

Fully Homomorphic Encryption over the Integers with Shorter Public Keys^{*}

Jean-Sébastien Coron¹, Avradip Mandal¹, David Naccache², and Mehdi Tibouchi^{1,2}

¹ Université du Luxembourg
6, rue Richard Coudenhove-Kalergi
L-1359 Luxembourg, Luxembourg
{jean-sebastien.coron, avradip.mandal}@uni.lu
² École normale supérieure
Département d'informatique, Groupe de cryptographie
45, rue d'Ulm, F-75230 Paris CEDEX 05, France
{david.naccache, mehdi.tibouchi}@ens.fr

Abstract. At Eurocrypt 2010 van Dijk *et al.* described a fully homomorphic encryption scheme over the integers. The main appeal of this scheme (compared to Gentry's) is its conceptual simplicity. This simplicity comes at the expense of a public key size in $\tilde{O}(\lambda^{10})$ which is too large for any practical system. In this paper we reduce the public key size to $\tilde{O}(\lambda^7)$ by encrypting with a quadratic form in the public key elements, instead of a linear form. We prove that the scheme remains semantically secure, based on a stronger variant of the approximate-GCD problem, already considered by van Dijk *et al.*

We also describe the first implementation of the resulting fully homomorphic scheme. Borrowing some optimizations from the recent Gentry-Halevi implementation of Gentry's scheme, we obtain roughly the same level of efficiency. This shows that fully homomorphic encryption can be implemented using simple arithmetic operations.

1 Introduction

Fully Homomorphic Encryption. An encryption scheme is homomorphic if it supports operations on encrypted data. For example RSA is multiplicatively homomorphic since $c_1 = m_1^e \bmod N$ and $c_2 = m_2^e \bmod N$ yield the encryption of $m_1 \cdot m_2$ without using the private key. Similarly, Paillier cryptosystem [13] is additively homomorphic because from $c_1 = g^{m_1} r^N \bmod N^2$ and $c_2 = g^{m_2} s^N \bmod N^2$ one can compute the encryption of $m_1 + m_2$.

In a breakthrough work Gentry described in 2009 the first encryption scheme that supports both addition and multiplication on ciphertexts, *i.e.* a fully homomorphic encryption scheme [6]. The construction proceeds by successive steps: first Gentry describes a “somewhat homomorphic” scheme that supports a limited number of additions and multiplications on ciphertexts. This is because every ciphertext has a noise component and any homomorphic operation applied to ciphertexts increases the noise in the resulting ciphertext. Once this noise reaches a certain threshold the resulting ciphertext does not decrypt correctly anymore; this limits the degree of the polynomial that can be applied to ciphertexts.

Secondly Gentry shows how to “squash” the decryption procedure so that it can be expressed as a low degree polynomial in the bits of the ciphertext and the secret key (equivalently a circuit of small depth). Then the breakthrough idea consists in evaluating this decryption polynomial not on the bits of the ciphertext and the secret key (as in regular decryption), but homomorphically on the *encryption* of those bits. Then instead of recovering the bit plaintext, one gets an encryption of this bit plaintext, *i.e.* yet another ciphertext for the same plaintext; see Figure 1 for an illustration. Now if the degree of the decryption polynomial is small enough, the resulting noise in this new ciphertext can be smaller than in the original ciphertext; this is called the “ciphertext refresh” procedure. Given two refreshed ciphertexts one can apply again the homomorphic operation (either addition or multiplication), which was not necessarily possible on the original ciphertexts because of the noise threshold. Using this “ciphertext refresh” procedure the number of permissible homomorphic operations becomes unlimited and we get a fully homomorphic encryption scheme.

^{*} An extended abstract of this paper will appear at CRYPTO 2011. This is the full version.

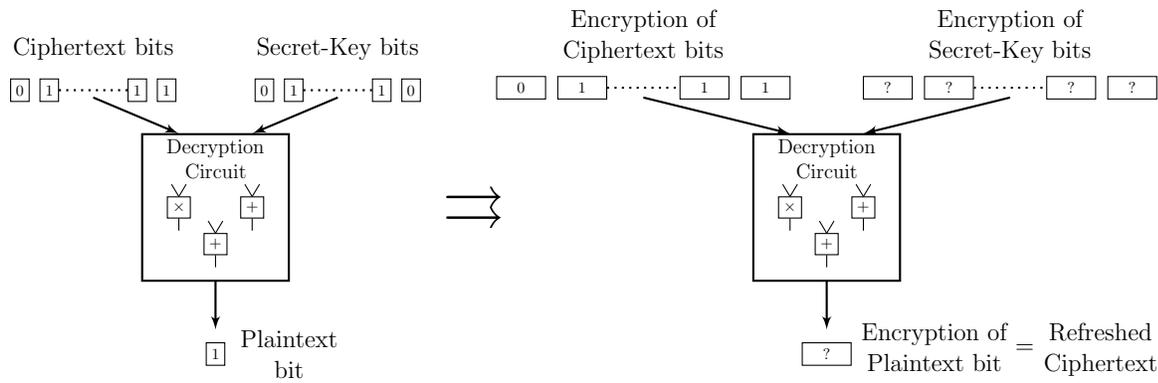


Fig. 1. The decryption circuit applied on the ciphertext bits and secret key bits (left), and the ciphertext refresh procedure with the decryption circuit applied homomorphically on the encryption of those bits (right).

The prerequisite for the “ciphertext refresh” procedure is that the degree of the polynomial that can be evaluated on ciphertexts exceeds the degree of the decryption polynomial (times two, since one must allow for a subsequent addition or multiplication of refreshed ciphertexts); this is called the “bootstrappability” condition. Once the scheme becomes bootstrappable it can be converted into a fully homomorphic encryption scheme by providing the encryption of the secret key bits inside the public key.

Based on Gentry’s approach, two different fully homomorphic schemes are known: Gentry’s scheme [6] based on ideal lattices and a scheme by van Dijk, Gentry, Halevi and Vaikuntanathan (DGHV) over the integers, that appeared at Eurocrypt 2010 [4].

Gentry’s scheme and its implementations. Gentry described in [6] a somewhat homomorphic encryption scheme that is similar to GGH [8, 15] over ideal lattices. To reduce the degree of the decryption polynomial, Gentry introduced the following transformation [6]: instead of using the original secret key, the decryption procedure uses a very sparse subset of values that adds up to the secret key; the full set of values is made part of the public key. To apply the new decryption procedure the original ciphertext must first be “expanded” using the full set of public values. This expanded ciphertext can then be decrypted with a low-degree polynomial in the bits of the new secret key (which are the characteristic vector of the sparse subset sum); this is called the “squashed decryption” procedure.

At PKC 2010 Smart and Vercauteren [17] made the first attempt to implement Gentry’s scheme using a variant based on principal ideal lattices and requiring that the determinant of the lattice be a prime number. However the authors of [17] could not obtain a bootstrappable scheme because that would have required a lattice dimension of at least $n = 2^{27}$, whereas due to the prime determinant requirement they could not generate keys for dimensions $n > 2048$.

Gentry and Halevi described in [7] the first implementation of Gentry’s scheme. The authors follow the same direction as Smart and Vercauteren, but for key generation they eliminate the requirement that the determinant is a prime. Additionally they present many clever optimizations. Four concrete parameter settings are provided, from a “toy” setting in dimension 512, to “small”, “medium” and “large” settings of dimensions 2048, 8192 and 32768, respectively. For the “large” setting public key size is 2.3 Gigabytes. The authors of [7] report that for an optimized implementation on a high-end workstation, key generation takes 2.2 hours, encryption takes 3 minutes, and ciphertext refresh takes 30 minutes.

The DGHV fully homomorphic scheme over the integers. At Eurocrypt 2010, van Dijk, Gentry, Halevi and Vaikuntanathan described a fully homomorphic encryption scheme over the integers [4]. As in Gentry’s scheme the authors first describe a somewhat homomorphic scheme supporting a limited number of additions and multiplications over encrypted bits. Then they apply Gentry’s “squash decryption”

technique to get a bootstrappable scheme and then Gentry’s “ciphertext refresh” procedure (see Fig. 1) to get a fully homomorphic scheme.

The main appeal of the scheme (compared to the original Gentry’s scheme) is its conceptual simplicity; namely all operations are done over the integers instead of ideal lattices. However the public-key was in $\tilde{\mathcal{O}}(\lambda^{10})$ which is too large for any practical system.

Our Contributions. In this paper we show how to reduce the public key size of the somewhat homomorphic scheme from $\mathcal{O}(\lambda^{10})$ down to $\mathcal{O}(\lambda^7)$. The idea consists in storing only a smaller subset of the public key and then generating the full public key on the fly by combining the elements in the small subset multiplicatively; we describe the new scheme in Section 3. In Section 4 we show that the new scheme is still semantically secure, but under a stronger variant of the approximate GCD assumption.

Our second contribution is to describe an implementation of the fully homomorphic DGHV scheme under our variant, using some of the optimizations from [7]. We use the refined analysis from [18] of the sparse subset sum problem; however we do not use the probabilistic decryption circuit from [18] because as in [7] the error probability is too high for our set of parameters. The main difficulty is to determine a secure set of concrete parameters; our approach is to implement the known attacks, measure their running time and extrapolate for large parameters; we can then fix the concrete parameters according to the desired level of security.

We obtain similar performances as the Gentry-Halevi implementation of Gentry’s scheme [7]. More precisely we use four security levels inspired by the levels from [7] (though they may not be directly comparable due to different notions of “security bits”): “toy”, “small”, “medium” and “large”, corresponding to 42, 52, 62 and 72 bits of security respectively. For “large” parameters, encryption and decryption take 3 minutes and 14 minutes respectively, with a public key size of 800 MBytes. Decryption is always close to instantaneous. This shows that fully homomorphic encryption can be implemented with a simple scheme.

2 The DGHV Scheme over the Integers.

In this section we first recall the somewhat homomorphic encryption scheme published by van Dijk, Gentry, Halevi and Vaikuntanathan at Eurocrypt 2010 [4]. The scheme is based on a set of public integers: $x_i = p \cdot q_i + r_i$, $0 \leq i \leq \tau$, where the integer p is secret.

Notation. We use the same notation as in [4]. For a real number x , we denote by $\lceil x \rceil$, $\lfloor x \rfloor$ and $\lceil x \rceil$ the rounding of x up, down, or to the nearest integer. For a real z and an integer p we denote the reduction of z modulo p by $[z]_p$ with $-p/2 < [z]_p \leq p/2$. We also denote $[z]_p$ by $z \bmod p$. We write $f(\lambda) = \tilde{\mathcal{O}}(g(\lambda))$ if $f(\lambda) = \mathcal{O}(g(\lambda) \log^k g(\lambda))$ for some $k \in \mathbb{N}$.

The scheme parameters. Given the security parameter λ , the following parameters are used:

- γ is the bit-length of the x_i ’s.
- η is the bit-length of secret key p .
- ρ is the bit-length of the noise r_i .
- τ is the number of x_i ’s in the public key.
- ρ' is a secondary noise parameter used for encryption.

For a specific η -bit odd integer p , we use the following distribution over γ -bit integers:

$$\mathcal{D}_{\gamma,\rho}(p) = \{ \text{Choose } q \leftarrow \mathbb{Z} \cap [0, 2^\gamma/p), r \leftarrow \mathbb{Z} \cap (-2^\rho, 2^\rho) : \text{Output } x = q \cdot p + r \}$$

KeyGen(1^λ). Generate a random odd integer p of size η bits. For $0 \leq i \leq \tau$ sample $x_i \leftarrow \mathcal{D}_{\gamma,\rho}(p)$. Relabel so that x_0 is the largest. Restart unless x_0 is odd and $[x_0]_p$ is even. Let $pk = (x_0, x_1, \dots, x_\tau)$ and $sk = p$.

$\text{Encrypt}(pk, m \in \{0, 1\})$. Choose a random subset $S \subseteq \{1, 2, \dots, \tau\}$ and a random integer r in $(-2^{\rho'}, 2^{\rho'})$, and output the ciphertext:

$$c = \left[m + 2r + 2 \sum_{i \in S} x_i \right]_{x_0}$$

$\text{Evaluate}(pk, C, c_1, \dots, c_t)$: given the circuit C with t input bits, and t ciphertexts c_i , apply the addition and multiplication gates of C to the ciphertexts, performing all the additions and multiplications over the integers, and return the resulting integer.

$\text{Decrypt}(sk, c)$. Output $m \leftarrow (c \bmod p) \bmod 2$. Note that since $c \bmod p = c - p \cdot \lfloor c/p \rfloor$ and p is odd, one can compute instead: $m \leftarrow [c]_2 \oplus [\lfloor c/p \rfloor]_2$.

