
COMPARISON-BASED ATTACKS AGAINST NOISE-FREE FULLY HOMOMORPHIC ENCRYPTION SCHEMES

We remark that this article is a preprint. The final version of this work is published in the proceedings of the 2018 International Conference on Information and Communication Security, October 29-31, 2018. DOI: http://dx.doi.org/10.1007/978-3-030-01950-1_11

Alessandro Barengi¹
alessandro.barengi@polimi.it

Nicholas Mainardi¹
nicholas.mainardi@polimi.it

Gerardo Pelosi¹
gerardo.pelosi@polimi.it

¹ Department of Electronics, Information and Bioengineering
Politecnico di Milano, IT

July 26, 2019

ABSTRACT

Homomorphic Encryption provides one of the most promising means to delegate computation to the cloud while retaining data confidentiality. We present a plaintext recovery attack against fully homomorphic schemes which have a polynomial time distinguisher for a given fixed plaintext, and rely on the capability of homomorphically compare a pair of encrypted integer values. We improve by a constant factor the computational complexity of an exhaustive search strategy, which is linear in the recovered plaintext value, and show that it significantly increases the number of recoverable plaintexts. We successfully validate our attack against two noise-free fully homomorphic encryption schemes, which fulfill the mentioned requisite and were claimed to be secure against plaintext recovery attacks.

Keywords FHE · Noise-free schemes · Plaintext Recovery Attack

1 Introduction

Fully Homomorphic Encryption (FHE) is a powerful cryptographic primitive, which allows to perform computation on encrypted data, retaining the correctness of the computation once the result is decrypted. The idea of FHE was first proposed by Rivest in 1978 with the name of *Privacy Homomorphisms* [21]. Designing a FHE scheme remained an open problem for three decades, during which only partially homomorphic encryption (PHE) schemes, which allow to perform only a set of operations (e.g., additions), or SomeWhat Homomorphic Encryption (SWHE) schemes, which allow to perform only a limited number of additions and multiplications, were proposed. In 2009 Gentry [11] proposed the first FHE scheme, allowing to perform an unbounded number of additions and multiplications on encrypted data. Despite the low computational efficiency, FHE has gained attention as it provides a way to outsource computation on private data to a third party such as a cloud-hosted service without revealing any information neither about the data involved in the computation nor about its result, since it is decrypted by the client alone, differently from other primitives such as Secure Multi-Party Computation [25] or Functional Encryption [1]. Since Gentry's seminal work, several schemes achieving better performances were proposed [6, 7, 10, 13], as well as new techniques to

speed up homomorphic computations, such as *batching* [12, 22]. Nevertheless, FHE schemes still have two practical concerns to be solved before wide adoption is possible: i) ciphertext expansion and ii) the computational overhead imposed on homomorphic operations to preserve the correctness of the decrypted result. Indeed, the ciphertext space of SWHE/FHE schemes is consistently larger than the plaintext one, therefore even a single operation on ciphertexts is quite time consuming. The preservation of the correctness of the decrypted result needs to cope with a certain amount of randomness, typically called *noise*, that is added to the ciphertext when processing it. The amount of noise cannot be too high, lest a decryption failure occurs. Unfortunately, each homomorphic operation, especially multiplication, increases the amount of noise in the ciphertexts. Therefore, after a while, the computation must be halted (as in SWHE schemes), or a procedure to refresh the ciphertext, i.e., decrease the noise without decrypting, must be run. Such a procedure, introduced by Gentry in his original scheme [11], is called *bootstrapping* and it allows to transform a SWHE scheme, satisfying certain constraints, in a FHE one. However, this procedure is quite cumbersome, and needs to be periodically performed, slowing down the overall computation. To overcome this burden, alternative noise management techniques have been proposed, such as modulus switching [5] and scale-invariant schemes [2, 3]. Acknowledging the difficulties in noise handling, some noise-free schemes have also been proposed: their ciphertexts have no noise, thus an unbounded number of operations can be performed without any costly noise management technique being involved. Nevertheless, while common noisy SWHE/FHE schemes are based on well-known and scrutinized mathematical problems, such as the Learning With Errors problem [17], noise-free schemes usually rely on less common algebraic trapdoors, which typically do not have widely scrutinized reductions to hard problems. Indeed, Liu in [16] proposed a noise-free FHE scheme, based on the *approximate greatest common divisor* problem, that was subsequently broken in [24]. Kipnis in [15] proposed a FHE scheme based on commutative rings, provably secure against ciphertext-only attacks; however knowing two plaintext-ciphertext pairs is sufficient to break the scheme [23]. Li in [14] proposed to employ non commutative rings to build FHE schemes, while Nuida [19] introduced a framework to construct FHE schemes based on group presentations obfuscated by Tietze transformations. The open challenge with schemes in [14, 19] resides in the definition of a mapping between integer plaintext values and the elements of the mentioned algebraic structures, without losing neither security guarantees nor homomorphic capabilities. Lastly, Wang in [24] introduced two noise-free octonion-based FHE schemes (called *OctoM* and *JordanM*) with trapdoors based on solving quadratic modular equations (with a composite modulus) and proved their security in a ciphertext-only scenario. Due to this property, they are, to the best of our knowledge, the only noise-free FHE schemes suitable for practical usage.

While in general the homomorphic capabilities of a cryptosystem do not weaken the security guarantees per se, they may increase the adversarial power, if combined with other vulnerabilities. The advantages provided by homomorphic capabilities to the attacker were discussed in [4], focusing on the so-called *linearly decryptable* schemes, i.e., cryptosystems whose decryption function can be expressed as a *dot product* between key and ciphertext values represented in a multi-dimensional vector space. Linearly decryptable schemes usually employ a significant amount of noise to hinder Known Plaintext Attacks (KPAs). Nevertheless, in [4] the authors shown that if the scheme can homomorphically evaluate the majority function, then a KPA becomes practically viable. Moreover, in [24] the authors introduced, for linearly decryptable schemes, an algorithm to determine if the plaintext corresponding to a given ciphertext is equal to the integer value 1. We remark that noise free *OctoM* and *JordanM* FHE schemes are linearly decryptable, and thus affected by the aforementioned issues.

