

# Block-Anti-Circulant Unbalanced Oil and Vinegar

Alan Szepieniec<sup>1,2</sup> and Bart Preneel<sup>1</sup>

<sup>1</sup> imec-COSIC KU Leuven, Belgium  
`first-name.last-name@esat.kuleuven.be`

<sup>2</sup> Nervos Foundation  
`alan@nervos.org`

**Abstract.** We introduce a new technique for compressing the public keys of the UOV signature scheme that makes use of block-anti-circulant matrices. These matrices admit a compact representation as for every block, the remaining elements can be inferred from the first row. This space saving translates to the public key, which as a result of this technique can be shrunk by a small integer factor. We propose parameters sets that take into account the most important attacks, and present performance statistics derived from a C implementation along with a comparison to LUOV.

**Keywords:** multivariate quadratic, post-quantum, unbalanced oil and vinegar

## 1 Introduction

Unbalanced Oil and Vinegar (UOV) is one of the longest-standing multivariate quadratic (MQ) signature schemes [10]. While the signatures are rather small, the public keys tend to be huge — they scale with the *cube* of the security parameter. Two notable improvements address this drawback in part.

First, the compression technique due to Petzoldt *et al.* allows most of the public key to be set arbitrarily; the remaining part is then computed with the secret key [14]. Since the arbitrary first part can be the output of a pseudo-random generator, the public key can be compressed to a short seed and the uncompressible second part.

Second, the field lifting technique due to Beullens and Preneel defines the public key over  $\mathbb{F}_2$  but solves the signature equation and produces a signature over an extension of  $\mathbb{F}_2$  [1]. As a result, the direct attack is more complex as it must be performed over a larger field; this allows a smaller number of equations for the same security level. At the same time, however, the public key admits a representation of just one bit for every polynomial coefficient as it was constructed that way.

We propose a third compression technique, relying on structured matrices to compactly represent objects of large size. In particular, the other rows of a circulant or anti-circulant matrix can be inferred from the first. Moreover, these matrices guarantee that  $B^T A B$  is anti-circulant if both  $A$  and  $B$  are, or if  $A$  is

anti-circulant and  $B$  is circulant. This property lends naturally to constructions of MQ public keys, where the matrix representation of the  $i$ th component's quadratic form can be presented as  $S^T F_i S$ . As a result, the public key consists of block-anti-circulant matrices if the matrices of the secret key are block-anti-circulant. It can therefore be represented compactly by the list of first rows of each component block.

The obvious question raised by this design concerns its impact on security. We analyze empirically the complexity of a direct algebraic attack. With respect to the UOV Reconciliation Attack [5], our analysis assumes pessimistically that a successful attack need only consider each block to be its own variable living in the quotient ring  $\mathbb{F}_q[x]/\langle x^\ell - 1 \rangle$ . Building on the insights gleaned from this empiricism and pessimistic analysis, we propose parameters for various security levels. Despite the conservative parameter choices, our compression technique achieves a notable size reduction of the public key and signatures — roughly half at all security levels compared to its immediate predecessor, LUOV.

## 2 Preliminaries

We use pythonic notation to slice submatrices from matrices:  $A_{[i:j,k:l]}$  represents the  $(j-i) \times (l-k)$  block of  $A$  whose upper left element has index  $(i, j)$ , with indices starting as they should at zero. Furthermore we denote by  $0_{[0:v,0:v]}$  the  $v \times v$  zero matrix.

A square matrix  $A$  is *anti-circulant*, and a square matrix  $B$  is *circulant*, if they are fully determined by their first rows  $(a_{\ell-1}, a_{\ell-2}, \dots, a_0)$  and  $(b_0, b_1, \dots, b_{\ell-1})$  via

$$A = \begin{pmatrix} a_{\ell-1} & a_{\ell-2} & \cdots & a_1 & a_0 \\ a_{\ell-2} & a_{\ell-3} & \cdots & a_0 & a_{\ell-1} \\ \vdots & \vdots & & \vdots & \vdots \\ a_1 & a_0 & \cdots & a_3 & a_2 \\ a_0 & a_{\ell-1} & \cdots & a_2 & a_1 \end{pmatrix} \quad \text{and} \quad B = \begin{pmatrix} b_0 & b_1 & \cdots & b_{\ell-2} & b_{\ell-1} \\ b_{\ell-1} & b_0 & \cdots & b_{\ell-3} & b_{\ell-2} \\ \vdots & \vdots & & \vdots & \vdots \\ b_2 & b_3 & \cdots & b_0 & b_1 \\ b_1 & b_2 & \cdots & b_{\ell-1} & b_0 \end{pmatrix}. \quad (1)$$

Circulant matrices are multiplication matrices of elements of the quotient ring  $R[x]/\langle x^\ell - 1 \rangle$ , where  $R$  is the base ring of the matrix. Denote by  $J$  the  $90^\circ$  degree rotation of the identity matrix, *i.e.*, with the ones on the perpendicular diagonal. Then left or right multiplication by  $J$  makes a circulant matrix anti-circulant and vice versa. We make use of the following lemmata.

