Computers, Materials & Continua

Tech Science Press

https://doi.org/10.32604/cmc.2026.084319

ARTICLE

Pareto-Based Multi-Objective Evaluation of Multivariate Signature Schemes in the NIST Standardization Process

Jian Zhang[1], **Seong-Min Cho**[2] **and Seung-Hyun Seo**[3,*]

[1]Department of Electrical and Electronic Engineering, Graduate School, Hanyang University, Seoul, Republic of Korea
[2]Electronics and Telecommunications Research Institute (ETRI), Daejeon, Republic of Korea
[3]Division of Electrical Engineering, Hanyang University ERICA, Ansan, Republic of Korea

*Corresponding Author: Seung-Hyun Seo. Email: seosh77@hanyang.ac.kr

Received: 20 April 2026; Accepted: 02 July 2026; Published: 13 August 2026

ABSTRACT: Classical digital signature schemes such as RSA and ECDSA are threatened by the development of quantum computers, prompting the need for post-quantum cryptography (PQC). Among the various PQC candidates, multivariate quadratic (MQ) signature schemes have attracted significant attention due to their small signature size and efficient signing and verification. Evaluating these schemes is challenging because compactness, computational efficiency, and effective security are objectives that often conflict with one another and cannot usually be optimized simultaneously. In this work, we adapt a systematic multi-objective evaluation framework to MQ-based signature schemes in the NIST additional digital signature standardization process. The framework is based on Pareto optimality and enables a structured comparison of trade-offs among public key size, signature size, computational performance, and effective security without relying on subjective weighting. By representing each scheme in a multi-dimensional objective space, the Pareto-based framework used in this work identifies non-dominated schemes. Using this framework, we conduct a comparative analysis of four representative MQ-based signature schemes—UOV, MAYO, QR-UOV, and SNOVA. We analyze their design principles, parameter choices, performance characteristics, and security against known cryptanalytic attacks. To improve reproducibility and reduce hardware-related bias, all schemes are executed and evaluated on a unified platform, and performance metrics are obtained through repeated experiments. The results reveal that no single scheme simultaneously optimizes all criteria, highlighting inherent trade-offs among MQ-based designs. This Pareto-based evaluation framework provides a structured and reproducible approach for evaluating PQC schemes and offers practical insights for selecting appropriate schemes under different system and performance requirements.

KEYWORDS: Post-quantum cryptography; multivariate signatures; Pareto optimality; performance evaluation; multi-objective analysis

1 Introduction

Digital signature schemes play a central role in modern computing systems, ensuring the authenticity and integrity of digital information. In modern mobile and distributed systems, digital signatures are widely used for device authentication, application integrity protection, and secure service provisioning, including mobile payments, secure messaging, and cloud access. As these systems continue to scale, practical deployment increasingly requires cryptographic solutions that balance security, computational efficiency, and resource consumption under real-world constraints such as latency, storage, and processing capability. In classical cryptography, widely deployed schemes such as RSA [1] and elliptic curve signatures [2] derive

their security from number-theoretic problems. However, with the advent of large-scale quantum computers, Shor's algorithm [3] makes these assumptions no longer reliable.

As mobile and distributed systems increasingly rely on public-key infrastructures (PKI) for secure communication, the transition to quantum-resistant digital signatures has become essential. Once quantum computers can efficiently solve integer factorization or discrete logarithm problems, the cryptographic foundations of mobile internet security will be at risk. Therefore, post-quantum digital signatures are vital to guarantee long-term security for mobile and cloud communication infrastructures. In this context, selecting suitable PQC signature schemes has become a critical challenge, as different designs exhibit significant trade-offs across multiple performance and security dimensions. To address this broader challenge, the National Institute of Standards and Technology (NIST) launched the Post-Quantum Cryptography (PQC) standardization project [4] in 2016.

After three rounds of evaluation, in July 2022, NIST announced its first set of selected algorithms [5]: ML-KEM (CRYSTALS-Kyber) [6] for key encapsulation, and ML-DSA (CRYSTALS-Dilithium) [7], FN-DSA (Falcon), together with SLH-DSA (SPHINCS+) [8] for digital signatures. Except for SLH-DSA, all selected schemes rely on the hardness of the lattice problems. In March 2025, NIST selected HQC (Hamming Quasi-Cyclic), based on error-correcting codes, as a new key encapsulation mechanism (KEM) algorithm for standardization [9]. HQC will serve as an alternative to ML-KEM, improving diversity in the underlying mathematical foundations, with its standard expected by 2027.

Despite these advancements, the reliance on lattice-based constructions has raised concerns regarding potential future breakthroughs that may weaken their security assumptions. To mitigate this risk, NIST initiated an additional digital signature standardization process [10], aiming to explore alternative hardness assumptions beyond the lattices. By June 2023, 40 schemes were accepted as first-round candidates, covering a wide range of cryptographic paradigms, including multivariate, code-based, isogeny-based, and MPC-in-the-head constructions. In 2024, NIST selected 14 candidates for the second round of the additional digital signature standardization process [11]. Subsequently, in 2026, NIST released the second-round status report and selected nine of them to advance to the third round [12]. Table 1 summarizes the second-round candidates, with the third-round candidates highlighted in bold.

Table 1: Digital signature candidates in the NIST additional digital signature process.

Category	Schemes				
Code-Based	CROSS	LESS			
Symmetric-Based	**FAEST**				
Lattice-Based	**HAWK**				
Isogeny-Based	**SQIsign**				
MPC-in-the-Head	Mirath	**MQOM**	PERK	RYDE	**SDitH**
Multivariate	**UOV**	**MAYO**	**QR-UOV**	**SNOVA**	

Note: Candidates shown in bold were selected to advance to the third round.

Among these candidates, multivariate public-key cryptosystems (MPKCs) [13] represent one of the most promising alternatives, due to their compact signatures and efficient verification. Their security relies on the hardness of solving systems of multivariate polynomial equations. The Oil and Vinegar (OV) scheme [14], proposed by Patarin in 1997, is one of the first and most widely studied multivariate signature schemes. However, the original OV parameterization was shown to be insecure by the Kipnis–Shamir attack [15] in 1998, which broke the scheme in polynomial time. To overcome the shortcomings, the Unbalanced Oil and

Vinegar (UOV) variant [16] was introduced and remains a strong candidate within the class of multivariate signature schemes.

Within the NIST additional digital signature standardization process, UOV [17], MAYO [18], QR-UOV [19], and SNOVA [20] are the four candidate schemes based on the MQ problem. They were selected as third-round candidates in 2026 according to the NIST second-round status report [12]. Because the present study was conducted before the release of official third-round specification documents and reference implementations, our parameter comparison and experimental evaluation are based on the official second-round specification documents and corresponding second-round implementation packages. All four schemes follow the same oil-vinegar design principle: randomly selecting vinegar variables and solving a linear system of oil variables to generate the signature. However, they differ in their structural modifications: MAYO uses a whipping transformation to compress the key size, QR-UOV represents the public key using block matrices over a quotient ring, and SNOVA is designed to operate over non-commutative rings to reduce the public key size and improve efficiency. A central challenge is the assessment of their concrete security, due both to the lack of formal security proofs for the underlying UOV framework and to the additional algebraic structure introduced by each variant.

The security of these candidates is usually evaluated with respect to known attacks, which are generally classified into two categories: key-recovery attacks (e.g., Kipnis–Shamir [15], intersection [21], and MinRank attacks [21]) and forgery attacks (e.g., collision and direct attacks). The key-recovery attacks aim to recover the secret key from the known public key, and the forgery attacks aim to forge a signature-message pair that passes the signature verification.

So far, there have been several attempts to attack these MQ-based signature schemes. In 2021, Beullens proposed a rectangular MinRank attack [21], which significantly reduced the security of Rainbow, a multi-layer UOV construction. Subsequent studies showed that MinRank-type techniques are also relevant to small-public-key UOV variants. Furue and Ikematsu demonstrated that the rectangular MinRank attack can be applied to MAYO and QR-UOV [22], while Suzuki et al. unified several key-recovery attacks within the Jacobian framework and proposed an extended rectangular MinRank attack against UOV and its variants [23]. Although the currently proposed parameter sets remain resistant to this extended attack, these results show that MinRank-derived bounds must be considered when evaluating the concrete security of UOV-based signatures.

During the second-round evaluation, the cryptanalysis of UOV-family schemes progressed further. Ran proposed a wedge-product attack exploiting the exterior-algebra structure of polar forms over characteristic-two fields, which reduces the estimated security of several UOV and MAYO parameter sets [24]. This line of analysis was later generalized to odd-characteristic UOV-family schemes using symmetric algebra, making it relevant to the analysis of QR-UOV, although it does not outperform existing key-recovery attacks for the proposed QR-UOV parameters [25]. Furue and Ikematsu further reformulated these attacks using p^{ℓ}-truncated polynomial rings and applied the framework to reconciliation and intersection attacks [26].

SNOVA has also received substantial cryptanalytic attention. Several works analyzed its additional algebraic structure, including the attacks of Beullens [27], Cabarcas et al. [28], Ikematsu and Akiyama [29], Li and Ding [30], and Nakamura et al. [31]. More recent attacks further exploit SNOVA's structured representation, including its block-ring structure [32]. In particular, the wedge-product attack adapted to SNOVA exploits its block-ring structure and reduces the estimated security of several proposed parameter sets [32]. Therefore, SNOVA's compact key and signature sizes should be interpreted together with its less stable security margin. Following the recent cryptanalysis of SNOVA, Ding et al. further reformulated SNOVA using ring-equation, whipping, and tensor representations, and proposed more flexible parameter choices with competitive key and signature sizes [33]. Since these parameters are post-Round-2 demonstrative parameters rather than

official submissions, they are discussed as a possible direction for future SNOVA updates rather than being included in our comparison.

These results indicate that the cryptanalysis of MQ-based signature schemes is still rapidly progressing, and their concrete security requires careful re-evaluation in the context of post-quantum standardization. Recent studies include not only classical attacks, but also quantum analyses of specific attack procedures. In the quantum setting, security should not be estimated merely by applying a uniform square-root reduction to all classical attacks. For example, Cho and Seo proposed a quantum rectangular MinRank attack targeting multivariate signature schemes based on the multi-layer UOV structure [34], while May et al. provided updated classical and quantum forgery estimates for MAYO, QR-UOV, and SNOVA based on improved algorithms for the underdetermined MQ problem [35]. In this work, we use recent classical cryptanalytic estimates to define the effective-security metric for the Pareto analysis, and discuss quantum-adjusted bounds separately to avoid combining classical and quantum cost models into a single scalar value.

Beyond mathematical cryptanalysis, recent studies have shown that MQ-based signatures may also be affected by implementation-level threats, including side-channel leakage and fault-injection attacks. For example, fault-injection attacks have been studied against MAYO by manipulating vinegar-related randomness, while fault analysis of SNOVA shows that permanent or transient faults during signature generation may leak information about the secret structure [36,37]. In parallel, recent work on masked Gaussian elimination has shown that the linear-system solving step used in UOV-based signing requires careful side-channel protection, and a recent systematization study further analyzes physical attacks and countermeasures for UOV, MAYO, QR-UOV, and SNOVA [38,39]. Although these physical attack vectors are not directly modeled in the present work, they indicate that physical-attack resilience could serve as an additional Pareto dimension in future deployment-oriented extensions of the adapted framework.

Taken together, the diversity of recent cryptanalytic results and deployment constraints shows that the evaluation of PQC signature schemes naturally takes place in a high-dimensional objective space. Public key size, signature size, computational cost, and effective security differ in scale, interpretation, and deployment relevance. Improvements in one dimension often lead to degradation in another. For example, a scheme may achieve fast signing and verification at the cost of large public keys, whereas another scheme may reduce key size by introducing additional algebraic structure and higher computational overhead. Moreover, different MQ-based solutions may be vulnerable to different types of mathematical or implementation-level attacks. Therefore, comprehensive evaluation requires not only reporting individual benchmark values, but also analyzing how these conflicting indicators jointly shape the trade-off space.

Despite extensive research on MQ signature schemes, existing evaluations [40] are typically conducted in a metric-wise manner, where parameters such as public key size, signature size, attack complexity, and computational performance are compared separately. However, such isolated comparisons do not fully capture how these conflicting indicators jointly shape the trade-off space. As a result, it remains difficult to systematically assess the overall efficiency-security balance of different schemes. This motivates a complementary evaluation perspective that explicitly models the trade-offs among multiple objectives.

