Graded Encoding Schemes from Obfuscation

Pooya Farshim[1,2], Julia Hesse[1,2,3], Dennis Hofheinz[3(✉)], and Enrique Larraia[4]

[1] DIENS, École normale supérieure, CNRS, PSL Research University, Paris, France
[2] Inria, Rocquencourt, France
[3] Karlsruhe Institute of Technology, Karlsruhe, Germany
dennis.hofheinz@kit.edu
[4] Royal Holloway, University of London, London, UK

Abstract. We construct a graded encoding scheme (GES), an approximate form of graded multilinear maps. Our construction relies on indistinguishability obfuscation, and a pairing-friendly group in which (a suitable variant of) the strong Diffie–Hellman assumption holds. As a result of this abstract approach, our GES has a number of advantages over previous constructions. Most importantly:

- We can *prove* that the multilinear decisional Diffie–Hellman (MDDH) assumption holds in our setting, assuming the used ingredients are secure (in a well-defined and standard sense). Hence, our GES does not succumb to so-called "zeroizing" attacks if the underlying ingredients are secure.
- Encodings in our GES do not carry any noise. Thus, unlike previous GES constructions, there is no upper bound on the number of operations one can perform with our encodings. Hence, our GES essentially realizes what Garg et al. (EUROCRYPT 2013) call the "dream version" of a GES.

Technically, our scheme extends a previous, non-graded approximate multilinear map scheme due to Albrecht et al. (TCC 2016-A). To introduce a graded structure, we develop a new view of encodings at different levels as polynomials of different degrees.

Keywords: Multilinear maps · Graded encoding schemes
Indistinguishability obfuscation

1 Introduction

The GGH candidate multilinear map. In 2013, Garg, Gentry, and Halevi (GGH) [22] proposed the first plausible construction of an (approximate) multilinear map (MLM). In a nutshell, an MLM is a map $e : \mathbb{G}^{\kappa} \longrightarrow \mathbb{G}_T$ (for groups $\mathbb{G}$ and $\mathbb{G}_T$) that is linear in each input. Of course, we are most interested in the case of "cryptographically interesting" groups $\mathbb{G}$ (in which, e.g., computing discrete logarithms is infeasible), non-trivial maps e (with non-trivial kernel),

M. Abdalla and R. Dahab (Eds.): PKC 2018, LNCS 10770, pp. 371–400, 2018.
https://doi.org/10.1007/978-3-319-76581-5_13

and preferably large values of κ. The surprising cryptographic consequences of such "cryptographically interesting" MLMs were already investigated in 2003 by Boneh and Silverberg [6], but an actual construction of an MLM remained elusive until the candidate construction of GGH.

Unfortunately, GGH only presented an "approximate" MLM in the following sense:

- Instead of group elements, their e inputs (and outputs) are *encodings*. An encoding is a non-unique representation of a group element, and there is no guarantee about which particular encoding the group operation (or e) outputs. However, every encoding allows to derive a "canonical form" that uniquely determines the encoded group element. (This canonical form allows no further operations, though.)
- Each encoding carries a "noise level" that increases with each operation. If the noise level grows beyond a certain threshold, no further operations are possible.

However, the GGH MLM also has an important *graded* property that allows to evaluate e partially, in a sense we will detail later. In particular this graded structure has made the GGH MLM tremendously useful: notable applications of *graded* MLMs include indistinguishability obfuscation [23], witness encryption [25], attribute-based encryption for general circuits [24], and constrained pseudorandom functions for general circuits [7]. Furthermore, graded MLMs enable a very powerful class of programmable hash functions [32], which in turn allows to implement random oracles in certain "algebraic" applications [20,33].

After GGH's MLM construction, several other (graded and approximate) MLM constructions have been proposed [15,16,28,34]. However, *all* of these constructions (including the original GGH scheme) succumb to cryptanalytic attacks [12–14,37]. In particular, currently there is no obvious way to instantiate schemes relying on multilinear maps, e.g., the schemes from [7,20,24,25,33].[1]

Graded MLMs. There is one (approximate) MLM construction of Albrecht, Farshim, Hofheinz, Larraia, and Paterson (AFHLP) [2] that does not fall victim to any of the mentioned cryptanalytic attacks on MLMs. However, this construction does not offer a *graded* MLM, and thus cannot be used to bootstrap, e.g., witness encryption. Graded MLMs are algebraic tools that can enable other algebraic tools such as multilinear Groth-Sahai proofs, or multilinear programmable hash functions. It is thus still an interesting open problem whether graded MLMs exist, and whether the results of [23] can be augmented to even show equivalence to indistinguishability obfuscation.

Our contribution. In this work, we construct graded, approximate MLMs that do not succumb to any of the known attacks. Technically, we extend the non-graded MLM construction from AFHLP [2] to a graded MLM. We prove

[1] We note, however, that the cryptographic *tasks* that the constructions from [7,25] aim to achieve can be directly achieved with indistinguishability obfuscation [1,23,42].

that the multilinear decisional Diffie–Hellman (MDDH) assumption [22] holds relative to our MLM, provided that the used ingredients are secure.

Interestingly, our MLM has two technical features that previous graded approximate MLMs do not have:

1. Our encodings do not carry any noise (although they are not unique). In particular, there is no limit on the number of operations that one can perform with our encodings.
2. The canonical forms derived from encodings allow further group operations (but no further pairings).

Our new MLM (when implemented with the indistinguishability obfuscator from [23,26]) currently forms the only plausible graded MLM, and thus the only plausible way to implement a number of MLM-based constructions [7,20,24,25,33].

Furthermore, our construction is generic and modular. In particular, we reduce the quest to develop a secure (graded) MLM to the quest for a secure indistinguishability obfuscator. This seems natural (and is standard in most areas of cryptography), but given the history of previous MLM candidates (which were based on complex algebraic or combinatorial assumptions), this is not an "understood feature" at all for MLMs.

In fact, taken together with recent constructions of indistinguishability obfuscation (iO) from multilinear maps (e.g., [3,23,35,36]), our result shows a (somewhat loose) equivalence of indistinguishability obfuscation (iO) and (graded and approximate) MLMs, in the presence of a pairing-friendly group. This equivalence is loose in the following sense. First, the assumptions on both ends of the equivalence do not match: some of these works (e.g., [23]) construct iO from MLMs which support very strong computational assumptions (much stronger than MDDH) or require asymmetric multilinear maps. On the other hand, we use iO to construct *symmetric* MLMs in which we can (at this point) only prove comparatively mild (though still useful) computational assumptions (such as MDDH). Still, there seems no inherent barrier to proving stronger computational assumptions for our construction, or to adapt our construction to asymmetric pairings, and we leave open to tighten this equivalence. Second, going through our equivalence suffers subexponential security loss. Namely, we require *probabilistic* indistinguishability obfuscation, which can be constructed from iO [11], but currently only through a sub-exponential reduction.

However, we note that such an equivalence would not be highly surprising given recent results on constructing iO from MLMs [3,35]. These works only require "one-shot" (but asymmetric) MLMs, and not even *graded* encodings as we construct them.

Related Work. Our work is closely related to [2], since the non-graded MLM there serves as a starting point for our graded MLM. We will summarize their construction in Sect. 4 and give an informal overview below.

Recently, Paneth and Sahai [39] have shown a near-equivalence of a suitable abstraction of MLMs with iO. Their result requires no computational assumptions at all, but also does not consider MLMs in our sense. In particular, they

construct an abstraction of a MLM that only admits restricted access to encodings similar to the one in [23]. Beyond the group operation and the multilinear map, efficient procedures for, e.g., uniform sampling, comparison or rerandomization of encodings, are not part of this abstraction. Conversely, our notion of a MLM, like the ones from [2,22], contains descriptions of efficient procedures for these tasks.

It would be interesting to see how the restricted MLMs of [39] can be used to instantiate the constructions from [5,8,20,33] directly, i.e., without making the detour via iO. However, since iO alone is not even known to imply one-way functions (see [29] for a discussion), this will require additional assumptions.

Pass et al. [40] give a security definition of graded MLMs that requires that whenever encodings are generically equivalent (that is, cannot be distinguished with generic operations alone), they should be computationally indistinguishable as encodings. They show that this MLMs which satisfy this strong assumption imply indistinguishability obfuscation. It is not clear, however, how to construct such strongly secure MLMs (without resorting to idealized models such as the generic group model).

1.1 The (Non-graded) Approximate Multilinear Map of AFHLP

Encodings. Since our own construction is an extension of the (non-graded) approximate MLM of [2], we first recall their work. Simplifying slightly, AFHLP encode a group element g^z (from a cyclic group $\mathbb{G}$ of order p) as

$$h \quad = \quad (g^z, c = \mathbf{Enc}((\alpha, \beta), pk), \pi),$$

where

- c is a homomorphic encryption (under some public key pk) of exponents $\alpha, \beta \in \mathbb{Z}_p$,
- π is a non-interactive zero-knowledge proof that these exponents represent z in the sense that $g^z = g^\alpha u^\beta$ for a publicly known group element u. (Hence, if we write $u = g^\omega$, we have $z = \alpha + \beta \cdot \omega$.)

Hence, AFHLP simply enhance the group element $g^z \in \mathbb{G}$ by an encrypted representation of its discrete logarithm z (and a suitable consistency proof). This added information will be instrumental in computing a multilinear map on many encodings. Note that since c and π will not be uniquely determined, there are many possible encodings of a $\mathbb{G}$-element g^z.

Addition. Encodings in the AFHLP construction can be added with an (obfuscated) public circuit **Add**. This circuit takes as input two encodings $h_1 = (g^{z_1}, c_1, \pi_1)$ and $h_2 = (g^{z_2}, c_2, \pi_2)$, and computes the new encoding $h_1 + h_2 = (g^z, c, \pi)$ as follows:

1. $g^z = g^{z_1 + z_2}$ is computed using the group operation in $\mathbb{G}$;
2. c is computed homomorphically from c_1 and c_2 (adding the encrypted exponent vectors (α_i, β_i));

3. the consistency proof π is computed using the decryption key sk as a witness to show that the resulting c indeed contains a valid representation of $z = z_1 + z_2$.

Here, only the computation of π requires secret information (namely, the decryption key sk). This secret information allows to derive a valid representation (α, β) of g^z. The most delicate part of the security proof from [2] is to argue that the obfuscated circuit knowing sk does not help in solving (a multilinear variant of) the decisional Diffie–Hellman problem.

The multilinear map. The AFHLP encodings can also be multiplied with an (obfuscated) public circuit **Mult**; this takes as input κ encodings $h_1, \ldots, h_\kappa$ with $h_i = (g^{z_i}, c_i, \pi_i)$, and outputs a single group element $g^{\prod_{i=1}^{\kappa} z_i}$. (Hence, elements from the target group $\mathbb{G}_T$ are trivially and uniquely encoded as $\mathbb{G}$-elements.) To compute $g^{\prod z_i}$ from the h_i, **Mult** first checks the validity of all proofs π_i, and then uses the decryption key sk to retrieve representations (α_i, β_i). If all π_i are verifying proofs, we may assume that $z_i = \alpha_i + \beta_i \cdot \omega$ (for $u = g^\omega$), so we can write

$$g^{\prod_{i=1}^{\kappa} z_i} = \prod_{i=0}^{\kappa} (g^{\omega^i})^{\gamma_i} \quad \text{for} \quad (\gamma_0, \ldots, \gamma_\kappa) = (\alpha_1, \beta_1) * \cdots * (\alpha_\kappa, \beta_\kappa), \tag{1}$$

where "$*$" denotes the convolution product of vectors.[2] The values g^{ω^i} (for $i \le \kappa$) are hardwired into **Mult**, so **Mult** can compute $g^{\prod z_i}$ through (1). Note that this way, **Mult** can compute a κ-linear map on encodings, but not a $(\kappa+1)$-linear map. This observation is the key to showing that the MDDH assumption holds in this setting. (Indeed, the MDDH assumption states that given $\kappa+1$ encodings $h_1, \ldots, h_{\kappa+1}$ as above, it is hard to distinguish $g^{\prod_{i=1}^{\kappa+1} z_i}$ from random.)

1.2 Our New Graded Encoding Scheme

Before proceeding any further, we briefly recall the notions of a graded multilinear map and a graded encoding scheme.

Graded maps. In a *graded* multilinear map setting, we have groups $\mathbb{G}_1, \ldots, \mathbb{G}_\kappa$, and (efficiently computable) bilinear maps $e_{i,j} : \mathbb{G}_i \times \mathbb{G}_j \longrightarrow \mathbb{G}_{i+j}$ for $i + j \le \kappa$. Hence, the $e_{i,j}$ also allow the evaluation of a multilinear map $e : \mathbb{G}_1^\kappa \longrightarrow \mathbb{G}_\kappa$ iteratively, e.g., through

$$e(g_1, \ldots, g_\kappa) := e_{1,\kappa-1}(g_1, e_{1,\kappa-2}(g_2, \cdots, e_{1,1}(g_{\kappa-1}, g_\kappa) \cdots)).$$

However, the $e_{i,j}$ also allow "partial" evaluation of e, which is the key to entirely new applications such as those in [7,23–25].

[2] Recall that the multiplication of polynomials can be implemented through the convolution product on the respective coefficient vectors. In particular, we have $\sum_{i=0}^{\kappa} \gamma_i X^i = \prod_{i=1}^{\kappa} (\alpha_i + \beta_i X)$.

