GGHLite: More Efficient Multilinear Maps from Ideal Lattices

Adeline Langlois[1], Damien Stehlé[1], Ron Steinfeld[2]

[1] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENS Lyon, INRIA, UCBL),
46 Allée d'Italie, 69364 Lyon Cedex 07, France.
[2] Clayton School of Information Technology, Monash University, Clayton, Australia.

Abstract. The GGH Graded Encoding Scheme [9], based on ideal lattices, is the first plausible approximation to a cryptographic multilinear map. Unfortunately, using the security analysis in [9], the scheme requires very large parameters to provide security for its underlying "encoding re-randomization" process. Our main contributions are to formalize, simplify and improve the efficiency and the security analysis of the re-randomization process in the GGH construction. This results in a new construction that we call GGHLite. In particular, we first lower the size of a standard deviation parameter of the re-randomization process of [9] from exponential to polynomial in the security parameter. This first improvement is obtained via a finer security analysis of the "drowning" step of re-randomization, in which we apply the *Rényi divergence* instead of the conventional *statistical distance* as a measure of distance between distributions. Our second improvement is to reduce the number of randomizers needed from $\Omega(n \log n)$ to 2, where n is the dimension of the underlying ideal lattices. These two contributions allow us to decrease the bit size of the public parameters from $O(\lambda^5 \log \lambda)$ for the GGH scheme to $O(\lambda \log^2 \lambda)$ in GGHLite, with respect to the security parameter λ (for a constant multilinearity parameter κ).

1 Introduction

Boneh and Silverberg [6] defined a *cryptographic κ-multilinear map e* as a map from $G_1 \times \ldots \times G_\kappa$ to G_T, all cyclic groups of order p, which enjoys three main properties: first, for any elements $g_i \in G_i$ for $i \leq \kappa$, $j \leq \kappa$ and $\alpha \in \mathbb{Z}_p$, we have $e(g_1, \ldots, \alpha \cdot g_j, \ldots, g_\kappa) = \alpha \cdot e(g_1, \ldots, g_\kappa)$; second, the map e is non-degenerate, i.e., if the g_i's are generators of their respective G_i's then $e(g_1, \ldots, g_\kappa)$ generates G_T; and third, there is no efficient algorithm to compute discrete logarithms in any of the G_i's. Bilinear maps ($\kappa = 2$) and multilinear maps have a lot of cryptographic applications, see [11,21,5] and [6,20,16,19], respectively. But unlike bilinear maps, built with pairings on elliptic curves, the construction of cryptographic multilinear maps was an open problem for several years. In [6], Boneh and

Silverberg studied the interest of such maps, and gave two applications: multipartite Diffie-Hellman key exchange and very efficient broadcast encryption. But they conjectured that multilinear maps will probably "come from outside the realm of algebraic geometry." In 2013, Garg, Gentry and Halevi [9] introduced the first "approximate" multilinear maps contruction, based on ideal lattices, and the powerful notion of *graded encoding scheme*. Based on their work, Coron, Lepoint and Tibouchi [7] recently described an alternative construction of graded encoding scheme.

We first give a high level description of the GGH graded encoding scheme [9]. If we come back to the definition of cryptographic multilinear maps, the authors of [9] notice that $\alpha \cdot g_i$ can be viewed as an "encoding" of the "plaintext" $\alpha \in \mathbb{Z}_q$. They consider the polynomial rings $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$ and $R_q = R/qR$ (replacing the exponent space $\mathbb{Z}_p$). They generate a small secret $g \in R$ and let $\mathcal{I} = \langle g\rangle$ be the principal ideal over R generated by g. They also sample a uniform $z \in R_q$ which stays secret. The "plaintext" is an element of $R/\mathcal{I}$, and is encoded via a division by z in R_q: to encode a coset of $R/\mathcal{I}$, return $[c/z]_q$, where c is an arbitrary small coset representative. In practice, as g is hidden, they give another public parameter y, which is an encoding of 1, and the encoding of the coset is computed as $[e \cdot y]_q$, where e is a small coset representative (possibly different from c). But, as opposed to multilinear maps, their graded encoding scheme uses the notion of *encoding level*: the plaintext e is a level-0 encoding, the encoding $[c/z]_q$ is a level-1 encoding, and at level i, an encoding of $e + \mathcal{I}$ is given by $[c/z^i]_q = [e \cdot y^i]_q$. These encodings are both additively and multiplicatively homomorphic, up to a limited number of operations. More precisely, a product of i level-1 encodings is a level-i encoding. One can multiply any number of encodings up to κ, instead of exactly κ in multilinear maps (the parameter κ is called the multilinearity parameter).

The authors of [9] introduced new hardness assumptions: the Graded Decisional Diffie-Hellman (GDDH) and its computational variant (GCDH). These are natural analogues of the Diffie-Hellman problems from group-based cryptography. To ensure their hardness, and hence the security of the cryptographic constructions, the second main difference with multilinear maps is the randomization of the encodings. The principle is as follows: first some level-1 encodings of 0, called $\{x_j = [b_j/z]_q\}_{j \le m_r}$, are given as part of the public parameters; then, to randomize a level-1 encoding $u' = [e \cdot y]_q$, one outputs $u = [u' + \sum_j \rho_j x_j]_q = [c/z]_q$ with $c = c' + \sum_j \rho_j b_j$, where the ρ_j's are sampled from a discrete Gaussian distribution over $\mathbb{Z}$ with deviation parameter σ^*. Without this re-randomization, the encoding u' of e allows e to be efficiently recovered using $u = [u'y^{-1}]_q$. Adding the re-randomization step prevents this division attack, but the statistical properties of the distribution of the re-randomized encoding u remain correlated to some extent with the original encoding u' (for instance, the center of the distribution of c is c', since the distribution of $\sum_j \rho_j b_j$ is known to be centered at 0). This property may allow other attacks that exploit this correlation. The question arises as to how to set the re-randomization parameter σ^* in order to guarantee security against such potential "statistical correlation" attacks – the larger the re-randomization parameters the smaller the correlation, and heuristically the more resistant the scheme is to such attacks. But increasing σ^* impacts the efficiency of the scheme.

In [9], the authors use a "drowning step" to solve this problem. This technique, also called "smudging,"was previously used in other applications [3,10,2,4].Generally, "drowning" consists in hiding a secret vector $\boldsymbol{s} \in \mathbb{Z}^n$ by adding a sufficiently large random noise $\boldsymbol{e} \in \mathbb{Z}^n$ to it, so that the distribution of $\boldsymbol{s} + \boldsymbol{e}$ becomes "almost independent" of $\boldsymbol{s}$. In all of the above applications, to achieve a security level 2^λ (where λ denotes the security parameter), the security analysis requires "almost independent" to be interpreted as "within statistical distance $2^{-\lambda}$ from a distribution that is independent of $\boldsymbol{s}$." In turn, this requirement implies the need for "exponential drowning," i.e., the ratio $\gamma = \|\boldsymbol{e}\|/\|\boldsymbol{s}\|$ between the magnitude of the noise and the magnitude of secret needs to be $2^{\Omega(\lambda)}$. Exponential drowning imposes a severe penalty on the efficiency of these schemes, as their security is related to γ-approximation lattice problems, whose complexity decreases exponentially with $\log \gamma$. As a result, the schemes require a lattice dimension n at least quadratic in λ and key length at least cubic in λ. In summary, the GGH re-randomization step, necessary for its security, is also a primary factor in its inefficiency.

Our contributions. First, we formalize the re-randomization security goal in the GGH construction, that is implicit in the work of [9]. A primary security goal of re-randomization is to guarantee security of the GDDH problem against statistical correlation attacks. Accordingly, we formulate a security goal that captures this security guarantee, by introducing a canonical variant of GDDH, called cGDDH. In this variant, the encodings of some elements are sampled from a canonical distribution whose statistical properties are independent of the encoded elements. Consequently, the canonical problems are by construction not subject to "statistical correlation" attacks. Our re-randomization security goal is formulated as the existence of an efficient computational reduction from the canonical problems to their corresponding non-canonical variants.