It should also be noted that the NIST additional digital signature standardization process already performs a comprehensive assessment of candidate schemes by considering security, performance, and standardization-related factors [11,12]. However, the goal of the NIST assessment is to support standardization decisions rather than to provide a formal multi-objective dominance analysis among schemes within a specific algebraic family. Similarly, benchmark-style evaluations usually report individual metrics such as key size, signature size, or running time, but they do not explicitly characterize how these metrics jointly determine non-dominance or application-dependent trade-offs. Therefore, the Pareto-based framework in

this work is intended to complement existing NIST and benchmark-style evaluations by providing an explicit multi-objective representation of the trade-offs among MQ-based candidates.

Beyond official standardization reports, recent work shows that PQC evaluation is moving beyond isolated algorithmic benchmarks toward deployment-oriented and framework-based assessment. For example, Post-Quantum Cryptographic Library Evaluation Operator (PQC-LEO) provides an evaluation framework for automating computational and networking performance measurements across x86 and ARM architectures [41], while recent cross-platform benchmarking studies evaluate NIST-selected PQC algorithms across heterogeneous computing environments [42]. Other studies further benchmark post-quantum signatures on RISC-V platforms [43] and evaluate NIST-standardized PQC algorithms on resource-constrained ARM Cortex-M0+ IoT devices, including memory and energy measurements [44]. These studies highlight the importance of architecture-dependent performance, memory footprint, energy consumption, communication overhead, and deployment constraints in practical PQC evaluation. However, they mainly provide benchmarking tools or empirical performance data and do not specifically analyze non-dominance relationships among MQ-based digital signature candidates.

To address this limitation, we introduce a Pareto-based multi-objective modeling framework for the evaluation of MQ signature schemes [45]. In this framework, each scheme is represented as a point in a multi-dimensional objective space, where public key size, signature size, and computational performance are minimized, and security—represented by the complexity of the best-known attack—is maximized. Rather than solving an optimization problem, the adapted framework performs a comparative analysis over a finite set of candidate schemes using Pareto dominance, allowing the identification of non-dominated solutions without introducing subjective weighting factors. To support a reproducible performance evaluation, we execute all considered schemes on a unified platform and measure their computational performance through repeated experiments. In particular, we evaluate three core operations of digital signatures—key generation, signing, and verification—by averaging the results over 1000 runs. The resulting measurements provide a common experimental baseline for the subsequent multi-objective analysis. We note that while the performance evaluation is conducted on an x86 platform for reproducibility and cross-scheme fairness, the implications of our results for mobile and embedded deployments are discussed in Section 3.

Based on this framework, we conduct a structured comparative analysis of MQ-based signature schemes across different security levels. The proposed approach enables a systematic exploration of trade-offs and provides deeper insights into the design space of multivariate cryptography.

The main contributions of this work can be summarized as follows:

- We present a comprehensive comparative analysis of four MQ-based signature schemes, covering their structural characteristics, parameter choices, performance, and resistance to known attacks.
- We formulate the comparison of MQ-based signature candidates as a high-dimensional multi-objective evaluation problem and apply a Pareto-based framework that jointly considers public key size, signature size, computational cost, and attack complexity. All performance measurements are obtained on a unified experimental platform to provide a reproducible basis for comparison.
- We identify Pareto-optimal schemes across different security levels and systematically reveal the structural trade-offs inherent in MQ-based designs.
- We provide a structured multi-objective perspective for evaluating PQC signature schemes, offering application-oriented insights to support future scheme design, parameter selection, and deployment-oriented decision-making.

The remainder of this paper is organized as follows. Section 2 introduces the necessary preliminaries. Section 3 presents a comparative analysis of multivariate signature schemes. Section 4 analyzes the

security against known attacks. Section 5 introduces the adapted Pareto-based evaluation framework and discusses the results. Finally, Section 6 concludes the paper.

2 Preliminaries

In this section, we introduce the concept of multivariate quadratic (MQ) maps and the associated MQ problem, which provides the computational hardness assumption underlying multivariate cryptography. We then describe the original Unbalanced Oil and Vinegar (UOV) scheme, which serves as the basis for the MQ-based schemes considered in this work. Finally, we explain the concept of Pareto optimality, which forms the adapted multi-objective evaluation framework.

2.1 Multivariate Quadratic Map and MQ Problem

Multivariate quadratic cryptography is based on the multivariate quadratic map. A multivariate quadratic map $\mathcal{P}$ with m components and n variables is defined as a sequence

$$p_1(x_1, \cdots, x_n), \cdots, p_m(x_1, \cdots, x_n) \tag{1}$$

where each p_i is a quadratic polynomial in n variables over a finite field $\mathbb{F}_q$.

Finding a preimage of a given vector $\mathbf{b} = (b_1, \cdots, b_m) \in \mathbb{F}_q^m$ under $\mathcal{P}$ is equivalent to solving the system of quadratic equations

$$\begin{cases} p_1(x_1, \cdots, x_n) = b_1 \\ \vdots \\ p_m(x_1, \cdots, x_n) = b_m. \end{cases} \tag{2}$$

This problem is known as the MQ problem, which is generally considered NP-hard over finite fields. To date, no polynomial-time algorithm is known for solving the MQ problem in the average case. The best-known classical algorithms for solving the MQ problem are based on Gröbner basis methods such as F_4/F_5 and the XL algorithm. In the quantum setting, no efficient quantum algorithm is known beyond the generic quadratic speedup provided by Grover's algorithm.

2.2 The Original UOV Signature Scheme

As all four digital signature candidates are variations of the original UOV scheme, it is necessary to describe the basic principle of the original UOV scheme.

The UOV scheme was introduced in 1999 by Kipnis et al. [16], and it is an optimized version of the balanced oil and vinegar (OV) scheme [14] in 1997. For cryptographic purposes with a public key, the UOV scheme treats a multivariate quadratic map $\mathcal{P} : \mathbb{F}_q^n \to \mathbb{F}_q^m$ as the public key, which should be well designed so that it vanishes on a hidden linear space $\mathcal{O} \subset \mathbb{F}_q^n$. This map should have a trapdoor structure to ensure that whoever has the trapdoor information can efficiently find the preimage $\mathcal{P}^{-1}$; otherwise, the solution will be computationally difficult.

The public map $\mathcal{P}$ consists of two multivariate polynomial maps: the central map $\mathcal{F}$ and the linear map $\mathcal{T}$, where $\mathcal{P} = \mathcal{F} \circ \mathcal{T}$. In the UOV scheme, $(\mathcal{F}, \mathcal{T})$ is the secret key, which includes the trapdoor details.

The central map $\mathcal{F}$ is defined as $\mathcal{F} : \mathbb{F}_q^n \to \mathbb{F}_q^m$:

$$f_k(x_1, \cdots, x_n) = \sum_{i=1}^{v} \sum_{j=i}^{n} \alpha_{i,j}^{(k)} x_i x_j, \tag{3}$$

where $v = n - m, 1 \le k \le m$. Let $(x_1, \cdots, x_n) \in \mathbb{F}_q^n$, the first v variables are called vinegar variables, the remaining m variables are called oil variables. The central map is constructed in such a way that there is no quadratic term consisting of the two oil variables. Therefore, when the vinegar variables are fully fixed, the central map is a linear system in the oil variables and can be efficiently solved. The linear map $\mathcal{T}$ is invertible and serves to hide the structure of the central map.

The original UOV signature scheme follows the hash-and-sign paradigm, as summarized in Algorithm 1.

Algorithm 1: Original UOV signature scheme

KeyGen (n, m, q)
Choose m polynomials $f_1, \ldots, f_m$ uniformly at random
Set the central map $\mathcal{F} = f_1, \ldots, f_m$
Choose an invertible linear map $\mathcal{T} : \mathbb{F}_q^n \to \mathbb{F}_q^n$ uniformly at random
Compute $\mathcal{P} = \mathcal{F} \circ \mathcal{T}$
pk = $\mathcal{P}$
sk = $(\mathcal{F}, \mathcal{T})$
return (pk, sk)
Sign(sk = $(\mathcal{F}, \mathcal{T})$, μ: message):
$t = \text{Hash}(\mu)$
Compute the preimage $u \in \mathbb{F}_q^n$ of t such that $\mathcal{F}(u) = t$
$s = \mathcal{T}^{-1}(u)$
$\sigma = s$
return σ
Verify(pk = $\mathcal{P}$, μ, $\sigma = s \in \mathbb{F}_q^n$):
$t = \text{Hash}(\mu)$
$t' = \mathcal{P}(s)$
return $(t == t')$

Although Algorithm 1 describes the original UOV construction, the NIST candidates considered in this work adopt different compact representations and algebraic modifications. Therefore, the scheme-specific asymptotic time and space complexities of UOV, MAYO, SNOVA, and QR-UOV are summarized in Section 3.1, where the concrete parameter notation of the evaluated candidates is introduced.

2.3 Pareto Optimality and Multi-Objective Comparison

In practical cryptographic evaluation, multiple performance and security criteria must be considered simultaneously. These criteria are often inherently conflicting; for example, reducing key size may negatively impact security, while increasing security typically requires larger parameters. As a result, it is generally not appropriate to aggregate all criteria into a single scalar objective without introducing subjective weighting.

To address this issue, we adopt a framework of multi-objective optimization based on Pareto optimality. Let $\mathcal{S} = \{S_1, S_2, \ldots, S_N\}$ denote a finite set of candidate schemes. Each scheme S_i is associated with a vector of objective values:

$$\mathbf{f}(S_i) = \big(f_1(S_i), f_2(S_i), \ldots, f_k(S_i)\big), \tag{4}$$

where each function f_ℓ represents a specific evaluation criterion.

For two schemes $S_i, S_j \in \mathcal{S}$, S_i is said to Pareto dominate S_j, denoted by $S_i \prec S_j$, if the following conditions hold:

$$\begin{cases} f_\ell(S_i) \le f_\ell(S_j), \quad \forall \ell \in \{1, \ldots, k\}, \\ \exists \ell \in \{1, \ldots, k\} \text{ such that } f_\ell(S_i) < f_\ell(S_j). \end{cases} \tag{5}$$

where all objectives are assumed to be expressed in a consistent direction (e.g., all minimized). This means that S_i is no worse than S_j in all criteria and strictly better in at least one.

A scheme S_i is Pareto-optimal if there does not exist any $S_j \in \mathcal{S}$ such that $S_j \prec S_i$. The set of all Pareto-optimal schemes forms the Pareto front:

$$\mathcal{P} = \left\{ S_i \in \mathcal{S} \mid \nexists S_j \in \mathcal{S},\ S_j \prec S_i \right\}. \tag{6}$$

The Pareto front characterizes the set of non-dominated solutions that represent optimal trade-offs among conflicting objectives. Unlike single-objective optimization, this framework identifies a set of equally optimal solutions, rather than a single optimum. This property makes Pareto analysis particularly suitable for cryptographic evaluations, where different schemes often exhibit trade-offs between efficiency metrics and security guarantees. By avoiding the use of weighted aggregation, Pareto optimality provides a structured and transparent way to compare candidate schemes under multiple criteria. It reduces the subjectivity associated with assigning numerical weights to different objectives, but it does not eliminate all modeling choices, since the selection of metrics and the construction of each metric still influence the resulting comparison.

In this work, Pareto optimality serves as the foundation of our multi-objective modeling framework. It enables a systematic and transparent comparison of MQ-based signature schemes without relying on subjective weighting, thereby revealing the inherent trade-offs among compactness, computational efficiency, and security.

3 Multivariate Signature Candidates in NIST Standardization Process

In this section, we provide a comparative analysis of the four MQ signature schemes selected in the NIST additional digital signature standardization process, namely UOV, MAYO, SNOVA, and QR-UOV. All these schemes are derived from the original UOV construction but adopt different design strategies to balance efficiency and security. This section not only analyzes their structural characteristics, but also provides experimentally measured performance data under a unified platform, which will serve as the basis for the multi-objective modeling framework introduced in Section 5.

3.1 Comparison of Characteristics

The four MQ-based signature candidates share the same underlying principle of the original UOV scheme but introduce distinct structural modifications to improve efficiency and reduce key size.