Unfortunately, we do not currently know how to implement such a "clean" graded multilinear map. Instead, all known graded MLM constructions work on encodings (i.e., non-unique representations of group elements). Such a construction is usually called a graded encoding scheme (GES). Following the GES notation, we will henceforth also call an encoding of a $\mathbb{G}_\ell$-element a *level-ℓ encoding*.

In the following, we will describe the main ideas for our GES.

Encodings in our scheme. In our GES, we generalize the linear representation of exponents in AFHLP to polynomials of higher degree. Additionally, we divide encodings into levels by restricting the maximum degree of the representing polynomial in each level. More formally, level-ℓ encodings take the form

$$h \;=\; (g^z, c = \mathbf{Enc}(P, pk), \pi, \ell),$$

where

- $g^z \in \mathbb{G}$ for a cyclic group $\mathbb{G}$ (that does not depend on ℓ) of prime order p,
- $P \in \mathbb{Z}_p[X]$ is a polynomial of degree up to ℓ, represented by its coefficient vector from $\mathbb{Z}_p^{\ell+1}$,
- c is the encryption (under a fully homomorphic encryption scheme) of P,
- π is a non-interactive zero-knowledge proof of the equality $g^z = g^{P(\omega)}$, where ω is defined through public values $u_0, \ldots, u_\kappa \in \mathbb{G}$ with $u_i = g^{\omega^i}$. (Hence, $g^z = g^{P(\omega)}$ is equivalent to $g^z = \prod_i u_i^{\gamma_i}$ for $P(X) = \sum_i \gamma_i X^i$.)

The encodings of AFHLP can be viewed as level-1 encodings in our scheme (with linear polynomials P).

Adding encodings. Encodings can be added using a public (obfuscated) circuit **Add** that proceeds similarly to the AFHLP scheme. In particular, **Add** adds the g^z and c parts of the input encodings homomorphically, and derives a consistency proof π with the decryption key sk as witness.

Multiplying encodings. The pairings $e_{i,j} : \mathbb{G}_i \times \mathbb{G}_j \longrightarrow \mathbb{G}_{i+j}$ are implemented over our encodings by (obfuscated) circuits $\mathbf{Mult}_{i,j}$. Circuit $\mathbf{Mult}_{i,j}$ takes as input two encodings $h_1 = (g^{z_1}, c_1, \pi_1, i)$ and $h_2 = (g^{z_2}, c_2, \pi_2, j)$ at levels i and j, respectively. The output of $\mathbf{Mult}_{i,j}$ is a level-$(i+j)$ encoding $h = (g^z, c, \pi, i+j)$, computed as follows:[3]

- g^z is computed as $g^z = g^{(P_1 \cdot P_2)(\omega)}$, where the polynomials P_1 and P_2 are extracted from c_1 and c_2 with sk, then multiplied to form $P := P_1 \cdot P_2 \in \mathbb{Z}_p[X]$, and finally used to compute

$$g^{(P_1 \cdot P_2)(\omega)} \;=\; g^{P(\omega)} \;=\; \prod_{\ell=0}^{i+j} u_\ell^{\gamma_\ell} \quad \text{for} \quad P(X) = \sum_{\ell=0}^{i+j} \gamma_\ell X^\ell.$$

  (Since the u_ℓ are public, this value can be computed as long as $i + j \leq \kappa$.)

[3] Since $\mathbf{Mult}_{i,j}$ can be used to multiply two encodings at level i as long as $2i \leq \kappa$, our GES can be viewed as *symmetric*. We note that we do not deal with the construction of generalized GES (see [22, Appendix A] for a definition).

- c is computed homomorphically from c_1 and c_2, as an encryption of the polynomial $P_1 \cdot P_2$.
- The consistency proof π (showing that indeed $g^z = g^{P(\omega)}$ for the polynomial P encrypted in c) is computed with the decryption key sk as witness.

The key insight needed to show that the MDDH assumption holds for our GES is the same as in AFHLP's non-graded, approximate MLM. Namely, observe that any $\mathbf{Mult}_{i,j}$ can only multiply encodings if $i + j \leq \kappa$. To compute the first component g^z of any "higher-level" encoding, knowledge of g^{ω^ℓ} for $\ell > i + j$ seems to be required. Under the SDDH assumption in $\mathbb{G}$, such g^{ω^ℓ} look random, even when given $u_0, \ldots, u_\kappa$. Of course, to turn this observation into a full proof, more work is required.

Neglected details. For a useful GES, it should be possible to generate encodings with "known discrete logarithm"; that is, we would like to be able to generate encodings for an externally given (or at least known) $z \in \mathbb{Z}_p$. For this reason, the standard way to generate encodings (at any level) is to set up P as a *constant* polynomial of the form $P(X) = z \in \mathbb{Z}_p$. (That is, we "reserve space" in c for polynomials P of degree ℓ in level-ℓ encodings, but, by default, use only constant polynomials.) For this type of encoding with "low-degree P," however, our security argument above does not apply. Rather, it requires that the degree of P increases at higher levels.

Hence, the central technical piece in our MDDH security proof will be a "switching theorem" that allows to replace a low-degree P in an encoding with an *equivalent* high-degree P' (that satisfies $P'(\omega) = P(\omega)$). The proof of this switching theorem is delicate, since it must work in a setting with (obfuscated) algorithms that use the decryption key sk. (Note that free access to sk would allow the retrieval of the used polynomial P from an encoding, and hence would prevent such a switching of polynomials.)

To this end, we will use *double encryptions* c (instead of the single encryption $c = \mathbf{Enc}(P, pk)$ described above), along with a Naor–Yung-style consistency proof in π. However, this consistency proof does not show equality of encryptions, but *equivalence* of encrypted representations P, P' in the sense of $P(\omega) = P'(\omega)$. This allows to switch representations without invalidating the consistency of the double encryption. As a result, the full consistency language used for π is considerably more complicated than the one sketched before. Additionally, the proof of our switching theorem requires a special and explicit "simulation trapdoor" and Groth–Sahai-style dual-mode proof systems.

We note that similar complications arose already in AFHLP's proof, and required similar measures. The main technical difference in our setting is that our multiplication circuits $\mathbf{Mult}_{i,j}$ output *encodings* (and not just group elements as in the multilinear map of AFHLP). Hence, our $\mathbf{Mult}_{i,j}$ circuits also need to construct consistency proofs π, which requires additional secrets (as witnesses) in the description of $\mathbf{Mult}_{i,j}$ and which entails additional steps in our switching theorem. (We give more details on the technical differences with AFHLP in the main body. However, we note that, in addition to providing a *graded* encoding scheme, we also provide simplified and tighter proofs.

Fortunately, the indistinguishability obfuscator from [23] requires only a relatively weak MLM variant and hence is not affected by the above-mentioned cryptanalyses.[4]

Assumptions. In summary, our construction uses a cyclic group in which the SDDH assumption holds, a probabilistic indistinguishability obfuscation scheme [11], a perfectly correct fully homomorphic encryption (FHE), a dual-mode non-interactive zero-knowledge proof systems, and a language with hard membership. All of these assumptions are implied by pairing-friendly SDDH groups (equipped with an asymmetric pairing) and sub-exponentially secure indistinguishability obfuscation (see [31]). We stress that plausible candidates for both ingredients exist (e.g., by combining [22,23] to an indistinguishability obfuscator candidate).

Road map. We first recall some preliminaries in Sect. 2 and the GES definition in Sect. 3. Section 4 recalls the AFHLP construction. We are then ready to present our GES construction in Sect. 5, and establish our central technical tool (the "switching theorem") in Sect. 6. We prove the hardness of MDDH in Sect. 7. In the appendices, we give a technical overview of AFHLP and the full proofs of the theorems from the main body of the paper.

2 Preliminaries

Notation. We denote the security parameter by $\lambda \in \mathbb{N}$ and assume that it is implicitly given to all algorithms in the unary representation 1^λ. By an algorithm we mean a stateless Turing machine. Algorithms are randomized unless stated otherwise, and PPT as usual stands for "probabilistic polynomial-time." In this paper, by a PPT algorithm we mean an algorithm that runs in polynomial time in the security parameter (rather than the total length of its inputs). Given a randomized algorithm $\mathcal{A}$ we denote the action of running $\mathcal{A}$ on input(s) $(1^\lambda, x_1, \ldots)$ with uniform random coins r and assigning the output(s) to $(y_1, \ldots)$ by $(y_1, \ldots) \leftarrow\!\!\$\ \mathcal{A}(1^\lambda, x_1, \ldots; r)$. For a finite set X, we denote its cardinality by $|X|$ and the action of sampling a uniformly random element x from X by $x \leftarrow\!\!\$\ X$. We write $[k] := \{1, \ldots, k\}$. Vectors are written in boldface $\mathbf{x}$, and slightly abusing notation, running algorithms on vectors of elements indicates component-wise operation. Throughout the paper $\perp$ denotes a special error symbol, and $\mathrm{poly}(\cdot)$ stands for a fixed (but unspecified) polynomial. A real-valued function $\mathrm{negl}(\lambda)$ is negligible if $\mathrm{negl}(\lambda) \in \mathcal{O}(\lambda^{-\omega(1)})$. We denote the set of all negligible functions by Negl. We use bracket notation for elements in $\mathbb{G}$, i.e., writing $[z]$ and $[z']$ for two elements g^z and $g^{z'}$ in $\mathbb{G}$ and $[z] + [z']$ for their product $g^z g^{z'}$.

Circuits. A polynomial-sized deterministic circuit family $\mathcal{C} := \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a sequence of sets $\mathcal{C}_\lambda$ of $\mathrm{poly}(\lambda)$-sized deterministic circuits (for a fixed polynomial $\mathrm{poly}(\lambda)$). We assume that for all $\lambda \in \mathbb{N}$ all circuits $\mathrm{C} \in \mathcal{C}_\lambda$ share a common

[4] A recent attack on MLMs (see [37]) tackles even the weak MLM security requirements the indistinguishability obfuscator from [23] has. However, the construction of [23] (resp., its MLM building block) can be suitably enhanced to thwart this attack [26].

input domain $(\{0,1\}^\lambda)^{a(\lambda)}$, where $a(\lambda)$ is the arity of the circuit family, and an output co-domain $\{0,1\}^\lambda$. A randomized circuit family is defined similarly except that the circuits also take random coins $r \in \{0,1\}^{\mathrm{rl}(\lambda)}$, for a polynomial $\mathrm{rl}(\lambda)$ specifying the length of necessary random coins. To make the coins used by a circuit explicit (e.g., to view a randomized circuit as a deterministic one) we write $\mathrm{C}(x;r)$.

2.1 Homomorphic Public-Key Encryption

Syntax.A homomorphic public-key encryption (PKE) scheme for a deterministic circuit family $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$ of arity at most $a(\lambda)$ is a tuple of ppt algorithms $\Pi := (\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec}, \mathbf{Eval})$ such that $(\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$ is a conventional public-key encryption scheme with message space $\{0,1\}^\lambda$ and **Eval** is a *deterministic* algorithm that on input a public key pk a circuit $\mathrm{C} \in \mathcal{C}_\lambda$ and ciphertexts $c_1, \ldots, c_n$ with $n \le a(\lambda)$ outputs a ciphertext c. Without loss of generality, we assume that secret keys of a homomorphic PKE scheme are the random coins used in key generation. This will allow us to check key pairs for validity.

Correctness and compactness. For the scheme $\Pi := (\mathbf{Gen}, \mathbf{Enc}, \mathbf{Dec})$, we require *perfect* correctness as a PKE scheme; that is, for any $\lambda \in \mathbb{N}$, any $m \in \{0,1\}^\lambda$, any $(sk, pk) \leftarrow_\$ \mathbf{Gen}(1^\lambda)$, and any $c \leftarrow_\$ \mathbf{Enc}(m, pk)$ we have that $\mathbf{Dec}(c, sk) = m$. We also require the FHE scheme to be fully compact in the following sense. For any $\lambda \in \mathbb{N}$, any $m_1, \ldots, m_n \in \{0,1\}^\lambda$ with $n \le a(\lambda)$, any $\mathrm{C} \in \mathcal{C}_\lambda$, any $(sk, pk) \leftarrow_\$ \mathbf{Gen}(1^\lambda)$ and any $c_i \leftarrow_\$ \mathbf{Enc}(m_i, pk)$ we have that $\mathbf{Eval}(pk, \mathrm{C}, c_1, \ldots, c_n)$ is in the range of $\mathbf{Enc}(\mathrm{C}(m_1, \ldots, m_n), pk)$.

A *fully* homomorphic encryption (FHE) scheme is a homomorphic PKE that correctly and compactly supports any circuit family containing polynomial-sized circuits of polynomial arity (for any a priori fixed polynomial bounds on the size and arity). In our constructions, full correctness and compactness are used to ensure that the outputs of the addition and multiplications circuits can be iteratively operated on. This in particular means that our GES is "noise-free" in the sense that its correctness is not affected by the number of operations operated on encodings.

A perfectly correct FHE scheme can be constructed from probabilistic indistinguishability obfuscation (and a re-randomizable public-key encryption scheme such as ElGamal), see [11]. (We note that the FHE scheme from [11] only enjoys perfect correctness when the obfuscator and encryption scheme are also perfectly correct.)