Our first main improvement to the GGH scheme relies on a new security analysis of the drowning step in the GGH re-randomization algorithm. We show that our re-randomization security goal can be satisfied *without* "exponential drowning," thus removing the main efficiency bottleneck. Namely, our analysis provides a re-randomization at security level 2^λ while allowing the use of a re-randomization deviation parameter σ^* that only drowns the norm of the randomness offset $r' \in \mathcal{I}$ (from the original encoding to be re-randomized) by a *polynomial* (or even constant) drowning ratio $\gamma = \lambda^{O(1)}$ (rather than $\gamma = 2^{\Omega(\lambda)}$, as needed in the analysis of [9]). However, our analysis only works for the search variant of the Graded Diffie-Hellman problem. Fortunately, we show that the two flagship applications of the GGH scheme – the N-party Key Agreement and the Attribute Based Encryption – can be modified to rely on this computational assumption (in the random oracle model).

Our second main improvement of the re-randomization process is to decrease m_r, the number of encodings of 0 needed, from $\Omega(n \log n)$ to 2. We achieve this result by presenting a new discrete Gaussian Leftover Hash Lemma (LHL) over algebraic rings. In [9], the authors apply the discrete Gaussian LHL from [1] to show that the distribution of the sum $\sum_{j \leq m_r} \rho_j r_j$ is close to a discrete Gaussian on the ideal $\mathcal{I}$. Our improvement consists in sampling the randomizers ρ_j as elements of the full n-dimensional ring R, rather than just from $\mathbb{Z}$. Since each randomizer now has n times more entropy than before, one may hope to obtain a similar LHL result as in [1] while reducing m_r by a factor $\approx n$. However, as the designers of the GGH scheme notice in [9, Se. 6.4], the proof techniques from [1] do not seem to immediately carry over to our "algebraic ring" LHL setting. Our new LHL over rings resolves this problem.

These contributions allow us to decrease the bit size of the public parameters from $O(\kappa^3 \lambda^5 \log(\kappa\lambda))$ for the GGH scheme to $O(\kappa^2 \lambda \log^2(\kappa\lambda))$ for GGHLite, for security level 2^λ for the graded Diffie-Hellman problem.

Technical overview. Our first main result is to reduce the size of the parameter σ^* in the re-randomization process. Technically, our improved analysis of drowning is obtained by using the *Rényi divergence* (RD) to replace the conventional statistical distance (SD) as a measure of distribution closeness. The RD was already exploited in a different context in [13, Claim 5.11], to show the hardness of Ring-LWE. Here, we use the RD to decrease the amount of drowning, by bounding the RD between a discrete Gaussian distribution and its offset.This suffices for relating the hardness of the search problems using these encoding distributions, even though the SD between the distributions is non-negligible. The technique does not seem to easily extend to the decision problems, as RD induces a multiplicative relationship between success probabilities, rather than an additive relationship as SD does.

Our second main result is a new LHL over the ring R. We now briefly explain this result and its proof. For a fixed $X = [x_1, x_2] \in R^2$, with each x_i sampled from $D_{R,s}$, our goal is to study the distribution $\widetilde{\mathcal{E}}_{X,s} = x_1 \cdot D_{R,s} + x_2 \cdot D_{R,s}$. In particular, we prove that $\widetilde{\mathcal{E}}_{X,s}$ is statistically close to $D_{\mathbb{Z}^n, sX^T}$. For this, we adapt the proof of the LHL in [1]: we follow a similar series of steps, but the proofs of these steps differ technically, as we exploit the ring structure.

We first show that $X \cdot R^2 = R$, except with some constant probability < 1. For this, we adapt a result from [23] on the probability that two Gaussian samples of R are coprime. Note that in contrast to the LHL over $\mathbb{Z}$ in [1], in our setting the probability that $X \cdot R^2 \neq R$ is non-negligible. This is unavoidable with the ring $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$, since each random element of R falls in the ideal $\langle x+1\rangle$ with probability $\approx 1/2$, both x_1 and x_2 (and hence the ideal they generate) get "stuck" in $\langle x+1\rangle$ with probability $\approx 1/4$. However, the probability of this bad event is bounded away from 1 by a constant and thus we only need a constant number of trials on average with random X's to obtain a good X by rejection.

Then, we define the orthogonal R-module $A_X = \{\boldsymbol{v} \in R^2 : X \cdot \boldsymbol{v} = 0\}$, and apply a directly adapted variant of [1, Le. 10] to show that if the parameter s is larger than the smoothing parameter $\eta_\varepsilon(A_X)$ (with A_X viewed as an integral lattice), then the SD between $\widetilde{\mathcal{E}}_{X,s}$ and the ellipsoidal Gaussian $D_{\mathbb{Z}^n, sX^T}$ is bounded by 2ε. We finally show that this condition on the smoothing parameter of A_X holds. For this, we observe that the Minkowski minima of the lattice A_X are equal, due to the R-module structure of A_X. This allows us to bound the last minimum from above using Minkowski's second theorem. A similar approach was previously used (e.g., in [12]) to bound the smoothing parameter of ideal lattices.

Notation. A function $f(\lambda)$ is said negligible if it is $\lambda^{-\omega(1)}$. For an integer q, we let $\mathbb{Z}_q$ denote the ring of integers modulo q. The notation $[\cdot]_q$ means that all operations within the square brackets are performed modulo q. We choose $n \geq 4$ as a power of 2, and let K and R respectively denote the polynomial ring $\mathbb{Q}[X]/\langle x^n + 1\rangle$ and $\mathbb{Z}[X]/\langle x^n + 1\rangle$. The rings K and R are isomorphic to the cyclotomic field of order $2n$ and its ring of integers, respectively. For an integer q, we let R_q denote the ring $\mathbb{Z}_q[x]/\langle x^n + 1\rangle \simeq R/qR$. For $z \in R$ we denote by $\mathrm{MSB}_\ell(z) \in \{0,1\}^{\ell \cdot n}$ the ℓ most-significant bits of each of the n coefficients of z. Vectors are denoted in bold. For $\boldsymbol{b} \in \mathbb{R}^d$ (resp. $g \in K$), we let $\|\boldsymbol{b}\|$ (resp. $\|g\|$) denote its Euclidean norm (resp. norm of its coefficient vector). The uniform distribution on finite set E is denoted by $U(E)$. The statistical distance (SD) between distributions D_1 and D_2 over a countable domain E is $\frac{1}{2}\sum_{x\in E}|D_1(x)-D_2(x)|$. For a function f over a countable domain E, we let $f(E)=\sum_{x\in E}f(x)$. Let $X\in\mathbb{R}^{m\times n}$ be a rank-n matrix and $U_X=\{\|X\boldsymbol{u}\|:\boldsymbol{u}\in\mathbb{R}^n,\|\boldsymbol{u}\|=1\}$. The smallest (resp. largest) singular value of X is denoted by $\sigma_n(X)=\inf(U_X)$ (resp. $\sigma_1(X)=\sup(U_X)$).

Remark. Due to lack of space, some contents have been postponed to the full version of this paper, available from the webpages of the authors.

2 Preliminaries

Lattices. We refer to [14,17] for introductions to the computational aspects of lattices. A d-dimensional *lattice* $\Lambda\subseteq\mathbb{R}^n$ is the set of all integer linear combinations $\sum_{i=1}^{d}x_i\boldsymbol{b}_i$ of some linearly independent vectors $\boldsymbol{b}_i\in\mathbb{R}^n$. The determinant $\det(\Lambda)$ is defined as $\sqrt{\det(B^TB)}$, where $B=(\boldsymbol{b}_i)_i$ is any such *basis* of Λ. For $i\le d$, the ith minimum $\lambda_i(\Lambda)$ is the smallest r such that Λ contains i linearly independent vectors of norms $\le r$.

Gaussian distributions. For a rank-n matrix $S\in\mathbb{R}^{m\times n}$ and a vector $\boldsymbol{c}\in\mathbb{R}^n$, the *ellipsoid* Gaussian distribution with parameter S and center $\boldsymbol{c}$ is defined as: $\forall\boldsymbol{x}\in\mathbb{R}^n,\rho_{S,\boldsymbol{c}}(x)=\exp(-\pi(\boldsymbol{x}-\boldsymbol{c})^T(S^TS)^{-1}(\boldsymbol{x}-\boldsymbol{c}))$. Note that $\rho_{S,\boldsymbol{c}}(x)=\exp(-\pi\|(S^T)^{\dagger}(\boldsymbol{x}-\boldsymbol{c})\|)$, where $X^{\dagger}$ denotes the pseudo-inverse of X. The *ellipsoid* discrete Gaussian distribution over a coset $\Lambda+z$ of a lattice Λ, with parameter S and center $\boldsymbol{c}$ is defined as: $\forall\boldsymbol{x}\in\Lambda+z,D_{\Lambda+z,S,\boldsymbol{c}}=\rho_{S,\boldsymbol{c}}(\boldsymbol{x})/\rho_{S,\boldsymbol{c}}(\Lambda)$.