In the MAYO scheme, the dimension of the oil space is less than the number of the quadratic equations. This design increases the difficulty of recovering the hidden oil space $\mathbf{O}$ from the public polynomials, allowing the use of smaller parameters without sacrificing security. To compensate for the resulting underdetermined system, MAYO incorporates a public matrix E into the linear system, ensuring solvability while maintaining compact key representation.

SNOVA adopts non-commutative ring operations to enrich the algebraic structure and reduce the public key size while still maintaining the advantage of short signatures. Instead of simply generalizing UOV to rings, SNOVA employs randomized matrix transformations before and after the core computation to

eliminate exploitable sparsity patterns. To further address the problem of large public key sizes, SNOVA also applies the technique of key-randomness alignment, where part of the randomness of the private key is shifted into the public key. Combined with the use of pseudorandom number generators seeded for multivariate cryptosystems over rings, this technique achieves a substantial reduction in key sizes while preserving security at the same level.

QR-UOV exploits block-matrix representations over a quotient ring $\mathbb{F}_q[x]/(f)$, allowing compact encoding of the public key. By representing elements as polynomial matrices and compressing them via coefficients, the storage cost is reduced significantly. To address the instability of these polynomial matrices under transposition, QR-UOV introduces an invertible matrix to preserve symmetry under transposition. The choice of the modulus polynomial f is critical, as reducible polynomials may introduce structural vulnerabilities. Overall, QR-UOV achieves smaller public keys while maintaining security through careful algebraic design.

Based on the design principles outlined above, Table 2 shows the recommended Round 2 parameter sets for the four MQ-based signature candidates. Here, n and m denote the total number of variables and equations, respectively. v and o represent the number of vinegar and oil variables; k is the whipping parameter in MAYO. l in SNOVA is the size of the non-commutative ring, while in QR-UOV is the block size parameter associated with the irreducible polynomial used in the quotient ring construction.

Table 2: Recommended round 2 parameter sets for the MQ-based signature candidates.

Sec. Level	Scheme	Field	n	m	v	o	k	l
I	UOV-Ip	$\mathbb{F}_{256}$	112	44	68	44	–	–
	UOV-Is	$\mathbb{F}_{16}$	160	64	96	64	–	–
	MAYO1	$\mathbb{F}_{16}$	86	78	78	8	10	–
	MAYO2	$\mathbb{F}_{16}$	81	64	64	17	4	–
	SNOVA-a	$\mathbb{F}_{16}$	54	17	37	17	–	2
	SNOVA-b	$\mathbb{F}_{16}$	33	8	25	8	–	3
	SNOVA-c	$\mathbb{F}_{16}$	29	5	24	5	–	4
	QR-UOV	$\mathbb{F}_{127}$	210	54	156	54	–	3
III	UOV	$\mathbb{F}_{256}$	184	72	112	72	–	–
	MAYO	$\mathbb{F}_{16}$	118	108	108	10	11	–
	SNOVA-a	$\mathbb{F}_{16}$	81	25	56	25	–	2
	SNOVA-b	$\mathbb{F}_{16}$	60	11	49	11	–	3
	SNOVA-c	$\mathbb{F}_{16}$	45	8	37	8	–	4
	SNOVA-d	$\mathbb{F}_{16}$	29	5	24	5	–	5
	QR-UOV	$\mathbb{F}_{127}$	306	78	228	78	–	3
V	UOV	$\mathbb{F}_{256}$	244	96	148	96	–	–
	MAYO	$\mathbb{F}_{16}$	154	142	142	12	12	–
	SNOVA-a	$\mathbb{F}_{16}$	108	33	75	33	–	2
	SNOVA-b	$\mathbb{F}_{16}$	81	15	66	15	–	3
	SNOVA-c	$\mathbb{F}_{16}$	70	10	60	10	–	4
	SNOVA-d	$\mathbb{F}_{16}$	35	6	29	6	–	5
	QR-UOV	$\mathbb{F}_{127}$	411	105	306	105	–	3

Table 3 summarizes the key generation, signing, and verification procedures of the four MQ-based signature candidates. In this table, *csk* and *cpk* denote the compact secret key and compact public key, respectively; σ represents the signature; $\mathbf{v}$ and $\mathbf{o}$ correspond to the vinegar and oil variables; $\mathbf{x}$ is an intermediate vector used for computing the oil variables; $\mathbf{O}$ denotes the oil space; $\mathbf{s}$ denotes the resulting signature vector; μ is the message to be signed; $salt$ is a random value providing additional randomness in the signing process; $H(\cdot)$ stands for a hash function; and P_i (or $P_i^{(3)}$, P_i^{22} depending on the scheme) denote the part of the quadratic polynomials forming the public map $\mathcal{P}$ (or $\mathcal{P}^*$) used during verification.

Table 3: Key generation, signing, and verification procedures of the MQ-based signature candidates.

Scheme	Key Generation	Signature σ	Verification
UOV	$csk = seed_{sk}$,	$\sigma = (\mathbf{s}, salt)$,	$t = H(\mu \Vert salt),\ t' = \mathcal{P}(s)$
	$cpk = (seed_{pk}, P_i^{(3)})$	$\mathbf{s} = \mathbf{v} + \mathbf{O} \cdot \mathbf{x}$	Check if $t = t'$
MAYO	$csk = seed_{sk}$,	$\sigma = (\mathbf{s}, salt)$,	$t = H(\mu \Vert salt),\ t' = \mathcal{P}^*(s)$
	$cpk = (seed_{pk}, P_i^{(3)})$	$\mathbf{s} = \mathbf{v} + \mathbf{O} \cdot \mathbf{x}$	Check if $t = t'$
SNOVA	$csk = (seed_{sk}, seed_{pk})$	$\sigma = (\mathbf{s}, salt)$,	$t = H(\mu \Vert salt),\ t' = \mathcal{P}(s)$
	$cpk = (seed_{pk}, P_i^{22})$	$\mathbf{s} = T^{-1} \cdot (\mathbf{v} \Vert \mathbf{o})$	Check if $t = t'$
QR-UOV	$csk = (seed_{sk}, seed_{pk})$	$\sigma = (\mathbf{s}, salt)$,	$t = H(\mu \Vert salt),\ t' = \mathcal{P}(s)$
	$cpk = (seed_{pk}, P_i^{22})$	$\mathbf{s} = T^{-1} \cdot (\mathbf{v} \Vert \mathbf{o})$	Check if $t = t'$

As shown in Table 3, all four schemes share a common theoretical foundation derived from the UOV framework. They follow the hash-and-sign paradigm: a message μ is first hashed with a random salt, vinegar variables are randomly chosen, and the oil variables are obtained by solving a linear system. The secret keys are generated from short random seeds, while the public keys consist of quadratic polynomials P_i. Verification is carried out by evaluating the public map $\mathcal{P}$ on the signature and checking consistency with the hash value. Therefore, all candidates adhere to the same design philosophy: compact key derivation from seeds, randomized signing through salt and vinegar variables, and verification via polynomial evaluation.

Despite this common foundation, the schemes differ in the way they optimize compactness, efficiency, and algebraic structure. The differences among the four candidates can be analyzed from three main perspectives: key generation, signing, and verification.

Key Generation. In the key generation process, all schemes utilize the seed of the key instead of the entire key to reduce the storage requirements. For the secret key, UOV and MAYO both rely on a single seed $seed_{sk}$ to derive the secret key, which ensures compactness and simplicity in secret key storage. By contrast, SNOVA and QR-UOV use two independent seeds $(seed_{sk}, seed_{pk})$ to form the private key. This design ensures independence between the randomness used for the secret and the public components, but at the cost of slightly larger secret keys. For the public key, the quadratic polynomials are often represented in block-matrix form. In UOV and MAYO, one uses

$$P_i = \begin{bmatrix} P_i^{(1)} & P_i^{(2)} \\ 0 & P_i^{(3)} \end{bmatrix}, \tag{7}$$

where 0 denotes a zero block. While in SNOVA and QR-UOV, it is shown as

$$P_i = \begin{bmatrix} P_i^{11} & P_i^{12} \\ P_i^{21} & P_i^{22} \end{bmatrix}. \tag{8}$$

To reduce the size of the public key, all four schemes store only the lower-right block (P_i^{22} in SNOVA/QR-UOV, and $P_i^{(3)}$ in UOV/MAYO) with $seed_{pk}$. Specifically, in UOV and MAYO, one first uses $seed_{sk}$ to generate $seed_{pk}$ and the oil space $\mathbf{O}$. Then, from $seed_{pk}$, the blocks $P_i^{(1)}$ and $P_i^{(2)}$ are derived. Finally, the lower-right block of the public key is computed by

$$P_i^{(3)} := \text{Upper}\Big(-\mathbf{O}^\top P_i^{(1)}\mathbf{O} - \mathbf{O}^\top P_i^{(2)} \Big), \tag{9}$$

which forms the compressed public key. In contrast, SNOVA and QR-UOV first use $seed_{sk}$ to generate a linear map T, and $seed_{pk}$ to generate the three blocks P_i^{11}, P_i^{12}, and P_i^{21}. The final public key P_i^{22} is then computed using these parts.

Signing. In the signing phase, both UOV and MAYO operate by first choosing random vinegar variables $\mathbf{v}$ and then forming a linear system in the intermediate variables $\mathbf{x}$ using the blocks $P_i^{(1)}$, $P_i^{(2)}$, and the matrix $\mathbf{O}$ derived from the $seed_{sk}$. Once $\mathbf{x}$ is obtained by solving this system, the signature vector is computed as

$$\mathbf{s} = \mathbf{v} + \mathbf{O} \cdot \mathbf{x}, \tag{10}$$

and combined with a random salt value to produce the final signature. A distinguishing feature of MAYO is the use of the so-called whipping transformation, which incorporates a fixed public matrix E into the construction of the linear system, thereby increasing the number of variables relative to equations and ensuring solvability. At the same time, this transformation makes the recovery of the hidden oil space significantly harder, thus allowing smaller parameters while still enabling a compact representation of the public map.

In contrast, the signing procedure of SNOVA and QR-UOV is more like the original UOV scheme. Instead of employing the oil space $\mathbf{O}$, these schemes directly assign random vinegar values and solve the resulting linear system to obtain the oil variables. The final signature vector is then derived by applying the linear map T to the concatenation of vinegar and oil variables, and combined with a random salt value to form the complete signature.

Verification. All schemes verify signatures by evaluating a public map on the signature and checking its consistency with the hashed message $H(\mu||salt)$. However, the form of the public map differs. UOV, SNOVA, and QR-UOV directly use $\mathcal{P}$, while MAYO employs a modified map $\mathcal{P}^*$ due to the whipping transformation by utilizing the public matrix E.

Asymptotic Time and Space Complexity. To provide a theoretical context for the empirical cycle counts, we summarize the asymptotic time and space complexity of the four evaluated MQ-based signature schemes. The analysis counts arithmetic operations over the corresponding base finite field and ignores fixed-size constants such as seed length, salt length, and hash output length. Standard Gaussian elimination is assumed for solving linear systems. For UOV, SNOVA, and QR-UOV, the number of oil variables is equal to the number of equations, i.e., $o = m$. For MAYO, o denotes the dimension of the oil space and is smaller than m; the number of vinegar variables is $v = n - o$. The parameters k and ℓ denote the MAYO whipping parameter and the ring parameter of SNOVA and QR-UOV, respectively. These asymptotic estimates provide a theoretical baseline for interpreting the measured cycle counts in Section 3.2.

The complexity estimates in Table 4 follow from the matrix dimensions in the key generation, signing, and verification algorithms of the corresponding schemes. In UOV, key generation is dominated by computing the lower-right public-key block $P_i^{(3)} = \mathrm{Upper}(-O^T P_i^{(1)} O - O^T P_i^{(2)})$ for all $i = 1, \ldots, m$, where $O \in \mathbb{F}_q^{v \times m}$, $P_i^{(1)} \in \mathbb{F}_q^{v \times v}$, and $P_i^{(2)} \in \mathbb{F}_q^{v \times m}$. This yields $O(m^2v^2 + m^3v)$. Signing constructs an $m \times m$ matrix L, computes $y = v^T P_i^{(1)} v$, and solves $Lx = t - y$ by Gaussian elimination, giving $O(mv^2 + vm^2 + m^3)$. Verification evaluates m quadratic forms in n variables, giving $O(mn^2)$.

