public key would be larger, however, and we therefore leave these additional terms out of the construction.

For D-JAMES, an attacker can also directly search for collisions in the hash function, with complexity $q^{n'/2}$. This bound is independent of the affine-subspace matching problem discussed above, and is the relevant generic bound reported in the Birthday column of Tables 7 and 8.

Hybrid. The second attack is the hybrid Gröbner basis attack introduced by Bettale *et al.* [4]. The attack consists in exhaustively guessing k variables and then solving the resulting system in the remaining $n-k$ variables using a Gröbner basis algorithm. Its complexity can be estimated as

$$\mathcal{O}\left(q^k \binom{n-k+d_{\text{reg}}}{d_{\text{reg}}}^{\omega}\right),$$

where d_{reg} denotes the degree of regularity and $\omega \approx 2.37$ is the linear algebra exponent. Following the complexity estimates of Faugère [13], we estimate the cost of the Gröbner basis computation using the corresponding binomial coefficient. We performed an exhaustive search on k with `Magma` to minimize that quantity.

Over $\mathbb{F}_2$, Joux and Vitse [16] exhibited another attack, with complexity

$$\tilde{\mathcal{O}}\left(2^{0.792\,m}\right).$$

For $q = 2$, we therefore need

$$m \geq \frac{\lambda}{0.792} \approx 1.26\,\lambda,$$

where λ denotes the target security level, *i.e.*, $m \geq 162$ for $\lambda = 128$ and $m \geq 324$ for $\lambda = 256$.

MinRank. We then consider the MinRank attacks described in Section 4. We rely on the complexity estimate derived by Bardet *et al.* using the support-minors method [1],

$$\mathcal{O}\left(d'(n-1)^4 \binom{2d'+1}{d'}^2\right),$$

where d' denotes the rank of the associated MinRank instance. For our construction, it is shown in [9] that this rank is

$$d' = d + 1 + a + r$$

when attacking T.
Introducing an IP modifier has a surprisingly large impact on the complexity of the MinRank-S attack. As explained in [9], for $r > 1$ the corresponding MinRank instance satisfies

$$d' \geq \frac{n}{2}.$$

The resulting security estimates are summarized in Tables 7 and 8. For each target security level, we report both the proposed Dragon parameter sets and their conservative No-Dragon counterparts.

Table 7: Complexity of attacks for the 128-bit security parameter sets. Complexities are given in $\log_2$ bit operations. For Dragon instances, the Birthday column is the hash-collision bound $q^{n'/2}$.

Instance	(q, m, n, a, r)	Birthday	Hybrid	MinRank-T	MinRank-S
Dragon	$(2, 162, 189, 27, 2)$	128	128	162	≥ 544
No Dragon	$(2, 256, 283, 27, 2)$	128	203	164	≥ 596
Dragon	$(4, 84, 105, 21, 2)$	128	133	136	≥ 285
No Dragon	$(4, 128, 149, 21, 2)$	128	202	138	≥ 326
Dragon	$(5, 73, 94, 21, 2)$	128	141	131	≥ 215
No Dragon	$(5, 111, 132, 21, 2)$	128	208	133	≥ 293
Dragon	$(13, 56, 77, 21, 2)$	129	140	131	≥ 179
No Dragon	$(13, 70, 91, 21, 2)$	129	173	132	≥ 208
Dragon	$(23, 53, 74, 21, 2)$	128	144	132	≥ 176
No Dragon	$(23, 57, 78, 21, 2)$	128	153	132	≥ 184

Table 8: Complexity of attacks for the 256-bit security parameter sets. Complexities are given in $\log_2$ bit operations. For Dragon instances, the Birthday column is the hash-collision bound $q^{n'/2}$.

Instance	(q, m, n, a, r)	Birthday	Hybrid	MinRank-T	MinRank-S
Dragon	$(2, 324, 390, 66, 2)$	256	256	322	≥ 814
No Dragon	$(2, 512, 578, 66, 2)$	256	405	325	≥ 1193
Dragon	$(4, 170, 223, 53, 2)$	256	269	268	≥ 476
No Dragon	$(4, 256, 309, 53, 2)$	256	405	270	≥ 649
Dragon	$(5, 153, 206, 53, 2)$	256	282	264	≥ 444
No Dragon	$(5, 221, 274, 53, 2)$	256	400	266	≥ 582
Dragon	$(13, 118, 171, 53, 2)$	257	282	264	≥ 307
No Dragon	$(13, 139, 192, 53, 2)$	257	329	265	≥ 417
Dragon	$(23, 110, 163, 53, 2)$	257	285	265	≥ 357
No Dragon	$(23, 114, 167, 53, 2)$	257	293	265	≥ 365

We deliberately chose parameter sets with a slightly higher security level than required to provide a safety margin against future improvements in known attacks. As usual, the parameter sets would have to be re-evaluated in the event of a fundamentally new attack.

