



HAL
open science

Fully Secure Functional Encryption for Inner Products, from Standard Assumptions

Shweta Agrawal, Benoît Libert, Damien Stehlé

► **To cite this version:**

Shweta Agrawal, Benoît Libert, Damien Stehlé. Fully Secure Functional Encryption for Inner Products, from Standard Assumptions. Crypto 2016, Aug 2016, Santa Barbara, United States. 10.1007/978-3-662-53015-3 . hal-01228559v2

HAL Id: hal-01228559

<https://inria.hal.science/hal-01228559v2>

Submitted on 15 Jun 2016 (v2), last revised 22 Nov 2016 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fully Secure Functional Encryption for Inner Products, from Standard Assumptions ^{*}

Shweta Agrawal¹, Benoît Libert², and Damien Stehlé²

¹ IIT Delhi, India

² ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France.

Abstract. Functional encryption is a modern public-key paradigm where a master secret key can be used to derive sub-keys SK_F associated with certain functions F in such a way that the decryption operation reveals $F(M)$, if M is the encrypted message, and nothing else. Recently, Abdalla *et al.* gave simple and efficient realizations of the primitive for the computation of linear functions on encrypted data: given an encryption of a vector \mathbf{y} over some specified base ring, a secret key $SK_{\mathbf{x}}$ for the vector \mathbf{x} allows computing $\langle \mathbf{x}, \mathbf{y} \rangle$. Their technique surprisingly allows for instantiations under standard assumptions, like the hardness of the Decision Diffie-Hellman (DDH) and Learning-with-Errors (LWE) problems. Their constructions, however, are only proved secure against *selective* adversaries, which have to declare the challenge messages M_0 and M_1 at the outset of the game.

In this paper, we provide constructions that provably achieve security against more realistic *adaptive* attacks (where the messages M_0 and M_1 may be chosen in the challenge phase, based on the previously collected information) for the same inner product functionality. Our constructions are obtained from hash proof systems endowed with homomorphic properties over the key space. They are (almost) as efficient as those of Abdalla *et al.* and rely on the same hardness assumptions.

In addition, we obtain a solution based on Paillier’s composite residuosity assumption, which was an open problem even in the case of selective adversaries. We also propose LWE-based schemes that allow evaluation of inner products modulo a prime p , as opposed to the schemes of Abdalla *et al.* that are restricted to evaluations of integer inner products of short integer vectors. We finally propose a solution based on Paillier’s composite residuosity assumption that enables evaluation of inner products modulo an RSA integer $N = p \cdot q$.

We demonstrate that the functionality of inner products over a prime field is powerful and can be used to construct bounded collusion FE for all circuits.

Keywords. Functional encryption, adaptive security, standard assumptions, DDH, LWE, extended LWE, composite residuosity.

^{*} This is the full version of a paper to appear in Crypto 2016.

1 Introduction

Functional encryption (FE) [55, 18] is a generalization of public-key encryption, which overcomes the all-or-nothing, user-based access to data that is inherent to public key encryption and enables fine grained, role-based access that makes it very desirable for modern applications. A bit more formally, given an encryption $\text{enc}(X)$ and a key corresponding to a function F , the key holder only learns $F(X)$ and nothing else. Apart from its theoretical appeal, the concept of FE also finds numerous applications. In cloud computing platforms, users can store encrypted data on a remote server and subsequently provide the server with a key SK_F which allows it to compute the function F of the underlying data without learning anything else.

In some cases, the message $X = (\text{IND}, M)$ consists of an index IND (which can be thought of as a set of descriptive attributes) and a message M , which is sometimes called “payload”. One distinguishes FE systems with public index, where IND is publicly revealed by the ciphertext but M is hidden, from those with private index, where IND and M are both hidden. Public index FE is popularly referred to as attribute based encryption.

A Brief History of FE. The birth of Functional Encryption can be traced back to Identity Based Encryption [56, 16] which can be seen as the first nontrivial generalization of Public Key Encryption. However, it was the work of Sahai and Waters [55] that coined the term Attribute Based Encryption, and the subsequent, natural unification of all these primitives under the umbrella of Functional Encryption took place only relatively recently [18, 48]. Constructions of public index FE have matured from specialized – equality testing [16, 12, 34], keyword search [15, 1, 43], boolean formulae [41], inner product predicates [43], regular languages [57] – to general polynomial-size circuits [33, 39, 17] and even Turing machines [36]. The journey of private index FE has been significantly more difficult, with inner product predicate constructions [43, 3] being the state of the art for a long time until the recent elegant generalization to polynomial-size circuits [40].

However, although private index FE comes closer than ever before to the goal of general FE, it falls frustratingly short. This is because all known constructions of private index FE only achieve *weak attribute hiding*, which severely restricts the function keys that the adversary can request in the security game – the adversary may request keys for functions f_i that do not decrypt the challenge ciphertext (IND^*, M^*) , i.e., $f_i(\text{IND}^*) \neq 0$ holds for all i . The most general notion of FE – private index, strongly attribute hiding – has been built for the restricted case of bounded collusions [38, 37] or using the brilliant, but ill-understood³ machinery of multi-linear maps [32] and indistinguishability obfuscation [32]. These constructions provide FE for general polynomial-size circuits and Turing

³ Indeed, the two candidate multi-linear maps [31, 23] put forth in 2013 were recently found to be insecure [22, 42].

machines [36], but, perhaps surprisingly, there has been little effort to build the general notion of FE ground-up, starting from smaller functionalities.

This appears as a gaping hole that begs to be filled. Often, from the practical standpoint, efficient constructions for a smaller range of functionalities, such as linear functions or polynomials, are extremely relevant, and such an endeavour will also help us understand the fundamental barriers that thwart our attempts for general FE. This motivates the question:

Can we build FE for restricted classes of functions, satisfying standard security definitions, under well-understood assumptions?

In 2015, Abdalla, Bourse, De Caro and Pointcheval [2] considered the question of building FE for linear functions. Here, a ciphertext C encrypts a vector $\mathbf{y} \in \mathcal{D}^\ell$ over some ring \mathcal{D} , a secret key for the vector $\mathbf{x} \in \mathcal{D}^\ell$ allows computing $\langle \mathbf{x}, \mathbf{y} \rangle$ and nothing else about \mathbf{y} . Note that this is quite different from the inner product predicate functionality of [43, 3]: the former computes the actual value of the inner product while the latter tests whether the inner product is zero or not, and reveals a hidden bit M if so. Abdalla *et al.* [2] showed, surprisingly, that this functionality allows for very simple and efficient realizations under standard assumptions like the Decision Diffie-Hellman (DDH) and Learning-with-Errors (LWE) assumptions [52]. The instantiation from DDH was especially unexpected since DDH is not known to easily lend itself to the design of such primitives.⁴ What enables this surprising result is that the functionality itself is rather limited – note that with ℓ queries, the adversary can reconstruct the entire message vector. Due to this, the scheme need not provide *collusion resistance*, which posits that no collection of secret keys for functions F_1, \dots, F_q should make it possible to decrypt a ciphertext that no individual such key can decrypt. Collusion resistance is usually the chief obstacle in proving security of FE schemes. On the contrary, for linear FE constructions, if two adversaries combine their keys, they do get a valid new key, but this key gives them a plaintext which could anyway be computed by their individual plaintexts. Hence, collusion is permitted by the functionality itself, and constructions can be much simpler. As we shall see below, linear FE is already very useful and yields many interesting applications, as we discuss in Appendix B.

More recently, Bishop, Jain and Kowalczyk [11] considered the same functionality as Abdalla *et al.* in the secret-key setting with the motivation of achieving function privacy.

While [11] considers adaptive adversaries, their construction requires bilinear maps and does not operate over standard DDH-hard groups. In the public-key setting, Abdalla *et al.* [2] only proved their schemes to be secure against *selective* adversaries, that have to declare the challenge messages M_0, M_1 of the semantic security game upfront, before seeing the master public key mpk . Selective security

⁴ And indeed, this unsuitability partially manifests itself in the limitation of message/function space of the aforementioned construction: message/function vectors must be short integer vectors, and the inner product is evaluated over the integers.

is usually too weak a notion for practical applications and is often seen as a stepping stone to proving full adaptive security. Historically, most flavors of functional encryption have been first realized for selective adversaries [12, 55, 41, 43, 32] before being upgraded to attain full security. Boneh and Boyen [13] observed that a standard complexity leveraging argument can be used to argue that a selectively-secure system is also adaptively secure. However, this argument is not satisfactory in general as the reduction incurs an exponential security loss in the message length. Quite recently, Ananth, Brakerski, Segev and Vaikuntanathan [7] described a generic method of building adaptively secure functional encryption systems from selectively secure ones. However their transformation is based on the existence of a *sufficiently expressive* selectively secure FE scheme, where sufficiently secure roughly means capable of evaluating a weak PRF. Since no such scheme from standard assumptions is known, their transformation does not apply to our case, and in any case would significantly increase the complexity of the construction, even if it did.

Our Results. In this paper, we describe fully secure functional encryption systems for the evaluation of inner products on encrypted data. We propose schemes that evaluate inner products of integer vectors, based on DDH, LWE and the Composite Residuosity hardness assumptions. Our DDH-based and LWE-based constructions for integer inner products are of efficiency comparable to those of Abdalla *et al.* [2] and rely on the same standard assumptions. Note that a system based on Paillier’s composite residuosity assumption was an open problem even for the case of selective adversaries, which we resolve in this work.

Additionally, we propose schemes that evaluate inner products modulo a prime p or a composite $N = pq$, based on the LWE and Composite Residuosity hardness assumptions. In contrast, the constructions of [2] must restrict the ring \mathcal{D} to the ring of integers, which is a significant drawback. Indeed, although their DDH-based realization allows evaluating $\langle \mathbf{x}, \mathbf{y} \rangle \bmod p$ when the latter value is sufficiently small, their security proof restricts the functionality to the computation of $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$.

The functionality of inner products over a prime field is powerful: we show that it can be bootstrapped all the way to yield a conceptually simple construction for bounded collusion FE for all circuits. The only known construction for general FE handling bounded collusions is by Gorbunov, Vaikuntanathan and Wee [38]. Our construction is conceptually simpler, albeit a bit more inefficient. Also, since it requires the inner product functionality over a prime field, it can only be instantiated with our LWE-based scheme for now.

1.1 Overview of techniques

We briefly summarize our techniques below.

Fully secure linear FE: hash proof systems. Our DDH-based construction and its security proof implicitly build on hash proof systems [25]. It involves

public parameters comprised of group elements $(g, h, \{h_i = g^{s_i} \cdot h^{t_i}\}_{i=1}^\ell)$, where g, h generate a cyclic group \mathbb{G} of prime order q , and the master secret key is $\text{msk} = (\mathbf{s}, \mathbf{t}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q^\ell$. On input of a vector $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_q^\ell$, the encryption algorithm computes $(g^r, h^r, \{g^{y_i} \cdot h_i^r\}_{i=1}^\ell)$ in such a way that a secret key of the form $SK_{\mathbf{x}} = (\langle \mathbf{s}, \mathbf{x} \rangle, \langle \mathbf{t}, \mathbf{x} \rangle)$ allows computing $g^{\langle \mathbf{y}, \mathbf{x} \rangle}$ in the same way as in [2]. Despite its simplicity and its efficiency (only one more group element than in [2] is needed in the ciphertext), we show that the above system can be proved fully secure using arguments – akin to those of Cramer and Shoup [24] – which consider what the adversary knows about the master secret key $(\mathbf{s}, \mathbf{t}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q^\ell$ in the information theoretic sense. The security proof is arguably simpler than its counterpart in the selective case [2]. As in all security proofs based on hash proof systems, it uses the fact that the secret key is known to the reduction at any time, which makes it simpler to handle secret key queries without knowing the adversary’s target messages $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_q^\ell$ in advance.

While our DDH-based realization only enables efficient decryption when the inner product $\langle \mathbf{x}, \mathbf{y} \rangle$ is contained in a sufficiently small interval, we show how to eliminate this restriction using Paillier’s cryptosystem in the same way as in [21, 20]. We thus obtain the first solution based on the Composite Residuosity assumption, which was previously an open problem (even in the case of selective security).

LWE-based fully secure linear FE. Our LWE-based construction builds on the dual Regev encryption scheme from Gentry, Peikert and Vaikuntanathan [34]. Its security analysis requires more work. The master public key contains a random matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. For simplicity, we restrict ourselves to plaintext vectors and secret key vectors with binary coordinates. Each vector coordinate $i \in \{1, \dots, \ell\}$ requires a master public key component $\mathbf{u}_i^T = \mathbf{z}_i^T \cdot \mathbf{A} \in \mathbb{Z}_q^n$, for a small norm vector $\mathbf{z}_i \in \mathbb{Z}^m$ made of Gaussian entries which will be part of the master secret key $\text{msk} = \{\mathbf{z}_i\}_{i=1}^\ell$. Each $\{\mathbf{u}_i\}_{i=1}^\ell$ can be seen as a syndrome in the GPV trapdoor function for which vector \mathbf{z}_i is a pre-image. Our security analysis will rely on the fact that each GPV syndrome has a large number of pre-images and, conditionally on $\mathbf{u}_i \in \mathbb{Z}_q^n$, each \mathbf{z}_i retains a large amount of entropy. In the security proof, this will allow us to apply arguments similar to those of hash proof systems [25] when we will generate the challenge ciphertext using $\{\mathbf{z}_i\}_{i=1}^\ell$. More precisely, when the first part $\mathbf{c}_0 \in \mathbb{Z}_q^m$ of the ciphertext is a random vector instead of an actual LWE sample $\mathbf{c}_0 = \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_0$, the action of $\{\mathbf{z}_i\}_{i=1}^\ell$ on $\mathbf{c}_0 \in \mathbb{Z}_q^m$ produces vectors that appear statistically uniform to any legitimate adversary. In order to properly simulate the challenge ciphertext using the master secret key $\{\mathbf{z}_i\}_{i=1}^\ell$, we use a variant of the extended LWE assumption [49] (eLWE) so as to have the (hint) values $\{\langle \mathbf{z}_i, \mathbf{e}_0 \rangle\}_{i=1}^\ell$ at disposal. One difficulty is that the reductions from LWE to eLWE proved in [6] and [19] handle a single hint vector \mathbf{z} . Fortunately, we extend the techniques of Brakerski *et al.* [19] using the gadget matrix from [44] to obtain a reduction from LWE to the multi-hint variant of eLWE that we use in the security proof. More specifically, we prove that the multi-hint variant of

eLWE remains at least as hard as LWE when the adversary obtains as many as $n/2$ hints, where n is the dimension of the LWE secret.