Table 4: Asymptotic time and space complexity of the evaluated MQ-based signature schemes.

Scheme	Operation	Time Complexity	Space Complexity
UOV	KeyGen	$O(m^2v^2 + m^3v)$	pk: $O(m^3)$
	Sign	$O(mv^2 + vm^2 + m^3)$	signature: $O(n)$
	Verify	$O(mn^2)$	
MAYO	KeyGen	$O(mv^2o + mvo^2)$	pk: $O(mo^2)$
	Sign	$O(k^2mv^2 + k^2m^2o + km^2o)$	signature: $O(kn)$
	Verify	$O(k^2mn^2)$	
SNOVA	KeyGen	$O(\ell^3o^2vn)$	pk: $O(m^3\ell^2)$
	Sign	$O(ov^2\ell^5 + o^2n\ell^6)$	signature: $O(n\ell^2)$
	Verify	$O(on^2\ell^5)$	
QR-UOV	KeyGen	$O(o^2vn/\ell)$	pk: $O(o^3/\ell)$
	Sign	$O(ov^2 + o^2v + o^3)$	signature: $O(n)$
	Verify	$O(on^2)$	

MAYO follows a similar compact Oil-and-Vinegar key-generation structure, but the oil-space dimension o is smaller than m. Therefore, computing $P_i^{(3)}$ for all m public equations costs $O(mv^2o + mvo^2)$. During signing, MAYO evaluates the whipped map $P^*(x_1, \ldots, x_k)$, which contains $O(k^2)$ quadratic terms. Constructing the corresponding linear system contributes $O(k^2mv^2 + k^2m^2o)$, and solving the $m \times ko$ system contributes $O(km^2o)$. Verification similarly evaluates $O(k^2)$ quadratic terms, each involving m quadratic forms in n variables, resulting in $O(k^2mn^2)$.

SNOVA generalizes the UOV structure over the matrix ring $R = \mathrm{Mat}_{\ell \times \ell}(\mathbb{F}_q)$. Since each ring element is an $\ell \times \ell$ matrix, a naive ring multiplication costs $O(\ell^3)$ base-field operations. Consequently, the key-generation costs $O(\ell^3o^2vn)$. In signing, the vinegar-only part contains o equations, v^2 vinegar pairs, and $O(\ell^2)$ internal summation terms, leading to $O(ov^2\ell^5)$. The construction and solution of the oil-variable linear system over the base field involve $o\ell^2$ unknowns and contribute $O(o^2n\ell^6)$. Verification evaluates the public map with o output equations, n^2 variable pairs, and $O(\ell^2)$ matrix-ring summation terms, giving $O(on^2\ell^5)$.

QR-UOV uses a quotient-ring representation with $v = \ell V$, $o = m = \ell M$, and $n = \ell N$. Its key generation computes the lower-right block P_i^{22} through block operations over the quotient ring. For each i, this costs $O(V^2M + VM^2) = O(VMN)$ operations over $\mathbb{F}_{q^\ell}$. Converting to base-field operations gives $O(o^2vn/\ell)$. Signing follows the UOV inversion structure by fixing vinegar variables and solving an $o \times o$ linear system, giving $O(ov^2 + o^2v + o^3)$. Verification evaluates o quadratic forms in n variables, resulting in $O(on^2)$.

These structural differences directly affect the position of each scheme in the later Pareto space. UOV retains the most direct oil-vinegar structure among the four schemes, which leads to efficient signing and

verification, but also results in a large public-key representation. MAYO modifies the UOV framework through the whipping transformation. This design substantially reduces public-key size and improves key-generation efficiency, but introduces additional operations in signing and verification. SNOVA uses non-commutative ring representations and key-randomness alignment to obtain compact public keys and signatures; however, the additional algebraic structure also increases computational overhead and makes the scheme more sensitive to recent structure-specific cryptanalysis. QR-UOV uses quotient-ring block representations to reduce public-key size compared with UOV, but this ring-based structure introduces additional arithmetic overhead and leads to a balanced rather than strongly optimized profile. Therefore, the design principles introduced in this section should be interpreted not only as algebraic constructions but also as structural sources of the trade-offs later observed in the Pareto analysis.

3.2 Comparison of Performance

Since the performance results reported in the specification documents of the four schemes [17–20] were obtained on different platforms, a direct comparison is difficult. We therefore performed our own measurements on a unified platform equipped with a 13th Gen Intel Core i7-1360P processor @ 5.0 GHz and 16 GB RAM, running Ubuntu 22.04. For each scheme and parameter set, we measured the cycle counts of key generation, signing, and verification. Each experiment was repeated 1000 times, and the average value was reported. All implementations were built with GCC 11.4.0 using the reference build configurations provided by the corresponding official Round-2 repositories, and the optimization level `-O3` was enabled for all schemes. Among architecture-specific compiler flags, `-march=native` was enabled for MAYO and QR-UOV according to their build configurations, while the SNOVA and UOV reference benchmark binaries were built without this option. The `-march=native` option allows the compiler to generate code specialized for the host processor by enabling available instruction-set extensions and microarchitecture-specific tuning. Therefore, it may affect cycle-count measurements through more efficient instruction selection, instruction scheduling, and vectorization. Thus, the evaluation uses a common hardware platform, operating system, compiler version, and optimization level, while preserving the reference build configuration of each official implementation. However, because architecture-specific compiler flags are not fully identical across all schemes, the reported cycle counts should be interpreted as measurements of the available official implementations under their stated reference build configurations, rather than as a fully compiler-flag-normalized comparison. A stricter compiler-normalized benchmark would require rebuilding and remeasuring all schemes under a consistent architecture-specific configuration, for example by enabling or disabling `-march=native` uniformly for all candidates.

We also examined the constant-time (CT) status of the tested implementations through the official specifications, repositories, and available documentation. The four implementations do not provide the same level of explicit CT documentation. For MAYO, the specification [18] states that all implementations except the Sage textbook implementation are protected against software-level side-channel attacks by avoiding secret-dependent data indexing and secret-dependent control flow. This provides the clearest CT-oriented implementation statement among the four schemes considered in this work, although it should not be interpreted as a formal CT verification report. For QR-UOV, the official Round 2 repository explicitly lists constant-time implementation of the rejection-sampling and linear-equation-solver routines as pending work. This suggests that the tested QR-UOV implementation should not be regarded as having an implementation-wide CT guarantee for these routines. For UOV and SNOVA, we did not identify an explicit implementation-wide CT guarantee or a formal CT verification report in the official documentation available to us. This does not prove that the implementations are non-constant-time; rather, it means that their CT status could not be definitively established from the available documentation. Accordingly, the reported

cycle counts in Table 5 should be interpreted as measurements of the available official implementations under the stated compilation environment, rather than as a fully constant-time-normalized or uniformly side-channel-protected benchmark.

Table 5: Round 2 parameter and performance comparison of the MQ-based signature candidates (sizes in bytes, performance in cycles).

Sec. Level	Scheme	$\|pk\|$	$\|sk\|$	$\|sign\|$	KeyGen	Sign	Verify
I	UOV-Ip	43,576	32	128	61,168k	1267k	175k
	UOV-Is	66,576	32	96	50,888k	813k	130k
	MAYO1	1420	24	454	1033k	3441k	1035k
	MAYO2	4912	24	186	1390k	2363k	566k
	SNOVA-a	9842	48	124	13,420k	17,164k	3583k
	SNOVA-b	2320	48	164.5	4437k	7976k	2482k
	SNOVA-c	1016	48	248	3548k	9088k	3469k
	QR-UOV	24,255	32	200	27,414k	3136k	2584k
III	UOV	189,232	32	200	391,159k	5180k	676k
	MAYO	2986	32	681	3330k	10155k	1937k
	SNOVA-a	31,266	48	178	62,407k	73,270k	11625k
	SNOVA-b	6005.5	48	286	26,746k	45,463k	10,597k
	SNOVA-c	4112	48	376	19,972k	41,539k	12,499k
	SNOVA-d	1578.5	48	378.5	7042k	22,836k	8606k
	QR-UOV	71,891	48	292	137,464k	10,939k	8694k
V	UOV	446,992	32	260	1,135,952k	11,726k	1692k
	MAYO	5554	40	964	7743k	21,565k	3163k
	SNOVA-a	71,890	48	232	189,592k	220,725k	27,274k
	SNOVA-b	15,203.5	48	380.5	86,852k	140,101k	25,919k
	SNOVA-c	8016	48	576	73,941k	143,855k	36,747k
	SNOVA-d	2716	48	453.5	14,285k	40,842k	14,462k
	QR-UOV	173,676	64	392	477,889k	26,587k	20,418k

Note: Blue values indicate the best result for each metric within the same security level, while red values indicate the least efficient result.

Table 5 presents the parameter sizes from the official Round-2 specification documents (version 2.0) of each scheme and performance results obtained under this unified experimental setup. To facilitate comparison, color coding is used in the table: blue text marks the best-performing results (smallest size or highest efficiency), while red text marks the least efficient results (largest size or lowest efficiency). This visual distinction helps highlight relative performance differences across schemes and parameter sets.

Key and Signature size. The four candidate schemes present a clear trade-off between public key compactness and signature size. UOV consistently produces the largest public keys because of its explicit representation of the quadratic system, while keeping the signature size moderate. In contrast, MAYO and SNOVA significantly reduce public key sizes through structural optimizations, such as the whipping transforms in MAYO and alignment of key randomness over a non-commutative ring in SNOVA. However, this compactness affects signatures in different ways: SNOVA achieves the shortest signatures, while MAYO achieves the largest signatures. QR-UOV provides a balanced design, producing smaller public keys than

UOV while maintaining moderate signature sizes. Overall, the results highlight a fundamental design paradox: schemes optimized for compact public keys typically produce larger signatures, while UOV maintains relatively small signatures at the expense of larger public keys.

Computational performance. The computational performance measured under the unified platform, equipped with a 13th Gen Intel Core i7-1360P processor (up to 5.0 GHz) and 16 GB of RAM, reveals distinct characteristics among the four schemes. For key generation, MAYO is consistently the most efficient scheme. Its cost is significantly lower than all other candidates, even at higher security levels. In contrast, UOV exhibits the highest key generation cost due to the explicit construction of large public maps. SNOVA and QR-UOV lie between these two extremes, with SNOVA generally more efficient than QR-UOV. For signing and verification, UOV consistently achieves the best performance. It remains the fastest scheme across all security levels, while MAYO shows moderate cost and QR-UOV slightly higher overhead. SNOVA exhibits the highest computational cost in both signing and verification.

Overall, the four schemes exhibit complementary performance characteristics: MAYO optimizes key generation, UOV optimizes signing and verification, SNOVA achieves compact representations at the cost of higher runtime overhead, and QR-UOV provides a balanced but non-dominant performance. These results provide a reproducible experimental dataset for the subsequent multi-objective analysis. In particular, the observed trade-offs among key size, signature size, and computational cost motivate the need for a multi-objective evaluation framework, which is formally developed in Section 5.

3.3 Discussion on Platform Generalizability

The performance experiments in this work were conducted on a unified x86 platform equipped with a 13th Gen Intel Core i7-1360P processor and 16 GB RAM. This setting provides a controlled and reproducible baseline for comparing the four MQ-based signature schemes on the tested x86 platform. However, as discussed in the introduction, an important deployment target for post-quantum digital signatures includes mobile devices, embedded systems, and IoT nodes. These platforms may differ substantially from the experimental platform in terms of microarchitecture, cache hierarchy, memory bandwidth, and available memory.

These architectural differences may affect the interpretation of the reported cycle counts, especially for memory-intensive schemes. Modern x86 processors typically provide relatively large cache hierarchies, high memory bandwidth, hardware prefetching, and aggressive out-of-order execution. These features can partially hide the cost of repeated accesses to large public keys, coefficient arrays, and matrix representations. As a result, schemes with large public keys or large working sets may appear more efficient on the x86 platform than they would on more constrained architectures.

This effect is particularly relevant for public-key-heavy schemes. For example, UOV at Security Level V has a public key size of approximately 447 KB. On an x86 processor with a relatively large cache hierarchy, repeated accesses to such data may be partly absorbed by the cache system. In contrast, on mobile or embedded processors with smaller caches and lower memory bandwidth, the same access pattern may lead to more frequent cache misses, higher memory traffic, and longer memory stalls. Consequently, the practical cost of memory-heavy key generation, signing, or verification routines may be larger than the cycle counts in Table 5.

