Partial Key Exposure Attacks on UOV and Its Variants

Yuki Seto[1] ✉, Hiroki Furue[2], and Atsushi Takayasu[1,3]

[1] The University of Tokyo, Japan
{✉sitoo,takayasu-a}@g.ecc.u-tokyo.ac.jp
[2] NTT Social Informatics Laboratories, Japan
hiroki.furue@ntt.com
[3] National Institute of Advanced Industrial Science and Technology, Japan

Abstract. In CRYPTO 2022, Esser et al. proposed a *partial key exposure attack* on several post-quantum cryptographic schemes including Rainbow which is a variant of *UOV*. The task of the attack is to recover a full secret key from its partial information such as a secret key with symmetric/asymmetric bit errors. One of the techniques Esser et al. developed is a *partial enumeration* that combines the standard algorithms to solve the MQ problem with enumeration. Although an efficient attack on Rainbow was proposed, UOV and its variants have still been paid much attention since UOV and its three variants, i.e., MAYO, QR-UOV and SNOVA, were selected as the Round 2 candidates of the additional call for digital signature schemes proposal by NIST. In this paper, we analyze partial key exposure attacks on UOV, MAYO, and QR-UOV. Although our proposed attacks use the partial enumeration, we refine their enumeration strategy. We employ two enumeration strategies and analyze the complexity of the proposed attacks. Then, we find a structural difference between UOV and its variants to resist partial enumeration. Specifically, the partial enumeration is effective if the number of vinegar variables is smaller than the number of equations and the order of a finite field is small. As a result, the proposed attack is the most effective on MAYO. While our attacks on UOV and QR-UOV are effective only when the symmetric error probabilities are 0.11 and 0.05, respectively, that on MAYO is effective even when the probability is close to 0.5.

Keywords: post-quantum cryptography · multivariate cryptography · UOV · partial key exposure attack.

1. Introduction

1.1 Background

UOV. To ensure the security of public key cryptosystems, the corresponding mathematical problems have to be computationally hard. In the cases of RSA [39] and elliptic curve cryptography [29,33], the prime factorization problem and the elliptic curve discrete logarithm problem have to be computationally hard.

However, Shor's quantum algorithm [40] can solve these problems in polynomial time. Therefore, post-quantum cryptosystems (PQC) that are believed to resist quantum attacks have been actively studied.

The National Institute of Standards and Technology (NIST) is currently working on a standardization project for post-quantum cryptography. In July 2022, NIST selected one encryption/key-establishment scheme and three digital signature schemes to be standardized. Subsequently, NIST announced an additional call for digital signature schemes in September 2022. Then, 14 algorithms were selected as Round 2 candidates in October 2024 from among 40 submissions.

Among the above 14 algorithms, *UOV* [5] and its variants are arguably strong candidates to be selected due to their compact signature sizes and efficient signing algorithms. UOV initially proposed by Kipnis et al. [28] is a form of multivariate quadratic (MQ) cryptosystems. In general, the MQ problem should be computationally hard to ensure the security of MQ cryptosystems. Since the MQ problem is NP-complete [23], MQ cryptosystems seem to possess a strong security guarantee on the surface. Nevertheless, various critical attacks on MQ cryptosystems including UOV variants have been proposed due to their specific structures. For example, although Rainbow [13] was a flagship variant of UOV and selected as a Round 3 candidate of the NIST PQC competition, Beullens [3] proposed an efficient attack on Rainbow. Therefore, the security of MQ cryptosystems has to be intensively studied.

A point to note is that there have been no critical attacks on UOV [28] for over twenty years. Then, seven schemes were submitted to the NIST call for additional signatures. Among them, four schemes, i.e., UOV, MAYO [4], QR-UOV [21] and SNOVA [44], were selected as Round 2 candidates. Since they are strong candidates to be standardized, we have to analyze the security from both theoretical and practical points of view.

Partial Key Exposure Attacks. To provide a strong security guarantee, it is theoretically interesting to analyze a *partial key exposure attack* in which the attacker obtains not only a public key but also some partial information of a secret key. In the case of RSA, it is widely known that half of the most significant bits of secret primes are sufficient to solve the factorization problem in polynomial time [10]. Similarly, there are several works [8,7,17,41,42] in which attackers obtain some consecutive bits of a secret exponent.

Heninger and Shacham [26] analyzed a more realistic scenario to capture practical side-channel leakages such as cold-boot attacks [24]. Specifically, they introduced an erasure model in which the attacker obtains random fractions of a secret key. Subsequently, the celebrated work has been followed by various papers [25,30,31,36] to analyze further realistic settings. Henecka et al. [25] studied an *error model* in which the attacker obtains an erroneous secret key. To be precise, the error model is also called a symmetric error model since each bit of a secret key is flipped with the same probability. Then, Paterson et al. [36] studied an asymmetric error model in which the attacker also obtains an erroneous

secret key, while the bit flip probabilities depend on whether the actual bit is 0 or 1.

Partial Key Exposure Attacks on UOV Variants. Although partial key exposure attacks on UOV and its variants were also analyzed, the attempts were not successful until recently. For example, Polanco's partial key exposure attack on Rainbow [38,43] can recover a full secret key only when the symmetric bit error probability is roughly 0.001. In CRYPTO 2022, Esser et al. [18] proposed partial key exposure attacks on several PQC schemes including Rainbow, and obtained impressive results. Their proposed attack on Rainbow for NIST security level I can recover a full secret key with 80-bit security when the symmetric (resp. asymmetric[4]) error probability[5] is 0.27 (resp. 0.54). Their proposed attack on Rainbow consists of two steps. Let $\mathcal{S} \circ \mathcal{F} \circ \mathcal{T}$ denote a public key of Rainbow, while the linear maps $\mathcal{S}$ and $\mathcal{T}$, and the quadratic map $\mathcal{F}$ are secret keys. The first step recovers $\mathcal{S}$ by solving the syndrome decoding problem. Then, the second step recovers $\mathcal{T}$ by applying *partial enumeration* that combines the standard algorithms to solve the MQ problem [1,2,6,9,11,12,15,16,20,27,28,32,34] with enumeration.

Since an efficient attack on Rainbow was proposed by Beullens [3], the goal of this paper is to analyze partial key exposure attacks on UOV, MAYO, and QR-UOV. For this purpose, we explain more about Esser et al.'s partial enumeration since the left linear map $\mathcal{S}$ is a specific structure of Rainbow. Briefly speaking, we can recover a secret key of UOV by solving the MQ problem with v variables. Since all parameters of UOV are determined so that the MQ problem is computationally infeasible, we want to utilize the partial information. If the symmetric bit error probability p is very small, an erroneous secret key is almost a correct secret key. Therefore, we can find the correct one by enumerating small errors. However, the approach is not effective if p becomes large. The partial enumeration combines these two naive approaches. Let v_{MQ} and v_{Enum} be parameters such that $v = v_{\mathrm{MQ}} + v_{\mathrm{Enum}}$. The partial enumeration first divides v variables into v_{MQ} and v_{Enum}. We assume that v_{Enum} coordinates of a secret key contain only small errors and enumerate them. Then, we recover the remaining v_{MQ} coordinates by solving the MQ problem with v_{MQ} variables. If we cannot recover a secret key, we divide v variables again. We can expect that the partial enumeration is also effective on UOV, MAYO, and QR-UOV due to the structural similarity to Rainbow. However, the results are not trivial since their structures are not completely the same.

