Fully Homomophic Encryption over the Integers Revisited

Jung Hee Cheon[1] and Damien Stehlé[2]

[1] SNU, Republic of Korea, [2] ENS de Lyon, France.
jhcheon@snu.ac.kr, damien.stehle@ens-lyon.fr

Abstract. Two main computational problems serve as security foundations of current fully homomorphic encryption schemes: Regev's Learning With Errors problem (LWE) and Howgrave-Graham's Approximate Greatest Common Divisor problem (AGCD). Our first contribution is a reduction from LWE to AGCD. As a second contribution, we describe a new AGCD-based fully homomorphic encryption scheme, which outperforms all prior AGCD-based proposals: its security does not rely on the presumed hardness of the so-called Sparse Subset Sum problem, and the bit-length of a ciphertext is only $\widetilde{O}(\lambda)$, where λ refers to the security parameter.

Keywords. Fully homomorphic encryption, approximate gcd, DGHV, LWE.

1 Introduction

Fully homomorphic encryption has been a major focus of interest in cryptography since Gentry's first proposal of a fully homomorphic encryption scheme [22, 21]. The security of Gentry's proposal relies on two hardness assumptions: some relatively ad-hoc problem involving lattices arising in algebraic number theory is assumed intractable, as is the Sparse Subset Sum Problem (SSSP), a variant of the subset sum problem in which the subset is constrained to be very small. The efficiency of Gentry's scheme was later improved [45, 46, 23], but soon two other design approaches were developed. The interest in Gentry's original design faded, as the latter approaches rely on better understood hardness assumptions and lead to more efficient instantiations.

Chronologically, the first alternative design was proposed by van Dijk, Gentry, Halevi and Vaikuntanathan [20]. They constructed a fully homomorphic scheme whose security relies on the hardness of SSSP as well as that of the Approximate Greatest Common Divisor problem (AGCD). The AGCD problem, introduced by Howgrave-Graham in [28], is to recover a secret integer p from many approximate multiples $q_i \cdot p + r_i$ of p (see Section 2 for a formal definition). The efficiency of the DGHV scheme has been improved in a series of works [18, 19, 29, 14, 17], and recently adapted to devise a graded encoding scheme serving as an approximation to cryptographic multilinear maps [16].

The other main family of fully homomorphic schemes was initiated by Brakerski and Vaikuntanathan, in [9, 10]. They proposed two fully homomorphic schemes with similar designs. One was relying on the hardness of the Learning With Errors problem (LWE) from [40, 41] while the other used the less understood Ring Learning With Errors problem from [31] to gain on the efficiency front. Note that in both cases, the SSSP hardness assumption is not required anymore. A series of subsequent works proposed efficiency and security improvements, as well as implementations [25, 24, 5, 6, 27, 12, 3].

The co-existence of these two main design strategies for fully homomorphic encryption is due to the combination of circumstances. On one hand, it is not known how the underlying hardness assumptions compare: there is no known reduction from AGCD and SSSP to LWE (or its ring variants), and reciprocally. On the other hand, both approaches seem to lead to implementations whose performances are relatively comparable.

Contributions. This work contains two main results that together lead to a better understanding of the relationship between the AGCD-based and LWE-based fully homomorphic encryption schemes.

Our first contribution is a reduction from LWE to a new and quite natural decision variant of AGCD. Informally, the goal is to distinguish between random approximate multiples $q_i p + r_i$ of a random p and integers uniformly chosen in an interval, with non-negligible distinguishing advantage and non-negligible probability over the choice of p. This AGCD variant is clearly no easier than the search variant considered in [20]. Our reduction implies that for certain distributions for p, the r_i's and the q_i's, AGCD is no easier than LWE. It may be combined with Regev's quantum reduction [40, 41] from the approximate variant of the shortest independent vectors problem (SIVP_γ) to LWE. Concretely, if we assume that SIVP_γ in dimension n with $\gamma = \mathrm{poly}(n)$ is exponentially hard to solve (quantumly) with respect to n, which is compatible with the state of the art algorithms for SIVP (see [34]), then AGCD is also exponentially hard to solve, even for bit-sizes of p, r_i, q_i that are quasi-linear in n.

Our second contribution is a fully homomorphic encryption scheme with security based on the hardness of our AGCD variant. In particular, the security does not rely on the presumed hardness of SSSP.[1] The scheme is a variant of the DGHV encryption scheme that embeds the plaintext message in the most significant bit modulo p of an AGCD sample: a ciphertext c corresponding to a plaintext m is of the form $c = qp + \lfloor p/2 \rceil m + r$. Parameters may be set so that security relies on the quantum hardness of SIVP_γ in dimension n with $\gamma = n^{O(\log n)}$, while the secret key, public/evaluation key and ciphertext expansion remain bounded as $\widetilde{O}(\lambda)$, $\widetilde{O}(\lambda^3)$ and $\widetilde{O}(\lambda)$, where λ is such that all known attacks require time $2^{\Omega(\lambda)}$.

Compared to DGHV, we obtain improved asymptotic efficiency and security solely relying on the hardness of LWE. The security and performance are quite similar to those of Brakerski's LWE-based scheme [5]. This is no coincidence, as both contributions build upon ideas from Brakerski's work.

Technical overview. The reduction from LWE to AGCD relies on several sub-reductions. We use the dimension-modulus trade-off for LWE from [8] and start from 1-dimensional LWE with an exponential modulus q. Informally, the goal is to distinguish from uniform the distribution $(a, a \cdot s + e) \in (\frac{1}{q}\mathbb{Z})/\mathbb{Z} \times \mathbb{R}/\mathbb{Z}$ with a uniform, e small and s an integer. We reduce this variant of LWE to another one where a is instead uniformly sampled in $\mathbb{R}/\mathbb{Z}$. The computational irrelevance of the discretization parameter q is implicit in [5, 8]: we go one small step further by simply removing it. We then reduce this one-dimensional scale-invariant variant of LWE to the problem considered by Regev in [38] and inspired from [1]. The problem consists in distinguishing from uniform samples of the form $(k + e)/s \in \mathbb{R}/\mathbb{Z}$, where k is uniformly sampled in $[0, s)$ and e is a small noise. A converse reduction was sketched in the appendix of [42], and our reduction was sketched by Oded Regev in a private communication [43]. We formalize the latter reduction. Our chain of reductions improves over the result of [38] in that for comparable hardness assumptions our reduction allows to take a bitsize for s that is the square root of that allowed by [38]. Finally, we scale and re-discretize samples $(k+e)/s \in \mathbb{R}/\mathbb{Z}$ to obtain a reduction from the latter problem to a decision variant of AGCD.

Our encryption scheme is inspired from that of [38] and the LWE-based Brakerski's fully homomorphic encryption scheme [5]. It is scale-invariant in the sense that it remains unchanged if we multiply both the secret key p and the ciphertext by the same quantity. It does not use a hidden greatest common divisor that is a square as in the Coron *et al.* scale invari-

[1] We still require a circular security assumption, like all known fully homomorphic encryption schemes.

ant version of the DGHV scheme [17]. Homomorphic addition comes without extra work. For homomorphic multiplication, we adapt the dimension-reduction technique from [9], that uses (invalid) encryptions of the bits of secret p (we assume that it is safe to publish these data, hence making a circular-security assumption). Finally, we bound the multiplicative depth of the decryption circuit is bounded as $O(\log \lambda)$ where λ refers to the security parameter. As the parameters may be set so that our homomorphic scheme supports this multiplicative depth, it is hence possible to bootstrap it [22], leading to a fully homomorphic encryption scheme. This allows us to circumvent the SSSP hardness assumption made in prior variants of the DGHV encryption scheme.

Finally, we propose a modification of our scheme in which the ciphertext bit-size is reduced. This is achieved by truncating the least significant bits of the ciphertext, which is made possible by the fact that the plaintext is not embedded into these. As a result, the ciphertext size is almost as low as $\gamma - \rho$, where γ is the bit-length of the AGCD samples and ρ is the bit-length of the AGCD noise. We remark that one can additionally use the technique of [19] to compress the public key.

Open problems. Our results show that the DGHV fully homomorphic encryption scheme [20] can be made to fit into the LWE landscape. The modified scheme asymptotically outperforms DGHV (and all subsequent variants), but its performance only matches Brakerski's LWE-based scheme [5]. Further, there exist recent LWE-based fully homomorphic encryption schemes with strengthened security [27, 11, 3],[2] and efficiency can be increased if one relies on the ring variant of LWE problem.

With this state of affairs, it may be tempting to drop the AGCD approach altogether. We prefer a more optimistic interpretation of our work. First, it gives greater confidence into the hardness of AGCD and simplifies AGCD-based encryption. This clearer landscape could serve as a firmer grounding for further developments. The AGCD problem can be seen as another way of expressing LWE: it may turn out to be more convenient for cryptographic design. Finally, the analogy does not seem to be complete: some variants of DGHV rely on a modification of AGCD in which a noiseless multiple of the secret integer p is published [18] (the security of these variants relies on an extra hardness assumption related to factoring). We are not aware of a similar problem in the LWE landscape.

Our scheme is relatively slow (compared to those based on Ring LWE), but several existing techniques could be exploited to accelerate it. For instance, it may be possible to pack more plaintexts into a single ciphertext, similarly to [14]. It may also be possible to refine the bootstrapping step rather than looking at decryption as a generic binary circuit. Finally, our variant with truncated ciphertexts raises the question of taking AGCD instances with small $(\gamma - \rho)$. To thwart attacks based on exhaustive search, we should have $\gamma - \rho \geq \lambda + \Omega(\log \lambda)$. If $\gamma - \rho \approx \lambda + \Omega(\log \lambda)$ turns out to be safe, then the ciphertext bit-sizes of our variant scheme based on truncation can be made quite small.

Road-map. We describe our LWE to AGCD reduction in Section 2. In Section 3, we describe our AGCD-based scheme, and we show in Section 4 how it may be extended into a fully homomorphic encryption scheme. Section 5 contains a modification of the scheme with smaller ciphertexts.

[2] Note that it may be possible to adapt these techniques to the AGCD framework. A DGHV variant was proposed in appendix of [27].

Notation. We use standard Landau notations. When manipulating reals, we in fact manipulate finite-precision approximations, with polynomially many bits of precision. If x is a real, then $\lfloor x \rceil$ refers to the nearest integer to x, rounding upwards in case of a tie. The notation log refers to the base-2 logarithm. We use the notation $(a_i)_{i=1,\ldots,k}$ or simply $(a_i)_i$ for a vector $(a_1, \ldots, a_k)$. Given $x, p \in \mathbb{R}$, we let $[x]_p$ denote the unique number in $(-p/2, p/2]$ that is congruent to x modulo p. The notation is extended to vectors $\boldsymbol{x} \in \mathbb{R}^n$ in the obvious way. We let $\mathbb{T}$ denote the torus $\mathbb{R}/\mathbb{Z}$. For an integer $q \geq 1$, we let $\mathbb{T}_q$ denote the set $\{0, 1/q, \ldots, (q-1)/q\}$ with addition modulo 1.

We use $a \leftarrow A$ to denote the operation of uniformly sampling an element a from a finite set A. When $\mathcal{D}$ is a distribution, the notation $a \leftarrow \mathcal{D}$ refers to sampling a according to distribution $\mathcal{D}$. We recall that the statistical distance between two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ with supports contained in a common measurable set is half the ℓ_1-norm of their difference.