**Lemma 1.** *Let  $A$  be circulant and  $B$  anti-circulant. Then  $AB$  and  $BA$  are anti-circulant.*

*Proof.* There must be elements  $a, b, b' \in R[x]/\langle x^\ell - 1 \rangle$  with multiplication matrices  $M_a, M_b$  and  $M_{b'}$  such that  $A = M_a$  and  $B = M_b J = J M_{b'}$ . Then  $AB = M_a M_b J = M_{ab} J$  and  $BA = J M_{b'} M_a = J M_{b'_a}$  are anti-circulant.  $\square$

**Lemma 2.** *The sum of circulant matrices is circulant. The sum of anti-circulant matrices is anti-circulant.*



a system of  $m$  equations of the form

$$\mathbf{x}_{[0:v]}^\top \left( F_{[0:v, v:(v+o)]}^{(i)} + F_{[v:(v+o), 0:v]}^{(i)\top} \right) \mathbf{x}_{[v:(v+o)]} = h_i - \mathbf{x}_{[0:v]}^\top F_{[0:v, 0:v]}^{(i)} \mathbf{x}_{[0:v]} \quad , \quad (3)$$

which is linear in the  $o = m$  oil variables  $\mathbf{x}_{[v:(v+o)]}$ . Solving this system completes  $\mathbf{x}$  and from this inverse the user computes the signature  $\mathbf{s} = S^{-1}\mathbf{x}$  straightforwardly.

### 3.2 Petzoldt's Compression Technique


Petzoldt's compression technique [14] rests on the observation that the composition with  $S$  is a *linear* action on the quadratic forms  $F^{(i)}$ . In particular, let  $\overrightarrow{F^{(i)}}$  denote the row-vector of all  $n(n+1)/2$  coefficients in accordance with any standard monomial order; then  $\overrightarrow{P^{(i)}} = \overrightarrow{F^{(i)}}A$  for some matrix  $A \in \mathbb{F}_q^{\frac{n(n+1)}{2} \times \frac{n(n+1)}{2}}$  whose coefficients are given by

$$A_{[\mathbf{mo}(i,j), \mathbf{mo}(r,s)]} = \begin{cases} S_{[r,i]}S_{[s,j]} + S_{[r,j]}S_{[s,i]} & \text{if } i \neq j \\ S_{[r,i]}S_{[s,i]} & \text{otherwise} \end{cases} \quad , \quad (4)$$

where  $\mathbf{mo} : \mathbb{N}^2 \rightarrow \mathbb{N}$  sends the pair  $(i, j)$  to the index of the monomial  $x_i x_j$  in the given monomial order.

As the  $o(o+1)/2$  oil coefficients are zero, the  $\overrightarrow{F^{(i)}}$  must live in a subspace of  $\mathbb{F}_q^{n(n+1)/2}$  of dimension  $n(n+1)/2 - o(o+1)/2$ . As a result, the  $\overrightarrow{P^{(i)}}$  must lie in a subspace of the same dimension. In particular, this means that the first  $v(v+1)/2 + ov$  coefficients of every  $\overrightarrow{P^{(i)}}$  can be set arbitrarily, after which the remaining  $o(o+1)/2$  coefficients are fixed as a function of  $S$ .

The public key, represented as a Macaulay matrix whose rows are  $\overrightarrow{P^{(i)}}$ , is thus divisible into two blocks, of dimensions  $m \times (v(v+1)/2 + vo)$ , and  $m \times o(o+1)/2$ , respectively. The first block can be generated by a pseudorandom generator, after which point the user can find the second only if he knows  $S$ . The public key can therefore be reduced to a short seed and the second block. Note that this size is independent of the number of vinegar variables.



**Fig. 1.** Petzoldt's compression technique.



matrix, or a circulant matrix, respectively. Our compression technique arises from the following observation.

**Theorem 1.** *Let  $A, C$  be block-circulant matrices, and  $B$  be a block-anti-circulant matrix, all with square blocks of height (and width)  $\ell$ . Then  $ABC$  is block-anti-circulant for blocks of the same size.*

*Proof.* The  $\ell \times \ell$  blocks of  $BC$  represent the sum of products of anti-circulant matrices with circulant ones. Via lemmata 1 and 2 one observes that these blocks are circulant. The same argument shows that the  $\ell \times \ell$  blocks of  $A(BC)$  are anti-circulant. The matrix  $ABC$  is thus block-anti-circulant.  $\square$

#### 4.1 Description

Let  $v = V \times \ell$ ,  $o = O \times \ell$  and  $N = O + V$ . We choose  $S$  to be block-circulant; this does not affect the overall shape (5) but does imply that the top right  $V \times O$  block must be block-circulant.

Likewise, the matrices  $F^{(i)}$  are chosen to be  $\ell \times \ell$  block-anti-circulant matrices  $F^{(i)}$  in the shape of (2). One observes via Thm. 1 that the matrices  $P^{(i)}$  are block-anti-circulant as well. These matrices can therefore be represented by only the first row of every block. This requires only  $N^2\ell$  elements per matrix as opposed to the highly redundant  $n^2 = N^2\ell^2$  elements associated with an explicit representation.