Contributions. We present a plaintext recovery attack, against FHE schemes having plaintexts in \mathbb{Z}_n , with $n > 2$, and where it is possible to devise an efficient algorithm able to determine if a generic ciphertext under a given key k is the encryption of a chosen plaintext m , which we denote as *m-distinguisher*. Although, to the best of our knowledge, such a distinguisher has been proposed for linearly decryptable schemes only, our attack will be applicable to any FHE scheme for which such a distinguisher can be found. Our attack, which is performed in a ciphertext-only scenario, leverages the capability to homomorphically compare two encrypted integer values, obtaining a computational complexity which is linear in the plaintext integer value being recovered and improving over an exhaustive search strategy by a significant constant factor. We successfully validate the proposed attack against two linearly decryptable noise free octonion-based cryptosystems [24] (*OctoM* and *JordanM*), which were claimed to be computationally secure in a ciphertext-only attack scenario. Furthermore, we apply our attack to retrieve enough plaintexts from the said schemes so that mounting a KPA to recover the key becomes viable.

2 Preliminaries

Definition 1 (Negligible Function). *A function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if, for every univariate positive polynomial, $poly(x) \in \mathbb{R}[x]$, there exists an integer $c > 0$ such that $\forall x > c$, $|\epsilon(x)| \leq \frac{1}{poly(x)}$.*

Definition 2 (Indicator Function). *Given a set S and a subset $A \subseteq S$, the indicator function of the elements of A over the ones included in S is defined as: $\mathbb{1}_A : S \rightarrow \{0, 1\}$, where $\mathbb{1}_A(x) = 1$ if $x \in A$, 0 otherwise.*

2.1 Homomorphic Encryption Algorithms

Our definition of Fully Homomorphic Encryption follows [8], without constraining the encryption function to deal with a single bit at time.

An homomorphic encryption (HE) scheme specifies three sets: \mathcal{M} , \mathcal{C} and \mathcal{F} . The set of plaintexts \mathcal{M} usually coincides with the set of integer values ranging from 0 to n , with $n \geq 2$, assumed to be the representatives of the residue classes modulo n , i.e., $\mathbb{Z}_n \equiv \mathbb{Z}/n\mathbb{Z}$. The ciphertext space \mathcal{C} includes elements with an algebraic representation that depends on the specific HE scheme at hand. The set of polynomials $\mathcal{F} \subseteq \mathbb{Z}_n[x_1, x_2, \dots, x_a]$, with $a \geq 1$ and degree greater or equal to zero, defines the functions that the HE scheme at hand allows to be evaluated. That is, each of these polynomials computes a function $f : \mathcal{M}^a \rightarrow \mathcal{M}$, $a \geq 1$ over the plaintexts, and is also referred to as an *arithmetic circuit* composed by gates performing multiplications and additions in \mathbb{Z}_n . We provide the definition of an HE scheme starting from an asymmetric HE scheme, and describe a symmetric one by difference.

Definition 3 (Public-key Homomorphic Encryption Scheme). *A public-key Homomorphic Encryption scheme is defined as a tuple of four polynomial time algorithms $\langle \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval} \rangle$:*

- **Key Generation.** $\langle sk, pk, evk \rangle \leftarrow \text{KeyGen}(1^\lambda)$ is a probabilistic algorithm that, given the security parameter λ , generates the secret key sk , the public key pk and the public evaluation key evk .
- **Encryption.** $c \leftarrow \text{Enc}(pk, m)$ is a probabilistic algorithm that, given a message $m \in \mathcal{M}$ and the public key pk , computes a ciphertext $c \in \mathcal{C}$
- **Decryption.** $m \leftarrow \text{Dec}(sk, c)$ is a deterministic algorithm that, given a ciphertext $c \in \mathcal{C}$ and the secret key sk , outputs a message $m \in \mathcal{M}$
- **Evaluation.** $c \leftarrow \text{Eval}(evk, f, c_1, c_2, \dots, c_a)$ is a probabilistic algorithm computing a ciphertext $c \in \mathcal{C}$, using an arithmetic circuit $f \in \mathcal{F}$ with $a \geq 1$ inputs, the ciphertexts c_1, c_2, \dots, c_a , and the evaluation key.

The following properties must hold:

- **Decryption Correctness.** $\forall m \in \mathcal{M} : \text{Dec}(sk, \text{Enc}(pk, m)) = m$.
- **Evaluation Correctness.** $\forall f \in \mathcal{F}, m_1, \dots, m_a \in \mathcal{M}$:
 $\Pr(\text{Dec}(sk, \text{Eval}(evk, f, c_1, \dots, c_a)) = f(m_1, \dots, m_a)) = 1 - \epsilon(\lambda)$,
 where $c_1 = \text{Enc}(pk, m_1) \wedge \dots \wedge c_a = \text{Enc}(pk, m_a)$ and $\epsilon(\lambda)$ is a negligible function of the security parameter of the scheme.
- **Compactness.** $\forall f \in \mathcal{F}, c_1, \dots, c_k \in \mathcal{C}$:
 $|\text{Eval}(evk, f, c_1, \dots, c_k)| \leq poly(\lambda)$, where $|\cdot|$ denotes the bit length of a ciphertext, while $poly(\cdot)$ denotes a positive univariate polynomial.