If X is a set of finite weight, we let $U(X)$ denote the uniform distribution over X. For a parameter $s > 0$, we let D_s denote the (continuous) Gaussian distribution of parameter s, i.e., the law over $\mathbb{R}$ with density function $x \mapsto \exp(-\pi x^2/s^2)/s$. We write $D_{\leq s}$ to refer to a $D_{s'}$ for some $s' \leq s$. If $\Lambda \subseteq \mathbb{R}^n$ is a full-rank lattice, $s > 0$ and $\boldsymbol{c} \in \mathbb{R}^n$, we let $D_{\Lambda,s,\boldsymbol{c}}$ denote the (discrete) Gaussian distribution with support Λ and density function $\boldsymbol{x} \mapsto \exp(-\pi\|\boldsymbol{x} - \boldsymbol{c}\|^2/s^2)/C$ with $C = \sum_{\boldsymbol{x} \in \Lambda} \exp(-\pi\|\boldsymbol{x} - \boldsymbol{c}\|^2/s^2)$. When $\boldsymbol{c} = \boldsymbol{0}$, we omit the last subscript. We recall a few properties of Gaussians in Appendix A.

We say that a distribution $\mathcal{D}$ over $\mathbb{Z}^n$ is (B, ε)-bounded if $\Pr_{x \leftarrow \mathcal{D}}[\|x\| \leq B]] \geq 1 - \varepsilon$. We say that $\mathcal{D}$ is (B, δ, ε)-contained if $\Pr_{x \leftarrow \mathcal{D}}[\|x\| \in [\delta B, B]] \geq 1 - \varepsilon$. For example, for all $\varepsilon \in (0, 1/2)$, the distribution $D_{\mathbb{Z}^n, r}$ is (B, ε)-bounded, with $B = O(r\sqrt{n \ln(n/\varepsilon)})$ (see [32]). If $r = \Omega(\sqrt{\ln(1/\varepsilon)})$, then the distribution $D_{\mathbb{Z}, r}$ is (B, δ, ε)-contained, with $B = O(r\sqrt{\ln(1/\varepsilon)})$ and $\delta = \varepsilon/\sqrt{\ln(1/\varepsilon)}$.

Throughout the paper, we let λ denote the security parameter: all known valid attacks against the cryptographic scheme under scope should require $2^{\Omega(\lambda)}$ bit operations to mount.

2 Hardness of Approximate GCD

We exhibit a reduction from the Learning With Errors problem (LWE) to a variant of the Approximate Greatest Common Divisor problem (AGCD). We first introduce the precise problems under scope.

We will consider the following decision variant of AGCD. The corresponding search variant (consisting in finding the unknown p) is frequent in the literature. There exists a (trivial) reduction from the search variant to the decision variant. Other decision variants of AGCD were considered in [19, 29, 17]. We believe that our decision variant of AGCD is more natural as it is less application-driven.

Definition 1 (AGCD). *Let $p, X \geq 1$, and ϕ a distribution over $\mathbb{Z}$ (that can depend on p). We define $A^{\mathrm{AGCD}}_{X,\phi}(p)$ as the distribution over $\mathbb{Z}$ obtained by sampling $q \leftarrow \mathbb{Z} \cap [0, X/p)$ and $r \leftarrow \phi$, and returning $x = q \cdot p + r$.*

Let $\mathcal{D}$ be a distribution over $\mathbb{Z} \cap [0, X)$. $\mathrm{AGCD}_{X,\phi}(\mathcal{D})$ consists in distinguishing, given arbitrarily many independent samples, between the uniform distribution over $\mathbb{Z} \cap [0, X)$ and the distribution $A^{\mathrm{AGCD}}_{X,\phi}(p)$ for a fixed $p \leftarrow \mathcal{D}$. We use the notation $\mathrm{AGCD}^m_{X,\phi}(\mathcal{D})$ to emphasize the number of samples m used by the eventual distinguisher.

We say that an algorithm $\mathcal{A}$ is an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{AGCD}_{X,\phi}(\mathcal{D})$ *if, with probability $\geq \varepsilon_2$ over the randomness of $p \leftarrow \mathcal{D}$, its distinguishing advantage between $A^{\mathrm{AGCD}}_{X,\phi}(p)$ and $U(\mathbb{Z} \cap [0, X))$ is $\geq \varepsilon_1$.*[3]

*For $\rho, \eta, \gamma \geq 1$, the (ρ, η, γ)-*AGCD *problem is* $\mathrm{AGCD}_{2^\gamma,\phi}(\mathcal{D})$ *with $\mathcal{D}$ the uniform distribution over η-bit prime integers and ϕ the uniform distribution over $\mathbb{Z} \cap (-2^\rho, 2^\rho)$.*

We will not rely on (ρ, η, γ)-AGCD in our constructions, but we recall it for comparison convenience with prior works. First, we do not need to impose that the secret p is prime. In fact, there is no known attack that exploits the factorization of p. Also, if the distribution $\mathcal{D}$ is sufficiently well-behaved, restricting $\mathcal{D}$ to prime integers (e.g., by rejection sampling) would result in a problem that is no easier, as the density of prime numbers is non-negligible. Second, we do not know how to reduce LWE to (ρ, η, γ)-AGCD. In particular, our reduction leads to distributions $\mathcal{D}$ and ϕ that are somewhat more cumbersome. However, from the perspective of cryptographic constructions, they may be used in the exact same manner as their (ρ, η, γ)-AGCD counterparts.

LWE was introduced by Regev [41]. We use the variant from [8].

Definition 2 (LWE). *Let $n, q \geq 1$, $\boldsymbol{s} \in \mathbb{Z}^n$ and ϕ a distribution over $\mathbb{R}$. We define $A^{\mathrm{LWE}}_{q,\phi}(\boldsymbol{s})$ as the distribution over $\mathbb{T}^n_q \times \mathbb{T}$ obtained by sampling $\boldsymbol{a} \leftarrow \mathbb{T}^n_q$ and $e \leftarrow \phi$, and returning $(\boldsymbol{a}, \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e)$.*

Let $\mathcal{D}$ be a distribution over $\mathbb{Z}^n$. $\mathrm{LWE}_{n,q,\phi}(\mathcal{D})$ *consists in distinguishing, given arbitrarily many independent samples, between $U(\mathbb{T}^n_q \times \mathbb{T})$ and $A^{\mathrm{LWE}}_{q,\phi}(\boldsymbol{s})$ for a fixed $\boldsymbol{s} \leftarrow \mathcal{D}$.*

In [41], Regev described a quantum reduction from several standard (worst-case) problems over n-dimensional Euclidean lattices to $\mathrm{LWE}_{n,p,D_{\leq\alpha}}(U((\mathbb{Z} \cap [0, p))^n))$, where the modulus p may be chosen as a polynomial in n and the parameter α may be set as $\mathrm{poly}(n)/\gamma$ with γ referring to the approximation factor of the considered lattice problem. The reduction assumes that $\alpha \geq \Omega(\sqrt{n}/q)$. Regev's reduction was partly dequantized in [36] and [8]. Further, a modulus-dimension trade-off was exhibited in [8]: in particular, if $q = \Omega(p^n)$ and $\alpha \leq \mathrm{poly}(n)\beta$, then $\mathrm{LWE}_{1,q,D_{\leq\alpha}}(U(\mathbb{Z} \cap [0, q)))$ is no easier than $\mathrm{LWE}_{n,p,D_{\leq\beta}}(U((\mathbb{Z} \cap [0, p))^n))$.

In [4], Applebaum *et al* gave an LWE self-reduction from secret distribution $U((\mathbb{Z} \cap [0, p))^n)$ to secret distribution $D_{\mathbb{Z}^n, O(\alpha p)}$ which reduces the distinguishing advantage from ε to $\Omega(\varepsilon)$, if $\alpha \geq \Omega(\sqrt{\ln(n/\varepsilon)}/p)$.

The main result of this section is the following.

Theorem 1. *Let $\alpha, \beta \in (0, 1)$, $X, B, m, q \geq 1$, and $\mathcal{D}$ a distribution over $\mathbb{Z}$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{AGCD}^m_{X,\lfloor D_{\leq\alpha}\rceil}(\lfloor X/\mathcal{D} \rceil)$. If $\mathcal{D}$ is $(B, \delta, \varepsilon_2/2)$-contained, $q \geq \Omega(\sqrt{\ln(m/\varepsilon_1)}B/\beta)$, $X \geq \Omega(mB^2/(\beta\varepsilon_1))$ and $\beta \leq O(\alpha\delta B/X)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2\delta\beta/\sqrt{\ln(m/\varepsilon_1)}))$-distinguisher for $\mathrm{LWE}^m_{1,q,D_{\leq\beta}}(\mathcal{D})$.*

Setting parameters in AGCD**.** We discuss a possible choice of secure parameters for AGCD.

Recall that there exists a (quantum) reduction from λ-dimensional lattice problems with approximation factors $\lambda^{\widetilde{O}(1)}$, to $\mathrm{LWE}_{1,q,D_{\leq\beta'}}(D_{\mathbb{Z},\sigma})$, for $q = 2^{\widetilde{O}(\lambda)}$, $\beta' = \lambda^{-\widetilde{\Omega}(1)}$ and $\sigma = O(\beta' q)$. We can hence reasonably assume that for these parameters, the time required to solve $\mathrm{LWE}_{1,q,D_{\leq\beta'}}(D_{\mathbb{Z},\sigma})$ is $2^{\widetilde{\Omega}(\lambda)}$, even for $\varepsilon_1, \varepsilon_2$ as small as $2^{-\widetilde{\Omega}(\lambda)}$.

[3] We do not explicitly focus on the distinguishing run-times, as our reductions almost preserve run-times.

We set $\beta = \sqrt{\lambda}\beta'$. As we can publicly increase the Gaussian noise of the LWE samples, $\mathrm{LWE}_{1,q,D_{\leq\beta}}(D_{\mathbb{Z},\sigma})$ is no easier than the latter variant. Now, Theorem 1 depends quite importantly on the potential smallness of samples from $\mathcal{D}$. We avoid such small values by rejection sampling. We define $\mathcal{D}$ as follows: Sample a fresh $s \leftarrow D_{\mathbb{Z},O(\sigma)}$ until $s \in (\sigma/2, 2\sigma)$. As a result, we can set $B = O(\sigma)$ and $\delta = \Omega(1)$. The condition on q in Theorem 1 is fulfilled. If we set $X = 2^\lambda\sigma^2$ and $\alpha = \Omega(\beta X/\sigma) \approx \beta 2^\lambda\sigma$, then all conditions are fulfilled, guaranteeing exponential hardness of $\mathrm{AGCD}^m_{X,\lfloor D_{\leq\alpha}\rceil}(\mathcal{D}^\star_\sigma)$, with $\mathcal{D}^\star_\sigma = \lfloor X/\mathcal{D}\rceil$.

To ease comparison with prior works, we define

$$\gamma = \log X,\ \eta = \log X - \log\sigma, \quad \text{and } \rho = \log\alpha + (\log\lambda)/2.$$

The bit-size of each AGCD sample $pq+r$ is $\approx\gamma$, the bit-size of the AGCD secret p is $\approx\eta$, and the bit-size of each noise term r is bounded by ρ, with probability exponentially close to 1 (in the analysis of the primitives, we will assume that each fresh noise r has magnitude $\leq 2^\rho$, hence forgetting about the unlikely event that one of the noises is bigger). With our choices of X and α, we have: $\gamma \approx \lambda + 2\log\sigma$, $\eta \approx \lambda + \log\sigma$ and $\rho \approx \eta + \log(\sqrt{\lambda}\beta)$.