Memory bandwidth is another important factor. Schemes based on large matrix representations or temporarily expanded forms may create working sets that are more easily handled on a high-bandwidth x86 platform than on a bandwidth-constrained embedded processor. For instance, although quotient-ring or block-matrix representations reduce storage requirements, their implementation may still involve

intermediate data movement and temporary buffers during computation. Such costs may become more visible on architectures with limited memory bandwidth.

For these reasons, the cycle counts reported in Table 5 should be interpreted as a platform-specific comparison under a controlled x86 environment, rather than as direct predictions of performance on mobile or embedded devices. A full cross-architecture evaluation, for example, on ARM-based mobile or embedded platforms, would be necessary to validate the practical deployment behavior of these schemes in the scenarios motivated in the introduction. We therefore consider cross-platform benchmarking an important direction for future work.

4 Security Analysis of MQ-Based Signature Candidates

This section analyzes the security of the four MQ-based signature candidates. We begin by outlining the main attack strategies that have been studied against MQ signatures, including both forgery attacks and key-recovery attacks. Then we compare the computational complexities of the schemes under these attacks, and consider recent advances in cryptanalysis and parameter updates.

4.1 Attack Strategies against MQ-Based Signatures

Table 6 summarizes the main known attacks considered against MQ-based signature schemes. These attacks can be broadly divided into two categories: forgery attacks, such as direct, collision, and underdetermined-MQ forgery attacks, and key-recovery attacks, including Kipnis–Shamir (KS), reconciliation, intersection, rectangular MinRank, wedge-product, symmetric-algebra, and truncated-polynomial-ring attacks.

Table 6: Known attacks against MQ-based signature schemes.

Attacks	Definition	Type	Algorithms	Target
Direct	Directly solve MQ problem	Forgery	Wiedemann XL, Gröbner basis (F4/F5)	UOV, MAYO, SNOVA, QR-UOV
Collision	Find hash collisions	Forgery	Memoryless claw finding	UOV, MAYO, SNOVA, QR-UOV
Kipnis-Shamir	Recover oil space via eigenvector structure	Key recovery	Wiedemann XL	UOV, MAYO, SNOVA, QR-UOV
Reconciliation	Recover oil space by solving quadratic constraints $y^{\top} P_i y = 0$	Key recovery	Wiedemann XL	MAYO, SNOVA, QR-UOV
Intersection	Find a vector in intersection of transformed oil spaces	Key recovery	Wiedemann XL	UOV, MAYO, SNOVA, QR-UOV
Rectangular MinRank	Solve low-rank linear combination of public matrices	Key recovery	Support Minors, Wiedemann XL	UOV, MAYO, SNOVA, QR-UOV

(Continued)

Table 6 (continued)

Attacks	Definition	Type	Algorithms	Target
Wedge-product	Recover oil space using exterior-algebra relations of polar forms	Key recovery	Sparse linear algebra	UOV, MAYO, SNOVA
Symmetric-algebra	Generalize wedge-type attacks using symmetric algebra	Key recovery	XL/linear algebra	QR-UOV
p^{ℓ}-truncated polynomial-ring	Reformulate reconciliation and intersection attacks using truncated polynomial rings	Key recovery	Wiedemann XL	UOV, MAYO, SNOVA, QR-UOV
Just-Guess	Solve underdetermined MQ instances by shifting the workload to exhaustive search	Forgery	Exhaustive search	MAYO, SNOVA, QR-UOV

The direct attack is the baseline strategy that attempts to solve the MQ system without exploiting the trapdoor. Typical methods include variants of the XL algorithm [46] and Gröbner basis techniques such as Faugère's F4 and F5.

A collision attack, also known as a Claw Finding attack, is a generic method that applies to all candidates. However, it does not use the structure of the schemes and attempts to find two distinct inputs that produce the same output on the public quadratic map. A straightforward approach to this problem is to store a large number of values from both sets until a collision occurs, namely keeping $\mathcal{P}(s_{i^*})$ and $\mathrm{Hash}(\mu \parallel \mathrm{salt}_{j^*})$ in memory for many choices of i and j. This technique has a significant memory requirement and is therefore considered less practical than memory-efficient alternatives, such as the memoryless claw finding attack, as described in the Round 2 specification of SNOVA [20].

The KS attack was first proposed in 1998 by Kipnis and Shamir [15] against the Oil and Vinegar problem. The central idea is to exploit the fact that the vectors of the oil space $\mathcal{O}$ appear as eigenvectors of a certain matrix derived from the public key. The usual complexity estimate is $O\left(q^{v-o-1} \cdot o^4\right)$. Here, the term q^{v-o-1} corresponds to the expected number of trials required to obtain a single valid vector, while the factor o^4 accounts for the computational cost of each such trial (see [15] for details).

Reconciliation attack was first described by Ding et al. [47] in 2008. It aims to recover vectors from the oil space $\mathcal{O}$ by solving a system of quadratic equations, thereby reconstructing the subspace $T^{-1}(\mathcal{O})$ and breaking signature schemes. The attack is equivalent to finding a nonzero vector $y \in T^{-1}(\mathcal{O})$ such that substituting y into the public quadratic forms yields zero. More concretely, if the public key is represented by quadratic form matrices P_i for $i = 1, \ldots, m$, the system to be solved is

$$y^{\top} P_i y = 0, \qquad i = 1, \ldots, m, \tag{11}$$

where y is an unknown vector in $T^{-1}(\mathcal{O})$. Thus, the reconciliation attack reduces to solving an MQ problem of m quadratic equations in $v = n - o$ variables.

In 2021, Beullens proposed a new key recovery attack called intersection attack [21]. It is similar to the KS attack, and in combination with the reconciliation attack. Its goal is to recover the oil subspace by finding a vector that lies in two differently transformed images of the oil space. The problem can be solved using algebraic solvers such as the XL algorithm, and a solution gives two vectors in the oil space; then the full oil subspace can be reconstructed efficiently.

Finally, the Rectangular MinRank attack [21] represents a more advanced algebraic method. It was first proposed for the Rainbow scheme. It tries to solve the MinRank problem to recover the oil space. Although originally proposed for Rainbow, it also applies to certain UOV-type schemes. It was noted in [48] that the rectangular MinRank attack can be applied to UOV when the condition $m > v$ holds. However, this condition is not satisfied in the second-round parameter sets of UOV [17], and thus the attack does not apply. Regarding other UOV-based variants, ref. [22] showed that the attack can be used against the initial version of MAYO. However, in the second version of MAYO [18], the parameters were adjusted to ensure $n - o \geq m$, thereby preventing this attack. More recently, Suzuki et al. unified the rectangular MinRank attack with the Kipnis–Shamir and singular-point attacks within the Jacobian framework, and proposed an extended rectangular MinRank attack by allowing a wider range of feasible target ranks [23]. Their complexity estimates show that the current parameter sets of UOV, MAYO, QR-UOV, and SNOVA remain above the claimed security levels against this attack, but the attack provides an important systematic bound for evaluating UOV-based variants.

Recent cryptanalysis has introduced several additional attacks that should also be considered in the security evaluation of UOV-based signature schemes. Ran proposed a wedge-product attack against UOV-family schemes over characteristic-two fields [24]. The attack exploits the fact that the polar forms of quadratic maps over characteristic two are alternating forms, and represents them in the exterior algebra. By deriving linear relations on the Plücker embedding of the secret oil space, the attack can recover the hidden subspace using sparse linear algebra. This attack reduces the estimated security of several UOV parameter sets and has a notable impact on MAYO2.

The wedge-product approach has also been adapted to SNOVA. Bros et al. exploited SNOVA's block-ring structure within the wedge-product framework [32]. Since SNOVA introduces additional algebraic structure to reduce the public-key size, its structured block representation creates extra linear dependencies that can be used in the attack. Therefore, the SNOVA block-ring wedge attack can be viewed as a SNOVA-specific instance of the broader wedge-product attack family, and it reduces the estimated security of several proposed SNOVA parameter sets.

Jin et al. later generalized this line of analysis to fields of odd characteristic using reduced symmetric algebra [25]. This extension is relevant to QR-UOV because QR-UOV is instantiated over odd-characteristic fields. However, for the proposed QR-UOV parameter sets, the resulting attack does not outperform the existing key-recovery attacks. Therefore, it mainly serves as an important confirmation that odd-characteristic UOV variants should also be included in recent algebraic security analyses.

Furue and Ikematsu further reformulated the exterior-algebra and symmetric-algebra approaches using p^{ℓ}-truncated polynomial rings [26]. This framework provides a unified interpretation of these algebraic key-recovery attacks and can be applied to reconciliation and intersection attacks. In particular, the resulting intersection attack reduces the estimated security of several UOV parameter sets, while the reconciliation attack under this framework also affects multiple SNOVA parameter sets.

In addition to key-recovery attacks, recent work has improved the analysis of forgery attacks against underdetermined MQ systems. May et al. proposed the Just-Guess algorithm, which transforms underdetermined MQ solving into a form dominated by exhaustive search and provides updated classical and quantum forgery estimates for MAYO, QR-UOV, and SNOVA [35]. Since the main workload becomes more search-oriented, the quantum version benefits more directly from Grover's quadratic speedup.

It is worth noting that algebraic solving techniques such as the XL algorithm and Gröbner basis methods are universal tools for MQ systems. They appear not only in direct attacks but also as building blocks in KS, reconciliation, intersection, and MinRank, underlining their central role in evaluating the hardness of MQ problems.

In the quantum case, known techniques such as Grover's algorithm can provide quadratic speedups for search-dominated subroutines. However, such speedups should not be applied uniformly to all classical attacks, because different attacks contain different algebraic and search components. For example, the Just-Guess attack provides concrete quantum forgery estimates for underdetermined MQ instances arising from MAYO, QR-UOV, and SNOVA [35]. Therefore, quantum-adjusted bounds are treated as separate post-quantum considerations in this work, rather than being obtained by simply halving all classical attack exponents or being merged into the classical effective-security metric.

Overall, these attacks highlight the main cryptanalytic strategies against MQ-based signatures. Direct and collision attacks do not exploit the trapdoor structure, yet they remain important baselines and often determine the effective security level for several parameter sets. Key-recovery methods such as Kipnis–Shamir, reconciliation, intersection, rectangular MinRank, wedge-product, symmetric-algebra, and truncated-polynomial-ring attacks exploit algebraic relations between oil and vinegar variables and strongly influence parameter design. Recent SNOVA-specific adaptations of wedge-product attacks and underdetermined-MQ attacks further show that additional structure introduced for key-size reduction may change the dominant security estimates. Therefore, the effective security metric used in the following comparison should be interpreted with respect to both the specification-based attacks and recent cryptanalytic updates.

4.2 Comparison of Complexity

To enable consistent evaluation, attack complexities in this section are presented using unified criteria across all schemes. All values are normalized to $\log_2$ scale and compared with respect to NIST security levels I, III, and V. The complexity estimates reported in the official Round-2 specification documents are used as baseline values [17–20]. To reflect recent cryptanalytic progress, we additionally include updated estimates from several recent attack families. MinRank-type estimates are updated using the extended rectangular MinRank attack [23]. Wedge-product and symmetric-algebra estimates are considered for UOV-family schemes, SNOVA, and odd-characteristic variants [24,25,32]. We also include estimates based on the p^{ℓ}-truncated polynomial-ring framework [26] and the Just-Guess forgery attack [35].

The first version of the second-round candidates in the NIST Additional Digital Signatures process was selected in late 2024. Shortly after the announcement, several cryptanalytic studies were published targeting these candidates, re-evaluating their security and identifying weaknesses in the original parameter sets.

Furue and Ikematsu [22] introduced a new cryptanalysis against UOV-based variants MAYO and QR-UOV. Their work showed that the rectangular MinRank attack, originally proposed for Rainbow, can also be applied to these schemes. The analysis estimated that the proposed parameters of MAYO and QR-UOV remain secure against this attack, though the complexity is close to that of the best-known algebraic methods. More recently, Suzuki et al. proposed an extended rectangular MinRank attack against UOV and

its variants [23]. Therefore, in this work, the MinRank column in Table 7 includes both rectangular MinRank and its extended variants where applicable.