Security. The IND-CPA security of a homomorphic PKE scheme is defined identically to a standard PKE scheme without reference to the **Dec** and **Eval** algorithms. Formally, we require that for any legitimate ppt adversary $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$,

$$\mathbf{Adv}^{\text{ind-cpa}}_{\Pi,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[\text{IND-CPA}^{\mathcal{A}}_{\Pi}(\lambda)\right] - 1 \in \text{Negl},$$

where game $\text{IND-CPA}^{\mathcal{A}}_{\Pi}(\lambda)$ is shown in Fig. 1 (left). Adversary $\mathcal{A}$ is legitimate if it outputs two messages of equal lengths.

$\underline{\text{IND-CPA}^{\mathcal{A}}_{\Pi}(\lambda)}$:	$\underline{\text{IND}^{\mathcal{A}}_{\mathbf{Obf}}(\lambda)}$:	$\underline{\text{Sel-IND}^{\mathcal{D}}_{\mathcal{A}}(\lambda)}$:
$(sk, pk) \leftarrow\!\!\$\ \mathbf{Gen}(1^\lambda)$	$(\mathrm{C}_0, \mathrm{C}_1, st) \leftarrow\!\!\$\ \mathcal{A}_1(1^\lambda)$	$(x, z) \leftarrow\!\!\$\ \mathcal{D}_1(1^\lambda)$
$(m_1, m_1, st) \leftarrow\!\!\$\ \mathcal{A}_1(pk)$	$b \leftarrow\!\!\$\ \{0,1\}$	$(\mathrm{C}_0, \mathrm{C}_1, st) \leftarrow\!\!\$\ \mathcal{A}_1(1^\lambda)$
$b \leftarrow\!\!\$\ \{0,1\}$	$\overline{\mathrm{C}} \leftarrow\!\!\$\ \mathbf{Obf}(1^\lambda, \mathrm{C}_b)$	$b \leftarrow\!\!\$\ \{0,1\};\ r \leftarrow\!\!\$\ \{0,1\}^{\mathrm{rl}(\lambda)}$
$c \leftarrow\!\!\$\ \mathbf{Enc}(m, pk)$	$b' \leftarrow\!\!\$\ \mathcal{A}_2(\overline{\mathrm{C}}, st)$	$y \leftarrow \mathrm{C}_b(x; r)$
$b' \leftarrow\!\!\$\ \mathcal{A}_2(c, st)$	Return $(b = b')$	$b' \leftarrow\!\!\$\ \mathcal{D}_2(y, \mathrm{C}_0, \mathrm{C}_1, st, z)$
Return $(b = b')$		Return $(b = b')$

Fig. 1. Left: IND-CPA security of a (homomorphic) PKE scheme. Middle: Indistinguishability security of an obfuscator. We require $\mathcal{A}_1$ to output two circuits of equal sizes. Right: Static-input (a.k.a. selective) X-IND property of $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$.

2.2 Obfuscators

Syntax and correctness. A ppt algorithm **Obf** is called an *obfuscator* for a (deterministic or randomized) circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$ if **Obf** on input the security parameter 1^λ and the description of a (deterministic or randomized) circuit $\mathrm{C} \in \mathcal{C}_\lambda$ of arity $a(\lambda)$ outputs a *deterministic* circuit $\overline{\mathrm{C}}$. For deterministic circuits, we require **Obf** to be perfectly correct in the sense the circuits C and $\overline{\mathrm{C}}$ are functionally equivalent; that is, that for all $\lambda \in \mathbb{N}$, all $\mathrm{C} \in \mathcal{C}_\lambda$, all $\overline{\mathrm{C}} \leftarrow\!\!\$\ \mathbf{Obf}(1^\lambda, \mathrm{C})$, and all $m_i \in \{0,1\}^\lambda$ for $i \in [a(\lambda)]$ we have that $\mathrm{C}(m_1, \ldots, m_{a(\lambda)}) = \overline{\mathrm{C}}(m_1, \ldots, m_{a(\lambda)})$. For randomized circuits, the authors of [11] define correctness via computational indistinguishability of the outputs of C and $\overline{\mathrm{C}}$. For our constructions we do *not* rely on this property and instead require that C and $\overline{\mathrm{C}}$ are functionally equivalent up to a change in randomness; that is, for all $\lambda \in \mathbb{N}$, all $\mathrm{C} \in \mathcal{C}_\lambda$, all $\overline{\mathrm{C}} \leftarrow\!\!\$\ \mathbf{Obf}(1^\lambda, \mathrm{C})$ and all $m_i \in \{0,1\}^\lambda$ for $i \in [a(\lambda)]$ there is an r such that $\overline{\mathrm{C}}(m_1, \ldots, m_{a(\lambda)}) = \mathrm{C}(m_1, \ldots, m_{a(\lambda)}; r)$. We note that the construction from [11] is correct in this sense as it relies on a correct indistinguishability obfuscator and a PRF to internally generate the required random coins.

Security. The security of an obfuscator **Obf** requires that for any legitimate ppt adversary $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$

$$\mathbf{Adv}^{\mathrm{ind}}_{\mathbf{Obf},\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[\text{IND}^{\mathcal{A}}_{\mathbf{Obf}}(\lambda)\right] - 1 \in \text{Negl},$$

where game IND is shown in Fig. 1 (middle). Depending on the adopted notion of legitimacy, different security notions for the obfuscator emerge; we consider the following one.

X-IND samplers [11]. Roughly speaking, the first phase of $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$ is an X-IND sampler if there is a set $\mathcal{X}$ of size at most X such that the circuits output by $\mathcal{A}$ are functionally equivalent outside $\mathcal{X}$, and furthermore within $\mathcal{X}$ the outputs of the circuits are computationally indistinguishable. Formally, let $X(\cdot)$ be a function such that $X(\lambda) \leq 2^\lambda$ for all $\lambda \in \mathbb{N}$. We call $\mathcal{A} := (\mathcal{A}_1, \mathcal{A}_2)$ an X-IND *sampler* if there are sets $\mathcal{X}_\lambda$ of size at most $X(\lambda)$ such that the following

two conditions hold: (1) For all (even unbounded) $\mathcal{D}$ the advantage function below is negligible.

$$\mathbf{Adv}^{\text{eq}}_{\mathcal{A},\mathcal{D}}(\lambda) := \Pr\left[(\mathrm{C}_0, \mathrm{C}_1, st) \leftarrow\!\!\$\ \mathcal{A}_1(1^\lambda); (x, r) \leftarrow\!\!\$\ \mathcal{D}(\mathrm{C}_0, \mathrm{C}_1, st) : \mathrm{C}_0(x; r) \neq \mathrm{C}_1(x; r) \wedge x \notin \mathcal{X}_\lambda\right]$$

(2) For all non-uniform PPT distinguishers $\mathcal{D} := (\mathcal{D}_1, \mathcal{D}_2)$ it holds that

$$X(\lambda) \cdot \mathbf{Adv}^{\text{sel-ind}}_{\mathcal{A},\mathcal{D}}(\lambda) := X(\lambda) \cdot \left(2 \Pr\left[\text{Sel-IND}^{\mathcal{D}}_{\mathcal{A}}(\lambda)\right] - 1\right) \in \text{NEGL},$$

where game $\text{Sel-IND}^{\mathcal{D}}_{\mathcal{A}}(\lambda)$ is shown in Fig. 1 (right). This game is named "static-input-IND" in [11]. and has a selective (or static) flavor since $\mathcal{D}_1$ chooses a differing-input x *before* it gets to see the challenge circuits. We call an obfuscator meeting this level of security a *probabilistic indistinguishability obfuscator* [11] and use **PIO** instead of **Obf** to emphasize this.

REMARK. We note that samplers that output two (possibly randomized) circuits $(\mathrm{C}_0, \mathrm{C}_1)$ for which the output distributions of $\mathrm{C}_0(x)$ and $\mathrm{C}_1(x)$ are identical on any input x, are Sel-IND-secure for any function $X(\lambda)$. The circuits samplers that we will use in our security proofs enjoy this property.

2.3 Dual-Mode NIZK Proof Systems

In our constructions we will be relying on special types of "dual-mode" non-interactive zero-knowledge (NIZK) proof systems. These systems have two common reference string (CRS) generation algorithms that produce indistinguishable CRSs in the "binding" and "hiding" modes. They are also perfectly complete in both modes, perfectly sound and extractable in the binding mode, and perfectly witness indistinguishable (WI) and perfectly zero knowledge (ZK) in the hiding mode. The standard prototype for such schemes are the pairing-based Groth–Sahai proofs [30], and using a generic NP reduction to the satisfiability of quadratic equations we can obtain a suitable proof system for any NP language.[5] We formalize the syntax and security of such proof systems next.

SYNTAX. A (group) setup algorithm $\mathbf{G}$ is a PPT Turing machine that on input 1^λ outputs gpk. A ternary relation $\mathbf{R}(gpk, x, w)$ is a deterministic algorithm that outputs 1 for true or 0 for false. A dual-mode extractable non-interactive zero-knowledge (NIZK) proof system Σ *for setup* $\mathbf{G}$ *and relation* $\mathbf{R}$ consists of six algorithms as follows. (1) $\mathbf{BCRS}(gpk)$ on input gpk in the support of $\mathbf{G}$ outputs a (binding) CRS crs and an extraction trapdoor td_e; (2) $\mathbf{HCRS}(gpk)$ on input gpk in the support of $\mathbf{G}$ outputs a (hiding) CRS crs and a simulation

[5] We note that extraction in Groth–Sahai proofs does not recover a witness for all types of statements. (Instead, for some types of statements, only g^{w_i} for a witness variable $w_i \in \mathbb{Z}_p$ can be recovered.) Here, however, we will only be interested in witnesses $w = (w_1, \ldots, w_n) \in \{0,1\}^n$ that are bit strings, in which case extraction always recovers w. (Extraction will recover g^{w_i} for all i, and thus all w_i too.).

trapdoor td_{zk}; (3) $\mathbf{Prove}(gpk, crs, x, w)$ on input gpk a first coordinate in the support of $\mathbf{G}$, a CRS crs, an instance x, and a witness w, outputs a proof π; (4) $\mathbf{Verify}(gpk, crs, x, \pi)$ on input gpk, crs, an instance x, and a proof π, outputs 1 for accept or 0 for reject; (5) $\mathbf{WExt}(td_e, x, \pi)$ on input an extraction trapdoor td_e, an instance x, and a proof π, outputs a witness w; and (6) $\mathbf{Sim}(td_{zk}, x)$ on input the simulation trapdoor td_{zk} and an instance x, outputs a simulated proof π.

We require the extractable dual-mode NIZK Σ for $(\mathbf{G}, \mathbf{R})$ to meet the following requirements.

CRS INDISTINGUISHABILITY. For $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, the two CRSs generated with $\mathbf{BCRS}(gpk)$ and $\mathbf{HCRS}(gpk)$ are computationally indistinguishable. Formally, we require the advantage of any PPT adversary $\mathcal{A}$ defined below to be negligible.

$$\mathbf{Adv}^{\mathrm{crs}}_{\Sigma,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[b \leftarrow\!\!\$ \{0,1\}; gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda); (crs_0, td_e) \leftarrow\!\!\$\ \mathbf{BCRS}(gpk); (crs_1, td_{zk}) \leftarrow\!\!\$\ \mathbf{HCRS}(gpk); b' \leftarrow\!\!\$\ \mathcal{A}(gpk, crs_b) : b = b'\right] - 1$$

PERFECT COMPLETENESS. For any $\lambda \in \mathbb{N}$, any $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, any $(crs, td_e) \leftarrow\!\!\$\ \mathbf{BCRS}(gpk)$, any (x, w) where it holds that $\mathbf{R}(gpk, x, w) = 1$, and any $\pi \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, x, w)$, it holds that $\mathbf{Verify}(gpk, crs, x, \pi) = 1$. We require this property to also hold for any choice of hiding CRS.

PERFECT SOUNDNESS UNDER $\mathbf{BCRS}$. For any $\lambda \in \mathbb{N}$, any $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, any CRS $(crs, td_e) \leftarrow\!\!\$\ \mathbf{BCRS}(gpk)$, any x where it holds that $\mathbf{R}(gpk, x, w) = 0$ for all $w \in \{0,1\}^*$, and any $\pi \in \{0,1\}^*$ we have that $\mathbf{Verify}(gpk, crs, x, \pi) = 0$.

PERFECT EXTRACTION UNDER $\mathbf{BCRS}$. For any $\lambda \in \mathbb{N}$, any $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, any CRS $(crs, td_e) \leftarrow\!\!\$\ \mathbf{BCRS}(gpk)$, any (x, π) with $\mathbf{Verify}(gpk, crs, x, \pi) = 1$, and any $w \leftarrow\!\!\$\ \mathbf{WExt}(td_e, x, \pi)$ we have that $\mathbf{R}(gpk, x, w) = 1$.

PERFECT WITNESS INDISTINGUISHABILITY UNDER $\mathbf{HCRS}$. For any $\lambda \in \mathbb{N}$, any $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, any $(crs, td_{zk}) \leftarrow\!\!\$\ \mathbf{HCRS}(gpk)$, and any (x, w_b) such that $\mathbf{R}(gpk, x, w_b) = 1$ for $b \in \{0,1\}$, the two distributions $\pi_b \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, x, w_b)$ are identical.