Proof overview. The proof of Theorem 1 consists of three sub-reductions. We first show that LWE is essentially equivalent to a variant of LWE that does not involve any discretization parameter q. That variant, which we name scale-invariant LWE (SILWE), is implicit in [5, 8]. We then show that SILWE is essentially equivalent to the problem studied in [39] (and inspired from [1]), which we name zero-dimensional LWE (ZDLWE). Finally, the third sub-reduction is from ZDLWE to AGCD.

In Appendix B, we give converse reductions for each one of the three sub-reductions. This implies that from the hardness viewpoint, AGCD and LWE are quite closely related.

2.1 Scale-invariant LWE

We consider the following LWE variant, in which the modulus q does not play a role anymore.

Definition 3 (Scale-Invariant LWE). *Let $n \geq 1$, $\boldsymbol{s} \in \mathbb{Z}^n$ and ϕ a distribution over $\mathbb{R}$. We define $A^{\mathrm{SILWE}}_\phi(\boldsymbol{s})$ as the distribution over $\mathbb{T}^n \times \mathbb{T}$ obtained by sampling $\boldsymbol{a} \leftarrow \mathbb{T}^n$ and $e \leftarrow \phi$, and returning $(\boldsymbol{a}, \langle\boldsymbol{a}, \boldsymbol{s}\rangle + e)$.*

Let $\mathcal{D}$ be a distribution over $\mathbb{Z}^n$. $\mathrm{SILWE}_{n,\phi}(\mathcal{D})$ consists in distinguishing, given arbitrarily many independent samples, between $U(\mathbb{T}^n \times \mathbb{T})$ and $A^{\mathrm{SILWE}}_\phi(\boldsymbol{s})$ for a fixed $\boldsymbol{s} \leftarrow \mathcal{D}$.

Lemma 1. *Let $\alpha, \beta \in (0, 1)$, $m, n, q, B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}^n$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{SILWE}^m_{n,D_{\leq\alpha}}(\mathcal{D})$. If $\mathcal{D}$ is $(B, \varepsilon_2/2)$-bounded, $q \geq \Omega(\sqrt{\ln(mn/\varepsilon_1)}B/\beta)$ and $\beta \leq O(\alpha)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2))$-distinguisher for $\mathrm{LWE}^m_{n,q,D_{\leq\beta}}(\mathcal{D})$.*

Proof. We map each input sample $(\boldsymbol{a}, b)$ for $\mathrm{LWE}_{n,q,D_{\leq\beta}}$ to an input sample $(\boldsymbol{a}', b')$ for $\mathrm{SILWE}_{n,D_{\leq\alpha}}$, as follows: Sample $\boldsymbol{f} \leftarrow D^n_r$ with $r = \Omega(\sqrt{\ln(mn/\varepsilon_1)}/q)$; set $\boldsymbol{a}' = \boldsymbol{a} + \boldsymbol{f}$ and $b' = b$. We show below that, with probability $\geq 1 - \varepsilon_2/2$ over the randomness of $\boldsymbol{s} \leftarrow \mathcal{D}$, this transformation maps the distributions $U(\mathbb{T}^n_q \times \mathbb{T})$ and $A^{\mathrm{LWE}}_{q,D_{\leq\beta}}(\boldsymbol{s})$ to distributions within statistical distances $O(\varepsilon_1/m)$ from $U(\mathbb{T}^n \times \mathbb{T})$ and $A^{\mathrm{SILWE}}_{D_{\leq\alpha}}(\boldsymbol{s})$, respectively.

By Lemma 10, the distribution of $\boldsymbol{f}$ mod 1 is within statistical distance $O(\varepsilon_1/m)$ from $U(\mathbb{T}^n)$, and the distribution of $\boldsymbol{f}$ conditioned on ($\boldsymbol{f}$ mod 1) is $D_{\mathbb{Z}^n/q,r}$. The former imples that $\boldsymbol{a}'$ is

uniformly distributed over $\mathbb{T}^n$. Now, we consider two cases. If b was uniformly distributed in $\mathbb{T}$ independently of $\boldsymbol{a}$, then b' is uniformly distributed in $\mathbb{T}$ independently of $\boldsymbol{a}'$. Now, assume that $b = \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e$ for some fixed $\boldsymbol{s}$ and $e \leftarrow D_{\leq\beta}$. Then $b' - \langle \boldsymbol{a}', \boldsymbol{s} \rangle = e - \langle \boldsymbol{f}, \boldsymbol{s} \rangle$. By Lemma 12, the distribution of $e - \langle \boldsymbol{f}, \boldsymbol{s} \rangle$ (conditioned on $\boldsymbol{a}'$) is within statistical distance $O(\varepsilon_1/m)$ from $D_{\leq\sqrt{\beta^2+\|\boldsymbol{s}\|^2 r^2}}$, assuming that $(1/r^2 + \|\boldsymbol{s}\|^2/\beta^2)^{-1/2} \geq \Omega(\sqrt{\ln(mn/\varepsilon_1)}/q)$. We have $\|\boldsymbol{s}\| \leq B$ with probability $\geq 1 - \varepsilon_2/2$ (over the randomness of $\boldsymbol{s}$). When this is the case, we obtain that $e - \langle \boldsymbol{f}, \boldsymbol{s} \rangle$ (conditioned on $\boldsymbol{a}'$) is within statistical distance $O(\varepsilon_1/m)$ from $D_{\leq\sqrt{\beta^2+\|\boldsymbol{s}\|^2 r^2}}$ (by using the condition on q and the definition of r). Finally, the assumptions on β, B and r ensure that $\sqrt{\beta^2 + \|\boldsymbol{s}\|^2 r^2} \leq \alpha$. □

2.2 Zero-dimensional LWE

We now show that SILWE is essentially equivalent to the problem studied by Regev in [39]. The latter may be viewed as a zero-dimensional variant of LWE, as the provided samples are from $\mathbb{T}$ rather than $\mathbb{T}_q^n \times \mathbb{T}$.

Definition 4 (Zero-Dimensional LWE**).** *Let $s \in \mathbb{Z}$ and ϕ a distribution over $\mathbb{R}$. We define $A_\phi^{\mathrm{ZDLWE}}(s)$ as the distribution over $\mathbb{T}$ obtained by sampling $k \leftarrow \mathbb{Z} \cap [0, s)$ and $e \leftarrow \phi$, and returning $[(k+e)/s]_1$.*

Let $\mathcal{D}$ be a distribution over $\mathbb{Z}$. $\mathrm{ZDLWE}_\phi(\mathcal{D})$ consists in distinguishing, given arbitrarily many independent samples, between $U(\mathbb{T})$ and $A_\phi^{\mathrm{ZDLWE}}(s)$ for a fixed $s \leftarrow \mathcal{D}$.

The following result and its proof are derived from [43].

Lemma 2. *Let $\alpha, \beta \in (0,1)$, $B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{ZDLWE}^m_{D_{\leq\alpha}}(\mathcal{D})$. If $\mathcal{D}$ is $(B, \delta, \varepsilon_2/2)$-contained and $\beta \leq O(\alpha)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2\delta\alpha/\sqrt{\ln(m/\varepsilon_1)}))$-distinguisher for $\mathrm{SILWE}^m_{1,D_{\leq\beta}}(\mathcal{D})$.*

Proof. We describe a reduction from SILWE to ZDLWE. Let $r = \Theta(\sqrt{\ln(m/\varepsilon_1)})$ (chosen to be able to use Lemma 10) and $\delta' \leq \Theta(\delta\alpha/\sqrt{\ln(m/\varepsilon_1)})$. The reduction produces a guess s' of the SILWE secret s by sampling $s' \leftarrow B\delta' \cdot (\mathbb{Z} \cap [0, \lceil 1/\delta' \rceil))$; then it maps any input sample (a, b) for $\mathrm{SILWE}_{1,D_{\leq\beta}}$ to an input sample y for $\mathrm{ZDLWE}_{D_{\leq\alpha}}$, by setting $y = [a - b/s']_1$.

This transformation maps $U(\mathbb{T} \times \mathbb{T})$ to $U(\mathbb{T})$. We now show that it maps $A^{\mathrm{SILWE}}_{D_{\leq\beta}}(s)$ to a distribution that is within statistical distance $O(\varepsilon_1/m)$ from $A^{\mathrm{ZDLWE}}_{D_{\leq\alpha}}(s)$, with probability $\Omega(\varepsilon_2\delta')$ over the choice of $s \leftarrow \mathcal{D}$. Thanks to the assumption on $\mathcal{D}$, we have that $|s| \in [\delta B, B]$, with probability $\geq 1 - \varepsilon_2/2$ over the randomness of s. The success probability of the $\mathrm{ZDLWE}_{D_{\leq\alpha}}$ distinguisher conditioned on that event is $\geq \varepsilon_2/2$. Further, with probability $\geq \Omega(\delta')$ over the choice of s', we have $|s' - s| \leq B\delta'$. We now assume that $|s| \in [\delta B, B]$, $|s' - s| \leq B\delta'$ and that the $\mathrm{ZDLWE}_{D_{\leq\alpha}}$ distinguisher suceeds. This event has weight $\Omega(\varepsilon_2\delta')$.

By Lemma 10, the distribution of a is within statistical distance $O(\varepsilon_1/m)$ of the distribution obtained by sampling $k \leftarrow \mathbb{Z} \cap [0, s)$ and $f \leftarrow D_r$, and returning $[(k+f)/s]_1$. With these notations, we have $b = f + e \bmod 1$ and $y = k/s + f(1/s - 1/s') - e/s' \bmod 1$. The distribution of $sy - k$ is within statistical distance $O(\varepsilon_1/m)$ of $D_{\leq\alpha'}$ with $\alpha' = ((\beta s/s')^2 + r^2(1 - s/s')^2)^{1/2}$. Thanks to the properties on s and s', we have $|s/s'| \leq O(1)$ and $|1 - s/s'| \leq O(\delta'/\delta)$. This leads to $\alpha' \leq O(\beta + r\delta'/\delta)$. The condition on β and the choice of δ' ensure that $\alpha' \leq \alpha$. □

Note that the hardness result obtained here via LWE is stronger than the one from [39]. Indeed, the present approach leads to a (quantum) reduction from standard n-dimensional lattice problems with polynomial approximation factors to ZDLWE with an s of bitsize $\widetilde{O}(n)$, whereas [39] leads to an s of bitsize $\widetilde{O}(n^2)$. However, the latter reduction is classical rather than quantum.

2.3 Reducing ZDLWE to AGCD

Lemma 3. *Let $\alpha, \beta \in (0,1)$, $X, B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{AGCD}^m_{X,\lfloor D_{\leq\alpha}\rceil}(\lfloor X/\mathcal{D}\rceil)$. If $\mathcal{D}$ is $(B, \delta, \varepsilon_2/2)$-contained, $X \geq \Omega(mB^2/(\alpha\varepsilon_1))$ and $\beta \leq O(\alpha\delta B/X)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2))$-distinguisher for $\mathrm{ZDLWE}^m_{D_{\leq\beta}}(\mathcal{D})$.*

Proof. Given an input sample y for $\mathrm{ZDLWE}_{D_{\leq\beta}}$, the reduction produces an input sample x for $\mathrm{AGCD}_{D_{\mathbb{Z},\leq\alpha}}$, as follows: Set $x = [\lfloor Xy\rceil]_X$.