The requirement on the evaluation correctness trivially states that by decrypting the output of the Eval algorithm we obtain the correct result of the computation homomorphically performed by Eval on the ciphertexts. In particular, the Eval algorithm evaluates a polynomial, defined over the plaintext space, in the sequence of input ciphertexts by replacing the modular additions and multiplications with the homomorphic operations Add and Mul, respectively, that are, in turn, two probabilistic polynomial time algorithms defined over the ciphertext space \mathcal{C} :

- **Homomorphic Addition.** $c \leftarrow \text{Add}(evk, c_1, c_2)$ computes a ciphertext $c \in \mathcal{C}$ such that $\text{Dec}(sk, c) = \text{Dec}(sk, c_1) + \text{Dec}(sk, c_2)$
- **Homomorphic Multiplication.** $c \leftarrow \text{Mul}(evk, c_1, c_2)$ computes a ciphertext $c \in \mathcal{C}$ such that $\text{Dec}(sk, c) = \text{Dec}(sk, c_1) \cdot \text{Dec}(sk, c_2)$

When defining a symmetric-key homomorphic encryption scheme, the only difference is the key generation algorithm $\text{KeyGen}(1^\lambda)$ outputting a tuple $\mathbf{k} = \langle sk, pk, evk \rangle$ with $sk = pk$. We recall the categorization of HE schemes depending on the specific choice of the set of functions \mathcal{F} which can be evaluated. Specifically, a PHE scheme exhibits a function $f \in \mathcal{F}$ defined via an arithmetic circuit including a single type of gate (an additive one or a multiplicative one). A SWHE scheme exhibits a function $f \in \mathcal{F}$ defined via an arithmetic circuit with a depth no higher than a fixed (scheme-dependent) threshold. Finally, a FHE scheme exhibits a function $f \in \mathcal{F}$ defined via an unconstrained arithmetic circuit.

2.2 Homomorphic Comparisons

One of the requirements to apply our attack is the existence of an algorithm able to determine if a generic ciphertext is an encryption of a fixed plaintext m . Therefore, we now provide a formal definition for this algorithm, which we refer to as m -distinguisher.

Definition 4 (m -distinguisher). *Let $(\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval})$ be a homomorphic encryption scheme with security margin λ , and let \mathcal{M} , and \mathcal{C} be the plaintext and ciphertext spaces, bounded by the generated key $\mathbf{k} = \langle sk, pk, evk \rangle$. Let $A_{\mathbf{k}}^m \subset \mathcal{C}$, be the set of ciphertexts corresponding to the encryption of a plaintext $m \in \mathcal{M}$, i.e.: $A_{\mathbf{k}}^m = \{c \in \mathcal{C} \text{ s.t. } \text{Dec}(sk, c) = m\}$.*

Given a plaintext $m \in \mathcal{M}$, an m -distinguisher is a deterministic polynomial time algorithm A_m taking as input a ciphertext $c \in \mathcal{C}$ and the public portion of \mathbf{k} (i.e., $\mathbf{k}_{\text{pub}} = \langle pk, evk \rangle$ for public-key systems and $\mathbf{k}_{\text{pub}} = \langle evk \rangle$ for symmetric ones), and computes the indicator function of the elements of $A_{\mathbf{k}}^m$ over the set of ciphertexts, $\mathbb{1}_{A_{\mathbf{k}}^m} : \mathcal{C} \rightarrow \{0, 1\}$, in such a way that

$$\frac{|\{c \in \mathcal{C} \text{ s.t. } A_m(c, \mathbf{k}_{\text{pub}}) = \mathbb{1}_{A_{\mathbf{k}}^m}(c)\}|}{|\mathcal{C}|} \geq 1 - \epsilon(\lambda),$$

where $\epsilon(\lambda)$ is a negligible function of the security margin of the system.

Given the existence of this m -distinguisher, our attack leverages the capability to homomorphically compare two encrypted integers. Therefore, we now present the main methods proposed in the literature to compute this functionality, including the one used in our attack. First of all, performing comparisons requires to homomorphically evaluate the *greater-than* function on a chosen integer interval.

Definition 5 (Greater-than Function). *Given a positive integer b and an interval of integers $D_t = \{0, 1, \dots, t-1\}$, with $t \geq 2$, the greater-than function $GT_{t,b} : D_t \times D_t \rightarrow \{b-1, b\}$ is defined as:*

$$GT_{t,b}(x, y) = \begin{cases} b & \text{if } x \geq y, \\ b-1 & \text{otherwise} \end{cases}$$

To the extent of evaluating this function with an HE scheme, we need to find a polynomial $f_{gt} \in \mathcal{F} \subseteq \mathbb{Z}_n[x, y]$, such that $f_{gt}(x, y) = GT_{t,b}(x, y)$, with $2 \leq t < n$, $1 \leq b < n$, with x, y being the representatives of residue classes modulo n , (i.e., $x, y \in \mathbb{Z}_n$) considered as integers less than t . Such a polynomial can be easily found if the plaintext space is \mathbb{Z}_2 : indeed, additions and multiplications become xor and and gates, while the input variables are the single-bit values in the binary encodings of x and y , and thus there are many logical circuits computing $GT_{t,b}(\cdot, \cdot)$ function.

Considering a plaintext ring $\mathcal{M} = \mathbb{Z}_n$, with $n > 2$, which is the case targeted in our work, finding an efficiently computable polynomial for the $GT_{t,b}(\cdot, \cdot)$ function is a challenging task. Çetin in [9] reports two methods to compute the $GT_{t,b}(\cdot, \cdot)$ function which do not require interaction between the secret key owner and the party who performs homomorphic evaluation. However, both of these methods are not suitable for our attack: indeed, the first one is not applicable to a composite module n ; the second method computes an approximation of the $GT_{t,b}(\cdot, \cdot)$, while our attack

needs an exact computation of this function.