PERFECT ZERO KNOWLEDGE UNDER $\mathbf{HCRS}$. For any $\lambda \in \mathbb{N}$, any $gpk \leftarrow\!\!\$\ \mathbf{G}(1^\lambda)$, any $(crs, td_{zk}) \leftarrow\!\!\$\ \mathbf{HCRS}(gpk)$, and any (x, w) such that $\mathbf{R}(gpk, x, w) = 1$, the two distributions $\pi_0 \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, x, w)$ and $\pi_1 \leftarrow\!\!\$\ \mathbf{Sim}(td_{zk}, x)$ are identical.

2.4 Languages with Hard Membership

In our proofs of security we also rely on languages for which the membership problem is hard and whose yes-instances have unique witnesses. Formally, such a language family is defined as a tuple of four algorithms $\Lambda := (\mathbf{Gen_L}, \mathbf{YesSam_L}, \mathbf{NoSam_L}, \mathbf{R_L})$ as follows. (1) $\mathbf{Gen_L}(1^\lambda)$ is randomized and on input the security parameter outputs a language key lk; (2) $\mathbf{YesSam_L}(lk)$ is randomized and on input the language key lk outputs a yes-instance y; (3) $\mathbf{NoSam_L}(lk)$ is

randomized and on input the language key lk outputs a no-instance y; and (4) $\mathbf{R_L}(lk, y, w)$ is deterministic and on input lk, an instance y and a witness w outputs 1 for true or 0 for false.

We require $\mathbf{R_L}$ to satisfy the following correctness requirements. For all $\lambda \in \mathbb{N}$, all $lk \leftarrow\!\!\$\ \mathbf{Gen_L}(1^\lambda)$ and all $y \leftarrow\!\!\$\ \mathbf{YesSam_L}(lk)$ there is a $w \in \{0,1\}^*$ such that $\mathbf{R_L}(lk, y, w) = 1$. For a given lk, we denote the set of yes-instance by $\mathcal{L}_{lk}$. For all $\lambda \in \mathbb{N}$, all $lk \leftarrow\!\!\$\ \mathbf{Gen_L}(1^\lambda)$ and all $y \leftarrow\!\!\$\ \mathbf{NoSam_L}(lk)$ there is no $w \in \{0,1\}^*$ such that $\mathbf{R_L}(lk, y, w) = 1$. We also require $\mathbf{R_L}$ to have unique witnesses: for all $\lambda \in \mathbb{N}$, all $lk \leftarrow\!\!\$\ \mathbf{Gen_L}(1^\lambda)$, all $y \leftarrow\!\!\$\ \mathbf{YesSam_L}(lk)$ and all $w, w' \in \{0,1\}^*$ if $\mathbf{R_L}(lk, y, w) = \mathbf{R_L}(lk, y, w') = 1$ then $w = w'$.

Finally, the language is required to have a hard membership problem in the sense that for any PPT adversary $\mathcal{A}$

$$\mathbf{Adv}^{\mathrm{mem}}_{\Lambda,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[b \leftarrow\!\!\$\ \{0,1\}; lk \leftarrow\!\!\$\ \mathbf{Gen_L}(1^\lambda); y_0 \leftarrow\!\!\$\ \mathbf{NoSam_L}(lk); y_1 \leftarrow\!\!\$\ \mathbf{YesSam_L}(lk); b' \leftarrow\!\!\$\ \mathcal{A}(lk, y_b) : b = b'\right] - 1 \in \text{NEGL}.$$

Such languages can be instantiated using the DDH problem as follows. Algorithm $\mathbf{Gen_L}(1^\lambda)$ outputs the description of a prime-order group $(\mathbb{G}, g, p, 1)$ as lk. Algorithm $\mathbf{YesSam_L}(lk)$ samples a Diffie–Hellman tuple (g^a, g^b, g^{ab}), and $\mathbf{NoSam_L}(lk)$ outputs a non-Diffie–Hellman tuple (g^a, g^b, g^c) for a random $c \neq ab \pmod p$ when $b = 0$. Relation $\mathbf{R_L}$ on instance (g_1, g_2, g_3) and witness $w = a$ checks if $g_1 = g^a$ and $g_3 = g_2^a$. The hardness of membership for this language family follows from the DDH assumption.

3 Graded Encoding Schemes

We start by recalling (a slight variant of) the definition of graded encoding systems from Garg, Gentry and Halevi (GGH) [22].

κ-GRADED ENCODING SYSTEM. Let R be a (non-trivial) commutative ring and $S := \{S_i^{(a)} \subset \{0,1\}^* \ : \ a \in R, 0 \le i \le \kappa\}$ a system of sets. Then (R, S) is called a *κ-graded encoding system* if the following conditions are met.

1. For each *level* $i \in \{0, \ldots, \kappa\}$ and for any $a_1, a_2 \in R$ with $a_1 \neq a_2$ we have that $S_i^{(a_1)} \cap S_i^{(a_2)} = \emptyset$.
2. For each level $i \in \{0, \ldots, \kappa\}$, the set $\{S_i^{(a)} \ : \ a \in R\}$ is equipped with a binary operation "+" and a unary operation "−" such that for all $a_1, a_2 \in R$ and every $u_1 \in S_i^{(a_1)}, u_2 \in S_i^{(a_2)}$ it holds that
$$u_1 + u_2 \in S_i^{(a_1+a_2)} \quad \text{and} \quad -u_1 \in S_i^{(-a_1)}.$$
Here $a_1 + a_2$ and $-a_1$ denote addition and negation is R.
3. For each two levels $i, j \in \{0, \ldots, \kappa\}$ with $i + j \le \kappa$, there is a binary operation "×" such that for all $a_1, a_2 \in R$ and every $u_1 \in S_i^{(a_1)}, u_2 \in S_j^{(a_2)}$ it holds that
$$u_1 \times u_2 \in S_{i+j}^{(a_1 \cdot a_2)}.$$
Here $a_1 \cdot a_2$ denotes multiplication in R.

The difference to the GGH definition is that we do not require the operations "+" and "×" to be associative or commutative. (Indeed, our upcoming construction does not satisfy these properties.) We are not aware of any applications that require the associativity or commutativity of *encodings*. However, we stress that the operations "+" and "×" must respect the ring operations from R. For instance, while we may have $(u_1 + u_2) + u_3 \neq u_1 + (u_2 + u_3)$ for some $u_i \in S_j^{(a_i)}$, both the left-hand and the right-hand sides lie in $S_j^{(a_1+a_2+a_3)}$.

Throughout the paper, we refer to an element $a \in R$ as an *exponent* and a bit string $u \in S_i^{(a)}$ as an *encoding* of a. Further, we write $S_i := \bigcup_{a \in R} S_i^{(a)}$ for the set of all level-i encodings.

We now define graded encoding *schemes* by introducing explicit algorithms for manipulating encodings of a graded encoding system.

κ-GRADED ENCODING SCHEME. Let (R, S) be a κ-graded encoding system. A *graded encoding scheme (GES)*

$$\Gamma = (\mathbf{Setup}, \mathbf{Eq}, \mathbf{Add}, \mathbf{Mult}, \mathbf{Sam}, \mathbf{Ext})$$

associated to (R, S) consists of the following PPT algorithms.

Setup$(1^\lambda, 1^\kappa)$**:** On input the security parameter 1^λ and the (multi)linearity 1^κ, it outputs parameters of Γ (which are assumed to be provided to all other algorithms). We note that this algorithm runs in time poly(λ) as long as κ is polynomial in λ.

Eq$_i(h_1, h_2)$**:** For $i \in \{0, \ldots, \kappa\}$ and two encodings $h_1 \in S_i^{(a)}$ and $h_2 \in S_i^{(b)}$, this deterministic algorithm outputs 1 if and only if $a = b$ in R.

Add$_i(h_1, h_2)$**:** This deterministic algorithm performs the "+" operation of (R, S) in level i. For $i \in \{0, \ldots, \kappa\}$ and encodings $h_1 \in S_i^{(a_1)}$ and $h_2 \in S_i^{(a_2)}$ this algorithm outputs an encoding in $h \in S_i^{(a_1+a_2)}$.

Mult$_{i,j}(h_1, h_2)$**:** This deterministic algorithm performs the "×" operation of (R, S). For $i, j \in \{0, \ldots, \kappa\}$ with $i + j \leq \kappa$ and encodings $h_1 \in S_i^{(a_1)}$ and $h_2 \in S_j^{(a_2)}$ this algorithm outputs an encoding in $S_{i+j}^{(a_1 \cdot a_2)}$.

Sam$_i(a)$**:** For $i \in \{0, \ldots, \kappa\}$ and $a \in R$, this probabilistic algorithm samples an encoding from $S_i^{(a)}$.

Ext$_i(h)$**:** For $i \in \{0, \ldots, \kappa\}$ and input $h \in S_i$, this deterministic algorithm outputs a bit string. Algorithm **Ext**$_i$ is required to respect membership in $S_i^{(a)}$, i.e., it outputs identical strings for any two encodings $h_1, h_2 \in S_i^{(a)}$.

Our definition of a GES essentially implements the "dream version" of GESs [22], but differs in two aspects:

- GGH do not permit sampling for specific values $a \in R$. (Instead, GGH provide an algorithm to sample a random a along with its encoding.)
- GGH's zero-testing algorithm is substituted with an equality test (through **Eq**$_i$) above. Our equality test must only work for *consistent* encodings from some $S_i^{(a)}$ and $S_i^{(b)}$. In contrast, the dream version of GGH requires that the set $S_i^{(0)}$ is efficiently recognizable.

4 Approximate Multilinear Maps

We recall the approximate multilinear maps due to AFHLP [2]. The authors construct both symmetric and asymmetric multilinear maps. Their symmetric construction can be seen as a starting point for our GES.

4.1 Syntax

We start with the syntax of multilinear group (MLG) schemes [2]. Informally, a κ-MLG scheme is a restricted form of a graded encoding scheme where encodings belong to levels 0, 1 and κ only and the **Mult** algorithm takes κ encodings at level 1 and outputs an encoding at level κ. We formalize MLG schemes in terms of a GES.

Symmetric MLG schemes. A symmetric κ-linear group scheme is a κ-graded encoding scheme associated to (R, S), where (R, S) is defined similarly to a κ-graded encoding system except that $S := \{S_i^{(a)} \subset \{0,1\}^* \ : \ a \in R, i \in \{0, 1, \kappa\}\}$ and the "$\times$" operation is redefined as a κ-ary map that for any $a_1, \ldots, a_\kappa \in R$ and any $u_1 \in S_1^{(a_1)}, \ldots, u_\kappa \in S_1^{(a_\kappa)}$ satisfies

$$u_1 \times \cdots \times u_\kappa \in S_\kappa^{(a_1 \cdots a_\kappa)}.$$

The associated **Mult** algorithm on inputs $h_i \in S_1^{(a_i)}$ for $i \in [\kappa]$ outputs an encoding in $S_\kappa^{(a_1 \cdots a_\kappa)}$. Algorithms **Eq**, **Add**, **Sam** and **Ext** are defined analogously and restricted to $i \in \{0, 1, \kappa\}$ only.

4.2 Overview of AFHLP

In a nutshell, [2] works with redundant encodings of elements h of the base group $\mathbb{G}$ of the form $h = g^{x_0}(g^\omega)^{x_1}$ where g^ω comes from an SDDH instance. Vector $\mathbf{x} = (x_0, x_1)$ *represents* element h. The set S_1 consists of all strings of the form (h, c_1, c_2, π) where $h \in \mathbb{G}$, ciphertext c_1 is a homomorphic encryption under public key pk_1 of a vector $\mathbf{x}$ representing h, ciphertext c_2 is a homomorphic encryption under a second public key pk_2 of another vector $\mathbf{y}$ also representing h, and π is a NIZK proof showing consistency of the two vectors $\mathbf{x}$ and $\mathbf{y}$. Here consistency means that the plaintexts vectors $\mathbf{x}$ and $\mathbf{y}$ underlying c_1 and c_2 encode the same group element h. Note that each element of the base group $\mathbb{G}$ is multiply represented in S_1, but that equality of elements in S_1 is easy to test (via checking the equality of first components).

Addition of two elements in S_1 is carried out by an obfuscation of a circuit $\mathrm{C}_{\mathrm{Add}}[sk_1, sk_2]$, which has the two secret keys hardwired in. The circuit checks the respective proofs, adds the group elements in $\mathbb{G}$ and uses the additive homomorphic property of the encryption scheme to combine ciphertexts. It then uses witness (sk_1, sk_2) to generate a NIZK proof showing equality of encodings. Note that the new encoding is as compact as the two input encodings.

The multilinear map on inputs $(h_i, c_{i,1}, c_{i,2}, \pi_i)$ for $1 \leq i \leq \kappa$ is computed using an obfuscation of a circuit $\mathrm{C}_{\mathrm{Map}}[sk_1, \omega]$, which has sk_1 and ω hardwired in. The circuit recovers the exponents of h_i in the form $(x_{i,1} + \omega \cdot x_{i,2})$ from $c_{i,1}$ via the decryption algorithm $\mathbf{Dec}(\cdot, sk_1)$. It then uses these to compute the group element $g^{\prod_i (x_{i,1}+\omega \cdot x_{i,2})}$, which is defined to be the output of **Mult**. (The target set S_κ is therefore $\mathbb{G}$, the base group.) The κ-linearity of **Mult** follows immediately from the form of the exponent. See the full version [19] for technical details.