Evaluation inner products modulo p . Our construction from the DDH assumption natively supports the computation of inner products modulo a prime p as long as the remainder $\langle \mathbf{x}, \mathbf{y} \rangle \bmod p$ falls in a polynomial-size interval. Under the Paillier and LWE assumptions, we first show how to compute integer inner products $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$. In a second step, we upgrade our Paillier and LWE-based systems so as to compute inner products modulo a composite $N = pq$ and a prime p , respectively, without leaking the actual value $\langle \mathbf{x}, \mathbf{y} \rangle$ over \mathbb{Z} .

Hiding anything but the remainder modulo N or p requires additional techniques. In the context of LWE-based FE, this is achieved by using an LWE modulus of the form $q = p \cdot p'$ and multiplying plaintexts by p' , so that an inner product modulo q over the ciphertext space natively translates into an inner product modulo p for the underlying plaintexts.

The latter plaintext/ciphertext manipulations do not solve another difficulty which arises from the discrepancy between the base rings of the master key and the secret key vectors: indeed, the master key consists of integer vectors, whereas the secret keys are defined modulo an integer. When the adversary queries a secret key vector $\mathbf{x} \in \mathbb{Z}_p^\ell$ (or \mathbb{Z}_N^ℓ), it gets the corresponding combination modulo p of the master key components. By making appropriate vector queries that are linearly dependent modulo p (and hence valid), an attacker could learn a combination of the master key components which is singular modulo p but invertible over the field of rational numbers: it would then obtain the whole master key! However, note that as long as the adversary only queries secret keys for $\ell - 1$ independent vectors over \mathbb{Z}_p^ℓ (or \mathbb{Z}_N^ℓ), there is no reason not to reveal more than $\ell - 1$ secret keys overall. In order to make sure that the adversary only obtains redundant information by making more than $\ell - 1$ queries, we assume that a trusted authority keeps track of all vectors \mathbf{x} for which secret keys were previously given out (more formally, the key generation algorithm is stateful).

Compiling Linear FE to Bounded Collusion General FE. We provide a conceptually simpler way to build q -query Functional Encryption for all circuits. The only known construction for this functionality was suggested by Gorbunov *et al.* in [38]. At a high level, the q -query construction by Gorbunov *et al.* is built in several layers, as follows:

1. They start with a single key FE scheme for all circuits, which was provided by [54].
2. The single FE scheme is compiled into a q -query scheme for NC_1 circuits. This is the most non-trivial part of the construction. They run N copies of the single key scheme, where $N = O(q^4)$. To encrypt, they encrypt the views of some N -party MPC protocol computing some functionality related to C , *à la* “MPC in the head”. For the MPC protocol, they use the BGW [9] semi-honest MPC protocol without degree reduction and exploit the fact that this protocol is completely non-interactive when used to compute bounded

degree functions. The key generator provides the decryptor with a subset of the single query FE keys, where the subsets are guaranteed to have small pairwise intersections. This subset of keys enables the decryptor to recover sufficiently many shares of $C(x)$ which enables her to compute $C(x)$ via polynomial interpolation. However, an attacker with q keys only learns a share x_i in the clear if two subsets of keys intersect, and due to small pairwise intersections, this does not occur often enough for him learn sufficiently many shares of x , hence, by the guarantees of secret sharing, input x remains hidden.

3. Finally, they bootstrap the q -query FE for NC_1 to a q -query FE for all circuits using computational randomized encodings [8]. They must additionally use cover-free sets to ensure that fresh randomness is used for each randomized encoding.

Our construction replaces steps 1 and 2 with a inner product modulo p FE scheme, and then uses step 3 as in [38]. Thus, the construction of single key FE in step 1 by Sahai and Seyalioglu, and the nontrivial “MPC in the head” of step 2 can both be replaced by the simple abstraction of an inner product FE scheme. For step 3, observe that the bootstrapping theorem of [38] provides a method to bootstrap an FE for NC_1 that handles q queries to an FE for all polynomial-size circuits that is also secure against q queries. The bootstrapping relies on the result of Applebaum *et al.* [8, Theorem 4.11] which states that every polynomial time computable function f admits a perfectly correct computational randomized encoding of degree 3. In more detail, let \mathcal{C} be a family of polynomial-size circuits. Let $C \in \mathcal{C}$ and let x be some input. Let $\tilde{C}(x, R)$ be a randomized encoding of C that is computable by a constant depth circuit with respect to inputs x and R . Then consider a new family of circuits \mathcal{G} defined by:

$$G_{C,\Delta}(x, R_1, \dots, R_S) = \left\{ \tilde{C}\left(x; \bigoplus_{a \in \Delta} R_a\right) : C \in \mathcal{C}, \Delta \subseteq [S] \right\},$$

for some sufficiently large S (quadratic in the number of queries q). As observed in [38], circuit $G_{C,\Delta}(\cdot, \cdot)$ is computable by a constant degree polynomial (one for each output bit). Given an FE scheme for \mathcal{G} , one may construct a scheme for \mathcal{C} by having the decryptor first recover the output of $G_{C,\Delta}(x, R_1, \dots, R_S)$ and then applying the decoder for the randomized encoding to recover $C(x)$.

However, to support q queries the decryptor must compute q randomized encodings, each of which needs fresh randomness. This is handled by hardcoding S random elements in the ciphertext and using random subsets $\Delta \subseteq [S]$ (which are cover-free with overwhelming probability) to compute fresh randomness $\bigoplus_{a \in \Delta} R_a$ for every query. The authors then conclude that bounded query FE for NC_1 suffices to construct a bounded query FE scheme for all circuits.

We observe that the ingredient required to bootstrap is *not FE for the entire circuit class* NC_1 but rather only the particular circuit class \mathcal{G} as described above. This circuit class, being computable by degree 3 polynomials, may be supported by a linear FE scheme, via linearization of the degree 3 polynomials! To illustrate,

let us consider FE secure only for a single key. Then, the functionality that the initial FE must support is exactly the randomized encoding of [8], which, indeed, is in NC_0 . Now, to support q queries, we must ensure that each key uses a fresh piece of randomness, and this is provided using a cover-free set family S as in [38] – the key generator picks a random subset $\Delta \subseteq [S]$ and sums up its elements to obtain i.i.d. randomness for the key being requested. To obtain a random element in this manner, addition over the integers does not suffice, we must take addition modulo p . Here, our inner product modulo p construction comes to our rescue!

Putting it together, the encryptor encrypts all degree 3 monomials in the inputs R_1, \dots, R_S and x_1, \dots, x_ℓ . Note that this ciphertext is polynomial in size. Now, for a given circuit C , the keygen algorithm samples some $\Delta \subseteq [S]$ and computes the symbolic degree 3 polynomials which must be released to the decryptor. It then provides the linear FE keys to compute the same. By correctness and security of Linear FE as well as the randomizing polynomial construction, the decryptor learns $C(x)$ and nothing else. The final notion of security that we obtain is non-adaptive simulation based security **NA-SIM** [48, 38], i.e. (poly, poly, 0) **SIM** security, where the adversary can request a polynomial number of pre-challenge keys, ask for polynomially sized challenge ciphertexts but may not request post-challenge keys. For more details, we refer the reader to Section 6. We note that the construction of [38] also achieves the stronger **AD-SIM** security, but for a scheme that supports only a *single* ciphertext and bounded number of keys. The bound on the number of ciphertexts is necessary due to a lower bound by [18]. The notion of single ciphertext, bounded key FE appears to be quite restrictive, hence we do not study **AD-SIM** security here.

We note that subsequent to our work, Agrawal and Rosen [5] used our adaptively secure mod p inner products FE scheme in a more sophisticated manner than we do here, to achieve ciphertext size that improves upon the construction of [38].

2 Background

In this section, we recall the hardness assumptions underlying the security of the schemes we will describe. The functionality and security definitions of functional and non-interactive controlled functional encryption schemes are given in Appendix A.

Our first scheme relies on the standard DDH assumption in ordinary (i.e., non-pairing-friendly) cyclic groups.

Definition 1. *In a cyclic group \mathbb{G} of prime order q , the **Decision Diffie-Hellman (DDH)** problem is to distinguish the distributions $D_0 = \{(g, g^a, g^b, g^{ab}) \mid g \leftarrow \mathbb{G}, a, b \leftarrow \mathbb{Z}_q\}$, $D_1 = \{(g, g^a, g^b, g^c) \mid g \leftarrow \mathbb{G}, a, b, c \leftarrow \mathbb{Z}_q\}$.*

A variant of our first scheme relies on Paillier’s composite residuosity assumption.

Definition 2 ([50]). *Let $N = pq$, for prime numbers p, q . The **Decision Composite Residuosity (DCR)** problem is to distinguish the distributions $D_0 := \{z = z_0^N \bmod N^2 \mid z_0 \leftarrow \mathbb{Z}_N^*\}$ and $D_1 := \{z \leftarrow \mathbb{Z}_{N^2}^*\}$.*

Our third construction builds on the Learning-With-Errors (LWE) problem, which is known to be at least as hard as certain standard lattice problems in the worst case [53, 19].

Definition 3. Let q, α, m be functions of a parameter n . For a secret $\mathbf{s} \in \mathbb{Z}_q^n$, the distribution $A_{q,\alpha,\mathbf{s}}$ over $\mathbb{Z}_q^n \times \mathbb{Z}_q$ is obtained by sampling $\mathbf{a} \leftarrow \mathbb{Z}_q^n$ and an $e \leftarrow D_{\mathbb{Z},\alpha,q}$, and returning $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e) \in \mathbb{Z}_q^{n+1}$. The Learning With Errors problem $\text{LWE}_{q,\alpha,m}$ is as follows: For $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, the goal is to distinguish between the distributions:

$$D_0(\mathbf{s}) := U(\mathbb{Z}_q^{m \times (n+1)}) \quad \text{and} \quad D_1(\mathbf{s}) := (A_{q,\alpha,\mathbf{s}})^m.$$

We say that a PPT algorithm \mathcal{A} solves $\text{LWE}_{q,\alpha}$ if it distinguishes $D_0(\mathbf{s})$ and $D_1(\mathbf{s})$ with non-negligible advantage (over the random coins of \mathcal{A} and the randomness of the samples), with non-negligible probability over the randomness of \mathbf{s} .

3 Fully secure functional encryption for inner products from DDH

In this section, we show that an adaptation of the DDH-based construction of Abdalla *et al.* [2] provides full security under the standard DDH assumption. Like [2], the scheme computes inner products over \mathbb{Z} as long as they land in a sufficiently small interval.

In comparison with the solution of Abdalla *et al.*, we only introduce one more group element in the ciphertext and all operations are just as efficient as in [2]. Our scheme is obtained by modifying [2] in the same way as Damgård's encryption scheme [26] was obtained from the Elgamal cryptosystem. The original DDH-based system of [2] encrypts a vector $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_q^\ell$ by computing $(g^r, \{g^{y_i} \cdot h_i^r\}_{i=1}^\ell)$, where $\{h_i = g^{s_i}\}_{i=1}^\ell$ are part of the master public key and $\text{sk}_{\mathbf{x}} = \sum_{i=1}^\ell s_i \cdot x_i \bmod q$ is the secret key associated with the vector $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}_q^\ell$. Here, we encrypt \mathbf{y} in the fashion of Damgård's Elgamal, by computing $(g^r, h^r, \{g^{y_i} \cdot h_i^r\}_{i=1}^\ell)$. The decryption algorithm uses secret keys of the form $\text{sk}_{\mathbf{x}} = (\sum_{i=1}^\ell s_i \cdot x_i, \sum_{i=1}^\ell t_i \cdot x_i)$, where $h_i = g^{s_i} \cdot h^{t_i}$ for each i and $\mathbf{s} = (s_1, \dots, s_\ell) \in \mathbb{Z}_q^\ell$ and $\mathbf{t} = (t_1, \dots, t_\ell) \in \mathbb{Z}_q^\ell$ are part of the master key msk .

The scheme and its security proof also build on ideas from the Cramer-Shoup cryptosystem [24, 25]. Analogously to the bounded-collusion-resistant IBE schemes of Goldwasser *et al.* [35], the construction can be seen as an applying a hash proof system [25] with homomorphic properties over the key space. It also bears similarities with the broadcast encryption system of Dodis and Fazio [28] in the way to use hash proof systems to achieve adaptive security.

Setup($1^\lambda, 1^\ell$): Choose a cyclic group \mathbb{G} of prime order $q > 2^\lambda$ with generators $g, h \leftarrow \mathbb{G}$. Then, for each $i \in \{1, \dots, \ell\}$, sample $s_i, t_i \leftarrow \mathbb{Z}_q$ and compute $h_i = g^{s_i} \cdot h^{t_i}$. Define $\text{msk} := \{(s_i, t_i)\}_{i=1}^\ell$ and

$$\text{mpk} := \left(\mathbb{G}, g, h, \{h_i\}_{i=1}^\ell \right).$$

Keygen(msk, \mathbf{x}): To generate a key for the vector $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}_q^\ell$, compute

$$\text{sk}_{\mathbf{x}} = (s_{\mathbf{x}}, t_{\mathbf{x}}) = (\sum_{i=1}^{\ell} s_i \cdot x_i, \sum_{i=1}^{\ell} t_i \cdot x_i) = (\langle \mathbf{s}, \mathbf{x} \rangle, \langle \mathbf{t}, \mathbf{x} \rangle).$$

Encrypt(mpk, \mathbf{y}): To encrypt a vector $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_q^\ell$, sample $r \leftarrow \mathbb{Z}_q$ and compute

$$C = g^r, \quad D = h^r, \quad \{E_i = g^{y_i} \cdot h_i^r\}_{i=1}^{\ell}.$$

Return $C_{\mathbf{y}} = (C, D, E_1, \dots, E_\ell)$.