7 From GeMSS to James

7.1 Performance Evaluation

To implement our algorithms, we modified the code of GeMSS submitted to NIST[8] to support D-James and James. The resulting implementation is publicly available at

https://github.com/alex-on-the-internet/D-James

The resulting benchmarks are reported in Table 9. Since the original implementation is highly optimized for $q = 2$, we restrict this comparison to this case; supporting $q > 2$ efficiently would require rewriting the underlying finite-field arithmetic. The large difference in signing time between James and GeMSS is largely attributable to the degree of the central polynomial: our construction uses degree 5, compared to degree 513 for GeMSS. For a more meaningful comparison, we also report the performance of RedGeMSS, whose central polynomial has degree 17.

To further distinguish the cost of the construction from implementation overhead, we also performed a controlled comparison with the GeMSS implementation using matched values of n and the degree of the central polynomial. In this setting, the two implementations exhibit comparable running times across key generation, signing, and verification. We do not include these measurements in Table 9, since the corresponding parameter sets are not standard GeMSS instances; they are only intended to provide a comparison of the implementations under comparable computational conditions.

Table 9: Benchmark results for D-James and James compared to GeMSS. Results are reported for the 128-bit and 256-bit security levels ($q = 2$). The reported values are median running times over 100 independent executions, measured on an Intel(R) Core(TM) Ultra 5 235U (2.00 GHz) with 8 GB of RAM.

	KeyGen		**Signature**		**Verify**	
Instance	**128**	**256**	**128**	**256**	**128**	**256**
D-James	49 ms	947 ms	0.81 ms	4.36 ms	63 μs	1.36 ms
James	112 ms	2.2 s	1.91 ms	8.39 ms	59 μs	781 μs
GeMSS	165 ms	1.9 s	3.43 s	18.1 s	94 μs	865 μs
RedGeMSS	50 ms	648 ms	27.2 ms	172 ms	95 μs	843 μs

[8] https://polsys.team.lip6.fr/Links/NIST/GeMSS.html

Independent implementation. The correctness and performance of our implementation were further assessed through an independent implementation of JAMES and D-JAMES in Python and Rust, available at https://github.com/mjosaarinen/xdjames. This implementation supports both $q = 2$ and $q > 2$, and provides an independent implementation of the parameter sets described in Section 6.1. Table 10 reports its signing performance for several values of q. These measurements provide an independent confirmation of both the implementation and the strong dependence of signing performance on q.

Table 10: Independent signing benchmarks (Python/Rust implementation) for D-JAMES and JAMES at the 128-bit and 256-bit security levels.

q	D-JAMES		JAMES	
	Sign 128	**Sign 256**	**Sign 128**	**Sign 256**
2	133 ms	2.4 s	432 ms	4.5 s
4	1.4 s	26.1 s	2.3 s	41.4 s
5	1.6 s	46.9 s	4.0 s	40.8 s
13	5.2 s	271 s	284 s	486 s
23	377 s	1026 s	472 s	2273 s

7.2 Signature Size Comparison

We now compare the signature size of our schemes with those of other post-quantum signature schemes. For each scheme, we select the parameter set yielding the shortest signature, thereby prioritizing signature size over public-key size. The comparison is reported in Table 11.
Our parameter sets target 128-bit security, whereas the NIST Level 1 parameter sets are defined according to the NIST security categories. The comparison is therefore not necessarily made at exactly the same security target. Nevertheless, we expect differences of this magnitude in the security target to result in only a moderate increase in signature size.

Table 11: Comparison with other post-quantum signature schemes. Sizes are given in bytes (B); sizes in bits (b) are also given for the shortest signatures.

Scheme	Security	Signature	Public key
JAMES-short	128-bit	30 B (238 b)	1.29 MB
D-JAMES-short	128-bit	20 B (156 b)	1.35 MB
G*e*MSS	SL1	33 B (258 b)	352 kB
UOV	SL1	96 B	66.5 kB
MAYO	SL1	186 B	4.91 kB
SPHINCS+	SL1	7.85 kB	32 B
Dilithium	SL2	2.42 kB	1.31 kB
Falcon	SL1	666 B	897 B
SNOVA	SL1	124 B	9.84 kB
Classical baseline			
ECDSA	128-bit	64 B	33 B

As can be seen, D-JAMES achieves the shortest signatures among the schemes considered, with a signature size of only 156 bits. In particular, its signature is substantially shorter than those of all other post-quantum signature schemes considered, as well as the classical ECDSA baseline. This comes at the cost of a significantly larger public key, which represents the main trade-off of our construction.