In the original paper, this construction is generalized to the asymmetric setting via representations of the form $g^{\langle \mathbf{x}, \boldsymbol{\omega} \rangle}$ with $\mathbf{x}, \boldsymbol{\omega} \in \mathbb{Z}_N^\ell$ for $\ell \in \{2, 3\}$ (where $\langle \mathbf{x}, \boldsymbol{\omega} \rangle$ denotes inner products modulo the base-group order). The special case $\boldsymbol{\omega} := (1, \omega)$ then gives an MLG scheme where MDDH is shown to be hard. We refer the reader to the original work [2] for the details.

5 The GES Construction

We now present our construction of a graded encoding scheme Γ according to the syntax introduced in Sect. 3. We will use the following ingredients in our construction. A similar set of building blocks were used in [2].

1. A group setup algorithm $\mathbf{Setup}_{\mathbb{G}}(1^\lambda)$ that samples (the description of) a group $\mathbb{G}$, along with a random generator g of $\mathbb{G}$ and the group order p and the identity element 1.[6] We implicitly assume efficient algorithms for checking group membership, performing the group operation, inversion, and randomly sampling group elements. We further assume a unique binary representation for every group element and a randomness extractor for this group.
2. A general-purpose probabilistic indistinguishability obfuscator **PIO** that we assume is secure against X-IND samplers.
3. A perfectly correct and IND-CPA-secure fully homomorphic PKE scheme Π with plaintext space $\mathbb{Z}_p^{\kappa+1}$.
4. An extractable dual-mode NIZK proof system Σ.
5. A language family Λ with hard membership problem and unique witnesses.

Given the above components, with formal syntax and security as defined in Sect. 2, our graded encoding scheme Γ consists of the algorithms detailed in the sections that follow. (See the introduction for an intuition.)

5.1 Setup

The **Setup** algorithm of Γ gets as input 1^λ and 1^κ. It samples parameters $pp_{\mathbb{G}} \leftarrow\!\!\$\ \mathbf{Setup}_{\mathbb{G}}(1^\lambda)$ with $pp_{\mathbb{G}} := (\mathbb{G}, g, p, 1)$, generates two encryption key pairs $(pk_j, sk_j) \leftarrow\!\!\$\ \mathbf{Gen}(1^\lambda)$ for $j = 1, 2$, and an element $\omega \leftarrow\!\!\$\ \in \mathbb{Z}_p$. We will refer to $\mathbb{G}$ as the *base group*. It sets

$$[\boldsymbol{\omega}] := ([\omega], \ldots, [\omega^\kappa]),$$

[6] It is conceivable that our security proofs also hold for non-prime p up to statistical defect terms related to randomization of elements modulo a composite number.

a vector of κ elements in the base group $\mathbb{G}$, with κ the number of desired levels It then samples $lk \leftarrow\!\!\$\ \mathbf{Gen}_{\mathbf{L}}(1^\lambda)$, and sets

$$gpk := (pp_{\mathbb{G}}, pk_1, pk_2, [\boldsymbol{\omega}], lk).$$

We define $\mathbf{G}(1^\lambda)$ to be the randomized algorithm that runs the above steps and outputs gpk. This algorithm will be used to define the NIZK proof system.

The **Setup** algorithm continues by generating a *binding* CRS $(crs', td_e) \leftarrow\!\!\$\ \mathbf{BCRS}(gpk)$, and also a *no-instance* of $\mathcal{L}_{lk}$ via $y \leftarrow\!\!\$\ \mathbf{NoSam}_{\mathbf{L}}(lk)$. It sets $crs := (crs', y)$. (The relation $\mathbf{R}$ that the NIZK should support will be defined shortly in Sect. 5.2.)

Finally, it constructs two obfuscated circuits $\overline{\mathrm{C}}_{\mathrm{Mult}}$ and $\overline{\mathrm{C}}_{\mathrm{Add}}$ of circuits $\mathrm{C}_{\mathrm{Mult}}$ and $\mathrm{C}_{\mathrm{Add}}$, which will be described in Sects. 5.3 and 5.4, respectively. **Setup** also selects a seed hk for a randomness extractor and outputs the scheme parameters

$$pp := (gpk, crs, hk, \overline{\mathrm{C}}_{\mathrm{Add}}, \overline{\mathrm{C}}_{\mathrm{Mult}}).$$

5.2 Encodings and Equality

Level-0 encodings. We treat algorithms for level-0 encodings separately in our construction as they behave somewhat differently to those from the other levels. For instance, when multiplied by other encodings, they do not result in an increase in encoding levels. The canonical choice for level-0 encodings is the ring $\mathbb{Z}_p$, which we adopt in this paper. These encodings, therefore, come with natural algorithms for generation, manipulation and testing of elements. Algorithm **Mult** when applied to inputs one of which is at level 0 corresponds to multiplication with the element in the zeroth level. The latter can in turn be implemented with a shift-and-add algorithm that employs the encoding addition **Add** of Sect. 5.3. We omit explicit mention of operations for level-0 encodings to ease notation and focus on the more interesting cases at levels 1 and above.[7]

Level-κ encodings. We set $S_\kappa := \mathbb{G}$ in our scheme and use the algorithms associated with $\mathbb{G}$ for generation, equality testing, and addition of encodings at level κ. Once again, we omit these operations from the addition circuit for clarity. The multiplication circuit can only be called on a level-κ together with a level-0 encoding, which we have already excluded. However, we still have to deal with outputs at level κ in **Mult**.

Other levels. For $0 < \ell < \kappa$ and $z \in \mathbb{Z}_p$, the encodings in $S_\ell^{(z)}$ consist of all tuples of the form

$$h := ([z], c_1, c_2, \pi, \ell),$$

[7] We mention that previous GESs used more complex level-0 encodings, and since their encodings were *noisy*, they allowed only a limited number of operations on each encoding. Hence, implementing **Mult** on level-0 inputs via shift-and-add could be too costly in their settings.

where c_1, c_2 are two ciphertexts in the range of $\mathbf{Enc}(\cdot, pk_1)$ and $\mathbf{Enc}(\cdot, pk_2)$, respectively,[8] and π is a verifying NIZK proof under crs' that:

(1) either c_1 and c_2 contain polynomials P_1 and P_2 of degree at most ℓ, such that $P_1(\omega) = P_2(\omega) = z$,
(2) or $y \in \mathcal{L}_{lk}$ (or both).

More formally, π must be a verifying proof that $(gpk, ([z], c_1, c_2, \ell))$ satisfies one relation $\mathbf{R}_1$ or $\mathbf{R}_2$ as follows.

Relation $\mathbf{R}_1$ on input gpk, an encoding $([z], c_1, c_2, \ell)$, and a witness $(P_1, P_2, r_1, r_2, sk_1, sk_2)$ accepts iff all of the following hold:

- $[z] \in \mathbb{G}$;
- both P_1 and P_2 are polynomials over $\mathbb{Z}_p$ of degree $\leq \ell$ (given by their coefficient vectors);
- both P_1 and P_2 represent z in the sense that $[z] = [P_1(\omega)]$ and $[z] = [P_2(\omega)]$;
- both c_i are encryptions of (or decrypt to) P_i in the following sense:

$$\begin{aligned} &\text{for both } i \in \{1,2\} \ : \ c_i = \mathbf{Enc}(P_i, pk_i; r_i) \\ &\qquad\qquad \vee \\ &\text{for both } i \in \{1,2\} \ : \ (pk_i, sk_i) = \mathbf{Gen}(sk_i) \wedge P_i = \mathbf{Dec}(c_i, sk_i). \end{aligned}$$

Note that there are two types of witnesses that can be used in proof generation for $\mathbf{R}_1$, namely (P_1, P_2, r_1, r_2) and (sk_1, sk_2).

Let $\mathbf{R_L}$ be the relation for the trapdoor language Λ. Relation $\mathbf{R}_2$, given gpk, an encoding, and a witness w_y, accepts iff $\mathbf{R_L}(lk, y, w_y)$ accepts. (Note that the output of $\mathbf{R}_2$ is independent of input encodings.) Hence, intuitively, $\mathbf{R}_2$ provides an explicit trapdoor to simulate consistency proofs (in case $y \in \mathcal{L}_{lk}$).

We define $\mathbf{R} := \mathbf{R}_1 \vee \mathbf{R}_2$ and assume that Σ is a proof system with respect to $(\mathbf{G}, \mathbf{R})$ with $\mathbf{G}$ as defined in Sect. 5.1.

Valid and consistent encodings. The following convention will be useful in the context of valid of encodings and the correctness of out scheme. We call an encoding h *valid* if the proof π verifies correctly under crs'. We write $\mathbf{Val}_\ell(h)$ iff h is valid and the level implicit in h matches ℓ. We call h *consistent* (with respect to gpk) if h is in the language defined by the first three conditions of relation $\mathbf{R}_1$ as well as the *first* clause of the disjunction above. (In particular, the corresponding ciphertexts c_i are possible outputs of $\mathbf{Enc}(P_i, pk_i)$; this implies that these ciphertexts behave as expected under the homomorphic evaluation algorithm **Eval**.) Note that consistency implies validity but the converse is not necessarily the case and hence a valid encoding may not lie in any S_ℓ. For example this would be the case if an "anomalous" ciphertext *decrypts* correctly to a valid representation, but does not lie in the range of **Enc**. Furthermore, validity can be publicly and efficiently checked, while this is *not* necessarily the

[8] This "honest-ciphertext-generation" condition is necessary for the (bi)linearity of our addition and multiplication algorithms. Unfortunately, this also prevents the sets $S_\ell^{(z)}$ from being efficiently recognizable.

case for consistency. We note, however, that if the encryption scheme does not allow for anomalous ciphertexts, our GES would also have efficiently recognizable encodings. We leave the construction of such FHE schemes as an open problem.

ALGORITHM **Eq**. The equality algorithm $\mathbf{Eq}_\ell$ returns 1 iff the first components of the inputs match. The correctness of this algorithm follows from the fact that the base group $\mathbb{G}$ has unique representations. (Recall from GES syntax that $\mathbf{Eq}_\ell$ is only required to work with respect to consistent encodings.)

POLYNOMIAL REPRESENTATIONS. A significant conceptual difference with the work of AFHLP is that we represent exponents in $\mathbb{Z}_p$ with polynomials instead of vectors. This generalization enables natural notion of levels corresponding to the degrees of the representing polynomials. We observe that a level-ℓ encoding h is not a *valid* level-ℓ' encoding if $\ell' \neq \ell$ as the perfectly sound proof π included in h depends on the instance and in particular on the level.

5.3 Addition

We now provide a procedure for adding two level-ℓ encodings $h = ([z], c_1, c_2, \pi, \ell)$ and $h' = ([z'], c_1', c_2', \pi', \ell)$ in S_ℓ. Conceptually, our addition circuit operates similarly to that of AFHLP. The main difference is that encodings contain polynomials and the levels. We exploit the structure of the base group as well as the homomorphic properties of the encryption scheme to "add together" the first and second components of the inputs. We then use (sk_1, sk_2) as a witness to generate a proof π'' that the new tuple is well formed. For technical reasons we check both the validity of h and h' (by checking π and π') and their consistency (using (sk_1, sk_2)).

Figure 2 details the operation of the addition circuit $\mathrm{C}_{\mathrm{Add}}$. A **PIO** of this circuit will be made public via the parameters pp. We emphasize that step 5,

CIRCUIT $\mathrm{C}_{\mathrm{Add}}[gpk, crs, sk_1, sk_2, td_e](\ell, h, h')$: // for $1 \leq \ell \leq \kappa - 1$

1. if $\neg(\mathbf{Val}_\ell(h) \wedge \mathbf{Val}_\ell(h'))$ then return $\perp$
2. parse $([z], c_1, c_2, \pi, \ell) \leftarrow h$ and $([z'], c_1', c_2', \pi', \ell) \leftarrow h'$
3. $[z''] \leftarrow [z] + [z']$; $c_1'' \leftarrow c_1 + c_1'$; $c_2'' \leftarrow c_2 + c_2'$
4. $P_1 \leftarrow \mathbf{Dec}(c_1, sk_1)$; $P_2 \leftarrow \mathbf{Dec}(c_2, sk_2)$
   $P_1' \leftarrow \mathbf{Dec}(c_1', sk_1)$; $P_2' \leftarrow \mathbf{Dec}(c_2', sk_2)$
5. if $[z] \neq [P_1(\omega)] \vee [z] \neq [P_2(\omega)] \vee [z'] \neq [P_1'(\omega)] \vee [z'] \neq [P_2'(\omega)]$ then
   5.1. $w_y' \leftarrow\!\!\$\ \mathbf{WExt}(td_e, ([z], c_1, c_2), \pi)$
   5.2. if $\neg \mathbf{R}_2(gpk, ([z], c_1, c_2, \ell), w_y')$ then return $\perp$
   5.3. $\pi'' \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, ([z''], c_1'', c_2''), w_y')$
6. else $\pi'' \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, ([z''], c_1'', c_2''), (sk_1, sk_2))$
7. return $([z''], c_1'', c_2'', \pi'', \ell)$

Fig. 2. The probabilistic circuit used to add encodings for levels $1 \leq \ell \leq \kappa - 1$. The checks at 5 are never passed in an honest execution of the protocol. We emphasize that the test in step 5 is implemented using the values $[\omega^i]$. The random coins needed for randomized operations are internally generated after obfuscating with **PIO**.

that is, the explicit consistency check, is never reached under a binding crs' (due to the perfect soundness of the proof system), but they may be reached with a hiding crs' later in the security analysis. Let us expand on this.

