Geometric Forgeries: Structural Cryptanalysis of MAYO

Massimo Ostuzzi

Ruhr University Bochum, Bochum, Germany

Abstract. MAYO is a post-quantum signature scheme based on UOV that uses the *whipping* transformation to achieve more compact public keys and retain efficient signing. We present the first structural forgery attack that exploits the geometry of the resulting multivariate quadratic (MQ) maps. Our starting point is the observation that whipping amplifies polar cancellations in the underlying MQ map, producing subspaces with low-dimensional quadratic and polar images. We combine this geometric insight with the Pseudo-Oil Attack (ASIACRYPT 2026) and a descent to the binary field into a structural forgery algorithm.
Independently, we introduce a multi-target reduction for homogeneous MQ maps. Projecting onto a quotient of the codomain reduces the number of equations, and homogeneity allows suitable solutions to be rescaled to the original targets.
Concretely, the combination of our forgery attacks lowers the estimated bit security of the Security Level I and III parameter sets of MAYO Round 2 by roughly 18 and 16 bits, respectively. In particular, the Level I parameter set falls 15 bits short of its required security level. For the recently released MAYO Round 3 parameters, our attack still reduces the estimated security by roughly 6, 4 and 2 bits for the new MAYO1, MAYO2 and MAYO3 respectively, bringing MAYO2 below the required security by 3 bits.

1 Introduction

The transition to post-quantum cryptography has turned concrete security from an asymptotic concern into an immediate design constraint. This is particularly visible in NIST's ongoing standardisation process for additional post-quantum signatures, where nine candidates have recently advanced to a third evaluation round. The central goal of this process is to diversify the mathematical assumptions on which future signature standards may rely [ABC+26].

Public key constructions usually derive their security from the presumed hardness of well-established mathematical problems. Yet, practical schemes are rarely built to rely on pristine random instances alone, and designers typically introduce additional algebraic structure to obtain more compact keys and more efficient computations. Accordingly, the security of those schemes rests on *structured variants* of hard mathematical problems. Prominent examples are the structured variants of LWE and SIS [Ajt96,Reg05,LPR10,LS15], called module LWE and module SIS, underlying the security of the standardised ML-KEM

[BDK+18,Nat24b] and ML-DSA [DKL+18,Nat24a]. This structure is a double-edged sword: it is, at the same time, what makes the scheme really practical, and what potentially creates new attack surface for cryptanalysis. Consequently, a meaningful security analysis must answer two separate questions: how hard is the underlying generic problem, and how much of that hardness survives in the structured instance used by the scheme?

MQ Signatures. Multivariate cryptography is another particularly clear example of this distinction. The central computationally hard problem is the *MQ problem* [FY79,HLY12], i.e. inverting a multivariate quadratic (MQ) map $\mathcal{P} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ on a codomain element $\mathbf{t} \in \mathbb{F}_q^m$, called *target*. Introduced by Kipnis, Patarin and Gouben [KPG99], Unbalanced Oil and Vinegar (UOV) uses a secret linear subspace O, called *oil subspace*, of dimension o such that $\mathcal{P}(O) = 0$. The point is that on O inverting the MQ map reduces to performing Gaussian elimination on a linear system of m equations in o variables. To ensure that a solution to this linear system exists with high probability and thus enable efficient signing, the size of the oil subspace must be at least as large as the number of equations, i.e. $o \geq m$. In practice, this prevents UOV's public key to be compact enough for many real world applications.

MAYO. In 2021, Beullens proposed a solution to this problem by introducing MAYO [Beu22,BCC+24], whose secret oil subspace has dimension significantly smaller than m. To retain the efficiency in signing, one then *whips up* the UOV MQ map $\mathcal{P}$ into an MQ map $\mathcal{P}^* : \mathbb{F}_q^\kappa \otimes \mathbb{F}_q^n \simeq \mathbb{F}_q^{n\kappa} \to \mathbb{F}_q^m$ for which $\mathbb{F}_q^\kappa \otimes O \simeq O^\kappa$ is an oil subspace. The clever insight is that the dimension of O^κ grows by a factor of κ to κo, while the dimension of the codomain remains m. In turn, the new condition for efficient signing is relaxed to $\kappa o \geq m$, and the corresponding public keys benefit from a significant size reduction.

Forgery Attacks. Forging a signature for MAYO reduces to finding a preimage under $\mathcal{P}^*$ of a random target $\mathbf{t} \in \mathbb{F}_q^m$. As reported in both the Round 2 [BCC+24] and the recent Round 3 [BCC+26] specifications submitted to the NIST competition, there is no known attack exploiting the structure of the MQ map $\mathcal{P}^*$ produced by the whipping up technique, and, consequently, the security of the scheme is estimated by considering generic attacks only. In particular, the most efficient forgery attack consists of a generic technique for inverting underdetermined MQ maps, leveraging the abundance of variables [TW12,Has21,Ost26,ACF+26]. Recent work [Ost26] in this line of research highlighted that these techniques actually consist in computing some surrogates of the oil subspace, called *pseudo-oil* subspaces. These can be computed more efficiently than recovering the actual oil subspace (i.e. than performing a key recovery attack), but, instead of allowing to invert the MQ map in fully polynomial time, they only provide a reduction in the (still exponential) forgery complexity. We refer to this forgery attack as *Pseudo-Oil Attack*.

In [Ost26], the author leaves as an interesting open problem whether the structure of the whipped up MQ map of MAYO exposes pseudo-oil structures that could accelerate forgeries. We answer this open question in the affirmative.

1.1 Technical Overview

For an MQ map $\mathcal{P}$, let us denote by $\mathcal{P}'$ its polar map. The most naive strategy for inverting $\mathcal{P}^*$ is fixing $n\kappa - m$ variables and run a Gröbner Basis algorithm on the resulting square system of m quadratic equations in m variables.

From Geometry To Forgery. The starting point of our improved forgery algorithm is an elementary observation. For any vector $\mathbf{u} \in \mathbb{F}_q^n$, consider the differential at $\mathbf{u}$

$$D_{\mathbf{u}} : \mathbf{x} \mapsto \mathcal{P}'(\mathbf{u}, \mathbf{x}).$$

The whipping transformation of MAYO creates the pair of subspaces

$$U = \mathbb{F}_q^\kappa \otimes \langle \mathbf{u} \rangle \simeq \langle \mathbf{u} \rangle^\kappa, \qquad \mathcal{O} = \mathbb{F}_q^\kappa \otimes \ker D_{\mathbf{u}} \simeq (\ker D_{\mathbf{u}})^\kappa,$$

which satisfy $\mathcal{P}^{*\prime}(U, \mathcal{O}) = 0$. Moreover, since polar maps are alternating in even characteristic, the span W of $\mathcal{P}^*(U)$ has unusually small dimension, at most κ. These geometric properties imply that, for $\mathbf{u} \in U$ and $\mathbf{v} \in \mathcal{O}$, we have

$$\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathcal{P}^*(\mathbf{u}) + \mathcal{P}^*(\mathbf{v}).$$

Consequently, reducing modulo W removes the U component from the inversion problem on $\mathcal{P}^*$. Once a solution $\mathbf{v} \in \mathcal{O}$ has been found, the residual target belongs to W and can be reached from U by linear algebra and coordinate-wise square roots. Therefore, inverting $\mathcal{P}^*$ restricted to $\mathcal{O}$ naturally splits into inverting a smaller MQ map and a linear system.

This observation can be seen in the broad context of the Pseudo-Oil Attack: U and $\mathcal{O}$ supplied by the whipping up are pseudo-oil subspaces in the sense of [Ost26], i.e. they are not as powerful as the real oil subspace O^κ over which inverting the whipped MQ map is fully polynomial time, but still they allow to reorganise the inversion problem into smaller algebraic systems. In this paper, we refine and expand this observation, and, by combining it with ideas from the Pseudo-Oil Attack, we turn geometric insights into structural forgery attacks that outperforms all generic techniques.

Multiple Targets. In addition, we introduce a simple, yet effective, generic multi-target technique. Indeed, MAYO is based on the *salted* hash-and-sign paradigm [DH76], meaning that an adversary can generate many instances of the MQ problem for the same message M by varying the salt in the hash function. Consequently, the security of MAYO reduces to a multi-target MQ inversion problem.

More generally, for a homogeneous quadratic map $\mathcal{P}$, d linearly independent targets span a subspace T that can be quotiented out from the codomain of $\mathcal{P}$,

leaving with an MQ map with $m - d$. By construction, a nonzero preimage of $0 \in \mathbb{F}_q^m$ through the projected MQ map is sent into the target span T: if its image lies on one of the punctured lines spanned by the d targets, homogeneity of $\mathcal{P}$ lets the attacker rescale the projected preimage to a preimage of that target. This gives an explicit tradeoff between a cheaper projected solve, i.e. with less equations, and a nonzero failure probability.

1.2 Our Contributions

We make the following contributions.

1. Starting from a geometric observation on whipped MQ maps, we present a structural forgery attack that reduces the size of the MQ problem by κ both in variables and equations.
2. Building on the first attack, we develop a more sophisticated structural attack and analyze two settings arising from different choices of *emulsifier matrices*, *structured* as in Round 2 and *unstructured*. We turn this additional geometry into a forgery attack by combining it with the Pseudo-Oil Attack and a descent to $\mathbb{F}_2$, which, to the best of our knowledge, is novel in this setting.
3. Independently of these structural attacks, we introduce a solver-agnostic multi-target reduction for homogeneous MQ maps. Despite its simplicity, this observation appears to have been overlooked in the MQ cryptanalysis literature and can be applied independently of the underlying MQ solver.
4. We then show that the our multi-target technique can be combined seamlessly with our structural forgery attack.
5. We optimize the resulting combined attack for the Round 2 MAYO parameter sets. Most notably, our forgery attack yields speedups by a factor of 2^{18} over the state-of-the-art generic attacks for the Security Level I parameter set, bringing it to 128 bits of security only, i.e. 15 bits below the required security level, see Table 1.

Table 1: Forging a Round 2 MAYO signature. Parameters below the required security level in red. The 113 bits for MAYO2 is the effect of the wedge attack [Ran26].