A more effective solution is proposed in [18]: the greater-than function is computed as $GT_{t,b}(x, y) = SIGN_{t,b}(x - y)$, where $SIGN_{t,b}(z)$ is a function defined over $\overline{D}_t \subseteq \mathbb{Z} = \{-t+1, \dots, 0, \dots, t-1\}$ such that $SIGN_{t,b}(z) = b$ if $z \geq 0$, $b - 1$ otherwise. The homomorphic evaluation of the function $SIGN_{t,b}(\cdot)$ requires polynomial $f_{sign} \in \mathcal{F} \subseteq \mathbb{Z}_n[z]$ fulfilling $f_{sign}(z \bmod n) = SIGN_{t,b}(z)$, with $2 \leq t \leq \frac{n}{2}$, $1 \leq b < n$ and $z \in \overline{D}_t$. In [18], the polynomial f_{sign} is computed applying the Lagrange interpolation formula to $2t - 1$ points having coordinates $(z, SIGN_{t,b}(z))$, with $z \in \overline{D}_t$, and considering a prime modulus, i.e., $n = p$.

As we are considering as a plaintext space the ring \mathbb{Z}_n with a generic modulus $n > 2$, we introduce an additional constraint on the integer t , formalized in Lemma 1, to extend the applicability of the aforementioned method to a generic ring \mathbb{Z}_n :

Lemma 1. *Given an integer $t \geq 2$, and a set $\overline{D}_t = \{-t + 1, \dots, 0, \dots, t - 1\}$, the polynomial $f(z) \in \mathbb{Z}_n[z]$, with $n > 2$, interpolating $2t - 1$ points having the x-coordinate ranging over all values in $z \in \overline{D}_t$ exists if $t \leq \frac{q}{2}$, where q is the smallest prime factor of n .*

In conclusion, by Lagrange interpolation we can obtain a polynomial $f_{sign} \in \mathcal{F} \subseteq \mathbb{Z}_n[z]$ which computes the function $SIGN_{t,b}(z), \forall z \in \overline{D}_t$, and then a polynomial $f_{gt} \in \mathcal{F} \subseteq \mathbb{Z}_n[x, y]$, computing the function $GT_{t,b}(x, y), \forall x, y \in D_t$, as $f_{gt}(x, y) = f_{sign}(x - y)$. Since $f_{gt} \in \mathcal{F}$, it can be homomorphically evaluated by the Eval algorithm of the HE scheme, by replacing addition and multiplications of the polynomial with corresponding homomorphic operations (Add and Mul) in the ciphertext space \mathcal{C} . In the following, we denote the the algorithm $\text{Eval}(evk, c_1, c_2, f_{gt})$ by $HGT_{t,b}(c_1, c_2)$, which takes as input two ciphertexts with corresponding plaintexts $m_1, m_2 \in D_t$, and outputs an encryption of $GT_{t,b}(m_1, m_2)$. In particular, since $GT_{t,b}$ is defined on the interval $D_t = \{0, \dots, t - 1\}$, $t \leq \frac{q}{2}$, with q being the smallest prime factor of n , then $c_1, c_2 \in \mathcal{C}_t = \{c \in \mathcal{C} \text{ s.t. } \text{Dec}(sk, c) < t\}$ is a sufficient condition for $\text{Dec}(sk, HGT_{t,b}(c_1, c_2)) = GT_{t,b}(m_1, m_2)$. The computational complexity required to interpolate $2t - 1$ points by applying the Lagrange formula is $O(t^2)$ operations in \mathbb{Z}_n ; while the evaluation of the polynomial $f_{sign} \in \mathbb{Z}_n[z]$, whose degree is at most $2t - 2$, has a computational complexity $O(t)$. From this, it is easy to note that the computational cost of the $HGT_{t,b}(\cdot, \cdot)$ algorithm is $O(t)$. We note that, while there are no current algorithms to compute $HGT_{t,b}(\cdot, \cdot)$ in less than $O(t)$, research efforts driven by the usefulness of a homomorphic comparison may lead to an improvement in this sense. Since our methodology relies on the computation of $HGT_{t,b}(\cdot, \cdot)$ as an atomic component, such improvements will positively affect the efficiency of our attack.

3 Attack Strategy

In the following we detail a plaintext recovery attack which takes as input a ciphertext and the publicly available evaluation key evk of the HE scheme at hand (either a FHE, or a SWHE capable of computing $HGT_{t,b}$). Since a FHE scheme must allow the evaluation of arbitrary polynomials, then it must provide to the evaluator a method to obtain encryptions of known values, preferably avoiding interaction with the key owner. In case the HE scheme is an asymmetric one, such ciphertexts can be directly obtained employing the public key encryption algorithm, while for a symmetric scheme, the encryption of any value can be obtained from a single encryption of $\hat{m}=1$ leveraging homomorphic operations. Thus, we assume, in case of a symmetric FHE scheme, that an encryption of $\hat{m}=1$ is embodied in the evaluation key evk to allow the computation of encryptions of known values by the evaluator. Such encryption of \hat{m} can be used also by the attacker to obtain the ciphertexts required to perform the attack.

Comparison-based Attack. The core idea is to perform a homomorphic binary search over the possible candidates for the value of the plaintext corresponding to the ciphertext at hand. To this end, a comparison function CMP , taking two ciphertexts as inputs and yielding an outcome in cleartext, is computed leveraging the homomorphic greater-than function $HGT_{t,b}$ (see Section 2). Since the result of the $HGT_{t,b}$ function is still encrypted, the m -distinguisher is employed to determine its actual (plaintext) value. The attacker can compute the comparison function CMP employing the aforementioned strategy as follows:

Definition 6 (Comparison Function). *Given the ciphertexts $c_1, c_2 \in \mathcal{C}_t$ and A_m the algorithm computing the m -distinguisher, where m is a fixed plaintext value, the function $CMP : \mathcal{C}_t \times \mathcal{C}_t \rightarrow \{1, 0, -1\}$ is computed as:*