This completes the description of the scheme. It is shown in [4] that the scheme is a somewhat homomorphic scheme and that it is semantically secure under the approximate-GCD assumption.

Definition 2.1 (Approximate GCD). *The (ρ, η, γ) -approximate-GCD problem is: For a random η -bit odd integer p , given polynomially many samples from $\mathcal{D}_{\gamma, \rho}(p)$, output p .*

Note that after one **Mult** operation $c \leftarrow c_1 \cdot c_2$ the ciphertext size doubles since there is no modular reduction involved. To reduce the ciphertext size after one **Mult** two techniques are described in [4]. The second and simpler technique consists in generating x_0 without noise, that is $x_0 = q_0 \cdot p$, and then reducing the ciphertext modulo x_0 . The scheme is still semantically secure under the (stronger) approximate-GCD assumption with error-free x_0 . While this problem seems easier to solve, as the adversary is given an exact multiple of p , no better attack is known against it than on the unmodified problem.

We recall the constraints on the scheme parameters [4]:

- $\rho = \omega(\log \lambda)$ to avoid brute force attack on the noise (see Section 6.1).
- $\eta \geq \rho \cdot \Theta(\lambda \log^2 \lambda)$ in order to support homomorphic operations for evaluating the “squashed decryption circuit” (see Section 5).
- $\gamma = \omega(\eta^2 \cdot \log \lambda)$ in order to thwart lattice-based attacks (see Section 6).
- $\tau \geq \gamma + \omega(\log \lambda)$ for the reduction to approximate GCD [4].
- $\rho' = \rho + \omega(\log \lambda)$ for the secondary noise parameter.

To satisfy these constraints the following parameter set is suggested in [4]: $\rho = \lambda$, $\rho' = 2\lambda$, $\eta = \tilde{O}(\lambda^2)$, $\gamma = \tilde{O}(\lambda^5)$ and $\tau = \gamma + \lambda$. The public key size is then $\tilde{O}(\lambda^{10})$. In practice the size of the x_i 's should be at least $\gamma = 2^{23}$ bits to prevent lattice attacks. The public key size is then at least 2^{46} bits, which is too large for any practical system.

3 Our Variant of the DGHV Scheme

3.1 Description

Our technique consists in working with integers x'_{ij} of the form $x'_{ij} = x_{i,0} \cdot x_{j,1} \bmod x_0$ for $1 \leq i, j \leq \beta$ where β is a new parameter. Then only 2β integers $x_{i,b}$ need to be stored in the public key in order to generate the $\tau = \beta^2$ integers x'_{ij} used for encryption. In other words we encrypt using a quadratic form in the public key elements instead of a linear form, which enables to reduce the public key size from τ down to roughly $2\sqrt{\tau}$ integers of γ bits.

Our technique requires to use an error-free x_0 , that is $x_0 = q_0 \cdot p$, since otherwise the error would grow too large. Additionally for encryption we consider a linear combination of the $x'_{i,j}$ with coefficients in $[0, 2^\alpha)$ instead of bits; this enables to further reduce the public key size.

KeyGen(1^λ). Generate a random prime $p \in [2^{\eta-1}, 2^\eta)$. Let $x_0 = q_0 \cdot p$ where q_0 is a random square free 2^λ -rough³ integer in $[0, 2^\gamma/p)$. Generate integers $x_{i,b}$ for $1 \leq i \leq \beta$ and $b \in \{0, 1\}$:

$$x_{i,b} = p \cdot q_{i,b} + r_{i,b}, \quad 1 \leq i \leq \beta, \quad 0 \leq b \leq 1 \quad (1)$$

where $q_{i,b}$ are random integers in $[0, q_0)$ and $r_{i,b}$ are integers in $(-2^\rho, 2^\rho)$. Let $sk = p$ and $pk = (x_0, x_{1,0}, x_{1,1}, \dots, x_{\beta,0}, x_{\beta,1})$.

Encrypt($pk, m \in \{0, 1\}$). Generate a random vector $\mathbf{b} = (b_{i,j})$ of size $\tau = \beta^2$ and with components in $[0, 2^\alpha)$. Generate a random integer r in $(-2^{\rho'}, 2^{\rho'})$. Output the ciphertext:

$$c = m + 2r + 2 \sum_{1 \leq i,j \leq \beta} b_{i,j} \cdot x_{i,0} \cdot x_{j,1} \quad \text{mod } x_0 \quad (2)$$

Evaluate and Decrypt: same as in the original scheme, except that ciphertexts are reduced modulo x_0 after addition and multiplication.

3.2 Constraints on the Parameters

The first three constraints are the same as in the original DGHV scheme:

- $\rho = \omega(\log \lambda)$ to avoid brute force attack on the noise (see Section 6.1).
- $\eta \geq (2\rho + \alpha) \cdot \Theta(\lambda \log^2 \lambda)$ in order to support homomorphic operations for evaluating the “squashed decryption circuit” (see Section 5).
- $\gamma = \omega(\eta^2 \cdot \log \lambda)$ in order to thwart lattice-based attacks (see Section 6).
- $\alpha \cdot \beta^2 \geq \gamma + \omega(\log \lambda)$ for the reduction to approximate GCD (see Section 4).
- $\rho' = 2\rho + \alpha + \omega(\log \lambda)$ for the secondary noise parameter (see Section 4).

To satisfy these conditions we can still take $\rho = \lambda$, $\eta = \tilde{O}(\lambda^2)$ and $\gamma = \tilde{O}(\lambda^5)$ as in the original scheme, and we can take $\alpha = \lambda$, $\beta = \tilde{O}(\lambda^2)$ and $\rho' = 4\lambda$. The main difference is that instead of having $\tau = \tilde{O}(\lambda^5)$ integers x_i 's, we now have only $2\beta = \tilde{O}(\lambda^2)$ integers x_i . Hence the public key size becomes $\tilde{O}(\lambda^7)$ instead of $\tilde{O}(\lambda^{10})$. In Section 7.6 we describe concrete parameters in order to resist all known attacks.

Remark 3.1. It is possible to generate q_0 as a uniformly random square free 2^λ -rough integer of suitable size in probabilistic polynomial time: it suffices to generate a uniformly random number with known factorization [1] and try again if it has small or repeated factors. However, this makes key generation rather unpractical. Alternatively, one can choose q_0 as the product of $(\gamma - \eta)/\lambda^2$ random primes, each of size λ^2 bits.⁴ This is faster, but the security of the scheme then depends on a slightly more convoluted, though no less plausible, computational assumption, to account for the different key distribution.

3.3 Correctness

We refer to Appendix A for the definition of correct homomorphic scheme with respect to a given circuit or circuit set. As in [4, 6] we define a *permitted circuit* as one where for any $i \geq 1$ and any set of integer inputs all less than $\tau^i \cdot 2^{i(\rho'+2)}$ in absolute value, the generalized circuit's output has absolute value at most $2^{i(\eta-3-n)}$ with $n = \lceil \log_2(\lambda + 1) \rceil$; we let $\mathcal{C}_\mathcal{E}$ be the set of permitted circuits. As in [4], we have (see proof in Appendix B):

Lemma 3.1. *The scheme from above is correct for $\mathcal{C}_\mathcal{E}$.*

³ An integer is said to be a -rough when it does not contain prime factors smaller than a . Note that for $a > 2$ such integer must be odd.

⁴ The reason we choose λ^2 -bit factors rather than λ is because factorization algorithms like ECM have a complexity subexponential in the size of factors, and can thus be used to extract λ -bit prime factors efficiently. In the implementation, to thwart this attack, it is safe to generate q_0 as a product of, say, 1000-bit primes.

Remark 3.2. Since “fresh” ciphertexts output by `Encrypt` have noise at most $\tau \cdot 2^{\rho'+2}$, the ciphertext output by `Evaluate` applied to a permitted circuit has noise at most $2^{\eta-3-n} < p/(4(\lambda+1))$. A bound of $p/2$ would suffice to ensure correct decryption, but this stronger bound will be useful to prove the correctness of the bootstrappable version of this scheme later on.

Remark 3.3. The definition of a permitted circuit doesn’t seem to give an easy criterion to determine whether a given computation is permitted. However, it is easy to give a sufficient condition on a multivariate polynomial f for the associated arithmetic circuit C to be permitted. If f is of degree d and if the sum of the absolute values of its coefficients is denoted by $\|f\|_1$, then $C \in \mathcal{C}_\mathcal{E}$ provided that:

$$d \leq \frac{\eta - 3 - n - \log \|f\|_1}{\rho' + 2 + 2 \log \beta}$$

Following [4], we refer to such polynomials f as *permitted polynomials*, and denote the set of these polynomials by $\mathcal{P}_\mathcal{E}$.

4 Security of our Variant

4.1 Overview

In this section we show that our variant is still semantically secure, but under the (stronger) error-free approximate GCD assumption. Our security proof follows the same strategy as in [4]: show that an adversary breaking the scheme’s semantic security can be converted into a LSB predictor for $z \bmod p$, where z is an integer such that $z \bmod p$ is small; this in turns enables to recover p in the approximate-GCD problem.

For this one must show that given $c \leftarrow \text{Encrypt}(pk, m)$, the distribution of $c' = [c + z]_{x_0}$ is essentially the same as $\text{Encrypt}(pk, m')$ with $m' = m \oplus [z \bmod p]_2$. In [4] this is done by showing that the distribution of $\lfloor c/p \rfloor = \sum_{i=1}^{\tau} b_i \cdot q_i$ where $\mathbf{b} \leftarrow \{0, 1\}^\tau$ is statistically close to uniform in \mathbb{Z}_{q_0} . For this [4] applies the leftover hash lemma on the hash function family $h(\mathbf{b}) = \sum_{i=1}^{\tau} b_i \cdot q_i \bmod q_0$ parametrized by the q_i ’s, which is clearly pairwise independent.

Similarly to prove the security of our variant we must apply the leftover hash lemma on the hash function family $h' : [0, 2^\alpha]^{\beta^2} \rightarrow \mathbb{Z}_{q_0}$ where:

$$h'(\mathbf{b}) = \sum_{1 \leq i, j \leq \beta} b_{i,j} \cdot q_{i,0} \cdot q_{j,1} \bmod q_0$$

The main difficulty is to show that h' is (almost) pairwise independent; as shown below this requires to study the zeroes of the corresponding quadratic form. We note that our result might be of independent interest since it enables to construct a universal hash function with a small memory footprint.

4.2 Leftover Hash Lemma

A family \mathcal{H} of hash functions $h : X \rightarrow Y$ is pairwise independent if for all $x \neq x'$ it holds that $\Pr_h[h(x) = h(x')] = 1/|Y|$. Since h' is not exactly pairwise independent we introduce a slightly more general definition:⁵

Definition 4.1. A family \mathcal{H} of hash functions $h : X \rightarrow Y$ is ε -pairwise independent if

$$\sum_{x \neq x'} \left(\Pr_{h \leftarrow \mathcal{H}} [h(x) = h(x')] - \frac{1}{|Y|} \right) \leq |X|^2 \cdot \frac{\varepsilon}{|Y|}$$

⁵ Note that this is quite different from “ ε -almost universal hash function families” in the sense of Wegman and Carter [20]. We need the collision probability $\Pr_{h \leftarrow \mathcal{H}} [h(x) = h(x')]$ to be at most $(1 + \varepsilon)/|Y|$ on average, with negligible ε ; $2/|Y|$ is not good enough.

The following lemma is a straightforward generalization of the usual leftover hash lemma.