Matrices that represent quadratic forms, such as  $F^{(i)}$  and  $P^{(i)}$ , are invariant under addition of skew-symmetric matrices. Over odd-characteristic fields<sup>3</sup> one can therefore always choose  $F^{(i)}$  and  $P^{(i)}$  to be symmetric, even when they are block-anti-circulant (but not necessarily when they are (block-)circulant). This reduces the storage requirement to  $N(N+1)\ell/2$  field elements, down from  $n(n+1)/2$ . For fields of even characteristic, upper-triangular matrix representatives of the quadratic forms are preferred, and in this case the same compression argument applies. However, this means that the  $\ell \times \ell$  blocks on the diagonal must be either identity or zero matrices.

We depart from the Macaulay matrix representation of the public key  $\mathbf{P}$  or of the secret map  $\mathbf{F}$  traditionally used in Petzoldt's compression technique. Instead, both  $\mathbf{P}$  and  $\mathbf{F}$  are represented as lists of symmetric block-anti-circulant matrices. Nevertheless, Petzoldt's compression technique still applies. The pseudorandom generator is used to generate the first row of every  $\ell \times \ell$  block in the upper-triangular part, except for the bottom-most  $O \times (O+1)/2$  blocks which are computed using  $S$ . Figure 2 elaborates.

More explicitly, let  $S = \left( \begin{array}{c|c} I_{[0:v,0:v]} & S' \\ \hline 0_{[0:o,0:v]} & I_{[0:o,0:o]} \end{array} \right)$  for some block-circulant  $v \times o$  matrix  $S'$ . The bottom right  $o \times o$  block of  $P^{(i)}$  is given by

$$P_{[v:n,v:n]}^{(i)} = S'^T F_{[0:v,0:v]}^{(i)} S' + F_{[v:n,0:v]}^{(i)} S' + S'^T F_{[0:v,v:n]}^{(i)} . \quad (6)$$

<sup>3</sup> We restrict focus to odd-characteristic fields because the use of even-characteristic fields induces a security degradation, as shown in Sect. 4.2.



**Fig. 2.** Petzoldt's compression technique with  $\ell \times \ell$  block-anti-circulant matrices.

The nonzero blocks of  $F^{(i)}$  are given by

$$F_{[0:v,0:v]}^{(i)} = P_{[0:v,0:v]}^{(i)} \quad (7)$$

$$F_{[0:v,v:n]}^{(i)} = -P_{[0:v,0:v]}^{(i)} S' + P_{[0:v,v:n]}^{(i)} \quad (8)$$

$$F_{[v:n,0:v]}^{(i)} = -S'^T P_{[0:v,0:v]}^{(i)} + P_{[v:n,0:v]}^{(i)}. \quad (9)$$

Altogether, if Petzoldt's technique is used in conjunction with our block-anti-circulant compression, then the public key is given by  $m\ell O(O+1)/2$  field elements and a short seed.

## 4.2 Security

This section evaluates to which extent the additional structure in the public key facilitates attacks; based on this analysis, we propose parameters later on. The following attacks are considered: Direct Algebraic Attack, Kipnis-Shamir Attack, and UOV Reconciliation Attack.

The Kipnis-Shamir Attack and the UOV Reconciliation Attack can be accelerated by performing arithmetic in the quotient ring  $\mathbb{F}_q[x]/\langle x^\ell - 1 \rangle$ . (We assume, optimistically from the point of view of the attacker, that the overhead of converting between circulant and anti-circulant matrices is negligible.) Arithmetic in the quotient ring can in turn be accelerated using the Chinese Remainder Theorem and the factorization  $\frac{\mathbb{F}_q[x]}{\langle x^\ell - 1 \rangle} \cong \frac{\mathbb{F}_q[x]}{\langle f_0(x) \rangle} \oplus \dots \oplus \frac{\mathbb{F}_q[x]}{\langle f_t(x) \rangle}$ , where  $\prod_{i=0}^t f_i(x) = x^\ell - 1$ . For the purpose of estimating attack complexity, we assume the cost is dominated by arithmetic over the largest component ring in this direct sum, associated with  $f_0$ , the largest-degree<sup>4</sup> irreducible factor of  $x^\ell - 1$ .

Embedded in this assumption is the assertion that if the attack should succeed over a smaller ring, say  $\mathbb{F}_q[x]/\langle f_j(x) \rangle$  with  $\deg(f_j) < \deg(f_0)$ , this success does not help the attacker. Indeed, if successful, such a partial attack outputs the representative of  $S$  in  $\mathbb{F}_q[x]/\langle f_j(x) \rangle$ . However, the attacker needs the matching representative in  $\mathbb{F}_q[x]/\langle f_0(x) \rangle$  for a complete attack, and this component is independent of the previous one.

<sup>4</sup> Or any one of the irreducible factors of largest degree, if there are more than one.