If y is uniformly distributed over $\mathbb{T}$, then so is x over $\mathbb{Z} \cap [0, X)$. Now, assume that $y = (k+e)/s$ for some fixed s (sampled from $\mathcal{D}$), $k \leftarrow \mathbb{Z} \cap [0, s)$ and $e \leftarrow D_{\leq\beta}$. Then $x = kp + r - \Delta$ with $p = \lfloor X/s\rceil$, $r = \lfloor Xe/s\rceil$ and $\Delta = \lfloor X(k+e)/s\rceil - k\lfloor X/s\rceil - r$. We have $|\Delta| \leq 2 + k \leq O(B)$, with probability $\geq 1 - \varepsilon_2/2$ over the choice of $s \leftarrow \mathcal{D}$. The distribution of r is $\lfloor D_{\alpha'}\rceil$ for some $\alpha' \in [X\beta/B, X\beta/(\delta B)]$ (where we used the fact that $|s| \geq \delta B$, which holds with high probability over the choice of s, by assumption on $\mathcal{D}$). We observe that the statistical distance between $\lfloor D_{\alpha'}\rceil$ and $\lfloor D_{\alpha'}\rceil - \Delta$ is $O(\Delta/\alpha') \leq O(B^2/(X\beta)) \leq O(\varepsilon_1/m)$.

It now suffices to show that the distribution of $k \leftarrow \mathbb{Z} \cap [0, s)$ is statistically close to the uniform distribution over $\mathbb{Z} \cap [0, X/p)$. A simple calculation shows that the statistical distance between the uniform distributions over these two intervals is $O(|X/p - s|/s)$. By definition of p, we have that $|ps - X| \leq s$ and hence $|X/(ps) - 1| \leq 1/p$. The latter is $O(B/X) \leq O(\varepsilon_1/m)$, thanks to the assumption on X. □

Theorem 1 is obtained by combining Lemmas 1, 2 and 3.

3 An AGCD-based additive homomorphic encryption scheme

In this section, we propose an additive homomorphic encryption (AHE) scheme whose security relies on the hardness of the AGCD problem. This scheme is similar to the DGHV encryption scheme [20], but the plaintext message is embedded into the ciphertext c as the most significant bit of $c \bmod p$ for the secret key p. It may be viewed as an AGCD adaptation of Regev's encryption scheme from [38].

We let ρ denote a bound on the bit-length of the error, η the bit-length of the secret greatest common divisor, and γ the bit-length of an AGCD sample. The parameter τ refers to the number of encryptions of zero contained in the public key.

Parameters. We set parameters such that they satisfy the following constraints.

- $\rho \geq \lambda$, to protect against the brute force attacks on the noise such as [13, 28].
- $\gamma \geq \Omega(\frac{\lambda}{\log\lambda}(\eta - \rho)^2)$ and $\gamma \leq \eta^2$, to thwart the lattice reduction attacks on AGCD such as the orthogonal lattice attacks [35, 20], Lagarias' simultaneous Diophantine approximation [30] and the Cohn-Heninger attack [15].

- η will be determined later to support correct decryption. For the moment, we only suppose that $\rho < \eta$.
- $\tau = \gamma + 2\lambda + 2$, to be able to use the leftover hash lemma in the security proof (see Subsection 3.3).

Note that there is some discrepancy with the conditions with prior works on AGCD. Part of it stems from the fact that we place ourselves in the context of sub-exponential attackers rather than polynomial-time attackers. Further, once adapted to this attacker setup, the condition corresponding to thwarting lattice attacks is $\gamma \geq \Omega(\lambda\eta^2)$ in prior works. In fact, that condition is too stringent: lattice attacks are tharted even if our (weaker) condition is satisfied. Moreover, our condition is compatible with the LWE to AGCD reduction when applied to exponentially intractable LWE parameters, as explained in Section 2. This has a significant impact on the asymptotic performance of the scheme, as γ may be set much smaller.

Concretely, we set $\rho = \lambda$, $\eta = \rho + L\log\lambda$ for an $L > 0$ to be chosen to provide desirable functionalities, $\gamma = \Omega(L^2\lambda\log\lambda)$ and $\tau = \gamma + 2\lambda + 2$. Note that the ciphertext size γ is quasi-linear in λ. Assume one wants to rely on the exponential hardness of lattice problems for approximation factors $n^{O(L)}$ for a small L. First, one has to set $n = \Omega(L\lambda)$. In that case, via the reduction from Section 2, one can set $\sigma = \Omega(L^2\lambda\log\lambda)$, and $\eta' = cL^2\lambda\log\lambda$ for some constant c, $\rho' = \eta' - L\log\lambda$, $\gamma' = 2\eta' + \lambda$ and $\tau' = \gamma' + 2\lambda + 2$.

3.1 The construction

The scheme **AHE** is defined as follows:

AHE.KeyGen(λ). Given a security parameter λ, determine parameters (X, σ, α) and (γ, η, ρ) providing security λ and decryption correctness (see analysis below). We refer to the discussion just after Theorem 1 for the relationship between the two parameter sets. Sample $p \leftarrow \mathcal{D}^{\star}_{\sigma}$ (of bitsize $\approx \eta$). For $0 \leq i \leq \tau$, sample $x_i \leftarrow A^{\mathrm{AGCD}}_{X,\lfloor D_\alpha \rceil}(p)$. Relabel so that x_0 is the largest and x_1 has an odd $\lfloor \frac{x_1}{p} \rceil$, and restart if we cannot find such an x_1. Output the secret key $sk = p$ and the public key $pk = (x_0, x_1, \ldots, x_\tau)$.

AHE.Enc$_{pk}(m)$. Given a message $m \in \{0,1\}$, uniformly sample a subset $S \subseteq \{1, 2, \ldots, \tau\}$, and output

$$c = \left[\sum_{i \in S} x_i + \left\lfloor \frac{x_1}{2} \right\rceil m \right]_{x_0}.$$

AHE.Add$_{x_0}(c_1, c_2)$. Given two ciphertexts c_1, c_2, output $c_{\mathbf{add}} = [c_1 + c_2]_{x_0}$.

AHE.Dec$_{sk}(c)$. Given a ciphertext c, output $m = \left[\left\lfloor \frac{2c}{p} \right\rceil \right]_2$.

Note that $[\lfloor x \rceil]_2$ may not be equal to $\lfloor [x]_2 \rceil$ for some $x \in \mathbb{R}$. In fact, the latter has value in $\{0, 1, -1\}$ while the former has value in $\{0, 1\}$. However, they are congruent modulo 2.

3.2 Correctness

We analyze the noise growth at encryption and addition, and provide a sufficient condition for decryption correctness.

Lemma 4 (Encryption noise). *Let* $(sk = p, pk = (x_0, \ldots, x_\tau)) \leftarrow \mathbf{AHE.KeyGen}(\lambda)$ *and* $c \leftarrow \mathbf{AHE.Enc}_{pk}(m)$ *for a message* $m \in \{0,1\}$*. Then*

$$c = r + \left\lfloor \frac{p}{2} \right\rceil m \mod p$$

for some r *with* $|r| \leq (2\tau + 1/2)(2^\rho - 1) + 1/2$*.*

Proof. Write $x_i = pq_i + r_i$ with $q_i \in \mathbb{Z}$ and $r_i = [x_i]_p$ for $0 \leq i \leq \tau$. We have $\lfloor \frac{x_1}{2} \rceil = \frac{pq_1}{2} + \frac{r_1}{2} + \delta$ for $|\delta| \leq 1/2$. Since q_1 is odd, we have, modulo p:

$$c \;=\; \sum_{i \in S} x_i + \left\lfloor \frac{x_1}{2} \right\rceil m - kx_0 \;=\; \sum_{i \in S} r_i - kr_0 + \left\lfloor \frac{p}{2} \right\rceil m + \left(\frac{r_1}{2} + \delta \right) m,$$

for some $k \in [0, \tau]$. Therefore, we have $c = r + \left\lfloor \frac{p}{2} \right\rceil m \bmod p$ for some r with $|r| \leq (2\tau + 1/2)(2^\rho - 1) + 1/2$. □

Lemma 5 (Addition noise). *Let* $(sk = p, pk = (x_0, \ldots, x_\tau)) \leftarrow \mathbf{AHE.KeyGen}(\lambda)$ *and* $c_i \leftarrow \mathbf{AHE.Enc}_{pk}(m_i)$ *with* $c_i = r_i + \left\lfloor \frac{p}{2} \right\rceil m_i \bmod p$ *for all* $i \in \{1, 2\}$*. If* $c_{\boldsymbol{add}} \leftarrow \mathbf{AHE.Add}_{x_0}(c_1, c_2)$*, then*

$$c_{\boldsymbol{add}} \;=\; r + \left\lfloor \frac{p}{2} \right\rceil [m_1 + m_2]_2 \mod p,$$

for some r *with* $|r| \leq |r_1 + r_2| + 2^\rho$*.*

Proof. We have, modulo p:

$$c_{\mathbf{add}} \;=\; c_1 + c_2 - \delta x_0 \;=\; r_1 + r_2 - \delta r_0 + \left\lfloor \frac{p}{2} \right\rceil [m_1 + m_2]_2 - \delta'$$

for some $\delta, \delta' \in [-1, 1]$. Hence we can write $c_{\mathbf{add}} = r + \left\lfloor \frac{p}{2} \right\rceil [m_1 + m_2]_2 \bmod p$ for some r with $|r| \leq |r_1 + r_2| + 2^\rho$. □

Lemma 6 (Decryption noise). *Let* p *a positive integer and* $m \in \{0,1\}$*. Given an integer* c*, we have*

$$\mathbf{AHE.Dec}_p(c) = m \quad \text{if } c = r + \left\lfloor \frac{p}{2} \right\rceil m \bmod p \quad \text{with } |r| < \frac{p}{4} - \frac{1}{2}.$$

Proof. Write $c = pq + r + \left\lfloor \frac{p}{2} \right\rceil m$. Then, for some $b \in \{0,1\}$:

$$\left\lfloor c \cdot \frac{2}{p} \right\rceil = \left\lfloor 2q + m + \frac{2r + b}{p} \right\rceil = 2q + m + \left\lfloor \frac{2r + b}{p} \right\rceil,$$

which is congruent to m modulo 2 when $|r| < \frac{p}{4} - \frac{1}{2}$. □

Theorem 2 (Correctness). *Let* $\ell \geq 1$*. Let* $(sk = p, pk = (x_0, \ldots, x_\tau)) \leftarrow \mathbf{AHE.KeyGen}(\lambda)$ *and* $c_i \leftarrow \mathbf{AHE.Enc}_{pk}(m_i)$ *for* $i = 1, \ldots, \ell$ *and* $m_i \in \{0,1\}$*. Let* $c = [\sum_{i=1}^{\ell} c_i]_{x_0}$*. Then we have*

$$\mathbf{AHE.Dec}_p(c) = \left[\sum_{i=1}^{\ell} m_i \right]_2 \quad \text{when} \quad \ell \leq \frac{2^{\eta - \rho}}{6(4\tau + 1)}.$$

In particular, a fresh ciphertext (i.e., with $\ell = 1$*) decrypts correctly if* $\eta - \rho \geq \log(24\tau + 6)$*.*

Proof. For $1 \leq i \leq \ell$, write $c_i = pq_i + r_i + \lfloor \frac{p}{2} \rceil m_i$ for some integers q_i, r_i, and m_i with $|r_i| \leq (2\tau + 1/2)(2^\rho - 1) + 1/2$ and $m_i \in \{0, 1\}$ by Lemma 4. Since $|\sum_{i=1}^{\ell} c_i| \leq \ell x_0/2$, there exists a $k \in \mathbb{Z} \cap [0, \ell/2]$ such that, modulo p:

$$c \;=\; \sum_{i=1}^{\ell} r_i - k r_0 + \left\lfloor \frac{p}{2} \right\rceil \left[\sum_{i=1}^{\ell} m_i \right]_2 \;+\; \left\lfloor \frac{p}{2} \right\rceil \left(\sum_{i=1}^{\ell} m_i - \left[\sum_{i=1}^{\ell} m_i \right]_2 \right).$$

So c is correctly decrypted when $r := \sum_{i=1}^{\ell} r_i - k r_0 + \frac{1}{2}\left(\sum_{i=1}^{\ell} m_i - [\sum_{i=1}^{\ell} m_i]_2 \right)$ is small. By applying Lemmas 4 and 5, we have $|r| \leq (3\ell/2)((2\tau + 1/2)(2^\rho - 1) + 1/2) + \ell/2$. It is less than $\frac{p}{4} - \frac{1}{2}$ if $\ell \leq \frac{2^{\eta-\rho}}{6(4\tau+1)}$. The proof may be completed by using Lemma 6. □

To guarantee correct decryption, it suffices to take $\eta \geq \rho + \log(24\tau + 6)$.