Decrypt($\text{mpk}, \text{sk}_{\mathbf{x}}, C_{\mathbf{y}}$): Given $\text{sk}_{\mathbf{x}} = (s_{\mathbf{x}}, t_{\mathbf{x}})$, compute

$$E_{\mathbf{x}} = \left(\prod_{i=1}^{\ell} E_i^{x_i} \right) / (C^{s_{\mathbf{x}}} \cdot D^{t_{\mathbf{x}}}).$$

Then, compute and output $\log_g(E_{\mathbf{x}})$.

The decryption algorithm requires to compute a discrete logarithm. This is in general too expensive. Like in [2], this can be circumvented by imposing that the computed inner product lies in an interval $\{0, \dots, L\}$, for some polynomially bounded integer L . Then, computing the required discrete logarithm may be performed in time $\tilde{O}(L^{1/2})$ using Pollard's kangaroo method [51]. As reported in [10], this can be reduced to $\tilde{O}(L^{1/3})$ operations by precomputing a table of size $\tilde{O}(L^{1/3})$. Note that even though the functionality is limited (decryption may not be performed efficiently for all key vectors and for all message vectors), while proving security we will let the adversary query any key vector in \mathbb{Z}_q^ℓ .

Before proceeding with the security proof, we would like to clarify that, although the scheme of [2] only decrypts values in a polynomial-size space, the usual complexity leveraging argument does not prove it fully secure via a polynomial reduction. Indeed, when ℓ is polynomial in λ , having the inner product $\langle \mathbf{y}, \mathbf{x} \rangle$ in a small interval does not mean that original vector $\mathbf{y} \in \mathbb{Z}_q^\ell$ lives in a polynomial-size universe. In Section 5, we show how to eliminate the small-interval restriction using Paillier's cryptosystem [50].

The security analysis uses similar arguments to those of Cramer and Shoup [24, 25] in that it exploits the fact that mpk does not reveal too much information about the master secret key. At some step, the challenge ciphertext is generated using msk instead of the public key and, as long as msk retains a sufficient amount of entropy from the adversary's view, it will perfectly hide which vector among $\mathbf{y}_0, \mathbf{y}_1$ is actually encrypted. The reason why we can prove adaptive security is the fact that, as usual in security proofs relying on hash proof systems [24, 25], the reduction knows the master secret key at any time. It can thus correctly answer all secret key queries without knowing the challenge messages $\mathbf{y}_0, \mathbf{y}_1$ beforehand.

The DDH-based scheme can easily be generalized so as to rely on weaker variants of DDH, like the Decision Linear assumption [14] or the Matrix DDH assumption [30].

Theorem 1. *The scheme provides full security under the DDH assumption. (The proof is given in Appendix E).*

4 Full security under the LWE assumption

We describe two LWE-based schemes: the first one for integer inner products of short integer vectors, the second one for inner products over a prime field \mathbb{Z}_p .

In both cases, the security relies on the hardness of a variant of the extended-LWE problem. The extended-LWE problem introduced by O’Neill, Peikert and Waters [49] and further investigated in [6, 19]. At a high level, the extended-LWE problem can be seen as $\text{LWE}_{\alpha, q}$ with a fixed number m of samples, for which some extra information on the LWE noises is provided: the adversary is provided a given linear combination of the noise terms. More concretely, the problem is to distinguish between the distributions

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}, \mathbf{z}, \langle \mathbf{e}, \mathbf{z} \rangle) \text{ and } (\mathbf{A}, \mathbf{u}, \mathbf{z}, \langle \mathbf{e}, \mathbf{z} \rangle),$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \alpha q}^m$, and \mathbf{z} is sampled from a specified distribution. Note that in [49], a noise was added to the term $\langle \mathbf{e}, \mathbf{z} \rangle$. The LWE to extended-LWE reductions from [6, 19] do not require such an extra noise term.

We will use a variant of extended-LWE for which multiple hints $(\mathbf{z}_i, \langle \mathbf{e}, \mathbf{z}_i \rangle)$ are given, for the same noise vector \mathbf{e} .

Definition 4 (Multi-hint extended-LWE). *Let q, m, t be integers, α be a real and τ be a distribution over $\mathbb{Z}^{t \times m}$, all of them functions of a parameter n . The multi-hint extended-LWE problem $\text{mheLWE}_{q, \alpha, m, t, \tau}$ is to distinguish between the distributions of the tuples*

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}, \mathbf{Z}, \mathbf{Z} \cdot \mathbf{e}) \text{ and } (\mathbf{A}, \mathbf{u}, \mathbf{Z}, \mathbf{Z} \cdot \mathbf{e}),$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, $\mathbf{e} \leftarrow D_{\mathbb{Z}, \alpha q}^m$, and $\mathbf{Z} \leftarrow \tau$.

A reduction from LWE to mheLWE is presented in Subsection 4.3.

4.1 Integer inner products of short integer vectors

In the description hereunder, we consider the message space $\mathcal{P} = \{0, \dots, P-1\}^\ell$, for some integer P and where $\ell \in \text{poly}(n)$ denotes the dimension of vectors to encrypt. Secret keys are associated with vectors in $\mathcal{V} = \{0, \dots, V-1\}^\ell$ for some integer V . As in the DDH case, inner products are evaluated over \mathbb{Z} . However, unlike our DDH-based construction, we can efficiently decrypt without confining inner product values within a small interval: here the inner product between the plaintext and key vectors belongs to $\{0, \dots, K-1\}$ with $K = \ell PV$, and it is possible to set parameters so that the scheme is secure under standard hardness assumptions while K is more than polynomial in the security parameter. We compute ciphertexts using a prime modulus q , with q significantly larger than K .

Setup $(1^n, 1^\ell, P, V)$: Set integers $m, q \geq 2$, real $\alpha \in (0, 1)$ and distribution τ over $\mathbb{Z}^{\ell \times m}$ as explained below. Set $K = \ell PV$. Sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ and $\mathbf{Z} \leftarrow \tau$. Compute $\mathbf{U} = \mathbf{Z} \cdot \mathbf{A} \in \mathbb{Z}_q^{\ell \times n}$. Define $\text{mpk} := (\mathbf{A}, \mathbf{U}, K, P, V)$ and $\text{msk} := \mathbf{Z}$.

Keygen(msk, \mathbf{x}): Given a vector $\mathbf{x} \in \mathcal{V}$, compute and return the secret key $\mathbf{z}_x := \mathbf{x}^T \cdot \mathbf{Z} \in \mathbb{Z}^m$.

Encrypt(mpk, \mathbf{y}): To encrypt a vector $\mathbf{y} \in \mathcal{P}$, sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}, \alpha q}^m$ and $\mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha q}^\ell$ and compute

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_0 \in \mathbb{Z}_q^m, \\ \mathbf{c}_1 &= \mathbf{U} \cdot \mathbf{s} + \mathbf{e}_1 + \left\lfloor \frac{q}{K} \right\rfloor \cdot \mathbf{y} \in \mathbb{Z}_q^\ell. \end{aligned}$$

Then, return $C := (\mathbf{c}_0, \mathbf{c}_1)$.

Decrypt($\text{mpk}, \mathbf{x}, \mathbf{z}_x, C$): Given a ciphertext $C := (\mathbf{c}_0, \mathbf{c}_1)$ and a secret key \mathbf{z}_x for $\mathbf{x} \in \mathcal{V}$, compute $\mu' = \langle \mathbf{x}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_x, \mathbf{c}_0 \rangle \bmod q$ and output the value $\mu \in \{-K + 1, \dots, K - 1\}$ that minimizes $|\lfloor \frac{q}{K} \rfloor \cdot \mu - \mu'|$.

Setting the parameters. Let B_τ be such that with probability $\geq 1 - n^{-\omega(1)}$, each row of sample from τ has norm $\leq B_\tau$. As explained just below, correctness may be ensured by setting

$$\alpha^{-1} \geq K^2 B_\tau \omega(\sqrt{\log n}) \quad \text{and} \quad q \geq \alpha^{-1} \omega(\sqrt{\log n}).$$

The choice of τ is driven by the reduction from LWE to mhlWE (as summarized in Theorem 4), and more precisely from Lemma 4 (another constraint arises from the use of Lemma 9 at the end of the security proof). We may choose $\tau = D_{\mathbb{Z}, \sigma_1}^{\ell \times m/2} \times (D_{\mathbb{Z}^{m/2}, \sigma_2, \delta_1} \times \dots \times D_{\mathbb{Z}^{m/2}, \sigma_2, \delta_\ell})$, where $\delta_i \in \mathbb{Z}^\ell$ denotes the i th canonical vector, and the standard deviation parameters satisfy $\sigma_1 = \Theta(\sqrt{n \log m} \max(\sqrt{m}, K))$ and $\sigma_2 = \Theta(n^{7/2} m^{1/2} \max(m, K^2) \log^{5/2} m)$.

To ensure security based on $\text{LWE}_{q, \alpha', m}$ in dimension $\geq c \cdot n$ for some $c \in (0, 1)$ via Theorems 2 and 4 below, one may further impose that $\ell \leq (1 - c) \cdot n$ and $m = \Theta(n \log q)$, to obtain $\alpha' = \Omega(\alpha / (n^6 K \log^2 q \log^{5/2} n))$. Note that $\text{LWE}_{q, \alpha', m}$ enjoys reductions from lattice problems when $q \geq \Omega(\sqrt{n} / \alpha')$.

Combining the security and correctness requirements, we may choose $\alpha' = 1 / ((n \log q)^{O(1)} \cdot K^2)$ and $q = \Omega(\sqrt{n} / \alpha')$, resulting in LWE parameters that make LWE resist all known attacks running in time 2^λ , as long as $n \geq \tilde{\Omega}(\lambda \log K)$.

Decryption correctness. To show the correctness of the scheme, we first observe that, modulo q :

$$\begin{aligned} \mu' &= \langle \mathbf{x}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_x, \mathbf{c}_0 \rangle \\ &= \lfloor q/K \rfloor \cdot \langle \mathbf{x}, \mathbf{y} \rangle + \langle \mathbf{x}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_x, \mathbf{e}_0 \rangle. \end{aligned}$$

Below, we show that the magnitude of the term $\langle \mathbf{x}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_x, \mathbf{e}_0 \rangle$ is $\leq \ell V B_\tau \alpha q \omega(\sqrt{\log n})$ with probability $\geq 1 - n^{-\omega(1)}$. Thanks to the choices of α and q , the latter upper bound is $\leq \lfloor q/K \rfloor / 4$, which suffices to guarantee decryption correctness.

Note that \mathbf{e}_1 is an integer Gaussian vector of dimension ℓ and standard deviation $\alpha q \geq \omega(\sqrt{\log n})$, and that $\|\mathbf{x}\| \leq \sqrt{\ell V}$. As a result, we have

that $|\langle \mathbf{x}, \mathbf{e}_1 \rangle| \leq \sqrt{\ell} V \alpha q \omega(\sqrt{\log n})$ holds with probability $1 - n^{-\omega(1)}$. Similarly, as $\|\mathbf{z}_x\| \leq \ell V B_\tau$, we obtain that $|\langle \mathbf{z}_x, \mathbf{e}_0 \rangle| \leq \ell V B_\tau \alpha q \omega(\sqrt{\log n})$ holds with probability $1 - n^{-\omega(1)}$.

Full security. In order to prove adaptive security of the scheme, we use the multi-hint extended-LWE from Definition 4. Before we provide the formal proof, we provide some intuition.

Intuition. Here we describe some challenges in proving adaptive security for our LWE construction. To begin we describe the approach used by Abdalla et. al. [2] in showing selective security for a similar construction. In the selective game, the adversary must announce the challenge vectors $\mathbf{y}_0, \mathbf{y}_1$ at the outset of the game. By definition of an admissible adversary, every query \mathbf{x}^i made must satisfy the property that $\langle \mathbf{x}^i, (\mathbf{y}_0 - \mathbf{y}_1) \rangle = 0$ (over \mathbb{Z}) for all i . For ease of exposition, consider challenge messages $\mathbf{y}_0, \mathbf{y}_1$ that only differ in the last co-ordinate. Then, the simulator knows at the very beginning of the game, the subspace within which all queries must lie. Since the secret key is structured as $(\mathbf{x}^i)^T \mathbf{Z}$, it suffices for the simulator to pick all but the final column of \mathbf{Z} in order to answer all legitimate key requests. It can set the public parameters by constructing all except one row of \mathbf{U} using its choice of \mathbf{Z} , and receiving the final \mathbf{u}_ℓ from the LWE oracle. Now the challenge ciphertext can be embedded along this dimension to argue security.

In the adaptive game however, the simulator cannot know in advance which subspace the adversary's queries will lie in, hence it must pick the entire master secret key \mathbf{Z} to answer key requests. Given that the simulator has no secrets, it is unclear how it may leverage the adversary. To handle this, our approach is to carefully analyze the entropy loss that occurs in the master secret \mathbf{Z} via that keys seen by the adversary. We show that despite seeing linear relations involving \mathbf{Z} , there is enough residual entropy left in the master secret so that the challenge ciphertext created using this appears uniform to the adversary.

To the best of our knowledge, this proof technique has not been used in prior constructions of LWE based FE systems, which mostly rely on a "punctured trapdoor" approach. This approach roughly provides the simulator with a trapdoor that can be used to answer key requests but vanishes w.h.p for the challenge. Our simulator does not use trapdoors, but relies on an argument about entropy leakage as described above. We now proceed with the formal proof.