Smoothing parameter. Introduced by [15], the *smoothing parameter* $\eta_\varepsilon(\Lambda)$ of an n-dimensional lattice Λ and a real $\varepsilon>0$ is defined as the smallest s such that $\rho_{1/s}(\Lambda^*\setminus\{0\})\le\varepsilon$. We use the following properties.

Lemma 2.1 ([15, Le. 3.3]). *Let Λ be an n-dimensional lattice and $\varepsilon>0$. Then $\eta_\varepsilon(\Lambda)\le\sqrt{\ln(2n(1+1/\varepsilon))/\pi}\cdot\lambda_n(\Lambda)$.*

Lemma 2.2 ([1, Le. 3]). *For a rank-n lattice Λ, constant $0<\varepsilon<1$, vector $\boldsymbol{c}$ and matrix S with $\sigma_n(S)\ge\eta_\varepsilon(\Lambda)$, if $\boldsymbol{x}$ is sampled from $D_{\Lambda,S,\boldsymbol{c}}$ then $\|\boldsymbol{x}\|\le\sigma_1(S)\sqrt{n}$, except with probability $\le\frac{1+\varepsilon}{1-\varepsilon}\cdot2^{-n}$.*

Algebraic number rings and ideal lattices. For $g,x\in R$, we let $[x]_g$ denote the reduction of x modulo the principal ideal $I=\langle g\rangle$ with respect to the $\mathbb{Z}$-basis $(g,x\cdot g,\ldots,x^{n-1}\cdot g)$, i.e., $[x]_g$ is the unique element of R in $\mathcal{P}_g=\{\sum_{i=0}^{n-1}c_ix^ig:c_i\in[-1/2,1/2)\cap\mathbb{R}\}$ such that $x-[x]_g\in\langle g\rangle$. The set $\mathcal{P}_g\cap R$ is a set of unique representatives of the cosets of I in R, that make up the quotient ring R/I. To use our improved drowning lemma in Section 4, we need a lower bound on the last singular value $\sigma_n(\mathrm{rot}(b))$ of the matrix $\mathrm{rot}(b) \in \mathbb{Z}^{n \times n}$ corresponding to the map $x \mapsto b \cdot x$ over R, for a Gaussian distributed $b \hookleftarrow D_{I,\sigma}$. In the following, and in the rest of the paper, we abuse notation and write b for this matrix.

Lemma 2.3 (Adapted from [23, Le. 4.1]). *Let $R = \mathbb{Z}^n[x]/(x^n+1)$ for n a power of 2. For any ideal $I \subseteq R$, $\delta \in (0,1)$, $t \geq \sqrt{2\pi}$ and $\sigma \geq \frac{t}{\sqrt{2\pi}} \cdot \eta_\delta(I)$, we have:*

$$\Pr_{b \hookleftarrow D_{I,\sigma}}\left[\|b^{-1}\| \geq \frac{t}{\sigma\sqrt{n/2}}\right] \leq \Pr_{b \hookleftarrow D_{I,\sigma}}\left[\sigma_n(b) \leq \frac{\sigma\sqrt{n/2}}{t}\right] \leq \frac{1+\delta}{1-\delta}\frac{n\sqrt{2\pi \mathrm{e}}}{t}.$$

3 GGH and its re-randomization procedure

In this section, we recall the Garg et al. scheme from [9], and its related hard problems. We then discuss the re-randomization step of the scheme and explain what should be expected from it, in terms of security. This security requirement is unclear in [9] and [1]. We formulate it precisely. This will drive our re-randomization design in the following sections.

3.1 The GGH scheme

We recall the GGH scheme in Figure 1. We present it here in a slightly more general form than [9]: we leave as a parameter the distribution χ_k of the re-randomization coefficients ρ_j for a level-k encoding (for any $k \leq \kappa$). In the original GGH scheme, we have $\chi_k = D_{\mathbb{Z},\sigma_k^*}$ for some σ_k^*'s, i.e., the ρ_j's are integers sampled from a discrete Gaussian distribution. Looking ahead, in Section 5, we analyze a more efficient variant, in which $\chi_k = D_{R,\sigma_k^*}$, so that the ρ_j's belong to R.

The aim of isZero is to test whether the input $u = [c/z^\kappa]_q$ is a level-κ encoding of 0 or not, i.e., whether $c = g \cdot r$ for some $r \in R$. The following conditions ensure correctness of isZero, when $\chi_k = D_{\mathbb{Z},\sigma_k^*}$ (for all $k \leq \kappa$): the first one implies that false negatives do not exist (if u is level-κ encoding of 0, then isZero(u) returns 1), whereas the second one implies that false positives occur with negligible probability.

$$q > \max((n\ell_{g^{-1}})^8, ((m_r+1) \cdot n\sigma_1^*\sigma')^{8\kappa}) \qquad (1)$$

$$q > (2n\sigma)^4. \qquad (2)$$

The aim of ext is to extract a quantity from its input $u = [c/z^\kappa]_q$ that depends only on the encoded value $[c]_g$, but not on the randomizers. To

- **Instance generation** $\mathsf{InstGen}(1^\lambda, 1^\kappa)$: Given security parameter λ and multilinearity parameter κ, determine scheme parameters n, q, m_r, σ, σ', $\ell_{g^{-1}}$, ℓ, based on the scheme analysis. Then proceed as follows:
  - Sample $g \hookleftarrow D_{R,\sigma}$ until $\|g^{-1}\| \leq \ell_{g^{-1}}$ and $\mathcal{I} = \langle g \rangle$ is a prime ideal. Define encoding domain $R_g = R/\langle g \rangle$.
  - Sample $z \hookleftarrow U(R_q)$.
  - Sample a level-1 encoding of 1: set $y = [a \cdot z^{-1}]_q$ with $a \hookleftarrow D_{1+I,\sigma'}$.
  - For $k \leq \kappa$, sample m_r level-k encodings of 0: set $x_j^{(k)} = [b_j^{(k)} \cdot z^{-k}]_q$ with $b_j^{(k)} \hookleftarrow D_{I,\sigma'}$ for all $j \leq m_r$.

    (Note that $a = 1 + gr_y$ and $b_j^{(k)} = gr_j^{(k)}$ for some $r_y, r_j^{(k)} \in R$.)
  - Sample $h \hookleftarrow D_{R,\sqrt{q}}$ and define the zero-testing parameter $p_{zt} = [\frac{h}{g} z^\kappa]_q \in R_q$.
  - Return public parameters $\mathrm{par} = (n, q, y, \{x_j^{(k)}\}_{j \leq m_r, k \leq \kappa})$ and p_{zt}.
- **Level-0 sampler** $\mathsf{samp}(\mathrm{par})$: Sample $e \hookleftarrow D_{R,\sigma'}$ and return e.

  (Note that $e = e_L + ge_H$ for some unique coset representative $e_L \in \mathcal{P}_g$, and some $e_H \in R$.)
- **Level-k encoding** $\mathsf{enc}_k(\mathrm{par}, e)$: Given level-0 encoding $e \in R$ and parameters par:
  - Encode e at level k: Compute $u' = [e \cdot y^k]_q$.
  - Re-randomize: Sample $\rho_j \hookleftarrow \chi_k$ for $j \leq m_r$ and return $u = [u' + \sum_{j=1}^{m_r} \rho_j x_j^{(k)}]_q$.

    (Note that $u' = [c'/z^k]_q$ with $c' \in e_L + I$ and $u = [(c' + \sum_j \rho_j b_j^{(k)})/z^k]_q$.)
- **Adding encodings** add: Given level-k encodings $u_1 = [c_1/z^k]_q$ and $u_2 = [c_2/z^k]_q$:
  - Return $u = [u_1 + u_2]_q$, a level-k encoding of $[c_1 + c_2]_g$.
- **Multiplying encodings** mult: Given level-k_1 encoding $u_1 = [c_1/z^{k_1}]_q$ and a level-k_2 encoding $u_2 = [c_2/z^{k_2}]_q$:
  - Return $u = [u_1 \cdot u_2]_q$, a level-$(k_1 + k_2)$ encoding of $[c_1 \cdot c_2]_g$.
- **Zero testing at level** κ $\mathsf{isZero}(\mathrm{par}, p_{zt}, u)$: Given a level-$\kappa$ encoding $u = [c/z^\kappa]_q$, return 1 if $\|[p_{zt} u]_q\|_\infty < q^{3/4}$ and 0 else.

