

Practical Multilinear Maps over the Integers

Jean-Sébastien Coron¹, Tancrede Lepoint^{2,3}, and Mehdi Tibouchi⁴

¹ University of Luxembourg

`jean-sebastien.coron@uni.lu`

² CryptoExperts, France

³ École Normale Supérieure, France

`tancrede.lepoint@cryptoexperts.com`

⁴ NTT Secure Platform Laboratories, Japan

`tibouchi.mehdi@lab.ntt.co.jp`

Abstract. Extending bilinear elliptic curve pairings to multilinear maps is a long-standing open problem. The first plausible construction of such multilinear maps has recently been described by Garg, Gentry and Halevi, based on ideal lattices. In this paper we describe a different construction that works over the integers instead of ideal lattices, similar to the DGHV fully homomorphic encryption scheme. We also describe a different technique for proving the full randomization of encodings: instead of Gaussian linear sums, we apply the classical leftover hash lemma over a quotient lattice. We show that our construction is relatively practical: for reasonable security parameters a one-round 7-party Diffie-Hellman key exchange requires about 25 seconds per party.

1 Introduction

Multilinear maps. Extending bilinear elliptic curve pairings to multilinear maps is a long-standing open problem. In 2003 Boneh and Silverberg showed two interesting applications of multilinear maps [BS03], namely multipartite Diffie-Hellman and very efficient broadcast encryption; however they were pessimistic about the existence of such maps.

The first plausible construction of multilinear maps has recently been described by Garg, Gentry and Halevi, based on ideal lattices [GGH13]. The main difference with bilinear pairings is that the encoding $a \cdot g$ of an element a is randomized (with some noise) instead of deterministic; only the final multilinear map $e(a_1 \cdot g, \dots, a_\kappa \cdot g)$ is a deterministic function of the a_i 's only. The construction has bounded degree with a maximum degree κ at most polynomial in the security parameter. Indeed, the encoding noise grows linearly with the degree, and when the noise reaches a certain threshold the encoding can become incorrect, as for ciphertexts in a somewhat homomorphic encryption scheme. The security of the construction relies on new hardness assumptions which are natural extensions of the Decisional Diffie-Hellman (DDH) assumption. To gain more confidence in their scheme the authors provide an extensive cryptanalytic survey. The authors focus on two applications: multipartite Diffie-Hellman and non-interactive zero-knowledge proof systems.

The construction from [GGH13] works in the polynomial ring $R = \mathbb{Z}[X]/(X^n + 1)$, where n is large enough to ensure security. One generates a secret short ring element $\mathbf{g} \in R$, generating a principal ideal $\mathcal{I} = \langle \mathbf{g} \rangle \subset R$. One also generates an integer parameter q and another random secret $\mathbf{z} \in R/qR$. One encodes elements of the quotient ring R/\mathcal{I} , namely elements of the form $\mathbf{e} + \mathcal{I}$ for some \mathbf{e} , as follows: a level- i encoding of the coset $\mathbf{e} + \mathcal{I}$ is an element of the form $u_k = [\mathbf{c}/\mathbf{z}^i]_q$, where $\mathbf{c} \in \mathbf{e} + \mathcal{I}$ is short. Such encodings can be both added and multiplied, as long as the norm of the numerators remain shorter than q ; in particular the product of κ encodings at level 1 gives an encoding at level κ . For such level- κ encodings one can then define a zero-testing parameter $\mathbf{p}_{zt} = [\mathbf{h}\mathbf{z}^\kappa/\mathbf{g}]_q$, for some small $\mathbf{h} \in R$. Then given a level- κ encoding $\mathbf{u} = [\mathbf{c}/\mathbf{z}^\kappa]$ one can compute $[\mathbf{p}_{zt} \cdot \mathbf{u}]_q = [\mathbf{h}\mathbf{c}/\mathbf{g}]_q$. When \mathbf{c} is an encoding of zero we have $\mathbf{c}/\mathbf{g} \in R$, which implies that $\mathbf{h}\mathbf{c}/\mathbf{g}$ is small in R , and therefore $[\mathbf{h}\mathbf{c}/\mathbf{g}]_q$ is small; this provides a way to test whether a level- κ encoding \mathbf{c} is an encoding of 0. For the same reason the high-order bits of $[\mathbf{p}_{zt} \cdot \mathbf{u}]_q = [\mathbf{h}\mathbf{c}/\mathbf{g}]_q$ only depend on the coset $\mathbf{e} + \mathcal{I}$ and not on the particular $\mathbf{c} \in \mathbf{e} + \mathcal{I}$; this makes it possible to extract a representation of cosets encoded at level κ , and eventually defines a degree- κ multilinear map for level-1 encodings.

Our contributions. Our main contribution is to describe a different construction that works over the integers instead of ideal lattices, similar to the DGHV fully homomorphic encryption scheme [DGHV10] and its batch variant [CCK⁺13]. Our construction offers the same flexibility as the original from [GGH13]; in particular it can be modified to support the analogue of asymmetric maps and composite-order maps. We provide a comparison with the original scheme, to better highlight the similarities. As in [GGH13], the security of our construction relies on new assumptions; it cannot be derived from “classical” assumptions such as the Approximate-GCD assumption used in [DGHV10]. We describe various possible attacks against our scheme; this enables to derive parameters for which our scheme remains secure against these attacks.

Our new construction works as follows: one first generates n secret primes p_i and publish $x_0 = \prod_{i=1}^n p_i$; one also generates n small secret primes g_i , and a random secret integer z modulo x_0 . A level- k encoding of a vector $\mathbf{m} = (m_i) \in \mathbb{Z}^n$ is then an integer c such that for all $1 \leq i \leq n$:

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

for some small random integers r_i ; the integer c is therefore defined modulo x_0 by CRT. It is clear that such encodings can be both added and multiplied modulo x_0 , as long as the numerators remain smaller than the p_i ’s. In particular the product of κ encodings c_j at level 1 gives an encoding at level κ where the corresponding vectors \mathbf{m}_j are multiplied componentwise. For such level- κ encodings one defines a zero-testing parameter p_{zt} with:

$$p_{zt} = \sum_{i=1}^n h_i \cdot (z^\kappa \cdot g_i^{-1} \pmod{p_i}) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}$$

for some small integers h_i . Given a level- κ encoding c as in (1), one can compute $\omega = p_{zt} \cdot c \pmod{x_0}$, which gives:

$$\omega = \sum_{i=1}^n h_i \cdot (r_i + m_i \cdot (g_i^{-1} \pmod{p_i})) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}.$$

Then if $m_i = 0$ for all i , since the r_i ’s and h_i ’s are small, we obtain that ω is small compared to x_0 ; this enables to test whether c is an encoding of 0 or not. Moreover for non-zero encodings the leading bits of ω only depend on the m_i ’s and not on the noise r_i ; for level- κ encodings this enables to extract a function of the m_i ’s only, which eventually defines as in [GGH13] a degree- κ multilinear map.¹

Our second contribution is to describe a different technique for proving the full randomization of encodings. Namely as in [GGH13] the Leftover Hash Lemma (LHL) cannot be directly applied since the encodings live in some infinite ring instead of a finite group. The solution in [GGH13] consists in using linear sums with Gaussian coefficients; it is shown in [AGHS12] that the resulting sum has a Gaussian distribution (over some lattice). As noted by the authors, this can be seen as a “leftover hash lemma over lattices”. In this paper we describe a different technique that does not use Gaussian coefficients; instead it consists in working modulo some lattice $L \subset \mathbb{Z}^n$ and applying the leftover hash lemma over the quotient \mathbb{Z}^n/L , which is still a finite group. Such technique was already used to prove the security of the batch variant of the DGHV fully homomorphic encryption scheme [CCK⁺13, CLT13]. Here we provide a more formal description: we clearly state our “LHL over lattices” so that it can later be applied as a black-box (as the corresponding Theorem 3 in [AGHS12]). The quotient lattice technique can independently be applied to the original encoding scheme from [GGH13], while the Gaussian sum technique from [AGHS12] is also applicable to ours.

Our third contribution is to describe the first implementation of multilinear maps. It appears that the basic versions of both [GGH13] and our scheme are rather unpractical, because of the huge public parameter size required to randomize the encodings. Therefore we use a simple optimization

¹ Technically for p_{zt} we use a *vector* of integers instead of a single integer (see Section 3).

that consists in storing only a small subset of the public elements and combining them pairwise to generate the full public-key. Such optimization was originally described in [GH11] for reducing the size of the encryption of the secret-key bits in the implementation of Gentry’s FHE scheme [Gen09]. It was also used in [CMNT11] to reduce the public-key size of the DGHV scheme; however as opposed to [CMNT11] our randomization of encodings becomes heuristic only, whereas in [CMNT11] the semantic security was still guaranteed. Thanks to this optimization our construction becomes relatively practical: for reasonable security parameters a multipartite Diffie-Hellman computation with 7 users requires less than 26 seconds, with a public parameter size of roughly 2.5 GBytes.

2 Definition of Randomized Encodings and Multilinear Maps

In this section we recall the setting used in [GGH13] for their notion of randomized encodings and multilinear maps, which they call *graded encoding schemes*. There are essentially two main differences with classical bilinear pairings (and their generalization to cryptographic multilinear maps as considered in [BS03]):

1. In bilinear pairings (and more generally cryptographic multilinear maps) we have a map $e: G^\kappa \rightarrow G_T$ that is linear with respect to all its κ inputs:

$$e(a_1 \cdot g, \dots, a_\kappa \cdot g) = \left(\prod_{i=1}^{\kappa} a_i \right) \cdot e(g, \dots, g). \quad (2)$$

One can view $a \cdot g$ as an “encoding” of the integer $a \in \mathbb{Z}_p$ over the group G of order p generated by g . The main difference in our setting is that encodings are now randomized. This means that an element $a \in R$ (where R is a ring that plays the role of the exponent space \mathbb{Z}_p) has many possible encodings; only the final multilinear map $e(a_1 \cdot g, \dots, a_\kappa \cdot g)$ is a deterministic function of the a_i ’s only, and not on the randomness used to encode a_i into $a_i \cdot g$.

2. The second main difference is that to every encoding is now associated a *level*. At level 0 we have the “plaintext” ring elements $a \in R$, at level 1 we have the encoding $a \cdot g$, and by combining k such encodings at level 1 one obtains a level- k encoding where the underlying elements a_i are homomorphically multiplied in R . The difference with “classical” cryptographic multilinear maps is that we can now multiply any (bounded) subset of encodings, instead of strictly κ at a time as with (2). For encodings at level κ we have a special zero-testing parameter \mathbf{p}_{zt} that can extract a deterministic function of the underlying ring elements. This enables to define a degree- κ multilinear map for encodings at level 1.

We recall in Appendix A the formal definition of a κ -Graded Encoding System from [GGH13]. For simplicity we only consider the symmetric case throughout the paper; we refer to [GGH13] for a more general framework that can handle the asymmetric case. We also recall in Appendix B the definition of the procedures for manipulating encodings and in Appendix C the corresponding hardness assumptions.

3 Our new Encoding Scheme

System parameters. The main parameters are the security parameter λ and the required multilinearity level $\kappa \leq \text{poly}(\lambda)$. Based on λ and κ , we choose the vector dimension n , the bit-size η of the primes p_i , the bit-size α of the primes g_i , the maximum bit-size ρ of the randomness used in encodings, and various other parameters that will be specified later; the constraints that these parameters must satisfy are described in the next section. For integers z, p we denote the reduction of z modulo p by $(z \bmod p)$ or $[z]_p$ with $-p/2 < [z]_p \leq p/2$.