Lemma 4.1 (Leftover hash lemma). *Let \mathcal{H} be a family of ε -pairwise independent hash functions. Suppose that $h \leftarrow \mathcal{H}$ and $x \leftarrow X$ are chosen uniformly and independently. Then $(h, h(x))$ is $(\frac{1}{2}\sqrt{|Y|/|X|} + \varepsilon)$ -uniform over $\mathcal{H} \times Y$.*

Proof. Let $\mathbf{p} \in \mathbb{R}^{\mathcal{H} \times Y}$ denote the probability vector corresponding to a random choice of $x \in X$ and $h \in \mathcal{H}$. We must show that:

$$|\mathbf{p} - \mathbf{1}|_1 \leq \sqrt{|Y|/|X|} + \varepsilon$$

where $\mathbf{1}$ corresponds to the uniform distribution. Observing that $\mathbf{p} - \mathbf{1}$ is orthogonal to $\mathbf{1}$, we apply the Cauchy-Schwarz inequality and the Pythagorean theorem to obtain:

$$|\mathbf{p} - \mathbf{1}|_1 \leq \sqrt{|H| \cdot |Y|} \cdot \|\mathbf{p} - \mathbf{1}\|_2 = \sqrt{|H| \cdot |Y|} \sqrt{\|\mathbf{p}\|_2^2 - \|\mathbf{1}\|_2^2}$$

We have $\|\mathbf{1}\|_2^2 = 1/(|H| \cdot |Y|)$ and:

$$\begin{aligned} \|\mathbf{p}\|_2^2 &= \sum_{(h',y) \in \mathcal{H} \times Y} \Pr_{(h,x) \leftarrow \mathcal{H} \times X} [h = h'; h(x) = y]^2 = \frac{1}{|H|^2} \sum_{(h,y) \in \mathcal{H} \times Y} \Pr_{x \leftarrow X} [h(x) = y]^2 \\ &= \frac{1}{|H|^2} \sum_{(h,y) \in \mathcal{H} \times Y} \left(\frac{1}{|X|^2} \sum_{(x,x') \in X^2} \mathbf{1}_{h(x)=y} \right)^2 = \frac{1}{|H|^2 \cdot |X|^2} \sum_{(x,x') \in X^2} \sum_{h \in \mathcal{H}} \mathbf{1}_{h(x)=h(x')} \\ &= \frac{1}{|H| \cdot |X|^2} \sum_{(x,x') \in X^2} \Pr_{h \leftarrow \mathcal{H}} [h(x) = h(x')] = \frac{1}{|H| \cdot |X|} + \frac{1}{|H| \cdot |X|^2} \sum_{x \neq x'} \Pr_{h \leftarrow \mathcal{H}} [h(x) = h(x')] \\ &\leq \frac{1}{|H| \cdot |X|} + \frac{1 + \varepsilon}{|H| \cdot |Y|} \end{aligned}$$

Hence, the statistical distance is bounded as:

$$|\mathbf{p} - \mathbf{1}|_1 \leq \sqrt{|H| \cdot |Y|} \sqrt{\frac{1}{|H| \cdot |X|} + \frac{1 + \varepsilon}{|H| \cdot |Y|} - \frac{1}{|H| \cdot |Y|}} = \sqrt{\frac{|Y|}{|X|} + \varepsilon}$$

as desired. \square

4.3 Proof of Pairwise Independence

Let q be an integer. Let \mathcal{H} be a hash function family from $\{0, \dots, 2^\alpha - 1\}^{\beta \times \beta}$ to \mathbb{Z}_q . The members $h \in \mathcal{H}$ are associated to elements $q_{i,0}, q_{i,1}$ of \mathbb{Z}_q for $1 \leq i \leq \beta$. For $\mathbf{b} \in \{0, \dots, 2^\alpha - 1\}^{\beta \times \beta}$, we let:

$$h(\mathbf{b}) = \sum_{1 \leq i, j \leq \beta} b_{ij} q_{i,0} q_{j,1} \pmod{q}$$

Lemma 4.2. *For an odd prime integer q , the hash function family \mathcal{H} is ε -pairwise independent, with:*

$$\varepsilon = \frac{1}{q} + \frac{\beta^2}{2^{\alpha\beta^2 - 2(\alpha+1)\beta}}$$

Proof. For each choice of $\mathbf{b} \neq \mathbf{b}'$, the probability $\Pr_{h \leftarrow \mathcal{H}} [h(\mathbf{b}) = h(\mathbf{b}')]$ can be expressed in terms of the number of zeros of a certain hyperbolic quadratic form in $\mathbb{Z}_q^{2\beta}$. More precisely let $A = (a_{ij})$ be the $\beta \times \beta$ matrix in $\mathbf{M}_\beta(\mathbb{Z}_q)$ given by $a_{ij} = b_{ij} - b'_{ij}$. We have:

$$\Pr_h [h(\mathbf{b}) = h(\mathbf{b}')] = \frac{1}{q^{2\beta}} \# \left\{ (u_1, \dots, u_\beta, v_1, \dots, v_\beta) \in \mathbb{Z}_q^{2\beta} : \sum_{1 \leq i, j \leq \beta} a_{ij} u_i v_j = 0 \right\}$$

Now the quadratic form $Q = \sum_{1 \leq i, j \leq \beta} a_{ij} u_i v_j$ has the matrix $\frac{1}{2} \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}$, which is clearly conjugate to $\frac{1}{2} \begin{pmatrix} 0 & J \\ J & 0 \end{pmatrix}$ where J is the canonical row echelon form of A . It follows that Q is the orthogonal sum of r hyperbolic planes, with r the rank of A . Hence, its number of zeros is well-known (see e.g. [10, Theorem 6.32] for the non-degenerate case, from which the general case follows immediately):

$$\#\left\{ (u_1, \dots, u_\beta, v_1, \dots, v_\beta) \in \mathbb{Z}_q^{2\beta} : \sum_{1 \leq i, j \leq \beta} a_{ij} u_i v_j = 0 \right\} = q^{2\beta-1} + q^{2\beta-r} - q^{2\beta-r-1}$$

In particular, we get:

$$\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] - \frac{1}{q} \leq \frac{1}{q^r}$$

This estimate is quite sufficient for our purposes, except in the case when $r = 1$. Therefore, we need to bound the number of pairs $(\mathbf{b}, \mathbf{b}')$ such that the corresponding matrix A is of rank 1. Noting that A has all its entries in $-2^\alpha + 1, \dots, 2^\alpha - 1$, it is enough to bound the cardinality of the set U_α of matrices of rank 1 in $M_\beta(\mathbb{Z}_q)$ with entries in that interval.

To do so, note that a matrix of rank 1 with a nonzero upper-left entry is entirely determined by its first line and its first column. If the entries are in $\{-2^\alpha + 1, \dots, 2^\alpha - 1\}$, this leaves $2^{\alpha+1} - 2$ choices for the upper-left entries and $(2^{\alpha+1} - 1)^{2\beta-2}$ choices for the remainder of the first line and the first column. Hence, there are less than $2^{2(\alpha+1)(\beta-1)}$ matrices in U_α with a nonzero upper-left entry (and usually much fewer, since not all first lines and first columns need to give rise to matrices with all their entries in the proper interval). The same argument can be applied to any other nonzero entry (i, j) , leading to the coarse bound:

$$|U_\alpha| < \beta^2 \cdot 2^{2(\alpha+1)\beta}$$

Now, the number of pairs $(\mathbf{b}, \mathbf{b}')$ such that the corresponding matrix A is of rank 1 is at most $|X| \cdot |U_\alpha|$, since for any choice of \mathbf{b} , there are at most $|U_\alpha|$ possible values of \mathbf{b}' such that A is in U_α . We can thus bound the value δ defined by:

$$\delta = \frac{|Y|}{|X|^2} \sum_{\mathbf{b} \neq \mathbf{b}'} \left(\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] - \frac{1}{|Y|} \right)$$

as required. Indeed:

$$\begin{aligned} \delta &= \frac{q}{|X|^2} \sum_{\mathbf{b} \neq \mathbf{b}'} \left(\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] - \frac{1}{q} \right) \leq \frac{q}{|X|^2} \left(\sum_{\substack{\mathbf{b} \neq \mathbf{b}' \\ A \notin U_\alpha}} \frac{1}{q^2} + \sum_{\substack{\mathbf{b} \neq \mathbf{b}' \\ A \in U_\alpha}} \frac{1}{q} \right) \\ &\leq \frac{q}{|X|^2} \left(\frac{|X|^2}{q^2} + \frac{|X| \cdot |U_\alpha|}{q} \right) \leq \frac{1}{q} + \frac{|U_\alpha|}{|X|} \leq \frac{1}{q} + \frac{\beta^2}{2^{\alpha\beta^2 - 2(\alpha+1)\beta}} \end{aligned}$$

which concludes the proof. \square

Corollary 4.1. *When q is a product of distinct primes greater than 2^α , the hash function family \mathcal{H} is ε -pairwise independent, with:*

$$\varepsilon = \frac{\log q}{\log p} \left(\frac{e}{p} + \frac{\beta^2 \cdot 2^{(\log q)/(\log p)}}{2^{\alpha\beta^2 - 2(\alpha+1)\beta}} \right) \quad \text{where } p \text{ is the smallest prime factor of } q.$$

Proof. The proof is largely similar to the previous one. See Appendix C for details.

4.4 Semantic Security

We are now ready to show that our variant is semantically secure under the (stronger) error-free approximate GCD assumption. The proof follows the same strategy as [4]; we refer to Appendix D for the details. For two specific integers p and q_0 , we define the modified distribution:

$$\mathcal{D}'_\rho(p, q_0) = \{ \text{Choose } q \leftarrow \mathbb{Z} \cap [0, q_0), r \leftarrow \mathbb{Z} \cap (-2^\rho, 2^\rho) : \text{Output } x = q \cdot p + r \}$$

Definition 4.2 (Error-free approximate GCD). *The (ρ, η, γ) -error-free-approximate-GCD problem is: For a random η -bit prime integer p , given $x_0 = q_0 \cdot p$ where q_0 is a random square free 2^λ -rough integer in $[0, 2^\gamma/p)$, and polynomially many samples from $\mathcal{D}'_\rho(p, q_0)$, output p .*

Theorem 4.1. *Let \mathcal{A} be an attacker with advantage ε against our variant encryption scheme with parameters $(\rho, \rho', \eta, \gamma, \tau = \beta^2)$ polynomial in the security parameter λ . There exists an algorithm \mathcal{B} for solving the (ρ, η, γ) -error-free-approximate-GCD problem that succeeds with probability at least $\varepsilon/2$. The running time of \mathcal{B} is polynomial in the running time of \mathcal{A} , λ and $1/\varepsilon$.*

5 Making the Scheme Fully Homomorphic

5.1 The Squashed Scheme

Gentry’s transformation to “squash the decryption” consists in adding to the public key some extra information about the secret key and use this extra information to “post process” the ciphertext. Then the post-processed ciphertext can be decrypted by a decryption polynomial of small degree. This requires to introduce an additional complexity assumption, namely the sparse subset-sum assumption.

We follow the description of [4]. Three more parameters κ , θ and Θ are added. Concretely, one uses $\theta = \lambda$, $\kappa = \gamma + 2 + \lceil \log_2(\theta + 1) \rceil$, and $\Theta = \tilde{O}(\lambda^3)$.⁶ One adds to the public key a set $\mathbf{y} = \{y_1, \dots, y_\Theta\}$ of rational numbers in $[0, 2)$ with κ bits of precision, such that there is a sparse subset $S \subset \{1, \dots, \Theta\}$ of size θ with $\sum_{i \in S} y_i \simeq 1/p \pmod{2}$. The expanded ciphertext is computed using the y_i ’s. The secret key sk is replaced by the indicator vector of the subset S .

However adding Θ elements y_i each of size κ bits would give a public key of size $\Theta \cdot \kappa = \tilde{O}(\lambda^8)$, instead of $\tilde{O}(\lambda^7)$ in our variant. Therefore instead of storing the y_i ’s in the public key as in [4], we generate the y_i ’s using a pseudo-random generator⁷ $f(\text{se})$. Then only the seed se and y_1 need to be stored in the public key, and the other y_i ’s can be recovered during ciphertext expansion by applying $f(\text{se})$ again. We obtain the following squashed scheme:

KeyGen. Generate $sk^* = p$ and pk^* as before. Set $x_p \leftarrow \lfloor 2^\kappa/p \rfloor$, choose at random a Θ -bit vector $\mathbf{s} = (s_1, \dots, s_\Theta)$ with Hamming weight θ with $s_1 = 1$, and let $S = \{i : s_i = 1\}$.

Initialize a system-wide pseudo-random number generator f with a random seed se , and use $f(\text{se})$ to generate integers $u_i \in [0, 2^{\kappa+1})$ for $2 \leq i \leq \Theta$. Then, set u_1 such that $\sum_{i \in S} u_i = x_p \pmod{2^{\kappa+1}}$. Set $y_i = u_i/2^\kappa$ and $\mathbf{y} = \{y_1, \dots, y_\Theta\}$. Hence each y_i is a positive number smaller than two, with κ bits of precision after the binary point. Also, $\lfloor \sum_{i \in S} y_i \rfloor_2 = (1/p) - \Delta_p$ for some $|\Delta_p| < 2^{-\kappa}$.

Output the secret key $sk = \mathbf{s}$ and public key $pk = (pk^*, \text{se}, y_1)$.