Figure 2 compares the signature size and public-key size of the proposed scheme with those of several parameter sets from UOV, MAYO, and SNOVA. For completeness, we also include RSA as a reference point.

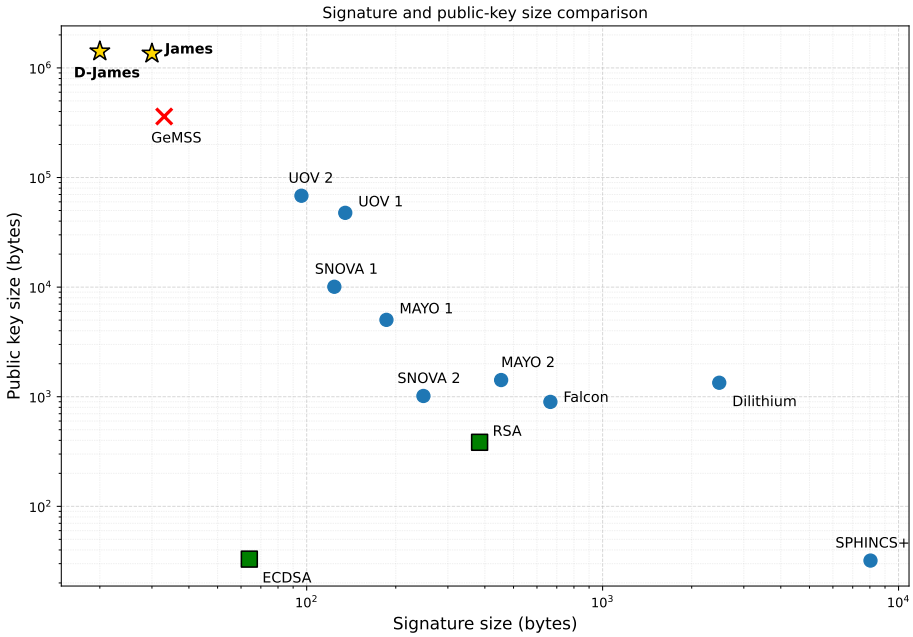

Fig. 2: Comparison of signature and public-key sizes.

We therefore do not claim that JAMES or D-JAMES provides a better overall trade-off than the other schemes considered. Rather, the purpose of this comparison is to demonstrate that our constructions achieve exceptionally short signatures, at the expense of larger public keys. This trade-off can be attractive in applications where signature size is the primary constraint.

8 Conclusion

The cryptanalysis of GeMSS, in particular through MinRank attacks, highlighted the difficulty of designing HFE-based signatures combining compact signatures and resistance against algebraic attacks. In contrast, UOV and its variants remain actively developed, as illustrated by several NIST submissions [6], [26], [3]. The recent results of [9] suggest that some HFE-based constructions can be repaired.

In this paper, we introduced new variants of $\mathrm{HFE}^{-}_{\mathrm{IP}}$, called JAMES and D-JAMES, which resist the known attacks considered in this work, including MinRank-S, while enabling extremely short signatures (for example, 156 bits at the 128-bit security level). They can be viewed as a relatively simple repair of GeMSS: the main modification is the introduction of the IP perturbation, which thwarts the MinRank-S attack.

Furthermore, we investigated the possibility of using Dragon-type public keys, which significantly reduce signature lengths. We also presented an analysis of differential attacks targeting the perturbation mechanism. The resulting post-quantum signature scheme, D-James, achieves competitive estimated performance and offers the shortest signatures among the public-key signature schemes considered, both classical and post-quantum, at the cost of a significantly larger public key. Related work explores the combination of Dragon terms with alternative perturbation mechanisms, also achieving ultra-short signatures [22].

Nevertheless, the numerous cryptanalytic advances that have emerged in multivariate cryptography over the years call for caution. Accordingly, we do not recommend deploying these signatures in security-critical applications until they have undergone further scrutiny and extensive analysis by the cryptographic community.

