Cryptanalysis of Two Candidate Fixes of Multilinear Maps over the Integers

Jean-Sébastien Coron[1], Tancrède Lepoint[2], and Mehdi Tibouchi[3]

[1] University of Luxembourg
jean-sebastien.coron@uni.lu
[2] CryptoExperts, France
tancrede.lepoint@cryptoexperts.com
[3] NTT Secure Platform Laboratories, Japan
tibouchi.mehdi@lab.ntt.co.jp

November 30, 2014

Abstract. Shortly following Cheon, Han, Lee, Ryu and Stehlé's attack against the multilinear map of Coron, Lepoint and Tibouchi (CLT), two independent approaches to thwart this attack have been proposed on the cryptology ePrint archive, due to Garg, Gentry, Halevi and Zhandry on the one hand, and Boneh, Wu and Zimmerman on the other. In this short note, we show that both countermeasures can be defeated in polynomial time using extensions of the Cheon *et al.* attack.

1 Introduction

Soon after Garg, Gentry and Halevi proposed a first candidate cryptographic multilinear map in [GGH13], multilinear maps and their applications became a very active area of research in the cryptologic community. However, all current candidate constructions [GGH13,CLT13,GGH14] are only approximate multilinear maps. In particular, they require a trusted setup to create public parameters from secret values; whoever learns these values can break all security notions related to multilinear maps.

Recently, Cheon, Han, Lee, Ryu and Stehlé described an attack using low-level encodings of zero [CHL+14] against Coron, Lepoint and Tibouchi's candidate multilinear map over the integers (CLT) [CLT13]. This attack recovers the secret factors of the public modulus in polynomial time, which breaks the construction completely.

A week after Cheon *et al.*'s attack was made public, two different approaches to patch the CLT multilinear map have been independently proposed by Garg, Gentry, Halevi and Zhandry [GGHZ14, Sec. 7][1], and Boneh, Wu and Zimmerman [BWZ14]. In essence, these countermeasures modify the form of the CLT encodings in an attempt to remove the multiplicative structure obtained during the zero-testing procedure in Cheon *el al.*'s attack.

Our Results. In this report, we show that both proposed countermeasures are insecure. They are susceptible to direct extensions of Cheon *et al.* attack, that will still recover the factorization of the public modulus (and hence all secret parameters) in polynomial time. Proof-of-concept implementations of our attacks are available at `https://github.com/coron/cltattack`.

2 The CLT Multilinear Map Scheme

In this section, we recall the multilinear map scheme (CLT) of Coron, Lepoint and Tibouchi [CLT13]. From a given security level λ and a target multilinearity level κ, a trusted setup phase outputs the public parameters from secret values. In particular, it generates n secret primes $p_1, \ldots, p_n$ and publishes the public modulus $x_0 = \prod_{i=1}^{n} p_i$ (where n is large enough to ensure correctness and

[1] The revised version of [GGHZ14] of November 12 2014 can be accessed from the cryptology ePrint archive.

security), and generates a random invertible $z \in \mathbb{Z}_{x_0}$ multiplicative mask. The (possibly secret) encoding space is a ring $R = \mathbb{Z}_{g_1} \times \cdots \times \mathbb{Z}_{g_n}$ for primes g_i's.

A level-k encoding c of some $m = (m_1, \ldots, m_n) \in R$ is an integer c such that, for all $1 \leqslant i \leqslant n$:

$$c \equiv \frac{r_i \cdot g_i + m_i}{z^k} \pmod{p_i} \tag{1}$$

where r_i is uniformly distributed from a bounded distribution with at least λ bits of entropy. The integer c is therefore defined modulo x_0 by CRT. Encodings can be homomorphically processed as long as the noises r_i's do not wrap modulo p_i: if c is a level-k encoding of $m \in R$ and c' is a level-k' encoding of $m' \in R$, then $c \cdot c' \bmod x_0$ is a level-$(k + k')$ encoding of $m \cdot m'$, and when $k = k'$, $c + c' \bmod x_0$ is a level-k encoding of $m + m'$.

For level-κ encodings, the trusted setup publishes a zero-testing parameter[2] p_{zt} such that

$$p_{zt} = \sum_{i=1}^{n} h_i \cdot [g_i^{-1} \cdot z^\kappa \bmod p_i] \cdot (x_0/p_i) \bmod x_0 \,,$$

where the h_i's are uniformly distributed from a bounded distribution with at least λ bits of entropy.