**Direct Attack.** A direct algebraic attack involves deploying Gröbner basis type algorithms [7,6,11,12] in order to solve for  $\mathbf{s} \in \mathbb{F}_q$  the system of multivariate quadratic polynomial equations given by  $(\mathbf{s}^\top P^{(i)} \mathbf{s})_{i=0}^{m-1} = \mathbf{h}$ , where  $\mathbf{h} = \mathbf{H}(d) \in \mathbb{F}_q^m$  is the hash of a target document. The question is whether the introduction of the blockwise anti-circulant structure in order to compress the public key decreases the complexity of such an attack. We implemented the scheme with and without block-anti-circulant compression in Magma in order to test empirically whether this is the case.

In particular, we instantiate two systems of polynomials:

1.  $m$  equations in  $n$  variables without block-anti-circulant compression; this corresponds to  $\ell = 1$ .
2.  $m$  equations in  $n = N \times \ell$  variables with block-anti-circulant compression; this corresponds to  $\ell > 1$ .

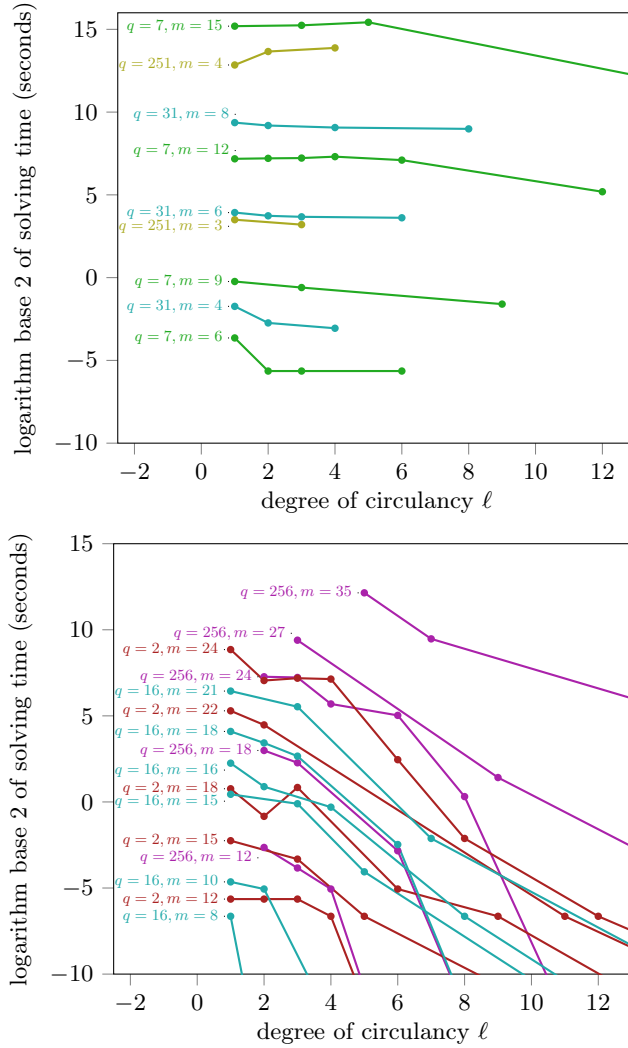
In both cases, the first  $n - m$  variables were assigned random values that still guarantee that a solution exists. Figure 3 shows the running time of these attacks as a function of  $\ell$ , for various values of  $(q, m)$ , as performed by Magma’s implementation of  $F_4$  on an eight core 2.9 GHz machine. The plots suggest that over fields of even characteristic, block-anti-circulant matrices come with a security degradation proportional to the degree of circulancy. In contrast, the security of the same construction but over fields of odd characteristic seems largely unaffected by the degree of circulancy, except possibly at the extremal point where  $\ell = m$ .

Given the correspondence between anti-circulant matrices and the ring  $\frac{\mathbb{F}_q[x]}{\langle x^\ell - 1 \rangle}$ , another natural question is whether arithmetic in this ring can help mount a direct attack. Solutions might be found in each component term of  $\frac{\mathbb{F}_q[x]}{\langle x^\ell - 1 \rangle} \cong \frac{\mathbb{F}_q[x]}{\langle f_0(x) \rangle} \oplus \frac{\mathbb{F}_q[x]}{\langle f_1(x) \rangle} \oplus \dots$  before being joined together using the Chinese Remainder Theorem. However, finding even one such solution still requires solving a system of  $m$  equations in  $N$  variables; and since  $N > m$ , the complexity of this task is already captured by Fig. 3.

**Kipnis-Shamir Attack.** The present proposal is not the first time circulant matrices have been considered in conjunction with UOV. Peng and Tang recently proposed choosing the secret quadratic forms  $F^{(i)}$  to have a specific structure such that during signature generation, the coefficient matrix becomes circulant [13]. This embedded structure not only shrinks the secret key, but it also speeds up signature generation. However, Hashimoto shows that this scheme is vulnerable to a Kipnis-Shamir attack, despite the numbers of vinegar and oil variables being unbalanced [9].