In the analysis, we need to specify how $\mathrm{C}_{\mathrm{Add}}$ behaves if it encounters valid inputs (in the sense the proofs pass NIZK verification), but nevertheless are inconsistent in the sense that at least one of encodings does not decrypt to a valid representation. Let us call such inputs *bad.*

With the knowledge of secret keys, such bad inputs can be recognized, and the natural choice would be to define $\mathrm{C}_{\mathrm{Add}}$ to abort when this is the case. With this choice, however, we run into the following problem. During the security proof we will set the addition circuit to answer all valid inputs (including bad ones) with simulated proofs. On the other hand, the original addition circuit rejects such inputs. (Furthermore, it cannot even simulate proofs for wrong statements, and hence cannot answer bad inputs with valid-looking proofs.)

On a high level, we would like to modify how $\mathrm{C}_{\mathrm{Add}}$ reacts on bad inputs so that it uses a NIZK simulation trapdoor on bad inputs. The difficulty with this strategy is that no such simulation trapdoor exists when the NIZK CRS is binding. Hence, we create our own NIZK trapdoor through an extra "OR branch" in the proved statement (akin to the Feige–Lapidot–Shamir transform). This gives us a little more flexibility in defining and using that trapdoor.

More specifically, recall that our CRS is of the form $crs = (crs', y)$ where crs' is a binding CRS for the dual-mode NIZK proof system, and y is a no-instance of $\mathcal{L}_{lk}$. However our actual means to fake proofs will be to switch y to a yes-instance and use a witness w_y to produce proofs. Specifically, in the security proof, we will eventually let $\mathrm{C}_{\mathrm{Add}}$ use a simulation trapdoor w_y (instead of a simulation trapdoor for the NIZK). The benefit of this is that $\mathrm{C}_{\mathrm{Add}}$ will know an extraction trapdoor td'_e (that of course only exists if the CRS crs' is in the binding mode) which it can use to extract a witness from a given proof π. Thus, whenever $\mathrm{C}_{\mathrm{Add}}$ encounters a bad input, it can extract a witness w'_y, which *must* at that point be a simulation trapdoor w_y. This simulation trapdoor w_y can then immediately be used to produce a fake proof π'' even upon bad inputs. In other words, $\mathrm{C}_{\mathrm{Add}}$ knows no simulation trapdoor a priori, but it can extract one from any simulated proof for a false statement.

The $\mathbf{Add}_\ell$ algorithm simply runs the obfuscated circuit on the input encodings and ℓ. The correctness of this algorithm follows from that of Π, the completeness of Σ and the correctness, in our sense, of the (probabilistic) obfuscator **PIO**. Note that FHE correctness is only guaranteed to hold with respect to ciphertexts that are in the range of encryption or evaluation (and not necessarily for anomalous ones that decrypt correctly). This, in particular, means that we cannot enlarge the set of encodings to contain all valid ones (as opposed to just consistent ones) to get efficient decidability of encoding sets as correctness can no longer be established. (See also remark on validity on page 18.) Note that full compactness ensures that the ciphertexts output by $\mathbf{Add}_\ell$ are in the range of encryption, and hence they can be further operated on with **Eval**.

5.4 Multiplication

Given two encodings $h = ([z], c_1, c_2, \pi, \ell)$ and $h' = ([z'], c'_1, c'_2, \pi', \ell')$ at levels ℓ and ℓ' respectively, the multiplication algorithms operates analogously to addition as follows. The corresponding circuit $\mathrm{C}_{\mathrm{Mult}}$ has both decryption keys and now also $\omega \in \mathbb{Z}_p$ hardwired in. After validity checks and decrypting the input ciphertexts, it performs the multiplication of the polynomials encrypted under c_i and c'_i homomorphically using a convolution operation on the coefficient vectors. However, it cannot obviously compute the element $[zz']$ in the base group $\mathbb{G}$. Suppose c_1 and c'_1 encrypt polynomials P and P' of degrees at most ℓ and ℓ' respectively and such that $[z] = [P(\omega)]$ and $[z'] = [P'(\omega)]$. The multiplication circuit uses the explicit knowledge of ω and polynomials P and P' to compute $[zz'] = [(P * P')(\omega)]$.[9] Circuit $\mathrm{C}_{\mathrm{Mult}}$ is shown in Fig. 3. Note that similarly to addition, step 6 performs explicit checks of consistency of encodings that will only be used in the analysis under a hiding crs'.

The correctness of these maps follows from the correctness of Π and **PIO**, and the completeness of Σ.

Enabling graded multiplication. The main difference between our circuit $\mathrm{C}_{\mathrm{Mult}}$ and that of [2] is that here we need to output auxiliary information (c_1, c_2, π) for multiplied encodings at output levels below κ. This information allows the multiplication algorithm to operate in a graded fashion as any output encoding by $\mathrm{C}_{\mathrm{Mult}}$ can be fed back into $\mathrm{C}_{\mathrm{Mult}}$ as long as it lies at a level

Circuit $\mathrm{C}_{\mathrm{Mult}}[gpk, crs, \omega, sk_1, sk_2, td_e](\ell, \ell', h, h')$: // for $1 \le \ell, \ell' \le \kappa - 1$

1. if $\neg(\mathbf{Val}_\ell(h) \wedge \mathbf{Val}_{\ell'}(h')) \vee \ell + \ell' > \kappa$ then return $\perp$
2. parse $([z], c_1, c_2, \pi, \ell) \leftarrow h$ and $([z'], c'_1, c'_2, \pi', \ell') \leftarrow h'$
3. $c''_1 \leftarrow c_1 * c'_1$; $c''_2 \leftarrow c_2 * c'_2$
4. $P_1 \leftarrow \mathbf{Dec}(c_1, sk_1)$; $P_2 \leftarrow \mathbf{Dec}(c_2, sk_2)$
   $P'_1 \leftarrow \mathbf{Dec}(c'_1, sk_1)$; $P'_2 \leftarrow \mathbf{Dec}(c'_2, sk_2)$
5. $z'' \leftarrow (P_1 * P'_1)(\omega)$
6. if $[z] \neq [P_1(\omega)] \vee [z] \neq [P_2(\omega)] \vee [z'] \neq [P'_1(\omega)] \vee [z'] \neq [P'_2(\omega)]$ then
   6.1. $w'_y \leftarrow_\$ \mathbf{WExt}(td_e, ([z], c_1, c_2), \pi)$
   6.2. if $\neg\mathbf{R}_2(gpk, ([z], c_1, c_2), w'_y)$ then return $\perp$
   6.3. $\pi'' \leftarrow_\$ \mathbf{Prove}(gpk, crs, ([z'']), c''_1, c''_2), w'_y)$
7. else $\pi'' \leftarrow_\$ \mathbf{Prove}(gpk, crs, ([z'']), c''_1, c''_2), (sk_1, sk_2))$
8. If $(\ell + \ell' = \kappa)$ then return $[z'']$ else return $([z''], c''_1, c''_2, \pi'', \ell + \ell')$

Fig. 3. Circuit used for multiplying encodings for levels $1 \le \ell, \ell' \le \kappa - 1$. Step 6 is never reached in an honest execution of the protocol with a binding crs. The random coins needed for randomized operations are internally generated after obfuscating with **PIO**.

[9] Observe that with the explicit knowledge of $P * P'$ and the powers $([\omega^i])_{1 \le i \le \kappa}$ it is also possible to compute $[zz']$ as long as $P * P'$ is of degree $\le \kappa$; this will be exploited in the security analysis in Sect. 7.

$\ell < \kappa$.[10] In order to enable $\mathrm{C}_{\mathrm{Mult}}$ to generate this auxiliary information, we use an encryption scheme that is also homomorphic with respect to multiplication in the plaintext ring. In contrast, AFHLP only rely on an additively homomorphic encryption scheme.

5.5 Sampling

Given polynomials P_1 and P_2 of degree at most ℓ and satisfying $P_1(\omega) = P_2(\omega) = z$ we can generate an encoding from $S_\ell^{(z)}$ by computing

$$\begin{aligned} h \leftarrow \big([z], c_1 = \mathbf{Enc}(P_1, pk_1; r_1), c_2 = \mathbf{Enc}(P_2, pk_2; r_2), \\ \pi = \mathbf{Prove}(gpk, crs, ([z]_i, c_1, c_2, \ell), (P_1, P_2, r_1, r_2); r), \ell\big). \end{aligned} \tag{2}$$

Hence, our sampling algorithm $\mathbf{Sam}_\ell(z)$ sets $P_1(X) = P_2(X) = z \in \mathbb{Z}_p$ and computes an encoding through (2). We call these the *canonical* encodings of z, independently of ℓ. We note that this procedure is that in [2] adapted to the generalized notion of polynomial representations.

5.6 Extraction

Since at each level ℓ the first component $[z]$ is unique for each set $S_\ell^{(z)}$, we may extract a uniform string from $h = ([z], c_1, c_2, \pi, \ell)$ for a uniform z by applying a randomness extractor seeded with hk to $[z]$.

6 Indistinguishability of Encodings

We show that a key property used by AFHLP in the analysis of their multilinear map [2, Theorem 5.3] is also exhibited by our graded scheme. Roughly speaking, this property states that for any given level ℓ, any two valid encodings of the same $\mathbb{Z}_p$-element are computationally indistinguishable. This claim is formalized via the κ-Switch game shown in Fig. 4. Note that in this game, we allow the adversary to not only choose the representation polynomials, but also let him see part of the private information not available through the public parameters, namely the exponent ω.

Theorem 1 (Encoding switch). *Let Γ be the GES constructed in Sect. 5 with respect to an X-*IND*-secure probabilistic obfuscator* **PIO**, *an* IND-CPA*-secure encryption scheme Π, a dual-mode NIZK proof system Σ, and a language family Λ. Then, encodings of the same ring element $z \in \mathbb{Z}_p$ are indistinguishable at all levels. More precisely, for any legitimate* PPT *adversary $\mathcal{A}$ there are* PPT *adversaries $\mathcal{B}_1$, $\mathcal{B}_2$, $\mathcal{B}_3$ and $\mathcal{B}_4$ of essentially the same complexity as $\mathcal{A}$ such that for all $\lambda \in \mathbb{N}$*

$$\mathbf{Adv}^{\kappa\text{-switch}}_{\Gamma,\mathcal{A}}(\lambda) \leq 3 \cdot \left(\mathbf{Adv}^{\mathrm{mem}}_{\Lambda,\mathcal{B}_1}(\lambda) + 6 \cdot \mathbf{Adv}^{\mathrm{ind}}_{\mathbf{PIO},\mathcal{B}_2}(\lambda) + \mathbf{Adv}^{\mathrm{crs}}_{\Sigma,\mathcal{B}_3}(\lambda)\right) + 2 \cdot \mathbf{Adv}^{\text{ind-cpa}}_{\Pi,\mathcal{B}_4}(\lambda).$$

[10] Recall that encodings at level κ can only be multiplied with level-0 encodings, i.e., with elements in $\mathbb{Z}_p$.

κ-Switch$^{\mathcal{A}}_{\Gamma}(\lambda)$:
$(pp;\omega) \leftarrow\!\!\$\ \mathbf{Setup}(1^\lambda, 1^\kappa)$ // ω generated within **Setup**
$((P_{0,1}, P_{0,2}), (P_{1,1}, P_{1,2}), \ell, st) \leftarrow\!\!\$\ \mathcal{A}_1(pp, \omega)$
$b \leftarrow\!\!\$\ \{0,1\};\ r_1, r_2 \leftarrow\!\!\$\ \{0,1\}^{\mathrm{rl}(\lambda)}$
$c_1 \leftarrow \mathbf{Enc}(P_{b,1}, pk_1; r_1);\ c_2 \leftarrow \mathbf{Enc}(P_{b,2}, pk_2; r_2)$
$\pi \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, ([P_{b,1}(\omega)], c_1, c_2, \ell), (P_{b,1}, P_{b,2}, r_1, r_2))$
$h_b \leftarrow ([P_{b,1}(\omega)], c_1, c_2, \pi, \ell)$
$b' \leftarrow\!\!\$\ \mathcal{A}_2(h_b, st)$
Return $(b = b')$

Fig. 4. Game formalizing the indistinguishability of encodings. (This game is specific to our construction Γ from Sect. 5.) An adversary is legitimate if it outputs polynomials such that $P_{0,1}(\omega) = P_{0,2}(\omega) = P_{1,1}(\omega) = P_{1,2}(\omega)$ of degree at most ℓ. We note that $\mathcal{A}$ gets explicit access to secret exponent ω generated at setup. Here $\mathrm{rl}(\lambda)$ is a polynomial indicating the length of the random coins used by the encryption algorithm.

The proof of this result follows largely that in [2] and we include it in the full version [19] of this paper. The main difference is that we have to deal with obfuscations of the new multiplication circuit.

Proof (Outline). We proceed via a sequence of 5 games, starting with κ-Switch and ending in a game where the challenge encoding is independent of the bit b. Figure 5 shows the steps used in the proof of the theorem. We use helper Lemma 1 for changing the addition and multiplication circuits to "forget" (one or both) the secret keys and the extraction trapdoor. We now justify each of these steps in more detail below. See the full version [19] of this paper for a full proof.