3.3 Security

We first recall the classical Leftover Hash Lemma (LHL) over finite sums modulo an integer as in [20].

Lemma 7. *Sample $x_1, \ldots, x_\tau \leftarrow \mathbb{Z}_{x_0}$ independently, sample $s_1, \ldots, s_\tau \leftarrow \{0, 1\}$, and set $y = \sum_{i=1}^{\tau} s_i x_i \bmod x_0$. Then $(x_1, \ldots, x_\tau, y)$ is within statistical distance $\frac{1}{2}\sqrt{x_0/2^\tau}$ from $U(\mathbb{Z}_{x_0}^{\tau+1})$.*

This LHL is used to show that the **AHE** ciphertext is computationally indistinguishable from a uniform integer modulo x_0, independently of the encrypted plaintext bit.

Theorem 3 (Security). *Under the assumptions that $\mathrm{AGCD}_{X,\lfloor D_\alpha \rceil}(\mathcal{D}_\sigma^\star)$ is hard and that $\tau \geq \log X + 2\lambda + 2$, the **AHE** scheme described above is IND-CPA secure.*

Proof. The key generation procedure produces independent $x_i \leftarrow A_{X,\lfloor D_\alpha \rceil}^{\mathrm{AGCD}}(p)$. With probability exponentially close to 1, there exists i such that x_i is not the largest, and $\lfloor x_i/p \rceil$ is odd.

Now, in the IND-CPA security experiment, we replace the sampling of the x_i's by $x_i \leftarrow U(\mathbb{Z} \cap [0, X))$, independently, for $i = 0, \ldots, \tau$. We still sort them so that x_0 is the largest, and x_1 is such that $\lfloor x_1/p \rceil$ is odd (we resample if we cannot find such an x_1). The resulting public key distribution is computationally indistinguishable from the genuine public key distribution, under the assumption that $\mathrm{AGCD}_{X,\lfloor D_\alpha \rceil}(\mathcal{D}_\sigma^\star)$ is hard.

With this modified key generation procedure, the distribution of $(x_i)_{2 \leq i \leq \tau}$ is within exponentially small statistical distance from $U((\mathbb{Z} \cap [0, x_0))^{\tau-1})$. Using Lemma 7 and the assumption on τ, the tuple $(x_2, \cdots, x_\tau, \sum_{i>1} s_i x_i \bmod x_0)$ is within exponentially small statistical distance from $U((\mathbb{Z} \cap [0, x_0))^\tau)$. As a result, the distribution of the challenge ciphertext in the IND-CPA experiment is within exponentially small statistical distance from $U(\mathbb{Z} \cap [0, x_0))$, independently of the underlying plaintext. In that experiment, the distinguishing advantage of the adversary is exponentially small. □

4 A scale-invariant AGCD-based FHE

In this section, we first extend the **AHE** scheme into a somewhat homomorphic scheme allowing a certain amount of homomorphic data manipulation, and then use Gentry's bootstrapping technique [22] to obtain a fully homomorphic encryption scheme.

We adapt some notations from [7] to our context. Let n be a positive integer. Given $x \in \mathbb{Z} \cap [0, 2^n)$ and $y \in \mathbb{R}$, define

$$\mathrm{BD}_n(x) = (x_0, x_1, \ldots, x_{n-1}) \in \{0,1\}^n \text{ with } x = \sum_{i=0}^{n-1} x_i 2^i$$
$$\mathcal{P}_n(y) = (y, 2y, \ldots, 2^{n-1}y) \in \mathbb{R}^n.$$

Then we can see that

$$\langle \mathrm{BD}_n(x), \mathcal{P}_n(y) \rangle = \sum_{i=1}^{n-1} x_i (2^i y) = xy.$$

We also recall the definition of a tensor product on the vector space $\mathbb{R}^n$

$$(u_1, \ldots, u_n) \otimes (v_1, \ldots, v_n) = (u_1v_1, u_1v_2, \ldots, u_1v_n, \ldots, u_nv_1, \ldots, u_nv_n),$$

and its relation with the inner product

$$\langle \boldsymbol{u} \otimes \boldsymbol{u}', \boldsymbol{v} \otimes \boldsymbol{v}' \rangle = \langle \boldsymbol{u}, \boldsymbol{v} \rangle \cdot \langle \boldsymbol{u}', \boldsymbol{v}' \rangle.$$

4.1 The construction

The scheme **SHE** is identical to **AHE**, except concerning the following two procedures. Its security is inherited from that of **AHE**.

SHE.MultKeyGen(sk). Let $p = sk$. For all $k \in \mathbb{Z} \cap [0, 2\gamma - 2]$, sample $q^*_{i,j}, r^*_{i,j}$ as in $A^{\mathrm{AGCD}}_{X,\lfloor D_\alpha \rceil}(p)$ and publish a vector $\boldsymbol{y} = (pq^*_{i,j} + r^*_{i,j})_{0 \le i,j < \gamma} + \frac{p}{2}([\mathcal{P}_\gamma(2/p)]_2 \otimes [\mathcal{P}_\gamma(2/p)]_2)$ as a multiplication key.

SHE.Mult$_{x_0,\boldsymbol{y}}(c_1, c_2)$. Given two ciphertexts c_1, c_2, output

$$c_{\mathbf{mult}} := [\langle \mathrm{BD}_\gamma(c_1) \otimes \mathrm{BD}_\gamma(c_2), \boldsymbol{y} \rangle]_{x_0}.$$

The (i,j) component of the γ^2-dimensional vector $\boldsymbol{y}$ is a *fake* encryption of $[2^{i+1}/p]_2 \cdot [2^{j+1}/p]_2$, because it is not decrypted into $[2^{i+1}/p]_2 \cdot [2^{j+1}/p]_2$.

4.2 Correctness

We now prove the correctness of the homomorphic multiplication procedure.

Lemma 8. *Let p be a positive integer. If $c = pq + r + \lfloor p/2 \rceil m \in \mathbb{Z} \cap [0, 2^\gamma - 2]$ with $q, r \in \mathbb{Z}$ and $m \in \{0,1\}$, then we have*

$$\langle \mathrm{BD}_\gamma(c), [\mathcal{P}_\gamma(2/p)]_2 \rangle = 2a + m + \varepsilon$$

for an integer a with $|a| \le (\gamma - \eta + 4)/2$ and a real ε with $|\varepsilon| < (2|r| + 1)/p$.

Proof. Let $\lfloor p/2 \rceil = (p+b)/2, b \in \{0,1\}$. Then, $2c/p = 2q + m + \varepsilon$, which is equal to $m + \varepsilon$ modulo 2 for $\varepsilon = (2r+b)/p$ with $|\varepsilon| \le (2|r|+1)/p$.

Since $\mathrm{BD}_\gamma(c)$ is an integer, we have, modulo 2:

$$\langle \mathrm{BD}_\gamma(c), [\mathcal{P}_\gamma(2/p)]_2 \rangle \;\equiv\; \langle \mathrm{BD}_\gamma(c), \mathcal{P}_\gamma(2/p) \rangle \;=\; 2c/p.$$

So, $\langle \mathrm{BD}_\gamma(c), [\mathcal{P}_\gamma(2/p)]_2 \rangle = 2a + m + \varepsilon$ for some integer a. Using $2/p + 2^2/p + \cdots + 2^{\eta-2}/p = 2(2^{\eta-2}-1)/p < 1$, we have

$$|\langle \mathrm{BD}_\gamma(c), [\mathcal{P}_\gamma(2/p)]_2 \rangle| \;\le\; \sum_{i=0}^{\gamma-1} \left| \left[\frac{2^{i+1}}{p} \right]_2 \right| \;\le\; \gamma - \eta + 3,$$

which implies $|a| \le (\gamma - \eta + 4)/2$.

□

Lemma 9 (Multiplication noise). *Let* $(sk = p, pk = (x_0, \ldots, x_\tau)) \leftarrow \mathbf{AHE.KeyGen}(\lambda)$ *and* $\boldsymbol{y} \leftarrow \mathbf{SHE.MultKeyGen}(sk)$. *Given* $c_1, c_2 \in \mathbb{Z} \cap (-x_0/2, x_0/2]$ *satisfying* $c_i = r_i + \lfloor \frac{p}{2} \rceil m_i \bmod p$ *for* $i \in \{0,1\}$, *we have*

$$[\langle \mathrm{BD}_\gamma(c_1) \otimes \mathrm{BD}_\gamma(c_2), \boldsymbol{y} \rangle]_{x_0} \;=\; pq + r + \left\lfloor \frac{p}{2} \right\rceil m_1 m_2$$

for some $q, r \in \mathbb{Z}$ *with* $|r| < (\gamma - \eta + 6)(|r_1| + |r_2|) + \gamma^2 \cdot 2^{\rho+1}$.