The circulancy in the scheme of Peng and Tang arises as a result of recycling oil-vinegar coefficients across the quadratic forms  $F^{(i)}$ . The algebraic relation that describes this recycling, is exactly the algebraic property that gives rise to the attack. If the  $F^{(i)}$  are chosen independently, the required relation does not hold and the attack fails — or rather, the attack works only with the exponential complexity  $O(q^{v-o})$  of regular unbalanced oil and vinegar.





**Fig. 3.** Running time of direct algebraic attack for odd and even characteristic.

The  $F^{(i)}$  in our construction do have structure, but do not have algebraic properties relating  $F^{(i)}$  for various  $i$ . The coefficient matrix obtained while generating a signature does not have a circulant or block-anti-circulant structure. The attack can be performed over the constituent terms of  $\mathbb{F}_q[x]/\langle x^\ell - 1 \rangle$ , after which the partial solutions are joined together with the Chinese Remainder Theorem. The number  $V$  of vinegar *blocks* must be chosen accordingly, *i.e.*, such that the targeted security level is reached by  $(q^{\deg(f_0)})^{V-O}$ , where  $f_0(x)$  is the largest

degree irreducible factor of  $x^\ell - 1$ . In fact, we consider  $(q^{\deg(f_0)})^{(V-O)/2}$  instead, to account for a speedup on quantum computers due to Grover's algorithm [8].

**UOV Reconciliation Attack.** The UOV Reconciliation Attack [5] is an algebraic key recovery attack that mounts a search for the matrix  $S$  by treating its elements as variables and solving the system of equations obtained by equating  $\left(S^{-1\top} P^{(i)} S^{-1}\right)_{[v:n, v:n]} = 0_{[0:o, 0:o]}$  for all  $i \in \{0, \dots, m-1\}$ . Ding *et al.* argue that the search can be decomposed into a series of steps of which the first dominates the complexity of the entire procedure [5]. This first step requires solving a system of  $m$  quadratic equations in  $v$  variables, originating from the number of polynomials, *i.e.*,  $m$ , and the number of unknowns in the rightmost column of  $S$ , *i.e.*,  $v$ . In the case of UOV where  $v > m$  it is tempting to use a result by Thomae and Wolf showing how to reduce solving a system of  $m$  quadratic equations in  $n = \alpha m$  variables to solving one of  $m - \lfloor \alpha \rfloor + 1$  equations in as many variables [15]. However, Beullens and Preneel argue that this reduction does not apply to this first step of the UOV Reconciliation Attack because it finds an arbitrary solution and not necessarily one that is consistent with the other steps [1]. Instead, Beullens and Preneel estimate the complexity of this attack as strictly larger than that of solving a system of  $v$  equations in  $v$  variables.

With respect to our construction, an attack performed over the quotient ring  $\frac{\mathbb{F}_q[x]}{\langle x^\ell - 1 \rangle} = \frac{\mathbb{F}_q[x]}{\langle f_0(x) \rangle} \oplus \dots$  suffices to break the scheme. In this case the attack represents a search for the  $V \times O$  unknown ring elements of the matrix  $S$ . In particular, the last column of  $S$  has only  $V = v/\ell$  unknowns. However, the number  $m$  of equations remains unaffected by this ring switch. Therefore, as long as  $V \geq m$ , we can argue that the complexity of the Reconciliation Attack is lower-bounded by solving a system of  $V$  equations in  $V$  variables over  $\mathbb{F}_q[x]/\langle f_0(x) \rangle$ .

### 4.3 Parameters and Comparison

We advise against using fields of even characteristic in light of the poor resilience of our block-anti-circulant compression against direct algebraic attacks, as shown in Fig. 3. However, we note that using odd characteristic fields does not preclude using the field lifting technique of Beullens and Preneel, although it does make it less effective. Denote by  $r$  the extension degree, *i.e.*, the signature equation is defined over  $\mathbb{F}_{q^r}$  instead of  $\mathbb{F}_q$ .

We estimate the complexity of algebraic system solving using the Wiedemann method [12] along with Groverized fixing of variables [3,1]. This makes for a complexity of

$$C_{m,n,k} = O\left(q^{k/2} \cdot \binom{n-k+2}{2} \binom{d_{reg}(k) + n - k}{n-k}\right)^2, \quad (10)$$

where  $k$  is the number of variables that are quantumly guessed, and the degree of regularity  $d_{reg}$  is given by the degree of the first non-positive term in the formal

power series expansion of

$$HS(z) = \frac{(1 - z^2)^m}{1 - z^n} . \quad (11)$$

To obtain one concrete number, we take the minimum of  $C_{m,n,k}$  over all  $k$  and pretend as though the constant hidden by the Landau notation is equal to 1.

Table 1 presents a selection of parameter sets designed to meet various target levels of post-quantum security, measured in terms of the base 2 logarithm of the best attack’s complexity. For convenience, it also offers comparisons with variants of UOV, namely:

- LUOV — UOV with Petzoldt’s compression technique and field lifting [1].
- PCT — UOV with Petzoldt’s compression technique [14].
- Plain — Plain UOV with no compression [10].