Theorem 2. *Assume that $\ell \leq n^{O(1)}$, $m \geq 4n \log_2 q$, $q > \ell K^2$ and τ is as described above. Then the functional encryption scheme above is fully secure, under the $\text{mheLWE}_{q,\alpha,m,\ell,\tau}$ hardness assumption.*

Proof. The proof proceeds with a sequence of games that starts with the real game and ends with a game in which the adversary's advantage is negligible. For each i , we call S_i the event that the adversary wins in Game i .

Game 0: This is the genuine full security game. Namely: the adversary \mathcal{A} is given the master public key mpk ; in the challenge phase, adversary \mathcal{A} comes

up with two distinct vectors $\mathbf{y}_0, \mathbf{y}_1 \in \mathcal{P}$ and receives an encryption C of \mathbf{y}_β for $\beta \leftarrow \{0, 1\}$ sampled by the challenger; when \mathcal{A} halts, it outputs $\beta' \in \{0, 1\}$ and S_0 is the event that $\beta' = \beta$. Note that any vector $\mathbf{x} \in \mathcal{V}$ queried by \mathcal{A} to the secret key extraction oracle must satisfy $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle$ over \mathbb{Z} if \mathcal{A} is a legitimate adversary.

Game 1: We modify the generation of $C = (\mathbf{c}_0, \mathbf{c}_1)$ in the challenge phase. Namely, at the outset of the game, the challenger picks $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}, \alpha q}^m$ (which may be chosen ahead of time) as well as $\mathbf{Z} \leftarrow \tau$. The master public key mpk is computed by setting $\mathbf{U} = \mathbf{Z} \cdot \mathbf{A} \bmod q$. In the challenge phase, the challenger picks a random bit $\beta \leftarrow \{0, 1\}$ and encrypts \mathbf{y}_β by computing (modulo q)

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_0, \\ \mathbf{c}_1 &= \mathbf{Z} \cdot \mathbf{c}_0 - \mathbf{Z} \cdot \mathbf{e}_0 + \mathbf{e}_1 + \lfloor q/K \rfloor \cdot \mathbf{y}_\beta, \end{aligned}$$

with $\mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha q}^\ell$. As the distribution of C is the same as in Game 0, we have $\Pr[S_1] = \Pr[S_0]$.

Game 2: We modify again the generation of $C = (\mathbf{c}_0, \mathbf{c}_1)$ in the challenge phase. Namely, the challenger picks $\mathbf{u} \leftarrow \mathbb{Z}_q^m$, sets $\mathbf{c}_0 = \mathbf{u}$ and computes \mathbf{c}_1 using \mathbf{c}_0, \mathbf{Z} and \mathbf{e}_0 as in Game 1.

Under the mheLWE hardness assumption with $t = \ell$, this modification has no noticeable effect on the behavior of \mathcal{A} . Below, we prove that $\Pr[S_2] \approx 1/2$, which completes the proof of the theorem.

Let $\mathbf{x}^i \in \mathcal{V}$ be the vectors corresponding to the secret key queries made by \mathcal{A} . As \mathcal{A} is a legitimate adversary, we have $\langle \mathbf{x}^i, \mathbf{y}_0 \rangle = \langle \mathbf{x}^i, \mathbf{y}_1 \rangle$ over \mathbb{Z} for each secret key query \mathbf{x}^i . Let $g \neq 0$ be the gcd of the coefficients of $\mathbf{y}_1 - \mathbf{y}_0$ and define $\mathbf{y} = (y_1, \dots, y_\ell) = \frac{1}{g}(\mathbf{y}_1 - \mathbf{y}_0)$. We have that $\langle \mathbf{x}^i, \mathbf{y} \rangle = 0$ (over \mathbb{Z}) for all i . Consider the lattice $\{\mathbf{x} \in \mathbb{Z}^\ell : \langle \mathbf{x}, \mathbf{y} \rangle = 0\}$: all the queries \mathbf{x}^i must belong to that lattice. Without loss of generality, we assume the n_0 first entries of \mathbf{y} are zero (for some n_0), and all remaining entries are non-zero. Further, the rows of the following matrix form a basis of a full-dimensional sublattice:

$$\mathbf{X}_{top} = \left(\begin{array}{c|cccccc} \mathbf{I}_{n_0} & & & & & & \\ \hline & -y_{n_0+2} & y_{n_0+1} & & & & \\ & & -y_{n_0+3} & y_{n_0+2} & & & \\ & & & \ddots & \ddots & & \\ & & & & & -y_\ell & y_{\ell-1} \end{array} \right) \in \mathbb{Z}^{(\ell-1) \times \ell}. \quad (4.1)$$

We may assume that through the secret key queries, the adversary learns exactly $\mathbf{X}_{top} \mathbf{Z}$, as all the queried vectors \mathbf{x}^i can be obtained as rational combinations of the rows of \mathbf{X}_{top} .

Let $\mathbf{X}_{bot} = \mathbf{y}^T \in \mathbb{Z}^{1 \times \ell}$. Consider the matrix $\mathbf{X} \in \mathbb{Z}_q^{\ell \times \ell}$ obtained by putting \mathbf{X}_{top} on top of \mathbf{X}_{bot} . We claim that \mathbf{X} is invertible modulo q . To see this, observe

that

$$\mathbf{X}\mathbf{X}^T = \left(\begin{array}{c|ccc|c} \mathbf{I}_{n_0} & & & & \\ \hline & y_{n_0+1}^2 + y_{n_0+2}^2 & -y_{n_0+1} \cdot y_{n_0+3} & & \\ & -y_{n_0+1} \cdot y_{n_0+3} & y_{n_0+2}^2 + y_{n_0+3}^2 & \ddots & \\ & & \ddots & \ddots & \ddots \\ & & & -y_{\ell-2} \cdot y_{\ell} & y_{\ell-1}^2 + y_{\ell}^2 \\ \hline & & & & \|\mathbf{y}\|^2 \end{array} \right)$$

It can be proved by induction that its determinant is

$$\det(\mathbf{X}\mathbf{X}^T) = \left(\prod_{k=n_0+2}^{\ell-1} y_k^2 \right) \cdot \|\mathbf{y}\|^4.$$

As each of the y_k 's is small and non-zero, they are all non-zero modulo prime q . Similarly, the integer $(\sum_{k=n_0+1}^{\ell} y_k^2)$ is non-zero and $< \ell P^2 < q$. This shows that $(\det \mathbf{X})^2 \neq 0 \pmod q$, which implies that \mathbf{X} is invertible modulo q .

In Game 2, we have $\mathbf{c}_1 = \mathbf{Z}\mathbf{u} - \mathbf{f} + \lfloor q/K \rfloor \cdot \mathbf{y}_\beta$, with $\mathbf{f} := -\mathbf{Z}\mathbf{e}_0 + \mathbf{e}_1$. We write:

$$\mathbf{c}_1 = \mathbf{X}^{-1} \cdot \mathbf{X} \cdot (\mathbf{Z}\mathbf{u} - \mathbf{f} + \lfloor q/K \rfloor \cdot \mathbf{y}_\beta) \pmod q.$$

We will show that the distribution of $\mathbf{X} \cdot \mathbf{c}_1 \pmod q$ is (almost) independent of β . As \mathbf{X} is (almost) independent of β and invertible over \mathbb{Z}_q , this implies that the distribution of \mathbf{c}_1 is (almost) independent of β and $\Pr[S_2] \approx 1/2$.

The first $\ell - 1$ entries of $\mathbf{X} \cdot \mathbf{c}_1$ do not depend on β because $\mathbf{X}_{top} \cdot \mathbf{y}_0 = \mathbf{X}_{top} \cdot \mathbf{y}_1 \pmod q$.

It remains to prove that the last entry of $\mathbf{X} \cdot \mathbf{c}_1 \pmod q$ is (almost) independent of β . For this, we show that the residual distribution of $\mathbf{X}_{bot}\mathbf{Z}$ given the tuple $(\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}\mathbf{Z})$ has high entropy. Using (a variant of) the leftover hash lemma with randomness $\mathbf{X}_{bot}\mathbf{Z}$ and seed \mathbf{u} , we will then conclude that given $(\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}\mathbf{Z})$, the pair $(\mathbf{u}, \mathbf{X}_{bot}\mathbf{Z}\mathbf{u})$ is close to uniform and hence statistically hides $\lfloor q/K \rfloor \cdot \mathbf{y}_\beta$ in \mathbf{c}_1 .

Write $\mathbf{A} = (\mathbf{A}_1^T | \mathbf{A}_2^T)^T$ with $\mathbf{A}_1, \mathbf{A}_2 \in \mathbb{Z}_q^{(m/2) \times n}$. Similarly, write $\mathbf{Z} = (\mathbf{Z}_1 | \mathbf{Z}_2)$ with $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathbb{Z}_q^{\ell \times (m/2)}$. Recall that by construction, every entry of \mathbf{Z}_1 is independently sampled from a zero-centered integer Gaussian of standard deviation parameter $\sigma_1 = \Theta(\sqrt{n \log m} \max(\sqrt{m}, K))$. Further, every entry of \mathbf{Z}_2 is independently sampled from a (not zero-centered) integer Gaussian of standard deviation parameter σ_2 that is larger than σ_1 .

Lemma 1. *Conditioned on $(\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}\mathbf{Z}_1)$, the row vector $\mathbf{X}_{bot}\mathbf{Z}_1$ is distributed as $\mathbf{c} + D_{\|\mathbf{y}\|^2 \mathbb{Z}^{m/2}, \|\mathbf{y}\|_{\sigma_1, -\mathbf{c}}}$ for some vector \mathbf{c} that depends only on $\mathbf{X}_{top}\mathbf{Z}_1$.*

Proof. Thanks to Lemma 9 in Appendix C, we have that $\mathbf{Z}_2\mathbf{A}_2$ is within $2^{-\Omega(n)}$ statistical distance to uniform. It hence statistically hides the term $\mathbf{Z}_1\mathbf{A}_1$ in

$\mathbf{Z}\mathbf{A} = \mathbf{Z}_1\mathbf{A}_1 + \mathbf{Z}_2\mathbf{A}_2$, and we obtain that given $(\mathbf{A}, \mathbf{Z}\mathbf{A})$, the distribution of each entry of \mathbf{Z}_1 is still $D_{\mathbb{Z}, \sigma_1}$.

Note that in $\mathbf{X}_{top}\mathbf{Z}_1$ and $\mathbf{X}_{bot}\mathbf{Z}_1$, matrices \mathbf{X}_{top} and \mathbf{X}_{bot} act in parallel on the columns of \mathbf{Z}_1 . To prove the claim, it suffices to consider the distribution of $\mathbf{X}_{bot}\mathbf{z}$ conditioned on $\mathbf{X}_{top}\mathbf{z}$, with \mathbf{z} sampled from $D_{\mathbb{Z}^\ell, \sigma_1}$. Let $\mathbf{b} = \mathbf{X}_{top}\mathbf{z} \in \mathbb{Z}^{\ell-1}$ and fix $\mathbf{z}_0 \in \mathbb{Z}^\ell$ arbitrary such that $\mathbf{b} = \mathbf{X}_{top}\mathbf{z}_0$. The distribution of \mathbf{z} given that $\mathbf{X}_{top}\mathbf{z} = \mathbf{b}$ is $\mathbf{z}_0 + D_{\Lambda, \sigma_1, -\mathbf{z}_0}$, with $\Lambda = \{\mathbf{x} \in \mathbb{Z}^\ell : \mathbf{X}_{top}\mathbf{x} = \mathbf{0}\}$. By construction of \mathbf{X} , we have that $\Lambda = \mathbb{Z}\mathbf{y}$. As a result, the conditional distribution of $\mathbf{X}_{bot}\mathbf{z}$ is $c + D_{\|\mathbf{y}\|^2\mathbb{Z}, \|\mathbf{y}\|_{\sigma_1, -c}}$ with $c = \langle \mathbf{y}, \mathbf{z}_0 \rangle \in \mathbb{Z}$. \square

Now, let us write $\mathbf{u} = (\mathbf{u}_1^T | \mathbf{u}_2^T)^T$ where $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{Z}_q^{m/2}$. We have $\mathbf{X}_{bot}\mathbf{Z}\mathbf{u} = \mathbf{X}_{bot}\mathbf{Z}_1\mathbf{u}_1 + \mathbf{X}_{bot}\mathbf{Z}_2\mathbf{u}_2$. Thanks to the claim above and to Lemma 9, we obtain that the distribution of $(\mathbf{u}_1, \langle D_{\|\mathbf{y}\|^2\mathbb{Z}^{m/2}, \|\mathbf{y}\|_{\sigma_1, -c}}, \mathbf{u}_1 \rangle)$ is within $2^{-\Omega(n)}$ statistical distance to uniform (note that $D_{\|\mathbf{y}\|^2\mathbb{Z}^{m/2}, \|\mathbf{y}\|_{\sigma_1, -c}} = \|\mathbf{y}\|^2 \cdot D_{\mathbb{Z}^{m/2}, \sigma_1 / \|\mathbf{y}\|, -c / \|\mathbf{y}\|^2}$, that $\|\mathbf{y}\|^2$ is invertible modulo q and that $\sigma_1 / \|\mathbf{y}\|$ satisfies the assumption of Lemma 9). This implies that given $(\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}\mathbf{Z})$, the pair $(\mathbf{u}, \mathbf{X}_{bot}\mathbf{Z}\mathbf{u})$ is close to uniform, which completes the security proof. \square

4.2 Inner products modulo a prime p

We now modify the LWE-based scheme above so that it enables secure functional encryption for inner products modulo prime p . The plaintext and key vectors now belong to \mathbb{Z}_p^ℓ .

Note that the prior scheme evaluates inner products over the integers and is insecure if ported as is to the modulo p setting. To see this, consider the following simple attack in which the adversary requests a single key \mathbf{x} so that integer inner product with the challenge messages \mathbf{y}_0 and \mathbf{y}_1 are different by a multiple of p . Since the functionality posits that the inner product evaluations only agree modulo p , this is an admissible query. However, since decryption is performed over \mathbb{Z}_q with q much larger than p , the adversary can easily distinguish. To prevent this attack, we scale the encrypted message by a factor of q/p (instead of $\lfloor q/K \rfloor$ as in the previous scheme): decryption modulo q forces arithmetic modulo p on the underlying plaintext.