Algorithm 1: Plaintext Recovery Attack

Input: ciphertext $c \in \mathcal{C}_s$, $\mathcal{C}_s = \{c \in \mathcal{C} \text{ s.t. } \text{Dec}(sk, c) < s\}$
Output: plaintext $m_c = \text{Dec}(sk, c)$, $m_c \in \mathbb{Z}_n$

```

1 begin
2   for  $i \leftarrow 1$  to  $\sigma$  do
3      $c_{gt} \leftarrow HGT_{t,m}(c, \text{Enc}(pk, (i-1)t))$ 
4     if  $A_{(m)}(c_{gt}, \mathbf{k}_{\text{pub}}) = 1$  then
5        $c_{gt} \leftarrow HGT_{t,m}(c, \text{Enc}(pk, it))^{-1} + \text{Enc}(pk, 1)$ 
6       if  $A_{(m)}(c_{gt}, \mathbf{k}_{\text{pub}}) = 1$  then
7          $m_c \leftarrow \text{BinarySearch}(c - \text{Enc}(pk, (i-1)t))$ 
8         if  $m_c \neq \perp$  then
9           return  $m_c + (i-1)t$ 

```

$$CMP(c_1, c_2) = \begin{cases} 1 & \text{if } A_m(HGT_{t,m}(c_1, c_2), \mathbf{k}_{\text{pub}}) = 1 \wedge A_m(c_1 - c_2 + c_m, \mathbf{k}_{\text{pub}}) \neq 1, \\ 0 & \text{if } A_m(HGT_{t,m}(c_1, c_2), \mathbf{k}_{\text{pub}}) = 1 \wedge A_m(c_1 - c_2 + c_m, \mathbf{k}_{\text{pub}}) = 1, \\ -1 & \text{otherwise} \end{cases}$$

with c_m being an encryption of m computed by the attacker.

Denoting with T_d the computational complexity of the m -distinguisher, we have that the time complexity of CMP is $T_{CMP} = O(t + 2T_d)$ as its execution involves at most two computations of the m -distinguisher plus one computation of the $HGT_{t,m}$ function, which has complexity $O(t)$. Leveraging the function CMP , the binary search strategy locates the value of the actual plaintext in the range D_t , which is t elements wide, with a computational cost of $O(T_{CMP} \cdot \log(t)) = O((t + 2T_d) \log(t))$.

Starting from the strategy which has just been described, we improve its effectiveness extending the range of the recoverable plaintexts. To this end, we split the set of recoverable plaintexts into $|D_t| = t$ sized chunks, find into which chunk the plaintext is likely to be contained, and proceed to retrieve it employing the binary search approach. We denote with D_s the set of recoverable plaintexts ($D_s = \{0, 1, \dots, s-1\}$, $s \leq n$), and with \mathcal{C}_s the set of ciphertexts obtained encrypting plaintexts in D_s , i.e.: $\mathcal{C}_s = \{c \in \mathcal{C} \text{ s.t. } \text{Dec}(sk, c) < s\}$. The recoverable message space D_s is split into σ chunks containing numerically consecutive plaintexts, with $\sigma = \lceil \frac{s}{t} \rceil$. For instance, the i -th chunk contains the plaintexts in $\{(i-1)t, \dots, it-1\}$.

Algorithm 1 shows how our improved attack is performed. It iterates over all the σ chunks, testing, for each one of it, if the plaintext m_c , corresponding to the input ciphertext c , may be contained in the chunk being scanned (lines 2–9). To this end, the algorithm starts by testing if m_c may be in a chunk $\{(i-1)t, \dots, it-1\}$ by verifying if $GT_{t,m}(m_c, (i-1)t) = m$ (lines 3–4). In case this test succeeds (line 4, case of the if being taken), Algorithm 1 proceeds to test also if m_c is smaller than the upper bound it of the chunk at hand, by verifying that $GT_{t,m}(m_c, it) = m - 1$ with an analogous approach (lines 5–6). We note that the answer of these tests are subject to potential false positives. Indeed, if $m_c \notin \{(i-1)t, \dots, it-1\}$, then $m_c - (i-1)t \notin \overline{D}_t$ or $m_c - (it) \notin \overline{D}_t$: thus, it means that the polynomial $f_{\text{sign}}(z) \in \mathbb{Z}_n[z]$, obtained by interpolating points whose x -coordinates range over \overline{D}_t , is evaluated on a point $z \notin \overline{D}_t$, thus yielding an outcome which is either outside the set $\{m-1, m\}$ or (by coincidence) inside it. Therefore, it may happen that $f_{gt}(m_c, (i-1)t) = f_{\text{sign}}(m_c - (i-1)t) = m$ and $f_{gt}(m_c, it) = f_{\text{sign}}(m_c - (it)) = m-1$ even if $m_c \notin \{(i-1)t, \dots, it-1\}$. In this case, the interval $\{(i-1)t, \dots, it-1\}$ is identified as a false positive. However, these false positive are filtered out later in the algorithm. If the tests at lines 3–6 succeed, then the current chunk may contain the plaintext m_c , and so Algorithm 1 attempts a plaintext recovery employing the binary search approach described in precedence over the current chunk (line 7). However, since the binary search is effective only under the assumption that the sought plaintext is in D_t , Algorithm 1 exploits the homomorphic properties of the scheme to subtract the value of the lower bound of the current chunk from m_c , working on its corresponding ciphertext c . In

¹This is erroneously $it - 1$ in the final version of the paper

this way, line 7 of the algorithm will actually retrieve the value of $m_c \bmod t$, provided that the chunk detection was not reporting a false positive. If a result is returned (line 8), the actual value of m_c is reconstructed adding back the lower bound of the current chunk to the value retrieved by the binary search (line 9).