Encrypt and Evaluate. Generate a ciphertext c^* as before. Then for $i \in \{1, \dots, \Theta\}$ set $z_i \leftarrow \lfloor c^* \cdot y_i \rfloor_2$, keeping only $n = \lceil \log_2(\theta + 1) \rceil$ bits of precision after the binary point for each z_i . Output both c^* and $\mathbf{z} = (z_1, \dots, z_\Theta)$.

Decrypt: Output $m \leftarrow \lfloor c^* - \lfloor \sum_i s_i z_i \rfloor \rfloor_2$.

⁶ We use $\Theta = \tilde{O}(\lambda^3)$ instead of $\Theta = \omega(\kappa \cdot \log \lambda)$ in [4] from a better analysis of the hardness of the SSSP problem (see Section 6.3).

⁷ Note that f doesn’t really need to be a cryptographically strong PRNG: all that is needed is that the sparse subset-sum problem remains hard when the subset is generated by f . Heuristically, this is a mild requirement. In our implementation, we use random numbers produced by the PRNG from the `glibc`.

For this one uses a simple grade-school addition of the numbers $a_i = s_i \cdot z_i$. As illustrated in Fig. 2 the bits of the a_i 's are arranged in Θ rows and $n + 1$ columns (one column before the binary point and n columns after). To see how this grade-school addition can be performed efficiently, first recall the following result from [4, §6.2].

Lemma 5.2. *Let $\mathbf{b} = (b_1, b_2, \dots, b_\Theta)$ be any binary vector, and denote its Hamming weight by W . Write the binary digits of W as $W = \overline{W_k \cdots W_1 W_0}^2$. Then the j -th bit W_j of W can be expressed as a binary polynomial of degree exactly 2^j in the b_i 's, namely the 2^j -th elementary symmetric polynomial:*

$$W_j = \sum_{\substack{I \subset \{1, \dots, \Theta\} \\ |I|=2^j}} \prod_{i \in I} b_i$$

Moreover, the bits W_0, W_1, \dots, W_j can be simultaneously computed by an arithmetic circuit of size $2^j \cdot \Theta$.

That the W_j 's are given by elementary symmetric polynomials is classical (see e.g. [2, Lemma 4]). Thus, to compute them, it suffices to find the top 2^j coefficients of the polynomial $(X - b_1)(X - b_2) \cdots (X - b_\Theta)$, which can be done iteratively with at most $2^j \cdot \Theta$ operations. We recall the dynamic programming algorithm from [4]:

Algorithm 1 Computation of the least significant bits W_0, \dots, W_j of the Hamming weight of $\mathbf{b} = (b_1, \dots, b_\Theta)$.

- | | |
|---|---|
| <p>1: $P_{0,0} \leftarrow 1; P_{k,0} \leftarrow 0$ for $k = 1, 2, \dots, 2^j$.</p> <p>2: for $\ell = 1$ to θ do</p> <p>3: for $k = 2^j$ down to 1 do</p> <p>4: $P_{k,\ell} \leftarrow b_\ell \cdot P_{k-1,\ell-1} + P_{k,\ell-1}$;</p> <p>5: end for</p> <p>6: end for</p> <p>7: $W_k \leftarrow P_{2^k,\theta}$ for $k = 0, 1, \dots, j$;</p> <p>8: return W_0, \dots, W_j.</p> | <p>$\triangleright P_{k,\ell}$ is the k-th symmetric polynomial in b_1, \dots, b_ℓ.</p> |
|---|---|
-

Then, the procedure to compute

$$Q = \left\lfloor \sum_{k=1}^{\Theta} a_k \right\rfloor$$

is as follows. We number the columns containing the a_k 's from left to right as 0 (before the binary point), $-1, -2, \dots, -n$.

As usual, grade-school addition starts from the rightmost column (column $-n$). Adding all Θ bits from that column produces a bit of result and a certain number of bits of carry. Since we are only interested in the $n + 1$ least significant bits of the sum, we only need to keep track of the result and the first n carry bits: this amounts to computing the rightmost bits $W_0^{(-n)}, W_1^{(-n)}, \dots, W_n^{(-n)}$ of the Hamming weight $W^{(-n)}$ of column $-n$, which can be done with at most $2^n \cdot \Theta$ multiplications according to the previous lemma.

Now, push carry bit $W_1^{(-n)}$ to column $-n + 1$, carry bit $W_2^{(-n)}$ to column $-n + 2$ and so on. We can then continue the grade-school addition process from column $-n + 1$, where we only need to compute the result and $n - 1$ carry bits, namely the bits $W_j^{(-n+1)}$ of the Hamming weight $W^{(-n+1)}$ of the column, including the possible carry bit from column $-n$. This amounts to at most $2^{n-1} \cdot (\Theta + 1)$ multiplications. When this is done, push the carry bits $W_1^{(-n+1)}, \dots, W_{n-1}^{(-n+1)}$ to columns $-n + 2, -n + 3, \dots, 0$ respectively, move to the next column and continue as before. This is illustrated in Figure 2 for $n = 4$.

The grade-school addition algorithm requires at most $2^n \cdot \Theta + 2^{n-1} \cdot (\Theta + 1) + 2^{n-2} \cdot (\Theta + 2) + \dots + 2 \cdot (\Theta + n - 1) \leq 2^n \cdot (\Theta + n - 1) \leq 4\theta \cdot \Theta$ multiplications and less than $\Theta + \theta^2$ additions between the resulting products.

Furthermore, we can compute the degree of any of the bits involved as a binary polynomial in the bits of the a_i 's. Clearly, $W_j^{(-n)}$ is of degree 2^j . Then we can see inductively from the shape of the elementary symmetric polynomials that for $k \geq 1$, $W_j^{(-n+k)}$ is of degree $2^j + 2^{k+1} - 3$. It follows that the degree of Q as a polynomial in the bits of the a_i 's is 2^n (given by the degree of $W_n^{(-n)}$, which is the highest of all). This is illustrated in Figure 2 for $n = 4$.

Finally, the parity of the integer closest to $\sum a_i$ is obtained by xor-ing the Hamming weights of column 0 and -1 , and the decryption m is then computed as the xor of the result with the least significant bit of c^* . As a result, m is expressed as a polynomial of degree 2^n in the bits of c^* and the a_i 's, with a coefficient vector of 1-norm less than 2Θ . Since a_i is just $s_i \cdot z_i$, this means that m is a polynomial f of the ciphertext bits and the secret key satisfying $d = \deg f = 2^{n+1}$ and $\|f\|_1 \leq 2\Theta$. In view of Remark 3.3, f is a permitted polynomial as long as $d \leq (\eta - 4 - n - \log_2 \Theta)/(\rho' + 2)$ which is satisfied by choosing η according to the constraint in Section 3.2. \square

6 Attacks

In this section we recall the known attacks. For each attack we provide an asymptotic analysis (as in [4]) and we also run the attacks in practice in order to derive concrete parameters for our implementation. We use four security levels inspired by the levels from [7]: “toy”, “small”, “medium” and “large”, corresponding to 42, 52, 62 and 72 bits of security respectively. Therefore our “high” level of security corresponds to the “high” level of security in [7]. For security parameter λ we wish to ensure that the best attack requires at least 2^λ clock cycles on a standard PC.

Note that we use the SAGE [14] interface to the `fp111` lattice reduction package [16] which is to our knowledge the fastest publicly available. However any progress in LLL implementations will require an increase of our security parameters.

6.1 Brute Force Attack on the Noise

The easiest attack is the brute force attack on the noise in the public key. Given $x_0 = q_0 \cdot p$ and $x_1 = q_1 \cdot p + r_1$ with $|r_1| < 2^\rho$, one can guess r_1 and compute $\gcd(x_0, x_1 - r_1)$ to recover p . The state of the art algorithm for computing GCD's is the Stehlé-Zimmermann algorithm [19] with time complexity $\tilde{O}(\gamma)$ for integers of γ bits. The attack complexity is then $2^\rho \cdot \tilde{O}(\gamma)$. Therefore the attack is thwarted if $\rho = \omega(\log \lambda)$.

A better attack [12] consists in computing $p = \gcd(x_0, \prod_{i=-2^\rho}^{2^\rho} (x_1 - i) [x_0])$. Using fast multiplication the asymptotic complexity is also $2^\rho \cdot \tilde{O}(\gamma)$.

Concrete parameters. In order to fix ρ we have run some experiments with GCD computations of huge integers. We first describe the results for the basic GCD attack, using the Stehlé-Zimmermann algorithm [19]. The result of practical experiments with the Sage library [14] are summarized in Table 1; they are consistent with the running times given in [19].

γ	2^{21}	2^{22}	2^{23}	2^{24}	2^{25}
Time (s)	2.7	6.7	16.4	39.2	89.9
\log_2 cycles	32.4	33.7	35.0	36.3	37.5

Table 1. Running time of a single GCD computation on two random integers of size γ bits, using the Sage library.

The better attack from [12] consists in computing

$$p = \gcd(x_0, \prod_{i=-2^\rho}^{2^\rho-1} (x_1 - i) [x_0])$$

In this attack the $2^{\rho+1}$ GCD computations are replaced by $2^{\rho+1} - 1$ modular multiplications of γ -bit integers and one single final GCD computation. The result of practical experiments with the Sage library [14] are summarized in Table 2.

γ	2^{21}	2^{22}	2^{23}	2^{24}	2^{25}
Time (s)	0.5	1.2	2.8	6.7	15.4
\log_2 cycles	30.0	31.3	32.5	33.7	34.9

Table 2. Running time of a single modular multiplication on two random integers of size γ bits, using the Sage library.

The size γ of the x_i 's depends on η which depends on ρ , so we cannot determine ρ directly. We refer to Table 4 in Section 7.6 for the concrete parameters. From the \log_2 cycles count in Table 2, one can check that the attack requires at least 2^λ clock cycles for each of the four levels of security.

6.2 Approximate-GCD Attack on the Public Key

We do not consider Coppersmith's attack since as shown in [4] it does not apply for the range of parameters. We consider the attack based on Nguyen and Stern's orthogonal lattice [11] (see Section B.1 in [4]). One considers the set of τ integers $x_i = p \cdot q_i + r_i$ and $x_0 = p \cdot q_0$. We consider the first t integers x_i and we consider a vector \mathbf{u} orthogonal to these first t integers x_i modulo x_0 , that is:

$$\sum_{i=1}^t u_i \cdot x_i = 0 \pmod{x_0}$$

This gives

$$\sum_{i=1}^t u_i \cdot r_i = 0 \pmod{p}$$

Now if the u_i 's are sufficiently small, since the r_i 's are small this equality will hold over \mathbb{Z} . This gives a vector \mathbf{u} orthogonal to the r_i 's; by collecting sufficiently many such vectors \mathbf{u} orthogonal to $\mathbf{r} = (r_i)$ one can recover \mathbf{r} and eventually the secret key p . In our analysis we consider that the attack succeeds if at least one vector \mathbf{u} orthogonal to \mathbf{r} in \mathbb{Z} has been obtained.

To find \mathbf{u} we build the lattice L of row vectors orthogonal to $\mathbf{x} = (x_1, \dots, x_t)$ modulo x_0 . From Minkowski's bound there exists a nonzero lattice vector of norm about $\sqrt{t} \cdot \det(L)^{1/t} \simeq 2^{\gamma/t}$. Since to get $\mathbf{u} \cdot \mathbf{r} = 0$ over \mathbb{Z} we must have $\|\mathbf{u}\|_\infty \leq 2^{\eta-\rho} < 2^\eta$ this gives the condition $t > \gamma/\eta$. However when the lattice dimension t is large, lattice reduction algorithms will not recover such a short vector but only an approximation.

As in [4] we use the following ‘‘rule of thumb’’ conjecture about lattice algorithms performance: there exists a constant μ such that for any k and any dimension t , one cannot find a $\mu^{t/k}$ approximation of the shortest vector in time smaller than 2^k . Since we must find a vector \mathbf{u} such that $\|\mathbf{u}\|_\infty \leq 2^{\eta-\rho}$, we need better than a $2^{\eta-\rho}$ approximation of the shortest vector. To get a 2^η approximation (which is not quite enough to recover \mathbf{u}), from $t > \gamma/\eta$ the time required is then at least 2^k where $k = (\log_2 \mu)\gamma/\eta^2$. We recover the asymptotic condition from [4]:

$$\gamma = \eta^2 \cdot \omega(\log \lambda)$$

To derive concrete parameters we have run some experiments with the LLL and BKZ-20 lattice reduction algorithms applied to the following lattice L :

$$L = \begin{bmatrix} 1 & & & & \frac{-x_1}{x_t}[x_0] \\ & 1 & & & \frac{-x_2}{x_t}[x_0] \\ & & \ddots & & \vdots \\ & & & 1 & \frac{-x_{t-1}}{x_t}[x_0] \\ & & & & x_0 \end{bmatrix} \quad (3)$$

We did not run BKZ with block size ≥ 25 since as observed in [5] the running seems to be exponential in the lattice dimension.