  (Note that $[p_{zt} \cdot u]_q = [hc/g]_q$.)
- **Extraction at level** κ $\mathsf{ext}(\mathrm{par}, p_{zt}, u)$: Given a level-$\kappa$ encoding $u = [c/z^\kappa]_q$, return $v = \mathrm{MSB}_\ell([p_{zt} \cdot u]_q)$.

  (Note that if $c = [c]_g + gr$ for some $r \in R$, then $v = MSB_\ell(\frac{h}{g}([c]_g + gr)) = MSB_\ell(\frac{h}{g}[c]_g + hr)$, which is equal to $MSB_\ell(\frac{h}{g}[c]_g)$, with probability $1 - \lambda^{-\omega(1)}$.)

Fig. 1. The GGH graded encoding scheme.

avoid trivial solutions, one requires that this extracted value has min-entropy $\geq 2\lambda$ (if that is the case, then one can obtain a uniform distribution on $\{0,1\}^{\lambda}$, using a strong randomness extractor). The following two inequalities guarantee these properties, when $\chi_k = D_{\mathbb{Z},\sigma_k^*}$ (for all k). The first one implies that $\varepsilon_{ext} = \Pr[\mathsf{ext}(u) \neq \mathsf{ext}(u')]$ is negligible, when u and u' encode the same value $[c]_g$, whereas the second one provides large min-entropy.

$$1/4 \log q - \log(\frac{2n}{\varepsilon_{ext}}) \ \geq \ \ell \ \geq \ \log(\frac{n\sigma}{8}). \tag{3}$$

3.2 The GDDH, GCDH and Ext-GCDH problems

The computational problems that are required to be hard for the GGH scheme depend on the application. Here we recall the definitions of the Graded Decisional and Computational Diffie-Hellman (GDDH and GCDH) problems from [9]. We introduce another natural variant that we call the Extraction Graded Computational Diffie-Hellman (Ext-GCDH), in which the goal is to compute the extracted string of a Diffie-Hellman encoding.

Definition 3.1 (GCDH/Ext-GCDH/GDDH). *The problems GCDH, Ext-GCDH and GDDH are defined as follows with respect to experiment of Figure 2:*[3]

- ***κ-graded CDH problem (GCDH):*** *On inputs* par*, p_{zt} and the u_i's of Step 2, output a level-κ encoding of $\prod_{i\geq 0} e_i + \mathcal{I}$, i.e., $w \in R_q$ such that $\|[p_{zt}(v_C - w)]_q\| \leq q^{3/4}$.*
- ***Extraction κ-graded CDH problem (Ext-GCDH):*** *On inputs* par*, p_{zt} and the u_i's of Step 2, output the extracted string for a level-κ encoding of $\prod_{i\geq 0} e_i + \mathcal{I}$, i.e., $w = \mathsf{ext}(\mathrm{par}, p_{zt}, v_C) = MSB_{\ell}([p_{zt} \cdot v_C]_q)$.*
- ***κ-graded DDH problem (GDDH):*** *Distinguish between v_D and v_R, i.e., between the distributions $\mathcal{D}_{DDH} = \{\mathrm{par}, p_{zt}, (u_i)_{0\leq i\leq\kappa}, v_D\}$ and $\mathcal{D}_R = \{\mathrm{par}, p_{zt}, (u_i)_{0\leq i\leq\kappa}, v_R\}$.*

Ext-GCDH is at least as hard as GDDH: given v_x with $x \in \{\mathrm{DDH}, \mathrm{R}\}$, use the Ext-GCDH oracle to compute $w = \mathsf{ext}(\mathrm{par}, p_{zt}, v_C)$. Nevertheless, we show (see full version) that it suffices for instantiating, in the random oracle model, at least some of the interesting applications of graded encoding schemes, at a higher efficiency than the instantiations of [9] based on GDDH.

[3] Note that we use a slightly different process from [9], by adding a re-randomization to the element v_D. Without it, there exists a "division attack" against GDDH.

```
Given parameters $\lambda, n, q, m_r, \kappa, \sigma'$, proceed as follows:

1. Run $\mathsf{InstGen}(1^n, 1^\kappa)$ to get
   $\mathrm{par} = (n, q, y, \{x_j^{(k)}\}_{j,k})$ and $p_{zt}$.
2. For $i = 0, \ldots, \kappa$:
    - Sample $e_i \hookleftarrow D_{R,\sigma'}$, $f_i \hookleftarrow D_{R,\sigma'}$,
    - Set $u_i = [e_i \cdot y + \sum_j \rho_{ij} x_j]_q$
      with $\rho_{ij} \hookleftarrow \chi_1$ for all $j$.
3. Set $u^* = \left[\prod_{i=1}^{\kappa} u_i\right]_q$.
4. Set $v_C = [e_0 u^*]_q$.
5. Sample $\rho_j \hookleftarrow \chi_\kappa$ for all $j$,
   set $v_D = [e_0 u^* + \sum_j \rho_j x_j^{(\kappa)}]_q$.
6. Set $v_R = [f_0 u^* + \sum_j \rho_j x_j^{(\kappa)}]_q$.
```

Fig. 2. The GGH security experiment.

```
Given parameters $\lambda, n, q, m_r, \kappa, (\sigma_k^*)_{k \le \kappa}$, proceed as follows:

1. Run $\mathsf{InstGen}(1^n, 1^\kappa)$ to get
   $\mathrm{par} = (n, q, y, \{x_j^{(k)}\}_{j,k})$ and $p_{zt}$.
   Write $x_j^{(k)} = [b_j^{(k)} z^{-k}]_q$ and
   $B^{(k)} = [b_1^{(k)}, \cdots, b_{m_r}^{(k)}] \in \mathcal{I}^{m_r}$.
2. For $i = 0, \ldots, \kappa$:
    - Sample $e_i \hookleftarrow U(R_g)$, $f_i \hookleftarrow U(R_g)$,
    - Set $u_i = [c_i z^{-1}]_q \hookleftarrow D_{can}^{(1)}(e_i)$
      with $c_i \hookleftarrow D_{\mathcal{I}+e_i, \sigma_1^*(B^{(1)})^T}$.
3. Set $u^* = \left[\prod_{i=1}^{\kappa} u_i\right]_q$.
4. Set $v_C = [e_0 u^*]_q$.
5. Set $v_D = [c_D \cdot z^{-\kappa}]_q \hookleftarrow D_{can}^{(\kappa)}(\prod_{i=0}^{\kappa} e_i)$,
   with $c_D \hookleftarrow D_{\mathcal{I}+\prod_{i=0}^{\kappa} e_i, \sigma_\kappa^*(B^{(\kappa)})^T}$.
6. Set $v_R = [c_R \cdot z^{-\kappa}]_q \hookleftarrow D_{can}^{(\kappa)}(f_0 \prod_{i=1}^{\kappa} e_i)$,
   with $c_R \hookleftarrow D_{\mathcal{I}+f_0 \prod_{i=1}^{\kappa} e_i, \sigma_\kappa^*(B^{(\kappa)})^T}$.
```

Fig. 3. The canonical security experiment.

3.3 The GGH re-randomization security requirement

The encoding re-randomization step in the GGH scheme is necessary for the hardness of the problems above. In [9], Garg et al. imposed the informal requirement that the re-randomization process "erases" the structure of the input encoding, while preserving the encoded coset. In setting parameters, they interpreted this requirement in the following natural way.

Definition 3.2 (Strong re-randomization security requirement). *Let $u' = [c'/z^k]_q$, with $c' = e_L + gr'$ be a fixed level-k encoding of $e_L \in R_g$, and let $u = [u' + \sum_j \rho_j x_k^{(j)}]_q = [c/z^k]_q$ with $c = e_L + gr$ and $r = r' + \sum_j \rho_j r_j^{(k)}$ be the re-randomized encoding, with $\rho_j \hookleftarrow \chi_k$ for $j \le m_r$. Let $D_u^{(k)}(e_L, r')$ denote the distribution of u (over the randomness of ρ_j's), parameterized by (e_L, r') and let $D_{\mathrm{can}}^{(k)}(e_L)$ denote some canonical distribution, parameterized by e_L, that is independent of r'. Then we say that the* strong *re-randomization security requirement is satisfied at level k with respect to $D_{\mathrm{can}}^{(k)}(e_L)$ and encoding norm $\gamma^{(k)}$ if $\Delta(D_u^{(k)}(e_L, r'), D_{\mathrm{can}}^{(k)}(e_L)) \le 2^{-\lambda}$ for any $u' = [c'/z^k]_q$ with $\|c'\| \le \gamma^{(k)}$.*

The authors of [9] argued that with $\chi_k = D_{\mathbb{Z}, \sigma_k^*}$ (for $k \le \kappa$) and a "drowning ratio" $\sigma_k^*/\|r'\|$ exponential in security parameter λ, the distribution $D_u^{(k)}(e_L, r')$ is within negligible statistical distance to the canonical distribution $D_{\text{can}}^{(k)}(e_L) = [D_{\mathcal{I}+e_L,\sigma_k^*(B^{(k)})^T} \cdot z^{-k}]_q$. This requirement may be stronger than needed. Accordingly, we now clarify the desired goal.