1.2 Our Contribution

In this paper, we propose partial key exposure attacks on UOV and its two variants, MAYO and QR-UOV. We follow Esser et al.'s partial enumeration.

[4] In their asymmetric setting, a bit flipping probability from 1 to 0 is variable, while that from 0 to 1 is a fixed probability 0.001.

[5] Although Esser et al. also analyzed the erasure setting, we focus on the error setting since the latter is more realistic and technically difficult.

Moreover, we refine the enumeration strategy. To be honest, since the description of Esser et al.'s enumeration strategy is not clear, we cannot follow it completely. Instead, we provide complete descriptions of our proposed two enumeration strategies. We analyze all cases and compare the qualities of these methods. Throughout the paper, we focus on parameter sets for NIST security level I (SL I) which achieves 143-bit security. We say that a partial key exposure attack is effective if the complexity is less than 143-bit.

Based on our estimates, we can conclude that MAYO is much weaker than UOV and QR-UOV against partial key exposure attacks. In particular, partial key exposure attacks on UOV (resp. QR-UOV) are effective only when the error probabilities are less than 0.11 (resp. 0.05), while those on MAYO are effective even when the error probability is close to 0.5. Since Esser et al. only attacked Rainbow by using partial enumeration, they could not provide structural observations on when the method is more effective. In contrast, since we attack UOV, MAYO, and QR-UOV, we can find two structural differences among the three schemes. At first, partial key exposure attacks are effective if the number of vinegar variables v is smaller than the number of equations m. To define the parameter v_{MQ} of partial enumeration, we cannot set an arbitrary value. In short, v_{MQ} has to be smaller than the number of equations m. If we set larger $v_{\mathrm{MQ}} > m$, the partial enumeration may output an incorrect secret key even if the enumeration step finds the correct solution. On the other hand, v_{MQ} is always smaller than v by definition. Therefore, we can set an arbitrary v_{MQ} if $v \leq m$ holds. Since MAYO satisfies the condition, our attack is effective on MAYO. However, QR-UOV also satisfies the condition, while a partial key exposure attack is not effective. The problem stems from the large order of the finite field used in schemes. Although the orders of UOV (resp. MAYO) are 256 (resp. 16), that of QR-UOV is 7^{10}. Therefore, the enumeration step becomes more time-consuming. Summarizing the above discussion, this is the first result to clarify the structure of UOV to resist partial key exposure attacks.

1.3 Organization of the Paper

In Section 2, we provide the preliminaries about the MQ problem, UOV variants, partial key exposure attacks, and partial enumeration. In Section 3, we propose partial key exposure attacks on UOV variants under the symmetric error setting and estimate their complexity. In section 4, we modify the attack from Section 3 and apply it to asymmetric errors.

2. Preliminaries

2.1 Multivariate Quadratic Problem

As explained in Section 1.1, UOV and its variants are constructed using the hardness of solving multivariate quadratic equations over a finite field. This problem is called the multivariate quadratic (MQ) problem and is defined as follows.

Definition 1 (MQ Problem). *Let n be the number of variables, m be the number of equations, and q be the order of the finite field. The MQ problem* $\mathrm{MQ}(n, m, q)$ *is a problem that takes as input a system of quadratic polynomials*

$$\mathcal{F} = (\mathcal{F}_1, \mathcal{F}_2, \ldots, \mathcal{F}_m) \in \mathbb{F}_q[x_1, \ldots, x_n]^m$$

over the finite field $\mathbb{F}_q$ to find one solution to the quadratic equations

$$\mathcal{F}(\mathbf{x}) = \mathbf{0}.$$

When $n \leq m$ (resp. $n > m$) holds, $\mathrm{MQ}(n, m, q)$ *is said to be* overdetermined *(resp.* underdetermined*).*

In the "very" overdetermined MQ problem where m is much greater than n, a random instance has no solution with high probability [22]. On the other hand, In underdetermined instances, the expected number of solutions increases exponentially [22].

The MQ problem is known to be NP-complete [23]. There are many studies [1,2,6,9,11,12,15,16,20,27,28,32,34] on solving the MQ problem, and various algorithms have been proposed so far. `CryptographicEstimators` [19] provides estimations of the complexity of these algorithms. Hereafter, we use the best complexity provided by `CryptographicEstimators` v1.4.0[6] as the estimated value of $\mathcal{C}_{\mathrm{MQ}}(n, m, q)$, where $\mathcal{C}_{\mathrm{MQ}}(n, m, q)$ represents the complexity of solving $\mathrm{MQ}(n, m, q)$.

2.2 Unbalanced Oil and Vinegar (UOV)

We explain the public and secret keys of UOV [5], which are of interest in partial key exposure attacks. UOV is parameterized by positive integers $n, m, v, o = n-v$ and an order of a finite field q, and its public/secret keys have the following format.

Secret Key. A (random) matrix

$$\mathbf{T} = \begin{pmatrix} \mathbf{t}_1 & \mathbf{t}_2 & \ldots & \mathbf{t}_o \end{pmatrix} \in \mathbb{F}_q^{v \times o}.$$

Public Key. A system of quadratic polynomials

$$\mathcal{P} = (\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_m) \in \mathbb{F}_q[x_1, \ldots, x_n]^m$$

which satisfies

$$\mathcal{P}(\mathbf{t}) = 0$$

for all $\mathbf{t} \in \mathrm{Span}\left\{ \begin{bmatrix} \mathbf{t}_1 \\ \mathbf{e}_1 \end{bmatrix}, \ldots, \begin{bmatrix} \mathbf{t}_o \\ \mathbf{e}_o \end{bmatrix} \right\}$ (the *secret subspace*), where $\mathbf{e}_i$ denotes the unit vector in $\mathbb{F}_q^o$ whose i-th component is 1.

Table 1: UOV parameter sets for NIST SL I

Name	n	v	o	m	q
`uov-Ip`	112	68	44	44	256
`uov-Is`	160	96	64	64	16

In UOV, the value of the parameter o is equal to the value of m. Table 1 shows the parameter sets of UOV [5] proposed for NIST security level I, which we consider in our attack estimations.

We explain an important result used in our partial key exposure attacks.

Theorem 1 (Theorem 1 in [37]). *When $n \leq 3m$, a UOV secret key can be obtained in polynomial time using the public key $\mathcal{P}$ and any non-zero vector included in the secret subspace.*

2.3 UOV Variants

In this paper, we develop partial key exposure attacks on UOV and two of its variants, MAYO and QR-UOV. These two variants both have the UOV parameters n, m, v, o, and q, and their public and secret keys have the same structure as those of UOV. We provide a brief overview of the distinguishing features of the key structures of these two UOV variants compared to UOV.

MAYO. MAYO [4] has an additional parameter k in addition to the UOV parameters. One notable feature of MAYO is that o is much smaller than m, which allows us to choose smaller parameters without reducing its security.

Table 2 shows the parameter sets of MAYO proposed for NIST security level I.