Security	Generic	This work
MAYO1	146	**128**
MAYO2	113	147
MAYO3	223	**207**
MAYO5	295	**293**

During the preparation of this manuscript, the MAYO team released the specifications for the Round 3 [BCC+26] of the NIST competition. One important modification of the scheme is the introduction of linear contributions in the definition of the whipped MQ map $\mathcal{P}^*$, making it inhomogeneous. This apparently small addition secures MAYO against our attack ideas. More precisely, the new linear terms increase the dimension of the output subspaces reducing the impact of our structural insight, and the inhomogeneity prevents the multi-target approach completely. We thus have the following additional contribution:

6. Applying our single-target structural forgery attack to Round 3 parameter sets, we reduce the bit security of MAYO1, MAYO2 and MAYO3 by $6, 4$ and 2 bits respectively, see Table 2. This brings MAYO2 below the required 143 bits of security.

Table 2: Forging a Round 3 MAYO signature. Parameters below the required security level in red.

Security	Generic	This work
MAYO1	153	**147**
MAYO2	144	**140**
MAYO3	223	**221**
MAYO5	295	314

Moreover, this work confirms that the choices operated for Round 3 are actually extremely beneficial for MAYO, boosting the scheme's security in a non-negligible way. To quantify this security gain, we will also estimate the cost of our attack against the new MAYO parameter sets as if there was no linear term and the same emulsifier choice as in Round 2.

AI Disclosure. We used ChatGPT Sol 5.6 Ultra to improve the fluency of some small portion of this manuscript. Moreover, we used the aforementioned model to improve our code, which was initially wrote by the author without the assitance of LLMs. The code can be found in the git repository `https://github.com/maxostuzzi/forging_mayo`. The authors take full responsability of the content of the manuscript and code.

Organisation. Section 2 recalls the prerequisites to understand our algorithms, such as MQ maps, polar forms, and the construction of MAYO's public keys. Section 3 presents and analyses the first simple forgery attack. Section 4 refines the geometric observation of Section 3, handling the two different emulsifier regimes. Section 5 explains and analyses how to turn the our new geometric insight into a forgery attack. Section 6 gives the multi-target reduction and its expected running time analysis. Finally, Section 7 concludes with the optimised concrete security estimates, both for Round 2 and Round 3.

2 Preliminaries

In this section, we recall useful concepts and conventions. In this work, we restrict ourselves to fields of even characteristic, since these are the ones used in MAYO.

Linear Algebra. Let V and W be vector spaces over $\mathbb{F}_q$. We use the notation $V \simeq W$ to indicate that V and W are isomorphic as vector spaces. We denote by 0 the zero vector, and the ambient vector space will be clear from the context. Let $\mathbf{v}_1, \ldots, \mathbf{v}_k \in V$. We denote by $\langle \mathbf{v}_1, \ldots, \mathbf{v}_k \rangle$ the linear subspace of V spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_k$, and, more generally, for a subset $S \subseteq V$, $\langle S \rangle$ is the linear subspace generated by S. If $W \subseteq V$, the quotient space V/W is defined as the set of cosets

$$V/W := \{\mathbf{v} + W : \mathbf{v} \in V\}.$$

Throughout the paper, we will often identify a coset $\mathbf{v} + W \in V/W$ with one of its representatives $\mathbf{v} \in V$, meaning that one can take any lift of $\mathbf{v} + W$ to V. We denote by $V \oplus W$ the direct sum of V and W.

All tensor products are taken over $\mathbb{F}_q$. Operationally, for the canonical basis $\mathbf{e}_1, \ldots, \mathbf{e}_k$ of $\mathbb{F}_q^k$ and any linear subspace $S \subseteq V$, we identify $\mathbb{F}_q^k \otimes S \simeq S^k$ via the isomorphism

$$\sum_{i=1}^{k} \mathbf{e}_i \otimes \mathbf{s}_i \mapsto (\mathbf{s}_1, \ldots, \mathbf{s}_k).$$

Thus $\mathbf{e}_i \otimes \mathbf{s}$ places $\mathbf{s}$ in the i-th block and zero in all other blocks.

When splitting the k blocks into groups of sizes a and $k-a$, we embed $\mathbb{F}_q^a$ and $\mathbb{F}_q^{k-a}$ into $\mathbb{F}_q^k$ as the first a and the last $k-a$ coordinate positions, respectively. Consequently, for subspaces $S, T \subseteq V$, we write

$$(\mathbb{F}_q^a \otimes S) \oplus (\mathbb{F}_q^{k-a} \otimes T) \simeq S^a \times T^{k-a}.$$

Whipped MQ Maps. Unless specified otherwise, we restrict our attention to *homogeneous* quadratic polynomials, i.e. polynomials in which every monomial has total degree two. These are the ones used in the Round 2 specifications of MAYO.

Let us fix positive integers $n, m \in \mathbb{Z}$. Let $V \subseteq \mathbb{F}_q^n$ and $W \subseteq \mathbb{F}_q^m$ be linear subspaces. A *multivariate quadratic map* (MQ map) is a function

$$\mathcal{P} : V \to W, \quad \mathbf{x} \mapsto (p_1(\mathbf{x}), \ldots, p_m(\mathbf{x}))^t,$$

where each p_i is a homogeneous quadratic polynomial in the variables $\mathbf{x} = (x_1, \ldots, x_n)$. The *polar map* $\mathcal{P}' : V \times V \to W$ associated to $\mathcal{P}$ is a bilinear map defined as

$$\mathcal{P}'(\mathbf{x}, \mathbf{y}) = \mathcal{P}(\mathbf{x} + \mathbf{y}) - \mathcal{P}(\mathbf{x}) - \mathcal{P}(\mathbf{y}).$$

We call $\mathcal{P}'(\mathbf{x}, \mathbf{y})$ the *polar pairing* of $\mathbf{x}$ and $\mathbf{y}$. Moreover, we say that $\mathbf{x}$ (resp. $X \subseteq V$ subspace) and $\mathbf{y}$ (resp. $Y \subseteq V$ subspace) are polar orthogonal if $\mathcal{P}'(\mathbf{x}, \mathbf{y}) = 0$ (resp. $\mathcal{P}'(X, Y) = 0$).

Lemma 1. *In even characteristic, all vectors are polar self-orthogonal.*

Proof. For all $\mathbf{x} \in V$, we have

$$\mathcal{P}'(\mathbf{x}, \mathbf{x}) = \mathcal{P}(2\mathbf{x}) - 2\mathcal{P}(\mathbf{x}) = 0.$$

□

Given an MQ map $\mathcal{P}$ as above, an integer $\kappa > 0$ and invertible matrices $E_i \in \mathrm{GL}_m(\mathbb{F}_q)$ for $1 \leq i \leq \kappa$ and $E_{ij} \in \mathrm{GL}_m(\mathbb{F}_q)$ for $1 \leq i < j \leq \kappa$, the corresponding *Whipped* MQ (WMQ) map is the MQ map $\mathcal{P}^* : \mathbb{F}_q^\kappa \otimes V \simeq V^\kappa \to W$ defined as

$$\mathcal{P}^*(\mathbf{x}_1, \ldots, \mathbf{x}_\kappa) = \sum_{i=1}^{\kappa} E_i \mathcal{P}(\mathbf{x}_i) + \sum_{1 \leq i < j \leq \kappa} E_{ij} \mathcal{P}'(\mathbf{x}_i, \mathbf{x}_j).$$

We call the integer κ *whipping factor*.

MQ Problem. Let $\mathcal{P} : V \to W$ be an MQ map, where $\dim V = n$ and $\dim W = m$. An instance of the $\mathsf{MQ}_q(n, m)$ *problem* consists of the tuple $(\mathcal{P}, \mathbf{t})$, with $\mathbf{t} \in W$ called *target*. The vector $\mathbf{s} \in V$ is a solution of the instance if $\mathcal{P}(\mathbf{s}) = \mathbf{t}$, i.e. $\mathbf{s}$ is a preimage of $\mathbf{t}$. The $\mathsf{MQ}_q(n, m)$ problem is called *overdetermined*, *square* or *underdetermined* when $n < m$, $n = m$ or $n > m$, respectively. We write $\mathsf{GB}_q(n, m)$ for the bit complexity of solving an instance of the $\mathsf{MQ}_q(n, m)$ problem in expectation using Gröbner Basis algorithms.

Forging MAYO Signatures. MAYO is a salted hash-and-sign signature scheme. The public key consists of a WMQ map $\mathcal{P}^* : \mathbb{F}_q^{n\kappa} \to \mathbb{F}_q^m$. Forging a MAYO signature for a message m amounts to solving the one-out-of-many $\mathsf{MQ}_q(n\kappa, m)$ problems with inputs $\{(\mathcal{P}^*, \mathbf{t}_i)\}_i$, where the $\mathbf{t}_i \in \mathbb{F}_q^m$ come from hashing the message m together with different salts. More precisely, a successful forgery consists of a vector $\mathbf{s} \in \mathbb{F}_q^{n\kappa}$ and an index j such that $\mathcal{P}^*(\mathbf{s}) = \mathbf{t}_j$. We call the single-target variant of this problem *WMQ problem*.

3 A First Simple Attack

Let $\mathcal{P} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ be a MQ map and $\kappa > 0$ the whipping factor. For convenience, we denote by $N = n\kappa$ the dimension of the domain of the WMQ map $\mathcal{P}^* : \mathbb{F}_q^N \to \mathbb{F}_q^m$, as defined in Section 2. In this section, we present a simple forgery attack that leverages the structure of the WMQ map. The main idea is that a polar orthogonality condition on the base MQ map $\mathcal{P}$ produces polar cancellations on the WMQ map $\mathcal{P}^*$.

Whipping Up a Polar Condition. Let $\mathbf{u}_1 \in \mathbb{F}_q^n$ be a random nonzero vector and define the *differential* at $\mathbf{u}_1$ as

$$D_{\mathbf{u}_1} : \mathbb{F}_q^n \to \mathbb{F}_q^m, \quad \mathbf{x} \mapsto \mathcal{P}'(\mathbf{u}_1, \mathbf{x}).$$

The map $D_{\mathbf{u}_1}$ is linear of rank $\operatorname{rk} D_{\mathbf{u}_1} \leq m$. Define $U = \mathbb{F}_q^\kappa \otimes \langle \mathbf{u}_1 \rangle$ and $\mathcal{O} = \mathbb{F}_q^\kappa \otimes \ker D_{\mathbf{u}_1}$ of dimensions κ and at least $\kappa \cdot \max\{1, n-m\}$ respectively, since $\mathbf{u}_1 \in \ker D_{\mathbf{u}_1}$. For the same reason, note that $U \subseteq \mathcal{O}$.