Given a level-κ encoding c as in Equation (1), one can compute $\omega = p_{zt} \cdot c \bmod x_0$, which gives:

$$\omega = \sum_{i=1}^{n} h_i \cdot \big(r_i + m_i \cdot (g_i^{-1} \bmod p_i)\big) \cdot (x_0/p_i) \bmod x_0 \,.$$

The multilinear map parameters are chosen so that if $m_i = 0$ for all i, using the bounds on the r_i's and h_i's, the integer ω is small compared to x_0; respectively if at least one of the $m_i \neq 0$, w will roughly be of the same size as x_0 due to the uncanceled g_i^{-1} in the expression above. This enables to test whether c is an encoding of $0 \in R$ or not.

Since the leading bits of ω only depend on the m_i's and not on the noises r_i's, this enables to extract from level-κ encodings a function of the m_i's only, which eventually defines a degree-κ multilinear map.

Public Sampling and Rerandomization. In order to allow users to publicly create encodings, the CLT public parameters contain ℓ level-0 encodings $x_1, \ldots, x_\ell$ of random ring elements (as defined in Equation (1)). The public sampling procedure then amounts in a subset-sum of these x_i's, and by the leftover-hash lemma one obtains a level-0 encoding of a nearly uniform $m \in R$.

Now, level-0 encodings can be turned in level-k encodings for any $k \leqslant \kappa$ using (powers of) a public level-1 encoding y of $1 = (1, \ldots, 1) \in R$. Without loss of generality and for simplicity, we assume $k = 1$ as in [CLT13]. To avoid a naive division after multiplying by y, the public parameters contains $\tau \geqslant n$ level-1 encodings $x'_1, \ldots, x'_\tau$ of $0 \in R$. Following a multiplication by y, the resulting encoding is rerandomized by a subset-sum of the x'_i's: the output of the rerandomization procedure is nearly independent from the randomness of the input by a left-over hash lemma over lattices [CLT13, Sec. 4].

3 The Cheon *et al.* Attack

Recently, Cheon *et al.* proposed an attack against the CLT scheme [CHL+14]. This attack makes use of low-level encodings of 0: if such encodings are made public, one can recover in polynomial time all values supposed to be kept secret. In particular, one can use the values x'_i's used by the rerandomization procedure. Therefore the Cheon *et al.* attack completely breaks the CLT scheme.

[2] In [CLT13], it publishes a vector of such elements. Since Cheon *et al.* attack, and ours, do not require more than one element (*e.g.* the first element of the vector), we omit such setting for simplicity.

Let us recall the attack briefly. We describe a slight simplification of [CHL+14] in which we use a single ciphertext c instead of two ciphertexts c_0 and c_1. This enables to obtain as eigenvalues directly the CRT components of c, instead of ratios of the CRT components of c_0 and c_1.

Let c be a level-0 encoding with $c = c_i \bmod p_i$. Let y be a level-1 encoding of $1 \in R$, let x'_j be level-1 encodings of $0 \in R$ with $x'_j = r'_{ij} \cdot g_i/z \bmod p_i$, and x_j be level-1 encodings where $x_j = x_{ij}/z \bmod p_i$.

For $1 \leqslant j, k \leqslant n$, we can compute:

$$\omega_{jk} = [(c \cdot x_j \cdot x'_k \cdot y^{\kappa-2}) \cdot p_{zt}]_{x_0} \tag{2}$$

and we have:

$$\begin{aligned}\omega_{jk} &= \sum_{i=1}^{n} h_i \cdot [(c \cdot x_j \cdot x'_k \cdot y^{\kappa-2}) \cdot z^{\kappa} \cdot g_i^{-1} \bmod p_i] \cdot (x_0/p_i) \\ &= \sum_{i=1}^{n} x_{ij} h'_i c_i r'_{ik} \bmod x_0 \end{aligned} \tag{3}$$

where $h'_i = h_i \cdot [y^{\kappa-2} \bmod p_i] \cdot (x_0/p_i)$. Equation (3) actually holds over the integers (instead of only modulo x_0), because the previous encoding is an encoding of 0, and therefore ω_{jk} is smaller than x_0. Therefore we can write:

$$\omega_{jk} = \sum_{i=1}^{n} x_{ij} h'_i c_i r'_{ik}$$

over the integers. We note that ω_{jk} is a quadratic form in the x_{ij}'s and the r'_{ik}'s. By spanning $1 \leqslant j, k \leqslant n$, one can construct a matrix $\boldsymbol{W}_c = (\omega_{jk})_{1 \leqslant j,k \leqslant n}$ such that

$$\boldsymbol{W}_c = \boldsymbol{X} \times \boldsymbol{C} \times \boldsymbol{R},$$

where $\boldsymbol{X} = (x_{ij} \cdot h'_i)_{1 \leqslant j,i \leqslant n}$ and $\boldsymbol{R} = (r'_{ik})_{1 \leqslant i,k \leqslant n}$ and $\boldsymbol{C} = \begin{pmatrix} c_1 & & & \\ & c_2 & & \\ & & \ddots & \\ & & & c_n \end{pmatrix}$.

Finally, one can publicly compute:

$$\boldsymbol{W} = \boldsymbol{W}_c \cdot \boldsymbol{W}_I^{-1} = \boldsymbol{X} \times \boldsymbol{C} \times \boldsymbol{X}^{-1},$$

where $I = 1$, which is a level-0 encoding of 1; this means that for W_I we take $c = 1$ in Equation (2). Since $\boldsymbol{C}$ is a diagonal matrix, by computing the eigenvalues of $\boldsymbol{W}$ one can recover the c_i's, and then the p_i's. Finally, Cheon *et al.* describe how to recover all the other secret values in [CHL+14].

Remark 1. One can also use the following optimization: perform all computations modulo a known prime q instead of over $\mathbb{Q}$. It suffices to choose q to be slightly larger than the c_i's, so that these components can be recovered from their value modulo q. Therefore we can perform all computations modulo a prime q of size $\mathcal{O}(\log_2 \max |c_i|)$ bits, instead of integers of size $\mathcal{O}(\log_2 x_0)$ bits.

4 Cryptanalysis of the Garg *et al.* Countermeasure

A transformation of CLT multilinear maps is described by Garg, Gentry, Halevi and Zhandry in [GGHZ14] in order to resist the Cheon *et al.* attack. The technique consists in embedding a CLT encoding into a matrix of encodings; the goal is to eliminate the native encodings of zero that enables the Cheon *et al.* attack. In this section we show that the countermeasure is insecure; namely we show an extension of the Cheon *et al.* attack that can recover all secret parameters in polynomial time, as in the original attack.

The Garg *et al.* countermeasure. Let c be a native level-i CLT encoding of some $m \in R$; then the level-i matrix encoding of the same $m \in R$ is a $2\kappa+1$ matrix $\boldsymbol{U}$ of the form:

$$\boldsymbol{U} = \left[\boldsymbol{T} \times \begin{bmatrix} \$ & 0 & \dots & 0 \\ 0 & \$ & \dots & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \dots & c \end{bmatrix} \times \boldsymbol{T}^{-1} \right]_{x_0}$$

where the '$'s represent native CLT encodings of random elements at level i, the 0's are native encodings of zero at level i, and $\boldsymbol{T}$ is a random $(2\kappa+1) \times (2\kappa+1)$ matrix modulo x_0, the same for all encodings. To allow for rerandomization of encodings, one publishes matrix encodings of 0. With enough such matrix encodings of 0, one can rerandomize $\boldsymbol{U}$, as in the original CLT scheme.

Zero-testing is done as follows. The CLT single-element zero-test parameter p_{zt} is replaced by two vectors $q_{zt} = (\boldsymbol{s}, \boldsymbol{t})$ defined as follows:

$$\boldsymbol{s} = [(\$ \dots \$\ 0 \dots 0\ \$) \times \boldsymbol{T}^{-1}]_{x_0} \quad \text{and} \quad \boldsymbol{t} = [\boldsymbol{T} \times (0 \dots 0\ \$ \dots \$\ \$)^T \times p_{zt}]_{x_0}$$

where 0 and '$' are level-0 native CLT encodings of zero and random elements. Given a matrix $\boldsymbol{U}$ as above at level κ, one computes:

$$\omega = \boldsymbol{s} \times \boldsymbol{U} \times \boldsymbol{t} \bmod x_0 \,.$$

This gives:

$$\omega = (\$ \times c + 0) \cdot p_{zt} \bmod x_0$$

where '$' is a level-0 native CLT encoding, and 0 is a level-κ native CLT encoding of zero. Since $(\$ \times c + 0)$ is a CLT encoding of 0 when c is (and whp is not when c is not), we get a zero-testing procedure for c, that is ω is small compared to x_0 if $\boldsymbol{U}$ is an encoding of zero (and whp is not when $\boldsymbol{U}$ is not).