Given a “random” lattice L of dimension n without a particularly short vector, one can expect that the shortest vector \mathbf{v} has a norm according to Minkowsky bound $\|\mathbf{v}\| \simeq \sqrt{n} \cdot (\det L)^{1/n}$. With lattice reduction algorithms, one gets a vector \mathbf{b} of norm $\|\mathbf{b}\| \simeq c^n \cdot (\det L)^{1/n}$ where c^n is called the Hermite factor.

We summarize in Table 3 the result of our practical experiments for estimating the Hermite factor c and the running time. Our values for the Hermite factor are approximately the same as in [5]. Experimentally the running time (expressed in number of clock cycles) can be approximated by:

$$T(n, \gamma) = b \cdot n^a \cdot \frac{\gamma}{n} \quad (4)$$

where the experimental coefficients a, b are given in Table 3.

	LLL	BKZ-20
$c = \text{Hermite factor}^{1/n}$	1.021	1.013
a, b	(5.0, 0.06)	(5.2, 0.36)

Table 3. Hermite factor and running time coefficients for LLL and BKZ-20 lattice reduction algorithms, using the SAGE implementation [14].

When applying lattice reduction on the lattice L given by (3), we obtain a short vector $\|\mathbf{u}\|$ of bit-size $\gamma/t + c_2 \cdot t$ instead of γ/t , where $c_2 = \log_2 c$ where the constant c is estimated in Table 3. Since we must have $\|\mathbf{u}\| \leq 2^{\eta-\rho} < 2^\eta$, this gives the condition $\gamma/t + c_2 \cdot t \leq \eta$.

We consider the equation $\eta = c_2 \cdot t + \gamma/t$ or equivalently:

$$c_2 \cdot t^2 - \eta \cdot t + \gamma = 0 \quad (5)$$

Let $\Delta = \eta^2 - 4\gamma c_2$ be the discriminant of this equation. If $\Delta < 0$, then the previous equation has no solution. Therefore given η one could fix γ such that $\Delta < 0$; this would give the condition $\gamma \geq \eta^2/(4c_2)$.

However we wish to use a smaller value for γ . For $\Delta > 0$ the smallest solution of equation (5) is:

$$t_{\min} = \frac{\eta - \sqrt{\eta^2 - 4c_2\gamma}}{2c_2} \quad (6)$$

This is the minimum lattice dimension for which a lattice reduction algorithm with Hermite constant $c_2 = \log_2 c$ can output a vector \mathbf{u} sufficiently short to be orthogonal to \mathbf{r} . From the dimension t_{\min} and γ one can then extrapolate the running time of the lattice reduction algorithm using the estimate in (4). We require that for security parameter λ we have:

$$T(t_{\min}, \gamma) = b \cdot (t_{\min})^{a-1} \cdot \gamma \geq 2^\lambda$$

for both lattice reduction algorithms LLL and BKZ-20. Therefore given η we must fix a large enough γ such that:

$$b \cdot \left(\frac{\eta - \sqrt{\eta^2 - 4c_2\gamma}}{2c_2} \right)^{a-1} \cdot \gamma \geq 2^\lambda$$

for both LLL and BKZ-20, with values $(a, b, c_2 = \log_2 c)$ from Table 3.

Since BKZ-20 outputs shorter vectors the value of t_{\min} is smaller for BKZ-20 than for LLL; however BKZ-20 takes more time than LLL. In our experiments we found that it was more advantageous to run LLL on larger dimensions. We refer to Table 4 in Section 7.6 for the concrete parameters

6.3 Lattice Attack on the Sparse Subset-sum Problem

We use the refined analysis from [18] of the sparse subset sum problem. The attacker must solve the equation:

$$\sum_{i=1}^{\theta} s_i \cdot u_i = x_p \pmod{2^\kappa}$$

where the u_i 's are known and the secret-key $\mathbf{s} = (s_1, \dots, s_\theta)$ is of small Hamming weight θ . We assume that the attack knows p and therefore x_p .

We consider the knapsack lattice of row vectors from [3]:

$$L = \begin{bmatrix} 2 & & & & u_1 \\ & 2 & & & u_2 \\ & & \ddots & & \vdots \\ & & & 2 & u_\theta \\ & & & & 2^\kappa \\ 1 & 1 & \dots & 1 & x_p \end{bmatrix}$$

We have that $(\pm 1, \dots, \pm 1, 0)$ is a short vector of the lattice of norm $\sqrt{\theta}$. The determinant of the lattice is $\det L \simeq 2^{\theta+\kappa} = 2^{\theta+\gamma+2}$. From Minkowsky's bound we can expect that the norm of the second shortest vector is $\simeq (\det L)^{1/\theta} \simeq 2^{\gamma/\theta}$.

Asymptotic analysis. To find the shortest vector we need better than a $2^{\gamma/\theta}$ approximation. From the lattice ‘‘Rule of Thumb’’ conjecture with the previous notations the time required is then at least 2^k with $k = (\log_2 \mu)\theta^2/\gamma$. Asymptotically the condition is therefore:

$$\theta^2 = \gamma \cdot \omega(\log \lambda)$$

Therefore with $\gamma = \tilde{O}(\lambda^5)$ we can take $\theta = \tilde{O}(\lambda^3)$.

Concrete parameters. As observed in [5] the minimum gaps for which LLL or BKZ can retrieve the shortest vector is proportional (up to a small constant) to the corresponding Hermite factor. Since the shortest vector has norm $\sqrt{\theta}$ and the second shortest vector has norm $\simeq 2^{\gamma/\theta}$, we obtain the following condition for finding the shortest vector: $\log_2 \sqrt{\theta} + c_2 \cdot \theta \leq \gamma/\theta$, where $c_2 = \log_2 c$ where the constant c is estimated in Table 3. Therefore to avoid the previous attack one could take the constraint $\theta^2 \geq \gamma/c_2$.

However we observed experimentally that running LLL and BKZ-20 against the knapsack lattice L takes approximately the same amount of time as running LLL and BKZ-20 against the lattice L in section 6.2, for the same lattice dimension $t = \theta$ and the same value of γ . In Section 6.2 the value γ has been determined such that the LLL and BKZ-20 running times both satisfy $T(t_{\min}, \gamma) \geq 2^\lambda$. Note that the value of t_{\min} for LLL is larger than for BKZ-20. Therefore we can take the value of t_{\min} from equation (6) that corresponds to LLL and use the weaker constraint $\theta = t_{\min}$. Then LLL (and therefore BKZ-20) will take at least 2^λ clock cycles on the lattice L of this section.

6.4 Birthday-like Attack on the Sparse Subset-sum Problem

There is a birthday-like exhaustive search attack against the sparse subset sum in the squashed scheme [7]. When using the optimization from [7] which consists in representing the secret key \mathbf{s} in θ boxes of $B = \Theta/\theta$ bits each, such that each box has a single 1-bit in it (see Section 7.2), the attack has complexity $(\Theta/\theta)^{(\theta-1)/2}$. Therefore asymptotically we must have $\theta \log \Theta = \omega(\log \lambda)$ and for concrete parameters the attack complexity must be $\geq 2^\lambda$.

7 Implementation of the Fully Homomorphic Scheme

7.1 Recryption

Now that decryption can be expressed as an arithmetic circuit of low depth, it is possible to achieve bootstrapping, i.e. to publicly evaluate the decryption circuit homomorphically on a ciphertext, which produces another ciphertext for the same message, but with possibly less noise (a “recryption”). This process, which is Gentry’s key idea [6] for achieving fully homomorphic encryption, is illustrated in Figure 1. The procedure that evaluates the decryption circuit homomorphically, called **Recrypt**, takes as input encryptions of the ciphertext bits, and encryptions of the secret key bits.

In the case of the DGHV scheme or of our variant, 0 and 1 are valid encryptions of themselves, so the ciphertext bits can be passed as is to the decryption circuit. However, encryptions of the secret key bits should also be made available publicly, so the key generation from §5.1 should be modified to include encryptions σ_i of the s_i ’s into the public key $pk = (pk^*, \mathbf{y}, \boldsymbol{\sigma})$. Then the **Recrypt** procedure is simply obtained by applying the decryption circuit to the ciphertext bits and the encrypted secret key bits.

Note that such ciphertexts σ_i are normally generated using the $x_{i,b}$ ’s from the public key, leading to σ_i ’s with noise of size ρ' . However since we are at key generation phase we can do better and let $\sigma_i = s_i + 2r'_i + p \cdot q'_i \pmod{x_0}$ for $1 \leq i \leq \Theta$, for random integers q'_i modulo q_0 and random integers r'_i in $(-2^\rho, 2^\rho)$. The resulting ciphertexts σ_i have the same distribution as regular ciphertexts but with noise ρ instead of ρ' . Since $\rho < \rho'$ this enables to reduce the size η of p required to achieve bootstrappability.

For the refreshed ciphertext to decrypt correctly, its noise must be small enough, and in fact small enough that a multiplication operation between refreshed ciphertexts still decrypts correctly. The ciphertext bits are noise-free encryptions of themselves and the encrypted secret key bits contain ρ bits of noise, so one must have $d \cdot \rho < \eta/2$, where d is the degree of the decryption circuit discussed in the previous section (or in fact, only half that degree, because only the degree with respect to the secret key bits matters; the contribution in the ciphertext bits z_i can be ignored as they are used directly and without noise).

7.2 Optimization of the Decryption Circuit

We use the optimization from [4] which consists in representing the secret key \mathbf{s} in θ boxes of $B = \Theta/\theta$ bits each, such that each box has a single 1-bit in it. This enables to obtain a grade-school addition algorithm that requires $\mathcal{O}(\theta^2)$ multiplications of bits instead of $\mathcal{O}(\Theta \cdot \theta)$.

More precisely we denote by $s_{k,i}$ the i -th secret key bit in box k , where $1 \leq k \leq \theta$ and $1 \leq i \leq B$. We use the same subscript notation for the expanded ciphertext with $z_{k,i}$. The decryption equation becomes:

$$m \leftarrow c^* - \left[\sum_{k=1}^{\theta} \left(\sum_{i=1}^B s_{k,i} z_{k,i} \right) \right] \pmod{2}$$

The observation from [7] is that the sum $q_k \stackrel{\text{def}}{=} \sum_{i=1}^B s_{k,i} z_{k,i}$ is obtained by adding B numbers, only one being non-zero. Therefore to compute the j -th bit of q_k it suffices to xor all the j -th bits of the numbers $s_{k,i} \cdot z_{k,i}$. When applying homomorphic decryption, this corresponds to simply adding modulo x_0 all the ciphertexts corresponding to these bits. See Figure 3 for an illustration of the decryption circuit.

The decryption equation is now $m \leftarrow c^* - \left\lfloor \sum_{k=1}^{\theta} q_k \right\rfloor \pmod{2}$ where the q_k 's are rational in $[0, 2)$ with n bits of precision after the binary point. As in [7] it suffices to perform a grade-school addition of the q_k 's. The bits of the q_k 's are arranged in θ rows and $n + 1$ columns (one column before the binary point and n columns after); see Figure 2 for an illustration with $n = 4$. In this case, note that the decryption polynomial is of degree $2^n - 1$ in the q_k 's (the degree of $W_1^{(-1)}$, in the notations of Section 5.2) instead of 2^n as before, because the n -th carry bit of the leftmost column $W_n^{(-n)}$ vanishes. Since we choose θ as a power of 2 minus 1, this means that the decryption polynomial is of degree θ in the secret key bits overall.

7.3 Compression of Encrypted Secret Key Bits

The modification of the public key described previously, to accommodate for the Recrypt procedure, has the problem of increasing public key size significantly. Namely the vector σ in the enlarged public key consists of $\Theta = \tilde{O}(\lambda^3)$ ciphertexts, each of size $\gamma = \tilde{O}(\lambda^5)$, so we obtain a public key size of $\tilde{O}(\lambda^8)$, instead of $\tilde{O}(\lambda^7)$ in the basic scheme.

To alleviate this problem, an additional compression trick is described in [7]. Instead of generating the secret key as a single bit vector $\mathbf{s} = (s_1, \dots, s_{\Theta})$, one uses two bit vectors $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$ of length $\sqrt{\Theta}$, and \mathbf{s} is then recovered on the fly during decryption as the following matrix with Θ entries:

$$s_{i,j} = s_i^{(0)} \cdot s_j^{(1)}$$

The bit vectors $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$ are chosen such that \mathbf{s} is of Hamming weight θ . This can be combined with the division in “boxes” described in §7.2: $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$ can be divided in $\sqrt{\theta}$ boxes of length $\sqrt{B} = \sqrt{\Theta}/\theta$, each of Hamming weight 1, yielding a natural division of \mathbf{s} in θ boxes of length B .