Gm.	crs'	y	$\mathrm{C}_{\mathrm{Add}}$ knows	$\mathrm{C}_{\mathrm{Mult}}$ knows	c_1 enc.	c_2 enc.	**Remark**
0	bind.	$\notin \mathcal{L}_{lk}$	sk_1, sk_2, td_e	$\underline{sk_1}, sk_2, td_e$	$P_{b,1}$	$P_{b,2}$	
1	hid.	$\in \mathcal{L}_{lk}$	w_y	$\underline{sk_1}, w_y$	$P_{b,1}$	$P_{b,2}$	Lemma 1 ($i = 1$)
2	hid.	$\in \mathcal{L}_{lk}$	w_y	$\underline{sk_1}, w_y$	$P_{b,1}$	$P_{1,2}$	IND-CPA wrt. pk_2
3	bind.	$\notin \mathcal{L}_{lk}$	sk_1, sk_2, td_e	$\underline{sk_1}, sk_2, td_e$	$P_{b,1}$	$P_{1,2}$	Lemma 1 (reverse, $i = 1$)
4	hid.	$\in \mathcal{L}_{lk}$	w_y	$\underline{sk_2}, w_y$	$P_{b,1}$	$P_{1,2}$	Lemma 1 ($i = 2$)
5	hid.	$\in \mathcal{L}_{lk}$	w_y	$\underline{sk_2}, w_y$	$P_{1,1}$	$P_{1,2}$	IND-CPA wrt. pk_1
							Encoding indep. of b

Fig. 5. Outline of the proof steps of Theorem 1. The underlined secret key in the "$\mathrm{C}_{\mathrm{Mult}}$ knows" column indicates the key that is used in decryption to construct $[z'']$. For instance, in Game_0, key sk_1 is used to obtain P_1 and P'_1, which are then used to compute $[z''] = [(P_1 * P'_1)(\omega)]$ within $\mathrm{C}_{\mathrm{Mult}}$.

Game_0**:** This is the κ-Switch game with a binding crs' and $y \notin \mathcal{L}_{lk}$. The addition and multiplication circuits are defined in Figs. 2 and 3, respectively.

Game_1**:** We change the public parameters so that they include a *hiding* crs', a yes instance y via $\mathbf{YesSam}_{\mathbf{L}}(lk)$ and obfuscations of circuits $\widehat{\mathrm{C}}_{\mathrm{Add}}$ and $\widehat{\mathrm{C}}^{(1)}_{\mathrm{Mult}}$ (see Fig. 6). Thus, the second circuit uses sk_1 to decrypt the first ciphertexts given as inputs. Observe that these circuits use the witness w_y to $y \in \mathcal{L}_{lk}$ to produce the output proofs π'', and therefore the *simultaneous* knowledge of decryption keys sk_1, sk_2 is no longer needed. The difference with the previous game can be bounded by our helper Lemma 1 with $i = 1$, where we rely on PIO security, CRS indistinguishability, and the membership problem.

Game_2**:** This game generates the second challenge ciphertext c_2 by encrypting polynomial $P_{1,2}$ *even when* $b = 0$. We bound this transition via the IND-CPA security of Π with respect to pk_2. The reduction will choose a first decryption key sk_1 and a witness w_y so as to be able to construct $\widehat{\mathrm{C}}^{(1)}_{\mathrm{Mult}}$. It will also generate a NIZK simulation trapdoor td_{zk} (recall the CRS is in the hiding mode) to construct simulated proofs π for the (inconsistent) challenge encoding h_b. Note that the perfect ZK property guarantees that these proofs are identically distributed to the real ones in Game_1.

Game_3**:** The public parameters are changed back to include a binding crs', a no-instance $y \notin \mathcal{L}_{lk}$ and a (PIO) obfuscation of the original circuits $\mathrm{C}_{\mathrm{Add}}$, $\mathrm{C}_{\mathrm{Mult}}$ with both decryption keys hardwired. The difference with the previous game is bounded again via Lemma 1 (in the reverse direction and with $i = 1$).

Game_4**:** This transitions is defined analogously to that introduced in Game_1 except that this time we invoke Lemma 1 with $i = 2$ and switch to circuits $\widehat{\mathrm{C}}_{\mathrm{Add}}$ and $\widehat{\mathrm{C}}^{(2)}_{\mathrm{Mult}}$. Observe that knowledge of sk_1 is no longer needed.

Game_5**:** This transitions is defined analogously to that introduced in Game_2. The only difference is that this game generates the *first* challenge ciphertext c_1 by encrypting $P_{1,1}$ even when $b = 0$.

Finally, note that the challenge encoding in Game_5 is independent of the random bit b and the advantage of any (even unbounded) adversary $\mathcal{A}$ is 0.

In the proof of Theorem 1, we need the next Lemma for changing the addition and multiplication circuits to "forget" (one or both) the secret keys and the extraction trapdoor. The proof can be found in the full version [19] of this paper.

Lemma 1 (Forgetting secret keys). *Let* Γ *be the GES from Sect. 5 with respect to an* X-IND*-secure probabilistic obfuscator* **PIO***, an* IND-CPA*-secure encryption scheme* Π*, a dual-mode NIZK proof system* Σ*, and a language family* Λ*. For* $i = 1, 2$*, consider the modified parameter generation algorithm* $\mathbf{Setup}^{(i)}$ *that samples a* yes-instance $y \in \mathcal{L}_{lk}$ *and outputs obfuscations of the circuits* $\widehat{\mathrm{C}}_{\mathrm{Add}}$ *and* $\widehat{\mathrm{C}}^{(i)}_{\mathrm{Mult}}$ *shown in Fig. 6. Let*

$$\mathbf{Adv}^{\kappa\text{-forget}}_{\Gamma,i,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[pp_0 \leftarrow_{\$} \mathbf{Setup}(1^\lambda, 1^\kappa);\ pp_1 \leftarrow_{\$} \mathbf{Setup}^{(i)}(1^\lambda, 1^\kappa);\right.$$
$$\left. b \leftarrow_{\$} \{0,1\}; b' \leftarrow_{\$} \mathcal{A}(pp_b)\ :\ b = b'\right] - 1.$$

CIRCUIT $\widehat{\mathrm{C}}_{\mathrm{Add}}[gpk, crs, w_y](\ell, h, h')$:
1. if $\neg(\mathbf{Val}_\ell(h) \wedge \mathbf{Val}_\ell(h'))$ then return $\perp$
2. parse $([z], c_1, c_2, \pi, \ell) \leftarrow h$, and $([z'], c_1', c_2', \pi', \ell) \leftarrow h'$
3. $[z''] \leftarrow [z] + [z']$; $c_1'' \leftarrow c_1 + c_1'$; $c_2'' \leftarrow c_2 + c_2'$
4. // omitted: depends on sk_1 and sk_2
5. $\pi'' \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, ([z''], c_1'', c_2'', \ell), w_y)$
6. // omitted: depends on sk_1 and sk_2
7. return $([z''], c_1'', c_2'', \pi'', \ell)$

CIRCUIT $\widehat{\mathrm{C}}^{(i)}_{\mathrm{Mult}}[gpk, crs, \omega, sk_i, w_y](\ell, \ell', h, h')$:
1. if $\neg(\mathbf{Val}_\ell(h) \wedge \mathbf{Val}_{\ell'}(h')) \vee \ell + \ell' > \kappa$ then return $\perp$
2. parse $([z], c_1, c_2, \pi, \ell) \leftarrow h$ and $([z'], c_1', c_2', \pi', \ell') \leftarrow h'$
3. $c_1'' \leftarrow c_1 \cdot c_1'$; $c_2'' \leftarrow c_2 \cdot c_2'$
4. $P_i \leftarrow \mathbf{Dec}(c_i, sk_i)$; $P_i' \leftarrow \mathbf{Dec}(c_i', sk_i)$ // depends on sk_i only
5. $z'' \leftarrow (P_i * P_i')(\omega)$
6. $\pi'' \leftarrow\!\!\$\ \mathbf{Prove}(gpk, crs, ([z''], c_1'', c_2'', \ell + \ell'), w_y)$
7. // omitted: depends on sk_1 and sk_2
8. If $(\ell + \ell' = \kappa)$ then return $[z'']$ else return $([z''], c_1'', c_2'', \pi'', \ell + \ell')$

Fig. 6. Top: Circuit $\widehat{\mathrm{C}}_{\mathrm{Add}}$ where witness w_y to $y \in \mathcal{L}_{lk}$ is used to produce π''. Note that the secret keys (sk_1, sk_2) or the extraction trapdoor td_e are no longer used by this circuit. Bottom: Circuits $\widehat{\mathrm{C}}^{(i)}_{\mathrm{Mult}}$ were only one key sk_i is used to decrypt P_i and P_i' and witness w_y to $y \in \mathcal{L}_{lk}$ is used to produce π''. The secret key sk_{3-i} and the extraction trapdoor td_e are not used by this circuit.

Then, for any $i \in \{1, 2\}$ *and any* PPT *adversary* $\mathcal{A}$ *there are* PPT *adversaries* $\mathcal{B}_1, \mathcal{B}_2$ *and* $\mathcal{B}_3$ *of essentially the same complexity as* $\mathcal{A}$ *such that for all* $\lambda \in \mathbb{N}$

$$\mathbf{Adv}^{\kappa\text{-forget}}_{\Gamma,i,\mathcal{A}}(\lambda) \leq \mathbf{Adv}^{\mathrm{mem}}_{\Lambda,\mathcal{B}_1}(\lambda) + 6 \cdot \mathbf{Adv}^{\mathrm{ind}}_{\mathbf{PIO},\mathcal{B}_2}(\lambda) + \mathbf{Adv}^{\mathrm{crs}}_{\Sigma,\mathcal{B}_3}(\lambda).$$

7 Hardness of MDDH

We are now ready to show that MDDH is hard for our GES. We improve [2] by providing a simpler and tighter proof of security. One corollary of our result is that there are no "zeroizing" attacks on our scheme as such attacks immediately lead to the break of MDDH [12,13,22]. We start by providing formal definition of MDDH as well as the strong DDH problem whose hardness we assume in our analyses.

THE q-SDDH PROBLEM [4,43]. For $q \in \mathbb{N}$ we say that the q-SDDH problem is hard for a group $\mathbb{G}$ if

$$\mathbf{Adv}^{q\text{-sddh}}_{\mathbb{G},\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[q\text{-SDDH}^{\mathcal{A}}_{\mathbb{G}}(\lambda)\right] - 1 \in \mathrm{NEGL},$$

where game q-SDDH$^{\mathcal{A}}_{\mathbb{G}}(\lambda)$ is shown in Fig. 7 (left). We note that this assumption can only hold in *asymmetric* pairing-friendly groups. (With such asymmetric pairings, we could then implement, e.g., the dual-mode NIZK proof system

q-SDDH$^{\mathcal{A}}_{\mathbb{G}}(\lambda)$	κ-MDDH$^{\mathcal{A}}_{\Gamma}(\lambda)$
$pp_{\mathbb{G}} \leftarrow\!\!\$\ \mathbf{Setup}_{\mathbb{G}}(1^\lambda)$	$pp \leftarrow\!\!\$\ \mathbf{Setup}(1^\lambda, 1^\kappa)$
$b \leftarrow\!\!\$\ \{0,1\}$	$b \leftarrow\!\!\$\ \{0,1\}$
$\omega, \tau_0 \leftarrow\!\!\$\ \mathbb{Z}_p$	$a_1, \ldots, a_{\kappa+1}, z \leftarrow\!\!\$\ \mathbb{Z}_p$
$\tau_1 \leftarrow \omega^{q+1} \pmod p$	$h_i \leftarrow\!\!\$\ \mathbf{Sam}_1(a_i)$
$b' \leftarrow\!\!\$\ \mathcal{A}(pp_{\mathbb{G}}, \{[\omega^i]\}_{i=1}^{q}, [\tau_b])$	$h_0^* \leftarrow\!\!\$\ \mathbf{Sam}_\kappa(z)$
Return $(b = b')$	$h_1^* \leftarrow \mathbf{Mult}(h_1, \ldots, h_\kappa)^{a_{\kappa+1}}$
	$b' \leftarrow\!\!\$\ \mathcal{A}(pp, \{h_i\}_{i=1}^{\kappa+1}, h_b^*)$
	Return $(b = b')$

Fig. 7. Left: The SDDH problem. Here $p = p(\lambda)$ denotes the group order implicit in pp. Right: The MDDH problem. The sampler algorithms output canonical encodings. The κ-ary algorithm **Mult** is defined by applying the 2-ary algorithm **Mult** of the scheme iteratively to inputs.

from [30].) It is not too difficult to show via re-randomization of the group generator that hardness of q-SDDH implies that of $(q-1)$-SDDH. We use this fact to simplify our theorem statement below.

The κ-MDDH problem [6,22]. For $\kappa \in \mathbb{N}$ we say that the κ-MDDH problem is hard for a GES Γ if

$$\mathbf{Adv}^{\kappa\text{-mddh}}_{\Gamma,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[\kappa\text{-MDDH}^{\mathcal{A}}_{\Gamma}(\lambda)\right] - 1 \in \text{Negl},$$

where game κ-MDDH$^{\mathcal{A}}_{\Gamma}(\lambda)$ is shown in Fig. 7 (middle).