Table 7: Complexity ($\log_2$ scale) of known attacks against MQ-based signature schemes.

SL	Scheme	Col.	Dir.	KS	Inter.	Recon.	MR	Wedge	Sym.	Trunc.	JG
I	UOV-Ip	191	145	218	166	–	275	141	140	128	–
	UOV-Is	143	165	154	176	–	374	176	175	159	–
	MAYO1	169	156	303	355	197	276	–	–	238	158
	MAYO2	141	155	211	228	167	210	113	–	113	–
	SNOVA-a	158	165	165	153	195	229	103	–	103	–
	SNOVA-b	166	171	209	221	187	220	110	–	124	–
	SNOVA-c	182	184	309	353	249	276	155	–	162	218
	QR-UOV	201	150	718	460	164	158	182	164	165	–
III	UOV	303	218	348	250	–	436	207	206	182	–
	MAYO	229	224	416	482	262	360	–	–	314	240
	SNOVA-a	222	234	253	221	288	337	140	–	140	–
	SNOVA-b	220	226	461	529	424	476	246	–	251	–
	SNOVA-c	279	287	469	506	365	395	214	–	225	–
	SNOVA-d	273	281	385	412	257	294	155	–	173	368
	QR-UOV	286	211	1056	653	231	219	249	231	231	–
V	UOV	399	278	445	312	–	567	258	257	223	–
	MAYO	298	296	546	629	334	459	–	–	397	333
	SNOVA-a	287	302	341	288	379	446	180	–	180	–
	SNOVA-b	293	302	617	690	558	622	313	–	325	–
	SNOVA-c	343	350	805	922	673	776	–	–	695	469
	SNOVA-d	323	334	465	494	310	346	181	–	201	–
	QR-UOV	380	279	1414	851	291	277	328	291	291	–

Note: Blue values indicate the lowest attack complexity for each scheme, corresponding to the most relevant known classical attack estimate among the listed attacks.

Beullens [27], Li and Ding [30] presented cryptanalysis of the SNOVA. Their studies demonstrated that several parameter sets originally proposed for SNOVA do not meet the intended NIST security levels. By exploiting structural properties of the scheme, they derived attacks with complexities lower than those claimed in the specification, showing that the original estimates were overly optimistic. These works highlighted the insufficiency of the first-round parameters and motivated the adjustments incorporated in the SNOVA Round 2 submission. Recent wedge-product attacks further exploit the block-ring structure of SNOVA and provide lower estimates for several Round-2 parameter sets [32].

In response to these findings, the authors of the four candidates provided updated specification packages in early 2025. The revised versions included parameter adjustments designed to restore the intended security margins. For example, the second version of MAYO modified its parameters to ensure that the relation $n - o \geq m$ holds, thereby preventing rectangular MinRank from applying effectively. SNOVA updated its parameter sets to resist defects caused by the structure.

Nevertheless, the security analysis and design of SNOVA have continued to evolve after the Round-2 submission. Following the wedge-product analysis of SNOVA, Ding et al. reformulated SNOVA using

ring-equation, whipping, and tensor representations, and proposed more flexible parameter choices with competitive key and signature sizes [33]. However, these parameter sets are post-Round-2 demonstrative parameters rather than official SNOVA parameters used in the main complexity comparison. Therefore, they are not included in the main comparison. Instead, they indicate that SNOVA's position in the objective space may change if future standardized versions adopt the reformulated framework.

The final security estimates and recent cryptanalytic updates are summarized in Table 7. The columns "Col.", "Dir.", "KS", "Inter.", and "Recon." correspond to specification-based attacks, where "MR" includes both rectangular MinRank and extended rectangular MinRank estimates when applicable. The columns "Wedge", "Sym.", "Trunc.", and "JG" denote the wedge-product attack, symmetric-algebra attack, p^{ℓ}-truncated polynomial-ring attack, and Just-Guess forgery estimate, respectively. The "Wedge" column includes both the characteristic-two wedge-product attack against UOV-family schemes and the SNOVA block-ring wedge attack. We use the blue text to mark the lowest complexity of each scheme, which represents the most relevant threat among all listed classical estimates.

At Security Level I, the dominant attack varies across schemes. For UOV-Ip, the lowest listed estimate is given by the p^{ℓ}-truncated polynomial-ring attack, whereas UOV-Is is mainly limited by the collision estimate. MAYO1 is limited by the direct attack, while MAYO2 has the lowest estimates from the wedge-product and p^{ℓ}-truncated polynomial-ring attacks. For SNOVA, wedge-product estimates dominate most parameter sets, with ties against truncated-polynomial-ring estimates in some cases. QR-UOV is mainly limited by the direct attack among the listed classical estimates.

At Security Level III, UOV obtains its lowest listed estimate from the p^{ℓ}-truncated polynomial-ring attack, while MAYO is still dominated by the direct attack. For SNOVA, SNOVA-a has the lowest estimates from the wedge-product and p^{ℓ}-truncated polynomial-ring attack, SNOVA-c and SNOVA-d are mainly limited by wedge-product estimates, and SNOVA-b is limited by the collision estimate. QR-UOV continues to be mainly limited by the direct attack among all attacks.

At Security Level V, UOV and MAYO follow the same dominant-attack pattern as in Level III: UOV is mainly limited by the p^{ℓ}-truncated polynomial-ring attack, whereas MAYO is limited by the direct attack. In contrast, QR-UOV obtains its lowest listed estimate from a rectangular MinRank attack, indicating that MinRank-type estimates become more relevant at this level. For SNOVA, SNOVA-a again has tied the lowest estimates from the wedge-product and p^{ℓ}-truncated polynomial-ring attack, SNOVA-d is mainly limited by the wedge-product estimate, and SNOVA-b and SNOVA-c are mainly limited by collision estimates.

Overall, the comparison shows that the dominant attack mechanism depends strongly on both the scheme design and the parameter set. For UOV, recent algebraic key-recovery estimates based on truncated polynomial rings become relevant across several security levels. For MAYO, the direct attack remains the main limiting factor for most parameter sets, although MAYO2 is also affected by wedge-type and truncated-polynomial-ring algebraic estimates. For QR-UOV, the direct attack gives the lowest listed estimate at Security Levels I and III, whereas a MinRank-derived bound becomes the lowest listed estimate at Security Level V. SNOVA exhibits the greatest sensitivity to structure-specific analysis, since wedge-product attacks exploiting its block-ring representation give the lowest or tied-lowest estimates for several parameter sets.

Therefore, the effective security of each parameter set is represented by the minimum classical attack complexity among the listed estimates. This scalar value provides a unified security objective for the multi-objective evaluation framework developed in Section 5.

Because this effective-security value is used as a single scalar objective, Table 7 reports classical attack complexities only. This avoids combining classical and quantum cost models into a single value.

Quantum estimates, such as the quantum variant of the Just-Guess attack, are therefore treated as separate post-quantum considerations rather than being merged into the classical effective-security metric.

5 Pareto-Based Multi-Objective Evaluation Framework

In this section, we introduce a Pareto-based multi-objective modeling framework to systematically analyze the trade-offs among MQ-based signature schemes. Based on the performance metrics obtained in Section 3 and the security analysis in Section 4, we represent each scheme in a multi-dimensional objective space and perform a Pareto-based comparative analysis to identify non-dominated schemes.

The purpose of this Pareto framework is not to make a standardization decision or to assign a single overall score. Instead, it makes the trade-off structure explicit by identifying which schemes cannot be dominated under the selected metrics. This is useful because PQC signature candidates often optimize different objectives, and a weighted score would require application-dependent preferences. By avoiding fixed weights, the proposed analysis exposes whether a scheme is dominated, which objectives prevent dominance, and how scheme selection changes under different deployment priorities. In this sense, the framework provides an additional layer of interpretation on top of existing security and performance evaluations.

We emphasize that Pareto dominance itself is a classical concept in multi-objective optimization. The contribution of this work is not to propose a new Pareto algorithm, but to adapt this concept into a reproducible and application-oriented evaluation framework for MQ-based signature candidates by jointly incorporating compactness, measured computational cost, and updated effective-security estimates.

5.1 Evaluation Metrics

In this work, we consider six evaluation metrics: public key size, signature size, key generation cost, signing cost, verification cost, and effective security. These metrics are selected to provide a comprehensive characterization of MQ-based signature schemes from both structural and computational perspectives.

Public key size and signature size are fundamental indicators of compactness, which is particularly important in resource-constrained environments such as mobile and embedded systems. MQ-based schemes are often motivated by their ability to produce short signatures, making these metrics central to practical deployment.

The computational costs of key generation, signing, and verification capture the runtime performance of each scheme. These three operations correspond to different usage scenarios: key generation reflects setup cost, while signing and verification determine the efficiency of online operations. By evaluating all three stages, we obtain a complete view of the computational behavior of each scheme.

The effective security is defined as the complexity of the most efficient known classical attack, representing the effective classical security margin. This definition follows standard cryptographic practice, where the strength of a scheme is evaluated against its best-known attack. In this work, the attack set used to compute effective security is not limited to the original specification-based estimates, but also includes the recent classical cryptanalytic updates summarized in Table 7.

We acknowledge that this effective security metric is a modeling choice. Different MQ-based schemes may be constrained by different types of attacks, such as direct forgery, collision attacks, MinRank-type attacks, wedge-product attacks, or structure-specific algebraic attacks. Reducing these heterogeneous attacks to a single scalar value inevitably loses some information about the attack surface. Nevertheless, using the minimum complexity among the best-known applicable attacks provides a conservative and comparable security indicator: it reflects the most serious currently known threat against each parameter set. To preserve

transparency, the individual attack estimates and the corresponding dominant attacks are reported separately in Table 7, while the scalar effective security value is used only as an input to the Pareto model.

Together, these six metrics provide a balanced and comprehensive basis for evaluating MQ-based signature schemes. They enable a systematic analysis of trade-offs among compactness, efficiency, and security without relying on subjective weighting.

5.2 Evaluation Model and Normalization

Let $\mathcal{S}$ denote the set of candidate schemes under a fixed security level. Each scheme $S \in \mathcal{S}$ is represented by a six-dimensional objective vector:

$$\mathbf{f}(S) = \left(f_1, f_2, f_3, f_4, f_5, f_6\right), \tag{12}$$

where the objectives correspond to public key size, signature size, key generation cost, signing cost, verification cost, and effective security.

To enable a fair comparison across heterogeneous evaluation criteria with different units and scales, all metrics are normalized using min–max scaling. Given a metric value x, its normalized value x' is computed based on the minimum and maximum values of that metric across all schemes within the same security level.

For metrics where smaller values indicate better performance (e.g., public key size, signature size, and computational cost), a reversed min–max normalization is applied:

$$x' = \frac{x_{\max} - x}{x_{\max} - x_{\min}}, \tag{13}$$

so that larger normalized values correspond to better performance.

For metrics where larger values indicate better performance (e.g., effective security), standard min–max normalization is used:

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}}. \tag{14}$$

After normalization, all metric values are mapped to the range $[0, 1]$, where a higher value consistently represents better performance. In addition, normalization is performed independently within each security level to avoid distortions caused by large differences in scale across levels. This ensures that comparisons remain meaningful within each security category.

5.3 Pareto Analysis across Security Levels

Pareto dominance is used to compare schemes in the normalized multi-objective space. A scheme is said to be Pareto-optimal if no other scheme performs at least as well in all metrics and strictly better in at least one metric.

The resulting Pareto front should therefore be interpreted with respect to the specific metrics and security estimates selected in this work. It is not an absolute ranking of the schemes, but a transparent representation of trade-offs under the stated modeling assumptions.

We apply this analysis to all schemes under each security level. To illustrate these analyses, Fig. 1 presents a heatmap of the normalized metrics. Each row corresponds to a scheme and each column corresponds to a metric. If a scheme dominated another, its corresponding row would exhibit values that are consistently higher across all columns. However, such a pattern is not observed. Instead, each scheme exhibits a mixture of strong and weak metrics, visually explaining the absence of dominance relationships.