Table 2: MAYO parameter sets for NIST SL I

Name	n	v	o	m	q	k
MAYO$_1$	66	58	8	64	16	9
MAYO$_2$	78	60	18	64	16	4

QR-UOV. QR-UOV has an additional parameter ℓ in addition to the UOV parameters. In QR-UOV, the secret key $\mathbf{T}$ and the matrices that represent each quadratic form $\mathcal{P}_i$ of the public key $\mathcal{P}$ have the structure of a block matrix composed of $\ell \times \ell$ matrices. Using a specific ℓ-degree polynomial f, each block

[6] v1.4.0 is the latest version as of October 2024.

represents an element of the quotient ring $\mathbb{F}_q[x]/(f)$. The polynomial f is publicly chosen from the irreducible polynomials to ensure the security against existing attacks [21]. In the secret key, an element $g \in \mathbb{F}_q[x]/(f)$ corresponds to a matrix $\Phi_g^f \in \mathbb{F}_q^{\ell \times \ell}$ that satisfies the following equation:

$$(1, x, \ldots, x^{\ell-1})\Phi_g^f = (g, xg, \ldots, x^{\ell-1}g). \tag{1}$$

For example, if $q = 7$ and $f = x^3 - 3x - 1$, a matrix corresponding to $g = 3 + 5x + 2x^2 \in \mathbb{F}_q[x]/(f)$ is

$$\Phi_g^f = \begin{pmatrix} 3 & 2 & 5 \\ 5 & 2 & 3 \\ 2 & 5 & 2 \end{pmatrix}.$$

Since an $\ell \times \ell$ matrix over $\mathbb{F}_q$ contains ℓ^2 elements of $\mathbb{F}_q$, while an element of $\mathbb{F}_q[x]/(f)$ can be represented by ℓ elements of $\mathbb{F}_q$, the public/secret keys can be compressed by replacing each block in the matrices with its corresponding element in $\mathbb{F}_q[x]/(f) \simeq \mathbb{F}_{q^\ell}$. The compressed public/secret keys also have the same structure as those of UOV with n/ℓ variables and m equations. This transformation is called the Pull-back method [21]. The parameter sets corresponding to this "compressed" key are shown in Table 3.

Table 3: Compressed QR-UOV parameter sets for NIST SL I

Name[7]	n	v	o	m	q
QR-UOV-Ia	84	74	10	100	7^{10}
QR-UOV-Ib	75	55	20	60	31^3
QR-UOV-Ic	67	60	7	70	31^{10}
QR-UOV-Id	70	52	18	54	127^3

2.4 Key Exposure Model

In our partial key exposure attacks, we consider the *error model* [18] as the key exposure model. In the error model, the attacker obtains an erroneous version of a secret key.

Definition 2 (Symmetric Error Model). *Let $N \in \mathbb{N}$ and $\mathbf{k} = (k_1, \ldots, k_N) \in \{0,1\}^N$ be a binary representation of a secret key. Further let $\boldsymbol{\varepsilon} = (\varepsilon_1, \ldots, \varepsilon_N) \in \{0,1\}^N$ be an error vector. The attacker obtains*

$$\tilde{\mathbf{k}} \coloneqq \mathbf{k} \oplus \boldsymbol{\varepsilon} = (k_1 \oplus \varepsilon_1, \ldots, k_N \oplus \varepsilon_N)$$

[7] The names of each parameter set are not defined in the QR-UOV submission [21], and are assigned in this paper.

where the binary operator $\oplus$ *represents the XOR operation. For each* $i \in \{1, \ldots, N\}$, $\varepsilon_i = 1$ *holds with an error probability* p.

Note that if the error probability p is greater than 0.5, the bit-inverted $\tilde{\mathbf{k}}$ can also be considered as a key obtained in the error model, with an error probability of $1 - p < 0.5$. Therefore, it is sufficient to consider only the case where $p \leq 0.5$.

In Definition 2, we considered a *symmetric error* setting, i.e., the probability of 0 flipping to 1 ($p_{0\to1}$), is equal to the probability of 1 flipping to 0 ($p_{1\to0}$). However, in practical side-channel attacks such as cold-boot attacks [24], these two probabilities might differ. We follow Esser et al. [18] and assume a symmetric error setting in our detailed explanation of the proposed attacks (Section 3) for simplicity. For asymmetric errors, we provide an overview and an analysis of the attack in Section 4.

2.5 Partial Enumeration

We explain the approach of partial enumeration used by Esser et al. [18] in their partial key exposure attack on Rainbow.

Under the condition of $n \leq 3m$, Theorem 1 shows that we can recover the whole UOV secret key $\mathbf{T}$ from only one column of $\mathbf{T}$. Our targets in this paper, UOV, MAYO, and QR-UOV (compressed key), satisfy this condition. Therefore, the goal of the partial key exposure attacks against them is to recover one column of $\mathbf{T}$ from an erroneous secret key $\tilde{\mathbf{T}}$.

From the definition of the UOV public key, the i-th column $\mathbf{t}_i$ of $\mathbf{T}$ is a solution to the system of m equations in v variables:

$$\mathcal{P}\begin{pmatrix}\mathbf{x}\\ \mathbf{e}_i\end{pmatrix} = \mathbf{0}. \tag{2}$$

There are two naive methods to recover $\mathbf{t}_i$ from a column $\tilde{\mathbf{t}}_i$ of $\tilde{\mathbf{T}}$:

(a) Enumerate some of the errors on $\tilde{\mathbf{T}}$ and check if it satisfies equation (2). Since the number of error bits is very close to its expected value with high probability, it is possible to obtain a solution of (2) with lower complexity. For example, when enumerating errors where only γ of the v elements in $\mathbb{F}_q$ contain errors, the number of enumerations becomes $\binom{v}{\gamma}(q-1)^\gamma$. However, when the number of errors is large, enumeration is inefficient compared to methods specialized for the MQ problem.

(b) Solve equation (2) as an instance of the MQ(v, m, q). The attacker can certainly obtain a solution in one process, but cannot utilize the information from the erroneous key so it is relatively inefficient when $\tilde{\mathbf{T}}$ does not contain many errors.

In order to take the advantages of both, Esser et al. [18] used an approach called *partial enumeration*, which is a hybrid approach of methods (a) and (b) in their partial key exposure attack on Rainbow. More precisely, they took an approach of enumerating errors on some elements of the column to be recovered like (a), and obtaining the remaining elements by solving the MQ problem like (b).

3. Our Attacks on Symmetric Errors

In this section, we propose partial key exposure attacks on UOV and its variants under the symmetric error setting, and estimate their complexity to compare the resistance of the schemes. As mentioned in Section 2.5, our proposed attacks are based on the concept of partial enumeration. We explain our attacks in Section 3.1, then explain the method of estimating the complexity and show our estimation results on the three schemes, UOV [5], MAYO [4], and QR-UOV [21] in Section 3.2.

3.1 Our Attacks

The integer b_q denotes the number of bits used to represent the elements of $\mathbb{F}_q$. Our attacks on UOV and its variants are parametrized by the integers v_{MQ}, $v_{\mathrm{Enum}} \coloneqq v - v_{\mathrm{MQ}}$ and the set of errors $\mathcal{E} \in \{0,1\}^{b_q v_{\mathrm{Enum}}}$ to be enumerated. Since we can recover the entire UOV secret key from a single column by Theorem 1, we focus on recovering a single column in this paper. Until a column is recovered, the proposed method repeats the partial enumeration process shown below for each column $\tilde{\mathbf{t}} \coloneqq \tilde{\mathbf{t}}_i$ of $\tilde{\mathbf{T}}$.