3.4 Our security goal: canonical assumptions

We formalize a re-randomization security goal to capture a security guarantee against "statistical correlation" attacks on GCDH/Ext-GCDH/GDDH. We define *canonical variants* cGCDH/Ext-cGCDH/cGDDH of GCDH/Ext-GCDH/GDDH, using Figure 3. The main difference with Figure 2 is that the encodings $u_i = [c_i/z]_q$ of the hidden elements e_i, are sampled from a canonical distribution $D_{\text{can}}^{(1)}(e_i)$, parameterized by e_i, whose statistical parameters are independent of the encoded coset e_i, so that it is "by construction" immune against statistical correlation attacks. In particular, in the canonical distribution $D_{\text{can}}^{(1)}(e_i)$ that we use, c_i is sampled from a discrete Gaussian distribution $D_{\mathcal{I}+e_i,\sigma_1^*(B^{(1)})^T}$ (over the choice of the randomization, for a fixed e_i), whose statistical parameters such as center (namely 0) and deviation matrix $\sigma_1^*(B^{(1)})^T$ are independent of e_i. The only dependence this distribution has on the encoded element e_i is via its support $\mathcal{I} + e_i$.

We believe the canonical problems are cleaner and more natural than the non-canonical variants, since they decouple the re-randomization aspect from the rest of the computational problem. As a further simplification, the canonical variants also have their level-0 elements e_i distributed uniformly on R_g (rather than as reductions mod $\mathcal{I}$ of Gaussian samples).

Definition 3.3 (cGCDH/Ext-cGCDH/cGDDH). *The canonical problems cGCDH, Ext-cGCDH and cGDDH are defined as follows with respect to the experiment of Figure 3 and canonical encoding distribution $D_{can}^{(k)}(e)$ (parameterized by encoding level k and encoded element e):*

- ***cGCDH:*** *On inputs* par*, p_{zt} and the u_i's, output $w \in R_q$ such that $\|[p_{zt}(v_C - w)]_q\| \le q^{3/4}$.*
- ***Ext-cGCDH:*** *On inputs* par*, p_{zt} and the u_i's, output:* $w = \mathsf{ext}(\text{par}, p_{zt}, v_C) = MSB_\ell([p_{zt} \cdot v_C]_q)$.
- ***cGDDH:*** *Distinguish between $\mathcal{D}_{DDH} = \{\text{par}, p_{zt}, (u_i)_{0\le i\le\kappa}, v_D\}$ and $\mathcal{D}_R = \{\text{par}, p_{zt}, (u_i)_{0\le i\le\kappa}, v_R\}$.*

Remark. One could consider alternative definitions of natural canonical encoding distributions besides the one we adopt here (see full paper for examples for which our results also apply).

Given the canonical problems on whose hardness we wish to rely, our security goal for re-randomization with respect to the GCDH (resp. Ext-GCDH/GDDH) problems can now be easily formulated: hardness of the latter should be implied by hardness of the former.

Definition 3.4 (Re-randomization security goal). *We say that the re-randomization security goal is satisfied with respect to GCDH (resp. Ext-GCDH/GDDH) if any adversary against GCDH (resp. Ext-GCDH/GDDH) with run-time $T = O(2^{\lambda})$ and advantage $\varepsilon = \Omega(2^{-\lambda})$ can be used to construct an adversary against cGCDH (resp. Ext-cGCDH/cGDDH) with run-time $T' = \mathrm{poly}(T, \lambda)$ and advantage $\varepsilon' = \Omega(\mathrm{poly}(\varepsilon, \lambda))$.*

4 Polynomial drowning via Rényi divergence

In this section, we present our first result towards our improvement of the GGH scheme re-randomization. It shows that one may reduce the re-randomization "drowning" ratio $\sigma_k^*/\|r'\|$ from exponential to polynomial in the security parameter λ. Although the SD between the re-randomized encoding distribution D_1 (essentially a discrete Gaussian with an added offset vector r') and the desired canonical encoding distribution D_2 (a discrete Gaussian without an added offset vector) is then non-negligible, we show that these encoding distributions are still sufficiently close with respect to an alternative closeness measure to the SD, in the sense that switching between them preserves the success probability of any search problem adversary receiving these encodings as input, up to a small multiplicative constant. This allows us to show that our re-randomization goal is satisfied for the search problems GCDH and Ext-GCDH.

Technically, the closeness measure we study is the *Rényi divergence* $R(D_1\|D_2)$ between the distributions D_1 and D_2, defined as the expected value of $D_1(r)/D_2(r)$ over the randomness of r sampled from D_1 (for brevity we will call $R(D_1\|D_2)$ the RD between D_1 and D_2). Intuitively, the RD is an alternative to SD as measure of distribution closeness, where we replace the *difference* between the distributions in SD, by the *ratio* of the distributions in RD. Accordingly, one may hope RD to have analogous properties to SD, where addition in the property of SD is replaced by multiplication in the analogous property of RD. Remarkably, this holds true in some sense, and we explore some of this below. In particular, a very important property of the SD is that for any two distributions D_1, D_2 on space X, and any event $E \subseteq X$, we have $D_1(E) \geq D_2(E) - \Delta(D_1, D_2)$. Lyubashevsky et al. [13] observed an analogous property of the RD that follows roughly the above intuition:

$D_1(E) \geq D_2(E)^2/R(D_1\|D_2)$. The latter property implies that as long as $R(D_1\|D_2)$ is bounded as $\mathrm{poly}(\lambda)$, any event of non-negligible probability $D_2(E)$ under D_2 will also have non-negligible probability $D_1(E)$ under D_1. We show that for our offset discrete Gaussian distributions D_1, D_2 above, we have $R(D_1\|D_2) = O(\mathrm{poly}(\lambda))$, if $\sigma_k^*/\|r'\| = \Omega(\mathrm{poly}(\lambda))$, as required for our re-randomization security goal.

The Rényi divergence (RD) and its properties. We review the RD [18,8] and some of its properties. For convenience, our definition of the RD is the exponential of the usual definition used in information theory [8], and coincides with a discrete version of the quantity R defined for continuous density functions in [13, Claim 5.11].

For any two discrete probability distributions P and Q such that $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$ over a domain X and $\alpha > 1$, we define the Rényi Divergence of orders α and ∞ by

$$R_\alpha(P\|Q) = \left(\sum_{x\in X} \frac{P(x)^\alpha}{Q(x)^{\alpha-1}}\right)^{\frac{1}{\alpha-1}} \quad \text{and} \quad R_\infty(P\|Q) = \max_{x\in X} \frac{P(x)}{Q(x)},$$

with the convention that the fraction is zero when both numerator and denominator are zero. A convenient choice for computations (as also used in [13]) is $\alpha = 2$, in which case we omit α. Note that $R_\alpha(P\|Q)^{\alpha-1} = \sum_x P(x) \cdot (P(x)/Q(x))^{\alpha-1} \leq R_\infty(P\|Q)^{\alpha-1}$. We list several properties of the RD that can be considered the multiplicative analogues of those of the SD. The following lemma is proven in the full version.