A related difficulty in adapting the previous LWE-based scheme to modular inner products is the distribution of the noise component after inner product evaluation. Ciphertexts are manipulated modulo q , which internally manipulates plaintexts modulo p . If implemented naively, the carries of the plaintext computations may spill outside of the plaintext slots and bias the noise components of the ciphertexts. This may result in distinguishing attacks. To handle this, we take q a multiple of p . This adds some technical complications, as \mathbb{Z}_q is hence not a field anymore.

A different attack is that the adversary may request keys for vectors that are linearly dependent modulo p but linearly independent over the integers. Note that with ℓ such queries, the attacker can recover the master secret key. To prevent this attack, we modify the scheme in that the authority is now stateful

and keeps a record of all key queries made so far, so that it can make sure that key queries that are linearly dependent modulo p remain so modulo q . We also take q a power of p to simplify the implementation of this idea.

We note that for our application to bounded query FE for all circuits, all queries will be linearly independent modulo p , hence we will not require a stateful keygen. For details, see Section 6.

We now describe our scheme for inner products modulo p .

Setup($1^n, 1^\ell, p$): Set integers $m, q = p^k$ for some integer k , real $\alpha \in (0, 1)$ and distribution τ over $\mathbb{Z}^{\ell \times m}$ as explained below. Sample $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}$ and $\mathbf{Z} \leftarrow \tau$. Compute $\mathbf{U} = \mathbf{Z} \cdot \mathbf{A} \in \mathbb{Z}_q^{\ell \times n}$. Define $\text{mpk} := (\mathbf{A}, \mathbf{U})$ and $\text{msk} := \mathbf{Z}$.

Keygen($\text{msk}, \mathbf{x}, \text{st}$): Given a vector $\mathbf{x} \in \mathbb{Z}_p^\ell$, and an internal state st , compute the secret key $\mathbf{z}_\mathbf{x}$ as follows. Recall that **Keygen** is a stateful algorithm with empty initial State st . At any point in the scheme execution, State st contains at most ℓ tuples $(\mathbf{x}_i, \bar{\mathbf{x}}_i, \mathbf{z}_i)$ where the \mathbf{x}_i 's are (a subset of the) key queries that have been made so far, and the $(\bar{\mathbf{x}}_i, \mathbf{z}_i)$'s are the corresponding secret keys. If \mathbf{x} is linearly independent from the \mathbf{x}_i 's modulo p , set $\bar{\mathbf{x}} = \mathbf{x} \in \mathbb{Z}^\ell$ (with coefficients in $[0, p)$), $\mathbf{z}_\mathbf{x} = \bar{\mathbf{x}}^T \cdot \mathbf{Z} \in \mathbb{Z}^m$ and add $(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}_\mathbf{x})$ to st . If $\mathbf{x} = \sum_i k_i \mathbf{x}_i \pmod p$ for some k_i 's in $[0, p)$, then set $\bar{\mathbf{x}} = \sum_i k_i \bar{\mathbf{x}}_i \in \mathbb{Z}^\ell$ and $\mathbf{z}_\mathbf{x} = \sum_i k_i \mathbf{z}_i \in \mathbb{Z}^m$. In both cases, return $(\bar{\mathbf{x}}, \mathbf{z}_\mathbf{x})$.

Encrypt(mpk, \mathbf{y}): To encrypt a vector $\mathbf{y} \in \mathbb{Z}_p^\ell$, sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}, \alpha q}^m$ and $\mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha q}^\ell$ and compute

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{A} \cdot \mathbf{s} + \mathbf{e}_0 \in \mathbb{Z}_q^m, \\ \mathbf{c}_1 &= \mathbf{U} \cdot \mathbf{s} + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y} \in \mathbb{Z}_q^\ell. \end{aligned}$$

Then, return $C := (\mathbf{c}_0, \mathbf{c}_1)$.

Decrypt($\text{mpk}, (\bar{\mathbf{x}}, \mathbf{z}_\mathbf{x}), C$): Given $C := (\mathbf{c}_0, \mathbf{c}_1)$ and a secret key $(\bar{\mathbf{x}}, \mathbf{z}_\mathbf{x})$ for $\mathbf{x} \in \mathbb{Z}_p^\ell$, compute $\mu' = \langle \bar{\mathbf{x}}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_\mathbf{x}, \mathbf{c}_0 \rangle \pmod q$ and output the value $\mu \in \mathbb{Z}_p$ that minimizes $|p^{k-1} \cdot \mu - \mu'|$.

Decryption correctness. Correctness derives from the following observation:

$$\begin{aligned} \mu' &= \langle \bar{\mathbf{x}}, \mathbf{c}_1 \rangle - \langle \mathbf{z}_\mathbf{x}, \mathbf{c}_0 \rangle \\ &= p^{k-1} \cdot (\langle \mathbf{x}, \mathbf{y} \rangle \pmod p) + \langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_\mathbf{x}, \mathbf{e}_0 \rangle \pmod q. \end{aligned}$$

By adapting the proof of the first LWE-based scheme, we can show that the magnitude of the term $\langle \bar{\mathbf{x}}, \mathbf{e}_1 \rangle - \langle \mathbf{z}_\mathbf{x}, \mathbf{e}_0 \rangle$ is $\leq \ell^2 p^2 B_\tau \alpha q \omega(\sqrt{\log n})$ with probability $\geq 1 - n^{-\omega(1)}$. This follows from the bound $\|\mathbf{z}_\mathbf{x}\| \leq \ell \|\bar{\mathbf{x}}\| \leq \ell^2 p^2 B_\tau$.

Setting the parameters. The main difference with the previous LWE-based scheme with respect to parameter conditions is the choice of q of the form $q = p^k$ instead of q prime. As explained just above, correctness may be ensured by setting

$$\alpha^{-1} \geq \ell^2 p^3 B_\tau \omega(\sqrt{\log n}) \quad \text{and} \quad q \geq \alpha^{-1} \omega(\sqrt{\log n}).$$

The choice of τ is driven by Lemma 2 below (the proof requires that σ_1 is large) and the reduction from LWE to mheLWE (as summarized in Theorem 4), and more precisely from Lemma 4. We may choose $\tau = D_{\mathbb{Z}, \sigma_1}^{\ell \times m/2} \times (D_{\mathbb{Z}^{m/2}, \sigma_2, \delta_1} \times \dots \times D_{\mathbb{Z}^{m/2}, \sigma_2, \delta_\ell})$, where $\delta_i \in \mathbb{Z}^\ell$ denotes the i th canonical vector, and the standard deviation parameters satisfy $\sigma_1 = \Theta(\sqrt{n \log m} \max(\sqrt{m}, K'))$ and $\sigma_2 = \Theta(n^{7/2} m^{1/2} \max(m, K'^2) \log^{5/2} m)$, with $K' = (\sqrt{\ell} p)^\ell$.

To ensure security based on $\text{LWE}_{q, \alpha', m}$ in dimension $\geq c \cdot n$ for some $c \in (0, 1)$ via Theorems 2 and 4 below, one may further impose that $\ell \leq (1 - c) \cdot n$ and $m = \Theta(n \log q)$, to obtain $\alpha' = \Omega(\alpha / (n^6 K' \log^2 q \log^{5/2} n))$. Remember that $\text{LWE}_{q, \alpha', m}$ enjoys reductions from lattice problems when $q \geq \Omega(\sqrt{n}/\alpha')$.

Note that the parameter conditions make the scheme efficiency degrade quickly when ℓ increases, as K' is exponential in ℓ . Assume that $p \leq n^{O(1)}$ and $\ell = \Omega(\log n)$. Then $\sigma_1, \sigma_2, 1/\alpha, 1/\alpha'$ and q can all be set as $2^{\tilde{O}(\ell)}$. To maintain security against all $2^{o(\lambda)}$ attacks, one may set $n = \tilde{\Theta}(\ell \lambda)$.

Theorem 3. *Assume that $\ell \leq n^{O(1)}$, $m \geq 4n \log_2 q$ and τ is as described above. Then the stateful functional encryption scheme above is fully secure, under the $\text{mheLWE}_{q, \alpha, m, \ell, \tau}$ hardness assumption.*

Proof. The sequence of games in the proof of Theorem 2 can be adapted to the modified scheme. The main difficulty is to show that in the adapted version of the last game, the winning probability is close to $1/2$. Let us recall that game in details.

Game 2': At the outset of the game, the challenger picks $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e}_0 \leftarrow D_{\mathbb{Z}, \alpha q}^m$ as well as $\mathbf{Z} \leftarrow \tau$. The master public key mpk is computed by setting $\mathbf{U} = \mathbf{Z} \cdot \mathbf{A} \bmod q$ and is provided to the adversary. In the challenge phase, adversary \mathcal{A} comes up with two distinct vectors $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_p^\ell$. The challenger picks a random bit $\beta \leftarrow \{0, 1\}$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$ and encrypts \mathbf{y}_β by computing (modulo q)

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{u}, \\ \mathbf{c}_1 &= \mathbf{Z} \cdot \mathbf{c}_0 - \mathbf{Z} \cdot \mathbf{e}_0 + \mathbf{e}_1 + p^{k-1} \cdot \mathbf{y}_\beta, \end{aligned}$$

with $\mathbf{e}_1 \leftarrow D_{\mathbb{Z}, \alpha q}^\ell$. Note that any vector $\mathbf{x} \in \mathbb{Z}_p^\ell$ queried by \mathcal{A} to the secret key extraction oracle must satisfy $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle \bmod p$ if \mathcal{A} is a legitimate adversary. Adversary \mathcal{A} is then given a secret key $(\bar{\mathbf{x}}, \mathbf{z}_\mathbf{x})$ as in the real scheme. When \mathcal{A} halts, it outputs $\beta' \in \{0, 1\}$ and wins in the event that $\beta' = \beta$.

Define $\mathbf{y} = \mathbf{y}_1 - \mathbf{y}_0 \in \mathbb{Z}_p^\ell$. Let $\mathbf{x}_i \in \mathbb{Z}_p^\ell$ be the vectors corresponding to the secret key queries made by \mathcal{A} . As \mathcal{A} is a legitimate adversary, we have $\langle \mathbf{x}_i, \mathbf{y} \rangle = 0 \bmod p$ for each secret key query \mathbf{x}_i .

We consider the view of the adversary after it has made exactly j key queries that are linearly independent modulo p , for each j from 0 up to $\ell - 1$. In fact, counter j may stop increasing before reaching $\ell - 1$, but without loss of generality, we may assume that it eventually reaches $\ell - 1$. We are to show by induction that for any j , the view of the adversary is almost independent of β . In particular, for

all $j < \ell - 1$, this implies that the $(j + 1)$ th linearly independent key query is almost (statistically) independent of β . It also implies, for $j = \ell - 1$, that the adversary's view through Game 2' is almost independent of β , which is exactly what we are aiming for. In what follows, we take $j \in \{0, \dots, \ell - 1\}$, and assume that state st is independent from β . We also assume that the j th private key query occurs after the challenge phase since the adversary's view is trivially independent of β before the generation of the challenge ciphertext.

At this stage, the state st contains exactly j tuples $(\mathbf{x}_i, \bar{\mathbf{x}}_i, \mathbf{z}_i)$, where the vectors $\{\mathbf{x}_i\}_{i=1}^j$ form a \mathbb{Z}_p -basis of a subspace of the $(\ell - 1)$ -dimensional vector space $\mathbf{y}^\perp := \{\mathbf{x} \in \mathbb{Z}_p^\ell : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod p\}$. From \mathbf{y} , we deterministically extend $\{\mathbf{x}_i\}_{i=1}^j$ into a basis of \mathbf{y}^\perp that is statistically independent of β . A way to interpret this is to imagine that the challenger makes dummy private key queries $\{\mathbf{x}_i\}_{i=j+1}^{\ell-1}$ for itself so as to get a full basis of \mathbf{y}^\perp and creates the corresponding $\{\bar{\mathbf{x}}_i\}_{i=j+1}^{\ell-1}$ in \mathbb{Z}^ℓ . We define $\mathbf{X}_{top} \in \mathbb{Z}^{(\ell-1) \times \ell}$ as the matrix whose i th row is $\bar{\mathbf{x}}_i$ for all i , including the genuine and dummy keys. Through the secret key queries, the adversary learns at most $\mathbf{X}_{top} \mathbf{Z} \in \mathbb{Z}^{(\ell-1) \times m}$.

Let $\mathbf{x}' \in \mathbb{Z}_p^\ell$ be a vector that does not belong to \mathbf{y}^\perp , and $\mathbf{X}_{bot} \in \mathbb{Z}^{1 \times \ell}$ be the canonical lift of $(\mathbf{x}')^T$ over the integers. Consider the matrix $\mathbf{X} \in \mathbb{Z}^{\ell \times \ell}$ obtained by putting \mathbf{X}_{top} on top of \mathbf{X}_{bot} . By construction, the matrix \mathbf{X} is invertible modulo p , and hence modulo $q = p^k$. Also, by induction and construction, $\mathbf{X} \in \mathbb{Z}^{\ell \times \ell}$ is statistically independent of $\beta \in \{0, 1\}$.