1. Divide indices $\{1, \ldots, v\}$ of the vector $\tilde{\mathbf{t}}$ randomly into two disjoint index sets: I_{Enum} and I_{MQ}, where $|I_{\mathrm{Enum}}| = v_{\mathrm{Enum}}$ and $|I_{\mathrm{MQ}}| = v_{\mathrm{MQ}}$.
2. Enumerate errors in $\mathcal{E}$ on I_{Enum}, then solve (2) as an instance of $\mathrm{MQ}(v_{\mathrm{MQ}}, m, q)$ to recover elements at I_{MQ} for each enumerated error. If a solution is found, output it as the recovered $\mathbf{t}_i$.

Hereafter, the term *I-part* will be used to denote the entire elements whose index is included in an index set I. Figure 1 shows an example of I_{Enum} and I_{MQ}.

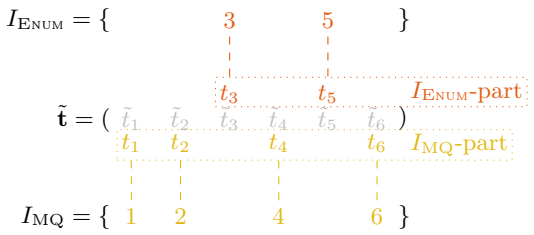

Fig. 1: An example of I_{Enum} and I_{MQ}

Step 2. successfully stops if $\mathcal{E}$ contains the actual error in the I_{Enum}-part of $\tilde{\mathbf{t}}$. v_{MQ} satisfies $0 \leq v_{\mathrm{MQ}} \leq \min\{v, m\}$ so that the MQ instance solved in Step 2. is overdetermined. This is because there is a high possibility that a column different from the original secret key will be obtained as the result of Step 2.. Because

of this condition, we have to enumerate errors on at least $v - m$ elements when $v > m$ holds, even if the error probability p is large and $\tilde{\mathbf{t}}$ contains numerous errors.

The overall process of the proposed method is shown in Algorithm 1.

Algorithm 1 Partial Key Exposure Attacks on UOV and its variants

Parameters $v_{\mathrm{MQ}}, v_{\mathrm{Enum}} \coloneqq v - v_{\mathrm{MQ}}, \mathcal{E}$
Input A UOV public key $\mathcal{P}$ and an erroneous UOV secret key $\tilde{\mathbf{T}}$
Output A recovered column $\mathbf{t}$ of $\tilde{\mathbf{T}}$ and its index i

loop
  for $i = 1, 2, \ldots, o$ **do**
    $\tilde{\mathbf{t}} \leftarrow \tilde{\mathbf{t}}_i$: Copy the i-th column of $\tilde{\mathbf{T}}$.
    $I_{\mathrm{MQ}} \xleftarrow{\$} \{I \subset \{1, 2, \ldots, v\} \mid |I| = v_{\mathrm{MQ}}\}$
    $I_{\mathrm{Enum}} \leftarrow \{1, 2, \ldots, v\} \setminus I_{\mathrm{MQ}}$
    for $\boldsymbol{\varepsilon} \in \mathcal{E}$ **do**
      Recover the error $\boldsymbol{\varepsilon}$ on I_{Enum}-part of $\tilde{\mathbf{t}}$.
      Solve an instance of $\mathrm{MQ}(v_{\mathrm{MQ}}, m, q)$ to recover the rest.
      if obtained $\mathbf{t}$ as a solution **then**
        return $\mathbf{t}, i$
      end if
    end for
  end for
end loop

It should be noted that the error subset $\mathcal{E}$ to be enumerated remains a matter of discretion. Since the complexity and the success probability of a single partial enumeration process depend on $\mathcal{E}$, we want to choose a "good" $\mathcal{E}$ for attacks. Intuitively, a smaller $\mathcal{E}$ leads to lower complexity, while including highly probable errors leads to a higher probability of success. We propose two different strategies to decide $\mathcal{E}$ and then compare them in the analysis part of this section.

EnumGamma Strategy. We use an integer $\gamma \in [0, v_{\mathrm{Enum}}]$ as an additional parameter and define $\mathcal{E}$ as the set of all errors on v_{Enum} elements where exactly γ of the v_{Enum} elements in $\mathbb{F}_q$ are erroneous. There are $\binom{v_{\mathrm{Enum}}}{\gamma}$ ways to choose γ elements from the v_{Enum} elements. For each choice of γ elements, there are $(2^{b_q} - 1)^\gamma$ candidates of bit errors. Thus, in one partial enumeration process, the MQ problem is solved $|\mathcal{E}| = \binom{v_{\mathrm{Enum}}}{\gamma}(2^{b_q} - 1)^\gamma$ times, and its complexity becomes $\binom{v_{\mathrm{Enum}}}{\gamma}(2^{b_q} - 1)^\gamma \cdot \mathcal{C}_{\mathrm{MQ}}(v_{\mathrm{MQ}}, m, q)$.

We can take advantage of the bias in the number of correct elements in $\tilde{\mathbf{t}}$ by using the EnumGamma strategy. We can set the number of elements with no enumeration $v_{\mathrm{Enum}} - \gamma$ depending on the expected number of elements with no errors, which leads the correct elements to be used in solution without modification as much as possible.

In order to derive the success probability, let ω be the number of erroneous elements contained in $\tilde{\mathbf{t}}$. The probability of successfully recovering $\tilde{\mathbf{t}}$ with a single partial enumeration is

$$\frac{\binom{\omega}{\gamma}\binom{v-\omega}{v_{\mathrm{Enum}}-\gamma}}{\binom{v}{v_{\mathrm{Enum}}}}$$

because the I_{Enum}-part contains exactly γ of ω erroneous elements and $v_{\mathrm{Enum}}-\gamma$ of $v-\omega$ correct elements in successful situations.

EnumGammaLimited Strategy. In addition to γ in EnumGamma, we also use two parameters, $w_{\min}$ and $w_{\max}$. We make $\mathcal{E}$ more "limited" than EnumGamma; We define $\mathcal{E}$ the set of errors where exactly γ of the v_{Enum} elements in $\mathbb{F}_q$ are erroneous and each of γ elements has error bits of at least $w_{\min}$ and at most $w_{\max}$. Since there are $\sum_{w=w_{\min}}^{w_{\max}}\binom{b_q}{w}$ candidates of errors contained in each of the γ erroneous elements, for each choice of γ elements, there are $\left(\sum_{w=w_{\min}}^{w_{\max}}\binom{b_q}{w}\right)^{\gamma}$ candidates of bit errors. The complexity of one partial enumeration process becomes $\binom{v_{\mathrm{Enum}}}{\gamma}\left(\sum_{w=w_{\min}}^{w_{\max}}\binom{b_q}{w}\right)^{\gamma}$.

We can take advantage of the bias in the number of error bits contained in each element by using the EnumGammaLimited strategy. For example, if p is very small, the number of elements with a few error bits will be relatively large. In such cases, we can effectively enumerate the errors by setting $w_{\max}$ to a small value.

Let ω be the number of erroneous elements contained in $\tilde{\mathbf{t}}$, and ω_{in} be the number of elements in $\tilde{\mathbf{t}}$ that have error bits of at least $w_{\min}$ and at most $w_{\max}$. The probability of successfully recovering $\tilde{\mathbf{t}}$ with a single partial enumeration is

$$\frac{\binom{\omega_{\mathrm{in}}}{\gamma}\binom{v-\omega}{v_{\mathrm{Enum}}-\gamma}}{\binom{v}{v_{\mathrm{Enum}}}}$$

because the I_{Enum}-part contains exactly γ of ω_{in} erroneous elements and $v_{\mathrm{Enum}}-\gamma$ of $v-\omega$ correct elements in the successful situations.