Lemma 2. *The subspace U is polar orthogonal to $\mathcal{O}$ with respect to $\mathcal{P}^{*\prime}$, i.e. $\mathcal{P}^{*\prime}(U, \mathcal{O}) = 0$. Moreover, let $W = \langle \operatorname{im} \mathcal{P}^*_{|U} \rangle$ be the span of the image of U through $\mathcal{P}^*$. Then $\dim W \leq \kappa$.*

Proof. The orthogonality of U and $\mathcal{O}$ follows directly from the definition of the whipped up MQ map $\mathcal{P}^*$ and the construction of U and $\mathcal{O}$. Regarding the dimension of W, by Lemma 1, the subspace W is spanned only by $E_i \mathcal{P}(\mathbf{u}_1)$, for $i \leq \kappa$, and the claim follows. □

Delivering the Attack. For simplicity, here we assume that $\dim W = \kappa$, which happens with high probability. A decrease in the dimension of W does not alter the effectiveness of the attack.

Assume that $\dim \mathcal{O} \geq m$, which is true for (almost) all parameter sets of MAYO. Let $\mathcal{O}_m$ be a m-dimensional subspace of $\mathcal{O}$ containing U, and restrict the MQ map $\mathcal{P}^*$ to such $\mathcal{O}_m$. Moreover, let V be a complement of U in $\mathcal{O}_m$. Recall that forging a MAYO signature reduces to solving the WMQ problem. Therefore, let $\mathbf{t} \in \mathbb{F}_q^m$ be a random target. The attack consists of two steps, which we call *Reduced* and *Finishing* steps. The goal is to find $\mathbf{u} \in U$ and $\mathbf{v} \in V$ such that $\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathbf{t}$.

The attack is based on the identity

$$\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathcal{P}^*(\mathbf{u}) + \mathcal{P}^*(\mathbf{v}), \tag{1}$$

which is missing the polar pairing of $\mathbf{u}$ and $\mathbf{v}$ because $V \subseteq \mathcal{O}$, see Lemma 2.

Consider the original MQ problem modulo W: this yields the $\mathsf{MQ}_q(m - \dim U, m - \dim W)$ problem

$$\mathcal{P}^*(\mathbf{v}) = \mathbf{t} \mod W, \quad \text{for } \mathbf{v} \in V.$$

The Reduced step consists in solving this MQ problem using a Gröbner Basis algorithm, thereby fixing a vector $\mathbf{v} \in V$.

The second step consists in solving the MQ problem

$$\mathcal{P}^*(\mathbf{u}) = \mathbf{t} - \mathcal{P}^*(\mathbf{v}), \quad \text{for } \mathbf{u} \in U. \tag{2}$$

Note that both the left and right side of the equation lie in W. Moreover, let us write $\mathbf{u} = \sum_i \alpha_i \mathbf{u}_1 \otimes \mathbf{e}_i$, where $\mathbf{e}_i$ is the i-th standard basis vector of $\mathbb{F}_q^\kappa$ and $\alpha_i \in \mathbb{F}_q$. Then, by Lemma 2, we have that

$$\mathcal{P}^*(\mathbf{u}) = \sum_{i=1}^{\kappa} \alpha_i^2 E_i \mathcal{P}(\mathbf{u}_1).$$

All the $E_i\mathcal{P}(\mathbf{u}_1) \in \mathbb{F}_q^m$ is a priori linearly independent. Therefore, solving the MQ problem in Equation (2) reduces to solving a linear system of equations and taking modular square-roots, which can be done efficiently. Let $\mathbf{u} \in U$ be a solution, and let us define the vector $\mathbf{s} = \mathbf{u} + \mathbf{v}$.

Proposition 1. *The vector $\mathbf{s} \in \mathbb{F}_q^N$ is a solution of the initial WMQ problem with input target $\mathbf{t}$.*

Proof. The claim trivially follows from the idenitity in Equation (1) and the fact that we have solved for $\mathbf{u} \in U$ satisfying Equation (2). □

This immediately implies the following corollary.

Corollary 1. *There exists a forgery attack for a MAYO public key with parameters (n, m, q, κ) that runs in expected time*

$$\mathsf{GB}_q(m - \kappa, m - \kappa),$$

ignoring the cost of solving the linear system in Equation (2).

Proof. One only needs to show that the solving strategy described above produces at least one solution in expectation. This is the case, since both steps amount to solve square systems of equations, admitting exactly one solution in expectation. In other words, in expectation there exists one $\mathbf{v} \in V$ that solves the first MQ problem, and one $\mathbf{u} \in U$ that solves the linear system. Therefore there exists one $\mathbf{s} \in \mathcal{O}_m$ solving the original WMQ problem, in expectation, which is a valid MAYO signature.

Due to its rigid dependence on the input parameter κ, this attack does not typically outperform existing approaches. Nevertheless, it exposes how the structure of the whipped up MQ map can be leveraged algorithmically to forge MAYO signatures. In particular, increasing the dimension of U would yield more efficient forgeries.

4 Enlarging the Subspace U

The attack from Section 3 could not beat generic forgery attacks because the dimension of U is not large enough. In this section, we show that it is possible to increase the dimension of U at the cost of a more complicated attack. We distinguish two different regimes, depending on the choice of the emulsifier matrices E_{ij}.

4.1 Unstructured Emulsifiers

We start with the *unstructured regime*: we assume that the emulsifier matrices E_{ij} are chosen such that the probability of nonzero low-rank linear combination is negligible, but without any additional structure, as suggested in [Beu22].

Let $\mathbf{u}_1 \in \mathbb{F}_q^n$ be any nonzero vector. We sample $\mathbf{u}_2 \in \ker D_{\mathbf{u}_1}$ randomly, such that $\mathbf{u}_2$ is linear independent with $\mathbf{u}_1$.

Remark 1. It would be tempting to define the vector space $\mathcal{O} = \mathbb{F}_q^\kappa \otimes \ker D_{\mathbf{u}_1,\mathbf{u}_2}$ and perform the attack of Section 3, where $D_{\mathbf{u}_1,\mathbf{u}_2} : \mathbb{F}_q^n \to (\mathbb{F}_q^m)^2$ be the *joint differential at* $\mathbf{u}_1$ *and* $\mathbf{u}_2$, i.e. the linear map defined as $D_{\mathbf{u}_1,\mathbf{u}_2}(\mathbf{x}) = (D_{\mathbf{u}_1}(\mathbf{x}), D_{\mathbf{u}_2}(\mathbf{x}))$. However, the dimension of $\mathcal{O}$ is generically $\max\{2\kappa, (n - 2m)\kappa\}$, which is $2\kappa < m$ for all MAYO's official parameter sets, meaning that $U = \mathcal{O}$ and the attack cannot be performed.

Let $\mathbf{u}_3 \in \ker D_{\mathbf{u}_2}$ be linearly independent from $\mathbf{u}_1$ and $\mathbf{u}_2$, and, for $0 \leq a \leq \kappa$, define the linear subspace

$$U = \mathbb{F}_q^a \otimes \langle \mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3 \rangle + \mathbb{F}_q^{(\kappa-a)} \otimes \langle \mathbf{u}_1, \mathbf{u}_2 \rangle \subseteq \mathbb{F}_q^{n\kappa},$$

which has dimension $u = 2\kappa + a$ by construction.

Let us now define $W = \langle \operatorname{im} \mathcal{P}^*_{|U} \rangle \subseteq \mathbb{F}_q^m$ and $R = \langle \operatorname{im} \mathcal{P}'^*_{|U\times U} \rangle \subseteq \mathbb{F}_q^m$, and denote by $w = \dim W$ and $r = \dim R$ their dimensions. By definition of the polar map, we have that $R \subseteq W$. In particular $r \leq w$.

Lemma 3. *It holds*

$$r \leq a + \binom{a}{2} + a(\kappa - a), \quad w \leq 2\kappa + 2a + \binom{a}{2} + a(\kappa - a).$$

Proof. Recall that, by Lemma 1, in even characteristic we have $\mathcal{P}'(\mathbf{x}, \mathbf{x}) = 0$ for all $\mathbf{x} \in \mathbb{F}_q^n$. Moreover, all the polar pairing combinations of $\mathbf{u}_1$ with $\mathbf{u}_2$ and $\mathbf{u}_2$ with $\mathbf{u}_3$ are zero, by construction. The only surviving ones are the pairings of $\mathbf{u}_1$ with $\mathbf{u}_3$. To count the number of these contributions, it is enough to count the number of ordered pairs $(i, j) \in \{1, \ldots, \kappa\} \times \{1, \ldots, a\}$ and subtract the repetition coming from the fact that in the definition of the WMQ map there is no term coupled with the matrix E_{ij} with $i > j$. Therefore, we have

$$r \leq \kappa a - \binom{a}{2} = a + \binom{a}{2} + a(\kappa - a).$$

Let now $\mathbf{v}_1, \ldots, \mathbf{v}_u$ be a basis of U. By definition of the polar map, we have

$$\langle \operatorname{im} \mathcal{P}^*_{|U} \rangle = \langle \mathcal{P}^*(\mathbf{v}_1), \ldots, \mathcal{P}^*(\mathbf{v}_u) \rangle + R.$$

Thus $w \leq u + r$, from which the claim follows. □

4.2 Structured Emulsifiers

We now consider the *structured regime*, corresponding to the choice of the emulsifier matrices in the Round 2 specifications of MAYO [BCC+24]. The goal is to use the structure of the emulsifier matrices to enlarge the subspace U while keeping the image subspace W small. We start by recalling the emulsifier schedule proposed in the Round 2 specifications of MAYO. First, identify $\mathbb{F}_q^m$ with the extension field $\mathbb{K} = \mathbb{F}_q[z]/(f(z))$, where f is an irreducible polynomial of degree m, and let ζ denote the residue class of z. We write $\varepsilon_{ij} \in \mathbb{K}$ for the field element

$\varepsilon_{1,1} = \zeta^{\kappa-1}$					
$\varepsilon_{1,2} = \zeta^{\kappa-2}$	$\varepsilon_{2,2} = \zeta^{2\kappa-2}$				
$\varepsilon_{1,3} = \zeta^{\kappa-3}$	$\varepsilon_{2,3} = \zeta^{2\kappa-3}$	$\varepsilon_{3,3} = \zeta^{3\kappa-4}$			
$\vdots$	$\vdots$	$\vdots$	$\ddots$		
$\varepsilon_{1,\kappa-1} = \zeta$	$\varepsilon_{2,\kappa-1} = \zeta^{\kappa+1}$	$\varepsilon_{3,\kappa-1} = \zeta^{2\kappa}$	$\cdots$	$\varepsilon_{\kappa-1,\kappa-1} = \zeta^{\binom{\kappa+1}{2}-2}$	
$\varepsilon_{1,\kappa} = 1$	$\varepsilon_{2,\kappa} = \zeta^{\kappa}$	$\varepsilon_{3,\kappa} = \zeta^{2\kappa-1}$	$\cdots$	$\varepsilon_{\kappa-1,\kappa} = \zeta^{\binom{\kappa+1}{2}-3}$	$\varepsilon_{\kappa,\kappa} = \zeta^{\binom{\kappa+1}{2}-1}$

Table 3: Emulsifier schedule in the Round 2 specifications of MAYO.

whose multiplication matrix is E_{ij}. Then, we have the structure summarised in Table 3.