**Table 1.** Proposed parameter sets and comparison to other variants of UOV.

scheme	parameters	$ pk $	$ sig $	sec. lvl.
Plain	$q = 256, v = 106, m = o = 53$	658.36 kB	159 bytes	128.85
PCT	$q = 256, v = 106, m = o = 53$	74.07 kB	159 bytes	128.85
LUOV	$q = 2, v = 296, m = o = 40, r = 68$	4.00 kB	2.79 kB	128.17
<b>BACUOV</b>	$q = 3, V = 49, O = 7, \ell = 7, r = 12$	2.34 kB	1.14 kB	129.32
Plain	$q = 256, v = 164, m = o = 82$	2.38 MB	246 bytes	191.89
PCT	$q = 256, v = 164, m = o = 82$	272.5 kB	246 bytes	191.89
LUOV	$q = 2, v = 444, m = o = 60, r = 84$	13.40 kB	5.16 kB	190.00
<b>BACUOV</b>	$q = 3, V = 76, O = 10, \ell = 7, r = 18$	6.58 kB	2.65 kB	192.08
Plain	$q = 256, v = 224, m = o = 112$	6.05 MB	336 bytes	256.50
PCT	$q = 256, v = 224, m = o = 112$	692.13 kB	336 bytes	256.50
LUOV	$q = 2, v = 600, m = o = 82, r = 90$	34.06 kB	7.49 kB	256.13
<b>BACUOV</b>	$q = 3, V = 104, O = 14, \ell = 7, r = 11$	17.59 kB	2.22 kB	256.68

Note that the choice  $q = 3$ , which minimizes the total size of public key and signature, is not represented in Fig. 3. In fact, this choice has a poor resilience against algebraic attack — its complexity decreases with increasing circulancy, albeit much slower than when  $q$  is even. Nevertheless, we argue that this subtle degradation is an artifact of the small coefficient field over which the system of equations is defined. In particular, extending this field by setting  $r > 1$  reduces the degradation or even halts it completely. Figure 4 shows a similar plot except for  $q = 3$  and various  $r$ . There is much less degradation when  $r = 3$  and it seems to vanish entirely for  $r = 5$ , which incidentally understates the recommended parameters by a large factor. A complete argument would run the toy attack for  $q = 3, r = 5$  and greater degrees of circulancy, but sadly this experiment is impossible with the available hardware and time.

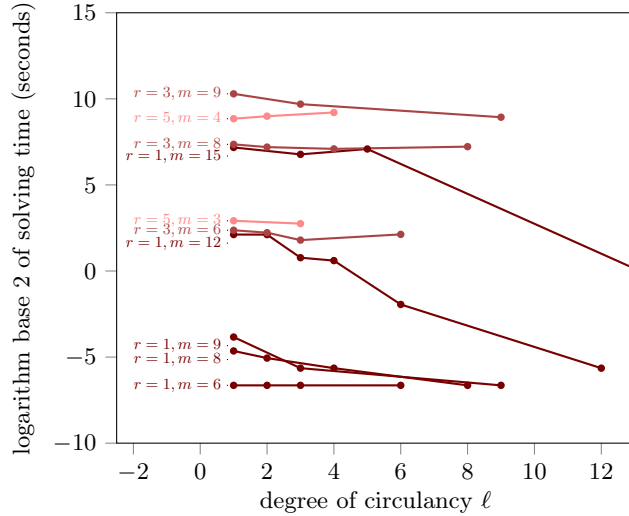


Fig. 4. Running time of direct algebraic attack for various  $r$  when  $q = 3$ .

#### 4.4 Implementation

A full working proof of concept implementation was developed in Sage, and a mildly optimized version in C for the purpose of comparison, see below. The direct attack timings were obtained from a Magma implementation that only generates block-anti-cyclic public keys but does not do compression of any kind. The security levels are estimated using a Sage script. All source code is available under the Community Research and Academic Programming License (CRAPL) from github: <https://github.com/aszepieniec/bacuov>.

For the purpose of an apples-to-apples comparison, we use the reference implementation of the round 2 NIST candidate LUOV [2] with the recommended parameter sets for a balanced public key and signature size. On the part of the block-anti-circulant scheme, we adapt the parameters to minimize the combined size of the public key and signature, subject to meeting the same security level as their LUOV counterpart. This implementation uses arithmetic over  $\mathbb{F}_q/\langle x^\ell - q \rangle$  as well as delayed modular reduction for various matrix operations, but exploits no parallelism. The performance numbers are given by the kilocycles (kc, first line), and milliseconds (ms, second line) in Table 2. These numbers are the average of 100 executions run on a 4-core Intel(R) Core(TM) i3-7100U CPU @ 2.40GHz with 3072 kB cache.