3.2 Estimation of Complexity

We follow Esser et al. [18] and assume that the error probability p is known under the error model. Therefore, the attacker can choose attack parameters depending on p. In this paper, for each p, we calculate the minimum complexity with respect to attack parameters as the estimated complexity. Intuitively, a higher v_{Enum} achieves the minimum complexity when the enumeration part contributes well to the reduction of complexity, while a lower v_{Enum} achieves the minimum when enumeration is not efficient on the other hand.

It is important to note that some errors may not be recoverable by our attack with the specified attack parameters. More precisely, if any I_{Enum}-part of

$\tilde{\mathbf{t}}_i$ cannot be recovered by enumerating errors in $\mathcal{E}$, our attack algorithm does not stop. For example, if we employ the EnumGamma strategy with parameter γ, and we obtain $\tilde{\mathbf{T}}$ that has more than $\gamma + v_{\mathrm{MQ}}$ erroneous elements in each column, then we cannot recover any of the column of $\tilde{\mathbf{T}}$ because we can recover at most $\gamma + v_{\mathrm{MQ}}$ elements in partial enumeration. Therefore, we only consider attack parameters for which the probability of obtaining such a $\tilde{\mathbf{T}}$ is less than 0.05 in our estimation. Following the work of Esser et al. [18], we randomly generate 20 pairs of $\mathbf{T}$ and $\tilde{\mathbf{T}}$ for each p and attack parameters, then estimate the attack complexity and take the average as the estimated complexity at the chosen p and parameters.

On QR-UOV, we focus in this paper on recovering the compressed key $\mathbf{T}'$ on $\mathbb{F}_q[x]/(f)$ which we explained in Section 2.3, rather than the original secret key $\mathbf{T}$ of QR-UOV. When we generate $\mathbf{T}$, the element of $\mathbb{F}_q[x]/(f)$ corresponding to each block of $\mathbf{T}$ is generated during the process [21]. Also, from equation (1), the coefficients of the element of $\mathbb{F}_q[x]/(f)$ corresponding to each block of $\mathbf{T}$ appear in the leftmost column. We consider in this paper that these values are obtained under the error model, and perform a partial key exposure attack on $\mathbf{T}'$.

We estimate the bit complexity of our attacks on the parameter sets for NIST security level I of UOV, MAYO, and QR-UOV with each of two strategies proposed in Section 3.1 to decide $\mathcal{E}$. Specifically, we are estimating for `uov-Ip`, MAYO1, and QR-UOV-Ia, respectively. Table 4 shows the specific values of these parameters and the bit complexity required for the most efficient key recovery or universal forgery attacks. As explained in Section 2.1, we use `CryptographicEstimators` [19] to estimate $\mathcal{C}_{\mathrm{MQ}}$.

Table 4: Parameters and the complexity of existing key recovery or universal forgery attacks of each parameter set

Name	n	v	o	m	q	Attack complexity	Attack type
`uov-Ip`	112	68	44	44	256	145	Direct
MAYO1	66	58	8	64	16	143	Reconciliation [14]
QR-UOV-Ia	84	74	10	100	7^{10}	148	Reconciliation [14]

In each plot of complexity, we marked with a dashed line the position of bit complexity = 143, which corresponds to the number of classical gates (2^{143}) required for normal attacks in NIST security level I [35]. Also we marked in each plot of parameters the position of $v_{\mathrm{Enum}} = 24$, which is the minimum value of v_{Enum} that we can choose for `uov-Ip`. The minimum value of v_{Enum} is 0 for the other two parameter sets.

EnumGamma. Figure 2a shows the estimated complexity of our attacks on `uov-Ip`, MAYO1, and QR-UOV-Ia with the EnumGamma strategy, while Figure 2b shows the parameters of our attacks that achieve the best complexity.

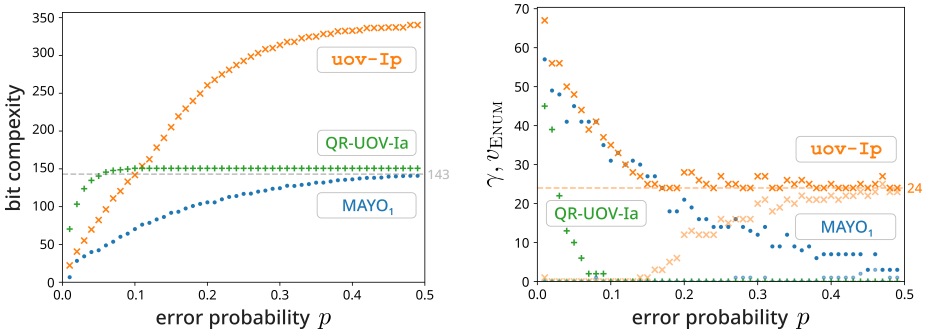

(a) Complexity (b) Parameters

The transparent marks in (b) show the value of γ.

Fig. 2: Bit complexity required and the best parameter v_{ENUM} and γ for the attack in Section 3.1 with EnumGamma

From the estimated complexity in Figure 2a, we find that the bit complexity already exceeds the requirements of NIST security level I (143) at $p = 0.06$ in the QR-UOV-Ia and at $p = 0.11$ in the `uov-Ip`. This fact shows that our attacks are not effective on `uov-Ip` and QR-UOV-Ia when p is large, because applying the normal attack (without partial key) is more efficient in such cases. On the other hand, the bit complexity of the attack against MAYO1 is gradually approaching 143 as p approaches 0.5, and it does not reach 143. Therefore, it is suggested that QR-UOV-Ia has the strongest resistance among the three schemes to our attacks with the EnumGamma strategy, followed by `uov-Ip`, while MAYO1 has less resistance. We list in Table 5 the maximum error probabilities where the attack bit complexity is 143 or less.

This order of resistance can also be found in Figure 2b. As previously stated in Section 3.2, we can measure the effect of combining enumeration by the value of v_{ENUM}. In QR-UOV-Ia, v_{ENUM} is always close to its lowest value 0 when p is larger than about 0.07. This suggests that enumeration of errors is not effective. Conversely, in MAYO1, even at $p \approx 0.5$, v_{ENUM} is higher than 0, indicating that the proposed method is more efficient than finding all elements by solving the MQ problem.

Table 5: Max error probability of attack bit complexity ≤ 143 in Figure 2a

	bit error prob.
`uov-Ip`	0.10
MAYO1	0.49
QR-UOV-Ia	0.05

EnumGammaLimited. Figure 3a shows the estimated complexity of our attacks on `uov-Ip`, MAYO1, and QR-UOV-Ia with the EnumGammaLimited strategy, while Figure 3b shows the parameters of our attacks that achieves the best complexity.

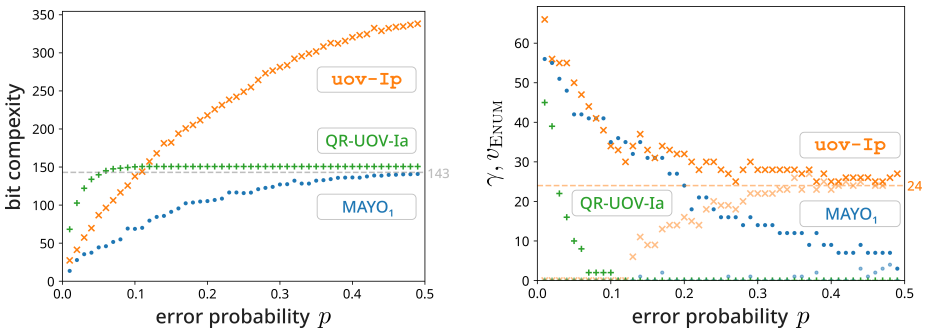

(a) Complexity (b) Parameters

The transparent marks in (b) shows the value of γ.