We now consider the time complexity of Algorithm 1 as a function of the value of the plaintext to be retrieved m_c . Algorithm 1 is expected to perform $\lceil \frac{m_c}{t} \rceil$ iterations of the outer loop. Each one of the iterations, save for the last one, will fail the membership tests with very high probability (false positives are unlikely), thus resulting in a computational effort which is $O(t + T_d)$ at each iteration. However, the overall worst-case complexity $T_a(m_c)$ of the improved plaintext recovery attack is:

$$O\left(\lceil \frac{m_c}{t} \rceil (t + T_d + T_{\text{BinarySearch}})\right) = O\left(\log(t)(m_c + \lceil \frac{m_c}{t} \rceil T_d)\right) \quad (1)$$

Therefore, our attack has a linear time complexity, which is the main reason why it is able to practically recover only ciphertexts whose corresponding plaintext is not too big. Indeed, by setting $t = 2^{20}$ (an upper bound imposed by the $O(t^2)$ computational cost of Lagrange interpolation), we see that, unless $T_d > 2^{23}$, recovering plaintexts as big as 2^{32} still retains a computational complexity $T_a(m_c) < 2^{40}$. Since many typical FHE scenarios involve computations on relatively small values (e.g. power consumption statistics from smart meters), we deem this plaintext recover capability effective enough to be worth considering.

To conclude the description of our attack, we now show the speed-up obtained by Algorithm 1 over an exhaustive search strategy leveraging only the m -distinguisher. This latter attack tries all plaintext values $x \in \mathbb{Z}_n$ in increasing order, with the recovered plaintext being the first x such that $A_m(\text{Enc}(pk, x) - c + \text{Enc}(pk, m)) = 1$. Denoting the value of the recovered plaintext as m_c , with this strategy we employ the m -distinguisher m_c times, therefore the complexity of this approach is $O(m_c T_d)$. The speed-up of our attack over this simple strategy can be computed as follows:

$$\frac{m_c T_d}{T_a(m_c)} = \frac{m_c T_d}{\log(t)(m_c + \lceil \frac{m_c}{t} \rceil T_d)} = \frac{m_c t T_d}{\log(t)(m_c t + m_c T_d)} = \frac{t T_d}{\log(t)(t + T_d)}$$

This calculation shows that our attack improves the exhaustive search strategy by a constant factor, thus without changing its asymptotic complexity. In particular, the speed-up depends on the values of t , chosen by the attacker, and T_d , given by the target scheme. Although this improvement may seem negligible, we will show, for the FHE schemes targeted by our attack, that the magnitude of the speed-up may be significant in practice, as it largely increases the number of recoverable plaintexts.

4 Two Case Studies

In this section, we evaluate our attack against two symmetric noise free FHE schemes [24], OctoM and JordanM. Although there is an efficient 1-distinguisher for these schemes, they were claimed to be secure against ciphertext-only adversaries aiming to recover either the plaintext or the secret key [24], making them a proper target for our attack.

4.1 Target Fully Homomorphic Encryption Schemes

We report a description of the two target symmetric FHE schemes, OctoM² and JordanM, focusing only on the details which are relevant for our attack: the characterization of the ciphertext space and the description of homomorphic operations. We refer the reader to [24] for further details. The plaintext space is the ring of integers \mathbb{Z}_n , with a composite $n > 2$, for both schemes. The ciphertext space of OctoM is the set of 8×8 matrices with entries in \mathbb{Z}_n , i.e., $\mathbb{Z}_n^{8 \times 8}$, while the one of JordanM is the set of 3×3 matrices with entries in $\mathbb{O}(\mathbb{Z}_n)$, where $\mathbb{O}(\mathbb{Z}_n)$ is the non commutative, non associative algebra of octonions whose support is the vector space \mathbb{Z}_n^8 . Both schemes perform a single matrix addition to homomorphically add two ciphertexts, while to homomorphically multiply two ciphertexts C_1, C_2 the two schemes employ different procedures:

²We find out that, to make OctoM multiplicatively homomorphic, some additional constraints are needed: they will be shown in the full version of the paper

- **OctoM Homomorphic Multiplication.** $C_{mul} = C_2 \cdot C_1 \cdot C_{-1}$, where \cdot denotes the matrix dot product and $C_{-1} \in \mathbb{Z}_n^{8 \times 8}$ is an encryption of the plaintext value $n - 1$, embodied in the evaluation key, evk .
- **JordanM Homomorphic Multiplication.** $C_{mul} = C_1 \star C_2$, where \star denotes the Jordan product.

4.2 Security Analysis

As acknowledged in [24], the decryption function of the target FHE schemes can be expressed as a dot product between the key and the ciphertext represented in a d -dimensional vector space defined over the ring \mathbb{Z}_n . For the two target FHE schemes, the ciphertext space dimension d is 64 for OctoM, since a ciphertext is an 8×8 matrix, while $d = 9 \cdot 8 = 72$ for JordanM, since the ciphertext matrix is a 3×3 matrix whose entries are elements of $\mathbb{O}(\mathbb{Z}_n)$. Both schemes are linearly decryptable (see Section 1) and thus vulnerable to KPAs: if the attacker has approximately d plaintexts/ciphertexts pairs, then a linear system of equations can be built to recover the key and decrypt any ciphertext. In addition, it is possible to efficiently compute a 1-distinguisher for any linearly decryptable scheme. We now describe the construction of this distinguisher, since we leverage it to perform our attack. Given a ciphertext C , represented as a d dimensional vector, consider the first $d + 1$ powers of C . Since the ciphertext space dimension is d , then these $d + 1$ ciphertexts are linearly dependent. Therefore, by definition, there are non trivial solutions to the system of d equations with $d + 1$ unknowns a_i defined as $\sum_{i=1}^{d+1} a_i C^i = 0$. Since the decryption function is linear and the encryption scheme is multiplicatively homomorphic, the following equality also holds: $\sum_{i=1}^{d+1} a_i m^i = 0$, where $m = \text{Dec}(sk, C)$. If $m = 1$, this equation becomes $\sum_{i=1}^{d+1} a_i 1^i = 0 \Rightarrow \sum_{i=1}^{d+1} a_i = 0$. Therefore, if the additional constraint $\sum_{i=1}^{d+1} a_i \neq 0$ is added to the system of equations $\sum_{i=1}^{d+1} a_i C^i = 0$, a solution is found if and only if $m \neq 1$. In conclusion, by looking at the solution of this system, we can determine if $m = 1$ or not. The computational complexity of this 1-distinguisher is $O(d^3)$, since solving a system of equations has cubic complexity. We remark that the system can be solved directly on ciphertexts, no knowledge about the plaintexts or the key is required. Therefore, this distinguisher is a particular case of Definition 4, since it does not employ k_{pub} , the publicly available portion of the cipher key. While the existence of the 1-distinguisher and of the KPA is acknowledged by designers of OctoM and JordanM too, their security analysis claims (Theorem 7 in [24]) that the hardness of solving quadratic equation systems in \mathbb{Z}_n (with a composite n) guarantees that no information about plaintexts can be inferred from ciphertexts. The proof of this claim is based on two reductions: first, the problem of finding the secret key is reduced to the problem of solving a system of multivariate quadratic modular equations, then the problem of recovering a plaintext is reduced to the problem of solving a univariate quadratic modular equation. These reductions state that if we can solve quadratic equations in \mathbb{Z}_n , then we can break the cryptosystems, but they do not state that recovering the secret key or a plaintext is as hard as solving quadratic equations in \mathbb{Z}_n . Thus, there is no logical contradiction between the existence of our attack and the hardness of solving quadratic equations over \mathbb{Z}_n (which is as hard as factoring n).