References

1. Bardet, M., Briaud, P., Bros, M., Gaborit, P., Tillich, J.: Revisiting algebraic attacks on MinRank and on the rank decoding problem. Des. Codes Cryptogr. **91**(11), 3671–3707 (2023). https://doi.org/10.1007/S10623-023-01265-X
2. Becker, A., Joux, A., May, A., Meurer, A.: Decoding random binary linear codes in $2^{n/20}$: How $1 + 1 = 0$ improves information set decoding. In: Pointcheval, D., Johansson, T. (eds.) Advances in Cryptology – EUROCRYPT 2012. Lecture Notes in Computer Science, vol. 7237, pp. 520–536. Springer (2012). https://doi.org/10.1007/978-3-642-29011-4_31
3. Benadjila, R., Bouillaguet, C., Feneuil, T., Rivain, M.: MQOM: MQ on my Mind Algorithm Specifications and Supporting Documentation (Version 2.0). Submission to the NIST Post-Quantum Cryptography Standardization Process (2025), https://mqom.org/index.html
4. Bettale, L., Faugère, J., Perret, L.: Hybrid approach for solving multivariate systems over finite fields. J. Math. Cryptol. **3**(3), 177–197 (2009). https://doi.org/10.1515/JMC.2009.009
5. Bettale, L., Faugère, J., Perret, L.: Cryptanalysis of Multivariate and Odd-Characteristic HFE Variants. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) Public Key Cryptography - PKC 2011 - 14th International Conference on Practice and Theory in Public Key Cryptography, Taormina, Italy, March 6-9, 2011. Proceedings. Lecture Notes in Computer Science, vol. 6571, pp. 441–458. Springer (2011). https://doi.org/10.1007/978-3-642-19379-8_27
6. Beullens, W.: MAYO: Practical Post-quantum Signatures from Oil-and-Vinegar Maps. In: AlTawy, R., Hülsing, A. (eds.) Selected Areas in Cryptography - 28th International Conference, SAC 2021, Virtual Event, September 29 - October 1, 2021, Revised Selected Papers. Lecture Notes in Computer Science, vol. 13203, pp. 355–376. Springer (2021). https://doi.org/10.1007/978-3-030-99277-4_17
7. Bouillaguet, C., Fouque, P., Macario-Rat, G.: Practical Key-Recovery for All Possible Parameters of SFLASH. In: Lee, D.H., Wang, X. (eds.) Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings. Lecture Notes in Computer Science, vol. 7073, pp. 667–685. Springer (2011). https://doi.org/10.1007/978-3-642-25385-0_36

8. Casanova, A., Faugère, J.C., Macario-Rat, G., Patarin, J., Perret, L., Ryckeghem, J.: GeMSS: A Great Multivariate Short Signature. Research report, UPMC - Paris 6 Sorbonne Universités ; INRIA Paris Research Centre, MAMBA Team, F-75012, Paris, France ; LIP6 - Laboratoire d'Informatique de Paris 6 (Dec 2017), https://inria.hal.science/hal-01662158
9. Cogliati, B., Macario-Rat, G., Patarin, J., Varjabedian, P.: State of the Art of HFE Variants - Is It Possible to Repair HFE with Appropriate Modifiers? In: Saarinen, M.O., Smith-Tone, D. (eds.) Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Oxford, UK, June 12-14, 2024, Proceedings, Part II. Lecture Notes in Computer Science, vol. 14772, pp. 144–167. Springer (2024). https://doi.org/10.1007/978-3-031-62746-0_7
10. Ding, J., Schmidt, D.: Cryptanalysis of HFEv and Internal Perturbation of HFE. In: Vaudenay, S. (ed.) Public Key Cryptography - PKC 2005, 8th International Workshop on Theory and Practice in Public Key Cryptography, Les Diablerets, Switzerland, January 23-26, 2005, Proceedings. Lecture Notes in Computer Science, vol. 3386, pp. 288–301. Springer (2005). https://doi.org/10.1007/978-3-540-30580-4_20
11. Dubois, V., Fouque, P., Shamir, A., Stern, J.: Practical Cryptanalysis of SFLASH. In: Menezes, A. (ed.) Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4622, pp. 1–12. Springer (2007). https://doi.org/10.1007/978-3-540-74143-5_1
12. Dubois, V., Granboulan, L., Stern, J.: Cryptanalysis of HFE with Internal Perturbation. In: Okamoto, T., Wang, X. (eds.) Public Key Cryptography - PKC 2007, 10th International Conference on Practice and Theory in Public-Key Cryptography, Beijing, China, April 16-20, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4450, pp. 249–265. Springer (2007). https://doi.org/10.1007/978-3-540-71677-8_17
13. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases without reduction to zero (F5). In: ISSAC '02. p. 75–83. Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/780506.780516
14. Faugère, J., Joux, A.: Algebraic Cryptanalysis of Hidden Field Equation (HFE) Cryptosystems Using Gröbner Bases. In: Boneh, D. (ed.) Advances in Cryptology - CRYPTO 2003, 23rd Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 2003, Proceedings. Lecture Notes in Computer Science, vol. 2729, pp. 44–60. Springer (2003). https://doi.org/10.1007/978-3-540-45146-4_3
15. Fusco, G., Bach, E.: Phase transition of multivariate polynomial systems. Math. Struct. Comput. Sci. **19**(1), 9–23 (2009). https://doi.org/10.1017/S0960129508007251
16. Joux, A., Vitse, V.: A crossbred algorithm for solving boolean polynomial systems. In: Kaczorowski, J., Pieprzyk, J., Pomykala, J. (eds.) Number-Theoretic Methods in Cryptology - First International Conference, NuTMiC 2017, Warsaw, Poland, September 11-13, 2017, Revised Selected Papers. Lecture Notes in Computer Science, vol. 10737, pp. 3–21. Springer (2017). https://doi.org/10.1007/978-3-319-76620-1_1
17. Kipnis, A., Shamir, A.: Cryptanalysis of the HFE Public Key Cryptosystem by Relinearization. In: Wiener, M.J. (ed.) Advances in Cryptology - CRYPTO '99, 19th Annual International Cryptology Conference, Santa Barbara, California, USA, Au-