Then, the public key only needs to contain encryptions of the bits of $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$, namely vectors $\sigma^{(0)}$ and $\sigma^{(1)}$ given by

$$\sigma_i^{(b)} = s_i^{(b)} + 2r'_{i,b} + p \cdot q'_{i,b} \pmod{x_0} \quad (1 \leq i \leq \sqrt{\Theta}, b = 0, 1)$$

with $r'_{i,b}$ and $q'_{i,b}$ chosen at random in $(-2^{\rho}, 2^{\rho})$ and $[0, q_0)$ respectively. This makes it possible to recover encryptions σ of the bits of \mathbf{s} on the fly at the time of recryption. This brings down public key size to about $\sqrt{\Theta} \cdot \gamma = \tilde{O}(\lambda^{6.5})$. Note on the other hand that this increases the noise in σ by a factor of 2 since the $\sigma_{i,j}$ are obtained as products of two ciphertexts; this implies that to keep bootstrappability the size η of p must be doubled.

Clearly, it is possible to compress the encrypted secret key bits further by using products of three $\sigma_i^{(b)}$'s or more, but this increases the noise even further. For our purposes, we have found the choice of two vectors $\sigma^{(b)}$ of encrypted secret key bits to be optimal.

7.4 Smaller Dimension for Knapsack Encryption

From the previous section the size of the public key in the full scheme is now about $(\beta + \sqrt{\Theta}) \cdot \gamma$ bits. The conditions from Section 3.2 imply that we must have $\beta = \tilde{O}(\lambda^2)$ to apply the leftover hash lemma. Since $\sqrt{\Theta} = \tilde{O}(\lambda^{1.5})$ we have that β is the bottleneck. Therefore in practice we would like to use a smaller β , for which the leftover hash lemma would not apply but no attack would work.

This implies that we must consider a lattice attack against the knapsack sum in the encryption algorithm. The analysis is the same as in Section 6.3, with $\tau = \beta^2$ instead of Θ . This gives the asymptotic condition $\tau^2 = \gamma \cdot \omega(\log \lambda)$ which for $\alpha < \tau$ is weaker than the condition $\alpha \cdot \tau \geq \gamma + \omega(\log \lambda)$ necessary for the reduction to the approximate GCD problem. Under this condition we can take $\tau = \tilde{O}(\lambda^3)$ instead of $\tau = \tilde{O}(\lambda^4)$ and therefore $\beta = \tilde{O}(\lambda^{1.5})$ instead of $\beta = \tilde{O}(\lambda^2)$. The public key size is then $(\beta + \sqrt{\Theta}) \cdot \gamma = \tilde{O}(\lambda^{6.5})$ instead of $\tilde{O}(\lambda^7)$.

7.5 Fully Homomorphic Scheme complete Description

For completeness we provide a complete description of the fully homomorphic scheme.

KeyGen(1^λ). Generate a random odd integer p of size η bits. Pick a number $q_0 \in [0, 2^\gamma/p)$ chosen as a product of random λ^2 -bit primes, and let $x_0 = q_0 \cdot p$. Generate β pairs $x_{i,0}, x_{i,1}$ of integers with for $1 \leq i \leq \beta$:

$$x_{i,b} = p \cdot q_{i,b} + r_{i,b}, \quad 1 \leq i \leq \beta, \quad 0 \leq b \leq 1 \quad (7)$$

where $q_{i,b}$ are random integers in $[0, q_0)$ and $r_{i,b}$ are integers in $(-2^\rho, 2^\rho)$. Let $pk^* = (x_0, x_{1,0}, x_{1,1}, \dots, x_{\beta,0}, x_{\beta,1})$.

Additionally generate random bit vectors $\mathbf{s}^{(0)}$ and $\mathbf{s}^{(1)}$ of length $\lceil \sqrt{\Theta} \rceil$, subject to the conditions that $s_1^{(0)} = s_1^{(1)} = 1$, that for each $k \in [0, \sqrt{\theta})$ and $b = 0, 1$, there is at most one nonzero bit among the $s_i^{(b)}$'s, $k \lfloor \sqrt{B} \rfloor < i \leq (k+1) \lfloor \sqrt{B} \rfloor$, with $B = \Theta/\theta$, and that $S = \{(i, j) : s_i^{(0)} \cdot s_j^{(1)} = 1\}$ contains exactly θ elements.

Initialize a system-wide pseudo-random number generator f with a random seed \mathbf{se} , and use $f(\mathbf{se})$ to generate integers $u_{i,j} \in [0, 2^{\kappa+1})$ for $1 \leq i, j \leq \lceil \sqrt{\Theta} \rceil$, $(i, j) \neq (1, 1)$. Then, set $u_{1,1}$ such that

$$\sum_{(i,j) \in S} u_{i,j} = x_p \pmod{2^{\kappa+1}}$$

where $x_p \leftarrow \lfloor 2^\kappa/p \rfloor$.

Compute encryptions $\sigma^{(b)}$ of the vectors $\mathbf{s}^{(b)}$, by picking, for each $i \in [1, \lceil \sqrt{\Theta} \rceil]$ and $b = 0, 1$, random integers $r'_{i,b} \in (-2^\rho, 2^\rho)$ and $q'_{i,b} \in [0, q_0)$, and setting:

$$\sigma_i^{(b)} = s_i^{(b)} + 2r'_{i,b} + p \cdot q'_{i,b} \pmod{x_0}$$

Finally, output the secret key as $sk = (\mathbf{s}^{(0)}, \mathbf{s}^{(1)})$, and the public key as $pk = (pk^*, \mathbf{se}, u_{1,1}, \sigma^{(0)}, \sigma^{(1)})$.

Encrypt($pk, m \in \{0, 1\}$). Choose a random integer matrix $\mathbf{b} = (b_{ij})_{1 \leq i, j \leq \beta} \in [0, 2^\alpha)^{\beta \times \beta}$ and a random integer r in $(-2^{\rho'}, 2^{\rho'})$. Output the ciphertext:

$$c^* = m + 2r + 2 \sum_{1 \leq i, j \leq \beta} b_{ij} \cdot x_{i,0} \cdot x_{j,1} \pmod{x_0}$$

Add(pk, c_1^*, c_2^*). Output $c_1^* + c_2^* \pmod{x_0}$.

Mult(pk, c_1^*, c_2^*). Output $c_1^* \cdot c_2^* \pmod{x_0}$.

Expand(pk, c^*). This procedure, ciphertext expansion, takes a ciphertext c^* and computes the associated \mathbf{z} -matrix. We can think of it as separate from either encryption or decryption, as it can be executed publicly given the ciphertext and public data. To do so, for every $1 \leq i, j \leq \sqrt{\Theta}$ first compute the random integer $u_{i,j}$ using the seeded pseudo-random number generator $f(\mathbf{se})$, then let $y_{i,j} = u_{i,j}/2^\kappa$ and compute $z_{i,j}$ given by:

$$z_{i,j} = \lfloor c^* \cdot y_{i,j} \rfloor_2$$

keeping only $n = \lceil \log_2(\theta + 1) \rceil$ bits of precision after the binary point. Define the matrix $\mathbf{z} = (z_{i,j})$. Output the expanded ciphertext $c = (c^*, \mathbf{z})$.

Decrypt(sk, c^*, \mathbf{z}). Output $m \leftarrow \lfloor c^* - \lfloor \sum_{i,j} s_i^{(0)} \cdot s_j^{(1)} \cdot z_{i,j} \rfloor \rfloor_2$.

Recrypt(pk, c^*, \mathbf{z}). Apply the decryption circuit to the expanded ciphertext \mathbf{z} and the encrypted secret key bits $\sigma_i^{(b)}$. Output the result as a refreshed ciphertext c_{new}^* .

7.6 Concrete Parameters

From the analysis of the known attacks in the previous section we are now ready to derive the concrete parameters for the four levels of security. For all four levels we take $\theta = 15$. In this case the degree of the decryption polynomial is $2\theta = 30$ when using the degree-2 compression of the encryption of the secret key bits. Since we must allow for an additional multiplication after Recrypt, the total degree is $d = 4 \cdot \theta = 60$. To allow for some margin we take $\eta = (d + 8)\rho = 68 \cdot \rho$. We obtain the parameters given in Table 4.

Parameters	λ	ρ	η	γ	β	Θ
Toy	42	16	1088	$1.6 \cdot 10^5$	12	144
Small	52	24	1632	$0.86 \cdot 10^6$	23	533
Medium	62	32	2176	$4.2 \cdot 10^6$	44	1972
Large	72	39	2652	$19 \cdot 10^6$	88	7897

Parameters	KeyGen	Encrypt	Expand	Decrypt	Recrypt	pk size
Toy	4.38 s	0.05 s	0.03 s	0.01 s	1.92 s	0.95 MB
Small	36 s	0.79 s	0.46 s	0.01 s	10.5 s	9.6 MB
Medium	5 min 9 s	10 s	8.1 s	0.02 s	1 min 20 s	89 MB
Large	43 min	2 min 57 s	3 min 55 s	0.05 s	14 min 33 s	802 MB

Table 4. Concrete parameters and corresponding timings, as measured using our implementation in Sage 4.5.3 [14] and GMP 4.3.2 [9], on a single core of a desktop computer with an Intel Core2 Duo E8500 CPU at 3.12 GHz. The public key is roughly $2(\beta + \sqrt{\Theta} + 1)\gamma$ bit long. Note that almost all the CPU time of key generation is spent in primality tests, to generate a rough q_0 .

Acknowledgments

We would like to thank Phong Q. Nguyen, Nigel P. Smart and the CRYPTO referees for helpful comments. The work described in this paper has been supported in part by the European Commission through the ICT program under contract ICT-2007-216676 ECRYPT II.

References

1. E. Bach, *How to generate factored random numbers*. SIAM J. Comput., vol. 17, 1988, pp. 179–193.
2. J. Boyar, R. Peralta and D. Pochuev, *On the multiplicative complexity of boolean functions over the basis $(\wedge, \oplus, 1)$* . Theor. Comput. Sci., vol. 235(1), 2000, pp. 43–57.
3. M. J. Coster, A. Joux, B. A. LaMacchia, A. M. Odlyzko, C.-P. Schnorr and J. Stern, *Improved low-density subset sum algorithms*. Computational Complexity, vol. 2, 1992, pp. 111–128.
4. M. van Dijk, C. Gentry, S. Halevi and V. Vaikuntanathan, *Fully homomorphic encryption over the integers*. In H. Gilbert (Ed.), *EUROCRYPT 2010*, LNCS, vol. 6110, Springer, 2010, pp. 24–43.
5. N. Gama and P. Q. Nguyen, *Predicting lattice reduction*. In N. P. Smart (Ed.), *EUROCRYPT 2008*, LNCS, vol. 4965, Springer, 2008, pp. 31–51.
6. C. Gentry, *A fully homomorphic encryption scheme*. Ph.D. thesis, Stanford University, 2009. Available at <http://crypto.stanford.edu/craig>.
7. C. Gentry and S. Halevi, *Implementing Gentry’s fully homomorphic encryption scheme*. In K. Paterson (Ed.), *EUROCRYPT 2011*, LNCS, Springer, 2011.
8. O. Goldreich, S. Goldwasser and S. Halevi, *Public-key cryptosystems from lattice reduction problems*. In B. S. Kaliski (Ed.), *CRYPTO ’97*, LNCS, vol. 1294, Springer, 1997, pp. 112–131.
9. T. Grandlung et al., *The GNU Multiple Precision arithmetic library* (Version 4.3.2), 2010, <http://gmp.lib.org>.
10. R. Lidl and H. Niederreiter, *Finite fields*. Encyclopedia of mathematics and its applications, vol. 20, Addison-Wesley, 1983.
11. P. Q. Nguyen and J. Stern, *The two faces of lattices in cryptology*. In J. H. Silverman (Ed.), *CaLC 2001*, LNCS, vol. 2146, Springer, pp. 146–180.
12. P. Q. Nguyen, personal communication.

