Improved Cryptanalysis of SNOVA

Ward Beulens

Eurocrypt 2025

04/05/2025

One-slide summary

SNOVA is **Multivariate signature scheme**. It is a **round-2 candidate** in the NIST additional digital signatures project. It has **good key and signature sizes** and **good performance**, at the cost of **a new ad-hoc design** using matrix rings.

We found that **SNOVA implicitly has the structure of a MAYO public map**, although with extra structure which leads to vulnerabilities.

We find **new forgery attacks which are up to a factor 2^{39} more efficient** than existing attacks.

Moreover, we demonstrate **a weak key attack, for one out of every 143k public keys in practice in a few minutes using** modest computational resources***.**

*My laptop, also known from other work.

Multivariate Crypto in 2nd Round of NIST On-Ramp

4 multivariate digital signatures in 2nd round of NIST on-ramp:

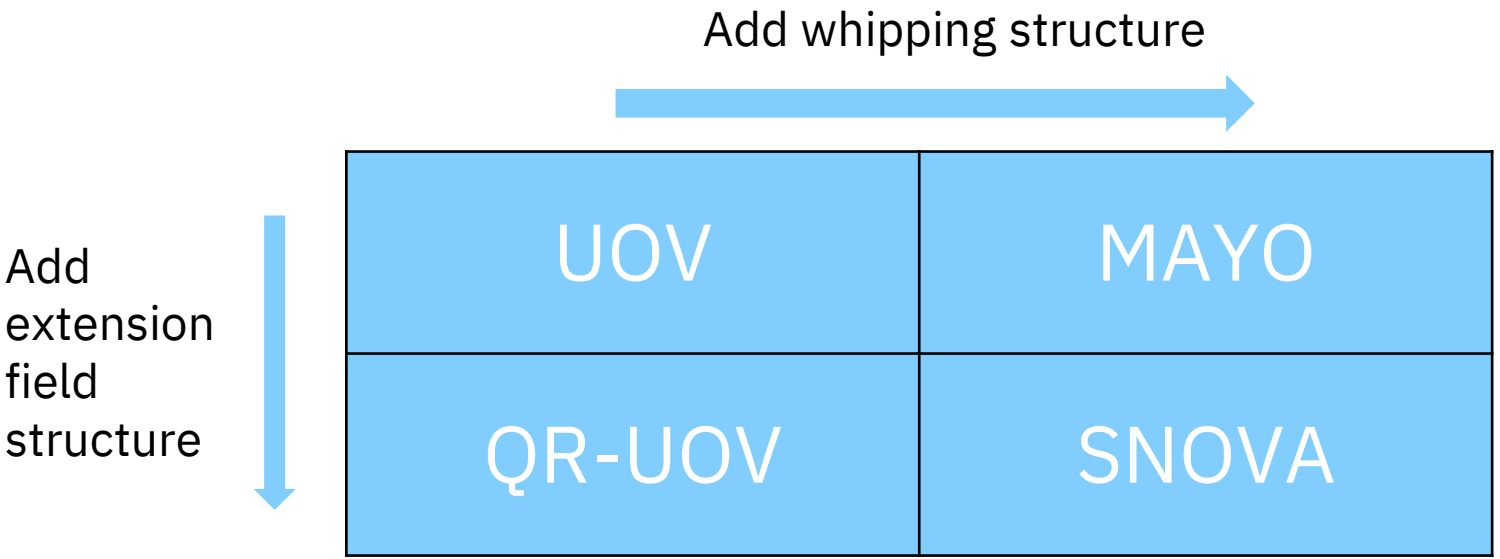

Rest of Talk:

1

UOV and MAYO whipping

2

SNOVA is a structured version of MAYO

3

New Attacks

Oil and Vinegar?

Multivariate quadratic maps as a trapdoor function.

A **multivariate quadratic map** is a sequence of m quadratic polynomials in n variables (over a finite field in MQ crypto). E.g., $m = n = 2$

$$p_1(x, y) = x + 5x^2 + 3xy$$
$$p_2(x, y) = x^2 + 5xy + 5y^2$$

Evaluation is efficient, but **sampling preimages is believed to be hard**, even for quantum computers.

Knowledge of a subspace O of dimension m on which all the polynomials vanish simultaneously (called an oil space in the MQ literature for historic reasons), **allows to sample preimages efficiently**.

This allows for a simple trapdoor-based digital signature scheme, called **Oil and Vinegar**

Parameters (NIST SL 1)

2 constraints:
- Finding oil space O should be hard
- It should be hard to solve $P(x) = y$ without O

Attacks:
Exponential in $n - 2m$
Exponential in m

	UOV-Ip	UOV-Is
# Variables n	112	160
# Equations m	44	64
Finite Field	GF(256)	GF(16)
Pk size	**44 KB**	**67 KB**
Signature size	128 B	96 B

$O(m^3)$
public key size

MAYO in a Nutshell

"Whip" map $P: \mathbb{F}_q^n \to \mathbb{F}_q^m$ with small space O "up" to a larger map $P^\star: \mathbb{F}_q^{kn} \to \mathbb{F}_q^m$, that vanishes on a larger oil space O^k.