In Game 2', we have $\mathbf{c}_1 = \mathbf{Z}\mathbf{u} - \mathbf{f} + p^{k-1} \cdot \mathbf{y}_\beta$, with $\mathbf{f} := -\mathbf{Z}\mathbf{e}_0 + \mathbf{e}_1$. We write:

$$\mathbf{c}_1 = \mathbf{X}^{-1} \cdot \mathbf{X} \cdot (\mathbf{Z}\mathbf{u} - \mathbf{f} + p^{k-1} \cdot \mathbf{y}_\beta) \pmod q.$$

We will show that the distribution of $\mathbf{X} \cdot \mathbf{c}_1 \pmod q$ is (almost) independent of β . As the matrix \mathbf{X} is independent of $\beta \in \{0, 1\}$ and invertible over \mathbb{Z}_q , this implies that the distribution of \mathbf{c}_1 is statistically independent of β (recall that \mathbf{X} is information-theoretically known to \mathcal{A} , which means that, if \mathbf{c}_1 carries any noticeable information on β , so does $\mathbf{X} \cdot \mathbf{c}_1 \pmod q$). This ensures that the winning probability in Game 2' is negligibly far from $1/2$.

First, the first $\ell - 1$ entries of $\mathbf{X} \cdot \mathbf{c}_1$ do not depend on β because we have the equality $p^{k-1} \cdot \mathbf{X}_{top} \cdot \mathbf{y}_0 = p^{k-1} \cdot \mathbf{X}_{top} \cdot \mathbf{y}_1 \pmod q$ by construction of \mathbf{X}_{top} .

It remains to prove that the last entry of $\mathbf{X} \cdot \mathbf{c}_1 \pmod q$ is (almost) independent of β . Let us write $\mathbf{A} = (\mathbf{A}_1^T | \mathbf{A}_2^T)^T$ with $\mathbf{A}_1, \mathbf{A}_2 \in \mathbb{Z}_q^{(m/2) \times n}$. Similarly, we also write $\mathbf{Z} = (\mathbf{Z}_1 | \mathbf{Z}_2)$ with $\mathbf{Z}_1, \mathbf{Z}_2 \in \mathbb{Z}^{\ell \times (m/2)}$. Recall that by construction, every entry of \mathbf{Z}_1 is independently sampled from a zero-centered integer Gaussian of standard deviation parameter $\sigma_1 = \Theta(\sqrt{n \log m} \max(\sqrt{m}, K'))$ with $K' = (\sqrt{\ell p})^\ell$. Further, every entry of \mathbf{Z}_2 is independently sampled from a (not zero-centered) integer Gaussian of standard deviation parameter σ_2 that is larger than σ_1 .

Lemma 2. *Conditioned on $(\mathbf{A}, \mathbf{Z}\mathbf{A}, \mathbf{X}_{top}\mathbf{Z}_1)$, the row vector $\mathbf{X}_{bot}\mathbf{Z}_1 \pmod p$ is within negligible statistical distance from the uniform distribution over $\mathbb{Z}_p^{m/2}$.*

Proof. Thanks to Lemma 9, we have that $\mathbf{Z}_2\mathbf{A}_2$ is within $2^{-\Omega(n)}$ statistical distance to uniform over $\mathbb{Z}_q^{(\ell-1) \times m}$. It hence statistically hides the term $\mathbf{Z}_1\mathbf{A}_1$ in

$\mathbf{ZA} = \mathbf{Z}_1\mathbf{A}_1 + \mathbf{Z}_2\mathbf{A}_2 \pmod q$, and we obtain that given $(\mathbf{A}, \mathbf{ZA})$, the distribution of each entry of \mathbf{Z}_1 is still $D_{\mathbb{Z}, \sigma_1}$.

Note that in $\mathbf{X}_{top}\mathbf{Z}_1$ and $\mathbf{X}_{bot}\mathbf{Z}_1$, matrices \mathbf{X}_{top} and \mathbf{X}_{bot} act in parallel on the columns of \mathbf{Z}_1 . To prove the claim, it suffices to consider the distribution of $\mathbf{X}_{bot}\mathbf{z}$ conditioned on $\mathbf{X}_{top}\mathbf{z}$, with \mathbf{z} sampled from $D_{\mathbb{Z}^\ell, \sigma_1}$. Let $\mathbf{b} = \mathbf{X}_{top}\mathbf{z} \in \mathbb{Z}^{\ell-1}$ and fix $\mathbf{z}_0 \in \mathbb{Z}^\ell$ arbitrary such that $\mathbf{b} = \mathbf{X}_{top}\mathbf{z}_0$. The distribution of \mathbf{z} given that $\mathbf{X}_{top}\mathbf{z} = \mathbf{b}$ is $\mathbf{z}_0 + D_{\Lambda, \sigma_1, -\mathbf{z}_0}$, with $\Lambda = \{\mathbf{x} \in \mathbb{Z}^\ell : \mathbf{X}_{top}\mathbf{x} = \mathbf{0}\}$ (where the equality holds over the integers). Note that Λ is a 1-dimensional lattice in \mathbb{Z}^ℓ .

We can write $\Lambda = \mathbf{y}' \cdot \mathbb{Z}$, for some $\mathbf{y}' \in \mathbb{Z}^\ell$. Note that there exists $\alpha \in \mathbb{Z}_p \setminus \{0\}$ such that $\mathbf{y}' = \alpha \cdot \mathbf{y} \pmod p$: otherwise, the vector \mathbf{y}'/p would belong to $\Lambda \setminus \mathbf{y}' \cdot \mathbb{Z}$, contradicting the definition of \mathbf{y}' . Further, we have $\|\mathbf{y}'\| = \det \Lambda \leq \det \Lambda'$, where Λ' is the lattice spanned by the rows of \mathbf{X}_{top} (see, e.g., [47], for properties on orthogonal lattices). Hadamard's bound implies that $\|\mathbf{y}'\| \leq (\sqrt{\ell}p)^{\ell-1}$.

By Lemma 7, the fact that $\sigma_1 \geq \sqrt{n}(\sqrt{\ell}p)^\ell$ implies that the distribution $(D_{\Lambda, \sigma_1, -\mathbf{z}_0} \pmod p\Lambda)$ is within $2^{-\Omega(n)}$ statistical distance from the uniform distribution over $\Lambda/p\Lambda \simeq \mathbf{y}\mathbb{Z}_p$. We conclude that the conditional distribution of $(\mathbf{X}_{bot}\mathbf{z} \pmod p)$ is within exponentially small statistical distance from the uniform distribution over \mathbb{Z}_p (here we use the facts that p is prime and that $\mathbf{X}_{bot}\mathbf{y} \neq 0 \pmod p$, by construction of \mathbf{X}_{bot}). \square

Now, if we write $\mathbf{u} = (\mathbf{u}_1^T | \mathbf{u}_2^T)^T$ for vectors $\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{Z}_q^{m/2}$, we have $\mathbf{X}_{bot}\mathbf{Z}\mathbf{u} = \mathbf{X}_{bot}\mathbf{Z}_1\mathbf{u}_1 + \mathbf{X}_{bot}\mathbf{Z}_2\mathbf{u}_2$. Thanks to Lemmas 2 and 10 (a variant of the leftover hash lemma modulo $q = p^k$), we obtain that conditioned on $(\mathbf{A}, \mathbf{ZA}, \mathbf{X}_{top}\mathbf{Z})$, the distribution of $(\mathbf{u}_1, \mathbf{X}_{bot}\mathbf{Z}_1\mathbf{u}_1)$ is within $2^{-\Omega(n)}$ statistical distance to uniform modulo q (here we used the assumption that $m \geq k + n/(\log p)$). This implies that given $(\mathbf{A}, \mathbf{ZA}, \mathbf{X}_{top}\mathbf{Z})$, the pair $(\mathbf{u}, \mathbf{X}_{bot}\mathbf{Z}\mathbf{u})$ is close to uniform, which completes the security proof. \square

4.3 Hardness of multi-hint extended-LWE

In this section, we prove the following theorem, which shows that for some parameters, the mheLWE problem is no easier than the LWE problem.

Theorem 4. *Let $n \geq 100$, $q \geq 2$, $t < n$ and m with $m = \Omega(n \log n)$ and $m \leq n^{O(1)}$. There exists $\xi \leq O(n^4 m^2 \log^{5/2} n)$ and a distribution τ over $\mathbb{Z}^{t \times m}$ such that the following statements hold:*

- *There is a reduction from $\text{LWE}_{q, \alpha, m}$ in dimension $n - t$ to $\text{mheLWE}_{q, \alpha \xi, m, t, \tau}$ that reduces the advantage by at most $2^{\Omega(t-n)}$,*
- *It is possible to sample from τ in time polynomial in n ,*
- *Each entry of matrix τ is an independent discrete Gaussian $\tau_{i,j} = D_{\mathbb{Z}, \sigma_{i,j}, \mathbf{c}_{i,j}}$ for some $\mathbf{c}_{i,j}$ and $\sigma_{i,j} \geq \Omega(\sqrt{mn \log m})$,*
- *With probability $\geq 1 - n^{-\omega(1)}$, all rows from a sample from τ have norms $\leq \xi$.*

Our reduction from LWE to mheLWE proceeds as the reduction from LWE to extended-LWE from [19], using the matrix gadget from [44] to handle the

multiple hints. We first reduce LWE to the following variant of LWE in which the first samples are noise-free. This problem generalizes the first-is-errorless LWE problem from [19].

Definition 5 (First-are-errorless LWE). *Let q, α, m, t be functions of a parameter n . The first-are-errorless LWE problem $\text{faelLWE}_{q, \alpha, m, t}$ is defined as follows: For $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, the goal is to distinguish between the following two scenarios. In the first, all m samples are uniform over $\mathbb{Z}_q^n \times \mathbb{Z}_q$. In the second, the first t samples are from $A_{q, \{0\}, \mathbf{s}}$ (where $\{0\}$ denotes the distribution that is deterministically zero) and the rest are from $A_{q, \alpha, \mathbf{s}}$.*

Lemma 3. *For any $n > t$, $m, q \geq 2$, and $\alpha \in (0, 1)$, there is an efficient reduction from $\text{LWE}_{q, \alpha, m}$ in dimension $n - t$ to $\text{faelLWE}_{q, \alpha, m, t}$ in dimension n that reduces the advantage by at most 2^{-n+t+1} .*

The proof, postponed to the appendices, is a direct adaptation of the one of [19, Le. 4.3].

In our reduction from faelLWE to mhelLWE , we use the following gadget matrix from [44, Cor. 10]. It generalizes the matrix construction from [19, Claim 4.6].

Lemma 4. *Let n, m_1, m_2 with $100 \leq n \leq m_1 \leq m_2 \leq n^{O(1)}$. Let $\sigma_1, \sigma_2 > 0$ be standard deviation parameters such that $\sigma_1 \geq \Omega(\sqrt{m_1 n \log m_1})$, $m_1 \geq \Omega(n \log(\sigma_1 n))$ and $\sigma_2 \geq \Omega(n^{5/2} \sqrt{m_1} \sigma_1^2 \log^{3/2}(m_1 \sigma_1))$. Let $m = m_1 + m_2$. There exists a probabilistic polynomial time algorithm that given n, m_1, m_2 (in unary) and σ_1, σ_2 as inputs, outputs $\mathbf{G} \in \mathbb{Z}^{m \times m}$ such that:*

- *The top $n \times m$ submatrix of \mathbf{G} is within statistical distance $2^{-\Omega(n)}$ of $\tau = D_{\mathbb{Z}, \sigma_1}^{n \times m_1} \times (D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_1} \times \dots \times D_{\mathbb{Z}^{m_2}, \sigma_2, \delta_n})^T$ with δ_i denoting the i th canonical unit vector,*
- *We have $|\det(\mathbf{G})| = 1$ and $\|\mathbf{G}^{-1}\| \leq O(\sqrt{nm_2} \sigma_2)$, with probability $\geq 1 - 2^{-\Omega(n)}$.*

Lemma 5. *Let $n, m_1, m_2, m, \sigma_1, \sigma_2, \tau$ be as in Lemma 4, and $\xi \geq \Omega(\sqrt{nm_2} \sigma_2)$. Let $q \geq 2$, $t \leq n$, $\alpha \geq \Omega(\sqrt{n}/q)$. Let τ_t be the distribution obtained by keeping only the first t rows from a sample from τ . There is a (dimension-preserving) reduction from $\text{faelLWE}_{q, \alpha, m, t}$ to $\text{mhelLWE}_{q, 2\alpha\xi, m, t, \tau_t}$ that reduces the advantage by at most $2^{-\Omega(n)}$.*

Proof. Let us first describe the reduction. Let $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^m \times \mathbb{Z}_q$ be the input, which is either sampled from the uniform distribution, or from distribution $A_{q, \{0\}, \mathbf{s}}^t \times A_{q, \alpha, \mathbf{s}}^{m-t}$ for some fixed $\mathbf{s} \leftarrow \mathbb{Z}_q^n$. Our objective is to distinguish between the two scenarios, using an mhelLWE oracle. We compute \mathbf{G} as in Lemma 4 and let $\mathbf{U} = \mathbf{G}^{-1}$. We let $\mathbf{Z} \in \mathbb{Z}^{t \times m}$ denote the matrix formed by the top t rows of \mathbf{G} , and let $\mathbf{U}' \in \mathbb{Z}^{m \times (m-t)}$ denote the matrix formed by the right $m - t$ columns of \mathbf{U} . By construction, we have $\mathbf{Z}\mathbf{U}' = \mathbf{0}$. We define $\mathbf{A}' = \mathbf{U} \cdot \mathbf{A} \bmod q$. We sample $\mathbf{f} \leftarrow D_{\alpha q(\xi^2 \mathbf{I} - \mathbf{U}'\mathbf{U}'^T)^{1/2}}$ (thanks to Lemma 4 and the choice of ξ , the matrix $\xi^2 \mathbf{I} - \mathbf{U}'\mathbf{U}'^T$ is positive definite). We sample \mathbf{e}' from $\{0\}^t \times D_{\alpha q}^{m-t}$ and

define $\mathbf{b}' = \mathbf{U} \cdot (\mathbf{b} + \mathbf{e}') + \mathbf{f}$. We then sample $\mathbf{c} \leftarrow D_{\mathbb{Z}^m - \mathbf{b}', \sqrt{2}\alpha\xi q}$, and define $\mathbf{h} = \mathbf{Z}(\mathbf{f} + \mathbf{c})$.