Lemma 4.1. *Let P_1, P_2, P_3 and Q_1, Q_2, Q_3 denote discrete distributions on a domain X and let $\alpha \in (1, \infty]$. Then the following properties hold:*

- ***Log. Positivity:*** $R_\alpha(P_1\|Q_1) \geq R_\alpha(P_1\|P_1) = 1$.
- ***Data Processing Inequality:*** $R_\alpha(P_1^f\|Q_1^f) \leq R_\alpha(P_1\|Q_1)$ *for any function f, where P_1^f (resp. Q_1^f) denotes the distribution of $f(y)$ induced by sampling $y \hookleftarrow P_1$ (resp. $y \hookleftarrow Q_1$).*
- ***Multiplicativity:*** *Let P and Q denote any two distributions of a pair of random variables (Y_1, Y_2) on $X \times X$. For $i \in \{1, 2\}$, assume P_i (resp. Q_i) is the marginal distribution of Y_i under P (resp. Q), and let $P_{2|1}(\cdot|y_1)$ (resp. $Q_{2|1}(\cdot|y_1)$) denote the conditional distribution of Y_2 given that $Y_1 = y_1$. Then we have:*
  - $R_\alpha(P\|Q) = R_\alpha(P_1\|Q_1) \cdot R_\alpha(P_2\|Q_2)$ *if Y_1 and Y_2 are independent.*
  - $R_\alpha(P\|Q) \leq R_\infty(P_1\|Q_1) \cdot \max_{y_1\in X} R_\alpha(P_{2|1}(\cdot|y_1)\|Q_{2|1}(\cdot|y_1))$.
- ***Weak Triangle Inequality:*** *We have:*

$$R_\alpha(P_1\|P_3) \leq \begin{cases} R_\alpha(P_1\|P_2) \cdot R_\infty(P_2\|P_3), \\ R_\infty(P_1\|P_2)^{\frac{\alpha}{\alpha-1}} \cdot R_\alpha(P_2\|P_3). \end{cases}$$

- R_∞ ***Triangle Inequality:*** *If* $R_\infty(P_1\|P_2)$ *and* $R_\infty(P_2\|P_3)$ *are defined, then* $R_\infty(P_1\|P_3) \leq R_\infty(P_1\|P_2) \cdot R_\infty(P_2\|P_3)$.
- ***Probability Preservation:*** *Let* $A \subseteq X$ *be an arbitrary event. Then* $Q_1(A) \geq P_1(A)^{\frac{\alpha}{\alpha-1}}/R_\alpha(P_1\|Q_1)$.

We note that the RD does not satisfy the (multiplicative) triangle inequality $R(P_1\|P_3) \leq R(P_1\|P_2) \cdot R(P_2\|P_3)$ in general (see [8]), but a weaker inequality holds if one of the pairs of distributions has a bounded R_∞ divergence, as shown above. We also observe that R_∞ *does* satisfy the triangle inequality.

For our re-randomization application, we are interested in the RD between two discrete Gaussians with the same deviation matrix S, that differ by some fixed offset vector d. The following result (proved in the full version) shows that their RD is $O(1)$ if $\sigma_n(S)/\|d\| = \Omega(1)$.

Lemma 4.2. *For any n-dimensional lattice Λ in $\mathbb{R}^n$ and matrix S, let P be the distribution $D_{\Lambda,S,w}$ and Q be the distribution $D_{\Lambda,S,z}$ for some fixed $w, z \in \mathbb{R}^n$. If $w, z \in \Lambda$, let $\varepsilon = 0$. Otherwise, fix $\varepsilon \in (0,1)$ and assume that $\sigma_n(S) \geq \eta_\varepsilon(\Lambda)$. Then $R(P\|Q) \leq \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^2 \cdot \exp\left(2\pi\|w-z\|^2/\sigma_n(S)^2\right)$.*

5 A discrete Gaussian leftover hash lemma over $\boldsymbol{R}$

In this section, we present our second main result for improving the GGH scheme re-randomization algorithm. Recall that the GGH algorithm re-randomizes a level-k encoding u' into $u = [u' + \sum_{j=1}^{m_r} \rho_j x_j^{(k)}]_q$, where the ρ_j's are sampled from $\chi_1 = D_{\mathbb{Z},\sigma_1^*}$ and $x_j^{(k)} = [b_j^{(k)}/z^k]_q = [gr_j^{(k)}/z^k]_q$. To show that the distribution of $\sum_{j=1}^{m_r} \rho_j b_j^{(k)}$ is close to a discrete Gaussian over $\mathcal{I}$, they then apply the discrete Gaussian LHL from [1, Th. 3], using $m_r = \Omega(n \log n)$ fixed elements $b_j^{(k)} \in \mathcal{I}$ that are published obliviously as randomizers "inside" the public zero-encodings $x_j^{(k)}$. We show that it suffices to sample 2 randomizers as elements of the full n-dimensional ring R, rather than just from $\mathbb{Z}$, i.e., we set $\chi_1 = D_{R,\sigma_1^*}$. Our proof follows the same high-level steps as the proof of [1, Th. 3], but differs technically, as explained in the introduction.

For a fixed $X = (x_1, x_2) \in R^2$, we define the distribution $\widetilde{\mathcal{E}}_{X,s} = x_1 D_{R,s} + x_2 D_{R,s}$ as the distribution induced by sampling $\boldsymbol{u} = (u_1, u_2) \in R^2$ from a discrete spherical Gaussian with parameter s, and outputting $y = x_1 u_1 + x_2 u_2$. We prove the following result on $\widetilde{\mathcal{E}}_{X,s}$.

Theorem 5.1. *Let $R = \mathbb{Z}[x]/\langle x^n + 1\rangle$ with n a power of 2 and $\mathcal{I} = \langle g\rangle \subseteq R$, for some $g \in R$. Fix $\varepsilon \in (0, 1/3)$, $X = (x_1, x_2) \in \mathcal{I}^2$ and $s > 0$ satisfying the conditions*

- ***Column span:*** $X \cdot R^2 = \mathcal{I}$.
- ***Smoothing:*** $s \geq \max(\|g^{-1}x_1\|_\infty, \|g^{-1}x_2\|_\infty) \cdot n \cdot \sqrt{\frac{2}{\pi}\log(2n(1 + 1/\varepsilon))}$.

Then, for all $x \in \mathcal{I}$ we have $\widetilde{\mathcal{E}}_{X,s}(x) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot D_{\mathcal{I},sX^T}(x)$. In particular, we have $\Delta(\widetilde{\mathcal{E}}_{X,s}, D_{\mathcal{I},sX^T}) \leq 2\varepsilon$. Finally, if $s \cdot \sigma_n(g^{-1}) \geq 7n^{1.5}\ln^{1.5}(n)$,[4] *$x_1, x_2 \hookleftarrow D_{\mathcal{I},s}$ and n grows to infinity, then the first condition holds with probability $\Omega(1)$.*

We prove this result for $g = 1$, and then we generalize to general g. First, we consider the column span condition.

Lemma 5.2 (Adapted from [23, Le. 4.2 and Le. 4.4]). *Let $S \in \mathbb{R}^{n\times n}$, and $\sigma_n(S) \geq 7n^{1.5}\ln^{1.5}(n)$. For n going to infinity, we have $\Pr_{x_1,x_2\hookleftarrow D_{R,S}}[X \cdot R^2 = R] \geq \Omega(1)$.*

Let $A_X \subseteq \{(v_1, v_2) \in R^2 : x_1v_1 + x_2v_2 = 0\}$ be the 1-dimensional R-module of vectors orthogonal to X. We view A_X as an n-dimensional lattice in $\mathbb{Z}^{2n}$, via the polynomial-to-coefficient-vector mapping.

Lemma 5.3 (Adapted from [1, Le. 10]). *Fix X such that $X \cdot R^2 = R$ and A_X as above. If $s \geq \eta_\varepsilon(A_X)$, then $\widetilde{\mathcal{E}}_{X,s}(z) \in [\frac{1-\varepsilon}{1+\varepsilon}, 1] \cdot D_{\mathbb{Z}^n,sX^T}(z)$ for any $z \in R$.*

We now study the quantity $\eta_\varepsilon(A_X)$. First, we show that all successive Minkowski minima of A_X are equal. This property is inherited from the “equal minima property” of ideal lattices in R.

Lemma 5.4. *Let X and A_X be as above. Then $\lambda_1(A_X) = \cdots = \lambda_n(A_X)$.*

Lemma 5.5. *Let X and A_X be as above. Let $s \geq \max(\|x_1\|_\infty, \|x_2\|_\infty)$. Then we have: $\eta_\varepsilon(A_X) \leq sn \cdot \sqrt{\frac{2}{\pi}\log(2n(1 + 1/\varepsilon))}$.*

Combining the above lemmas, we get Theorem 5.1 for $g = 1$. The general case is proved as follows. The injective map $y \mapsto g \cdot y$ on R takes the distribution $\widetilde{\mathcal{E}}_{\overline{X},s}$ with $\overline{X} = g^{-1} \cdot X$ to the distribution $\widetilde{\mathcal{E}}_{X,s}$, while it takes $D_{R,s\overline{X}^T}$ to $D_{\mathcal{I},sX^T}$, with $I = \langle g\rangle$. The conditions $X \cdot R^2 = \mathcal{I}$ and

[4] By abuse of notation, we identify $g^{-1} \in K$ with the linear map over $\mathbb{Q}^n$ obtained by applying the polynomial-to-coefficient-vector mapping to the map $r \mapsto g^{-1}r$.