$P: \mathbb{F}_q^n \to \mathbb{F}_q^m$

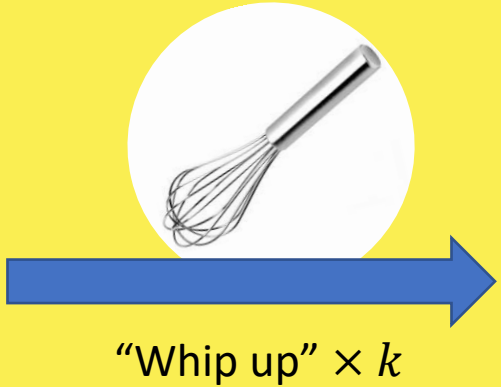

"Whip up" $\times k$

$P^\star: \mathbb{F}_q^{kn} \to \mathbb{F}_q^m$

Whipping Oil-and-Vinegar: Attempt 1

Let $P^{\star}(x_1, \ldots, x_k) = P(x_1) + P(x_2) + \cdots + P(x_k)$.

Then $P^{\star}$: $\mathbb{F}_q^{kn} \to \mathbb{F}_q^m$ vanishes on a large oil space

$$O^k = \{ (o_1, \ldots, o_k) \mid o_1, \ldots, o_k \in O \}$$

So, if $\dim(O^k) = ko \geq m$, then we can sample preimages for $P^{\star}$.

However, this simple approach is not preimage resistant. ☹

Whipping Oil-and-Vinegar

Choose matrices $\boldsymbol{E}_{i,j}$ for all $0 \leq i \leq j \leq k$ and set

$$P^{\star}(x_1, \ldots, x_k) = \sum_i \boldsymbol{E}_{ii} P(x_i) + \sum_{i<j} \boldsymbol{E}_{ij} P'(x_i, x_j)$$

Matrices $\boldsymbol{E}_{ij}$ are chosen such that non-trivial linear combinations of them have full rank, to prevent more efficient forgery attacks!

SNOVA (a.k.a. supernova)

“**S**imple **N**on-Commutative U**OV** with randomness **A**lignment”

Variant of UOV with smaller public keys. E.g., for SL 1:

1016 Byte public key

248 Byte signatures

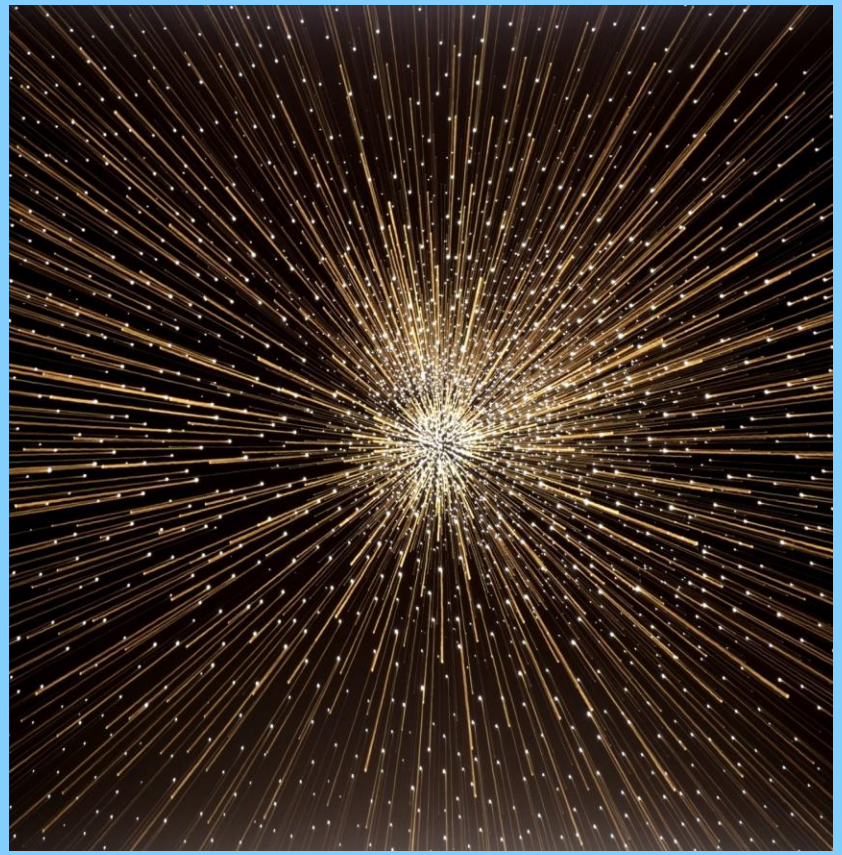

Structure of a SNOVA public map

Public key:

m structured bilinear forms P_i in n variables that vanish on an oil space of dimension o.