4.3 Breaking Target FHE Schemes with Our Attack

As already discussed in Section 3, since a FHE scheme must allow the evaluation of arbitrary polynomials, then it must provide to the evaluator a method to obtain encryptions of known values, preferably avoiding interaction with the key owner. Since no method to provide this capability was proposed for the considered FHE schemes, we assume that an encryption of 1 is embodied in the evaluation key to provide this capability to users of the FHE scheme. It is worth noting that an encryption of 1 for the OctoM scheme is not necessary in the evaluation key, since it can be computed by squaring the encryption of -1 already provided in the evaluation key. In the instantiation of our attack against JordanM, there is a relevant issue related to the fact that, as outlined in Definition 4, the m -distinguisher may have a wrong outcome on a negligible portion of the ciphertexts. Unfortunately, this portion is not negligible for several ciphertexts being used in our attack. The problem arises because of two random values, denoted as r, s , which are employed to randomize the encryption. In particular, we find out two relevant facts about these values³: first, if $s = 1 \vee r = 1$, then the 1-distinguisher will classify the ciphertext as an encryption of 1 independently from the plaintext value m ; secondly, the evaluation correctness property (see Definition 3) holds not only for the message m ,

³Proofs are omitted for the sake of brevity. They will be included in the full version of the paper as long as a fully detailed description of JordanM and OctoM cryptosystems

but for r, s too. As a consequence, a ciphertext C obtained through homomorphic evaluation of $GT_{t,1}$ function has only four possible assignments to its random values r, s , which are $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$, since the image of $GT_{t,1}$ is $\{0, 1\}$. Therefore, while for a generic ciphertext $Pr(r = 1 \vee s = 1)$ is negligible, and thus this issue is not relevant for the reliability of the 1-distinguisher, for a ciphertext C obtained through homomorphic evaluation of $GT_{t,1}$, the same probability is 0.75, and thus the outcome of the 1-distinguisher, which is fed with C in our attack, is likely to be erroneous.

To overcome this issue, we devise a ciphertext refreshing procedure, which employs the available encryption of 1 to compute a new ciphertext C' as $C + C - \text{Enc}(pk, 1) * C$, having the same plaintext m , but random values r', s' which are highly likely to be different from 1, since they depend on the random values chosen for the encryption of 1. Therefore, in our attack we employ a slightly tailored version of the distinguisher, whose output is equal to $A_1(C) \cdot A_1(C')$. By using this distinguisher, we can perform our attack on both the target FHE schemes to recover plaintexts. Then, after d plaintexts have been recovered, we can perform the KPA and recover the key, breaking the schemes.

We can now estimate the computational complexity of T_d for the target FHE schemes. For linearly decryptable schemes, T_d , the computational complexity of the 1-distinguisher is $O(d^3)$, with $d = 64$ for OctoM and $d = 72$ for JordanM, which means that $T_d = O(2^{19})$ for both schemes. We need to invoke it twice to increase its reliability, therefore the computational complexity for this distinguisher, denoted by T'_d , is $O(2^{20})$. Given this estimation, we can see that it is practical to recover plaintext values as big as 2^{32} , which is expected to be enough for a significant number of ciphertexts in FHE applications. The computational effort required to recover a plaintext value $m_c = 2^{32}$, can be computed as follows (see equation 1 in Section 3), assuming $T'_d = 2^{20}$ and $t = 2^{20}$.

$$T_a(2^{32}) \leq 2^{32} \log(2^{20}) + \left\lceil \frac{2^{32}}{2^{20}} \right\rceil 2^{20} \log(2^{20}) = 2^{32} \cdot 20 + 2^{32} \cdot 20 \leq 2^{38}$$