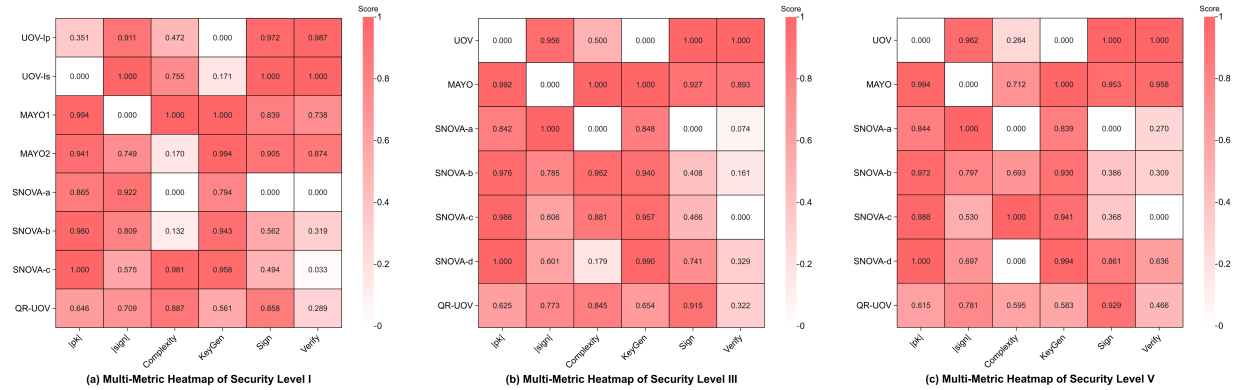

Figure 1: Visualization of trade-offs in MQ-based signature schemes.

The results show that, for all three security levels, no scheme is strictly dominated by another. In other words, the Pareto front coincides with the entire set of candidates. This outcome follows directly from the structure of the evaluation metrics. For each scheme, there exists at least one metric in which it outperforms the others, while exhibiting weaker performance in other dimensions. As a result, no scheme achieves simultaneous superiority across all objectives.

The heatmap further illustrates how the normalized metrics vary across security levels. In the effective-security dimension, MAYO1 and SNOVA-c obtain relatively high normalized scores at Security Level I, while SNOVA-a and SNOVA-b show lower scores, reflecting reduced effective-security estimates under recent structure-specific attacks. At Security Level III, MAYO achieves the highest normalized effective-security score, whereas SNOVA-a and SNOVA-d show lower scores. At Security Level V, SNOVA-c has the highest normalized effective-security score, while SNOVA-a and SNOVA-d remain weaker in this dimension. These observations indicate that compactness-oriented designs, especially SNOVA parameter sets, do not uniformly translate into stronger effective-security positions across all security levels.

This non-dominance pattern is observed consistently across all security levels, suggesting that the trade-offs are partly associated with structural design differences among the schemes. However, the result should not be over-interpreted, since dominance relationships become increasingly rare in high-dimensional multi-objective evaluation. Thus, the absence of dominance should be understood primarily as evidence of diverse metric specialization under the selected objectives.

This metric specialization appears differently across the four schemes. UOV occupies an extreme point in the objective space: it provides very strong local signing and verification performance, but pays for this advantage with very large public keys and high key-generation cost. MAYO represents a different trade-off by substantially reducing public-key size and key-generation cost, while accepting higher signing overhead and larger signatures. SNOVA is located near the compactness-oriented region of the design space, but its position is weakened by higher computational cost and by reduced effective security scores for several parameter sets under recent structure-specific attacks. QR-UOV, in contrast, does not optimize any single objective but remains close to the middle of the objective space, which explains why it behaves as a balanced but non-dominating candidate.

This observation also highlights a limitation of interpreting the Pareto front as a direct ranking. Since the present model contains six objectives, even a scheme with weak performance in several dimensions can remain non-dominated if it is sufficiently strong in at least one dimension. Therefore, the fact that all schemes lie on the Pareto front should not be interpreted as meaning that they are equally preferable in practice. Rather, it indicates that no scheme can be eliminated without first specifying deployment priorities.

From an application perspective, scheme selection should therefore be guided by system requirements rather than by a universal ranking. UOV is well suited to latency-sensitive applications where fast signing and verification dominate, and public keys can be cached or distributed offline. MAYO is advantageous when key-generation efficiency and compact public keys are important, such as in large-scale key deployment or scenarios where public-key distribution dominates communication cost. SNOVA is attractive for compact public-key and signature representations, but its deployment should account for the reduced and less stable security margins indicated by recent structure-specific cryptanalysis. QR-UOV offers a compromise solution for systems requiring balanced behavior without extreme optimization in any single metric.

The practical implications of the Pareto results can be further understood by considering representative deployment scenarios. In IoT and resource-constrained devices, public-key storage, communication bandwidth, and memory footprint are often more restrictive than raw CPU performance. In such settings, schemes with compact public keys, such as MAYO or compact SNOVA parameter sets, may be more attractive than UOV when public-key storage or distribution is the dominant bottleneck, even if their signing or verification costs are higher. However, the preferred scheme may change if the main bottleneck is signature transmission, verification latency, memory usage during computation, or energy consumption. For mobile authentication and secure messaging, both online latency and communication overhead are important. A scheme with fast signing and verification is advantageous for interactive authentication, but large public keys may increase transmission cost during certificate exchange or initial key registration. Therefore, UOV may be suitable when public keys are cached or distributed offline, whereas MAYO, SNOVA, or QR-UOV may be preferable when public keys are transmitted frequently. In cloud security or server-side authentication, storage and bandwidth constraints may be less severe, while verification throughput, key-management cost, and large-scale deployment efficiency become more important. In this case, UOV's fast verification can be beneficial for high-throughput verification workloads, whereas MAYO's efficient key generation and compact public keys may be useful for systems requiring frequent key provisioning. For applications requiring balanced behavior across several constraints, QR-UOV provides a compromise profile without extreme weakness in a single deployment dimension.

This application-level interpretation should also consider network transmission overhead. The six metrics used in the present Pareto model capture local computational cost, compactness, and effective security, but they do not explicitly model protocol-level communication cost. In real-world deployment scenarios such as TLS handshakes, IPsec, secure messaging, and device authentication, public keys and signatures may need to be transmitted over a network, and the resulting communication cost depends on bandwidth, round-trip latency, fragmentation, certificate-chain size, and caching assumptions. Therefore, public key size and signature size should not be interpreted merely as storage metrics, but also as factors that may influence deployment-level communication efficiency.

This issue is particularly relevant for UOV, whose public key size reaches approximately 447 KB at Security Level V, as shown in Table 5. Although UOV achieves the best signing and verification performance in local cycle-count measurements, transmitting such a large public key may offset this computational advantage in bandwidth-constrained environments. Conversely, compact-key schemes such as MAYO and SNOVA may be more attractive in communication-constrained deployments, even when their local computational costs are higher. In high-bandwidth environments where public keys are cached or transmitted infrequently, the impact of public-key size may be less pronounced, and local computation may become the dominant factor. Thus, practical efficiency should be regarded as deployment-dependent and should account for both local computation and protocol-level communication cost.

We note that incorporating network transmission cost as an additional objective in the Pareto framework would require deployment-specific assumptions, such as available bandwidth, round-trip latency,

certificate-chain overhead, fragmentation behavior, and key-caching policy. These assumptions may vary substantially across use cases and are therefore not included as a formal objective in the present model. Instead, we treat network overhead as a complementary qualitative consideration and identify a network-aware extension of the framework as an important direction for future work.

We emphasize that the Pareto analysis reflects the objective values measured from the available official implementations under the experimental setup described in Section 3.2. Since these implementations do not provide uniform, implementation-wide CT guarantees, CT hardening or masking may change the measured computational cost across key generation, signing, and verification. This limitation does not invalidate the Pareto-based framework itself, because the framework can incorporate any consistently measured objective values. However, it may affect the exact computational coordinates and, in some cases, the resulting dominance relationships. Therefore, the current results should not be interpreted as a comparison of uniformly CT-hardened or uniformly side-channel-protected implementations.

In addition, the performance comparison preserves the official reference build configurations of the tested implementations. Since architecture-specific compiler flags such as `-march=native` are not applied uniformly across all schemes, compiler-level optimization may affect the exact computational coordinates in the Pareto space. Therefore, the Pareto results should be interpreted as reference-implementation-based measurements under the stated build environment, rather than as a fully compiler-flag-normalized comparison. A fully normalized performance comparison would require rebuilding all implementations under a consistent set of architecture-specific compiler options.

The Pareto analysis in this work is based on the official Round-2 parameter sets to ensure a consistent comparison across UOV, MAYO, QR-UOV, and SNOVA. Recently reformulated SNOVA parameters may improve compactness and change SNOVA's position in the objective space, but they are not included in the main Pareto computation because they are post-Round-2 demonstrative parameters rather than official Round-2 submissions [33].

Overall, the adapted framework provides a systematic and transparent approach for evaluating MQ-based signature schemes. Even in the absence of strict dominance, it remains meaningful by explicitly revealing the trade-offs among compactness, efficiency, and security, thereby offering deeper insight into the design space.

6 Conclusion

In this paper, we presented a systematic comparative study of four MQ-based signature schemes, namely UOV, MAYO, SNOVA, and QR-UOV, within the context of the NIST additional digital signature standardization process. By integrating structural analysis, performance evaluation under a unified platform, and security assessment based on the best-known attacks, we provided a comprehensive view of their design characteristics and practical behavior.

To move beyond conventional metric-wise comparison, we introduced a Pareto-based multi-objective evaluation framework that jointly considers six key metrics, including public key size, signature size, computational cost, and effective security. This framework enables the analysis of trade-offs without relying on subjective weighting and provides a unified way to compare schemes across multiple conflicting criteria.

The security metric used in this framework is based on the minimum classical attack complexity among the updated estimates considered in this work, including both specification-based attacks and recent classical cryptanalytic results. This allows the Pareto analysis to reflect recent developments in MinRank-type attacks, wedge-product attacks, symmetric-algebra attacks, p^{ℓ}-truncated polynomial-ring attacks, and Just-Guess forgery estimates, while avoiding the direct combination of classical and quantum cost models.

The experimental results show that no scheme is strictly dominated by another under the six-dimensional evaluation model. That is, all schemes lie on the Pareto front, indicating that each design achieves a distinct balance among compactness, efficiency, and security. This finding highlights that MQ-based signature schemes are inherently characterized by structural trade-offs and that no single scheme can be considered universally optimal.

From a practical perspective, the results suggest that scheme selection should be guided by application requirements. Schemes such as UOV are well-suited for latency-sensitive scenarios due to their efficient signing and verification. MAYO is attractive in settings where compact public keys and efficient key generation are important. SNOVA remains attractive for compact representations, but its deployment should be considered together with its less stable security margin under recent structure-specific attacks. QR-UOV provides a balanced alternative when moderate performance across multiple dimensions is desired. These observations indicate that the adapted framework can support scenario-dependent selection, including IoT and embedded devices where storage and bandwidth are constrained, mobile authentication where latency and communication cost must be jointly considered, and cloud-side verification where throughput and key-management efficiency may dominate.

Overall, the adapted framework complements traditional comparison approaches by explicitly revealing the absence of a globally optimal solution and providing deeper insight into the design space of MQ-based signature schemes. The results also show that compactness-oriented designs should be evaluated together with their evolving cryptanalytic landscape, since additional algebraic structure introduced for compression may affect concrete security estimates.

Recent reformulation work on SNOVA further suggests that compact UOV-based designs may continue to evolve after new cryptanalytic results. Since the reformulated SNOVA parameters are post-Round-2 demonstrative parameters rather than official Round-2 submissions, they were not included in the main Pareto comparison. Evaluating such post-Round-2 parameter updates in a consistent multi-objective framework remains an important direction for future work [33].

Several limitations of the current framework should also be noted. First, the Pareto results are conditional on the selected metrics and modeling assumptions. Although the framework avoids assigning subjective numerical weights, the choice of objectives still affects the resulting comparison. In particular, the effective-security metric summarizes different attack types by the minimum known classical attack complexity. This provides a conservative and comparable indicator, but it cannot fully describe the different attack surfaces of the schemes. Second, the current model focuses on compactness, local computational cost, and mathematical attack complexity. It does not formally include deployment-dependent factors such as network latency, certificate-chain overhead, key-caching policy, memory footprint, energy consumption, or implementation-level resistance to side-channel and fault-injection attacks. Third, the performance measurements were obtained on a single x86-based platform using available official implementations. Since mobile, embedded, and IoT devices may have different cache hierarchies, memory bandwidth, and energy constraints, the measured cycle counts should be interpreted as platform-specific results rather than architecture-independent performance conclusions. Moreover, the tested implementations do not provide uniform implementation-wide constant-time guarantees, and architecture-specific compiler flags such as `-march=native` were not uniformly applied across all reference build configurations. Therefore, the current results should not be viewed as a fully constant-time-normalized, uniformly side-channel-protected, or fully compiler-flag-normalized comparison.