Proof. We have

$$\boldsymbol{y} = (pq^*_{i,j} + r^{**}_{i,j})_{0 \le i,j < \gamma} \;+\; \frac{p}{2} \left([\mathcal{P}_\gamma(2/p)]_2 \otimes [\mathcal{P}_\gamma(2/p)]_2 \right),$$

for some $r^{**}_{i,j} \in r^*_{i,j} + [-1/2, 1/2]$ for all i, j. We now use Lemma 8:

$$\begin{aligned}
&\langle \mathrm{BD}_\gamma(c_1) \otimes \mathrm{BD}_\gamma(c_2), \boldsymbol{y} \rangle \\
&= \langle \mathrm{BD}_\gamma(c_1) \otimes \mathrm{BD}_\gamma(c_2), (pq^*_{i,j} + r^{**}_{i,j})_{i,j} \rangle + \frac{p}{2} \langle \mathrm{BD}_\gamma(c_1), [\mathcal{P}_\gamma(2/p)]_2 \rangle \cdot \langle \mathrm{BD}_\gamma(c_2), [\mathcal{P}_\gamma(2/p)]_2 \rangle \\
&= \sum_{(i,j) \in J} \left(pq^*_{i,j} + r^{**}_{i,j} \right) + \frac{p}{2} (m_1 + \varepsilon_1 + 2a_1)(m_2 + \epsilon_2 + 2a_2),
\end{aligned}$$

for some index set $J \subseteq [0, \gamma)^2$, and some $a_1, a_2 \in \mathbb{Z}$, $\varepsilon_1, \varepsilon_2 \in \mathbb{R}$ that satisfy $|a_1|, |a_2| \le (\gamma - \eta + 4)/2$, $|\varepsilon_1| < (2|r_1|+1)/p$ and $|\varepsilon_2| < (2|r_2|+1)/p$. Since $\frac{p}{2}((m_1 + 2a_1)(m_2 + 2a_2) - m_1 m_2)$ is a multiple of p, we have that, for some integer q

$$[\langle \mathrm{BD}_\gamma(c_1) \otimes \mathrm{BD}_\gamma(c_2), \boldsymbol{y} \rangle]_{x_0} = pq + r + \left\lfloor \frac{p}{2} \right\rceil m_1 m_2,$$

where

$$r = \sum_{(i,j) \in J} r^{**}_{i,j} + \frac{p}{2} (\varepsilon_2 (m_1 + 2a_1) + \varepsilon_1 (m_2 + 2a_2) + \varepsilon_1 \varepsilon_2) - \frac{1}{2} m_1 m_2 - k r_0$$

for some $k \in [-1, \gamma^2]$. Therefore, we have $|r| < \gamma^2 \cdot 2^{\rho+1} + (\gamma - \eta + 6)(|r_1| + |r_2|)$. Note that r is an integer because all of $\mathrm{BD}_\gamma(c_1)$, $\mathrm{BD}_\gamma(c_2)$ and $\boldsymbol{y}$ has only integer components. □

Let $c_i \leftarrow \mathbf{SHE.Enc}_{pk}(m_i)$ with $c_i = r_i + \lfloor \frac{p}{2} \rceil m_i \bmod p$ for $i \in \{1,2\}$, $c_{\mathbf{add}} \leftarrow \mathbf{SHE.Add}_{pk}(c_1, c_2)$ and $c_{\mathbf{mult}} \leftarrow \mathbf{SHE.Mult}_{pk}(c_1, c_2)$. From Lemmas 5 and 9, we can see that

$$c_{\mathbf{add}} = r_{\mathbf{add}} + \left\lfloor \frac{p}{2} \right\rceil [m_1 + m_2]_2 \bmod p$$
$$c_{\mathbf{mult}} = r_{\mathbf{mult}} + \left\lfloor \frac{p}{2} \right\rceil [m_1 m_2]_2 \bmod p,$$

with $|r_{\mathbf{add}}| \le |r_1| + |r_2| + 2^\rho$ and $|r_{\mathbf{mult}}| \le (\gamma - \eta + 6)(|r_1| + |r_2|) + \gamma^2 \cdot 2^{\rho+1}$. Both the addition and multiplication in our scheme increase noise only *additively*.

Definition 5. *A scheme* **HE** *is L-homomorphic if for any depth L binary circuit $\mathcal{C}$ and any set of inputs $m_1, \ldots, m_\ell \in \{0,1\}$, it holds that*

$$\mathbf{HE.Dec}_{sk}(\mathbf{HE.Eval}_{evk}(\mathcal{C}, (c_1, \ldots, c_\ell)) = \mathcal{C}(m_1, \ldots, m_\ell)$$

with probability $\ge 1 - \lambda^{-\omega(1)}$, where $(pk, evk, sk) \leftarrow \mathbf{HE.KeyGen}(\lambda)$ and $c_i \leftarrow \mathbf{HE.Enc}_{pk}(m_i)$ for all $i \le \ell$.

Theorem 4. *The scheme* **SHE** *is L-homomorphic if*

$$\eta - \rho \ge L(1 + \log(\gamma - \eta + 6)) + 3 + \log\left(2\tau + \frac{\gamma^2}{\gamma - \eta + 6}\right).$$

Proof. For each $i \in [0, L]$, let c_i be a ciphertext with $c_i = r_i + \lfloor \frac{p}{2} \rceil m_i \bmod p$ after the evaluation of the i-th level gates. Let R_i be a bound on the noise magnitude $|r_i|$. First, we have $R_0 = (2\tau + 1/2)(2^\rho - 1) + 1/2$ by Lemma 4. By Lemmas 5 and 9, we have that $R_{i+1} := 2(\gamma - \eta + 6)R_i + \gamma^2 \cdot 2^{\rho+1}$ is a valid level $(i+1)$ bound (for all $i \ge 0$). By solving the recurrence equation, we obtain

$$R_L \;\le\; \left(2\tau + \frac{\gamma^2}{\gamma - \eta + 6}\right) 2^{\rho+1} 2^L (\gamma - \eta + 6)^L - \frac{\gamma^2 \cdot 2^{\rho+1}}{\gamma - \eta + 6},$$

which is at most $\frac{p}{4} - \frac{1}{2}$ if $(\eta - \rho)$ satisfies the condition. In that case, any ciphertext after evaluation of any circuit of depth L can be correctly decrypted. □

Combining with $\rho = \lambda$, $\gamma = \Omega(\frac{\lambda}{\log \lambda}(\eta - \rho)^2)$, $\tau = \gamma + \Omega(\lambda)$, we may take $\gamma = \Theta(\lambda L^2 \log \lambda)$. Note that the ciphertext size γ is quasi-linear in the security parameter λ.

4.3 Bootstrapping

We provide a bound on the multiplicative depth of the decryption circuit corresponding to $\mathbf{AHE.Dec}_p(c) = [\lfloor 2c/p \rceil]_2$.

We take an approximation z to $2/p$ such that $|z - \frac{2}{p}| < 2^{-(\gamma+\eta)}$. Write $z = \sum_{i=0}^{\gamma+\eta} z_{-i} 2^{-i}$ for $z_{-i} \in \{0,1\}$ for each i. As $c \in [0, 2^\gamma - 2]$, we have $|cz - 2c/p| < 2^{-\eta}$. Therefore, we have $[\lfloor cz \rceil]_2 = m$ when $c = pq + r + \lfloor \frac{p}{2} \rceil m$ and $|r| < \frac{p}{4} - \frac{1}{2}$. Since the η most significant bits of z are zero, the most expensive step in decryption consists in adding up to γ integers of bit-lengths $\le 2\gamma$. This can be implemented with a binary circuit of $O(\log \gamma)$ depth.

By Theorem 4 and [22], **SHE** is bootstrappable and may be turned into a fully homomorphic encryption scheme when $\eta - \rho = \Omega(\log^2 \gamma)$. For bootstrapping we publish encryptions of z_i's as a bootstrapping key. This requires space $O(\gamma^2)$.

5 Truncation of ciphertexts

Since the message bit is embedded into the most significant bit of the ciphertext modulo p, some least significant bits of the ciphertext are irrelevant to decryption correctness. However as we truncate more least significant bits of a ciphertext, decryption failure probability increases slightly at first and becomes overwhelming after some point between ρ and η bits of truncation.

Let $N = 2^\nu$ for a positive integer $\nu < \gamma$. Given a ciphertext c, define $\hat{c} = \lfloor c/N \rceil$ so that $|N\hat{c} - c| \le N/2$. In the following, the quantity $\hat{c}$ can play a similar role to that of the corresponding ciphertext c in each component of **SHE**, for appropriate ν.

1. In the encryption stage, given $\hat{c} := [\sum_{i\in S} \hat{x}_i + \lfloor \frac{\hat{x}_1}{2} \rceil m]_{\hat{x}_0}$ for some $S \subseteq \mathbb{Z} \cap [1, \tau]$, we have, for some integer k:

$$N\hat{c} = \sum_{i\in S} x_i + \lfloor \frac{x_1}{2} \rceil m - kx_0 + N\sum_{i\in S}(\hat{x}_i - x_i/N) + Nm\left(\lfloor \frac{\hat{x}_1}{2} \rceil - \lfloor \frac{x_1}{2} \rceil /N\right) - kN(\hat{x}_0 - x_0/N),$$

   which is equivalent to $r + \lfloor \frac{p}{2} \rceil m$ modulo p for $|r| \le (2^{\rho+1} + N)(\tau + 1)$.
2. In the decryption stage, given $\hat{c} = \lfloor c/N \rceil$ we have

$$\left[\left\lfloor \frac{2\hat{c}}{p/N} \right\rceil\right]_2 = m$$

   if $|r| < (p - N)/4 - 1/2$ when $c = r + \lfloor \frac{p}{2} \rceil m \bmod p$.
3. In the addition stage, given $\hat{c}_1$ and $\hat{c}_2$, we define $\hat{c}_{\mathbf{add}} = [\hat{c}_1 + \hat{c}_2]_{\hat{x}_0}$. Then we can show (similarly to Lemma 5):

$$N\hat{c}_{\mathbf{add}} \;=\; r + \left\lfloor \frac{p}{2} \right\rceil [m_1 + m_2]_2 \bmod p$$

   for $|r| \le |r_1| + |r_2| + (2^\rho + \frac{3}{2}N)$ when $c_i = pq_i + r_i + \lfloor \frac{p}{2} \rceil m$ for $i \in \{1, 2\}$.
4. In the multiplication stage, we use $\hat{\boldsymbol{y}}$ to be the vector of bit-length $(\gamma - \nu)^2$ obtained by removing all entries (i, j) of $\boldsymbol{y}$ such that $i < \nu$ or $j \le \nu$. Given $\hat{c}_1$ and $\hat{c}_2$, we set

$$\hat{c}_{\mathbf{mult}} = [\mathrm{BD}_{\hat{\gamma}}(\hat{c}_1) \otimes \mathrm{BD}_{\hat{\gamma}}(\hat{c}_2), \hat{\boldsymbol{y}}]_{\hat{x}_0},$$

   where $\hat{\gamma} = \gamma - \nu$. We can show that for all $i \in \{1, 2\}$, we have $\langle \mathrm{BD}_{\hat{\gamma}}(\hat{c}_i), [\mathcal{P}_{\hat{\gamma}}(2N/p)]_2 \rangle = \frac{2\hat{c}_i}{p/N} = 2a_i + m_i + \varepsilon_i$ for $a_i \in \mathbb{Z}$ with $|a_i| \le (\hat{\gamma} - \eta + 4)/2$, $m_i \in \{0, 1\}$ and $\varepsilon_i \in \mathbb{R}$ with $|\varepsilon_i| < \frac{2(|r|+N)}{p}$. Thus $\hat{c}_{\mathbf{mult}} = pq + r + \lfloor \frac{p}{2} \rceil m_1 m_2$ for some $q, r \in \mathbb{Z}$ satisfying

$$|r| < (\hat{\gamma} - \eta + 6)(|r_1| + |r_2|) + \hat{\gamma}^2 \cdot 2^{\rho+1},$$

   when $c_i = r_i + \lfloor \frac{p}{2} \rceil m_i \bmod p$ for $i \in \{1, 2\}$.
5. In the bootstrapping stage, we take $z \in 2^{-(\gamma+\eta-\rho)}$ to be an approximation of $2/p$ with $|z - \frac{2}{p}| < 2^{-(\gamma+\eta-\rho)}$. We have

$$\left\| N\hat{c}z - \frac{2}{p}c \right\| \le \left(c + \frac{1}{N}\right)\left(\frac{2}{p} + 2^{-(\gamma+\eta-\rho)}\right) \le \frac{(2^{\rho+1} + N)}{p}.$$

   Combining with $\frac{2c}{p} = 2q + m + \frac{2r+m}{p}$ for $c = pq + r + \lfloor \frac{p}{2} \rceil m$, we have $[\lfloor N\hat{c}z \rceil]_2 = m$ if $(2r + m + 2^{\rho+1} + N)/p < 1/2$. It is satisfied when $|r| < p/4 - 2^\rho - (N + 1)/2$. If $N \le p/4$, we have a similar homomorphic capacity of Theorem 4 and so **SHE** becomes bootstrappable similarly. In this case, however, the decryption can be done with a binary circuit of $O(\log(\gamma - \rho))$ depth.