In our scheme a level- k encoding of a vector $\mathbf{m} = (m_i) \in \mathbb{Z}^n$ is an integer c such that for all $1 \leq i \leq n$:

$$c \equiv \frac{r_i \cdot g_i + m_i}{z^k} \pmod{p_i} \quad (3)$$

where the r_i 's are ρ -bit random integers (specific to the encoding c), with the following secret parameters: the p_i 's are η -bit prime integers, the g_i 's are α -bit primes, and the denominator z is a random (invertible) integer modulo $x_0 = \prod_{i=1}^n p_i$. The integer c is therefore defined by CRT modulo x_0 , where x_0 is made public. Since the p_i 's must remain secret, the user cannot encode the vectors $\mathbf{m} \in \mathbb{Z}^n$ by CRT directly from (3); instead one includes in the public parameters a set of ℓ level-0 encodings x'_j of random vectors $\mathbf{a}_j \in \mathbb{Z}^n$, and the user can generate a random level-0 encoding by computing a random subset-sum of those x'_j 's.

Remark 1. From (3) each integer m_i is actually defined modulo g_i . Therefore our scheme encodes vectors \mathbf{m} from the ring $QR = \mathbb{Z}_{g_1} \times \dots \times \mathbb{Z}_{g_n}$.

Instance generation: $(\text{params}, \mathbf{p}_{zt}) \leftarrow \text{InstGen}(1^\lambda, 1^\kappa)$. We generate n secret random η -bit primes p_i and publish $x_0 = \prod_{i=1}^n p_i$. We generate a random invertible integer z modulo x_0 . We generate n random α -bit prime integers g_i and a secret matrix $\mathbf{A} = (a_{ij}) \in \mathbb{Z}^{n \times \ell}$, where each component a_{ij} is randomly generated in $[0, g_i) \cap \mathbb{Z}$. We generate an integer y , three sets of integers $\{x_j\}_{j=1}^\tau$, $\{x'_j\}_{j=1}^\ell$ and $\{\Pi_j\}_{j=1}^n$, a zero-testing vector \mathbf{p}_{zt} , and a seed s for a strong randomness extractor, as described later. We publish the parameters $\text{params} = (n, \eta, \alpha, \rho, \beta, \tau, \ell, y, \{x_j\}_{j=1}^\tau, \{x'_j\}_{j=1}^\ell, \{\Pi_j\}_{j=1}^n, s)$ and \mathbf{p}_{zt} .

Sampling level-zero encodings: $c \leftarrow \text{samp}(\text{params})$. We publish as part as our instance generation a set of ℓ integers x'_j , where each x'_j encodes at level-0 the column vector $\mathbf{a}_j \in \mathbb{Z}^n$ of the secret matrix $\mathbf{A} = (a_{ij}) \in \mathbb{Z}^{n \times \ell}$. More precisely using the CRT modulo x_0 we generate integers x'_j 's such that:

$$1 \leq j \leq \ell, \quad x'_j \equiv r'_{ij} \cdot g_i + a_{ij} \pmod{p_i} \quad (4)$$

where the r'_{ij} 's are randomly generated in $(-2^\rho, 2^\rho) \cap \mathbb{Z}$.

Our randomized sampling algorithm $\text{samp}(\text{params})$ works as follows: we generate a random binary vector $\mathbf{b} = (b_j) \in \{0, 1\}^\ell$ and output the level-0 encoding

$$c = \sum_{j=1}^{\ell} b_j \cdot x'_j \pmod{x_0}.$$

From Equation (4) this gives $c \equiv (\sum_{j=1}^{\ell} r'_{ij} b_j) \cdot g_i + \sum_{j=1}^{\ell} a_{ij} b_j \pmod{p_i}$; as required the output c is a level-0 encoding:

$$c \equiv r_i \cdot g_i + m_i \pmod{p_i} \quad (5)$$

of some vector $\mathbf{m} = \mathbf{A} \cdot \mathbf{b} \in \mathbb{Z}^n$ which is a random subset-sum of the column vectors \mathbf{a}_j . We note that for such level-0 encodings we get $|r_i \cdot g_i + m_i| \leq \ell \cdot 2^{\rho+\alpha}$ for all i .

The following Lemma states that, as required, the distribution of \mathbf{m} is statistically close to uniform over $QR = \mathbb{Z}_{g_1} \times \dots \times \mathbb{Z}_{g_n}$; the proof is based on applying the LHL over the set QR (see Appendix D).

Lemma 1. *Let $c \leftarrow \text{samp}(\text{params})$ and write $c \equiv r_i \cdot g_i + m_i \pmod{p_i}$. Let $\ell \geq n \cdot \alpha + 2\lambda$. The distribution of $(\text{params}, \mathbf{m})$ is statistically close to the distribution of $(\text{params}, \mathbf{m}')$ where $\mathbf{m}' \leftarrow QR$.*

Encodings at higher levels: $c_k \leftarrow \text{enc}(\text{params}, k, c)$. To allow encoding at higher levels, we publish as part of our instance-generation a level-one random encoding of $\mathbf{1}$, namely an integer y such that:

$$y \equiv \frac{r_i \cdot g_i + 1}{z} \pmod{p_i}$$

for random integers $r_i \in (-2^\rho, 2^\rho) \cap \mathbb{Z}$; as previously the integer y is computed by CRT modulo x_0 .

Given a level-0 encoding c of $\mathbf{m} \in \mathbb{Z}^n$ as given by (5), we can then compute a level-1 encoding of the same \mathbf{m} by computing $c_1 = c \cdot y \pmod{x_0}$. Namely we obtain as required:

$$c_1 \equiv \frac{r_i^{(1)} \cdot g_i + m_i}{z} \pmod{p_i} \quad (6)$$

for some integers $r_i^{(1)}$, and we get $|r_i^{(1)} \cdot g_i + m_i| \leq \ell \cdot 2^{2(\rho+\alpha)}$ for all i . More generally to generate a level- k encoding we compute $c_k = c_0 \cdot y^k \pmod{x_0}$.

In multipartite Diffie-Hellman key-exchange every user keeps a private level-0 encoding c and publishes a level-1 encoding of the same underlying (unknown) \mathbf{m} ; however one cannot publish $c_1 = c \cdot y \pmod{x_0}$ directly since the private level-0 encoding c could be recovered directly from $c = c_1/y \pmod{x_0}$. Instead the level-1 encoding c_1 must first be re-randomized into a new level-1 encoding c'_1 whose distribution does not depend on the original c as long as it encodes the same \mathbf{m} .

Re-randomization: $c' \leftarrow \text{reRand}(\text{params}, i, c)$. To allow re-randomization of encodings at level one, we publish as part of our instance-generation a set of n integers Π_j which are all level-1 random encodings of zero:

$$1 \leq j \leq n, \quad \Pi_j \equiv \frac{\varpi_{ij} \cdot g_i}{z} \pmod{p_i}.$$

The matrix $\mathbf{\Pi} = (\varpi_{ij}) \in \mathbb{Z}^{n \times n}$ is a diagonally dominant matrix generated as follows: the non-diagonal entries are randomly and independently generated in $(-2^\rho, 2^\rho) \cap \mathbb{Z}$, while the diagonal entries are randomly generated in $(n2^\rho, n2^\rho + 2^\rho) \cap \mathbb{Z}$.

We also publish as part of our instance-generation a set of τ integers x_j , where each x_j is a level-1 random encoding of zero:

$$1 \leq j \leq \tau, \quad x_j \equiv \frac{r_{ij} \cdot g_i}{z} \pmod{p_i}$$

and where the column vectors of the matrix $(r_{ij}) \in \mathbb{Z}^{n \times \tau}$ are randomly and independently generated in the half-open parallelepiped spanned by the columns of the previous matrix $\mathbf{\Pi}$; see Appendix K for a concrete algorithm.

Given as input a level-1 encoding c_1 as given by (6), we randomize c_1 with a random subset-sum of the x_j 's and a linear combination of the Π_j 's:

$$c'_1 = c_1 + \sum_{j=1}^{\tau} b_j \cdot x_j + \sum_{j=1}^n b'_j \cdot \Pi_j \pmod{x_0} \quad (7)$$

where $b_j \leftarrow \{0, 1\}$, and $b'_j \leftarrow [0, 2^\mu) \cap \mathbb{Z}$. The following Lemma shows that as required the distribution of c'_1 is nearly independent of the input (as long as it encodes the same \mathbf{m}). This follows essentially from our ‘‘leftover hash lemma over lattices’’; see Appendix F.

Lemma 2. *Let $c \leftarrow \text{samp}(\text{params})$, $c_1 \leftarrow \text{enc}(\text{params}, 1, c)$, and $c'_1 \leftarrow \text{reRand}(\text{params}, 1, c_1)$. Write $c'_1 \equiv (r'_i \cdot g_i + m_i)/z \pmod{p_i}$ for all $1 \leq i \leq n$, and $\mathbf{r} = (r_1, \dots, r_n)^T$. Let $\tau \geq n \cdot (\rho + \log_2(2n)) + 2\lambda$ and $\mu \geq \rho + \alpha + \lambda$. The distribution of $(\text{params}, \mathbf{r})$ is statistically close to that of $(\text{params}, \mathbf{r}')$, where $\mathbf{r}' \in \mathbb{Z}^n$ is randomly generated in the half-open parallelepiped spanned by the column vectors of $2^\mu \mathbf{\Pi}$.*

Writing $c'_1 \equiv (r'_i \cdot g_i + m_i)/z \pmod{p_i}$, and using $|r_{ij} \cdot g_i| \leq 2n2^{\rho+\alpha}$ for all i, j , we obtain $|r'_i \cdot g_i + m_i| \leq \ell 2^{2(\rho+\alpha)} + 2n\tau 2^{\rho+\alpha} + 2n2^{\mu+\rho+\alpha}$. Using $\mu \geq \rho + \alpha + \lambda$ this gives $|r'_i \cdot g_i + m_i| \leq 4n2^{\mu+\rho+\alpha}$.

Adding and Multiplying Encodings. It is clear that one can homomorphically add encodings. Moreover the product of κ level-1 encodings u_i can be interpreted as an encoding of the product. Namely given level-one encodings u_j of vectors $\mathbf{m}_j \in \mathbb{Z}^n$ for $1 \leq j \leq \kappa$, with $u_j \equiv (r_{ij} \cdot g_i + m_{ij})/z \pmod{p_i}$, we simply let:

$$u = \prod_{j=1}^{\kappa} u_j \pmod{x_0}.$$

This gives:

$$u \equiv \frac{\prod_{j=1}^{\kappa} (r_{ij} \cdot g_i + m_{ij})}{z^{\kappa}} \equiv \frac{r_i \cdot g_i + \left(\prod_{j=1}^{\kappa} m_{ij} \right)}{z^{\kappa}} \pmod{p_i}$$

for some $r_i \in \mathbb{Z}$. This is a level- κ encoding of the vector \mathbf{m} obtained by componentwise product of the vectors \mathbf{m}_i , as long as $\prod_{j=1}^{\kappa} (r_{ij} \cdot g_i + m_{ij}) < p_i$ for all i . When computing the product of κ level-1 encodings from `reRand` and one level-0 encoding from `samp` (as in multipartite Diffie-Hellman key exchange), we obtain from previous bounds $|r_i| \leq (4n2^{\mu+\rho+\alpha})^{\kappa} \cdot \ell \cdot 2^{\rho+1}$ for all i .