Fig. 3: Bit complexity required and the best parameter v_{Enum} and γ for the attack in Section 3.1 with EnumGammaLimited

From the estimated complexity in Figure 3a, we find that our attack with EnumGammaLimited is effective on MAYO1 and not effective on `uov-Ip` and QR-UOV-Ia as with our attack with EnumGamma. The notable difference between the results of EnumGamma and EnumGammaLimited is the reduction of complexity on `uov-Ip` at $p \geq 0.1$. This difference slightly reduces the maximum error probabilities shown in Table 6 where the attack bit complexity is 143 or less. The attack complexity on MAYO1 and QR-UOV-Ia is nearly equivalent to that of the attack on each parameter set with EnumGamma.

Table 6: Max error probability of attack bit complexity ≤ 143 in Figure 3a

	bit error prob.
`uov-Ip`	0.11
MAYO1	0.49
QR-UOV-Ia	0.05

Summary and Analysis of Estimation. The two results of our estimation show that both `uov-Ip` and QR-UOV-Ia are much more resistant to our attacks than MAYO1.

The resistance of `uov-Ip` can be explained by the inequality $v > m$. In `uov-Ip`, the number of variables v contained in the column to be recovered is about 20 to 30 more than the number of equations m. In our partial enumeration algorithm, v_{MQ} was determined to be less than or equal to m so that $\mathrm{MQ}(v_{\mathrm{MQ}}, m, q)$ becomes overdetermined. Therefore, even with a large p and numerous errors, we have to enumerate errors on at least $v - m$ elements. As a result, the complexity exceeds that of the usual attack by a large amount. In fact, for `uov-Ip`, the number of elements that need to be enumerated for errors is $v - m = 24$, then the number of enumerations required for its full enumeration is $q^{v-m} = 256^{24} = 2^{192}$, which well explains that the bit complexity exceeds the existing attack by about 200 in the range where p is close to 0.5.

Since $v > m$ does not hold for QR-UOV-Ia, there should be another reason for the resistance of QR-UOV-Ia. The very large value of q may be the cause of the resistance. If q is large, the attacker has to enumerate a huge number of errors even if γ is small. For this reason, increasing the value of γ to increase the number of elements in the error enumeration is not a good idea in QR-UOV.

Moreover, when q is large, there is a high probability that a single element contains at least one error bit. When $\gamma = 0$, our attack does nothing on the I_{ENUM}-part of the column of the erroneous secret key. Therefore, for recovery, the I_{ENUM}-part must not contain any errors. Thus, it is not a good idea to make v_{ENUM} large either. As a result, unless p is very small, our attack requires the same complexity as solving the MQ problem to recover the entire private key.

4. Our Attacks on Asymmetric Errors

4.1 Overview of the Attacks

In the case of asymmetric errors, the Hamming weight of an erroneous bit vector (e.g., an erroneous element) provides information about the number of errors. For instance, consider a scenario where $p_{0\to1} \ll p_{1\to0}$, and the original (error-free) bit vector consists of an equal number of 0s and 1s. In this case, the probability that the bit 1 obtained by the attacker is erroneous is very low, while the probability that the bit 0 is erroneous is relatively high. This information allows the attacker to prioritize the elements with fewer error bits and include them in the I_{ENUM}-part.

We modify the partial enumeration process described in Section 3.1 to take advantage of this information about erroneous elements. We show our modified process below:

1. Divide indices $\{1, \ldots, v\}$ of the vector $\tilde{\mathbf{t}}$ randomly into two disjoint index sets: I_{ENUM} and I_{MQ}, where $|I_{\mathrm{ENUM}}| = v_{\mathrm{ENUM}}$ and $|I_{\mathrm{MQ}}| = v_{\mathrm{MQ}}$. *The probability of getting each division depends on the Hamming weight of I_{ENUM}-part of $\tilde{\mathbf{t}}$.*
2. Enumerate errors in $\mathcal{E}$ on I_{ENUM}, then solve (2) as an instance of $\mathrm{MQ}(v_{\mathrm{MQ}}, m, q)$ to recover elements at I_{MQ} for each enumerated error. If a solution is found, output it as the recovered $\mathbf{t}_i$.

Intuitively, we want the I_{ENUM}-part of $\tilde{\mathbf{t}}$ to contain elements that are unlikely to be erroneous. We propose in this paper to determine the weight of selecting each division as the inverse of the probability of the I_{ENUM}-part containing an error.

4.2 Estimation of Complexity

In this paper, we estimate the bit complexity of our attacks for asymmetric errors on MAYO1. Following previous works [18,24], we assume that $p_{0\to1}$ is very small in our analysis, and fix $p_{0\to1}$ at 10^{-3}. Also, we fix $\mathcal{E} = \{00\cdots00\}$ for simplicity, which implies that we don't enumerate errors on the I_{ENUM}-part. This is because the best attack parameter γ under the symmetric error setting is almost always 0 in the attack on MAYO1.

Figure 4 shows the estimated complexity of our attacks. We find that our attack for asymmetric errors are always effective on MAYO1.

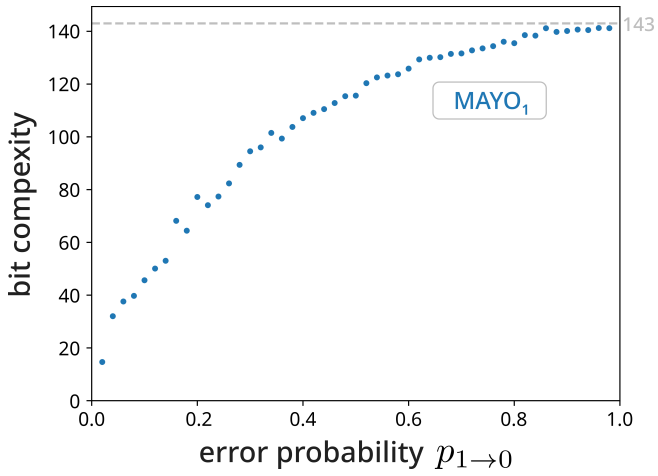

Fig. 4: Bit complexity required for our attack on MAYO1 under the asymmetric error setting with fixing $p_{0\to1} = 10^{-3}$

From the estimated complexity in Figure 4, we find that our attacks are effective on MAYO1 even in the case of asymmetric errors.

5. Conclusion

In this paper, we proposed partial key exposure attacks on UOV and its variants, MAYO and QR-UOV, based on the attack on Rainbow by Esser et al. [18]. We employed two enumeration strategies to refine the partial enumeration technique initially proposed by Esser et al. Also, we estimated the complexity of

our attacks on uov-Ip, MAYO1 and QR-UOV-Ia. Our estimation demonstrated that our attacks are most effective on MAYO1, while uov-Ip and QR-UOV-Ia have stronger resistance. Our results indicate that our attacks are particularly effective when $v < m$ and q is small.