Our Attack. We consider a matrix encoding $\boldsymbol{C}$ of zero at level 0, and we write:

$$\boldsymbol{C} = \left[\boldsymbol{T} \times \begin{bmatrix} \$ & 0 & \dots & 0 \\ 0 & \$ & \dots & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \dots & c \end{bmatrix} \times \boldsymbol{T}^{-1} \right]_{x_0} = [\boldsymbol{T} \times \boldsymbol{C}^* \times \boldsymbol{T}^{-1}]_{x_0} \,. \tag{4}$$

We consider two other matrices $\boldsymbol{X}$ and $\boldsymbol{R}$ which are encodings of 0 at level $\kappa - 1$ and 1 respectively, so that we can compute:

$$\omega = \boldsymbol{s} \times \boldsymbol{X} \times \boldsymbol{C} \times \boldsymbol{R} \times \boldsymbol{t} \bmod x_0 \,.$$

Writing $\boldsymbol{x} = \boldsymbol{s} \times \boldsymbol{X} \times \boldsymbol{T}$ and $\boldsymbol{r} = \boldsymbol{T}^{-1} \times \boldsymbol{R} \times \boldsymbol{t}$, we get:

$$\omega = \boldsymbol{x} \times \boldsymbol{C}^* \times \boldsymbol{r} \bmod x_0 \,.$$

Since $\boldsymbol{C}$ is an encoding of 0, we have that the integer ω is small compared to x_0.

We note that ω is a quadratic form in $\boldsymbol{x}$ and $\boldsymbol{r}$. Moreover, we know that in the original CLT scheme, ω is a linear form in the n CRT components c_i of a CLT encoding c at level κ, with $c = c_i / z^\kappa \bmod p_i$. We can therefore expand the previous vectors and matrices from dimension $2\kappa+1$ to dimension $(2\kappa+1) \cdot n$, and write:

$$\omega = \hat{\boldsymbol{x}} \times \hat{\boldsymbol{C}}^* \times \hat{\boldsymbol{r}}$$

where the $(i \cdot (2\kappa+1) + j)$-th coefficient of $\hat{\boldsymbol{x}}$ is $\boldsymbol{x}_j \bmod p_i$; similarly $\hat{\boldsymbol{C}}^*$ is a square matrix of dimension $(2\kappa+1) \cdot n$ which is block-diagonal, with the n sub-matrices $\boldsymbol{C}^* \bmod p_i$ on the diagonal.

Now by applying the Cheon *et al.* attack, we can recover the characteristic polynomial of $\hat{\boldsymbol{C}}^*$ over $\mathbb{Z}$. Namely as in the Cheon *et al.* attack, instead of using single vectors $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{r}}$, we can use $(2\kappa+1)\cdot n$ such vectors, so that we obtain a matrix:

$$W_{\boldsymbol{C}} = \hat{\boldsymbol{X}} \times \hat{\boldsymbol{C}}^* \times \hat{\boldsymbol{R}}\,.$$

As in the Cheon *et al.* attack we do this twice, once with $W_{\boldsymbol{C}}$ and once with $W_{\boldsymbol{I}}$, where $\boldsymbol{I}$ is the identity matrix. We can then compute:

$$W_{\boldsymbol{C}} \cdot W_{\boldsymbol{I}}^{-1} = \hat{\boldsymbol{X}} \times \hat{\boldsymbol{C}}^* \times \boldsymbol{X}^{-1}\,.$$

We can therefore compute over $\mathbb{Z}$ the characteristic polynomial $f(x)$ of $\hat{\boldsymbol{C}}^*$.

If the matrix $\boldsymbol{C}^*$ in (4) was diagonal, then the eigenvalues of $W_{\boldsymbol{C}} \cdot W_{\boldsymbol{I}}^{-1}$, which are the same as the eigenvalues of $\hat{\boldsymbol{C}}^*$, would give the CRT components modulo p_i of all native CLT encodings in the diagonal of $\boldsymbol{C}^*$, from which one could recover the p_i's and all secret values.