Zero Testing. `isZero(params, $\mathbf{p}_{zt}, u_{\kappa}$)` $\stackrel{?}{=} 0/1$. We can test equality between encodings by subtracting them and testing for zero. To enable zero testing we randomly generate an integer matrix $\mathbf{H} = (h_{ij}) \in \mathbb{Z}^{n \times n}$ such that \mathbf{H} is invertible in \mathbb{Z} and both $\|\mathbf{H}^T\|_{\infty} \leq 2^{\beta}$ and $\|(\mathbf{H}^{-1})^T\|_{\infty} \leq 2^{\beta}$, for some parameter β specified later; here $\|\cdot\|_{\infty}$ is the operator norm on $n \times n$ matrices with respect to the ℓ^{∞} norm on \mathbb{R}^n . The generation of \mathbf{H} is discussed in Appendix L. We then publish as part of our instance generation the following zero-testing vector $\mathbf{p}_{zt} \in \mathbb{Z}^n$:

$$(\mathbf{p}_{zt})_j = \sum_{i=1}^n h_{ij} \cdot (z^{\kappa} \cdot g_i^{-1} \pmod{p_i}) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}. \quad (8)$$

To determine whether a level- κ encoding c is an encoding of zero or not, we compute the vector $\boldsymbol{\omega} = c \cdot \mathbf{p}_{zt} \pmod{x_0}$ and test whether $\|\boldsymbol{\omega}\|_{\infty}$ is small:

$$\text{isZero}(\text{params}, \mathbf{p}_{zt}, c) = \begin{cases} 1 & \text{if } \|c \cdot \mathbf{p}_{zt} \pmod{x_0}\|_{\infty} < x_0 \cdot 2^{-\nu} \\ 0 & \text{otherwise} \end{cases}$$

for some parameter ν specified later.

Namely for a level- κ ciphertext c we have $c \equiv (r_i \cdot g_i + m_i)/z^{\kappa} \pmod{p_i}$ for some $r_i \in \mathbb{Z}$; therefore for all $1 \leq i \leq n$ we can write:

$$c = q_i \cdot p_i + (r_i \cdot g_i + m_i) \cdot (z^{-\kappa} \pmod{p_i}) \quad (9)$$

for some $q_i \in \mathbb{Z}$. Therefore combining (8) and (9), we get:

$$(\boldsymbol{\omega})_j = (c \cdot \mathbf{p}_{zt} \pmod{x_0})_j = \sum_{i=1}^n h_{ij} \cdot (r_i + m_i \cdot (g_i^{-1} \pmod{p_i})) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}. \quad (10)$$

Therefore if $m_i = 0$ for all $1 \leq i \leq n$, then $\|\boldsymbol{\omega}\|_{\infty} = \|c \cdot \mathbf{p}_{zt} \pmod{x_0}\|_{\infty}$ is small compared to x_0 when the r_i 's are small enough, i.e. a limited number of additions/multiplications on encodings has been performed. Conversely if $m_i \neq 0$ for some i we show that $\|\boldsymbol{\omega}\|_{\infty}$ must be large. More precisely we prove the following lemma in Appendix G.

Lemma 3. *Let n, η, α and β be as in our parameter setting. Let ρ_f be such that $2\beta + \alpha + \rho_f \leq \eta - 8$, and let $\nu = \eta - \beta - \rho_f - 3 \geq \alpha + \beta + 5$. Let c be such that $c \equiv (r_i \cdot g_i + m_i)/z^{\kappa} \pmod{p_i}$ for all $1 \leq i \leq n$, where $0 \leq m_i < g_i$ for all i . Let $\mathbf{r} = (r_i)_{1 \leq i \leq n}$ and assume that $\|\mathbf{r}\|_{\infty} < 2^{\rho_f}$. If $\mathbf{m} = 0$ then $\|\boldsymbol{\omega}\|_{\infty} < x_0 \cdot 2^{-\nu-2}$. Conversely if $\mathbf{m} \neq 0$ then $\|\boldsymbol{\omega}\|_{\infty} \geq x_0 \cdot 2^{-\nu+2}$.*

Extraction. $sk \leftarrow \text{ext}(\text{params}, \mathbf{p}_{zt}, u_\kappa)$. This part is essentially the same as in [GGH13]. To extract a random function of the vector \mathbf{m} encoded in a level- κ encoding c , we multiply c by the zero-testing parameter \mathbf{p}_{zt} modulo x_0 , collect the ν most significant bits of each of the n components of the resulting vector, and apply a strong randomness extractor (using the seed s from params):

$$\text{ext}(\text{params}, \mathbf{p}_{zt}, c) = \text{Extract}_s(\text{msbs}_\nu(c \cdot \mathbf{p}_{zt} \bmod x_0))$$

where msbs_ν extracts the ν most significant bits of the result. Namely if two encodings c and c' encode the same $\mathbf{m} \in \mathbb{Z}^n$ then from Lemma 3 we have $\|(c - c') \cdot \mathbf{p}_{zt} \bmod x_0\|_\infty < x_0 \cdot 2^{-\nu-2}$, and therefore we expect that $\boldsymbol{\omega} = c \cdot \mathbf{p}_{zt} \bmod x_0$ and $\boldsymbol{\omega}' = c' \cdot \mathbf{p}_{zt} \bmod x_0$ agree on their ν most significant bits, and therefore extract to the same value.²

Conversely if c and c' encode different vectors then by Lemma 3 we must have $\|(c - c') \cdot \mathbf{p}_{zt} \bmod x_0\|_\infty > x_0 \cdot 2^{-\nu+2}$, and therefore the ν most significant bits of the corresponding $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ must be different. This implies that for random $\mathbf{m} \in QR = \mathbb{Z}_{g_1} \times \cdots \times \mathbb{Z}_{g_n}$ the min-entropy of $\text{msbs}_\nu(c \cdot \mathbf{p}_{zt} \bmod x_0)$ when c encodes \mathbf{m} is at least $\log_2 |QR| \geq n(\alpha - 1)$. Therefore we can use a strong randomness extractor to extract a nearly uniform bit-string of length $\lfloor \log_2 |QR| \rfloor - \lambda$.

This concludes the description of our multilinear encoding scheme. In Appendix I we provide a comparison with the original scheme from [GGH13].

3.1 Setting the Parameters

The system parameters must satisfy the following constraints:

- The bit-size ρ of the randomness used for encodings must satisfy $\rho = \Omega(\lambda)$ to avoid brute force attack on the noise, including the improved attack from [CN12] with complexity $\tilde{O}(2^{\rho/2})$.
- The bit-size α of the primes g_i must be large enough so that the order of the group $QR = \mathbb{Z}_{g_1} \times \cdots \times \mathbb{Z}_{g_n}$ does not contain small prime factors; this is required to prove the security of the multipartite Diffie-Hellman Key Exchange (see Appendix J). One can take $\alpha = \lambda$.
- The parameter n must be large enough to thwart lattice-based attacks on the encodings, namely $n = \omega(\eta \log \lambda)$; see Section 5.1.
- The number ℓ of level-0 encodings x'_j for `samp` must satisfy $\ell \geq n \cdot \alpha + 2\lambda$ in order to apply the leftover hash lemma; see Appendix D.
- The number τ of level-1 encodings x_j must satisfy $\tau \geq n \cdot (\rho + \log_2(2n)) + 2\lambda$ in order to apply our “leftover hash lemma over lattices”. For the same reason the bit-size μ of the linear sum coefficients must satisfy $\mu \geq \alpha + \rho + \lambda$; see Lemma 2.
- The bit-size η of the primes p_i must satisfy $\eta \geq 2\beta + \alpha + \rho_f + 8$, where ρ_f is the maximum bit size of the randoms r_i a level- κ encoding (see Lemma 3). When computing the product of κ level-1 encodings and an additional level-0 encoding (as in a multipartite Diffie-Hellman key exchange with $N = \kappa + 1$ users), one obtains $\rho_f = \kappa \cdot (\mu + \rho + \alpha + \log_2 n + 2) + \rho + \log_2 \ell + 1$ (see previous Section).
- The number ν of most significant bits to extract can then be set to $\nu = \eta - \beta - \rho_f - 3$ (see Lemma 3).

3.2 Security of our Construction

As in [GGH13] the security of our construction relies on new assumptions that do not seem to be reducible to more classical assumptions. Namely as in [GGH13] one can make the assumption that solving the Graded DDH problem (GDDH) recalled in Appendix C is hard in our scheme. This enables to prove the security of the one-round N -way Diffie-Hellman key exchange protocol [GGH13]. Ideally one would like to reduce such assumption to a more classical assumption, such as the Approximate-GCD assumption, but that does not seem feasible. Therefore to gain more confidence in our scheme we describe various attacks in Section 5.

² As in [GGH13] it can be that all components $\boldsymbol{\omega}$ and $\boldsymbol{\omega}'$ differ by an additive factor x_0 ; however this happens with negligible probability in λ .

4 Another Leftover Hash Lemma over Lattices

As mentioned in the introduction, to prove the full randomization of encodings (Lemma 2) one cannot apply the classical Leftover Hash Lemma (LHL) because the noise in the encodings live in some infinite ring instead of a finite group. In [GGH13] the issue was solved by using linear sums with Gaussian coefficients. Namely the analysis in [AGHS12] shows that the resulting sum has a Gaussian distribution (over some lattice). As noted by the authors this technique can be seen as a “leftover hash lemma over lattices”. Such technique would be applicable to our scheme as well.

In this section we describe an alternative technique (without Gaussian coefficients) that can also be seen as a “leftover hash lemma over lattices”. It consists in working modulo a lattice $L \subset \mathbb{Z}^n$ and applying the classical leftover hash lemma over the finite group \mathbb{Z}^n/L . This technique was already used in [CLT13, CCK⁺13] to prove the security of a batch extension of the DGHV scheme. In this paper we provide a more formal description; namely we clearly state our “LHL over lattices” so that it can later be applied as a black-box (as the corresponding Theorem 3 in [AGHS12]). Our quotient lattice technique can independently be applied to the original encoding scheme from [GGH13].

4.1 Classical Leftover Hash Lemma

We first recall the classical Leftover Hash Lemma. We say that the distributions D_1, D_2 over a finite domain X are ε -statistically close if the statistical distance $\Delta(D_1, D_2) = \frac{1}{2} \sum_{x \in X} |D_1(x) - D_2(x)|$ is smaller than ε . A distribution D is ε -uniform if its statistical distance from the uniform distribution is at most ε . A family \mathcal{H} of hash functions from X to Y , both finite sets, is said to be pairwise-independent if for all distinct $x, x' \in X$, $\Pr_{h \leftarrow \mathcal{H}}[h(x) = h(x')] = 1/|Y|$.

Lemma 4 (Leftover Hash Lemma [HILL99]). *Let \mathcal{H} be a family of pairwise hash functions from X to Y . Suppose that $h \leftarrow \mathcal{H}$ and $x \leftarrow X$ are chosen uniformly and independently. Then, $(h, h(x))$ is $\frac{1}{2} \sqrt{|Y|/|X|}$ -uniform over $\mathcal{H} \times Y$.*

One can then deduce the following Lemma for random subset sums over a finite abelian group.

Lemma 5. *Let G be a finite abelian group. Set $x_1, \dots, x_m \leftarrow G$ uniformly and independently, set $s_1, \dots, s_m \leftarrow \{0, 1\}$, and set $y = \sum_{i=1}^m s_i x_i \in G$. Then (x_1, \dots, x_m, y) is $1/2 \sqrt{|G|/2^m}$ -uniform over G^{m+1} .*

Proof. We consider the following hash function family \mathcal{H} from $\{0, 1\}^m$ to G . Each member $h \in \mathcal{H}$ is parameterized by the elements $(x_1, \dots, x_m) \in G^m$. Given $\mathbf{s} \in \{0, 1\}^m$, we define $h(\mathbf{s}) = \sum_{i=1}^m s_i \cdot x_i \in G$. The hash function family is clearly pairwise independent. Therefore by Lemma 4, $(h, h(x))$ is $\frac{1}{2} \sqrt{|G|/2^m}$ -uniform over G^{m+1} . \square

4.2 Leftover Hash Lemma over Lattices

Let $L \subset \mathbb{Z}^n$ be a lattice of rank n of basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$. Then every $\mathbf{x} \in \mathbb{Z}^n$ can be uniquely written as:

$$\mathbf{x} = \xi_1 \mathbf{b}_1 + \dots + \xi_n \mathbf{b}_n$$

for some real numbers ξ_i . Moreover, for every vector $\mathbf{x} \in \mathbb{Z}^n$ there is a unique $\mathbf{a} \in L$ such that:

$$\mathbf{y} = \mathbf{x} - \mathbf{a} = \xi'_1 \mathbf{b}_1 + \dots + \xi'_n \mathbf{b}_n$$

where $0 \leq \xi'_i < 1$; we write $\mathbf{y} = \mathbf{x} \bmod \mathbf{B}$. Therefore each vector of \mathbb{Z}^n/L has a unique representative in the half-open parallelepiped defined by the previous equation.