In the above observations, we can see that encryption noise and addition noise are almost the same when $\nu \leq \rho$ and the decryption and the bootstapping work similarly when $\nu < \eta - 1$. For multiplication, as ν grows, the multiplication error decreases. Hence truncating ciphertexts by ρ bits results in similar performance, but with reduced ciphertext bit-length. In that setup, the bit-size of ciphertext becomes $\gamma - \rho$.

The known attacks on AGCD do not say much on the complexity of AGCD when $\gamma - \rho$ is small. A naive attack is as follows. Given $c = pq + r$, we first guess the $\gamma - \eta$ bits of q and then compute $\lfloor \frac{c}{q} \rceil = p + \lfloor \frac{r}{q} \rceil$. Since $\frac{r}{q} < 2^{\rho-(\gamma-\eta)}$, we can obtain the $\gamma - \rho$ most significant bits of p. This is significant. To avoid this attack, we need to set $\gamma - \eta \geq \lambda$. In that case, the ciphertext size is $\approx \gamma - \rho = (\gamma - \eta) + (\eta - \rho) \geq \lambda + \Omega(L \log \lambda)$.

Note that this truncation method is different from decreasing the bit-size ρ of the noise. If ρ is set smaller, then the ciphertext bit-length γ should be increased to resist lattice-based attacks, i.e., it must satisfy $\gamma \geq \Omega(\frac{\lambda}{\log \lambda}(\eta - \rho)^2)$. If we reduce ρ and η simultaneously, resistance against the lattice-based attacks can be maintained, but the scheme becomes susceptible to exhaustive search on the noise components r_i.

Acknowledgments. The authors thank Miran Kim, Adeline Langlois, Yongsoo Song and Ron Steinfeld for helpful discussions. We thank Oded Regev for his sketch of the SILWE to ZDLWE reduction [43]. The first author was supported by the ICT R&D program of MSIP/IITP [No. 10047212]. The second author was supported by the ERC Starting Grant ERC-2013-StG-335086-LATTAC.

References

1. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. of STOC*, pages 284–293. ACM, 1997.
2. M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proc. of STOC*, pages 601–610. ACM, 2001.
3. J. Alperin-Sheriff and C. Peikert. Faster bootstrapping with polynomial error. In *Proc. of CRYPTO*, volume 8616 of *LNCS*, pages 297–314. Springer, 2014.
4. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of CRYPTO*, volume 5677 of *LNCS*, pages 595–618. Springer, 2009.
5. Z. Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In *Proc. of CRYPTO*, volume 7417 of *LNCS*, pages 868–886. Springer, 2012.
6. Z. Brakerski, C. Gentry, and S. Halevi. Packed ciphertexts in LWE-based homomorphic encryption. In *Proc. of PKC*, volume 7778 of *LNCS*, pages 1–13. Springer, 2013.
7. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *Proc. of ITCS*, pages 309–325. ACM, 2012.
8. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *Proc. of STOC*, pages 575–584. ACM, 2013.
9. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proc. of FOCS*, pages 97–106. IEEE Computer Society Press, 2011.
10. Z. Brakerski and V. Vaikuntanathan. Fully homomorphic encryption from Ring-LWE and security for key dependent messages. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 505–524. Springer, 2011.
11. Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput*, 43(2):831–871, 2014.
12. Z. Brakerski and V. Vaikuntanathan. Lattice-based FHE as secure as PKE. In *Proc. of ITCS*, pages 1–12. ACM, 2014.
13. Y. Chen and P. Q. Nguyen. Faster algorithms for approximate common divisors: Breaking fully-homomorphic-encryption challenges over the integers. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 502–519. Springer, 2012.

14. J. H. Cheon, J.-S. Coron, J. Kim, M. S. Lee, T. Lepoint, M. Tibouchi, and A. Yun. Batch fully homomorphic encryption over the integers. In *Proc. of EUROCRYPT*, volume 7881 of *LNCS*, pages 315–335. Springer, 2013.
15. H. Cohn and N. Heninger. Approximate common divisors via lattices. Cryptology ePrint Archive, Report 2011/437, 2011. `http://eprint.iacr.org/`.
16. J.-S. Coron, T. Lepoint, and M. Tibouchi. Practical multilinear maps over the integers. In *Proc. of CRYPTO*, volume 8042 of *LNCS*, pages 476–493. Springer, 2013.
17. J.-S. Coron, T. Lepoint, and M. Tibouchi. Scale-invariant fully homomorphic encryption over the integers. In *Proc. of PKC*, volume 8383 of *LNCS*, pages 311–328. Springer, 2014.
18. J.-S. Coron, A. Mandal, D. Naccache, and M. Tibouchi. Fully homomorphic encryption over the integers with shorter public keys. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 487–504. Springer, 2011.
19. J.-S. Coron, D. Naccache, and M. Tibouchi. Public key compression and modulus switching for fully homomorphic encryption over the integers. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 446–464. Springer, 2012.
20. M. van Dijk, C. Gentry, S. Halevi, and V. Vaikuntanathan. Fully homomorphic encryption over the integers. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 24–43. Springer, 2010.
21. C. Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. `http://crypto.stanford.edu/craig`.
22. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. of STOC*, pages 169–178. ACM, 2009.
23. C. Gentry and S. Halevi. Implementing Gentry's fully-homomorphic encryption scheme. In *Proc. of EUROCRYPT*, volume 6632 of *LNCS*, pages 129–148. Springer, 2011.
24. C. Gentry, S. Halevi, and N. P. Smart. Fully homomorphic encryption with polylog overhead. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 465–482. Springer, 2012.
25. C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In *Proc. of CRYPTO*, volume 7417 of *LNCS*, pages 850–867. Springer, 2012.
26. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008. Full version available at `http://eprint.iacr.org/2007/432.pdf`.
27. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Proc. of CRYPTO*, volume 8042 of *LNCS*, pages 75–92. Springer, 2013.
28. N. A. Howgrave-Graham. Approximate integer common divisors. In *Proc. of CALC*, volume 2146 of *LNCS*, pages 51–66. Springer, 2001.
29. J. Kim, M. S. Lee, A. Yun, and J. H. Cheon. CRT-based fully homomorphic encryption over the integers. Cryptology ePrint Archive, Report 2013/057, 2013. `http://eprint.iacr.org/`.
30. J. C. Lagarias. The computational complexity of simultaneous diophantine approximation problems. *SIAM J. Comput.*, 14(1):196–209, 1985.
31. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.
32. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 700–718. Springer, 2012.
33. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
34. D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography, D. J. Bernstein, J. Buchmann, E. Dahmen (Eds)*, pages 147–191. Springer, 2009.
35. P. Q. Nguyen and J. Stern. The two faces of lattices in cryptology. In *Proc. of CALC*, volume 2146 of *LNCS*, pages 146–180. Springer, 2001.
36. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
37. C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 80–97. Springer, 2010.
38. O. Regev. New lattice based cryptographic constructions. In *Proc. of STOC*, pages 407–416. ACM, 2003.
39. O. Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004.
40. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.
41. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

42. O. Regev. The learning with errors problem, 2010. Invited survey in CCC 2010, available at `http://www.cims.nyu.edu/~regev/`.
43. O. Regev. Private communication, 2012.
44. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor. Comput. Science*, 53:201–224, 1987.
45. N. P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Proc. of PKC*, volume 6056 of *LNCS*, pages 420–443. Springer, 2010.
46. D. Stehlé and R. Steinfeld. Faster fully homomorphic encryption. In *Proc. of ASIACRYPT*, volume 6477 of *LNCS*, pages 377–394. Springer, 2010.

A Some useful lemmas on lattice Gaussians

The first statement of the following lemma is a special case of [33, Le. 4.1]. The second statement can be obtained by a simple calculation exploiting [40, Claim 3.8].

Lemma 10. *Let $r, \varepsilon > 0$ such that $r \geq \Omega(\sqrt{\ln(1/\varepsilon)})$. Then the distribution $D_r \bmod 1$ is within statistical distance $O(\varepsilon)$ from $U(\mathbb{T})$. If $x \leftarrow D_r$, then the distribution of x conditioned on $x \bmod 1$ is within statistical distance $O(\varepsilon)$ of $D_{\mathbb{Z},r}$.*

Lemma 11 (Special case of [26, Cor. 2.8]). *Let $q \geq 1$ and $r, \varepsilon > 0$ such that $r \geq \Omega(\sqrt{\ln(1/\varepsilon)})$. Then the distribution $D_{\mathbb{Z}/q,r} \bmod 1$ is within statistical distance $O(\varepsilon)$ from $U(\mathbb{T}_q)$.*

Lemma 12 (Special case of [40, Cor. 3.10]). *Let $n \geq 1, \boldsymbol{z} \in \mathbb{R}^n$ and $r, \varepsilon, \alpha > 0$ with $(1/r^2 + \|\boldsymbol{z}\|^2/\alpha^2)^{-1/2} \geq \Omega(\sqrt{\ln(n/\varepsilon)})$. Then the distribution of $\langle D^n_{\mathbb{Z},r}, \boldsymbol{z}\rangle + D_\alpha$ is within statistical distance $O(\varepsilon)$ from D_β with $\beta = (\alpha^2 + \|\boldsymbol{z}\|^2 r^2)^{1/2}$.*

Lemma 13 (Special case of [37, Th. 3.1]). *Let $r, q, s, \varepsilon > 0$ such that $r \geq \Omega(\sqrt{\ln(1/\varepsilon)}/q)$. Sample $x \leftarrow D_s$ and $k \leftarrow D_{\mathbb{Z}/q,r,x}$. Then the distribution of k is within statistical distance $O(\varepsilon)$ from $D_{\mathbb{Z}/q,(r^2+s^2)^{1/2}}$ and, conditioned on k, the distribution of x is within statistical distance $O(\varepsilon)$ from $k\frac{1}{(1+r^2/s^2)^{1/2}} + D_{(r^{-2}+s^{-2})^{-1/2}}$.*

We also use the following result, which can be derived from Lemmas 10, 11 and 13.

Lemma 14. *Let $r, q, \varepsilon > 0$ such that $r \geq \Omega(\sqrt{\ln(1/\varepsilon)}/q)$. Sample $x \leftarrow \mathbb{T}$ and $k \leftarrow D_{\mathbb{Z}/q,r,x}$. Then $k \bmod 1$ is uniformly distributed over $\mathbb{T}_q$ and the distribution of x conditioned on k is within statistical distance $O(\varepsilon)$ from $k + D_r \bmod 1$.*

Proof. Let $s \geq \Omega(\sqrt{\ln(1/\varepsilon)}/q)$. By Lemma 10, the uniform distribution over $\mathbb{T}$ is within statistical distance $O(\varepsilon)$ from the distribution $D_s \bmod 1$. By Lemma 13, the distribution of k is within statistical distance $O(\varepsilon)$ from $D_{\mathbb{Z}/q,(r^2+s^2)^{1/2}} \bmod 1$ and conditioned on k the distribution of x is within statistical distance $O(\varepsilon)$ from $k\frac{1}{(1+r^2/s^2)^{1/2}} + D_{(r^{-2}+s^{-2})^{-1/2}} \bmod 1$. The proof can be completed by using Lemma 11 and letting s tend to infinity. □

B Converse results for Lemmas 1, 2 and 3

Lemma 15. *Let $\alpha, \beta \in (0,1)$, $m, n, q, B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}^n$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{LWE}^m_{n,q,D_{\leq\beta}}(\mathcal{D})$. If $\mathcal{D}$ is $(B, \varepsilon_2/2)$-bounded, and $\alpha \leq \beta - \Omega(\sqrt{\ln(mn/\varepsilon_1)}B/q)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2))$-distinguisher for $\mathrm{SILWE}^m_{n,D_{\leq\alpha}}(\mathcal{D})$.*

Proof. The reduction architecture is similar to the one of Lemma 1. We map each input sample $(\boldsymbol{a}, b)$ for $\mathrm{SILWE}_{n,D_{\leq\alpha}}$ to an input samples $(\boldsymbol{a}', b')$ for $\mathrm{LWE}_{n,q,D_{\leq\beta}}$ as follows: Sample $\boldsymbol{a}' \leftarrow D_{\mathbb{Z}^n/q,r,\boldsymbol{a}}$ with $r = \Omega(\sqrt{\ln(mn/\varepsilon_1)})/q$ (for this, we independently sample each coordinate $a'_j \leftarrow D_{\mathbb{Z}/q,r,a_j}$); set $b' = b$.