In particular, note from Table 3 that row i equals row $i+1$ multiplied by ζ, except for the last missing column.

$$\begin{aligned} &\varepsilon_{\kappa\kappa} = \zeta^2 \varepsilon_{\kappa-1,\kappa}, \quad \text{and} \\ &\varepsilon_{i,\kappa-1} = \zeta \varepsilon_{i,\kappa}, \quad \text{for } 1 \le i \le \kappa - 1. \end{aligned} \tag{3}$$

As in the unstructured regime, choose linearly independent vectors $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{F}_q^n$ satisfying

$$\mathcal{P}'(\mathbf{u}_1, \mathbf{u}_2) = 0.$$

Consider their joint differential

$$D_{\mathbf{u}_1,\mathbf{u}_2} : \mathbb{F}_q^n \to \mathbb{K}^2, \qquad \mathbf{v} \mapsto \big(\mathcal{P}'(\mathbf{u}_1, \mathbf{v}), \mathcal{P}'(\mathbf{u}_2, \mathbf{v})\big).$$

The image $S = \operatorname{im} D_{\mathbf{u}_1,\mathbf{u}_2} \subseteq \mathbb{K}^2$ generically has dimension $n-2$, since $\mathbf{u}_1, \mathbf{u}_2 \in \ker D_{\mathbf{u}_1,\mathbf{u}_2}$. If this is not the case, we can resample $\mathbf{u}_1$ and $\mathbf{u}_2$. In the following, we use the convention that multiplication by ζ acts component-wise on $\mathbb{K}^2$. Consequently, define

$$I = S \cap \zeta^{-1} S = \{\mathbf{y} \in S : \zeta \mathbf{y} \in S\} \subseteq \mathbb{K}^2.$$

By Grassmann's inequality in the $2m$-dimensional $\mathbb{F}_q$-space $\mathbb{K}^2$, we obtain a lower bound on the dimension of I:

$$\dim I \ge 2(n-2) - 2m = 2(n-m) - 4.$$

Let $J \subseteq I$ be an a-dimensional subspace. The free parameter $a \le 2(n-m)-4$ determines how much we want enlarge the special subspace U from Section 3.

Since for $\mathbf{y} \in I$ both $\mathbf{y}$ and $\zeta\mathbf{y}$ belong to S by definition, we can choose $\mathbb{F}_q$-linear maps

$$A_0, A_1 : J \to \mathbb{F}_q^n$$

satisfying

$$D_{\mathbf{u}_1,\mathbf{u}_2}(A_0(\mathbf{y})) = \mathbf{y}, \qquad D_{\mathbf{u}_1,\mathbf{u}_2}(A_1(\mathbf{y})) = \zeta\mathbf{y}. \tag{4}$$

Lemma 4. *For all $\mathbf{x} \in \langle \mathbf{u}_1, \mathbf{u}_2 \rangle$, it holds*

$$\zeta\mathcal{P}'(\mathbf{x}, A_0(\mathbf{y})) = \mathcal{P}'(\mathbf{x}, A_1(\mathbf{y}))$$

Proof. Equation (4) for $\mathbf{y} \in J$ implies that

$$\zeta D_{\mathbf{u}_1,\mathbf{u}_2}(A_0(\mathbf{y})) = D_{\mathbf{u}_1,\mathbf{u}_2}(A_1(\mathbf{y})).$$

Spelling out this relation yields

$$\begin{aligned} \zeta\mathcal{P}'(\mathbf{u}_1, A_0(\mathbf{y})) &= \mathcal{P}'(\mathbf{u}_1, A_1(\mathbf{y})), \quad \text{and} \\ \zeta\mathcal{P}'(\mathbf{u}_2, A_0(\mathbf{y})) &= \mathcal{P}'(\mathbf{u}_2, A_1(\mathbf{y})), \end{aligned}$$

and the claim follows from bilinearity of the polar map. □

Writing $\mathbf{e}_1, \ldots, \mathbf{e}_\kappa$ for the canonical basis of $\mathbb{F}_q^\kappa$, define

$$\mathcal{A} : J \to \mathbb{F}_q^N, \quad \mathbf{y} \mapsto \mathbf{e}_{\kappa-1} \otimes A_0(\mathbf{y}) + \mathbf{e}_\kappa \otimes A_1(\mathbf{y}).$$

We can now define our linear subspace

$$U = \big(\mathbb{F}_q^\kappa \otimes \langle \mathbf{u}_1, \mathbf{u}_2 \rangle\big) \oplus \operatorname{im}\mathcal{A} \tag{5}$$

and denote its dimension by u. As in the unstructured regime, we define $W = \langle \operatorname{im} \mathcal{P}^*_{|U} \rangle \subseteq \mathbb{F}_q^m$ and $R = \langle \operatorname{im} \mathcal{P}^{*\prime}_{|U\times U} \rangle \subseteq \mathbb{F}_q^m$, and denote by $w = \dim W$ and $r = \dim R$ their dimensions.

Lemma 5. *Assuming that $m > 2$, the following hold*

$$u = 2\kappa + a, \qquad r \le 2a + \binom{a}{2}, \qquad w \le 2\kappa + 3a + \binom{a}{2}.$$

Proof. By definition, the map A_0 is injective, see Equation (4). In particular, this implies that $\mathcal{A}$ is injective and $\dim \operatorname{im} \mathcal{A} = a$.

Moreover, the sum in Equation (5) is direct. Indeed, if $\mathcal{A}(\mathbf{y}) \in \mathbb{F}_q^\kappa \otimes \langle \mathbf{u}_1, \mathbf{u}_2 \rangle$ for some $\mathbf{y} \in J$, then $A_0(\mathbf{y}) \in \ker D_{\mathbf{u}_1,\mathbf{u}_2}$. By Equation (4), this implies that $\mathbf{y} = D_{\mathbf{u}_1,\mathbf{u}_2}(A_0(\mathbf{y})) = 0$. Consequently, we have $u = 2\kappa + a$, proving the first of the three claims.

Let now $\mathbf{x} \in \langle \mathbf{u}_1, \mathbf{u}_2 \rangle$. For $1 \le i \le \kappa - 1$,

$$\begin{aligned} \mathcal{P}^{*\prime}\big(\mathbf{e}_i \otimes \mathbf{x}, \mathcal{A}(\mathbf{y})\big) &= \varepsilon_{i,\kappa-1}\mathcal{P}'(\mathbf{x}, A_0(\mathbf{y})) + \varepsilon_{i,\kappa}\mathcal{P}'(\mathbf{x}, A_1(\mathbf{y})) \\ &= \varepsilon_{i,\kappa}\big(\zeta\mathcal{P}'(\mathbf{x}, A_0(\mathbf{y})) + \mathcal{P}'(\mathbf{x}, A_1(\mathbf{y}))\big) \\ &= 0, \end{aligned}$$

where the first equality follows from Equation (3) and the second one by Lemma 4. Thus all these polar pairings cancel. For the last copy, the surviving contribution is

$$\begin{aligned}\mathcal{P}^{*\prime}\big(\mathbf{e}_\kappa \otimes \mathbf{x}, \mathcal{A}(\mathbf{y})\big) &= (\varepsilon_{\kappa-1,\kappa} + \zeta\varepsilon_{\kappa\kappa})\mathcal{P}'(\mathbf{x}, A_0(\mathbf{y})) \\ &= (\varepsilon_{\kappa-1,\kappa}(1+\zeta^3))\mathcal{P}'(\mathbf{x}, A_0(\mathbf{y})),\end{aligned} \tag{6}$$

which is nonzero for $m > 2$, provided the polar pairing is nonzero. Let $\varepsilon = \varepsilon_{\kappa-1,\kappa}(1+\zeta^3)$.

Now, write $\mathbf{y} = (y_1, y_2) \in \mathbb{K}^2$ and define

$$Y(J) = \langle y_1, y_2 : \mathbf{y} = (y_1, y_2) \in J\rangle \subseteq \mathbb{F}_q^m.$$

Equation (6) shows that the mixed polar pairings between $\mathbb{F}_q^\kappa \otimes \langle \mathbf{u}_1, \mathbf{u}_2\rangle$ and $\operatorname{im}\mathcal{A}$ span $\varepsilon Y(J)$, whose dimension is at most $2a$. The polar products between pairs of vectors in $\operatorname{im}\mathcal{A}$ span a subspace of dimension at most $\binom{a}{2}$, because the polar map is alternating in even characteristic. Finally, all polar products inside $\mathbb{F}_q^\kappa \otimes \langle \mathbf{u}_1, \mathbf{u}_2\rangle$ vanish. Hence we have the second claim

$$r \le 2a + \binom{a}{2}.$$

Let now $\mathbf{v}_1, \ldots, \mathbf{v}_u$ be a basis of U. By definition of the polar map, we have

$$\langle \operatorname{im}\mathcal{P}^*_{|U}\rangle = \langle \mathcal{P}^*(\mathbf{v}_1), \ldots, \mathcal{P}^*(\mathbf{v}_u)\rangle + R.$$

Therefore $w \le u + r$, and hence

$$w \le 2\kappa + 3a + \binom{a}{2}.$$

□

5 Extracting a Solution

In Table 4, we summarise the dimension of U and the bounds on the dimensions of $R \subseteq W$ in the unstructured and structured regimes.

Table 4: Dimensions of subspaces in unstructured and structured regimes depending on a.

Subspace	Unstructured	Structured
U	$2\kappa + a$	$2\kappa + a$
R	$a + \binom{a}{2} + a(\kappa - a)$	$2a + \binom{a}{2}$
W	$2\kappa + 2a + \binom{a}{2} + a(\kappa - a)$	$2\kappa + 3a + \binom{a}{2}$

Remark 2. In the following, we use the bounds in Table 4 as exact equalities. In practice, subspaces W and R of smaller dimension only make our forgery attack more efficient.