The $(\kappa, m, n, r_0, r_1, l)$-RANK problem [18]. For $\kappa, m, n, r_0, r_1 \in \mathbb{N}$ and a level function $l : [m] \times [n] \longrightarrow [\kappa]$, we say that the $(\kappa, m, n, r_0, r_1, l)$-RANK problem is hard for a GES Γ if

$$\mathbf{Adv}^{(\kappa,m,n,r_0,r_1,l)\text{-rank}}_{\Gamma,\mathcal{A}}(\lambda) := 2 \cdot \Pr\left[(\kappa, m, n, r_0, r_1, l)\text{-RANK}^{\mathcal{A}}_{\Gamma}(\lambda)\right] - 1 \in \text{Negl},$$

where game $(\kappa, m, n, r_0, r_1, l)$-RANK$^{\mathcal{A}}_{\Gamma}(\lambda)$ is shown in Fig. 7 (right).

7.1 Hardness of MDDH

Recall that the GES of Sect. 5 represents an element $z \in \mathbb{Z}_p$ at level ℓ with polynomials P_1 and P_2 of degree at most ℓ such that $P_j(\omega) = z$.

Theorem 1 (κ-SDDH $\Longrightarrow$ κ-MDDH)**.** *Let Γ be the GES constructed in Sect. 5 with respect to a base group $\mathbb{G}$ and an X-IND-secure probabilistic obfuscator* **PIO**.

Then, assuming the κ-SDDH assumption (see Fig. 7) holds in $\mathbb{G}$, and using our switching lemma, the κ-MDDH assumption holds in Γ.

More specifically, for any $\kappa \in \mathbb{N}$ and any PPT *adversary $\mathcal{A}$ there are* PPT *adversaries $\mathcal{B}_1$, $\mathcal{B}_2$ and $\mathcal{B}_3$ of essentially the same complexity as $\mathcal{A}$ such that for all $\lambda \in \mathbb{N}$*

$$\mathbf{Adv}^{\kappa\text{-mddh}}_{\Gamma,\mathcal{A}}(\lambda) \leq (\kappa+1) \cdot \mathbf{Adv}^{\kappa\text{-switch}}_{\Gamma,\mathcal{B}_1}(\lambda) + \mathbf{Adv}^{\text{ind}}_{\mathbf{PIO},\mathcal{B}_2}(\lambda) + \mathbf{Adv}^{\kappa\text{-sddh}}_{\mathbb{G},\mathcal{B}_3}(\lambda).$$

Proof (Outline). We provide a simpler proof compared to that of [2, Theorem 6.2] at the expense of relying on the slightly stronger κ-SDDH (instead of the $(\kappa-1)$-SDDH) problem. At a high level, our reduction has two steps: (1) Switch *all* encodings from polynomials of degree 0 to those of degree 1; and (2) Randomize the κ-MDDH challenge using the κ-SDDH instance. The key difference with the proof of [2, Theorem 6.2] is that we no longer need to carry out a two-step process to randomize the exponent of the MDDH challenge. In particular, we do not change the implementation of the multiplication circuit according to a κ-SDDH challenge. We outline the proof along a sequence of $\kappa+5$ games here and leave the full details to the full version [19].

Game_0**:** This is the κ-MDDH problem (Fig. 7, middle). We use $P_{i,1}$ and $P_{i,2}$ to denote the canonical degree-zero representation polynomials of a_i as generated by the sampler $\mathbf{Sam}_1(a_i)$.

Game_1–$\text{Game}_{\kappa+1}$**:** In these games we gradually switch the polynomials representations for level-1 encodings h_i for $1 \le i \le \kappa+1$ so that they take the form

$$P_{i,1}(X) = P_{i,2}(X) = X + a_i - \omega.$$

These polynomials are still valid and their degrees are *exactly* 1. Hence when multiplied together, the resulting polynomial will be of degree $s(\kappa+1)$. Each of these hops can be bounded via the κ-Switch game via Theorem 1.

$\text{Game}_{\kappa+2}$**:** This game only introduces a conceptual change: a_i for $1 \le i \le \kappa+1$ are generated as $a_i + \omega$. The distributions of these values are still uniform and the exponent of the MDDH challenge when $b=1$ is now

$$z_1 = \prod_{i=1}^{\kappa+1} (a_i + \omega),$$

which is a polynomial in ω of degree κ.

$\text{Game}_{\kappa+3}$**:** In this game we replace $\mathrm{C}_{\mathrm{Mult}}$ with $\mathrm{C}^*_{\mathrm{Mult}}$, a circuit that uses the implicit values $[\omega^i]$ for $0 \le i \le \kappa$ in steps 5 and 6. (Note that $[P(\omega)]$ can be computed using $[\omega^i]$ when the coefficients of P are explicitly known.) This change does not affect the functionality of the multiplication circuit and hence we can bound this hope via PIO security. As a result, the explicit knowledge ω is no longer needed to generate the multiplication circuit.

$\text{Game}_{\kappa+4}$**:** In this game, we replace $[\omega^\kappa]$ with a random value $[\sigma]$ in challenge preparation. (Note that level-κ encodings correspond to the base group.) We can bound this hop via the κ-SDDH game.

In the final game the challenge exponent (when $b=1$) is fully randomized. This means that the challenge is independent of b in $\text{Game}_{\kappa+4}$, which concludes the proof.

7.2 Downgrading Attacks

It might appear that our GES could be subject to a "downgrading" attack as follow. Start with any consistent encoding h at level ℓ whose representation

polynomial is of degree 0. Then "maul" h into an encoding at a lower level $\ell' < \ell$ by simply changing ℓ to ℓ' in h. Then use this malleability to attack, say, MDDH where challenge encodings are canonical and of degree 0 (see Sect. 5.5).

What is crucial and prevents this downgrade attack is the proof system. The consistency proof π proves that the encrypted values correspond to a polynomial P of degree up to ℓ such that $P(\omega) = z$. Note that this statement *depends on* ℓ. Hence, a proof for a level-2 encoding cannot be "reused" for a level-1 encoding, as in the attack: a single proof will not necessarily pass against two different statements even if they both have the same witness. In order to downgrade, the proof would have to be changed.

Indeed, suppose that one had a method for changing a proof π_2 of a level-2 encoding to a proof π_1 of the level-1 encoding (that is derived by simply omitting encrypted coefficients, as in a downgrading attack). Consider what happens if one start with equivalent level-2 encoding (in the sense of our switching lemma) with degree-2 polynomials P. Then, the statement that π_1 proves becomes false, so any such attack would contradict the soundness of the proof system.

Acknowledgments. We thank the anonymous reviewers for their helpful comments, and Kenny Paterson and Geoffroy Couteau for useful discussions. Pooya Farshim was supported in part by grant ANR-14-CE28-0003 (Project EnBid). Dennis Hofheinz was supported by ERC grant 724307, and by DFG grants HO 4534/2-2 and HO 4534/4-1. Enrique Larraia was supported by EPSRC grant EP/L018543/1.

References

1. Abusalah, H., Fuchsbauer, G., Pietrzak, K.: Constrained PRFs for unbounded inputs. In: Sako, K. (ed.) CT-RSA 2016. LNCS, vol. 9610, pp. 413–428. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29485-8_24
2. Albrecht, M.R., Farshim, P., Hofheinz, D., Larraia, E., Paterson, K.G.: Multilinear maps from obfuscation. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016. LNCS, vol. 9562, pp. 446–473. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-49096-9_19
3. Ananth, P., Sahai, A.: Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In: Coron, J.-S., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 152–181. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-56620-7_6
4. Boneh, D., Boyen, X.: Short signatures without random oracles. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 56–73. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24676-3_4
5. Boneh, D., Lewi, K., Raykova, M., Sahai, A., Zhandry, M., Zimmerman, J.: Semantically secure order-revealing encryption: multi-input functional encryption without obfuscation. In: Oswald and Fischlin (eds.) [38], pp. 563–594
6. Boneh, D., Silverberg, A.: Applications of multilinear forms to cryptography. Contemp. Math. **324**, 71–90 (2003)
7. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-42045-0_15

8. Boneh, D., Waters, B., Zhandry, M.: Low overhead broadcast encryption from multilinear maps. In: Garay and Gennaro [21], pp. 206–223
9. Canetti, R., Garay, J.A. (eds.): CRYPTO 2013, Part I. LNCS, vol. 8042. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4
10. Canetti, R., Garay, J.A. (eds.): CRYPTO 2013, Part II. LNCS, vol. 8043. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40084-1
11. Canetti, R., Lin, H., Tessaro, S., Vaikuntanathan, V.: Obfuscation of probabilistic circuits and applications. In: Dodis and Nielsen [17], pp. 468–497
12. Cheon, J.H., Han, K., Lee, C., Ryu, H., Stehlé, D.: Cryptanalysis of the multilinear map over the integers. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 3–12. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_1
13. Coron, J.-S., Gentry, C., Halevi, S., Lepoint, T., Maji, H.K., Miles, E., Raykova, M., Sahai, A., Tibouchi, M.: Zeroizing without low-level zeroes: new MMAP attacks and their limitations. In: Gennaro and Robshaw [27], pp. 247–266
14. Coron, J.-S., Lee, M.S., Lepoint, T., Tibouchi, M.: Cryptanalysis of GGH15 multilinear maps. In: Robshaw and Katz [41], pp. 607–628
15. Coron, J.-S., Lepoint, T., Tibouchi, M.: Practical multilinear maps over the integers. In: Canetti and Garay [9], pp. 476–493
16. Coron, J.-S., Lepoint, T., Tibouchi, M.: New multilinear maps over the integers. In: Gennaro and Robshaw [27], pp. 267–286
17. Dodis, Y., Nielsen, J.B. (eds.): TCC 2015, Part II. LNCS, vol. 9015. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46497-7
18. Escala, A., Herold, G., Kiltz, E., Ràfols, C., Villar, J.: An algebraic framework for Diffie-Hellman assumptions. In: Canetti and Garay [10], pp. 129–147
19. Farshim, P., Hesse, J., Hofheinz, D., Larraia, E.: Graded encoding schemes from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2018/011 (2015)
20. Freire, E.S.V., Hofheinz, D., Paterson, K.G., Striecks, C.: Programmable hash functions in the multilinear setting. In: Canetti and Garay [9], pp. 513–530
21. Garay, J.A., Gennaro, R. (eds.): CRYPTO 2014, Part I. LNCS, vol. 8616. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44371-2
22. Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 1–17. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-38348-9_1
23. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS, pp. 40–49. IEEE Computer Society Press, October 2013
24. Garg, S., Gentry, C., Halevi, S., Sahai, A., Waters, B.: Attribute-based encryption for circuits from multilinear maps. In: Canetti and Garay [10], pp. 479–499
25. Garg, S., Gentry, C., Sahai, A., Waters, B.: Witness encryption and its applications. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC, pp. 467–476. ACM Press, June 2013
26. Garg, S., Mukherjee, P., Srinivasan, A.: Obfuscation without the vulnerabilities of multilinear maps. Cryptology ePrint Archive, Report 2016/390 (2016)
27. Gennaro, R., Robshaw, M. (eds.): CRYPTO 2015, Part I. LNCS, vol. 9215. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-47989-6
28. Gentry, C., Gorbunov, S., Halevi, S.: Graph-induced multilinear maps from lattices. In: Dodis and Nielsen [17], pp. 498–527

29. Goldwasser, S., Rothblum, G.N.: On best-possible obfuscation. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 194–213. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-70936-7_11
30. Groth, J., Sahai, A.: Efficient non-interactive proof systems for bilinear groups. In: Smart, N. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 415–432. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78967-3_24
31. Groth, J., Sahai, A.: Efficient noninteractive proof systems for bilinear groups. SIAM J. Comput. **41**(5), 1193–1232 (2012)
32. Hofheinz, D., Kiltz, E.: Programmable hash functions and their applications. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 21–38. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_2
33. Hohenberger, S., Sahai, A., Waters, B.: Full domain hash from (leveled) multilinear maps and identity-based aggregate signatures. In: Canetti and Garay [9], pp. 494–512
34. Langlois, A., Stehlé, D., Steinfeld, R.: GGHLite: more efficient multilinear maps from ideal lattices. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 239–256. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_14
35. Lin, H.: Indistinguishability obfuscation from DDH on 5-linear maps and locality-5 PRGs. Cryptology ePrint Archive, Report 2016/1096 (2016)
36. Lin, H., Tessaro, S.: Indistinguishability obfuscation from bilinear maps and block-wise local PRGs. Cryptology ePrint Archive, Report 2017/250 (2017)
37. Miles, E., Sahai, A., Zhandry, M.: Annihilation attacks for multilinear maps: cryptanalysis of indistinguishability obfuscation over GGH13. In: Robshaw and Katz [41], pp. 629–658
38. Oswald, E., Fischlin, M. (eds.): EUROCRYPT 2015, Part II. LNCS, vol. 9057. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46803-6
39. Paneth, O., Sahai, A.: On the equivalence of obfuscation and multilinear maps. Cryptology ePrint Archive, Report 2015/791 (2015)
40. Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation from semantically-secure multilinear encodings. In: Garay and Gennaro [21], pp. 500–517
41. Robshaw, M., Katz, J. (eds.): CRYPTO 2016, Part II. LNCS, vol. 9815. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53008-5
42. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (eds.) 46th ACM STOC, pp. 475–484. ACM Press, May/June 2014
43. Zhang, F., Safavi-Naini, R., Susilo, W.: An efficient signature scheme from bilinear pairings and its applications. In: Bao, F., Deng, R., Zhou, J. (eds.) PKC 2004. LNCS, vol. 2947, pp. 277–290. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24632-9_20