Future works will focus on developing partial key exposure attacks on SNOVA, which aims to compare the resistance of all Round 2 candidate UOV variants to such attacks. Additionally, we may find better ways to determine the set of errors to improve the efficiency of our attacks, although theoretical analysis is expected to be very challenging.

Acknowledgments. This research was partially supported by JST CREST Grant Number JPMJCR2113, Japan, and JSPS KAKENHI Grant Numbers JP22KJ0554, JP24K02939, Japan.

References

1. Bardet, M., Faugère, J.C., Salvy, B., Spaenlehauer, P.J.: On the complexity of solving quadratic boolean systems. Journal of Complexity **29**(1), 53–75 (2013). https://doi.org/https://doi.org/10.1016/j.jco.2012.07.001, https://www.sciencedirect.com/science/article/pii/S0885064X12000611
2. Bettale, L., Faugere, J.C., Perret, L.: Hybrid approach for solving multivariate systems over finite fields. Journal of Mathematical Cryptology **3**(3), 177–197 (2009)
3. Beullens, W.: Breaking Rainbow takes a weekend on a laptop. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part II. Lecture Notes in Computer Science, vol. 13508, pp. 464–479. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15979-4_16
4. Beullens, W., Campos, F., Celi, S., Hess, B., Kannwischer, M.J.: MAYO. NIST CSRC (2023)
5. Beullens, W., Chen, M.S., Ding, J., Gong, B., Kannwischer, M.J., Patarin, J., Peng, B.Y., Schmidt, D., Shih, C.J., Tao, C., Yang, B.Y.: UOV: Unbalanced Oil and Vinegar. NIST CSRC (2023)
6. Björklund, A., Kaski, P., Williams, R.: Solving systems of polynomial equations over GF(2) by a parity-counting self-reduction. In: Baier, C., Chatzigiannakis, I., Flocchini, P., Leonardi, S. (eds.) ICALP 2019: 46th International Colloquium on Automata, Languages and Programming. Leibniz International Proceedings in Informatics (LIPIcs), vol. 132, pp. 26:1–26:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Patras, Greece (Jul 9–12, 2019). https://doi.org/10.4230/LIPIcs.ICALP.2019.26
7. Blömer, J., May, A.: New partial key exposure attacks on RSA. In: Boneh, D. (ed.) Advances in Cryptology – CRYPTO 2003. Lecture Notes in Computer Science, vol. 2729, pp. 27–43. Springer, Berlin, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2003). https://doi.org/10.1007/978-3-540-45146-4_2
8. Boneh, D., Durfee, G., Frankel, Y.: An attack on RSA given a small fraction of the private key bits. In: Ohta, K., Pei, D. (eds.) Advances in Cryptology – ASIACRYPT'98. Lecture Notes in Computer Science, vol. 1514, pp. 25–34. Springer, Berlin, Heidelberg, Germany, Beijing, China (Oct 18–22, 1998). https://doi.org/10.1007/3-540-49649-1_3

9. Bouillaguet, C., Chen, H.C., Cheng, C.M., Chou, T., Niederhagen, R., Shamir, A., Yang, B.Y.: Fast exhaustive search for polynomial systems in $\mathbb{F}_2$. In: Mangard, S., Standaert, F.X. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2010. Lecture Notes in Computer Science, vol. 6225, pp. 203–218. Springer, Berlin, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–20, 2010). https://doi.org/10.1007/978-3-642-15031-9_14
10. Coppersmith, D.: Finding a small root of a bivariate integer equation; factoring with high bits known. In: Maurer, U.M. (ed.) Advances in Cryptology – EUROCRYPT'96. Lecture Notes in Computer Science, vol. 1070, pp. 178–189. Springer, Berlin, Heidelberg, Germany, Saragossa, Spain (May 12–16, 1996). https://doi.org/10.1007/3-540-68339-9_16
11. Courtois, N., Goubin, L., Meier, W., Tacier, J.D.: Solving underdefined systems of multivariate quadratic equations. In: Naccache, D., Paillier, P. (eds.) PKC 2002: 5th International Workshop on Theory and Practice in Public Key Cryptography. Lecture Notes in Computer Science, vol. 2274, pp. 211–227. Springer, Berlin, Heidelberg, Germany, Paris, France (Feb 12–14, 2002). https://doi.org/10.1007/3-540-45664-3_15
12. Courtois, N., Klimov, A., Patarin, J., Shamir, A.: Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In: Preneel, B. (ed.) Advances in Cryptology – EUROCRYPT 2000. Lecture Notes in Computer Science, vol. 1807, pp. 392–407. Springer, Berlin, Heidelberg, Germany, Bruges, Belgium (May 14–18, 2000). https://doi.org/10.1007/3-540-45539-6_27
13. Ding, J., Chen, M.S., Petzoldt, A., Schmidt, D., Yang, B.Y., Kannwischer, M., Patarin, J.: Rainbow. NIST CSRC (2020)
14. Ding, J., Yang, B.Y., Chen, C.H.O., Chen, M.S., Cheng, C.M.: New differential-algebraic attacks and reparametrization of Rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 08: 6th International Conference on Applied Cryptography and Network Security. Lecture Notes in Computer Science, vol. 5037, pp. 242–257. Springer, Berlin, Heidelberg, Germany, New York, NY, USA (Jun 3–6, 2008). https://doi.org/10.1007/978-3-540-68914-0_15
15. Dinur, I.: Cryptanalytic applications of the polynomial method for solving multivariate equation systems over GF(2). In: Canteaut, A., Standaert, F.X. (eds.) Advances in Cryptology – EUROCRYPT 2021, Part I. Lecture Notes in Computer Science, vol. 12696, pp. 374–403. Springer, Cham, Switzerland, Zagreb, Croatia (Oct 17–21, 2021). https://doi.org/10.1007/978-3-030-77870-5_14
16. Dinur, I.: Improved algorithms for solving polynomial systems over GF(2) by multiple parity-counting. In: Marx, D. (ed.) 32nd Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2550–2564. ACM-SIAM, Virtual Conference (Jan 10–13, 2021). https://doi.org/10.1137/1.9781611976465.151
17. Ernst, M., Jochemsz, E., May, A., de Weger, B.: Partial key exposure attacks on RSA up to full size exponents. In: Cramer, R. (ed.) Advances in Cryptology – EUROCRYPT 2005. Lecture Notes in Computer Science, vol. 3494, pp. 371–386. Springer, Berlin, Heidelberg, Germany, Aarhus, Denmark (May 22–26, 2005). https://doi.org/10.1007/11426639_22
18. Esser, A., May, A., Verbel, J.A., Wen, W.: Partial key exposure attacks on BIKE, Rainbow and NTRU. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022, Part III. Lecture Notes in Computer Science, vol. 13509, pp. 346–375. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15982-4_12