Now the matrix $\boldsymbol{C}^*$ is not diagonal, hence we proceed as follows. Since $\hat{\boldsymbol{C}}^*$ is block-diagonal with the sub-matrices $\boldsymbol{C}^* \bmod p_i$ on the diagonal, the characteristic polynomial $f(x)$ of $\boldsymbol{C}^*$ is the product of the characteristic polynomials $f_i(x)$ of the sub-matrices $\boldsymbol{C}^* \bmod p_i$. Therefore we compute the factorization of $f(x)$ in $\mathbb{Z}$:

$$f(x) = \prod_{i=1}^{n} f_i(x)\,.$$

This can be done in polynomial time. Since f_i is also the characteristic polynomial of the matrix $\boldsymbol{C} \bmod p_i$, by the Cayley-Hamilton theorem we have:

$$f_i(\boldsymbol{C}) = \boldsymbol{0} \bmod p_i$$

for all $1 \leqslant i \leqslant n$. Therefore the secret p_i's can be recovered by taking the gcd of one non-zero coefficient of the matrix $f_i(\boldsymbol{C})$ with x_0.

5 Cryptanalysis of Boneh *et al.* Countermeasure

In a recent paper, Boneh, Wu and Zimmerman described another transformation of CLT multilinear maps in order to resist the Cheon *et al.* attack [BWZ14]. In this section we show that this transformation is also insecure.

The Boneh *et al.* Countermeasure. For simplicity we describe the countermeasure in the symmetric setting. We also assume that the message space $\mathbb{Z}_N$ is the direct product $\mathbb{Z}_{g_1} \times \cdots \times \mathbb{Z}_{g_n}$; that is we take $\Theta = 1$ in [BWZ14] —our attack naturally extends to any Θ. Following [BWZ14] we denote an encoding of an element $x \in \mathbb{Z}_N$ by $[x_1, \ldots, x_n]$ where $x = x_i \bmod g_i$ for all $1 \leqslant i \leqslant n$.

The transformation creates a new multilinear map ZMM with the same domain $\mathbb{Z}_N$. For this it uses CLT with 2 additional slots, therefore a total of $n+2$ slots, with a domain $\mathbb{Z}_{N'}$ such that $N' = N \cdot g_{n+1} \cdot g_{n+2}$. In the new multilinear map, an encoding of an element $x \in \mathbb{Z}_N$ is the pair of native CLT encodings with $n+2$ slots:

$$c = (x_L, x_R) = \big([x_1, \ldots, x_n, \zeta, \nu_L],\, [\eta_1, \ldots, \eta_n, \zeta, \nu_R] \big)$$

where the scalars ζ, ν_L, ν_R and $\eta_1, \ldots, \eta_n$ are chosen at random in the appropriate rings $\mathbb{Z}_{g_i}$.

To zero-test such encoding c, two additional native CLT encodings are made public:

$$t_L = [1, \ldots, 1, 1, 0] \quad \text{and} \quad t_R = [0, \ldots, 0, 1, 0]\,.$$

We want to test whether $(x_1, \ldots, x_n) = (0, \ldots, 0)$. For this we compute:

$$\omega = p_{zt} \cdot (x_L \cdot t_L - x_R \cdot t_R)\,.$$

We see from the structure of the native CLT encodings x_L, x_R, t_L and t_R that $x_L \cdot t_L - x_R \cdot t_R$ is an encoding of 0 iff $x = 0$. Therefore ω is small compared to x_0 if c is a ZMM encoding of zero (and whp is not when c is not).

Our Attack. We see that ω is still a $2(n+2)$ linear form in the CRT components of the CLT encodings x_L and x_R. Therefore the attack of Cheon *et al.* is easily extended by using matrices of dimension $2(n+2)$ instead of n. This enables to recover all secret parameters.

References

[BWZ14] Dan Boneh, David J. Wu, and Joe Zimmerman. Immunizing multilinear maps against zeroizing attacks. Cryptology ePrint Archive, Report 2014/930, 2014. `http://eprint.iacr.org/`.

[CHL+14] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. Cryptology ePrint Archive, Report 2014/906, 2014. `http://eprint.iacr.org/`.

[CLT13] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2013.

[GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.

[GGH14] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. Cryptology ePrint Archive, Report 2014/645, 2014. `http://eprint.iacr.org/`.

[GGHZ14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Fully secure functional encryption without obfuscation. Cryptology ePrint Archive, Report 2014/666, 2014. `http://eprint.iacr.org/`.