Future work can improve the framework in several directions. First, cross-architecture benchmarking on ARM-based mobile, embedded, and IoT platforms would help determine whether the observed Pareto positions remain stable under constrained memory and cache conditions. Second, deployment-oriented

metrics such as memory footprint, energy consumption, network bandwidth consumption, communication latency, and protocol-specific transmission cost should be incorporated when concrete application assumptions are available. Third, implementation-security metrics, including constant-time verification, masking overhead, leakage-test results, and fault-detection capability, could be introduced as additional Pareto dimensions to capture side-channel and fault-injection resilience. Finally, as MQ-based candidates continue to evolve in response to new cryptanalytic results, future evaluations should update the effective-security estimates and include post-Round-2 parameter revisions once they become official or sufficiently mature for fair comparison.

Acknowledgement: Not applicable.

Funding Statement: This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2019-II190033, Study on Quantum Security Evaluation of Cryptography based on Computational Quantum Complexity).

Author Contributions: The authors confirm contribution to the paper as follows: Conceptualization, Jian Zhang and Seong-Min Cho; methodology, Jian Zhang; validation, Jian Zhang and Seong-Min Cho; formal analysis, Jian Zhang; investigation, Jian Zhang; data curation, Jian Zhang; writing—original draft preparation, Jian Zhang; writing—review and editing, Seong-Min Cho and Seung-Hyun Seo; visualization, Jian Zhang; supervision, Seung-Hyun Seo; project administration, Seung-Hyun Seo; funding acquisition, Seung-Hyun Seo. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: All data supporting the results of this study are included within the article.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Rivest RL, Shamir A, Adleman L. A method for obtaining digital signatures and public-key cryptosystems. Commun ACM. 1978;21(2):120–6. doi:10.21236/ada606588.
2. Johnson D, Menezes A, Vanstone S. The elliptic curve digital signature algorithm (ECDSA). Int J Inf Secur. 2001;1(1):36–63. doi:10.1007/s102070100002.
3. Shor PW. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM Rev. 1999;41(2):303–32. doi:10.1137/s0036144598347011.
4. National Institute of Standards and Technology. Post-quantum cryptography [Internet]. [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization.
5. National Institute of Standards and Technology. PQC standardization process: announcing four candidates to be standardized, plus fourth round candidates [Internet]. 2022 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/News/2022/pqc-candidates-to-be-standardized-and-round-4.
6. National Institute of Standards and Technology. FIPS 203: module-lattice-based key-encapsulation mechanism standard [Internet]. 2024 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/pubs/fips/203/final.
7. National Institute of Standards and Technology. FIPS 204: module-lattice-based digital signature standard [Internet]. 2024 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/pubs/fips/204/final.
8. National Institute of Standards and Technology. FIPS 205: stateless hash-based digital signature standard [Internet]. 2024 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/pubs/fips/205/final.
9. National Institute of Standards and Technology. Status report on the fourth round of the NIST post-quantum cryptography standardization process [Internet]. 2025 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/pubs/ir/8545/final.

10. National Institute of Standards and Technology. Post-quantum cryptography: additional digital signature schemes [Internet]. [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/projects/pqc-dig-sig.
11. National Institute of Standards and Technology. NIST announces 14 candidates to advance to the second round of the additional digital signatures for the post-quantum cryptography standardization process [Internet]. 2024 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/news/2024/pqc-digital-signature-second-round-announcement.
12. Alagic G, Apon D, Cooper D, Dang Q, Liu YK, Miller C, et al. Status report on the second round of the additional digital signature schemes for the NIST post-quantum cryptography standardization process. Gaithersburg, MD, USA: National Institute of Standards and Technology; 2026. Report No.: NIST IR 8610.
13. Ding J, Petzoldt A, Schmidt DS. Multivariate public key cryptosystems. 2nd ed. Cham, Switzerland: Springer; 2020.
14. Patarin J. The oil and vinegar signature scheme. In: The Dagstuhl Workshop on Cryptography; 1997 Sep 22–26; Dagstuhl, Germany.
15. Kipnis A, Shamir A. Cryptanalysis of the oil and vinegar signature scheme. In: Advances in Cryptology–CRYPTO'98. Lecture Notes in Computer Science. Vol. 1462. Berlin, Germany: Springer; 1998. p. 257–66.
16. Kipnis A, Patarin J, Goubin L. Unbalanced oil and vinegar signature schemes. In: Advances in Cryptology–EUROCRYPT'99. Lecture Notes in Computer Science. Vol. 1592. Berlin, Germany: Springer; 1999. p. 206–22.
17. Beullens W, Chen MS, Ding J, Gong B, Kannwischer MJ, Patarin J, et al. UOV: unbalanced oil and vinegar [Internet]. 2025 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/uov-spec-round2-web.pdf.
18. Beullens W, Campos F, Celi S, Hess B, Kannwischer MJ. MAYO [Internet]. 2025 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/mayo-spec-round2-web.pdf.
19. Furue H, Ikematsu Y, Hoshino F, Takagi T, Kosuge H, Yamakoshi K, et al. QR-UOV [Internet]. 2025 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/qruov-spec-round2-web.pdf.
20. Wang LC, Chou CY, Ding J, Kuan YL, Leegwater JA, Li MS, et al. SNOVA [Internet]. 2025 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-2/spec-files/snova-spec-round2-web.pdf.
21. Beullens W. Improved cryptanalysis of UOV and rainbow. In: Advances in Cryptology–EUROCRYPT 2021; 2021 Oct 17–21; Zagreb, Croatia.
22. Furue H, Ikematsu Y. A new cryptanalysis against UOV-based variants MAYO, QR-UOV and VOX. IEICE Trans Fundam Electron Commun Comput Sci. 2025;108(3):174–82. doi:10.1587/transfun.2024cip0003.
23. Suzuki T, Furue H, Ito T, Nakamura S, Uchiyama S. An extended rectangular MinRank attack against UOV and its variants. In: Proceedings of the International Workshop on Security. Cham, Switzerland: Springer; 2025. p. 111–30.
24. Ran L. Wedges, oil, and vinegar: an analysis of UOV in the exterior algebra. In: Advances in Cryptology–EUROCRYPT 2026. Lecture Notes in Computer Science. Vol. 16544. Cham, Switzerland: Springer; 2026. p. 363–88.
25. Jin Y, Pan Y, He X, Gong B, Ding J. Security analysis on UOV families with odd characteristics: using symmetric algebra. In: Public-Key Cryptography–PKC 2026. Lecture Notes in Computer Science. Vol. 16551. Cham, Switzerland: Springer; 2026. p. 428–54.
26. Furue H, Ikematsu Y. Key recovery attacks on UOV using p^{ℓ}-truncated polynomial rings. Bellevue, WA, USA: IACR; 2026. Report No.: 2026/298.
27. Beullens W. Improved cryptanalysis of SNOVA. In: Advances in Cryptology–EUROCRYPT 2025; 2025 May 4–8; Madrid, Spain.
28. Cabarcas D, Li P, Verbel J, Villanueva-Polanco R. Improved attacks for SNOVA by exploiting stability under a group action. In: Advances in Cryptology–CRYPTO 2025; 2025 Aug 17–21; Santa Barbara, CA, USA.
29. Ikematsu Y, Akiyama R. Revisiting the security analysis of SNOVA. In: Proceedings of the 11th ACM Asia Public-Key Cryptography Workshop, APKC '24; 2024 Jul 1–5; Singapore, Singapore. New York, NY, USA: ACM; 2024. p. 54–61.

30. Li P, Ding J. Cryptanalysis of the SNOVA signature scheme. In: Post-Quantum Cryptography; 2024 Jun 12–14; Oxford, UK.
31. Nakamura S, Tani Y, Furue H. Lifting approach against the SNOVA scheme. IEICE Trans Fundam Electron Commun Comput Sci. 2025;E108-A(10):1373–81. doi:10.1587/transfun.2024eap1124.
32. Bros M, Le TH, Lichtinger J, Minaud B, Perlner R, Smith-Tone D, et al. Exploiting SNOVA's structure in the wedge product attack. Bellevue, WA, USA: IACR; 2026. Report No.: 2026/237.
33. Ding J, Guo H, Kuan Y-L, Leegwater JA, Li P, Tseng P-E, et al. Reformulating the SNOVA signature scheme. Bellevue, WA, USA: IACR; 2026. Report No.: 2026/659.
34. Cho SM, Seo SH. Quantum rectangular MinRank attack on multi-layer UOV signature schemes. Sci Rep. 2024;14(1):16340. doi:10.1038/s41598-024-66841-0.
35. May A, Ostuzzi M, Ressler H. Just guess: improved (quantum) algorithm for the underdetermined MQ problem. In: Advances in Cryptology–EUROCRYPT 2026. Lecture Notes in Computer Science. Vol. 16544. Cham, Switzerland: Springer; 2026. p. 334–62.
36. Banegas G, Villanueva-Polanco R. A fault analysis on SNOVA. arXiv:2509.12879. 2025.
37. Jendral S, Dubrova E. MAYO key recovery by fixing vinegar seeds. Bellevue, WA, USA: IACR; 2024. Report No.: 2024/1550.
38. Norga Q, Kundu S, Ojha UK, Ganguly A, Karmakar A, Verbauwhede I. Masking Gaussian elimination at arbitrary order with application to multivariate- and code-based PQC. In: Topics in Cryptology–CT-RSA 2025. Lecture Notes in Computer Science. Vol. 15598. Cham, Switzerland: Springer; 2025. p. 249–72.
39. Aulbach T, Campos F, Kämer J. SoK: on the physical security of UOV-based signature schemes. In: Post-Quantum Cryptography–PQCrypto 2025. Lecture Notes in Computer Science. Vol. 15154. Cham, Switzerland: Springer; 2025. p. 199–231.
40. Alagic G, Apon D, Cooper D, Dang Q, Liu YK, Miller C, et al. Status report on the third round of the NIST post-quantum cryptography standardization process. Gaithersburg, MD, USA: NIST; 2022.
41. Turino C, Buchanan WJ, Lo O, Thuemmler C. PQC-LEO: an evaluation framework for post-quantum cryptographic algorithms. In: 2025 IEEE 7th International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications (TPS-ISA); Piscataway, NJ, USA: IEEE; 2025. p. 237–47.
42. Abbasi M, Cardoso F, Váz P, Silva J, Martins P. A practical performance benchmark of post-quantum cryptography across heterogeneous computing environments. Cryptography. 2025;9(2):32. doi:10.3390/cryptography9020032.
43. Algar-Fernandez J, Villacís-Vanegas A, Amaro-Aular Y, Cano MD. Benchmarking post-quantum signatures and KEMs on RISC-V. Computers. 2026;15(2):116.
44. Chhetri R. Benchmarking NIST-standardised ML-KEM and ML-DSA on ARM Cortex-M0+: performance, memory, and energy on the RP2040. arXiv:2603.19340. 2026.
45. Marler RT, Arora JS. Survey of multi-objective optimization methods for engineering. Struct Multidiscip Optim. 2004;26(6):369–95. doi:10.1007/s00158-003-0368-6.
46. Courtois N, Klimov A, Patarin J, Shamir A. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In: Advances in Cryptology–EUROCRYPT 2000; 2000 May 14–18; Bruges, Belgium.
47. Ding J, Yang BY, Chen CHO, Chen MS, Cheng CM. New differential-algebraic attacks and reparametrization of rainbow. In: Applied Cryptography and Network Security–ACNS 2008; 2008 Jun 3–6; New York, NY, USA.
48. Ikematsu Y, Furue H, Akiyama R. New security analysis for UOV-based signature candidates with small public key size [Internet]. Gaithersburg, MD, USA: NIST; 2024 [cited 2026 Jul 1]. Available from: https://csrc.nist.gov/csrc/media/Events/2024/fifth-pqc-standardization-conference/documents/papers/new-security-analysis-for-uov.pdf.