Let $h \geq \max\{u, w\}$ be an integer and let us fix a h-dimensional subspace H such that $H \supseteq W$. Let us define the linear subspace

$$\mathcal{K}_H = \{\mathbf{x} \in \mathbb{F}_q^N : \mathcal{P}^{*\prime}(U, \mathbf{x}) = 0 \mod H\}, \tag{7}$$

and note that $U \subseteq \mathcal{K}_H$. Moreover, we have the lower bound on its dimension

$$N_H = \dim \mathcal{K}_H \geq \dim \mathbb{F}_q^N - u(\dim \mathbb{F}_q^m / H) = N - u(m-h).$$

For our attack to work, we require that its dimension is at least m. Thus, we have the following enabling condition

$$N - m - u(m-h) \geq 0.$$

Recall that forging a MAYO signature reduces to solving an instance of the WMQ problem, i.e. finding a vector $\mathbf{s} \in \mathbb{F}_q^N$ such that $\mathcal{P}^*(\mathbf{s}) = \mathbf{t}$, for a random target $\mathbf{t} \in \mathbb{F}_q^m$. We now explain how to leverage the properties of the subspaces we have built in order to forge.

Let us denote by V a complement of U in $\mathcal{K}_H$, i.e. $\mathcal{K}_H = U \oplus V$. The goal is to find $\mathbf{v} \in V$ and $\mathbf{u} \in U$ such that $\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathbf{t}$.

The attack is based on the definition of the polar map

$$\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \underbrace{\mathcal{P}^*(\mathbf{u})}_{\in W} + \mathcal{P}^*(\mathbf{v}) + \underbrace{\mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v})}_{\in H}. \tag{8}$$

In particular, we have

$$\mathcal{P}^*(\mathbf{u}) = 0 \mod W, \quad \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) = 0 \mod H, \tag{9}$$

and the fact that, modulo R, the equations corresponding to $\mathcal{P}^*(\mathbf{u})$ crucially contain no mixed terms, only squares. This is because, by definition, R is the subspace generated by the image of the polar map on $U \times U$.

Our solving strategy can be parsed in four steps, which we call *Reduced, OV, Square-Only* and *Finishing* steps.

5.1 Reduced Step

Reducing the identity in Equation (8) modulo H, we obtain the smaller MQ problem

$$\mathcal{P}^*(\mathbf{v}) = \mathbf{t} \mod H, \tag{10}$$

which is an $\mathsf{MQ}(\dim \mathcal{K}_H - u, m-h)$ problem. One can solve this using a Gröbner Basis algorithm, and, in expectation, this MQ problem admits $q^{\dim \mathcal{K}_H - u - (m-h)}$

solutions. In order to eventually obtain one full solution for the original WMQ problem, we enumerate need to enumerate q^{h-u} of these.

Note that, in general, this MQ problem is a very underdetermined one. Accordingly, it is possible to adapt the techniques in the Pseudo-Oil Attack [Ost26] to the MQ map underlying the MQ problem in Equation (10)

$$\mathcal{K}_H/U \to \mathbb{F}_q^m/H, \quad \mathbf{x}+U \mapsto \mathcal{P}^*(\mathbf{x})+H, \tag{11}$$

to solve it more efficiently.

Adapting the Pseudo-Oil Attack. The idea is to find a $(m-u)$-dimensional subspace $V' \subseteq V$ such that the MQ problem in the Reduced step naturally splits on V' as a sequence of $p \geq 1$ smaller MQ problems.

For an integer $p \geq 1$, let $b_1, \ldots, b_p \in \mathbb{Z}$ be positive integers, and let $k, \ell \in \mathbb{Z}$ be non-negative integers. We denote

$$B_i = \sum_{j=i}^{p} b_j, \qquad Q = \sum_{j=1}^{p} b_j^2, \quad L = B_1 + \ell, \qquad b_0 = m - h - k - B_1.$$

Then, we have the following theorem, adapted from [Ost26].

Theorem 1 ([Ost26], adapted). *Assume that the following inequalities hold:*

$$\begin{aligned} b_0 \geq \ell, \qquad & b_0(L+1) \leq N_H - u + 1, \\ N_H - m + 1 \ \geq \ & B_1(m-h-k) - \frac{B_1^2 + Q}{2}. \end{aligned} \tag{12}$$

Under mild heuristics, there exists a subspace $V' \subseteq V$ that can be computed in expected time

$$\mathcal{T}_{V'} = b_0\, \mathsf{GB}_q(L, L) + \sum_{i=1}^{p-1} b_i\, \mathsf{GB}_q(B_{i+1}, B_{i+1}) + \mathrm{poly}(n, m, \log q)$$

such that there exists a heuristic algorithm that returns all the solutions in V' for the MQ problem in Equation (10) *in expected time*

$$\mathcal{T} = q^{k+h-u} \left(\sum_{i=1}^{p} \mathsf{GB}_q(b_i, b_i) + \mathsf{GB}_q(m-h-k-L, m-h-L) \right).$$

Proof. Under the heuristic in [Ost26][Heuristic 1], we apply the techniques from the Pseudo-Oil Attack to the MQ map Equation (11), which is well-defined because $\mathcal{P}^*(U) = 0 \mod H$ and by definition of $\mathcal{K}_H$. In particular, the expected number of solutions is $q^{\dim \mathcal{K}_H - u - (m-h)}$. Differently from [Ost26] where finding one solution is enough, we need to find q^{h-u} solutions. This can be achieved by enlarging the guessing subspace of the Pseudo-Oil Attack, going from dimension k to dimension $k+h-u$, and consequently enumerating q^{k+h-u} vectors.

Accordingly, computing the subspace $V' \subseteq V$ can be done with the *Enriched Centrifugation* algorithm proposed in [Ost26], whose expected runtime coincides with the one in the theorem statement. Such subspace exists provided the inequalities in Equation (12) hold: the first two exactly coincide with the inequalities of the Pseudo-Oil Attack adapted to the MQ map defined in Equation (11), while the last one accounts for the larger guessing subspace of dimension $k+h-u$ that we require.

Computing all the solutions in V' for the MQ problem in Equation (10) can be done with the *Enriched Extraction* algorithm from the Pseudo-Oil Attack: the claimed expected runtime coincides with the expected runtime of q^{h-u} iterations of the Enriched Extraction algorithm, where each iteration produces one solution in expectation, according to [Ost26][Theorem 2]. □

5.2 OV Step

Assume that $\mathbf{v} \in V$ has been fixed such that it solves the MQ problem in Equation (10). In the OV step, we consider the MQ problem in $\mathcal{P}^*(\mathbf{u}+\mathbf{v}) = \mathbf{t}$ modulo W. This yields the following equation to solve

$$\mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) = \mathbf{t} - \mathcal{P}^*(\mathbf{v}) \mod W. \tag{13}$$

It is a linear system of equations in $\mathbf{u} \in U$. More precisely, since $\mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) \in H$ and $\mathbf{t} - \mathcal{P}^*(\mathbf{v}) \in H$, it consists of $\dim H/W = h - w$ equations in u variables.

Let σ_{OV} be the rank of the coefficient matrix of the linear system in Equation (13). Then

$$0 \leq \sigma_{\mathrm{OV}} \leq \min\{u, h-w\}.$$

Gaussian elimination either detects an inconsistency or computes the affine solution space, which we call U_{OV}. In the former case, we discard the current Reduced solution $\mathbf{v}$ and continue with the next one. Otherwise,

$$\dim U_{\mathrm{OV}} = u - \sigma_{\mathrm{OV}} \geq \max\{0, u-h+w\}.$$

In particular, $u-h+w < 0$ means that the OV linear system is overdetermined.

5.3 Square-Only step

As $\mathbf{v} \in V$ is fixed and $\mathbf{u} \in U$ is partially fixed, let us consider the MQ problem $\mathcal{P}^*(\mathbf{v}+\mathbf{u}) = \mathbf{t}$ modulo R, where now $\mathbf{u} \in U_{\mathrm{OV}}$. This is

$$\mathcal{P}^*(\mathbf{u}) + \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) = \mathbf{t} - \mathcal{P}^*(\mathbf{v}) \mod R. \tag{14}$$

Note that both $\mathcal{P}^*(\mathbf{u})$ and $\mathbf{t} - \mathcal{P}^*(\mathbf{v}) - \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v})$ lie in W, since $\mathbf{u} \in U_{\mathrm{OV}}$, so this MQ problem consists of $\dim W/R = w - r$ equations only. Therefore, this is a $\mathsf{MQ}(u - \sigma_{\mathrm{OV}}, w - r)$ problem. By definition of R, we have the following lemma.

Lemma 6. *The equations corresponding to Equation* (14) *only contain squares and linear terms in* $\mathbf{u}$.

Proof. Let $\mathbf{z}_1, \ldots, \mathbf{z}_u$ be a basis of U, and let $\mathbf{u} = \sum_{i=1}^{u} x_i \mathbf{z}_i$. Then, by definition of R, we have that

$$\mathcal{P}^*(\mathbf{u}) = \mathcal{P}^* \left(\sum_i x_i \mathbf{z}_i \right) = \sum_i x_i^2 \mathcal{P}^*(\mathbf{z}_i) = \mathbf{M}\mathbf{x}^2 \mod R,$$

where $\mathbf{M} \in \mathbb{F}_q^{(w-r)\times u}$ is the matrix whose columns are the $\mathcal{P}^*(\mathbf{z}_i)$ and $\mathbf{x}^2 = (x_1^2, \ldots, x_u^2)$. The linear terms come from the polar map

$$\mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) = \sum_i x_i \mathcal{P}^{*\prime}(\mathbf{z}_i, \mathbf{v}) = \mathbf{N}\mathbf{x},$$

where $\mathbf{N} \in \mathbb{F}_q^{(w-r)\times u}$ is the matrix whose columns are $\mathcal{P}^{*\prime}(\mathbf{z}_i, \mathbf{v})$. □

Now, let $e \geq 1$ be such that $q = 2^e$. Given the property modulo R in Lemma 6, we descend the equations in Equation (14) to a linear system over $\mathbb{F}_2$, with $e(w-r)$ equations in $e(u - \sigma_{\mathrm{OV}})$ variables.

We perform Gaussian elimination on this binary linear system and denote by U_2 the binary affine subspace of solutions. If this linear system is inconsistent, we backtrack to the OV step. Let ρ_2 be the rank of the binary linear system. Then, the dimension of U_2 is $e(u - \sigma_{\mathrm{OV}}) - \rho_2$.