13. P. Paillier, *Public-key cryptosystems based on composite degree residuosity classes*. In J. Stern (Ed.), *EUROCRYPT '99*, LNCS, vol. 1592, Springer, 1999, pp. 223–238.
14. W.A. Stein et al., *Sage Mathematics Software (Version 4.5.3)*, The Sage Development Team, 2010, <http://www.sagemath.org>.
15. D. Micciancio, *Improving lattice based cryptosystems using the Hermite normal form*. In J. H. Silverman (Ed.), *CaLC 2001*, LNCS, vol. 2146, Springer, pp. 126–145
16. X. Pujol, D. Stehlé et al., *fp111 lattice reduction library*, <http://perso.ens-lyon.fr/xavier.pujol/fp111/>.
17. N. P. Smart and F. Vercauteren, *Fully homomorphic encryption with relatively small key and ciphertext sizes*. In P. Q. Nguyen and D. Pointcheval (Eds.), *PKC 2010*, LNCS, vol. 6056, Springer, 2010, pp. 420–443.
18. D. Stehlé and R. Steinfeld, *Faster fully homomorphic encryption*. In M. Abe (Ed.), *ASIACRYPT 2010*, LNCS, vol. 6477, Springer, 2010, pp. 377–394.
19. D. Stehlé and P. Zimmermann, *A binary recursive GCD algorithm*. In D. A. Buell, *ANTS-VI*, LNCS, vol. 3076, Springer, 2004, pp. 411–425.
20. M.N. Wegman and J.L. Carter, *New hash functions and their use in authentication and set equality*, *Journal of Computer and System Sciences*, vol. 22(3), 1981, pp. 265–279.

A Definition

We recall the definition from [4, 6]. A homomorphic public-key encryption scheme \mathcal{E} has an additional algorithm `Evaluate` which takes as input the public key pk , a circuit C with t input bits and t ciphertexts c_i , and outputs another ciphertext c .

Definition A.1 (Correct Homomorphic Decryption). *The scheme $\mathcal{E} = (\text{KeyGen}, \text{Encrypt}, \text{Decrypt}, \text{Evaluate})$ is correct for a given t -input circuit C if, for any key-pair (sk, pk) output by $\text{KeyGen}(\lambda)$, any t plaintext bits m_1, \dots, m_t and any ciphertexts $\mathbf{c} = (c_1, \dots, c_t)$ with $c_i \leftarrow \text{Encrypt}(pk, m_i)$, it holds that $\text{Decrypt}(sk, \text{Evaluate}(pk, C, \mathbf{c})) = C(m_1, \dots, m_t)$.*

B Proof of Lemma 3.1

Given a ciphertext c output by $\text{Encrypt}(pk, m)$:

$$c = m + 2r + 2 \sum_{1 \leq i, j \leq \beta} b_{i,j} \cdot x_{i,0} \cdot x_{j,1} \pmod{x_0}$$

where $|r| < 2^{\rho'}$ and $b_{i,j} \in [0, 2^\alpha)$, we have:

$$c = m + 2r + 2 \sum_{1 \leq i, j \leq \beta} b_{i,j} \cdot r_{i,0} \cdot r_{j,1} \pmod{p}$$

This gives using $\rho' \geq 2 \cdot \rho + \alpha$:

$$|c \pmod{p}| \leq 2^{\rho'+1} + 2\tau \cdot 2^{2\rho+\alpha} \leq \tau \cdot 2^{\rho'+2} \quad (8)$$

Let $C \in \mathcal{C}_{\mathcal{E}}$ be a permitted circuit with t inputs and let C' be the corresponding circuit operating over the integers rather than modulo 2. Let $c_i \leftarrow \text{Encrypt}(pk, m_i)$. We have:

$$c \pmod{p} = C'(c_1, \dots, c_t) \pmod{p} = C'(c_1 \pmod{p}, \dots, c_t \pmod{p}) \pmod{p} \quad (9)$$

From (8) and the definition of permitted circuits, we obtain:

$$|C'(c_1 \pmod{p}, \dots, c_t \pmod{p})| \leq 2^{\eta-4} \leq p/8$$

Therefore $C'(c_1 \pmod{p}, \dots, c_t \pmod{p}) \pmod{p} = C'(c_1 \pmod{p}, \dots, c_t \pmod{p})$, which implies from (9):

$$c \pmod{p} = C'(c_1 \pmod{p}, \dots, c_t \pmod{p})$$

and eventually:

$$[c \pmod{p}]_2 = [C'(c_1 \pmod{p}, \dots, c_t \pmod{p})]_2 = [C'([c_1 \pmod{p}]_2, \dots, [c_t \pmod{p}]_2)]_2$$

which gives $[c \pmod{p}]_2 = C(m_1, \dots, m_t)$ which concludes the proof.

C Proof of Corollary 4.1

Again, given $\mathbf{b} \neq \mathbf{b}'$, we have to estimate the number of zeros N of the quadratic form $\sum_{1 \leq i, j \leq \beta} a_{ij} u_i v_j \pmod q$ where $a_{ij} = b_i - b'_i$.

Now, write q as the product of primes $p_1 \cdots p_k$ with $p_1 < \cdots < p_k$. Then the Chinese remainder theorem ensures that $N = N_1 \cdots N_k$ where N_ℓ is the number of zeros of the quadratic form mod p_ℓ . As we have previously seen, this satisfies:

$$N_\ell \leq p_\ell^{2\beta-1} + p_\ell^{2\beta-r_\ell}$$

where r_ℓ is the rank of the matrix $A = (a_{ij}) \pmod{p_\ell}$. In particular, if the reduction of $A \pmod{p_\ell}$ is of rank ≥ 2 for all ℓ , we obtain:

$$\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] = \frac{N_1 \cdots N_k}{q^{2\beta}} \leq \prod_\ell \left(\frac{1}{p_\ell} + \frac{1}{p_\ell^2} \right) \leq \frac{1}{q} \left(1 + \frac{1}{p} \right)^k \leq \frac{1}{q} \left(1 + \frac{e \cdot k}{p} \right)$$

where $p = p_1$ is the smallest prime factor of q (and $e = \exp(1)$).

On the other hand, the bound is less convenient if $r_\ell \leq 1$ for some ℓ . Note that $r_\ell = 0$ cannot happen, since in view of the condition $p_\ell > 2^\alpha$, $a_{ij} = 0 \pmod{p_\ell}$ implies $b_{ij} = b'_{ij}$. So we only have to consider the case of rank 1. If $r_\ell = 1$ for some ℓ , we have:

$$\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] \leq \prod_\ell \frac{2}{p_\ell} = \frac{2^k}{q}$$

Furthermore, the set $U_\alpha^{(\ell)}$ of matrices of rank 1 in $\mathbf{M}_\beta(\mathbb{Z}_{p_\ell})$ with entries in $\{-2^\alpha + 1, \dots, 2^\alpha - 1\}$ is bounded as before: $|U_\alpha^{(\ell)}| < \beta^2 \cdot 2^{2(\alpha+1)\beta}$ for all ℓ . Therefore, the set U_α matrices in $\mathbf{M}_\beta(\mathbb{Z}_q)$ with entries in the proper interval and a rank 1 reduction mod p_ℓ for some ℓ contains at most k times that many matrices:

$$|U_\alpha| < k \cdot \beta^2 \cdot 2^{2(\alpha+1)\beta}$$

Again, we can then bound the value δ given by:

$$\delta = \frac{|Y|}{|X|^2} \sum_{\mathbf{b} \neq \mathbf{b}'} \left(\Pr_h[h(\mathbf{b}) = h(\mathbf{b}')] - \frac{1}{|Y|} \right)$$

as follows:

$$\begin{aligned} \delta &\leq \frac{q}{|X|^2} \left(\sum_{\substack{\mathbf{b} \neq \mathbf{b}' \\ A \notin U_\alpha}} \frac{e \cdot k}{pq} + \sum_{\substack{\mathbf{b} \neq \mathbf{b}' \\ A \in U_\alpha}} \frac{2^k - 1}{q} \right) \leq \frac{q}{|X|^2} \left(\frac{|X|^2}{q} \cdot \frac{e \cdot k}{p} + \frac{|X| \cdot |U_\alpha| \cdot 2^k}{q} \right) \\ &\leq \frac{e \cdot k}{p} + \frac{|U_\alpha| \cdot 2^k}{|X|} \leq k \left(\frac{e}{p} + \frac{\beta^2}{2^{\alpha\beta^2 - 2(\alpha+1)\beta - k}} \right) \end{aligned}$$

It remains to observe that $k \leq (\log q)/(\log p)$ to conclude. \square

D Proof of Theorem 4.1

We follow the same proof strategy as in [4]. Consider an attacker \mathcal{A} against our encryption scheme with non-negligible advantage ε for parameters $(\rho, \rho', \eta, \gamma, \tau)$. In other words, \mathcal{A} takes as input a public key and a ciphertext (produced by **KeyGen** and **Encrypt** of our variant), and outputs the correct plaintext bit with probability at least $1/2 + \varepsilon$. Using \mathcal{A} we mimic the method from [4] to construct an algorithm \mathcal{B} to solve the (ρ, η, γ) -error-free-approximate-GCD problem.

\mathcal{B} is given a γ -bit number of the form $p \cdot q_0$, where p is a random η -bit odd integer and q_0 is a random integer from the interval $[0, 2^\gamma/p)$. In addition, \mathcal{B} is allowed to sample $\mathcal{D}'_\rho(p, q_0)$ polynomially many times. Its goal is to recover p with non-negligible probability.

At the cost of a loss in success probability at most polynomial in λ , we will assume that q_0 is a 2^λ -rough square free integer, as this property is satisfied by a polynomial fraction of integers in $[0, 2^\gamma/p)$.

The algorithm \mathcal{B} is then as follows.

Step 1: Creating a public Key Let $x_0 = p \cdot q_0$ as given to \mathcal{B} . We assume that q_0 is a randomly chosen integer in the interval $[0, 2^\gamma/p)$, as well as square-free and 2^λ -rough. We also sample $x_{i,0}, x_{i,1} \leftarrow \mathcal{D}'_\rho(p, q_0)$ for $1 \leq i \leq \beta$. Clearly, $pk = (x_0, x_{1,0}, x_{1,1} \cdots, x_{\beta,0}, x_{\beta,1})$ follows the same distribution as the public key generated by KeyGen.

Step 2: High Accuracy LSB predictor For any integer z , let $q_p(z)$ and $r_p(z)$ denote the quotient and remainder of z with respect to p . Given any $z \in [0, 2^\gamma)$ that is at most 2^ρ away from a multiple of p , the following subroutine tells us the least significant bit of $q_p(z)$ with high probability.

Learn-LSB(z, pk):

Input: $z \in [0, 2^\gamma)$ with $r_p(z) < 2^\rho$, a public key $pk = (x_0, x_{1,0}, x_{1,1} \cdots, x_{\beta,0}, x_{\beta,1})$

Output: The least significant bit of $q_p(z)$

1. For $k = 1$ to $\text{poly}(\lambda)/\varepsilon$ do:
 2. Choose noise $r_k \leftarrow (-2^{\rho'}, 2^{\rho'})$, a bit $m_k \leftarrow \{0, 1\}$ and $b_{i,j}^{(k)} \leftarrow [0, 2^\alpha)$ for $1 \leq i, j \leq \beta$.
 3. Set $c_k = [z + m_k + 2r_k + 2 \sum_{i,j} b_{i,j}^{(k)} \cdot x_{i,0} \cdot x_{j,1}]_{x_0}$
 4. Call \mathcal{A} to get a prediction of $a_k \leftarrow \mathcal{A}(pk, c_k)$
 5. Set $b_k = a_k \oplus \text{parity}(z) \oplus m_k$
6. Output the majority vote among the b_k 's.

Step 3: Binary GCD Using *Learn-LSB*(z, pk) we build a binary GCD algorithm $GCD(z_1, z_2)$ as in [4], to find out the odd part of $\text{gcd}(q_p(z_1), q_p(z_2))$ for any $z_1 = q_p(z_1)p + r_p(z_1)$ and $z_2 = q_p(z_2)p + r_p(z_2)$ where $r_p(z_i) \ll p$.

Step 4: Recovery of p To recover p , the solver \mathcal{B} draws $z_1^*, z_2^* \leftarrow \mathcal{D}_{\gamma,\rho}(p)$. With probability at least $\zeta(2)^{-1} = \frac{6}{\pi^2} \approx 0.61$, we will have coprime $q_p(z_1^*)$ and $q_p(z_2^*)$. Hence $GCD(z_1^*, z_2^*)$ will return $\tilde{z} = 1 \cdot p + r$ with $|r| < 2^\rho$ with high probability. Once we have \tilde{z} , we call $GCD(z_1^*, \tilde{z})$ to find out $q_p(z_1)$. Now \mathcal{B} answers $p = \lfloor z_1^*/q_p(z_1) \rfloor$.