19. Esser, A., Verbel, J., Zweydinger, F., Bellini, E.: `CryptographicEstimators`: a software library for cryptographic hardness estimation. Cryptology ePrint Archive, Paper 2023/589 (2023), https://eprint.iacr.org/2023/589
20. Faugère, J.C.: A new efficient algorithm for computing gröbner bases without reduction to zero (f5). In: Proceedings of the 2002 International Symposium on Symbolic and Algebraic Computation. p. 75–83. ISSAC '02, Association for Computing Machinery, New York, NY, USA (2002). https://doi.org/10.1145/780506.780516, https://doi.org/10.1145/780506.780516
21. Furue, H., Ikematsu, Y., Hoshino, F., Takagi, T., Yasuda, K., Miyazawa, T., Saito, T., Nagai, A.: QR-UOV. NIST CSRC (2023)
22. Fusco, G., Bach, E.: Phase transition of multivariate polynomial systems. In: Cai, J.Y., Cooper, S.B., Zhu, H. (eds.) Theory and Applications of Models of Computation. pp. 632–645. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
23. Garey, M.R., Johnson, D.S.: Computers and Intractability; A Guide to the Theory of NP-Completeness. W. H. Freeman & Co., USA (1990)
24. Halderman, J.A., Schoen, S.D., Heninger, N., Clarkson, W., Paul, W., Calandrino, J.A., Feldman, A.J., Appelbaum, J., Felten, E.W.: Lest we remember: Cold boot attacks on encryption keys. In: van Oorschot, P.C. (ed.) USENIX Security 2008: 17th USENIX Security Symposium. pp. 45–60. USENIX Association, San Jose, CA, USA (Jul 28 – Aug 1, 2008)
25. Henecka, W., May, A., Meurer, A.: Correcting errors in RSA private keys. In: Rabin, T. (ed.) Advances in Cryptology – CRYPTO 2010. Lecture Notes in Computer Science, vol. 6223, pp. 351–369. Springer, Berlin, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 2010). https://doi.org/10.1007/978-3-642-14623-7_19
26. Heninger, N., Shacham, H.: Reconstructing RSA private keys from random key bits. In: Halevi, S. (ed.) Advances in Cryptology – CRYPTO 2009. Lecture Notes in Computer Science, vol. 5677, pp. 1–17. Springer, Berlin, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2009). https://doi.org/10.1007/978-3-642-03356-8_1
27. Joux, A., Vitse, V.: A crossbred algorithm for solving boolean polynomial systems. In: Kaczorowski, J., Pieprzyk, J., Pomykała, J. (eds.) Number-Theoretic Methods in Cryptology. pp. 3–21. Springer International Publishing, Cham (2018)
28. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced Oil and Vinegar signature schemes. In: Stern, J. (ed.) Advances in Cryptology – EUROCRYPT'99. Lecture Notes in Computer Science, vol. 1592, pp. 206–222. Springer, Berlin, Heidelberg, Germany, Prague, Czech Republic (May 2–6, 1999). https://doi.org/10.1007/3-540-48910-X_15
29. Koblitz, N.: Elliptic curve cryptosystems. Mathematics of computation **48**(177), 203–209 (1987)
30. Kunihiro, N., Honda, J.: RSA meets DPA: Recovering RSA secret keys from noisy analog data. In: Batina, L., Robshaw, M. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2014. Lecture Notes in Computer Science, vol. 8731, pp. 261–278. Springer, Berlin, Heidelberg, Germany, Busan, South Korea (Sep 23–26, 2014). https://doi.org/10.1007/978-3-662-44709-3_15
31. Kunihiro, N., Shinohara, N., Izu, T.: Recovering RSA secret keys from noisy key bits with erasures and errors. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013: 16th International Conference on Theory and Practice of Public Key Cryptography. Lecture Notes in Computer Science, vol. 7778, pp. 180–197. Springer, Berlin, Heidelberg, Germany, Nara, Japan (Feb 26 – Mar 1, 2013). https://doi.org/10.1007/978-3-642-36362-7_12

32. Lokshtanov, D., Paturi, R., Tamaki, S., Williams, R.R., Yu, H.: Beating brute force for systems of polynomial equations over finite fields. In: Klein, P.N. (ed.) 28th Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 2190–2202. ACM-SIAM, Barcelona, Spain (Jan 16–19, 2017). https://doi.org/10.1137/1.9781611974782.143
33. Miller, V.S.: Use of elliptic curves in cryptography. In: Williams, H.C. (ed.) Advances in Cryptology – CRYPTO'85. Lecture Notes in Computer Science, vol. 218, pp. 417–426. Springer, Berlin, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 1986). https://doi.org/10.1007/3-540-39799-X_31
34. Miura, H., Hashimoto, Y., Takagi, T.: Extended algorithm for solving underdefined multivariate quadratic equations. In: Gaborit, P. (ed.) Post-Quantum Cryptography - 5th International Workshop, PQCrypto 2013. pp. 118–135. Springer, Berlin, Heidelberg, Germany, Limoges, France (Jun 4–7, 2013). https://doi.org/10.1007/978-3-642-38616-9_8
35. NIST: Call for Additional Digital Signature Schemes for the Post-Quantum Cryptography Standardization Process (2022)
36. Paterson, K.G., Polychroniadou, A., Sibborn, D.L.: A coding-theoretic approach to recovering noisy RSA keys. In: Wang, X., Sako, K. (eds.) Advances in Cryptology – ASIACRYPT 2012. Lecture Notes in Computer Science, vol. 7658, pp. 386–403. Springer, Berlin, Heidelberg, Germany, Beijing, China (Dec 2–6, 2012). https://doi.org/10.1007/978-3-642-34961-4_24
37. Pébereau, P.: One vector to rule them all: Key recovery from one vector in UOV schemes. In: Saarinen, M.J., Smith-Tone, D. (eds.) Post-Quantum Cryptography - 15th International Workshop, PQCrypto 2024, Part II. pp. 92–108. Springer, Cham, Switzerland, Oxford, UK (Jun 12–14, 2024). https://doi.org/10.1007/978-3-031-62746-0_5
38. Polanco, R.V.: Cold boot attacks on post-quantum schemes. Ph.D. thesis, Royal Holloway, University of London (2019)
39. Rivest, R.L., Shamir, A., Adleman, L.: A method for obtaining digital signatures and public-key cryptosystems. Commun. ACM **21**(2), 120–126 (Feb 1978). https://doi.org/10.1145/359340.359342, https://doi.org/10.1145/359340.359342
40. Shor, P.W.: Algorithms for quantum computation: Discrete logarithms and factoring. In: 35th Annual Symposium on Foundations of Computer Science. pp. 124–134. IEEE Computer Society Press, Santa Fe, NM, USA (Nov 20–22, 1994). https://doi.org/10.1109/SFCS.1994.365700
41. Takayasu, A., Kunihiro, N.: Partial key exposure attacks on RSA: Achieving the boneh-durfee bound. In: Joux, A., Youssef, A.M. (eds.) SAC 2014: 21st Annual International Workshop on Selected Areas in Cryptography. Lecture Notes in Computer Science, vol. 8781, pp. 345–362. Springer, Cham, Switzerland, Montreal, QC, Canada (Aug 14–15, 2014). https://doi.org/10.1007/978-3-319-13051-4_21
42. Takayasu, A., Kunihiro, N.: A tool kit for partial key exposure attacks on RSA. In: Handschuh, H. (ed.) Topics in Cryptology – CT-RSA 2017. Lecture Notes in Computer Science, vol. 10159, pp. 58–73. Springer, Cham, Switzerland, San Francisco, CA, USA (Feb 14–17, 2017). https://doi.org/10.1007/978-3-319-52153-4_4
43. Villanueva-Polanco, R.: Cold boot attacks on LUOV. Applied Sciences **10**(12) (2020). https://doi.org/10.3390/app10124106, https://www.mdpi.com/2076-3417/10/12/4106
44. Wang, L.C., Chou, C.Y., Ding, J., Kuan, Y.L., Li, M.S., Tseng, B.S., Tseng, P.E., Wang, C.C.: SNOVA. NIST CSRC (2023)