5.4 Finishing Step

In this last step, we consider the MQ problem $\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathbf{t}$ without modulo subspaces. Using the definition of the polar map, this is

$$\mathcal{P}^*(\mathbf{u}) + \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) = \mathbf{t} - \mathcal{P}^*(\mathbf{v}).$$

Descending to $\mathbb{F}_2$ and restricting to U_2, we get a $\mathsf{MQ}(e(u - \sigma_{\mathrm{OV}}) - \rho_2, er, 2)$ problem, which we solve using a binary Gröbner Basis algorithm, thereby finding a complete $\mathbf{u} \in U$ such that $\mathbf{s} = \mathbf{u} + \mathbf{v}$ solves the original WMQ problem.

5.5 Correctness and Runtime

In order to prove the correctness and compute the expected complexity of our algorithm, one needs to understand the distributions of the systems of equations arising during the extraction process.

We start by noting that the binary linear system of equation coming from Equation (14) is not uniformly random. More precisely, we have the following lemma.

Lemma 7. *The matrix* $\mathbf{M}$ *is invertible. In particular, we have*

$$(e-1)(u - \sigma_{\mathrm{OV}}) \leq \rho_2 \leq e(u - \sigma_{\mathrm{OV}}).$$

Proof. We borrow the notation from Lemma 6. Since the subspace W is defined as the span of the image of U through $\mathcal{P}^*$ modulo R, $\mathbf{M}$ has column rank $w-r$. Moreover, since $w-r=u$ and there are u rows, the matrix $\mathbf{M}$ is invertible. Let us rewrite the MQ map of Equation (14) as

$$\mathcal{L}(\mathbf{x}) = \mathbf{M}\mathbf{x}^2 + \mathbf{N}\mathbf{x} \mod R.$$

Then, we have that the rank is $\rho_2 = e(u-\sigma_{\mathrm{OV}}) - \dim \ker_{\mathbb{F}_2} \mathcal{L}$. To prove our claim, it is therefore enough to show that $\dim \ker_{\mathbb{F}_2} \mathcal{L}$ is at most $u-\sigma_{\mathrm{OV}}$. Let $\{\mathbf{k}_i\}_i$ be an $\mathbb{F}_q$-basis of $\langle \mathcal{L}^{-1}(0)\rangle$, and let $\mathbf{k} = \sum_{i=1}^{r} \alpha_i \mathbf{k}_i$. Then, since $\mathbf{M}\mathbf{k}_i^2 + \mathbf{N}\mathbf{k}_i = 0$ for all i, we have

$$\begin{aligned} 0 = \mathcal{L}(\mathbf{k}) &= \mathbf{M}\sum_i \alpha_i^2 \mathbf{k}^2 + \sum_i \alpha_i B\mathbf{k}_i \\ &= \mathbf{M}\sum_i \alpha_i^2 \mathbf{k}^2 + A\sum_i \alpha_i \mathbf{k}_i^2 \\ &= \mathbf{M}\sum_i^r (\alpha_i^2 - \alpha_i)\mathbf{k}_i^2. \end{aligned}$$

Since $\mathbf{M}$ is invertible and the $\mathbf{k}_i$ are linearly independent, we have that $\alpha_i^2 - \alpha_i = 0$ for all i, which implies that $\mathbf{k}$ is binary linear combinations of the $\mathbf{k}_i$. Consequently, we have

$$0 \leq \dim \ker_{\mathbb{F}_2} \mathcal{L} = \dim \langle \mathcal{L}^{-1}(0)\rangle \leq u - \sigma_{\mathrm{OV}},$$

and the claim follows. □

To estimate the rank distribution, we now need the following heuristic.

Heuristic 1 *The matrices* $\mathbf{M}$ *and* $\mathbf{N}$ *are distributed uniformly at random and independently in* $\mathrm{GL}_u(\mathbb{F}_q)$ *and* $\mathbb{F}_q^{u\times u}$*, respectively.*

As for the rest of the system of equations, we heuristically assume they are uniformly distributed.

Heuristic 2 *The distribution of every system of equations considered in the Reduced, OV and Finishing steps is uniformly random on the set of system of equations with the correct number of variables and equations.*

Theorem 2 (Correctness and Number of Solutions). *Assume that Heuristic 1 and 2 hold. Let* $V' \subseteq V$ *be the linear subspace computed in Theorem 1, and denote by* $\mathcal{O} = U \oplus V'$*. For a target* $\mathbf{t} \in \mathbb{F}_q^m$*, define*

$$\mathcal{V}_{\mathbf{t}} = \{\mathbf{v} \in V' : \mathcal{P}^*(\mathbf{v}) = \mathbf{t} \bmod H\}$$

and, for $\mathbf{v} \in \mathcal{V}_{\mathbf{t}}$*,*

$$\mathcal{U}_{\mathbf{t},\mathbf{v}} = \{\mathbf{u} \in U : \mathcal{P}^*(\mathbf{u}+\mathbf{v}) = \mathbf{t}\}.$$

The extraction procedure returns exactly

$$\{\mathbf{u}+\mathbf{v} : \mathbf{v} \in \mathcal{V}_{\mathbf{t}},\ \mathbf{u} \in \mathcal{U}_{\mathbf{t},\mathbf{v}}\} = (\mathcal{P}^*_{|\mathcal{O}})^{-1}(\mathbf{t}).$$

For a uniform target $\mathbf{t}$*, one has*

$$\mathbb{E}[|\mathcal{V}_{\mathbf{t}}|] = q^{h-u}, \qquad \mathbb{E}\big[|\mathcal{U}_{\mathbf{t},\mathbf{v}}| \mid \mathbf{v} \in \mathcal{V}_{\mathbf{t}}\big] = q^{u-h}.$$

In particular, the restriction to $\mathcal{O}$ *has one preimage on average.*

Proof. Every element of $\mathcal{O}$ can be uniquely decomposed as $\mathbf{u} + \mathbf{v}$, with $\mathbf{u} \in U$ and $\mathbf{v} \in V'$. Recall that by definition of polar map we have

$$\mathcal{P}^*(\mathbf{u} + \mathbf{v}) = \mathcal{P}^*(\mathbf{v}) + \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v}) + \mathcal{P}^*(\mathbf{u}).$$

Successively projecting this identity along $H \supseteq W \supseteq R$ gives exactly the equations solved in the Reduced, OV, Square-Only, and Finishing steps, which proves the first assertion.

There are q^{m-u} choices of $\mathbf{v}$, and a uniform target has a uniform image in the $(m-h)$-dimensional quotient $\mathbb{F}_q^m/H$. Hence the expected size of the fiber $\mathcal{V}_{\mathbf{t}}$ is $q^{m-u}/q^{m-h} = q^{h-u}$. Conditioned on a fixed $\mathbf{v}$ coming from the Reduced step, $\mathbf{t} - \mathcal{P}^*(\mathbf{v})$ is uniform in H. Averaging the fibers of the map $\mathbf{u} \mapsto \mathcal{P}^*(\mathbf{u}) + \mathcal{P}^{*\prime}(\mathbf{u}, \mathbf{v})$ from a set of size q^u to one of size q^h gives q^{u-h}. Therefore, the total number of expected solutions produced by the extraction procedure is one. □

Let us denote by τ_σ and $\pi_{\rho|\sigma}$ the probabilities $\Pr[\sigma_{\mathrm{OV}} = \sigma]$ and $\Pr[\rho_2 = \rho | \sigma_{\mathrm{OV}} = \sigma]$, respectively.

Theorem 3 (Expected Complexity). *Assume the same hypotheses of Theorem 1, and retain the notation for* $\mathcal{T}_{V'}$ *and* $\mathcal{T}$*.*

Assume that Heuristic 1 and 2 hold. The expected complexity of the Square-Only and Finishing steps conditioned on the rank of the OV system to be σ *is*

$$\mathcal{T}_{\mathrm{SOF}|\sigma} = \mathrm{GE}_2(e(u-\sigma), e(w-r)) + \sum_{\rho} \pi_{\rho|\sigma} 2^{\rho - e(w-r)} \mathsf{GB}_2(e(u-\sigma) - \rho, er).$$

Then, there exists an algorithm that forges a valid MAYO signature for the parameters (n, m, q, κ) *in expected time*

$$\mathcal{T}_{V'} + \mathcal{T} + q^{h-u} \left(\mathrm{GE}_q(u, h-w) + \sum_{\sigma} \tau_\sigma q^{\sigma - (h-w)} \mathcal{T}_{\mathrm{SOF}|\sigma} \right).$$

Proof. Since our strategy finds all existing solutions to the WMQ problem restricted to $U \oplus V$, to prove the claim one just needs to compute the expected complexity of a full run of our algorithm. By Theorem 1, constructing V' and enumerating the Reduced solutions cost $\mathcal{T}_{V'} + \mathcal{T}$ in expectation. By Theorem 2, the expected number of solutions to the Reduced step is q^{h-u}.

Let us call *iteration* the sequential run of the OV, Square-Only and Finishing step. Theorem 2 shows that in expectation one needs to perform q^{h-u} iterations in order to obtain one MAYO forged signature. In the OV step, one has to solve

a linear system over $\mathbb{F}_q$ with u variables and $h - w$ equations, whose cost is $\mathrm{GE}_q(u, h - w)$.

The cost of the Square-Only and Finishing steps must be averaged over the probability τ_σ of rank σ for the OV linear system and its consistency probability $q^{\sigma-(h-w)}$. Accordingly, assuming $\sigma_{\mathrm{OV}} = \sigma$, the binary linear system in the Square-Only step can be solved via Gaussian elimination in expected time $\mathrm{GE}_2(e(u - \sigma), e(w - r))$. Now, the expected cost of the Finishing step must be averaged over the rank and consistency probabilities $\pi_{\rho|\sigma}$ and $2^{\rho-e(w-r)}$ of the Square-Only binary linear system. Assuming that $\rho_2 = \rho$, the expected complexity of the Finishing step is $\mathsf{GB}_2(e(u - \sigma) - \rho, er)$.

□

6 Leveraging Multiple Targets

In this section, we study the multi-target variant of the MQ problem, MMQ problem for short. Let $\mathcal{P} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ be an MQ map, and let $\mathbf{t}_i$, $i \geq 1$, be an indexed stream of independent uniformly random elements of $\mathbb{F}_q^m$. Given $\mathcal{P}$ and sequential access to this target stream, the *multi-target MQ problem* asks to compute $(I, \mathbf{s}) \in \mathbb{N}_{>0} \times \mathbb{F}_q^n$ such that $\mathcal{P}(\mathbf{s}) = \mathbf{t}_I$. We will write for short $\mathsf{MMQ}_q(n, m)$ to indicate an instance of the MMQ problem.