We denote by $D_{\mathbf{B}}$ the distribution obtained by generating a random element in the quotient \mathbb{Z}^n/L and taking its unique representative in the half-open parallelepiped generated by the basis \mathbf{B} . Given a basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$ and $\mu \in \mathbb{Z}^*$ we denote by $\mu \mathbf{B}$ the basis $(\mu \mathbf{b}_1, \dots, \mu \mathbf{b}_n)$. We are now ready to state our “Leftover Hash Lemma over Lattices”, which we prove in Appendix E.

Lemma 6. Let $L \subset \mathbb{Z}^n$ be a lattice of rank n of basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$. Let \mathbf{x}_i for $1 \leq i \leq m$ be generated independently according to the distribution $\mathcal{D}_{\mathbf{B}}$. Set $s_1, \dots, s_m \leftarrow \{0, 1\}$ and $t_1, \dots, t_n \leftarrow \mathbb{Z} \cap [0, 2^\mu)$. Let $\mathbf{y} = \sum_{i=1}^m s_i \mathbf{x}_i + \sum_{i=1}^n t_i \mathbf{b}_i$ and $\mathbf{y}' \leftarrow \mathcal{D}_{2^\mu \mathbf{B}}$. Then the distributions $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$ and $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}')$ are ε -statistically close, with $\varepsilon = mn \cdot 2^{-\mu} + 1/2\sqrt{|\det L|/2^m}$.

We also show that the previous distribution $\mathcal{D}_{2^\mu \mathbf{B}}$ is not significantly modified when a small vector $\mathbf{z} \in \mathbb{Z}^n$ is added and the operator norm of the corresponding matrix B^{-1} is upper-bounded; see Appendix H.

Lemma 7. Let $L \subset \mathbb{Z}^n$ be a full-rank lattice of basis $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$, and let $B \in \mathbb{Z}^{n \times n}$ be the matrix whose column vectors are the \mathbf{b}_i 's. For any $\mathbf{z} \in \mathbb{Z}^n$, the distribution of $\mathbf{z} + \mathcal{D}_{2^\mu \mathbf{B}}$ is ε -statistically close to that of $\mathcal{D}_{2^\mu \mathbf{B}}$, where $\varepsilon = 2^{-\mu} \cdot (\|\mathbf{z}\|_\infty \cdot \|B^{-1}\|_\infty + 1)$.

4.3 Re-Randomization of Encodings

We are now ready to apply our ‘‘LHL over lattices’’ to prove the full randomization of encodings as stated in Lemma 2. Namely the re-randomization equation (7) can be rewritten in vector form as:

$$\mathbf{r}' = \mathbf{r} + \mathbf{X} \cdot \mathbf{b} + \mathbf{\Pi} \cdot \mathbf{b}'$$

where $\mathbf{b} \leftarrow \{0, 1\}^\tau$ and $\mathbf{b}' \leftarrow ([0, 2^\mu) \cap \mathbb{Z})^n$, and the columns of the matrix $\mathbf{X} \in \mathbb{Z}^{n \times \tau}$ are uniformly and independently generated in the parallelepiped spanned by the columns of the matrix $\mathbf{\Pi} \in \mathbb{Z}^{n \times n}$. Therefore to conclude it suffices to apply Lemma 6 and Lemma 7, using additionally an upper bound on $\|\mathbf{\Pi}^{-1}\|_\infty$. For this we use the fact that $\mathbf{\Pi}$ has been generated as a diagonally dominant matrix.

Given a matrix $B = (b_{ij}) \in \mathbb{R}^{n \times n}$, we let $A_i(B) = \sum_{k \neq i} |b_{ik}|$. A matrix $B = (b_{ij}) \in \mathbb{R}^{n \times n}$ is said to be *diagonally dominant* if $|b_{ii}| > A_i(B)$ for all i . We recall the following fact for diagonally dominant matrices [Var75], and refer to Appendix F for a full proof of Lemma 2.

Fact 1 Let $B = (b_{ij}) \in \mathbb{R}^{n \times n}$ be a diagonally dominant matrix. Then the matrix B is invertible and $\|B^{-1}\|_\infty \leq \max_{i=1, \dots, n} (|b_{ii}| - A_i(B))^{-1}$ where $\|\cdot\|_\infty$ is the operator norm on $n \times n$ matrices with respect to the ℓ^∞ norm on \mathbb{R}^n .

5 Attacks against our Multilinear Scheme

5.1 Lattice Attack on the Encodings

We first describe a lattice attack against level-0 encodings. We consider an element $x_0 = \prod_{i=1}^n p_i$ and a set of τ integers $x_j \in \mathbb{Z}_{x_0}$ such that:

$$x_j \bmod p_i = r_{ij}$$

where $r_{ij} \in (-2^\rho, 2^\rho) \cap \mathbb{Z}$. We want to estimate the complexity of the classical orthogonal lattice attack for recovering (some of) the noise values r_{ij} .

This attack works by considering the integer vector formed by a subset of the x_j 's, say $\mathbf{x} = (x_j)_{1 \leq j \leq t}$ for some $n < t \leq \tau$, and relating the lattice of vectors orthogonal to $\mathbf{x} \bmod x_0$ to the lattice of vectors orthogonal to each of the corresponding noise value vectors $\mathbf{r}_i = (r_{ij})_{1 \leq j \leq t}$.

More precisely, let $L \subset \mathbb{Z}^t$ the lattice of vectors \mathbf{u} such that:

$$\mathbf{u} \cdot \mathbf{x} \equiv 0 \pmod{x_0}.$$

Clearly, L contains $x_0 \mathbb{Z}^t$ so it is of full rank t . Moreover, we have

$$\text{vol}(L) = [\mathbb{Z}^t : L] = x_0 / \gcd(x_0, x_1, \dots, x_t) = x_0.$$

As a result, we heuristically expect the successive minima of L to be around $\text{vol}(L)^{1/t} \approx 2^{n\cdot\eta/t}$, and hence applying lattice reduction should yield a reduced basis $(\mathbf{u}_1, \dots, \mathbf{u}_t)$ with vectors of length $\|\mathbf{u}_k\| \approx 2^{n\cdot\eta/t+\alpha t}$ for some constant $\alpha > 0$ depending on the lattice reduction algorithm ($2^{\alpha t}$ is the Hermite factor).

Now, a vector $\mathbf{u} \in L$ satisfies $\mathbf{u} \cdot \mathbf{x} \equiv 0 \pmod{x_0}$, so for each $i \in \{1, \dots, n\}$, reducing modulo p_i gives:

$$\mathbf{u} \cdot \mathbf{r}_i \equiv 0 \pmod{p_i}.$$

In particular, if \mathbf{u} is short enough to satisfy $\|\mathbf{u}\| \cdot \|\mathbf{r}_i\| < p_i$, this implies $\mathbf{u} \cdot \mathbf{r}_i = 0$ in \mathbb{Z} . As a result, if we have:

$$2^{n\cdot\eta/t+\alpha t} \cdot 2^\rho < 2^\eta, \tag{11}$$

we expect the vectors $(\mathbf{u}_1, \dots, \mathbf{u}_{t-n})$ from the previous lattice reduction step to be orthogonal to the \mathbf{r}_i 's, and hence computing the rank n orthogonal lattice to the lattice spanned by those vectors should reveal the \mathbf{r}_i 's.

Since t must be greater than n for the attack to apply, condition (11) implies in particular that:

$$\alpha < \frac{\eta - \rho}{n}.$$

Since a Hermite factor of $2^{\alpha t}$ is achieved in time $2^{\Omega(1/\alpha)}$ (usually by carrying out BKZ reduction with block size $\beta = \omega(1/\alpha)$, in which each block is BKZ-reduced in time exponential in β , see e.g. [HPS11]), we obtain that this orthogonal lattice attack has a complexity exponential in n . In fact, with $\gamma = \eta \cdot n$, we get a time complexity of $2^{\Omega(\gamma/\eta^2)}$, similar to [DGHV10, §5.2].

5.2 Hidden Subset Sum Attack on Zero Testing

One can consider a similar attack on the zero-testing parameters. The zero-testing vector $\boldsymbol{\omega} = c \cdot \mathbf{p}_{zt} \pmod{x_0}$ corresponding to an encoding c of zero can be written as:

$$\boldsymbol{\omega} = \sum_{i=1}^n \left(r_i \cdot \prod_{i' \neq i} p_{i'} \right) \mathbf{h}_i = \sum_{i=1}^n R_i \mathbf{h}_i,$$

where the \mathbf{h}_i 's are the lines of the zero-testing matrix \mathbf{H} . This gives an expression of $\boldsymbol{\omega}$ as a linear combination of the unknown vectors \mathbf{h}_i , so we can think of an approach similar to the hidden subset sum attack of Nguyen and Stern [NS99] to recover the unknown \mathbf{h}_i 's or the R_i 's.

Such an approach, like the one from the previous section, would first involve computing the lattice $L \subset \mathbb{Z}^n$ of vectors \mathbf{u} orthogonal to $\boldsymbol{\omega}$ modulo x_0 , and hoping that the vectors in a reduced basis of L are short enough that they must necessarily be orthogonal to each of the \mathbf{h}_i 's in \mathbb{Z}^n .

However, L is a lattice of rank n and volume $x_0 \approx 2^{n\cdot\eta}$, so we expect its shortest vectors to be of length roughly 2^η . The attack can then only work if such a short vector \mathbf{u} is necessarily orthogonal to each of the \mathbf{h}_i 's in \mathbb{Z}^n . Equivalently, this will happen if the vector $\mathbf{v} = (\mathbf{u} \cdot \mathbf{h}_1, \dots, \mathbf{u} \cdot \mathbf{h}_n) \in \mathbb{Z}^n$, which we know is orthogonal to (R_1, \dots, R_n) modulo x_0 , is significantly shorter than the shortest vector in the lattice L' of vectors orthogonal to (R_1, \dots, R_n) modulo x_0 . But again L' is of rank n and volume x_0 , so its shortest vectors is of length about 2^η , and hence \mathbf{v} will typically not be shorter.

Therefore, the Nguyen–Stern hidden subset sum attack does not apply to our setting.

5.3 Attacks on the Inverse Zero Testing Matrix

A related observation is that if $\boldsymbol{\omega} = c \cdot \mathbf{p}_{zt} \pmod{x_0}$ denotes again the zero-testing vector corresponding to an encoding c of zero, and $\mathbf{T} = (\mathbf{H}^{-1})^T$ is the inverse zero-testing matrix, then, by Equation (10):

$$\mathbf{T}\boldsymbol{\omega} = (R_1, \dots, R_n) \in \mathbb{Z}^n,$$

where, as above, $R_i = r_i \cdot (x_0/p_i)$. In particular, if the lines of \mathbf{T} are written as \mathbf{t}_i , we get that for each $i \in \{1, \dots, n\}$:

$$p_i \mathbf{t}_i \cdot \boldsymbol{\omega} \equiv 0 \pmod{x_0}.$$

Thus, since $p_i \mathbf{t}_i$ is relatively small, we might hope to recover it as a short vector in the lattice of vectors orthogonal to $\boldsymbol{\omega}$ modulo x_0 . However, as we have seen previously, we expect a reduced basis of that lattice to have vectors of length roughly 2^η , whereas $p_i \mathbf{t}_i$ is of length about $2^{\eta+\beta}$. Therefore, provided that, say, β is super-logarithmic, there are exponentially many linear combinations of the reduced basis vectors under the target length, and hence we cannot hope to recover $p_i \mathbf{t}_i$ in that fashion.