Trapdoor function:

mk quadratic polynomials in nk variables vanishing on a space of dimension ok.

$$P_i^* = \sum_{\alpha=1}^{k} A_\alpha \cdot U^T \cdot Q_{\alpha,1}^{\otimes n/k} \cdot P_i \cdot Q_{\alpha,2}^{\otimes n/k} \cdot U \cdot B_\alpha$$

Green = $k \times k$ matrices Generated from seed
Black = Public key
Red = $n \times k$ matrix of variables

SNOVA = MAYO-style whipping of a structured UOV map

Theorem: A SNOVA public map is of the form

$$P^*(\boldsymbol{x_1}, \dots, \boldsymbol{x_k}) = \sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{E_{ij}} \cdot B\left(\boldsymbol{x_i}, \boldsymbol{x_j}\right),$$

where $B(x, y)$ is a bilinear map that vanishes on some oil space, and where $\boldsymbol{E}_{ij}$ are some matrices derived from $A_\alpha, B_\alpha, Q_{\alpha,1}, Q_{\alpha,2}$. Moreover, $\boldsymbol{E}_{ij}$ are block-diagonal matrices, each with n identical $k \times k$ blocks on the diagonal.

This is similar to a MAYO public map, except that B has extra algebraic structure, and that the $\boldsymbol{E}_{ij}$ matrices can have rank-deficient linear combinations.

Attack on SNOVA:

To forge signatures, we want to compute preimages for the SNOVA public map:

$$P^*(\boldsymbol{x_1}, \dots, \boldsymbol{x_k}) = \sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{E}_{ij} \cdot B\big(\boldsymbol{x_i}, \boldsymbol{x_j}\big),$$

Our strategy: Choose $v_1, v_2, \dots, v_k$ at random, and find a solution of the form

$$x_1 = v_1 + \alpha_1 x, \dots, x_k = v_k + \alpha_k x.$$

For a vector x and scalars $\alpha_1, \dots, \alpha_k$. If we plug this into P^* we get

$\sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{E}_{ij} \cdot B\big(\boldsymbol{v_i} + \boldsymbol{\alpha_i x}, \boldsymbol{v_j} + \boldsymbol{\alpha_j x}\big) = (\sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{\alpha_i \alpha_j E}_{ij}) \cdot B(\boldsymbol{x}, \boldsymbol{x})$ + affine function of x.

Step 1:

Find $\alpha_1, \dots, \alpha_k$ such that $\sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{\alpha_i \alpha_j E}_{ij}$ has minimal rank r (exhaustive search is good enough).

Step 2:

Solve for x, which now is a system of r quadratic equations and $m - r$ affine equations.

Cost of the Attack: SL1

Attack is most efficient for the $k = 2$, SL1 parameters:

In this case the $\boldsymbol{E}_{ij}$ matrices have 17 identical 4×4 matrices on the diagonal. So rank is multiple of 17.

Rank	Probability of achieving this rank	Cost of attack
3×17	100%	2^{137}
2×17	$1/500$	2^{97}
1×17	$1/142000$	2^{45} (A few minutes on my laptop)

Countermeasures

Proposal in paper:

Do the whipping properly MAYO-style. I.e., replace the E_{ij} matrices in

$$P^*(\boldsymbol{x_1}, \ldots, \boldsymbol{x_k}) = \sum_{\boldsymbol{0<i,j\leq k}} \boldsymbol{E}_{ij} \cdot B(\boldsymbol{x_i}, \boldsymbol{x_j})$$

By matrices that only have full rank non-trivial linear combinations. This simplifies the construction and completely blocks the attacks.

Extra benefit: Whipping parameter becomes independent of the algebraic structure in the underlying structured UOV instance. Extra flexibility for parameter selection.

Round 2 SNOVA submission:

New, more complex, trapdoor function:

$$\sum_{\alpha=0}^{l^2+l-1} \sum_{j=1}^{n} \sum_{k=1}^{n} A_{i,\alpha} \cdot U_j^t (Q_{i,\alpha,1} P_{i',jk} Q_{i,\alpha,2}) U_k \cdot B_{i,\alpha}$$

This still has the MAYO-whipping structure with E_{ij} matrices that have linear combinations with rank deficiencies, so our attack still applies. However, the E_{ij} don't have a block diagonal structure, so the attack is less powerful.

Estimates of the cost of the attack rely on heuristics about the rank behaviour of linear combinations of $\boldsymbol{E}_{ij}$ matrices.

Conclusions

SNOVA can be seen as a structured version of UOV (not equivalent to QR-UOV) combined with a weak version of MAYO whipping.

We give new forgery attacks against SNOVA, that exploit the weakness of the whipping used in SNOVA.

For some parameter sets the attacks are very serious. The attacks run in a few minutes on a standard laptop and work for a 1/143000 fraction of randomly generated keys.

Attack is mitigated in the 2nd round submission of SNOVA.

Conclusions

SNOVA can be seen as a structured version of UOV (not equivalent to QR-UOV) combined with a weak version of MAYO whipping.

We give new forgery attacks against SNOVA, that exploit the weakness of the whipping used in SNOVA.

For some parameter sets the attacks are very serious. The attacks run in a few minutes on a standard laptop and work for a 1/143000 fraction of randomly generated keys.

Attack is mitigated in the 2nd round submission of SNOVA.

¿ Questions ?