Finally, the reduction calls the mheLWE oracle on input $(\mathbf{A}', \mathbf{b}' + \mathbf{c}, \mathbf{Z}, \mathbf{h})$, and outputs the reply.

Correctness is obtained by showing that distribution $A_{q, \{0\}, \mathbf{s}}^t \times A_{q, \alpha, \mathbf{s}}^{m-t}$ is mapped to the mheLWE “LWE” distribution and that the uniform distribution is mapped to the mheLWE “uniform” distribution, up to $2^{-\Omega(n)}$ statistical distances (we do not discuss these tiny statistical discrepancies below). The proof is identical to the reduction analysis in the proof of [19, Le. 4.7]. \square

Theorem 4 is obtained by combining Lemmas 3, 4 and 5.

5 Constructions Based on Paillier

In this section, we show how to remove the main limitation of our DDH-based system which is its somewhat expensive decryption algorithm. To this end, we use Paillier’s cryptosystem [50] and the property that, for an RSA modulus $N = pq$, the multiplicative group $\mathbb{Z}_{N^2}^*$ contains a subgroup of order N (generated by $(N + 1)$) in which the discrete logarithm problem is easy. We also rely on the observation [21, 20] that combining the Paillier and Elgamal encryption schemes makes it possible to decrypt without knowing the factorization of $N = pq$.

5.1 Computing Inner Products over \mathbb{Z}

In the following scheme, key vectors \mathbf{x} and message vectors \mathbf{y} are assumed to be of bounded norm $\|\mathbf{x}\| \leq X$ and $\|\mathbf{y}\| \leq Y$, respectively. The bounds X and Y are chosen so that $X \cdot Y < N$, where N is the composite modulus of Paillier’s cryptosystem. Decryption allows to recover $\langle \mathbf{x}, \mathbf{y} \rangle \bmod N$, which is exactly $\langle \mathbf{x}, \mathbf{y} \rangle$ over the integers, thanks to the norm bounds. The security proof further requires that $\ell Y^2 < N$ and we thus assume $X, Y < (N/\ell)^{1/2}$.

Setup($1^\lambda, 1^\ell, X, Y$): Choose safe prime numbers $p = 2p' + 1$, $q = 2q' + 1$ with sufficiently large primes $p', q' > 2^{l(\lambda)}$, for some polynomial l , and compute $N = pq > XY$. Then, sample $g' \leftarrow \mathbb{Z}_{N^2}^*$ and compute $g = g'^{2N} \bmod N^2$, which generates the subgroup of $(2N)$ th residues in $\mathbb{Z}_{N^2}^*$ with overwhelming probability. Then, sample an integer vector $\mathbf{s} = (s_1, \dots, s_\ell)^T \leftarrow D_{\mathbb{Z}^\ell, \sigma}$ with discrete Gaussian entries of standard deviation $\sigma > \sqrt{\lambda} \cdot N^{5/2}$ and compute $h_i = g^{s_i} \bmod N^2$. Define

$$\text{mpk} := \left(N, g, \{h_i\}_{i=1}^\ell, Y \right)$$

and $\text{msk} := (\{s_i\}_{i=1}^\ell, X)$. The prime numbers p, p', q, q' are no longer needed.

Keygen(msk, \mathbf{x}): To generate a key for the vector $\mathbf{x} = (x_1, \dots, x_\ell) \in \mathbb{Z}^\ell$ with $\|\mathbf{x}\| \leq X$, compute $\text{sk}_{\mathbf{x}} = \sum_{i=1}^\ell s_i \cdot x_i$ over \mathbb{Z} .

Encrypt(mpk, \mathbf{y}): To encrypt a vector $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}^\ell$ with $\|\mathbf{y}\| \leq Y$, sample $r \leftarrow \{0, \dots, \lfloor N/4 \rfloor\}$ and compute

$$\begin{aligned} C_0 &= g^r \bmod N^2, \\ C_i &= (1 + y_i N) \cdot h_i^r \bmod N^2, \quad \forall i \in \{1, \dots, \ell\}. \end{aligned}$$

Return $C_{\mathbf{y}} = (C_0, C_1, \dots, C_\ell) \in \mathbb{Z}_{N^2}^{\ell+1}$.

Decrypt(mpk, $\text{sk}_{\mathbf{x}}$, $C_{\mathbf{y}}$): Given $\text{sk}_{\mathbf{x}} \in \mathbb{Z}$, compute

$$C_{\mathbf{x}} = \left(\prod_{i=1}^{\ell} C_i^{x_i} \right) \cdot C_0^{-\text{sk}_{\mathbf{x}}} \bmod N^2.$$

Then, compute and output $\log_{(1+N)}(C_{\mathbf{x}}) = \frac{C_{\mathbf{x}} - 1 \bmod N^2}{N}$.

As in previous constructions (including those of [2]), our security proof requires inner products to be evaluated over \mathbb{Z} , although the decryptor technically computes $\langle \mathbf{x}, \mathbf{y} \rangle \bmod N$. The reason is that, since secret keys are computed over the integers, our security proof only goes through if the adversary is restricted to only obtain secret keys for vectors \mathbf{x} such that $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle$ over \mathbb{Z} .

Theorem 5. *The scheme provides full security under the DCR assumption. (The proof is available in Appendix F).*

5.2 A Construction for Inner Products over \mathbb{Z}_N

Here, we show that our first scheme can be adapted in order to compute the inner product $\langle \mathbf{y}, \mathbf{x} \rangle \bmod N$ instead of computing it over \mathbb{Z} . To do this, a first difficulty is that, as in our LWE-based system, private keys are computed over the integers and the adversary may query private keys for vectors that are linearly dependent over \mathbb{Z}_N^ℓ but independent over \mathbb{Z}^ℓ . This problem is addressed as previously, by having the authority keep track of all previously revealed private keys. As in our LWE-based construction over \mathbb{Z}_p , we also need to increase the size of private keys (by a factor $\approx \ell$) because we have to use a different information-theoretic argument in the last step of the security proof.

Specifically, in the proof of Theorem 5, we only had to consider the conditional distribution of $\langle \mathbf{s}, \mathbf{y}_0 - \mathbf{y}_1 \rangle \bmod N$, where $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}^\ell$ are the two adversarially-chosen vectors and $\mathbf{s} \in \mathbb{Z}^\ell$ is the master key. Here, we need to make sure that, conditionally on the adversary's view, reducing all coordinates the secret key $\mathbf{s} \in \mathbb{Z}^\ell$ modulo N induces a statistically uniform distribution over \mathbb{Z}_N^ℓ . One of the reasons is that we are no longer guaranteed that $\langle \mathbf{y}_0 - \mathbf{y}_1, \mathbf{y}_0 - \mathbf{y}_1 \rangle \neq 0 \bmod N$.

Setup($1^\lambda, 1^\ell$): Choose safe prime numbers $p = 2p' + 1$, $q = 2q' + 1$ with sufficiently large primes $p', q' > 2^{l(\lambda)}$, for some polynomial l , and compute $N = pq$. Then, sample $g' \leftarrow \mathbb{Z}_{N^2}^*$ and compute $g = g'^{2N} \bmod N^2$, which generates the subgroup of $(2N)$ th residues in $\mathbb{Z}_{N^2}^*$ with overwhelming probability.

Then, sample an integer vector $\mathbf{s} = (s_1, \dots, s_\ell)^T \leftarrow D_{\mathbb{Z}^\ell, \sigma}$ with discrete Gaussian entries of standard deviation $\sigma > \sqrt{\lambda}(\sqrt{\ell}N)^{\ell+1}$ and compute $h_i = g^{s_i} \bmod N^2$. Define

$$\text{mpk} := \left(N, g, \{h_i\}_{i=1}^\ell \right)$$

and $\text{msk} := \{\bar{s}_i\}_{i=1}^\ell$.

Keygen($\text{msk}, \mathbf{x}, \text{st}$): To generate the j th secret key $\text{sk}_{\mathbf{x}}$ for a vector $\mathbf{x} \in \mathbb{Z}_N^\ell$ using the master secret key msk and an (initially empty) internal state st , a stateful algorithm is used. At any time, st contains at most ℓ tuples $(\mathbf{x}_i, \bar{\mathbf{x}}_i, \mathbf{z}_{\mathbf{x}_i})$ where the $(\bar{\mathbf{x}}_i, \mathbf{z}_{\mathbf{x}_i})$'s are the previously revealed secret keys and the \mathbf{x}_i are the corresponding vectors.

- If \mathbf{x} is linearly independent from the \mathbf{x}_i 's modulo N , set $\bar{\mathbf{x}} = \mathbf{x} \in \mathbb{Z}^\ell$ (with coefficients in $[0, N)$), $\mathbf{z}_{\mathbf{x}} = \langle \mathbf{s}, \mathbf{x} \rangle \in \mathbb{Z}$ and add $(\mathbf{x}, \bar{\mathbf{x}}, \mathbf{z}_{\mathbf{x}})$ to st .
- If $\mathbf{x} = \sum_i k_i \mathbf{x}_i \bmod N$ for some coefficients $\{k_i\}_{i \leq j-1}$ in \mathbb{Z}_N , then compute $\bar{\mathbf{x}} = \sum_i k_i \cdot \bar{\mathbf{x}}_i \in \mathbb{Z}^\ell$ and $\mathbf{z}_{\mathbf{x}} = \sum_i k_i \cdot \mathbf{z}_{\mathbf{x}_i} \in \mathbb{Z}^m$.

In either case, return $\text{sk}_{\mathbf{x}} = (\bar{\mathbf{x}}, \mathbf{z}_{\mathbf{x}})$.

Encrypt(mpk, \mathbf{y}): To encrypt a vector $\mathbf{y} = (y_1, \dots, y_\ell) \in \mathbb{Z}_N^\ell$, sample $r \leftarrow \{0, \dots, \lfloor N/4 \rfloor\}$ and compute

$$\begin{aligned} C_0 &= g^r \bmod N^2, \\ C_i &= (1 + y_i N) \cdot h_i^r \bmod N^2, \quad \forall i \in \{1, \dots, \ell\}. \end{aligned}$$

Return $C_{\mathbf{y}} = (C_0, C_1, \dots, C_\ell) \in \mathbb{Z}_{N^2}^{\ell+1}$.

Decrypt($\text{mpk}, \text{sk}_{\mathbf{x}}, C_{\mathbf{y}}$): Given $\text{sk}_{\mathbf{x}} = (\bar{\mathbf{x}}, \mathbf{z}_{\mathbf{x}}) \in \mathbb{Z}^\ell \times \mathbb{Z}$ with $\bar{\mathbf{x}} = (\bar{x}_1, \dots, \bar{x}_\ell)$, compute

$$C_{\mathbf{x}} = \left(\prod_{i=1}^{\ell} C_i^{\bar{x}_i} \right) \cdot C_0^{-\mathbf{z}_{\mathbf{x}}} \bmod N^2.$$

Then, compute and output $\log_{(1+N)}(C_{\mathbf{x}}) = \frac{C_{\mathbf{x}} - 1 \bmod N^2}{N}$.

From a security standpoint, the following result is proved in Appendix G.

Theorem 6. *The above stateful scheme provides full security under the DCR assumption.*

6 Bootstrapping Linear FE to Efficient Bounded FE for all circuits

In this section, we describe how to compile our Linear FE scheme, denoted by LinFE which computes linear functions modulo p (for us $p = 2$), into a bounded collusion FE scheme for all circuits, denoted by BddFE . The underlying scheme LinFE is assumed to be AD-IND secure, which, by [48], is equivalent to

non-adaptive simulation secure NA-SIM, since linear functions are “preimage sampleable”. We refer the reader to [48] for more details.

Let \mathcal{C} be a family of polynomial-size circuits. Let $C \in \mathcal{C}$ and let \mathbf{x} be some input. Let $\tilde{C}(\mathbf{x}, R)$ be a randomized encoding of C that is computable by a constant depth circuit with respect to inputs x and R (see [8]). Then consider a new family of circuits \mathcal{G} defined by:

$$G_{C,\Delta}(x, R_1, \dots, R_S) = \left\{ \tilde{C}\left(x; \bigoplus_{a \in \Delta} R_a\right) : C \in \mathcal{C}, \Delta \subseteq [S] \right\},$$

for some S to be chosen below. As observed in [38, Section 6], circuit $G_{C,\Delta}(\cdot, \cdot)$ is computable by a constant degree polynomial (one for each output bit). Given an FE scheme for \mathcal{G} , one may construct a scheme for \mathcal{C} by having the decryptor first recover the output of $G_{C,\Delta}(\mathbf{x}, R_1, \dots, R_S)$ and then applying the decoder for the randomized encoding to recover $C(\mathbf{x})$.

Note that to support q queries the decryptor must compute q randomized encodings, each of which needs fresh randomness. As shown above, this is handled by hardcoding sufficiently many random elements in the ciphertext and taking a random subset sum of these to generate fresh random bits for each query. As in [38], the parameters are chosen so that the subsets form a cover-free system, so that every random subset yields fresh randomness (with overwhelming probability).

In more details, we let the set S, v, m be parameters to the construction. Let Δ_i for $i \in [q]$ be a uniformly random subset of S of size v . To support q queries, we identify the set $\Delta_i \subseteq S$ with query i . If $v = O(\lambda)$ and $S = O(\lambda \cdot q^2)$ then the sets Δ_i are cover-free with high probability. For details, we refer the reader to [38, Section 5]. We now proceed to describe our construction. Let $L \triangleq (\ell + S \cdot m)^3$, where $m \in \text{poly}(\lambda)$ is the size of the random input in the randomized encoding and ℓ is the length of the messages to be encrypted.

BddFE.Setup($1^\lambda, 1^\ell$): Upon input the security parameter λ and the message space $\mathcal{M} = \{0, 1\}^\ell$, invoke $(\text{mpk}, \text{msk}) = \text{LinFE.Setup}(1^\lambda, 1^L)$ and output it.