In a nutshell, we show that access to multiple targets can reduce the effective number of equations of an MQ problem. After choosing d linearly independent targets, we quotient the codomain by their span and solve the resulting system of $m - d$ quadratic equations. This makes each MQ instance cheaper to solve, at the cost of a lower probability that a solution of the projected MQ problem yields a preimage of one of the original targets. This technique can be applied to any MQ solving algorithm, provided the underlying MQ map is homogeneous.

Eventually, we show that this multi-target technique is compatible with the structural forgery attack presented in Sections 4 and 5.

In the following, assume we are given on input an instance $(\mathcal{P}, \{\mathbf{t}_i\}_{i\geq 1})$ of the $\mathsf{MMQ}_q(n, m)$ problem.

6.1 Reducing the Number of Equations

Fix $1 \leq d \leq m$. By querying the target stream and discarding linearly dependent targets, we select d targets that are linearly independent. Since this requires $O(d) = O(m)$ stream queries in expectation, in the following we assume without loss of generality that the first d targets in the stream are linearly independent. Define the d-dimensional *target subspace* $T = \langle \mathbf{t}_1, \ldots, \mathbf{t}_d \rangle \subseteq \mathbb{F}_q^m$. Then, let us consider the $\mathsf{MQ}_q(n, m - d)$ map

$$\mathcal{P}_T : \mathbb{F}_q^n \to \mathbb{F}_q^m / T, \quad \mathbf{x} \mapsto \mathcal{P}(\mathbf{x}) + T.$$

Consider the $\mathsf{MQ}_q(n, m-d)$ problem with input $(\mathcal{P}_T, 0)$. Then, every nonzero solution $\mathbf{s} \in \mathbb{F}_q^n$ satisfies $\mathcal{P}(\mathbf{s}) \in T$. Our multi-target algorithm relies on the following heuristic.

Heuristic 3 *Let $\mathbf{s} \in \mathbb{F}_q^n$ be a nonzero solution to the $\mathsf{MQ}_q(n, m-d)$ problem with input $(\mathcal{P}_T, 0)$. Then, $\mathcal{P}(\mathbf{s}) \in T$ is a uniformly random element in T.*

We have the followig elementary lemma about the probability that $\mathcal{P}(\mathbf{s})$ lands in one of the punctured lines spanned by the targets $\mathbf{t}_i$.

Lemma 8. *Let $\mathbf{s} \in \mathbb{F}_q^n$ be a solution of the $\mathsf{MQ}_q(n, m-d)$ problem with input $(\mathcal{P}_T, 0)$. Under Heuristic 3, there exists $\lambda \in \mathbb{F}_q$ and an index $I \leq d$ such that $(I, \lambda\mathbf{s})$ solves the instance $(\mathcal{P}, \{\mathbf{t}_i\}_{i\geq 1})$ of the $\mathsf{MMQ}_q(n, m)$ problem with probability*

$$p_d = \begin{cases} \frac{d(q-1)}{q^d} & \text{if } q \text{ is even,} \\ \frac{d(q-1)}{2q^d} & \text{if } q \text{ is odd.} \end{cases}$$

Proof. Under Heuristic 3, the probability that $\mathcal{P}(\mathbf{s}) = \mathbf{t} = \alpha\mathbf{t}_I$ for some $\alpha \in \mathbb{F}_q^\times$ and $I \leq d$ is clearly $\frac{d(q-1)}{q^d}$. In this case, define $\lambda \in \mathbb{F}_q$ such that $\lambda^2 = \alpha^{-1}$, which exists with probability 1 and $1/2$ when q is even and odd, respectively. Then, we have

$$\mathcal{P}(\lambda\mathbf{s}) = \lambda^2\mathcal{P}(\mathbf{s}) = \alpha^{-1}\mathbf{t} = \mathbf{t}_I.$$

□

6.2 Multi-Target Solver

We give in Algorithm 1 the pseudocode for our multi-target solver-agnostic algorithm.

Algorithm 1: Multi-Target Solver

Input: An instance $(\mathcal{P}, \{\mathbf{t}_i\}_{i\geq 1})$ of the $\mathsf{MMQ}_q(n, m)$ problem; the number of targets $1 \leq d \leq m$; a probabilistic algorithm MQSolver that returns random nonzero solutions of MQ problems.

Output: A solution to the input MMQ problem.

```
T ← ⟨t_1, ..., t_d⟩.
Define P_T by quotienting the codomain of P by T.
while true do
    s ← MQSolver(P_T, 0)                          // s ≠ 0
    t ← P(s)
    if λ²t = t_I for some I ≤ d and λ ∈ F_q^× then
        return (I, λs)
```

Theorem 4. *Let $(\mathcal{P}, \{\mathbf{t}_i\}_{i\geq 1})$ be an instance of the of the $\mathsf{MMQ}_q(n, m)$ problem. Let T be the target subspace and $1 \leq d \leq m$ its dimension, and let MQSolver be a probabilistic algorithm that solves instances of the $\mathsf{MQ}_q(n', m')$ problem for*

random nonzero solutions in expected time $\mathcal{T}_{n',m',q}$. Then, under Heuristic 3, Algorithm 1 solves the instance of the $\mathsf{MMQ}_q(n,m)$ problem in expected time

$$\mathrm{poly}(n,m,\log q) + \frac{\mathcal{T}_{n,m-d,q} + \mathrm{poly}(n,m,\log q)}{p_d}.$$

Proof. Every output $\mathbf{s}$ of MQSolver satisfies $\mathcal{P}_T(\mathbf{s}) = 0$, which is equivalent to $\mathcal{P}(\mathbf{s}) \in T$. Under Heuristic 3, the success probability in each iteration of the while loop is p_d, as computed in Lemma 8. Hence, the expected number of invocations of MQSolver is $1/p_d$. If an iteration succeeds, the algorithm finds $I \leq d$ and $\lambda \in \mathbb{F}_q^\times$ described in Lemma 8 in time $\mathrm{poly}(n,m,\log q)$. □

Ignoring the polynomial overhead, the reduction is beneficial whenever

$$\frac{\mathcal{T}_{n,m-d,q}}{p_d} < \mathcal{T}_{n,m,q}.$$

The parameter d therefore captures the tradeoff between solving a system with fewer equations and repeating the projected solve more often.

6.3 Compatibility with the Structural Forgery Attack

In this section, we show that the multi-target technique is compatible with our structural forgery attack presented in Sections 4 and 5. Let $T \subseteq \mathbb{F}_q^m$ be the d-dimensional target space, $\mathcal{P}^*$ be the WMQ map and write

$$\mathcal{P}_T^*(\mathbf{x}) = \mathcal{P}^*(\mathbf{x}) \bmod T.$$

Thus, solving $\mathcal{P}_T^*(\mathbf{x}) = 0$ is equivalent to finding $\mathbf{x}$ such that $\mathcal{P}^*(\mathbf{x}) \in T$. Since reduction modulo T is linear, it commutes with polarisation, i.e.

$$\mathcal{P}_T^{*\prime}(\mathbf{x},\mathbf{y}) = \mathcal{P}^{*\prime}(\mathbf{x},\mathbf{y}) \bmod T.$$

Let U, W, R be as in Section 4. After reduction modulo T, the corresponding spaces are

$$W_T = W \bmod T = (W+T)/T, \qquad R_T = R \bmod T = (R+T)/T.$$

In particular, $R_T \subseteq W_T$, and

$$w_T = \dim W_T = w - \dim(W \cap T), \qquad r_T = \dim R_T = r - \dim(R \cap T).$$

Hence all cancellations used in the construction of U remain valid modulo T, and the upper bounds on w and r remain upper bounds on w_T and r_T. Moreover, by sampling more targets, one can enforce $w_T = w$ and therefore $r_T = r$.

Choose an h_T-dimensional subspace $H_T \subseteq \mathbb{F}_q^m/T$ containing W_T, for $h_T \geq \max\{u, w_T\}$. Then extraction of a solution therefore proceeds through the same sequence of quotient spaces, now associated with the flag

$$R_T \subseteq W_T \subseteq H_T \subseteq \mathbb{F}_q^m/T.$$

7 Security Estimates

In this section, we estimate the classical bit complexity of forging signatures for the homogeneous Round 2 construction of MAYO [BCC+24] and the inhomogeneous Round 3 construction. Our git repository[1] contains an optimisation script that allows to reproduce the results that follow.

7.1 MAYO Round 2

We compare four attacks: the best generic attack to date [Has21,Ost26,ACF+26], its combination with our multi-target technique, our structural attack against a single target, and the combination of the multi-target and structural forgery. This comparison allows us to distinguish the speedup gained from the whipping structure and that of access to multiple targets.

The complexity formula for the structural forgery is the one in Theorem 3. In particular, we used the CryptographicEstimator library [EVZB24] to estimate the cost of the Gröbner Basis algorithms used in the Reduced and Finishing steps. For the Reduced steps, optimised the quantities $k, \ell, b_0, b_1, \ldots, b_p$ of the Pseudo-Oil Attack using dynamic programming. For the expected complexity of Gaussian elimination, we used the formula

$$\mathrm{GE}_q(n,m) = \begin{cases} 0 & \text{if } m = 0, \\ n' m \min\{n', m\}^{\omega-2} (\log q)^\theta & \text{else,} \end{cases}$$

where $n' = \max\{1, n\}$, $2 \leq \omega \leq 3$ the matrix multiplication exponent, and $0 \leq \theta \leq 2$ is the field to bit operation conversion exponent. They can be set in the optimisation script, and by default we use $\omega = 2.81$ and $\theta = 2$.

For the multi-target algorithm, the complexity formula is the one in Theorem 4. For the forgery attack combining the structural insights and the multi-target attack, the optimisation algorithm follows the description in Section 6.3.

Our optimisation script also prints a complete breakdown of the complexity, showing how the complexity distributed among the differents steps of our algorithm. We report in Table 5 the estimates for our attacks, compared with generic single-instance forgeries.

Security Level I and III. For the first Security Level I parameter set, the structural insight alone reduces the estimated forgery cost by roughly 9 bits relative to the generic baseline. The multi-target technique alone gives a reduction of approximately 6 bits. Combining both techniques yields an estimated cost of 2^{128} bit operations, improving over the generic baseline by a 18. Note that the multi-target and structural improvements are not additive, since the balance of optimal attack parameters changes when combining the two techniques.