By Lemma 14, we have that the distribution of $\boldsymbol{a}'$ is within statistical distance $O(\varepsilon_1/m)$ from $U(\mathbb{T}_q^n)$, and that conditioned on $\boldsymbol{a}'$, the distribution of $\boldsymbol{f} := \boldsymbol{a} - \boldsymbol{a}'$ is within statistical distance $O(\varepsilon_1/m)$ from D_r. As a result, the transformation maps the uniform distribution over $\mathbb{T}^n \times \mathbb{T}$ to a distribution within statistical distance $O(\varepsilon_1/m)$ from the uniform distribution over $\mathbb{T}_q^n \times \mathbb{T}$. Further, if $b = \langle \boldsymbol{a}, \boldsymbol{s}\rangle + e$ for some fixed $\boldsymbol{s}$ and $e \leftarrow D_{\leq\alpha}$, then $b' = b = \langle \boldsymbol{a}', \boldsymbol{s}\rangle + \langle \boldsymbol{f}, \boldsymbol{s}\rangle + e$. Conditioned on $\boldsymbol{a}'$, the distribution of $\langle \boldsymbol{f}, \boldsymbol{s}\rangle + e$ is within statistical distance $O(\varepsilon_1/m)$ of $D_{\leq\sqrt{\alpha^2+\|\boldsymbol{s}\|^2 r^2}}$. We have $\|\boldsymbol{s}\| \leq B$ with probability $\geq 1 - \varepsilon_2/2$ over the randomness of $s \leftarrow \mathcal{D}$. This allows to complete the proof. □

Lemma 16 (Adapted from [42, App. A]). *Let $\alpha, \beta \in (0,1)$, $B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{SILWE}^m_{1,D_{\leq\beta}}(\mathcal{D})$. If $\mathcal{D}$ is $(B, \delta, \varepsilon_2/2)$-contained and $\alpha \leq O(\beta)$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2\delta\alpha/\sqrt{\ln(m/\varepsilon_1)}))$-distinguisher for $\mathrm{ZDLWE}^m_{D_{\leq\alpha}}(\mathcal{D})$.*

Proof. We reduce ZDLWE to SILWE. Let $r = \Theta(\sqrt{\ln(m/\varepsilon_1)})$ (chosen to be able to use Lemma 10) and $\delta' = \Theta(\delta\alpha/\sqrt{\ln(m/\varepsilon_1)})$. The reduction produces a guess s' of the ZDLWE secret s by sampling $s' \leftarrow B\delta'\cdot(\mathbb{Z}\cap[0, \lceil 1/\delta'\rceil))$; then it maps any input sample y for $\mathrm{ZDLWE}_{D_{\leq\alpha}}$ to an input sample (a, b) for $\mathrm{SILWE}_{1,D_{\leq\beta}}$, as follows: Sample $f \leftarrow D_r$; set $a = [y + f/s']_1$ and $b = [f]_1$.

We assume that $|s| \in [\delta B, B]$, $|s' - s| \leq B\delta'$ and that the $\mathrm{SILWE}_{1,D_{\leq\beta}}$ distinguisher succeeds. As in the proof of Lemma 2, this event has weight $\Omega(\varepsilon_2\delta')$.

Assume that y is uniformly distributed in $\mathbb{T}$. By Lemma 10, the distribution of b is within statistical distance $O(\varepsilon_1/m)$ from uniform, independently of y. Therefore, the distribution of the pair (a, b) is within statistical distance $O(\varepsilon_1/m)$ from uniform.

Now, assume that $y = (k + e)/s$ with $k \leftarrow \mathbb{Z} \cap [0, s)$ and $e \leftarrow D_{\leq\alpha}$. We have $a = (k + e)/s + f/s'$ and $b - as = [-e + f(1 - s/s')]_1$. Let $f' = fs/s'$. By Lemma 10, the distribution of $a' := (k + f')/s$ is within statistical distance $O(\varepsilon_1/m)$ from uniform and the distribution of f' conditioned on a' is within statististical distance $O(\varepsilon_1/m)$ of $D_{\mathbb{Z},r|s/s'|}$. The assumption of Lemma 10 holds because $r|s/s'| \geq \Omega(r) \geq \Omega(\sqrt{\ln(m/\varepsilon_1)})$, thanks to the choices of δ' and r. By Lemma 12, the distribution of $b - as = -e + f'(s'/s - 1)$ is within statistical distance $O(\varepsilon_1/m)$ of $D_{\leq\beta'}$ with $\beta' = \sqrt{\alpha^2 + r^2(1 - s/s')^2}$, assuming that $((s'/(rs))^2 + (s'/s - 1)^2/\alpha^2)^{-1/2} \geq \Omega(\sqrt{\ln(m/\varepsilon_1)})$. As $|s/s'| \geq \Omega(1)$ and $|s'/s - 1| \leq O(\delta'/\delta)$, the definitions of r and δ' imply that the latter condition holds. Further, as $|1 - s/s'| \leq O(\delta'/\delta)$, we have that $\beta' \leq O(\alpha + r\delta'/\delta)$. This completes the proof. □

Lemma 17. *Let $\alpha, \beta \in (0,1)$, $X, B \geq 1$ and $\mathcal{D}$ a distribution over $\mathbb{Z}$. Assume that there exists an $(\varepsilon_1, \varepsilon_2)$-distinguisher for $\mathrm{ZDLWE}^m_{D_{\leq\beta}}(\lfloor X/\mathcal{D}\rceil)$. If $\mathcal{D}$ is $(B, \delta, \varepsilon_2/2)$-contained, $X \geq \Omega(mB/\varepsilon_1)$, $\alpha \geq \Omega(Bm/(\delta\varepsilon_1))$ and $\beta \geq \alpha B/X$, then there exists an $(\Omega(\varepsilon_1), \Omega(\varepsilon_2))$-distinguisher for $\mathrm{AGCD}^m_{X,\lfloor D_{\leq\alpha}\rceil}(\mathcal{D})$.*

Proof. Given an input sample x for $\mathrm{AGCD}_{X,\lfloor D_\alpha\rceil}$, the reduction produces an input sample y for $\mathrm{ZDLWE}_{D_{\leq\alpha}}$ as follows: Sample $f \leftarrow \mathbb{T}$; set $y = (x + f)/X$. If x is uniformly distributed over $\mathbb{Z} \cap [0, X)$, then so is y over $\mathbb{T}$.

Now, assume that $x = qp + r$ for some fixed $p = \lfloor X/s \rceil$ (with s sampled from $\mathcal{D}$), $q \leftarrow \mathbb{Z} \cap [0, X/p)$ and $r \leftarrow \lfloor D_{\leq\alpha} \rceil$. We have $y = (q + e + \Delta)/s$ with $s = \lfloor X/p \rceil$, $e = (r + f)s/X$ and $\Delta = \varepsilon qs/X$ for some $\varepsilon \in [-1/2, 1/2]$.

As $\alpha \geq \Omega(\sqrt{\ln(m/\varepsilon_1)})$, the distribution of e is within statistical distance $O(\varepsilon_1/m)$ from $D_{\leq\alpha|s|/X}$. Further, we have $|\Delta| \leq |s|/p \leq O(B^2/X)$ (assuming that $|s| \leq B$). Thanks to the assumtions on $\mathcal{D}$ and α, we have $|\Delta| \leq O(\frac{\varepsilon_1}{m}\frac{\alpha|s|}{X})$, and the term $e + \Delta$ is within statistical distance $O(\varepsilon_1/m)$ from $D_{\leq\alpha|s|/X}$. Note that $\alpha B/X \leq \beta$.

It now suffices to show that the distributions $U(\mathbb{Z} \cap [0, s))$ and $U(\mathbb{Z} \cap [0, X/p))$ are within statistical distance $O(\varepsilon_1/m)$. The proof is identical to that of Lemma 3. □

C Orthogonal lattice attack on AGCD

Let us recall the orthogonal lattice attack on AGCD in [20].

Suppose we are given samples $(x_i = pq_i + r_i)_{1\leq i\leq m}$ from $A^{\mathrm{AGCD}}_{X,\lfloor D_\alpha \rceil}(p)$. Let 2^ρ be an upper bound on the magnitudes of the of r_i's. Consider the integral lattice $\mathcal{L}$ generated by the rows of the following $m \times (m + 1)$ matrix:

$$\begin{bmatrix} x_1 & 2^\rho & & & \\ x_2 & & 2^\rho & & \\ \vdots & & & \ddots & \\ x_m & & & & 2^\rho \end{bmatrix}$$

Define the vector $\boldsymbol{u} := (1, -\frac{r_1}{2^\rho}, \ldots, -\frac{r_d}{2^\rho})$. For any element $\boldsymbol{v} \in \mathcal{L}$, we have $\langle \boldsymbol{u}, \boldsymbol{v} \rangle \equiv 0 \bmod p$. Further, if $\|\boldsymbol{v}\|_1 < p$, then we have $|\langle \boldsymbol{u}, \boldsymbol{v} \rangle| \leq \|\boldsymbol{v}\|_1 < p$, since each component of $\boldsymbol{u}$ is at most 1. That is, we have $\langle \boldsymbol{u}, \boldsymbol{v} \rangle = 0$ over $\mathbb{Z}$. Hence if we find m linearly independent vectors $\boldsymbol{v}$ in $\mathcal{L}$ with $\|\boldsymbol{v}\|_1 < p$, we can recover $\boldsymbol{u}$ and hence find p from $\gcd(x_1 - r_1, \ldots, x_m - r_m)$ with overwhelming probability.

The lattice $\mathcal{L}$ has determinant $\approx 2^{\gamma+(m-1)\rho}$. Assuming that all minima are almost equal, their norms are $\approx 2^{\gamma/m+\rho(m-1)/m}$. In time 2^λ, lattice reduction [44, 2] allows to find m linearly independent lattice vectors of norms $\approx \lambda^{O(m/\lambda)} \cdot 2^{\rho(m-1)/m+\gamma/m}$. The optimal choice for m is $\approx \Theta(\sqrt{\gamma\lambda/\log\lambda})$, which leads to vector norms that are $\approx 2^{O(\sqrt{\gamma\log\lambda/\lambda})+\rho}$. The attack is thwarted if $\gamma \geq \Omega(\frac{\lambda}{\log\lambda}(\eta - \rho)^2)$.

The Simultaneous Diophantine Approximation algorithm for AGCD (see [20]) has similar performance.