$\overline{X} \cdot R^2 = R$ are equivalent. The smoothing condition is satisfied for $\overline{X}$ by the choice of s. Thus we can apply Theorem 5.1 with $g = 1$ to $\widetilde{\mathcal{E}}_{\overline{X},s}$, and conclude by applying the mapping M_g to get the general case of Theorem 5.1. For the very last statement of Theorem 5.1, it suffices to observe that $D_{\mathcal{I},s} = g \cdot D_{R,s(g^{-1})^T}$.[5] □

6 Our improved GGH grading scheme: GGHLite

We are now ready to describe our simpler and more efficient variant of the GGH grading scheme, that we call GGHLite. The scheme is summarized in Figure 4. The modifications from the original GGH scheme consist in:

- Using $m_r = 2$ re-randomization elements x_1, x_2 in the public key, sampling the randomizers ρ_1, ρ_2 from a discrete Gaussian D_{R,σ_1^*} over the whole ring R (rather than from $\mathbb{Z}$), applying our algebraic ring variant of the LHL from Section 5.
- Saving an exponential factor $\approx 2^\lambda$ in the re-randomization parameter σ_1^* by applying the RD bounds from Section 4.

In terms of re-randomization security requirement, we relax the strong SD-based requirement on the original GGH scheme to the following weaker RD-based requirement on GGHLite.

Definition 6.1 (Weak re-randomization security requirement). *Using the notations of Definition 3.2, we say that the* weak *re-randomization security requirement is satisfied at level k with respect to $D_{\mathrm{can}}^{(k)}(e_L)$ and encoding norm $\gamma^{(k)}$ if $R(D_u^{(k)}(e_L, r') \| D_{\mathrm{can}}^{(k)}(e_L)) = O(\mathrm{poly}(\lambda))$ for any $u' = [c'/z^k]_q$ such that $\|c'\| \leq \gamma^{(k)}$.*

We summarize GGHLite in Figure 4, which only shows the algorithms differing from those in the GGH scheme of Figure 1.

Choice of σ, $\ell_{g^{-1}}$ and σ', ℓ_b. The upper bound $\ell_{g^{-1}}$ on $\|g^{-1}\|$ in the rejection test of InstGen can be chosen as small as possible while keeping the rejection probability p_g bounded from 1. According to Lemma 2.3 with $t = 2\sqrt{2\pi \mathrm{e}} n p_g^{-1}$ and $\delta = 1/3$, one can choose

$$\ell_{g^{-1}} = 4\sqrt{\pi \mathrm{e} n}/(p_g \sigma) \quad \text{and} \quad \sigma \geq 2n\sqrt{\mathrm{e}\ln(8n)/\pi}/p_g, \tag{4}$$

to achieve $p_g < 1$. Note that the same choices apply to the GGH scheme: here we have a rigorous bound on p_g instead of the heuristic arguments

[5] With the same abuse of notation as in the previous footnote, for the term $(g^{-1})^T$.

- **Instance generation** $\mathsf{InstGen}(1^\lambda, 1^\kappa)$: Given security parameter λ and multilinearity parameter κ, determine scheme parameters n, q, $m_r = 2$, σ, σ', $\ell_{g^{-1}}$, ℓ_b, ℓ, based on the scheme analysis. Then proceed as follows:
  - Sample $g \hookleftarrow D_{R,\sigma}$ until $\|g^{-1}\| \leq \ell_{g^{-1}}$ and $\mathcal{I} = \langle g \rangle$ is a prime ideal.
  - Sample $z \hookleftarrow U(R_q)$.
  - Sample a level-1 encoding of 1: $y = [a \cdot z^{-1}]_q$ with $a \hookleftarrow D_{1+I,\sigma'}$.
  - For $k \leq \kappa$:
    - Sample $B^{(k)} = (b_1^{(k)}, b_2^{(k)})$ from $(D_{I,\sigma'})^2$. If $\langle b_1^{(k)}, b_2^{(k)} \rangle \neq \mathcal{I}$, or $\sigma_n(\mathrm{rot}(B^{(k)})) < \ell_b$, then re-sample.
    - Define level-k encodings of 0: $x_1^{(k)} = [b_1^{(k)} \cdot z^{-k}]_q$, $x_2^{(k)} = [b_2^{(k)} \cdot z^{-k}]_q$.
  - Sample $h \hookleftarrow D_{R,\sqrt{q}}$ and define the zero-testing parameter $p_{zt} = [\frac{h}{g} z^\kappa]_q \in R_q$.
  - Return public parameters $\mathrm{par} = (n, q, y, \{(x_1^{(k)}, x_2^{(k)})\}_{k \leq \kappa})$ and p_{zt}.
- **Level-k encoding** $\mathsf{enc}_k(\mathrm{par}, e)$: Given level-0 encoding $e \in R$ and parameters par:
  - Encode e at level k: Compute $u' = [e \cdot y^k]_q$.
  - Return $u = [(u' + \rho_1 \cdot x_1^{(k)} + \rho_2 \cdot x_2^{(k)})/z^k]_q$, with $\rho_1, \rho_2 \hookleftarrow D_{R,\sigma_k^*}$.

Fig. 4. The new algorithms of our $\mathsf{GGHLite}$ scheme.

for estimating in $\|g^{-1}\|$ in [9]; however, as in [9], we do not have a rigorous bound on the probability that $\mathcal{I}$ is prime conditioned on this choice.

Let p_b be the rejection probability for the lower bound ℓ_b on $\sigma_n(B^{(k)})$ in the rejection test of $\mathsf{InstGen}$. To keep p_b away from 1, we use that $\sigma_n(B^{(k)})^2 = \min_{u \in K, \|u\|=1} \sum_{i=1,2} \|u \cdot b_i^{(k)}\|^2 \geq \sum_{i=1,2} \sigma_n(b_i^{(k)})^2$. Applying Lemma 2.3 with $t = 2\sqrt{2\pi e} n p_b^{-1}$ and $\delta = 1/3$, we get that $\sigma_n(b_i^{(k)}) > \frac{p_b}{8\sqrt{\pi e n}} \cdot \sigma'$, except with probability $\leq p_b$ for $i \in \{1, 2\}$ if $\sigma' \geq \frac{t}{\sqrt{2\pi}} \eta_{1/3}(\mathcal{I})$, where $\eta_{1/3}(\mathcal{I}) \leq \sqrt{\ln(8n)/\pi} \cdot \|g\|$ by Lemma 2.1. Therefore, we can choose

$$\ell_b = \frac{p_b}{2\sqrt{\pi e n}} \cdot \sigma' \quad \text{and} \quad \sigma' \geq 2n^{1.5}\sigma\sqrt{e \ln(8n)/\pi}/p_b. \tag{5}$$

Zero-testing and extraction correctness. The correctness conditions for zero-testing and correctness remain the same as conditions (2), (3) for the original GGH scheme. The only modification needed is for condition (1), because in $\mathsf{GGHLite}$, $m_r = 2$ and $\rho_j \in R$ so $\|\rho_j b_j^{(1)}\| \leq \sqrt{n}\|\rho_j\|\|b_j^{(1)}\|$. Accordingly, condition (1) is replaced by:

$$q > \max\left((n\ell_{g^{-1}})^8, (3 \cdot n^{1.5}\sigma^*\sigma')^{8\kappa}\right). \tag{6}$$

Security. We state our improved re-randomization security reduction for $\mathsf{GGHLite}$, that works with much smaller parameters than GGH. To our knowledge, it is the first security proof in which the RD is used to replace the SD in a sequence of games, using the RD properties from Section 4 to combine the bounds on changes between games. This allows us to gain the benefits of RD over SD, for both the drowning and smoothing aspects. Namely, with $\varepsilon_d, \varepsilon_\rho, \varepsilon_e$ in Theorem 6.2 set as large as $O(\log \lambda/\kappa)$, our weak security requirement of Definition 6.1 is satisfied (the RD between real and canonical encoding distributions is bounded by the quantity $R = \mathrm{poly}(\lambda)$ in Theorem 6.2), and our re-randomization goal for Ext-GCDH is achieved (whereas the strong requirement of Definition 3.2 is not satisfied).