This completes the description of algorithm \mathcal{B} . It remains to show that *Learn-LSB*(z, pk) is a reliable oracle for $[q_p(z)]_2$. As in [4], this follows from the fact that the distribution of the ‘‘ciphertext’’ c_k at line 3 is statistically close to the distribution of a valid encryption of the bit $[r_p(z)]_2 \oplus m_k$, for all but a negligible fraction of the public keys; see Lemma D.1 below. Then as shown in [4] if \mathcal{A} has advantage ε in guessing the encrypted bit under pk , then for a fraction at least $\varepsilon/2$ of the primes $p \in [2^{\eta-1}, 2^\eta)$ the adversary has advantage at least $\varepsilon/2$, and for a fixed such p the adversary has advantage at least $\varepsilon/4$ for a fraction at least $\varepsilon/4$ of the corresponding public keys. From Lemma D.1 below this gives an advantage at least $\varepsilon/4 - \text{negl}$ for \mathcal{A} at step 4 of *Learn-LSB*; then from the majority vote *Learn-LSB*(z, pk) will return the correct answer with overwhelming probability, and \mathcal{B} will recover p . For such p 's this holds for a fraction at least $\varepsilon/4 - \text{negl}$ of the public keys. Thus for such p 's it suffices to repeat algorithm \mathcal{B} with a new random public key for $4/\varepsilon \cdot \omega(\log \lambda)$ times to recover p with overwhelming probability. The overall success probability of \mathcal{B} is therefore the corresponding fraction of such p 's, which is at least $\varepsilon/2$. This completes the proof of Theorem 4.1. \square

Lemma D.1. Fix the parameters $(\rho, \rho', \eta, \gamma, \tau)$ and the secret key $sk = p$. Let $pk = (x_0, x_{1,0}, x_{1,1}, \dots, x_{\beta,0}, x_{\beta,1})$ be generated at random as in the KeyGen. For every integer $x^* \in [0, 2^\gamma)$ that is at most 2^ρ away from a multiple of p , consider the following distribution

$$\mathcal{C}_{pk}(x^*) = \left\{ \text{For } 1 \leq i, j \leq \beta, b_{i,j} \leftarrow [0, 2^\alpha), r \leftarrow (-2^{\rho'}, 2^{\rho'}) : \text{output } c' = [x^* + 2r + 2 \cdot \sum_{i,j} b_{i,j} \cdot x_{i,0} \cdot x_{j,1}]_{x_0} \right\}.$$

Then with overwhelming probability (over the choice of sk, pk), every distribution $\mathcal{C}_{pk}(x^*)$ is statistically close to the distribution $\text{Encrypt}(pk, m = [x^*]_2)$.

Proof. The proof is the same as in Lemma 4.3 of [4], except that we use our generalized version of Leftover hash lemma (Lemma 4.1) and Corollary 4.1. Writing $c' = q'p + 2r' + m$, the idea is to show separately that q' and r' are distributed as in the scheme.

Regarding the noise r' , since $\rho' = 2\rho + \alpha + \omega(\log \lambda)$, the additional noise r added in the distribution is drawn from a set of cardinality superpolynomially larger than the noise in the public key elements and in the integer x^* . This ensures that the distribution of r' is statistically close to the uniform distribution in $(-2^{\rho'}, 2^{\rho'})$, as in regular encryption.

On the other hand, with our choice of parameters (in particular $\alpha\beta^2 = \gamma + \omega(\log \lambda)$) and with q_0 being a 2^λ -rough integer, Corollary 4.1, together with our generalization of the leftover hash lemma (Lemma 4.1), implies that the statistical distance between the distribution of $q_p(c) = \sum_{i,j} b_{i,j} \cdot q_{i,0} \cdot q_{j,1} \bmod q_0$ for a regular ciphertext c and the uniform distribution in \mathbb{Z}_{q_0} is negligible. But q' is the sum mod q_0 of such a value $q_p(c)$ with a random number, so its distribution is also statistically close to uniform.

Thus, the distributions of both q' and r' are statistically close to the distributions of $q_p(c)$ and $r_p(c)$ for $c = \text{Encrypt}(pk, m)$, which concludes the proof. \square

E Proof of Lemma 5.1

Fix a permitted polynomial $P(x_1, \dots, x_t) \in \mathcal{P}_E$, an arithmetic circuit C that computes P , and t fresh ciphertexts c_1, \dots, c_t that encrypt the input to C . Then, let $c^* = \text{Evaluate}(pk, C, c_1, \dots, c_t)$.

As recalled in §3, the correct decrypted message m is given by $[c^* - [c^*/p]]_2$, so what we need to show is that $[c^*/p] = [\sum_i s_i z_i] \bmod 2$.

To see this, recall first from Remark 3.2 that c^*/p is at most $1/(4(\theta + 1))$ away from the nearest integer $[c^*/p]$. On the other hand, recall that each z_i is computed as $[c^* \cdot y_i]_2$ keeping only n bits of precision after the binary point, so that we can write $z_i = [c^* \cdot y_i]_2 + \delta_i$ with $|\delta_i| < 2^{-(n+1)} \leq 1/(2(\theta + 1))$. Thus, we get, as in [4]:

$$\begin{aligned} \left[(c^*/p) - \sum s_i z_i \right]_2 &= \left[(c^*/p) - \sum s_i [c^* \cdot y_i]_2 - \sum s_i \delta_i \right]_2 \\ &= \left[(c^*/p) - c^* \cdot \left[\sum s_i y_i \right] - \sum s_i \delta_i \right]_2 \\ &= \left[(c^*/p) - c^* \cdot (1/p - \Delta_p) - \sum s_i \delta_i \right]_2 \\ &= \left[c^* \cdot \Delta_p - \sum s_i \delta_i \right]_2 \end{aligned}$$

This allows us to bound the difference mod 2:

$$\left| \left[(c^*/p) - \sum s_i z_i \right]_2 \right| \leq |c^* \cdot \Delta_p| + \left| \sum s_i \delta_i \right| < 2^{\gamma-\kappa} + \theta \cdot \frac{1}{2(\theta + 1)}$$

Therefore, the distance mod 2 of $\sum s_i z_i$ to $[c^*/p]$ is strictly less than $1/(4(\theta + 1)) + 2^{\gamma-\kappa} + \theta/(2(\theta + 1))$. Now $2^{\gamma-\kappa} = 2^{-n-2} \leq 1/(4(\theta + 1))$, so the total distance is strictly less than $1/2$. This concludes the proof. \square

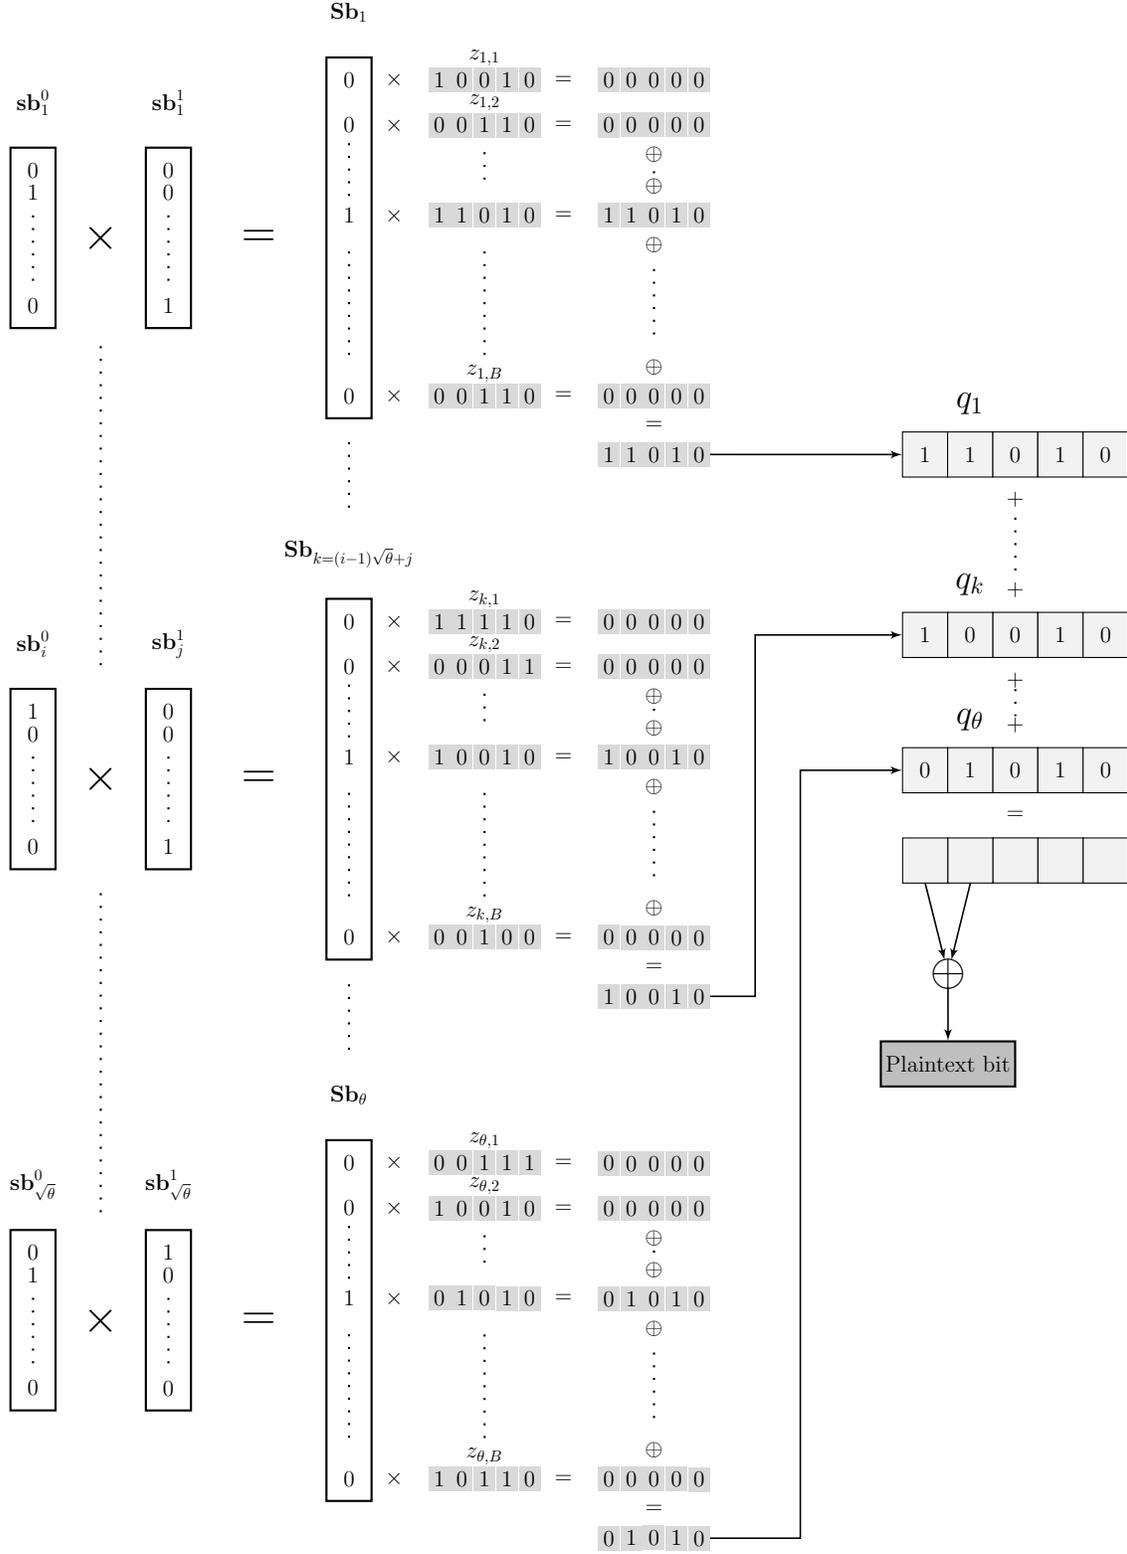


Fig. 3. Illustration of the decryption Circuit. $\mathbf{s}^{(0)}$, $\mathbf{s}^{(1)}$ are divided into $\sqrt{\theta}$ boxes of Hamming weight 1, namely $\mathbf{sb}_1^0, \dots, \mathbf{sb}_{\sqrt{\theta}}^0$ and $\mathbf{sb}_1^1, \dots, \mathbf{sb}_{\sqrt{\theta}}^1$. As a result \mathbf{s} gets divided into θ boxes of Hamming weight 1, namely $\mathbf{Sb}_1, \dots, \mathbf{Sb}_\theta$.