For Security Level III, our multi-target, structural and combined algorithms improve over the generic approach by $5, 9$ and 16 bits respectively.

[1] `https://github.com/maxostuzzi/forging_mayo`

Table 5: Comparison for the complexity of forging a Round 2 MAYO signature. Parameters below the required security level in red.

			This work		
Security	(n, m, q, κ)	Generic	Multi-Target	Structure	Combined
Level I	$(86, 78, 16, 10)$	146	140	135	**128**
Level I	$(81, 64, 16, 4)$	113	147	157	150
Level III	$(118, 108, 16, 11)$	223	218	214	**207**
Level V	$(154, 142, 16, 12)$	295	**293**	307	300

For the alternative Level I parameter set, our methods do not improve over the generic techniques. Indeed, this parameter set is affected by the key-recovery wedge attack [Ran26].

Security Level V. For Security Level V, only our multi-target approach is able to outperform the generic attack, by only 2 bits. This can be explained as a consequence of two different factors: the designer's choice for the whip factor κ and the underdeterminedness of the underlying generic MQ problem.

Indeed, our algorithm strongly benefits from a large κ, since it automatically allows for higher-dimensional U. Since the choice whip factor does not significantly grow with the security level, this explains why our structural attack is not as effective as for levels I and III.

One could still wonder why the multi-target does not help as much as for levels I and III. The reason is that, since this technique is generic, the only factor impacting its performance is the underdeterminedness of the underlying MQ problems. As noted in [MOR26,Ost26], the correct underdeterminedness measure is the *u-rate*, defined for MAYO instances as

$$\rho = \frac{2n\kappa}{m(m+1)}.$$

While for levels I and III the u-rates are 0.28 and 0.22 respectively, the u-rate of level V is 0.18.

7.2 MAYO Round 3

As anticipated in the introduction, the MAYO team has recently released the new Round 3 specification documents. The changes affect the effectiveness of our attack techniques. The most important modifications for us consists in:

- adding a linear part L to the whipped up construction,
- changing the emulsifiers schedule.

In particular, the linear map L is defined as

$$L(\Lambda, \mathbf{x}_1, \ldots, \mathbf{x}_\kappa) = \sum_i E_{0i} \Lambda \mathbf{x}_i,$$

where $\Lambda \in \mathbb{F}_q^{m \times n}$ is a linear map comes from essentially hashing the message and the salt and the E_{0i} matrices are new emulsifiers matrices. Consequently, the new WMQ map for Round 3 is

$$\mathcal{P}^*(\mathbf{x}_1, \ldots, \mathbf{x}_\kappa) = \sum_{i=1}^{\kappa} E_{ii}\mathcal{P}(\mathbf{x}_i) + \sum_{1 \le i < j \le \kappa} E_{ij}\mathcal{P}'(\mathbf{x}_i, \mathbf{x}_j) + L(\Lambda, \mathbf{x}_1, \ldots, \mathbf{x}_\kappa).$$

Multiple Targets. The linear part of the WMQ map trivially secures MAYO against our multi-target strategy. Indeed, our technique heavily relies on the WMQ being homogeneous to be able to rescale a preimage of a target line and hit the exact target. One could think of applying our multi-target technique on the restricted WMQ map

$$\ker L \xrightarrow{\mathcal{P}^*} \mathbb{F}_q^m,$$

in order to make the linear L part vanish automatically. However, this is incompatible with multi-target techniques, since the L linear map depends on Λ, which in turn directly depends on the salt and the message, i.e. is not fixed for multiple targets.

Structural Forgery. The same linear map L, together with the new choice of emulsifiers, reduces the impact of our structural observation, too.

Let U and $u = 2\kappa + a$ be as defined in Section 4. Every new linear term $E_{0i}\Lambda\mathbf{u}_i$ contributes as a new linearly independent direction in the codomain spaces W. Moreover, the new emulsifiers choice prevents us from using the structured regime.

Regarding R, we have that $r = \dim R$ coincides with the one given in the unstructured regime of Table 4, since the linear map does not contribute to the polar map $\mathcal{P}^{*\prime}$.

To slightly reduce the dimension of W, it is possible to choose $\mathbf{u}_1 \in \ker \Lambda$. The dimension of W therefore increases by the dimension $\ell_\Lambda = L(U)$ over the one reported in Table 4 for the unstructured regime. More precisely, it is easy to verify that $\ell_\Lambda \le \kappa + a$, and therefore,

$$w \le 3\kappa + 3a + \binom{a}{2} + a(\kappa - a).$$

Moreover, as mentioned in [Ost26], the Pseudo-Oil Attack can be performed on inhomogeneous MQ problems with little modification that do not alter the runtime.

We then perform our structural forgery attack with these new dimensions for U, W, R. The results against the newly proposed parameter sets of Round 3 are reported in Table 6. In the table, we also report the complexity of our attack as if there was no linear term and the old choice of emulsifiers, i.e. for the combined structural and multi-target algorithm. This table proves that adding a linear part to the WMQ map and changing the choice of the emulsifiers was, in

retrospection, a very good decision. Nevertheless, our structural attack is able to improve over generic forgeries by 6, 4 and 2 bits for respectively the new Security Level I and III parameter sets.

Table 6: Comparison for the complexity of forging a Round 3 MAYO signature. Parameters below the required security level in red.

Security	(n, m, q, κ)	Generic	R3 Structural	R2 Combined
Level I	$(88, 80, 16, 10)$	153	**147**	135
Level I	$(86, 64, 16, 5)$	144	**140**	134
Level III	$(118, 108, 16, 11)$	223	**221**	207
Level V	$(154, 142, 16, 12)$	**295**	314	300

8 Conclusions

In this work, we present two novel attacks. The first one is the first structural forgery attack that exploits the geometry of WMQ maps in MAYO. The second one is a generic multi-target attack that exploits the homogeneity of the MQ map and that can be coupled with most MQ solving algorithms. Their combination significantly impacts the security of MAYO as submitted to the second round of the NIST competition. Against the very recently released Round 3 MAYO, our structural insights still partially apply, while the multi-target technique is completely blocked since the underlying MQ maps are not homogeneous. Therefore, the impact of the attack is mitigated to only 6 and 4 bits of security loss for the two Security Level I sets. Yet, our attacks bring the second Security Level I set below the required 143 bits of security.

Independently from the actual security losses, this work is the first to show that the whip up construction of MAYO creates a geometric structure that can be algorithmically exploited. We believe that our observations might be extended, improved and generalised in the future to obtain even more efficient forgery attacks for MAYO.

References

ABC+26. Gorjan Alagic, Maxime Bros, Pierre Ciadoux, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Yi-Kai Liu, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Hamilton Silberg, Daniel Smith-Tone, and Noah Waller. Status report on the second round of the additional digital signature schemes for the nist post-quantum cryptography standardization process. NIST Interagency/Internal Report NIST IR 8610, National Institute of Standards and Technology, 2026.

ACF+26. Hideki Asanuma, Yilong Chen, Hiroki Furue, Kosuke Sakata, and Tsuyoshi Takagi. Improved complexity estimates for underdetermined MQ systems via generalized variable partitioning. Cryptology ePrint Archive, Paper 2026/1054, 2026.

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996.

BCC+24. Ward Beullens, Fabio Campos, Sofía Celi, Basil Hess, and Matthias J. Kannwischer. MAYO. Technical report, National Institute of Standards and Technology, 2024. available at `https://csrc.nist.gov/Projects/pqc-dig-sig/round-2-additional-signatures`.

BCC+26. Ward Beullens, Fabio Campos, Sofía Celi, Basil Hess, and Matthias J. Kannwischer. MAYO: Specification document, round 3 version, 2026.

BDK+18. Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Kyber: A CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy*, pages 353–367. IEEE Computer Society Press, April 2018.

Beu22. Ward Beullens. MAYO: Practical post-quantum signatures from oil-and-vinegar maps. In Riham AlTawy and Andreas Hülsing, editors, *SAC 2021*, volume 13203 of *LNCS*, pages 355–376. Springer, Cham, September / October 2022.

DH76. Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.

DKL+18. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR TCHES*, 2018(1):238–268, 2018.

EVZB24. Andre Esser, Javier A. Verbel, Floyd Zweydinger, and Emanuele Bellini. {SoK: CryptographicEstimators - a Software Library for Cryptographic Hardness Estimation}. In *AsiaCCS*. {ACM}, 2024.

FY79. Aviezri S. Fraenkel and Yaacov Yesha. Complexity of problems in games, graphs and algebraic equations. *Discrete Applied Mathematics*, 1(1–2):15–30, 1979.

Has21. Yasufumi Hashimoto. An improvement of algorithms to solve under-defined systems of multivariate quadratic equations. Cryptology ePrint Archive, Paper 2021/1045, 2021.

HLY12. Yun-Ju Huang, Feng-Hao Liu, and Bo-Yin Yang. Public-key cryptography from new multivariate quadratic assumptions. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 190–205. Springer, Berlin, Heidelberg, May 2012.

KPG99. Aviad Kipnis, Jacques Patarin, and Louis Goubin. Unbalanced Oil and Vinegar signature schemes. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 206–222. Springer, Berlin, Heidelberg, May 1999.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010.

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *DCC*, 75(3):565–599, 2015.

MOR26. Alexander May, Massimo Ostuzzi, and Henrik Ressler. Just guess: Improved (quantum) algorithm for the underdetermined MQ problem. LNCS, pages 334–362, 2026.

Nat24a. National Institute of Standards and Technology. Module-Lattice-Based Digital Signature Standard. Federal Information Processing Standards Publication FIPS 204, National Institute of Standards and Technology, 2024.

Nat24b. National Institute of Standards and Technology. Module-Lattice-Based Key-Encapsulation Mechanism Standard. Federal Information Processing Standards Publication FIPS 203, National Institute of Standards and Technology, 2024.

Ost26. Massimo Ostuzzi. Pseudo-oil subspaces and the geometry of underdetermined MQ problems. Cryptology ePrint Archive, Paper 2026/1122, 2026.

Ran26. Lars Ran. Wedges, oil, and vinegar: An analysis of uov in the exterior algebra. In *Advances in Cryptology – EUROCRYPT 2026: 45th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Rome, Italy, May 10–14, 2026, Proceedings, Part IV*, page 363–388, Berlin, Heidelberg, 2026. Springer-Verlag.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

TW12. Enrico Thomae and Christopher Wolf. Solving underdetermined systems of multivariate quadratic equations revisited. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 156–171. Springer, Berlin, Heidelberg, May 2012.