Conversely, recovering a plaintext as big as $m_c = 2^{32}$ via an exhaustive search strategy has a computational cost of $O(m_c T_d) = 2^{32} \cdot 2^{19} = 2^{51}$ (note that with this strategy we do not need to invoke twice the 1-distinguisher, thus we can use T_d instead of T'_d). Indeed, the speed-up of our attack is: $\frac{t T_d}{\log(t)(t + T'_d)} \geq \frac{2^{39}}{2^5 \cdot 2^{21}} = 2^{13}$. This result shows that the improvement of our attack is not negligible: Considering a computational effort fixed a-priori, the number of plaintexts recoverable by our attack is 2^{13} times bigger than the number of plaintexts recoverable by the exhaustive search strategy (when $t = 2^{20}$). For instance, with a computational cost bounded by 2^{38} , our attack can recover plaintexts up to 2^{32} , while the exhaustive search can recover plaintexts up to 2^{19} . We successfully implemented the OctoM and JordanM cryptosystems as well as the described plaintext recovery attack in Python 2.7, with the intent to verify the effectiveness of the proposed attack. In practice, the security level of the targeted cryptoscheme, affects the computational effort to perform the homomorphic operations as well as the modular arithmetic operations needed to evaluate a m -distinguisher. Therefore, such a dependency from the security level and/or the parameter sizes of the cryptoscheme is included in the computational complexity formulae of both our attack and the exhaustive search one as the same multiplicative factor (which has been omitted in the previous treatment). We remark that independently from the security margin, when the plaintext values are bounded (e.g., less than 2^{32}) our method largely improves the practicality of their derivation employing only ciphertext material.

5 Conclusions

We present a new type of plaintext recovery attack based on the capability of homomorphically evaluating the comparison between two encrypted integers and assuming the existence of an efficient algorithm to determine if a generic ciphertext is an encryption of a known value m . Although the computational cost of our attack is linear in the value of the plaintext being recovered, it significantly improves the number of recoverable plaintexts w.r.t. an exhaustive search strategy, which, in turn, might mean recovering a vast portion of ciphertexts in a FHE application scenario.

Acknowledgment

This work was supported in part by the European Community under grant agreement no. 688201 (M²DC – EU H2020)

References

- [1] Dan Boneh, Amit Sahai, and Brent Waters. Functional Encryption: Definitions and Challenges. In *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, 2011.
- [2] Joppe W. Bos, Kristin E. Lauter, Jake Loftus, and Michael Naehrig. Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme. In *IMACC 2013*, volume 8308 of *LNCS*, pages 45–64. Springer, 2013.
- [3] Zvika Brakerski. Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP. In *CRYPTO 2012*, volume 7417 of *LNCS*, pages 868–886. Springer, 2012.
- [4] Zvika Brakerski. When Homomorphism Becomes a Liability. In *TCC 2013*, volume 7785 of *LNCS*, pages 143–161. Springer, 2013.
- [5] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS 2012*, pages 309–325. ACM, 2012.
- [6] Zvika Brakerski and Vinod Vaikuntanathan. Efficient Fully Homomorphic Encryption from (Standard) LWE. In *FOCS 2011*, pages 97–106. IEEE CS, 2011.
- [7] Zvika Brakerski and Vinod Vaikuntanathan. Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages. In *CRYPTO 2011*, volume 6841 of *LNCS*, pages 505–524. Springer, 2011.
- [8] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 43(2):831–871, 2014.
- [9] Gizem S. Çetin, Yarkin Doröz, Berk Sunar, and William J. Martin. An investigation of complex operations with word-size homomorphic encryption. *ePrint Archive*, (1195), 2015. <https://eprint.iacr.org/2015/1195.pdf>.
- [10] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster Fully Homomorphic Encryption: Bootstrapping in Less Than 0.1 Seconds. In *ASIACRYPT 2016*, volume 10031 of *LNCS*, pages 3–33, 2016.
- [11] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC 2009*, pages 169–178. ACM, 2009.
- [12] Craig Gentry, Shai Halevi, and Nigel P. Smart. Fully Homomorphic Encryption with Polylog Overhead. In Pointcheval and Johansson [20], pages 465–482.
- [13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based. In *CRYPTO 2013*, volume 8042 of *LNCS*, pages 75–92. Springer, 2013.
- [14] Licheng Wang Jing Li. Noise-Free Symmetric Fully Homomorphic Encryption Based on Non-Commutative Rings. IACR ePrint Archive, Report 2015/641, 2015. <https://eprint.iacr.org/2015/641>.
- [15] Aviad Kipnis and Eliphaz Hibshoosh. Efficient Methods for Practical Fully Homomorphic Symmetric-key Encryption, Randomization and Verification. *IACR ePrint Archive*, 2012:637, 2012.
- [16] Dongxi Liu. Practical Fully Homomorphic Encryption without Noise Reduction. *IACR ePrint Archive*, 2015:468, 2015.

- [17] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In Pointcheval and Johansson [20], pages 700–718.
- [18] Harika Narumanchi, Dishant Goyal, Nitesh Emmadi, and Praveen Gauravaram. Performance Analysis of Sorting of FHE Data: Integer-Wise Comparison vs Bit-Wise Comparison. In *AINA 2017*, pages 902–908. IEEE CS, 2017.
- [19] Koji Nuida. A Simple Framework for Noise-Free Construction of Fully Homomorphic Encryption from a Special Class of Non-Commutative Groups. *IACR ePrint Archive*, 2014:97, 2014.
- [20] David Pointcheval and Thomas Johansson, editors. *EUROCRYPT 2012*, volume 7237 of *LNCS*. Springer, 2012.
- [21] R L Rivest, L Adleman, and M L Dertouzos. On Data Banks and Privacy Homomorphisms. *Foundations of Secure Computation*, Academia Press, 1978.
- [22] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic SIMD operations. *Des. Codes Cryptography*, 71(1):57–81, 2014.
- [23] Boaz Tsaban and Noam Lifshitz. Cryptanalysis of the MORE symmetric key fully homomorphic encryption scheme. *J. Mathematical Cryptology*, 9(2):75–78, 2015.
- [24] Yongge Wang and Qutaibah M. Malluhi. Privacy Preserving Computation in Cloud Using Noise-Free Fully Homomorphic Encryption (FHE) Schemes. In *ESORICS 2016*, volume 9878 of *LNCS*, pages 301–323. Springer, 2016.
- [25] Andrew Chi-Chih Yao. Protocols for Secure Computations (Extended Abstract). In *FOCS 1982*, pages 160–164. IEEE CS, 1982.