A more sophisticated attack based on the same observation uses the result of Hermann and May [HM08] on solving linear equations modulo an unknown factor of a public modulus. In our case, the components t_{ij} of \mathbf{t}_i form a small solution (smaller than 2^β) of the linear equation:

$$\sum_{j=1}^n \omega_j \cdot t_{ij} \equiv 0 \pmod{x_0/p_i}$$

modulo the unknown factor x_0/p_i of x_0 . The technique of Hermann–May can thus recover \mathbf{t}_i and factor x_0 provided that β is small enough relative to x_0 . For sufficiently large n , by [HM08, Theorem 4], this should be possible as long as $\beta/\eta < 1$. However, as noted by the authors, the complexity of the attack is exponential in n (it involves reducing a lattice of dimension $\Omega(\exp n)$), and there does not seem to be a way to approximate the solution in polynomial time, so the attack does not apply to our setting even though we do choose $\beta < \eta$.

6 Optimizations and Implementation

In this section we describe an implementation of our scheme in the one-round N -way Diffie-Hellman key exchange protocol; we recall the protocol in Appendix J as described in [BS03,GGH13].

We note that without optimizations the size of the public parameters makes our scheme completely unpractical; this is also the case in [GGH13]. Namely for sampling we need to store at least $n \cdot \alpha$ encodings (resp. $n \cdot \rho$ encodings for re-randomization), each of size $n \cdot \eta$ bits; the public-key size is then at least $n^2 \cdot \eta \cdot \alpha$ bits. With $n \simeq 10^4$, $\eta \simeq 10^3$ and $\alpha \simeq 80$, the public-key size would be at least 1 TB.³ Therefore we describe three heuristic optimizations to reduce the memory requirement:

1. Non-uniform sampling: for the sampling algorithm we use a small number of encodings ℓ only; this implies that the sampling cannot be proved uniform anymore.
2. Quadratic re-randomization: we only store a small subset of encodings which are later combined pairwise to generate the full set of encodings. This implies that the randomization of encodings becomes heuristic only.
3. Integer p_{zt} : we use a single integer p_{zt} instead of a vector \mathbf{p}_{zt} with n components. An encoding c of zero still gives a small integer $\omega = p_{zt} \cdot c \pmod{x_0}$, but the converse does not necessarily hold anymore.

6.1 Non-uniform Sampling

For sampling level-zero encodings we use a smaller value for ℓ , the number of encodings x_j in the public parameters. There is a simple meet-in-the-middle attack with complexity $\mathcal{O}(2^{\ell/2})$; therefore

³ In [GGH13] the following approximate setting is suggested: $n = \tilde{\mathcal{O}}(\kappa\lambda^2)$, $q = 2^{n/\lambda}$ and $m = \mathcal{O}(n^2)$. The public-key size contains at least m encodings of size $n \log_2 q$ bits each. Taking exactly $n = \kappa\lambda^2$ and $m = n^2$, the public-key size is then $m \cdot n \cdot (n/\lambda) = n^4/\lambda = \kappa^4\lambda^7$. With $\kappa = 6$ and $\lambda = 80$, we get a public-key size of 3100 TB.

we take $\ell = 2\lambda$. In this case the sampling cannot be proved uniform in $QR = \mathbb{Z}_{g_1} \times \cdots \times \mathbb{Z}_{g_n}$ anymore. However this does not seem to make the GDDH problem easier.

Note also that for such small ℓ given a level-0 encoding c , one can efficiently recover the coefficients of the subset sum with LLL, since this is a subset-sum problem with density $\ell/(\eta \cdot n) \ll 1$; however this does not give an attack, as in GDDH such level-0 encoding c is not available.⁴

6.2 Quadratic Re-randomization

To reduce the parameters size we use a simple optimization that consists in storing only a small subset of the public elements and combining them pairwise to generate the full public-key. Such optimization was originally described in [GH11] for reducing the size of the encryption of the secret-key bits in the implementation of Gentry's FHE scheme [Gen09]. It was also used in [CMNT11] to reduce the public-key size of the DGHV scheme; however as opposed to [CMNT11] our randomization of encodings becomes heuristic only, whereas in [CMNT11] the semantic security was still guaranteed.

For re-randomization we only store $\Delta = \lfloor \sqrt{n} \rfloor$ encoding $x_j^{(0)}$ at level 0 and also Δ encodings $x_j^{(1)}$ at level 1. The $x_j^{(0)}$ encode random $\mathbf{m}_j \in QR$, while the $x_j^{(1)}$ are encodings of $\mathbf{0}$. Then by pairwise multiplication we can generate $\Delta^2 \simeq n$ randomization elements at level 1, which are all encodings of $\mathbf{0}$. More precisely we have for $b = 0, 1$ and $1 \leq j \leq \Delta$:

$$x_j^{(b)} \equiv \frac{r_{ij}^{(b)} \cdot g_i + (1 - b) \cdot f_{ij}}{z^b} \pmod{p_i},$$

where $r_{ij}^{(b)}$ are random ρ -bit integers, and f_{ij} are random integers modulo g_i .

Given a level-1 encoding c_1 , we randomize it using a random subset-sum of pairwise products of the previous encodings:

$$c'_1 = c_1 + \sum_{i,j=1}^{\Delta} \alpha_{ij} \cdot x_i^{(0)} \cdot x_j^{(1)} \pmod{x_0},$$

where the α_{ij} 's are random bits; note that we don't use the encodings Π_j anymore. As a further optimization, we can use as in [CMNT11] a sparse vector α_{ij} , with small Hamming weight θ . There is a meet-in-the-middle attack of complexity $\mathcal{O}(n^{\theta/2})$. In our implementation we take $\theta = 16$; the reRand operation then becomes very efficient.

6.3 Zero-Testing Element

Instead of generating a zero-testing vector \mathbf{p}_{zt} with n components, we publish a zero-testing element p_{zt} which is a single integer:

$$p_{zt} = \sum_{i=1}^n h_i \cdot (z^\kappa \cdot g_i^{-1} \pmod{p_i}) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}$$

where the h_i 's are random β -bit integers. Therefore, we obtain a single integer $\omega = p_{zt} \cdot c \pmod{x_0}$, with

$$\omega = \sum_{i=1}^n h_i \cdot (r_i + m_i \cdot (g_i^{-1} \pmod{p_i})) \cdot \prod_{i' \neq i} p_{i'} \pmod{x_0}.$$

As previously if $\|\mathbf{r}\|_\infty < 2^{\rho'}$ we still have $|\omega| < x_0 \cdot 2^{-\nu}$. However the converse is no longer true: we can have $|\omega| < x_0 \cdot 2^{-\nu}$ for an encoding of a non-zero vector \mathbf{m} . This implies that two encodings of different vectors can now extract to the same value. While it is actually easy to generate such collisions using LLL, this does not seem to give an attack against the GDDH problem.

⁴ Alternatively one could use the same quadratic encryption technique as in [CMNT11]; namely the technique in [CMNT11] applies not only modulo an integer q but in any finite group; in that case the sampling could still be proved uniform in QR .

6.4 Parameters and Timings

We have implemented a one-round N -way Diffie-Hellman key exchange protocol with $N = 7$ users, in C++ using the Gnu MP library [Gt13] to perform operations on large integers. We refer to Appendix J for a description of the protocol. We provide our concrete parameters and the resulting timings in Table 1, for security parameters ranging from 52 to 80 bits.⁵

Table 1. Parameters and timings to instantiate a *one-round 7-way Diffie-Hellman key exchange protocol of a 160-bit session key* with $\ell = 160$, $\beta = 80$, $\alpha = 80$ and $N = 7$ (i.e. $\kappa = 6$), on a 16-core computer (Intel(R) Xeon(R) CPU E7-8837 at 2.67GHz) using GMP 5.1.1. Note that the Setup was parallelized on the 16 cores to speed-up the process while the other steps ran on a single core.

Instantiation	λ	n	η	Δ	ρ	pk size	Setup (once)	Publish (per party)	KeyGen (per party)
Small	52	540	1779	23	41	24 MB	4.5s	0.16s	0.10s
Medium	62	2100	1974	45	56	125 MB	22.3s	0.58s	0.62s
Large	72	8355	2182	91	72	700 MB	1292s	3.6s	4.0s
Extra	80	26265	2351	162	85	2.5 GB	19820s	12s	14s

The timings above show that our scheme is relatively practical, as the KeyGen phase of the multipartite Diffie-Hellman protocol requires only a few seconds per user; however the parameter size is still very large even with our optimizations.

References

- [AGHS12] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian Leftover Hash Lemma over infinite domains. Cryptology ePrint Archive, Report 2012/714, 2012. <http://eprint.iacr.org/>.
- [Bab86] László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [BS03] D. Boneh and A. Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324:71–90, 2003.
- [CCK⁺13] Jung Hee Cheon, Jean-Sébastien Coron, Jinsu Kim, Moon Sung Lee, Tancrede Lepoint, Mehdi Tibouchi, and Aaram Yun. Batch fully homomorphic encryption over the integers. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, Lecture Notes in Computer Science. Springer, 2013. To appear. Merged version of [CLT13] and [KLYC13].
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Batch fully homomorphic encryption over the integers. Cryptology ePrint Archive, Report 2013/036, 2013. <http://eprint.iacr.org/>.
- [CMNT11] Jean-Sébastien Coron, Avradip Mandal, David Naccache, and Mehdi Tibouchi. Fully homomorphic encryption over the integers with shorter public keys. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 487–504. Springer, 2011.
- [CN12] Yuanmi Chen and Phong Nguyen. Faster algorithms for approximate common divisors: Breaking fully-homomorphic-encryption challenges over the integers. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 502–519. Springer, 2012.
- [DGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2010.
- [Gen09] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. <http://crypto.stanford.edu/craig>.
- [GGH13] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices and applications. In *EUROCRYPT 2013*, Lecture Notes in Computer Science. Springer, 2013. Cryptology ePrint Archive, Report 2012/610. <http://eprint.iacr.org/>.
- [GH11] Craig Gentry and Shai Halevi. Implementing Gentry’s fully-homomorphic encryption scheme. In Kenneth Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 129–148. Springer, 2011.

⁵ More precisely for a security level λ we expect that the best attack requires at least 2^λ clock cycles. Note that we can take $\rho < 2\lambda$ because the GCD attack from [CN12] has a large polynomial factor in front of $2^{\rho/2}$.