gust 15-19, 1999, Proceedings. Lecture Notes in Computer Science, vol. 1666, pp. 19–30. Springer (1999). https://doi.org/10.1007/3-540-48405-1_2
18. Matsumoto, T., Imai, H.: Public Quadratic Polynomial-Tuples for Efficient Signature-Verification and Message-Encryption. In: Günther, C.G. (ed.) Advances in Cryptology - EUROCRYPT '88, Workshop on the Theory and Application of of Cryptographic Techniques, Davos, Switzerland, May 25-27, 1988, Proceedings. Lecture Notes in Computer Science, vol. 330, pp. 419–453. Springer (1988). https://doi.org/10.1007/3-540-45961-8_39
19. May, A., Meurer, A., Thomae, E.: Decoding random linear codes in $\tilde{\jmath}(2^{0.054n})$. In: Lee, D.H., Wang, X. (eds.) Advances in Cryptology – ASIACRYPT 2011. Lecture Notes in Computer Science, vol. 7073, pp. 107–124. Springer (2011). https://doi.org/10.1007/978-3-642-25385-0_6
20. Patarin, J.: Asymmetric Cryptography with a Hidden Monomial. In: Koblitz, N. (ed.) Advances in Cryptology - CRYPTO '96, 16th Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 1996, Proceedings. Lecture Notes in Computer Science, vol. 1109, pp. 45–60. Springer (1996). https://doi.org/10.1007/3-540-68697-5_4
21. Patarin, J., Macario-Rat, G., Bros, M., Koussa, E.: Ultra-Short Multivariate Public Key Signatures. IACR Cryptol. ePrint Arch. **2020**, 914 (2020)
22. Patarin, J., Vacek, J.: Ultra short signatures with Dragon $\mathrm{HFE}_{\mathrm{LL}}$'. IACR Cryptol. ePrint Arch. **2026**, 404 (2026), https://eprint.iacr.org/2026/404
23. Tao, C., Petzoldt, A., Ding, J.: Efficient Key Recovery for All HFE Signature Variants. In: Malkin, T., Peikert, C. (eds.) Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part I. Lecture Notes in Computer Science, vol. 12825, pp. 70–93. Springer (2021). https://doi.org/10.1007/978-3-030-84242-0_4
24. Vates, J., Smith-Tone, D.: Key Recovery Attack for All Parameters of HFE-. In: Lange, T., Takagi, T. (eds.) Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings. Lecture Notes in Computer Science, vol. 10346, pp. 272–288. Springer (2017). https://doi.org/10.1007/978-3-319-59879-6_16
25. Wagner, D.A.: A generalized birthday problem. In: Yung, M. (ed.) Advances in Cryptology - CRYPTO 2002, 22nd Annual International Cryptology Conference, Santa Barbara, California, USA, August 18-22, 2002, Proceedings. Lecture Notes in Computer Science, vol. 2442, pp. 288–303. Springer (2002). https://doi.org/10.1007/3-540-45708-9_19
26. Wang, L., Tseng, P., Kuan, Y., Chou, C.: A simple noncommutative UOV scheme. IACR Cryptol. ePrint Arch. **2022**, 1742 (2022), https://eprint.iacr.org/2022/1742