Theorem 6.2 (Security of GGHLite). *Let $\varepsilon_d, \varepsilon_\rho, \varepsilon_e \in (0, 1/2)$ and $\kappa \leq 2^n$. Suppose that the following conditions are satisfied for* GGHLite*:*

– ***LHL Smoothing:***

$$\sigma_1^* \geq n^{1.5} \cdot \ell_{g^{-1}} \cdot \sigma \cdot \sqrt{2\log(4n \cdot \varepsilon_\rho^{-1})/\pi}. \tag{7}$$

– ***Offset "Drowning:"***

$$\sigma_1^* \geq n^{1.5} \cdot (\sigma')^2 \cdot \sqrt{2\pi\varepsilon_d^{-1}}/\ell_b. \tag{8}$$

– samp ***Uniformity Smoothing:***

$$\sigma' \geq \sigma \cdot \sqrt{n \ln(4n \cdot \varepsilon_e^{-1})/\pi}. \tag{9}$$

Then, if A is an adversary against the (non-canonical) Ext-GCDH problem for GGHLite *with run-time T and advantage ε, then A is also an adversary against the canonical problem Ext-cGCDH for* GGHLite *with $T' = T$ and advantage*

$$\varepsilon' \geq (\varepsilon - 2^{-\Omega(n)})^2/R \quad \textit{with} \quad R = 2^{O(\kappa\cdot(\varepsilon_d+\varepsilon_\rho+\varepsilon_e+2^{-n}))}. \tag{10}$$

In particular, there exist ε_d, ε_e, ε_ρ bounded as $O(\log \lambda/\kappa)$ such that the re-randomization security goal in Definition 3.4 is satisfied by GGHLite *with respect to problem Ext-GCDH.*

7 Parameter settings

In Table 1, we summarize asymptotic parameters for GGHLite to achieve 2^λ security for the underlying Ext-GCDH problem, assuming the hardness of the canonical Ext-cGCDH problem, and to satisfy the zero-testing/extraction correctness conditions with error probability $\lambda^{-\omega(1)}$. For simplicity, we assume that $\kappa = \omega(1)$. For comparison, we also show the corresponding parameters for GGH. The “Condition” column lists the conditions that determine the corresponding parameter in the case of GGHLite. For security of the canonical Ext-cGCDH problem, we assume (as in [9]) that the best attack is the one described in [9, Se. 6.3.3], whose complexity is dominated by the cost of solving γ-SVP (the Shortest lattice Vector Problem with approximation factor γ) for the lattice $\mathcal{I}$, with γ set at $\approx q^{3/8}$ to get a sufficiently short multiple of g. By the lattice reduction “rule of thumb,” to make this cost 2^{λ}, we need to set

$$n = \Omega(\lambda \log q). \qquad (11)$$

Table 1. Asymptotic parameters.

Parameter	GGHLite	GGH[9]	Condition
m_r	2	$\Omega(n \log n)$	LHL: Th. 5.1
σ	$O(n \log n)$	$O(n \log n)$	Eq. (4)
$\ell_{g^{-1}}$	$O(1/\sqrt{n \log n})$	$O(1/\sqrt{n \log n})$	Eq. (4)
$\varepsilon_d, \varepsilon_e, \varepsilon_\rho$	$O(\kappa^{-1})$	$O(2^{-\lambda}\kappa^{-1})$	Eq. (10)
σ'	$\widetilde{O}(n^{2.5})$	$\widetilde{O}(n^{1.5}\sqrt{\lambda})$	Eq. (5)
σ_1^*	$\widetilde{O}(n^{4.5}\sqrt{\log \kappa})$	$\widetilde{O}(2^{\lambda} n^{4.5}(\lambda + \log \kappa))$	Drown: Eq. (8)
ε_{ext}	$O(\lambda^{-\omega(1)})$	$O(\lambda^{-\omega(1)})$	
q	$\widetilde{O}((n^{8.5}\sqrt{\log \kappa})^{8\kappa})$	$\widetilde{O}((2^{\lambda} n^{8} \lambda^{1.5})^{8\kappa})$	Corr.: Eq. (6)
n	$O(\kappa\lambda \log \lambda)$	$O(\kappa\lambda^2)$	SVP: Eq. (11)
$\lvert\mathsf{enc}\rvert$	$O(\kappa^2 \lambda \log^2(\kappa\lambda))$	$O(\kappa^2\lambda^3)$	$O(n \log q)$
$\lvert\mathrm{par}\rvert$	$O(\kappa^3 \lambda \log^2(\kappa\lambda))$	$O(\kappa^3 \lambda^5 \log(\kappa\lambda))$	$O(m_r \kappa n \log q)$

When $\kappa = \mathrm{poly}(\log \lambda)$, the dimension n, encoding length $\lvert\mathsf{enc}\rvert$ and public parameters length $\lvert\mathrm{par}\rvert$ in our scheme GGHLite are all asymptotically close to optimal, namely quasi-linear in the security parameter λ, versus quadratic (resp. cubic and quintic) in λ for GGH [9]. Thus we expect GGHLite's public parameters and encodings to be orders of magnitudes shorter than GGH for typical $\lambda \approx 100$.

Acknowledgments. We thank Vadim Lyubashevsky for useful discussions. This work has been supported in part by ERC Starting Grant ERC-2013-StG-335086-LATTAC, an Australian Research Fellowship (ARF) from the Australian Research Council (ARC), and ARC Discovery Grants DP0987734 and DP110100628.

References

1. S. Agrawal, G. Gentry, S. Halevi, and A. Sahai. Discrete gaussian leftover hash lemma over infinite domains. In *Proc. of ASIACRYPT*, volume 8269 of *LNCS*,

pages 97–116. Springer, 2013.

2. J. Alperin-Sheriff and C. Peikert. Circular and KDM security for identity-based encryption. In *Proc. of PKC*, volume 7293 of *LNCS*, pages 334–352. Springer, 2012.
3. G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In *Proc. of EUROCRYPT*, pages 483–501, 2012.
4. A. Banerjee, C. Peikert, and A. Rosen. Pseudorandom functions and lattices. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 719–737. Springer, 2012.
5. D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615, 2003.
6. D. Boneh and A. Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324:71–90, 2003.
7. J.-S. Coron, T. Lepoint, and M. Tibouchi. Practical multilinear maps over the integers. In *Proc. of CRYPTO*, pages 476–493, 2013.
8. T. van Erven and P. Harremoës. Rényi divergence and Kullback-Leibler divergence. *CoRR*, abs/1206.2459, 2012.
9. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *Proc. of EUROCRYPT*, volume 7881 of *LNCS*, pages 1–17. Springer, 2013.
10. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. of STOC*, pages 169–178. ACM, 2009.
11. A. Joux. A one round protocol for tripartite Diffie-Hellman. In *Proc. of ANTS*, volume 1838 of *LNCS*, pages 385–394. Springer, 2000.
12. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
13. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.
14. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective.* Kluwer Academic Press, 2002.
15. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
16. C. Papamanthou, R. Tamassia, and N. Triandopoulos. Optimal authenticated data structures with multilinear forms. In *Proc. of Pairing*, pages 246–264, 2010.
17. O. Regev. Lecture notes of *lattices in computer science*, taught at the Computer Science Tel Aviv University. Available at `http://www.cims.nyu.edu/~regev/`.
18. A. Rényi. On measures of entropy and information. In *Proc. of the Fourth Berkeley Symposium on Math. Statistics and Probability*, volume 1, pages 547–561, 1961.
19. R. Rothblum. On the circular security of bit-encryption. In *Proc. of TCC*, pages 579–598, 2013.
20. M. Rückert and D. Schröder. Aggregate and verifiably encrypted signatures from multilinear maps without random oracles. In *Proc. of ISA*, pages 750–759, 2009.
21. R. Sakai, K. Ohgishi, and M. Kasahara. Cryptosystems based on pairing. *SCIS*, 2000.
22. D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal. lattices. In *Proc. of EUROCRYPT*, volume 6632 of *LNCS*, pages 27–47. Springer, 2011. Conference version of [23].
23. D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure standard worst-case problems over ideal lattices, 2013. Full version of [22], available at `http://perso.ens-lyon.fr/damien.stehle/NTRU.html`.