- [Gt13] Torbjörn Granlund and the GMP development team. *GNU MP: The GNU Multiple Precision Arithmetic Library*, 5.1.1 edition, 2013. <http://gmplib.org/>.
- [HILL99] Johan Håstad, Russel Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28:12–24, 1999.
- [HM08] Mathias Herrmann and Alexander May. Solving linear equations modulo divisors: On factoring given any bits. In Josef Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *Lecture Notes in Computer Science*, pages 406–424. Springer, 2008.
- [HPS11] Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 447–464. Springer, 2011.
- [KLYC13] Jinsu Kim, Moon Sung Lee, Aaram Yun, and Jung Hee Cheon. CRT-based fully homomorphic encryption over the integers. Cryptology ePrint Archive, Report 2013/057, 2013. <http://eprint.iacr.org/>.
- [NS99] Phong Q. Nguyen and Jacques Stern. The hardness of the hidden subset sum problem and its cryptographic implications. In Michael J. Wiener, editor, *CRYPTO '99*, volume 1666 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 1999.
- [Pri51] G. Baley Price. Bounds for determinants with dominant principal diagonal. *Proceedings of the American Mathematical Society*, 2(3):497–502, 1951.
- [Sto96] Arne Storjohann. Near optimal algorithms for computing Smith Normal Forms of integer matrices. In Erwin Engeler, B. F. Caviness, and Yagati N. Lakshman, editors, *ISSAC '96*, pages 267–274. ACM, 1996.
- [Var75] James M. Varah. A lower bound for the smallest singular value of a matrix. *Linear Algebra and its Applications*, 11(1):3–5, 1975.

A Graded Encoding System

Definition 1 (κ -Graded Encoding System [GGH13]). A κ -Graded Encoding System for a ring R is a system of sets $\mathcal{S} = \{S_v^{(\alpha)} \in \{0, 1\}^* : v \in \mathbb{N}, \alpha \in R\}$, with the following properties:

1. For every $v \in \mathbb{N}$, the sets $\{S_v^{(\alpha)} : \alpha \in R\}$ are disjoint.
2. There are binary operations $+$ and $-$ (on $\{0, 1\}^*$) such that for every $\alpha_1, \alpha_2 \in R$, every $v \in \mathbb{N}$, and every $u_1 \in S_v^{(\alpha_1)}$ and $u_2 \in S_v^{(\alpha_2)}$, it holds that $u_1 + u_2 \in S_v^{(\alpha_1 + \alpha_2)}$ and $u_1 - u_2 \in S_v^{(\alpha_1 - \alpha_2)}$ where $\alpha_1 + \alpha_2$ and $\alpha_1 - \alpha_2$ are addition and subtraction in R .
3. There is an associative binary operation \times (on $\{0, 1\}^*$) such that for every $\alpha_1, \alpha_2 \in R$, every v_1, v_2 with $0 \leq v_1 + v_2 \leq \kappa$, and every $u_1 \in S_{v_1}^{(\alpha_1)}$ and $u_2 \in S_{v_2}^{(\alpha_2)}$, it holds that $u_1 \times u_2 \in S_{v_1 + v_2}^{(\alpha_1 \cdot \alpha_2)}$ where $\alpha_1 \cdot \alpha_2$ is multiplication in R .

B Multilinear map Procedures

We also recall the definition of the procedures for manipulating encodings. As previously we consider only the symmetric case; we refer to [GGH13] for the general case.

Instance Generation. The randomized $\text{InstGen}(1^\lambda, 1^\kappa)$ takes as inputs the parameters λ and κ , and outputs $(\text{params}, \mathbf{p}_{zt})$, where params is a description of a κ -Graded Encoding System as above, and \mathbf{p}_{zt} is a zero-test parameter.

Ring Sampler. The randomized $\text{samp}(\text{params})$ outputs a “level-zero encoding” $a \in S_0^{(\alpha)}$ for a nearly uniform element $\alpha \in R$. Note that the encoding a does not need to be uniform in $S_0^{(\alpha)}$.

Encoding. The (possibly randomized) $\text{enc}(\text{params}, a)$ takes as input a level-zero encoding $a \in S_0^{(\alpha)}$ for some $\alpha \in R$, and outputs the level-one encoding $u \in S_1^{(\alpha)}$ for the same α .

Re-Randomization. The randomized $\text{reRand}(\text{params}, i, u)$ re-randomizes encodings relative to the same level i . Specifically, given an encoding $u \in S_v^{(\alpha)}$, it outputs another encoding $u' \in S_v^{(\alpha)}$. Moreover for any two $u_1, u_2 \in S_v^{(\alpha)}$, the output distributions of $\text{reRand}(\text{params}, i, u_1)$ and $\text{reRand}(\text{params}, i, u_2)$ are nearly the same.

Addition and negation. Given params and two encodings relative to the same level, $u_1 \in S_i^{(\alpha_1)}$ and $u_2 \in S_i^{(\alpha_2)}$, we have $\text{add}(\text{params}, u_1, u_2) \in S_i^{(\alpha_1 + \alpha_2)}$ and $\text{neg}(\text{params}, u_1) \in S_i^{(-\alpha_1)}$. Below we write $u_1 + u_2$ and $-u_1$ as a shorthand for applying these procedures.

Multiplication. For $u_1 \in S_i^{(\alpha_1)}$ and $u_2 \in S_j^{(\alpha_2)}$, we have $\text{mul}(\text{params}, u_1, u_2) = u_1 \times u_2 \in S_{i+j}^{(\alpha_1 \cdot \alpha_2)}$.

Zero-test. The procedure $\text{isZero}(\text{params}, \mathbf{p}_{zt}, u)$ outputs 1 if $u \in S_\kappa^{(0)}$ and 0 otherwise.

Extraction. The procedure extracts a random function of ring elements from their level- κ encoding. Namely $\text{ext}(\text{params}, \mathbf{p}_{zt}, u)$ outputs $s \in \{0, 1\}^\lambda$, such that:

1. For any $\alpha \in R$ and $u_1, u_2 \in S_\kappa^{(\alpha)}$, $\text{ext}(\text{params}, \mathbf{p}_{zt}, u_1) = \text{ext}(\text{params}, \mathbf{p}_{zt}, u_2)$.
2. The distribution $\{\text{ext}(\text{params}, \mathbf{p}_{zt}, u) : \alpha \in R, u \in S_\kappa^{(\alpha)}\}$ is nearly uniform over $\{0, 1\}^\lambda$.

This concludes the definition of the procedures. In [GGH13] the authors consider a slightly relaxed definition of isZero and ext , where isZero can still output 1 even for some non-zero encoding u with negligible probability, and ext can extract different outputs when applied to encodings of the same elements, also with negligible probability; see [GGH13] for the corresponding definitions.

C Hardness Assumptions

We recall the hardness assumptions for multilinear maps from [GGH13]; as previously we consider only the symmetric case and refer to [GGH13] for the general case. In this symmetric case given a set of $\kappa + 1$ level-one encodings of random elements, it should be unfeasible to distinguish a level- κ encoding of their product from random.

Graded DDH (GDDH). Let \mathcal{GE} be a graded encoding scheme consisting of all the routines above. For an adversary \mathcal{A} and parameters λ, κ we consider the following process:

1. $(\text{params}, \mathbf{p}_{zt}) \leftarrow \text{InstGen}(1^\lambda, 1^\kappa)$
2. Choose $a_j \leftarrow \text{samp}(\text{params})$ for all $1 \leq j \leq \kappa + 1$
3. Set $u_j \leftarrow \text{reRand}(\text{params}, 1, \text{enc}(\text{params}, 1, a_j))$ for all $1 \leq j \leq \kappa + 1$ // encodings at level 1
4. Choose $b \leftarrow \text{samp}(\text{params})$ // encoding at level 0
5. Set $\tilde{u} = a_{\kappa+1} \times \prod_{i=1}^{\kappa} u_i$ // encoding of the right product at level κ
6. Set $\hat{u} = b \times \prod_{i=1}^{\kappa} u_i$ // encoding of a random product at level κ

The GDDH distinguisher is given as input the $\kappa + 1$ level-one encodings u_j and either \tilde{u} (encoding the right product) or \hat{u} (encoding a random product), and must decide which is the case. The GDDH assumption states that the advantage of any efficient adversary is negligible in the security parameter.

D Proof of Lemma 1

Write $x'_j \equiv r'_{ij} \cdot g_i + a_{ij} \pmod{p_i}$ for $1 \leq j \leq \ell$ and $1 \leq i \leq n$. Each component a_{ij} is randomly generated in $[0, g_i) \cap \mathbb{Z}$; therefore the column vectors \mathbf{a}_j of the matrix (a_{ij}) are randomly and independently generated in QR . By the leftover hash lemma over finite groups (Lemma 5), we have that $(\mathbf{a}_1, \dots, \mathbf{a}_\ell, \mathbf{m})$ with $\mathbf{m} = \mathbf{A} \cdot \mathbf{b}$ is ε -uniform over $QR^{\ell+1}$, where

$$\varepsilon = \frac{1}{2} \sqrt{\frac{|QR|}{2^\ell}} \leq 2^{(\alpha \cdot n - \ell)/2}.$$

Therefore by taking $\ell \geq n \cdot \alpha + 2\lambda$ we obtain that the distribution of $(\text{params}, \mathbf{m})$ is statistically close to the distribution of $(\text{params}, \mathbf{m}')$ for $\mathbf{m}' \leftarrow QR$. \square

E Proof of Lemma 6

We consider the intermediate variable:

$$\mathbf{y}'' = \left(\sum_{i=1}^m s_i \mathbf{x}_i \bmod \mathbf{B} \right) + \sum_{i=1}^n t_i \mathbf{b}_i. \quad (12)$$

We proceed in two steps. Firstly by applying the leftover hash lemma over the finite abelian group $G = \mathbb{Z}^n/L$, we obtain that the distributions $(\mathbf{x}_1, \dots, \mathbf{x}_m, \sum_{i=1}^m s_i \mathbf{x}_i \bmod \mathbf{B})$ and $(\mathbf{x}_1, \dots, \mathbf{x}_m, \boldsymbol{\psi})$ are ε_1 -statistically close, where $\boldsymbol{\psi} \leftarrow \mathcal{D}_{\mathbf{B}}$ and

$$\varepsilon_1 = 1/2\sqrt{|G|/2^m} = 1/2\sqrt{|\det(L)|/2^m}.$$

This implies that the distributions $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}'')$ and $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}')$ are also ε_1 -statistically close.

Secondly we write:

$$\sum_{i=1}^m s_i \mathbf{x}_i \bmod \mathbf{B} = \sum_{i=1}^m s_i \mathbf{x}_i - \sum_{j=1}^n \chi_j \mathbf{b}_j \quad (13)$$

where $\chi_j \in \mathbb{Z}$ for all j . We also write $\mathbf{x}_i = \sum_j \xi_{ij} \mathbf{b}_j$ where by definition $0 \leq \xi_{ij} < 1$ for all i, j . This gives:

$$\sum_{i=1}^m s_i \mathbf{x}_i \bmod \mathbf{B} = \sum_{i=1}^m s_i \sum_{j=1}^n \xi_{ij} \mathbf{b}_j - \sum_{j=1}^n \chi_j \mathbf{b}_j = \sum_{j=1}^n \left(\sum_{i=1}^m s_i \xi_{ij} - \chi_j \right) \mathbf{b}_j,$$

which implies $0 \leq \sum_{i=1}^m s_i \xi_{ij} - \chi_j < 1$ for all j , and therefore $0 \leq \chi_j \leq m$ for all j . Combining Equations (12) and (13) we have:

$$\mathbf{y}'' = \sum_{i=1}^m s_i \mathbf{x}_i + \sum_{i=1}^n (t_i - \chi_i) \mathbf{b}_i,$$

where as shown above $0 \leq \chi_i \leq m$ for all i . This implies that the distributions $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$ and $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}'')$ are ε_2 -statistically close, with $\varepsilon_2 = mn2^{-\mu}$. Therefore the distributions $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y})$ and $(\mathbf{x}_1, \dots, \mathbf{x}_m, \mathbf{y}')$ are $(\varepsilon_1 + \varepsilon_2)$ -statistically close, which proves the Lemma. \square

F Proof of Lemma 2

We use the following fact for diagonally dominant matrices [Pri51].