## 5 Conclusion

We propose to introduce a block-anti-circulant structure into the secret and private keys of the UOV signature scheme. While the addition of structure may

**Table 2.** Performance comparison between LUOV and block-anti-circulant UOV.

scheme	parameters	$ pk $	$ sig $	sec. lvl.	Keygen	Sign	Verify
LUOV	$q = 2, r = 48$ $m = 43, v = 222$	5 kB	1.57 kB	NIST II	14 795 kc 6.165 ms	54 831 kc 22.846 ms	35 748 kc 14.895 ms
<b>BACUOV</b>	$q = 3, V = 56$ $O = 8, \ell = 7, r = 12$	3.45 kB	1.31 kB	NIST II	682 436 kc 284.323 ms	332 498 kc 138.527 ms	593 693 kc 247.355 ms
<b>BACUOV</b>	$q = 7, V = 60$ $O = 5, \ell = 13, r = 5$	4.64 kB	1.55 kB	NIST II	2 685 130 kc 1118.648 ms	611 928 kc 254.956 ms	1 265 715 kc 527.364 ms
LUOV	$q = 2, r = 64$ $m = 61, v = 302$	14.1 kB	2.84 kB	NIST IV	40 039 kc 16.683 ms	163 638 kc 68.182 ms	90 331 kc 37.638 ms
<b>BACUOV</b>	$q = 3, V = 84$ $O = 11, \ell = 7, r = 16$	8.69 kB	2.60 kB	NIST IV	2 354 156 kc 980.803 ms	1 402 365 kc 584.275 ms	2 452 899 kc 1021.961 ms
<b>BACUOV</b>	$q = 7, V = 76$ $O = 7, \ell = 11, r = 16$	8.70 kB	5.35 kB	NIST IV	4 801 991 kc 2000.665 ms	2 260 414 kc 941.775 ms	3 740 856 kc 1558.526 ms
LUOV	$q = 2, r = 80$ $m = 76, v = 363$	27.1 kB	4.29 kB	NIST V	176 100 kc 73.374 ms	527 341 kc 219.723 ms	248 874 kc 103.696 ms
<b>BACUOV</b>	$q = 3, V = 104$ $O = 14, \ell = 7, r = 16$	17.60 kB	2.42 kB	NIST V	5 096 796 kc 2123.658 ms	2 804 193 kc 1168.403 ms	4 758 999 kc 1982.899 ms
<b>BACUOV</b>	$q = 7, V = 97$ $O = 9, \ell = 11, r = 8$	17.95 kB	3.41 kB	NIST V	12 019 482 kc 5007.974 ms	2 399 406 kc 999.722 ms	5 401 021 kc 2250.364 ms

accelerate some attacks, we argue that it is possible to either offset this acceleration or block it entirely by choosing parameters appropriately. The resulting public key is smaller than the variant of UOV that uses only Petzoldt’s compression trick by a factor  $\ell$  which determines the block size. For typical values of this parameter, *i.e.* between 7 and 13, the resulting public keys are tens of kilobytes in size for all security levels.

With respect to the certificate metric  $|pk| + |sig|$  (*i.e.*, the size of a link in a chain of signatures and public keys in a certificate), our scheme represents a small improvement over LUOV. As a result of this small improvement our scheme achieves the smallest combined size of public key and signature across all MQ signature schemes. While the size difference with respect to LUOV is marginal at the 128 bit security level, this difference increases noticeably for higher security levels and thus provides empirical evidence of the improved scaling behavior promised by the insertion of an anti-circulant structure.

The comparison between UOV with block-anti-circulant structure and LUOV shows that the bandwidth improvement comes at a significant performance penalty. In some cases, the block-anti-circulant algorithms are up to  $70\times$  slower than their LUOV counterparts. Nevertheless, it should be noted that the algorithms stand to benefit from instruction-level parallelism, which the current implementation does not employ. In contrast, bitwise parallelism is native to the fields over which LUOV operates. We expect aggressive optimization to shrink this performance penalty significantly, but ultimately leave this task to future work. Regardless, the smaller bandwidth requirement may justify the computational overhead depending on the context. The present construction provides the

protocol designer with a greater flexibility in choosing parameters, thus enabling him to better finetune the cryptosystem to the constraints of his problem.

Performance is not the only penalty associated with introducing a block-anti-circulant structure. Indeed, the security argument hinges on two new assumptions. First, a Kipnis-Shamir or UOV Reconciliation attack that exploits the block-anti-circulant structure is dominated by the cost of arithmetic in the largest ring in the decomposition  $\frac{\mathbb{F}_q[x]}{\langle x^\ell - 1 \rangle} \cong \frac{\mathbb{F}_q[x]}{\langle f_0(x) \rangle} \oplus \frac{\mathbb{F}_q[x]}{\langle f_1(x) \rangle} \oplus \dots$ . Second, the block-anti-circulant structure does not speed up a direct algebraic attack for large enough fields of odd order. These two new assumptions are in addition to the assumption introduced by LUOV, namely that defining the public key over a subfield does not speed up a direct algebraic attack. We invite the community to help scrutinize these assumptions.