Fact 2 Let $B = (b_{ij}) \in \mathbb{R}^{n \times n}$ be a diagonally dominant matrix. Then

$$\prod_{i=1}^n (|b_{ii}| - \Lambda_i(B)) \leq |\det B| \leq \prod_{i=1}^n (|b_{ii}| + \Lambda_i(B))$$

Proof (of Lemma 2). We write $c_i \equiv (r_i^{(1)} \cdot g_i + m_i)/z \bmod p_i$ for all $1 \leq i \leq n$ and define $\mathbf{r}^{(1)} = (r_i^{(1)}) \in \mathbb{Z}^n$. We also write $x_j \equiv r_{ij} \cdot g_i/z \pmod{p_i}$ and $\Pi_j \equiv \varpi_{ij} \cdot g_i/z \pmod{p_i}$ and define the matrix $\mathbf{X} = (r_{ij}) \in \mathbb{Z}^{n \times \tau}$ and $\mathbf{\Pi} = (\varpi_{ij}) \in \mathbb{Z}^{n \times n}$. From the re-randomization equation (7), we can write:

$$\mathbf{r} = \mathbf{r}^{(1)} + \mathbf{X} \cdot \mathbf{b} + \mathbf{\Pi} \cdot \mathbf{b}',$$

where $\mathbf{b} \leftarrow \{0, 1\}^\tau$, and $\mathbf{b}' \leftarrow ([0, 2^\mu] \cap \mathbb{Z})^n$.

Since at instance generation the columns of \mathbf{X} are generated uniformly and independently in the parallelepiped spanned by the columns of $\mathbf{\Pi}$, applying our “leftover hash lemma over

lattices" (Lemma 6) we obtain that the distribution of $(\text{params}, \mathbf{r})$ is ε_1 -close to the distribution of $(\text{params}, \mathbf{r}^{(1)} + \mathcal{D}_{2^\mu \mathbf{\Pi}})$, with

$$\varepsilon_1 = \tau n 2^{-\mu} + \frac{1}{2} \sqrt{|\det \mathbf{\Pi}|/2^\tau}.$$

Since $\mathbf{\Pi}$ is a diagonally dominant matrix, we obtain from Fact 2

$$|\det \mathbf{\Pi}| \leq \prod_{i=1}^n (|\varpi_{i,i}| + \Lambda_i(\mathbf{\Pi})) \leq (2n2^\rho)^n \leq 2^{n(\rho + \log_2(2n))}.$$

which gives $\varepsilon_1 \leq n\tau 2^{-\mu} + 2^{(n(\rho + \log_2(2n)) - \tau)/2}$. Therefore given the constraints $\tau \geq n \cdot (\rho + \log_2(2n)) + 2\lambda$ and $\mu \geq \rho + \alpha + \lambda$ we have that $\varepsilon_1 = \text{negl}(\lambda)$.

Now, using Lemma 7 we obtain that the distribution of $(\text{params}, \mathbf{r})$ is $(\varepsilon_1 + \varepsilon_2)$ -close to that of $(\text{params}, \mathcal{D}_{2^\mu \mathbf{\Pi}})$ for

$$\varepsilon_2 = 2^{-\mu} \left(\|\mathbf{r}^{(1)}\|_\infty \cdot \|\mathbf{\Pi}^{-1}\|_\infty + 1 \right).$$

We consider the initial level-0 encoding c and write $c \equiv r_j^{(c)} \cdot g_j + m_j \pmod{p_j}$ and $y \equiv r_j^{(y)} \cdot g_j + 1 \pmod{p_j}$, and define

$$\mathbf{r}^{(c)} = (r_1^{(c)}, \dots, r_n^{(c)})^T \quad \text{and} \quad \mathbf{r}^{(y)} = (r_1^{(y)}, \dots, r_n^{(y)})^T.$$

Since $c_1 = c \cdot y \pmod{x_0}$, we have

$$\|\mathbf{r}^{(1)}\|_\infty = \|(r_j^{(c)} + r_j^{(c)} r_j^{(y)} g_j + m_j r_j^{(y)})_{j=1, \dots, n}\|_\infty \leq \|\mathbf{r}^{(c)}\|_\infty \cdot \|\mathbf{r}^{(y)}\|_\infty \cdot 2^{\alpha+2}.$$

Therefore $\|\mathbf{r}^{(1)}\|_\infty \leq 2^{2\rho + \log_2(\ell) + \alpha + 2}$. Now, by Fact 1, we have

$$\|\mathbf{\Pi}^{-1}\|_\infty \leq \frac{1}{\min_{i=1, \dots, n} (|\varpi_{i,i}| - \Lambda_i(\mathbf{\Pi}))} \leq \frac{1}{n2^\rho - (n-1)2^\rho} \leq 2^{-\rho}.$$

This gives $\varepsilon_2 \leq 2^{-\mu} + 2^{\rho + \log_2(\ell) + \alpha + 2 - \mu}$. With the constraint $\mu \geq \rho + \alpha + \lambda$, Lemma 2 is proved. \square

G Proof of Lemma 3

We consider the vector $\mathbf{R} = (R_i)_{1 \leq i \leq n}$ where:

$$R_i = ((r_i + m_i \cdot g_i^{-1}) \pmod{p_i}) \cdot (x_0/p_i) \tag{14}$$

Equation (10) can then be written:

$$\boldsymbol{\omega} = \mathbf{H}^T \cdot \mathbf{R} \pmod{x_0} \tag{15}$$

If $\mathbf{m} = 0$ then we have $R_i = r_i \cdot x_0/p_i$ for all i , which gives using $p_i \geq 2^{\eta-1}$ for all i :

$$\|\mathbf{R}\|_\infty \leq \|\mathbf{r}\|_\infty \cdot \max_{1 \leq i \leq n} (x_0/p_i) \leq \|\mathbf{r}\|_\infty \cdot x_0 \cdot 2^{-\eta+1}$$

Since by definition $-p/2 < (z \pmod{p}) \leq p/2$, we have $|z \pmod{p}| \leq |z|$ for any z, p ; therefore we obtain from (15) using $\|\mathbf{r}\|_\infty < 2^{\rho_f}$

$$\|\boldsymbol{\omega}\|_\infty = \|\mathbf{H}^T \cdot \mathbf{R} \pmod{x_0}\|_\infty \leq \|\mathbf{H}^T \cdot \mathbf{R}\|_\infty \leq \|\mathbf{H}^T\|_\infty \cdot \|\mathbf{R}\|_\infty < x_0 \cdot 2^{\beta + \rho_f - \eta + 1} = x_0 \cdot 2^{-\nu - 2}$$

Conversely assume that $\|\boldsymbol{\omega}\|_\infty < x_0 \cdot 2^{-\nu+2}$. From (15) we have:

$$\mathbf{R} \equiv (\mathbf{H}^{-1})^T \cdot \boldsymbol{\omega} \pmod{x_0} \tag{16}$$

From (14) we have $\|\mathbf{R}\|_\infty < x_0/2$. Moreover

$$\|(\mathbf{H}^{-1})^T \cdot \boldsymbol{\omega}\|_\infty \leq \|(\mathbf{H}^{-1})^T\|_\infty \cdot \|\boldsymbol{\omega}\|_\infty \leq x_0 \cdot 2^{\beta-\nu+2} < x_0/2. \quad (17)$$

This shows that Equation (16) holds in \mathbb{Z} and not only modulo x_0 ; therefore we must have

$$\|\mathbf{R}\|_\infty \leq x_0 \cdot 2^{\beta-\nu+2}.$$

Letting $v_i = (r_i + m_i \cdot g_i^{-1}) \bmod p_i$ for $1 \leq i \leq n$, this gives $|v_i| \cdot (x_0/p_i) \leq x_0 \cdot 2^{\beta-\nu+2}$, and therefore $|v_i| \leq p_i \cdot 2^{\beta-\nu+2}$ for all i . We have $g_i \cdot (v_i - r_i) \equiv m_i \pmod{p_i}$; we show that the equality actually holds over \mathbb{Z} . Namely for all i we have $|m_i| < g_i < p_i/2$ and:

$$|g_i \cdot (v_i - r_i)| \leq |g_i| \cdot (|v_i| + |r_i|) \leq p_i \cdot 2^{\alpha+\beta-\nu+2} + 2^{\alpha+\rho_f} \leq p_i/8 + p_i/8 < p_i/2$$

which implies that the equality holds over \mathbb{Z} . Therefore $m_i \equiv 0 \pmod{g_i}$ for all i , which implies $\mathbf{m} = \mathbf{0}$. This proves Lemma 3. \square

H Proof of Lemma 7

Let $\mathbf{u} \leftarrow \mathcal{D}_{2^\mu \mathbf{B}}$ and $\mathbf{u}'' \leftarrow \mathcal{D}_{\mathbf{z}+2^\mu \mathbf{B}}$. We can write:

$$\begin{aligned} \mathbf{u} &= \mathbf{v} + \sum_{i=1}^n s_i \mathbf{b}_i \\ \mathbf{u}'' &= \mathbf{z} + \mathbf{v} + \sum_{i=1}^n s_i \mathbf{b}_i \end{aligned}$$

where $\mathbf{v} \leftarrow \mathcal{D}_{\mathbf{B}}$ and $s_i \leftarrow [0, 2^\mu) \cap \mathbb{Z}$. We consider the intermediate variable:

$$\mathbf{u}' = ((\mathbf{z} + \mathbf{v}) \bmod \mathbf{B}) + \sum_{i=1}^n s_i \mathbf{b}_i.$$

The distribution of \mathbf{u} and \mathbf{u}' are clearly the same. Let $\boldsymbol{\psi} = \mathbf{z} + \mathbf{v}$. We have:

$$\boldsymbol{\psi} \bmod \mathbf{B} = \boldsymbol{\psi} - \mathbf{B} \cdot \lfloor \mathbf{B}^{-1} \cdot \boldsymbol{\psi} \rfloor = \boldsymbol{\psi} - \sum_{i=1}^n t_i \mathbf{b}_i$$

where $\mathbf{t} = \lfloor \mathbf{B}^{-1} \cdot \boldsymbol{\psi} \rfloor$. This gives:

$$\mathbf{u}' = \mathbf{z} + \mathbf{v} + \sum_{i=1}^n (s_i - t_i) \mathbf{b}_i.$$

We have $\mathbf{t} = \lfloor \mathbf{B}^{-1} \cdot \mathbf{z} + \mathbf{B}^{-1} \cdot \mathbf{v} \rfloor$. Since \mathbf{v} is in the half-open parallelepiped we have that the components of $\mathbf{B}^{-1} \cdot \mathbf{v}$ are in $[0, 1)$, which gives:

$$\|\mathbf{t}\|_\infty \leq \|\mathbf{B}^{-1} \cdot \mathbf{z}\|_\infty + 1 \leq \|\mathbf{B}^{-1}\|_\infty \cdot \|\mathbf{z}\|_\infty + 1.$$

Therefore the variables \mathbf{u}' and \mathbf{u}'' are ε -statistically close, with $\varepsilon = 2^{-\mu} (\|\mathbf{B}^{-1}\|_\infty \cdot \|\mathbf{z}\|_\infty + 1)$. This proves the Lemma. \square

I Comparison with GGH Multilinear Maps

In this section we rewrite our scheme using exactly the same notations as in [GGH13] whenever possible, to better highlight the similarities.

The construction in [GGH13] works in the polynomial ring $R = \mathbb{Z}[X]/(X^n + 1)$, where n is large enough to ensure security. One generates a secret short ring element $\mathbf{g} \in R$, generating a principal ideal $\mathcal{I} = \langle \mathbf{g} \rangle \subset R$. One also generates an integer parameter q and another random secret $\mathbf{z} \in R/qR$. One encodes elements of the quotient ring R/\mathcal{I} , namely elements of the form $\mathbf{e} + \mathcal{I}$ for some \mathbf{e} , as follows: a level- i encoding of the coset $\mathbf{e} + \mathcal{I}$ is an element of the form $u_k = [\mathbf{c}/\mathbf{z}^i]_q$, where $\mathbf{c} \in \mathbf{e} + \mathcal{I}$ is short. Such encodings can be both added and multiplied, as long as the norm of the numerators remain shorter than q ; in particular the product of κ encodings at level 1 gives an encoding at level κ . For such level- κ encodings one can then define a zero-testing parameter $\mathbf{p}_{zt} = [\mathbf{h}\mathbf{z}^\kappa/\mathbf{g}]_q$, for some small $\mathbf{h} \in R$. Then given a level- κ encoding $\mathbf{u} = [\mathbf{c}/\mathbf{z}^\kappa]$ one can compute $[\mathbf{p}_{zt} \cdot \mathbf{u}]_q = [\mathbf{h}\mathbf{c}/\mathbf{g}]_q$; when \mathbf{c} is an encoding of zero we have $\mathbf{c}/\mathbf{g} \in R$, which implies that $\mathbf{h}\mathbf{c}/\mathbf{g}$ is small in R , and therefore $[\mathbf{h}\mathbf{c}/\mathbf{g}]_q$ is small; this provides a way to test whether a level- κ encoding \mathbf{c} is an encoding of 0.

In our construction one could write $R = \mathbb{Z}^n$, and define a secret short ring element $\mathbf{g} \in R$, generating a principal ideal $\mathcal{I} = \langle \mathbf{g} \rangle \subset R$, which gives $\mathcal{I} = (g_i \mathbb{Z})_{1 \leq i \leq n}$. We also generate a ring element $\mathbf{p} \in R$ and let the principal ideal $\mathcal{J} = \langle \mathbf{p} \rangle \subset R$, which gives $\mathcal{J} = (p_i \mathbb{Z})_{1 \leq i \leq n}$. We let $q := x_0 = \prod_{i=1}^n p_i$ and for convenience we denote by $[\mathbf{u}]_q$ the CRT isomorphism from R/\mathcal{J} to \mathbb{Z}_q . As in [GGH13], in our scheme a level- i encoding of the coset $\mathbf{e}_{\mathcal{I}} = \mathbf{e} + \mathcal{I}$ is an element of the form $u = [\mathbf{c}/\mathbf{z}^i]_q$ where $\mathbf{c} \in \mathbf{e}_{\mathcal{I}}$ is short. Such encodings can be both added and multiplied, by working over the integers via the CRT isomorphism $[\cdot]_q$.

However, we cannot apply the zero-testing procedure from [GGH13] in a straightforward way. Namely one could define the zero-testing parameter $p_{zt} = [\mathbf{h}\mathbf{z}^\kappa/\mathbf{g}]_q$ as in [GGH13] where $\mathbf{h} \in \mathbb{Z}^n$ is a relatively small ring element. As in [GGH13] given a level- κ encoding $u = [\mathbf{c}/\mathbf{z}^\kappa]_q$ one would compute the element:

$$\omega = p_{zt} \cdot u = \left[\frac{\mathbf{h}\mathbf{z}^\kappa}{\mathbf{g}} \cdot \frac{\mathbf{c}}{\mathbf{z}^\kappa} \right]_q = \left[\frac{\mathbf{h}\mathbf{c}}{\mathbf{g}} \right]_q. \quad (18)$$

As in [GGH13] if u is an encoding of 0 then \mathbf{c} is a multiple of \mathbf{g} over \mathbb{Z}^n hence $\mathbf{c}/\mathbf{g} \in \mathbb{Z}^n$ is short and therefore the vector $\mathbf{h}\mathbf{c}/\mathbf{g} \in \mathbb{Z}^n$ is short. However this does not imply that the corresponding integer ω obtained by CRT in (18) is small, and we do not have a simple way of identifying integers whose reductions modulo the unknown p_i 's are small.

Instead, we can define a slightly different notation: we consider the following additive homomorphism $R/\mathcal{J} \rightarrow \mathbb{Z}_q$

$$\mathbf{u} \rightarrow \{\mathbf{u}\}_q = \sum_{i=1}^n u_i \cdot \prod_{i' \neq i} p_{i'} \bmod q$$

where as before $q = x_0 = \prod_{i=1}^n p_i$ and we define the zero-testing parameter:

$$p_{zt} := \{\mathbf{h}\mathbf{z}^\kappa/\mathbf{g}\}_q$$

As in [GGH13] given a level- κ encoding $u = [\mathbf{c}/\mathbf{z}^\kappa]_q$ one can compute the element:

$$\omega = p_{zt} \cdot u \bmod q = \left\{ \frac{\mathbf{h}\mathbf{z}^\kappa}{\mathbf{g}} \right\}_q \cdot \left[\frac{\mathbf{c}}{\mathbf{z}^\kappa} \right]_q = \left\{ \frac{\mathbf{h}\mathbf{z}^\kappa}{\mathbf{g}} \cdot \frac{\mathbf{c}}{\mathbf{z}^\kappa} \right\}_q = \left\{ \frac{\mathbf{h}\mathbf{c}}{\mathbf{g}} \right\}_q$$

As in [GGH13] if u is an encoding of 0 then \mathbf{c} is a multiple of \mathbf{g} over \mathbb{Z}^n hence $\mathbf{c}/\mathbf{g} \in \mathbb{Z}^n$ is short and therefore the vector $\mathbf{h}\mathbf{c}/\mathbf{g} \in \mathbb{Z}^n$ is short; this time, this implies that $\omega = \{\mathbf{h}\mathbf{c}/\mathbf{g}\}_q$ is a short integer.

J One-Round N -Way Diffie-Hellman Key Exchange Protocol

In 2003, Boneh and Silverberg showed how to perform a multipartite Diffie-Hellman key exchange using multilinear maps [BS03]. As in [GGH13], our scheme can be used to instantiate a one-round

N -party Diffie-Hellman key exchange protocol in the common reference string model, under the GDDH assumption with $N = \kappa + 1$ users.

Consider N parties wishing to set up a shared secret key s using a one-round protocol (i.e. each party broadcasts one value to all other parties). We recall the construction from [GGH13]:

Setup($1^\lambda, 1^N$). Output $(\text{params}, \mathbf{p}_{zt}) \leftarrow \text{InstGen}(1^\lambda, 1^\kappa)$ as the public parameter, with $\kappa = N - 1$.

Publish(params, i). Each party i samples a random $c_i \leftarrow \text{samp}(\text{params})$ as a secret key, and publish as the public key the corresponding level-1 encoding $c'_i \leftarrow \text{reRand}(\text{params}, 1, \text{enc}(\text{params}, 1, c_i))$.

KeyGen($\text{params}, \mathbf{p}_{zt}, i, c_i, \{c'_j\}_{j \neq i}$). Each party i computes $\tilde{c}_i = c_i \cdot \prod_{j \neq i} c'_j$, and uses the extraction routine to compute the key $s \leftarrow \text{ext}(\text{params}, \mathbf{p}_{zt}, \tilde{c}_i)$.

Theorem 3 ([GGH13]). *The protocol described above is a one-round N -way Diffie-Hellman key exchange protocol if the GDDH assumption holds for the underlying encoding scheme.*

K Uniform Sampling of a Parallelepiped

In order to sample a uniformly random element in the half-open parallelepiped defined by the column vectors $\boldsymbol{\varpi}_j$ of matrix $\mathbf{II} \in \mathbb{Z}^{n \times n}$, one can proceed as follows.

First, compute the Smith Normal Form for \mathbf{II} . This is easily done in polynomial time, and can be done with near optimal complexity using Storjohann's algorithm [Sto96]. This yields a basis $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ of \mathbb{Z}^n and positive integers d_1, \dots, d_n such that $(d_1 \mathbf{b}_1, \dots, d_n \mathbf{b}_n)$ is a basis of the lattice L generated by the columns of \mathbf{II} .

Then, if we pick integers x_1, \dots, x_n at random such that x_i is uniformly distributed in $\{0, \dots, d_i - 1\}$, then clearly, the vector $\mathbf{x} = x_1 \mathbf{b}_1 + \dots + x_n \mathbf{b}_n$ is uniformly distributed modulo L .

To get a uniformly distributed vector in the half-open parallelepiped defined by the $\boldsymbol{\varpi}_j$'s, it then suffices to apply Babai's round-off algorithm [Bab86], i.e. write \mathbf{x} as a rational linear combination $\xi_1 \boldsymbol{\varpi}_1 + \dots + \xi_n \boldsymbol{\varpi}_n$ of the $\boldsymbol{\varpi}_j$'s and return the vector \mathbf{y} given by:

$$\mathbf{y} = \sum_{j=1}^n (\xi_j - \lfloor \xi_j \rfloor) \boldsymbol{\varpi}_j.$$

That vector is congruent to \mathbf{x} modulo L , so it is also in \mathbb{Z}^n and uniformly distributed modulo L , and it belongs to the parallelepiped by construction, so it is indeed a uniformly distributed element of the parallelepiped.

L Generation of the Matrix \mathbf{H}

For the construction of zero-testing parameters, we need to pick, with sufficient entropy, an invertible matrix $\mathbf{H} \in \mathbb{Z}^{n \times n}$ in such a way that both its operator norm and the norm of its inverse are not too large, namely $\|\mathbf{H}\|_\infty \leq 2^\beta$ and $\|\mathbf{H}^{-1}\|_\infty \leq 2^\beta$. In Section 3 the bounds must actually hold for \mathbf{H}^T , so we take the transpose of the resulting matrix.

For that purpose, we propose the following approach. For any matrix A of size $\lceil n/2 \rceil \times \lceil n/2 \rceil$ with coefficients in $\{-1, 0, 1\}$, define $\mathbf{H}_A \in \mathbb{Z}^{n \times n}$ as:

$$\mathbf{H}_A = \begin{pmatrix} I_{\lceil n/2 \rceil} & A \\ 0 & I_{\lceil n/2 \rceil} \end{pmatrix}.$$

Each line of \mathbf{H}_A has at most $1 + \lceil n/2 \rceil$ non zero coefficients, each in $\{-1, 0, 1\}$, so we clearly have $\|\mathbf{H}_A\| \leq 1 + \lceil n/2 \rceil$. Moreover, \mathbf{H}_A is invertible with $\mathbf{H}_A^{-1} = \mathbf{H}_{(-A)}$, so that the operator norm of its inverse admits a similar bound.

Similarly, the transpose \mathbf{H}'_A of \mathbf{H}_A also satisfies $\|\mathbf{H}'_A\| \leq 1 + \lceil n/2 \rceil$ (in fact, the slightly better bound by $1 + \lfloor n/2 \rfloor$ also holds) and so does its inverse.

Now let:

$$\beta' = \left\lfloor \frac{\beta}{\lceil \log_2(1 + \lceil n/2 \rceil) \rceil} \right\rfloor,$$

and generate β' uniformly random matrices A_i of size $\lfloor n/2 \rfloor \times \lceil n/2 \rceil$ with coefficients in $\{-1, 0, 1\}$; then pick \mathbf{H}_i randomly as either \mathbf{H}_{A_i} or its transpose for each $i \in \{1, \dots, \beta'\}$, and finally compute \mathbf{H} as the product of the \mathbf{H}_i 's. Then, since operator norms are sub-multiplicative, we have:

$$\|\mathbf{H}\|_\infty \leq \prod_{i=1}^{\beta'} \|\mathbf{H}_i\|_\infty \leq (1 + \lceil n/2 \rceil)^{\beta'} \leq 2^\beta,$$

and \mathbf{H}^{-1} satisfies the same bound. The set of matrices \mathbf{H} obtained in this manner is not very simple to describe but it is exponentially large.