We close by posing two questions. The first is prompted by the observation that the inserted structure by which public key compression is achieved, is highly specific. We used block-anti-circulant structure because it is straightforward and simultaneously compatible with both the construction of  $P_{[v:n,v:n]}^{(i)}$  from  $P_{[0:v,0:v]}^{(i)}$  and  $P_{[0:v,v:n]}^{(i)}$ , and with the canonical representation of quadratic forms as symmetric matrices. Nevertheless, it might be possible that an alternative to circulant and anti-circulant matrices is also compatible with the necessary arithmetic, possibly at the expense of a less straightforward instantiation. For instance, instead of using the matrices of multiplication of polynomials modulo  $x^\ell - 1$ , one might opt for the same matrices of multiplication but modulo an irreducible polynomial. The advantage of this alternative structure would be the impossibility of decomposing the resulting ring into smaller components. However, the question remains whether this alternative algebra is compatible with the construction of  $P_{[v:n,v:n]}^{(i)}$  from  $P_{[0:v,0:v]}^{(i)}$  and  $P_{[0:v,v:n]}^{(i)}$ , and with the symmetric matrix representation of quadratic forms — or if it is not, which compromises can still confer a net benefit.

Lastly, an interesting question is raised by our empirical results: why is there a significant security degradation associated with a larger degree of circulantcy specifically for fields of characteristic two? We conjecture that this degradation is related to the impossibility of representing quadratic forms over an even characteristic field by symmetric matrices. As a result, a block-anti-circulant representation of such a quadratic form necessarily contains blocks of zeros on its diagonal, thus greatly reducing the number of nonzero coefficients.

**Acknowledgements.** This work was supported in part by the Research Council KU Leuven: C16/15/058. In addition, this work was supported by the European Commission through the Horizon 2020 research and innovation programme under grant agreement H2020-DS-LEIT-2017-780108 FENTEC, by the Flemish Government through FWO SBO project SNIPPET S007619N and by the IF/C1 on Cryptanalysis of post-quantum cryptography. Alan Szepieniec was supported by a doctoral grant from Flemish Agency for Innovation and Entrepreneurship (VLAIO, formerly IWT) and is supported by Nervos Foundation. Lastly, the authors would like to thank Ward Beullens for useful feedback.

## References

1. Beullens, W., Preneel, B.: Field lifting for smaller UOV public keys. In: Patra, A., Smart, N.P. (eds.) INDOCRYPT 2017. LNCS, vol. 10698, pp. 227–246. Springer (2017)
2. Beullens, W., Preneel, B., Szepieniec, A., Vercauteren, F.: LUOV Signature Scheme proposal for NIST PQC Project (Round 2 version) <https://github.com/WardBeullens/LUOV>
3. Chen, M., Hülsing, A., Rijneveld, J., Samardjiska, S., Schwabe, P.: From 5-pass  $MQ$ -based identification to  $MQ$ -based signatures. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 135–165 (2016)
4. Czypek, P., Heyse, S., Thomae, E.: Efficient implementations of MQPKS on constrained devices. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 374–389. Springer (2012)
5. Ding, J., Yang, B., Chen, C.O., Chen, M., Cheng, C.: New differential-algebraic attacks and reparametrization of rainbow. In: Bellovin, S.M., Gennaro, R., Keromytis, A.D., Yung, M. (eds.) ACNS 2008. LNCS, vol. 5037, pp. 242–257 (2008)
6. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases without reduction to zero ( $F_5$ ). In: ISSAC 2002. pp. 75–83. ACM (2002)
7. Faugère, J.C.: A new efficient algorithm for computing Gröbner bases ( $F_4$ ). Journal of Pure and Applied Algebra 139(1-3), 61–88 (1999)
8. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Miller, G.L. (ed.) ACM STOC 1996. pp. 212–219. ACM (1996)
9. Hashimoto, Y.: On the security of Circulant UOV/Rainbow. IACR Cryptology ePrint Archive 2018, 947 (2018), <https://eprint.iacr.org/2018/947>
10. Kipnis, A., Patarin, J., Goubin, L.: Unbalanced oil and vinegar signature schemes. In: Stern, J. (ed.) EUROCRYPT '99. Lecture Notes in Computer Science, vol. 1592, pp. 206–222. Springer (1999)
11. Mohamed, M.S.E., Cabarcas, D., Ding, J., Buchmann, J.A., Bulygin, S.:  $MXL_3$ : An efficient algorithm for computing Gröbner bases of zero-dimensional ideals. In: Lee, D.H., Hong, S. (eds.) ICISC 2009. LNCS, vol. 5984, pp. 87–100. Springer (2009)
12. Mohamed, W.S.A., Ding, J., Kleinjung, T., Bulygin, S., Buchmann, J.: PWXL: A parallel Wiedemann-XL algorithm for solving polynomial equations over  $GF(2)$ . In: Cid, C., Faugère, J. (eds.) Conference on Symbolic Computation and Cryptography. pp. 89–100 (2010)
13. Peng, Z., Tang, S.: Circulant UOV: a new UOV variant with shorter private key and faster signature generation. TIS 12(3), 1376–1395 (2018)
14. Petzoldt, A., Buchmann, J.A.: A multivariate signature scheme with an almost cyclic public key. IACR Cryptology ePrint Archive 2009, 440 (2009), <http://eprint.iacr.org/2009/440>
15. Thomae, E., Wolf, C.: Solving underdetermined systems of multivariate quadratic equations revisited. In: Fischlin, M., Buchmann, J.A., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 156–171. Springer (2012)