BddFE.KeyGen(msk, C): Upon input the master secret key and a circuit C , do:

1. Sample a uniformly random subset $\Delta \subseteq S$ of size v .
2. Express $C(\mathbf{x})$ by $G_{C,\Delta}(\mathbf{x}, R_1, \dots, R_S)$, which in turn can be expressed as a sequence of degree 3 polynomials P_1, \dots, P_k , where $k \in \text{poly}(\lambda)$.
3. Linearize each polynomial P_i and let P'_i be its vector of coefficients. Note that the ordering of the coefficients can be arbitrary but should be public.
4. Output $\text{BddFE.SK}_C = \{\text{SK}_i = \text{LinFE.KeyGen}(\text{LinFE.msk}, P'_i)\}_{i \in [k]}$.

BddFE.Enc(\mathbf{x}, mpk): Upon input the public key and the plaintext \mathbf{x} , do:

1. Sample $R_1, \dots, R_S \leftarrow \{0, 1\}^m$.
2. Compute all symbolic monomials of degree 3 in the variables x_1, \dots, x_ℓ and $R_{i,j}$ for $i \in [S], j \in [m]$. The number of such monomials is $L = (\ell + S \cdot m)^3$. Arrange them according to the public ordering and denote the resulting vector by \mathbf{y} .
3. Output $\text{CT}_{\mathbf{x}} = \text{LinFE.Enc}(\text{LinFE.mpk}, \mathbf{y})$.

BddFE.Dec(mpk, $\text{CT}_{\mathbf{x}}$, SK_C): Upon input a ciphertext $\text{CT}_{\mathbf{x}}$ for vector \mathbf{x} , and a secret key $\text{SK}_C = \{\text{SK}_i\}_{i \in [k]}$ for circuit C , do the following:

1. Compute $G_{C, \Delta}(\mathbf{x}, R_1, \dots, R_S) = \{P_i(\mathbf{Y})\}_{i \in [k]} = \{\text{LinFE.Dec}(\text{CT}_{\mathbf{x}}, \text{SK}_i)\}_{i \in [k]}$.
2. Run the decoder for the randomized encoding to recover $C(\mathbf{x})$ from $G_{C, \Delta}(\mathbf{x}, R_1, \dots, R_S)$.

Correctness follows from the correctness of LinFE and the correctness of randomized encodings.

Security. The definition for q -NA-SIM security is provided in Appendix A. We proceed to describe our simulator **Bdd.Sim**. Let **RE.Sim** be the simulator guaranteed by the security of randomized encodings and **LinFE.Sim** be the simulator guaranteed by the security of the LinFE scheme.

Simulator Bdd.Sim($\{C_i, C_i(\mathbf{x}), \text{SK}_i\}_{i \in [q^*]}$): The simulator **Bdd.Sim** receives the secret key queries C_i , the corresponding (honestly generated) secret keys SK_i and the values $C_i(\mathbf{x})$ for $i \in [q^*]$ where $q^* \leq q$, and must simulate the ciphertext $\text{CT}_{\mathbf{x}}$. It proceeds as follows:

1. Sample $\Delta_1, \dots, \Delta_q \subseteq S$, of size v each.
2. For each $i \in [q^*]$, invoke **RE.Sim**($C_i(x)$) to learn $G_{C_i}(\mathbf{x}, \hat{R}_i)$ for some \hat{R}_i chosen by the simulator. Interpret

$$\hat{R}_i = \bigoplus_{a \in \Delta_i} R_a \quad \text{and} \quad G_{C_i, \Delta_i}(\mathbf{x}, R_1, \dots, R_S) = G_{C_i}(\mathbf{x}, \hat{R}_i) = (P_1(\mathbf{Y}), \dots, P_k(\mathbf{Y})).$$

3. Let $\text{CT}_{\mathbf{x}} = \text{LinFE.Sim}(\{G_{C_i, \Delta_i}, G_{C_i, \Delta_i}(\mathbf{x}, R_1, \dots, R_S), \text{SK}_i\}_{i \in [q^*]})$ and output it.

The correctness of **Bdd.Sim** follows from the correctness of **RE.Sim** and **LinFE.Sim**.

A last remaining technicality is that the most general version of our construction for FE for inner product modulo p is stateful. This is because a general adversary against LinFE may request keys that are linearly dependent modulo p but linearly independent over the integers, thus learning new linear relations in the master secret. This forces the simulator (and hence the key generator) to maintain a state.

However, in our application, we can make do with a stateless variant, since all the queries will be linearly independent over \mathbb{Z}_2 . To see this, note that in the above application of LinFE, each query is randomized by a unique random set Δ_i . Recall that by cover-freeness, the element $\bigoplus_{a \in \Delta_i} R_a$ must contain at least one fresh random element, say R^* , which is not contained by $\bigcup_{j \neq i} \Delta_j$. Stated a bit differently, if we consider the query vectors of size L , then cover-freeness implies that no query vector lies within the linear span of the remaining queries made by the adversary. For any query Q , there is at least one position $j \in [L]$ so that this position is nonzero in the L vector representing Q but zero for all other vectors. Hence the query vectors are linearly independent over \mathbb{Z}_2 , for which case, our construction of Section 4.2 is stateless.

Acknowledgements. We thank Fabrice Benhamouda and Hoeteck Wee for helpful discussions. This work has been supported in part by ERC Starting Grant ERC-2013-StG-335086-LATTAC. Part of this work was also funded by the “Programme Avenir Lyon Saint-Etienne de l’Université de Lyon” in the framework of the programme “Investissements d’Avenir” (ANR-11-IDEX-0007).

References

1. M. Abdalla, M. Bellare, D. Catalano, E. Kiltz, T. Kohno, T. Lange, J. Malone-Lee, G. Neven, P. Paillier, and H. Shi. Searchable encryption revisited: Consistency properties, relation to anonymous IBE, and extensions. In *Proc. of Crypto*, volume 3621 of *LNCS*, pages 205–222. Springer, 2005.
2. M. Abdalla, F. Bourse, A. De Caro, and D. Pointcheval. Simple functional encryption schemes for inner products. In *Proceedings of PKC*, volume 9020 of *LNCS*, pages 733–751. Springer, 2015.
3. S. Agrawal, D. M. Freeman, and V. Vaikuntanathan. Functional encryption for inner product predicates from learning with errors. In *Proc. of Asiacrypt*, volume 7073 of *LNCS*, pages 21–40. Springer, 2011.
4. S. Agrawal, S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption: New perspectives and lower bounds. In *Proc. of CRYPTO*, volume 8043 of *LNCS*, pages 500–518. Springer, 2013.
5. S. Agrawal and A. Rosen. Online-offline functional encryption for bounded collusions. Cryptology ePrint Archive, Report 2016/361, 2016. <http://eprint.iacr.org/>.
6. J. Alperin-Sheriff and C. Peikert. Circular and KDM security for identity-based encryption. In *Proceedings of PKC*, volume 7293 of *LNCS*, pages 334–352. Springer, 2012.
7. P. Ananth, Z. Brakerski, G. Segev, and V. Vaikuntanathan. The trojan method in functional encryption: From selective to adaptive security, generically. In *Proc. of CRYPTO*, LNCS, pages 657–677. Springer, 2015.
8. B. Applebaum, Y. Ishai, and E. Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.
9. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proc. of STOC*, pages 1–10. ACM, 1988.
10. D. Bernstein and T. Lange. Computing small discrete logarithms faster. In *Proc. of Indocrypt’12*, volume 7668 of *LNCS*, pages 317–338. Springer, 2012.
11. A. Bishop, A. Jain, and L. Kowalczyk. Function-hiding inner product encryption. In *Asiacrypt 2015*, volume 9452 of *LNCS*, pages 470–491. Springer, 2015.
12. D. Boneh and X. Boyen. Efficient selective-id secure identity-based encryption without random oracles. In *Proc. of Eurocrypt*, volume 3027 of *LNCS*, pages 223–238. Springer, 2004.
13. D. Boneh and X. Boyen. Efficient selective identity-based encryption without random oracles. *Journal of Cryptology*, 24(4):659–693, 2011.
14. D. Boneh, X. Boyen, and H. Shacham. Short group signatures. In *Proc. of Crypto*, volume 3152 of *LNCS*, pages 41–55. Springer, 2004.
15. D. Boneh, G. Di Crescenzo, R. Ostrovsky, and G. Persiano. Public key encryption with keyword search. In *Proc. of Eurocrypt*, volume 3027 of *LNCS*, pages 506–522. Springer, 2004.

16. D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. *SIAM J. Comput.*, 32(3):586–615 (electronic), 2003.
17. D. Boneh, C. Gentry, S. Gorbunov, S. Halevi, V. Nikolaenko, G. Segev, V. Vaikuntanathan, and D. Vinayagamurthy. Fully key-homomorphic encryption, arithmetic ABE and compact garbled circuits. In *Proc. of Eurocrypt*, volume 8441 of *LNCS*, pages 533–556. Springer, 2014.
18. D. Boneh, A. Sahai, and B. Waters. Functional encryption: Definitions and challenges. In *Proc. of TCC*, volume 6597 of *LNCS*, pages 253–273. Springer, 2011.
19. Z. Brakerski, A. Langlois, C. Peikert, Regev. O., and D. Stehlé. On the classical hardness of learning with errors. In *Proc. of STOC*, pages 575–584. ACM, 2013.
20. E. Bresson, D. Catalano, and D. Pointcheval. A simple public-key cryptosystem with a double trapdoor decryption mechanism and its applications. In *Proc. of Asiacrypt 2003*, volume 2894 of *LNCS*, pages 37–54. Springer, 2003.
21. J. Camenisch and V. Shoup. Practical verifiable encryption and decryption of discrete logarithms. In *CRYPTO*, pages 126–144, 2003.
22. J.-H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé. Cryptanalysis of the multilinear map over the integers. In *Proc. of EUROCRYPT*, volume 9056 of *LNCS*, pages 3–12. Springer, 2015.
23. J.-S. Coron, T. Lepoint, and M. Tibouchi. Practical multilinear maps over the integers. In *Proc. of CRYPTO*, volume 8042 of *LNCS*, pages 476–493, 2013.
24. R. Cramer and V. Shoup. A practical public key cryptosystem provably secure against adaptive chosen ciphertext attack. In *Proc. of CRYPTO*, volume 1462 of *LNCS*, pages 13–25. Springer, 1998.
25. R. Cramer and V. Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In *Proc. of Eurocrypt*, volume 2332 of *LNCS*, pages 45–64. Springer, 2002.
26. I. Damgård. Towards practical public key systems secure against chosen ciphertext attacks. In *Proc. of CRYPTO*, volume 576 of *LNCS*, pages 445–456. Springer, 1991.
27. A. De Caro, V. Iovino, A. Jain, A. O’Neill, O. Paneth, and G. Persiano. On the achievability of simulation-based security for functional encryption. In *Crypto’13*, volume 8043 of *LNCS*, pages 519–535. Springer, 2013.
28. Y. Dodis and N. Fazio. Public key trace and revoke scheme secure against adaptive chosen ciphertext attack. In *PKC 2003*, volume 2567 of *LNCS*, pages 100–115. Springer, 2003.
29. L. Ducas, A. Durmus, T. Lepoint, and V. Lyubashevsky. Lattice signatures and bimodal gaussians. In *Proc. of CRYPTO*, volume 8042 of *LNCS*, pages 40–56. Springer, 2013.
30. A. Escala, G. Herold, E. Kiltz, C. Ràfols, and G. Villard. An algebraic framework for Diffie-Hellman assumptions. In *Proc. of Crypto*, volume 8043 of *LNCS*, pages 129–147. Springer, 2013.
31. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices and applications. In *Proc. of Eurocrypt*, LNCS. Springer, 2013.
32. S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Proc. of FOCS*, pages 40–49, 2013.
33. S. Garg, C. Gentry, S. Halevi, A. Sahai, and B. Waters. Attribute-based encryption for circuits from multilinear maps. In *Proc. of Crypto*, volume 8043 of *LNCS*, pages 479–499. Springer, 2013.
34. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.

35. S. Goldwasser, A. Lewko, and D. Wilson. Bounded-collusion IBE from key homomorphism. In *Proceedings of TCC*, volume 7194 of *LNCS*, pages 564–581. Springer, 2012.
36. S. Goldwasser, Y. Tauman Kalai, R. Popa, V. Vaikuntanathan, and N. Zeldovich. How to run Turing machines on encrypted data. In *CRYPTO (2)*, pages 536–553, 2013.
37. S. Goldwasser, Y. Tauman Kalai, R. Popa, V. Vaikuntanathan, and N. Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proc. of STOC*, pages 555–564. ACM Press, 2013.
38. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Functional encryption with bounded collusions via multi-party computation. In *Proc. of Crypto*, volume 7417 of *LNCS*, pages 162–179. Springer, 2012.
39. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In *Proc. of STOC*, pages 545–554. ACM Press, 2013.
40. S. Gorbunov, V. Vaikuntanathan, and H. Wee. Predicate encryption for circuits from LWE. In *Proc. of CRYPTO*, LNCS. Springer, 2015.
41. V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *Proc. of ACM-CCS'06*, pages 89–98. ACM Press, 2006.
42. Y. Hu and H. Jia. Cryptanalysis of GGH map. In *Proc. of EUROCRYPT 2016*, volume 9665 of *LNCS*, pages 537–565. Springer, 2016.
43. J. Katz, A. Sahai, and B. Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Proc. of EUROCRYPT*, volume 4965 of *LNCS*, pages 146–162. Springer, 2008.
44. S. Ling, D. H. Phan, D. Stehlé, and R. Steinfeld. Hardness of k -LWE and Applications in Traitor Tracing. In *Proc. of CRYPTO*, volume 8616 of *LNCS*, pages 315–334. Springer, 2014.
45. D. Micciancio and P. Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 465–484. Springer, 2011.
46. D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. of EUROCRYPT*, volume 7237 of *LNCS*, pages 700–718. Springer, 2012.
47. P. Q. Nguyen and J. Stern. Merkle-Hellman revisited: A cryptanalysis of the Qu-Vanstone cryptosystem based on group factorizations. In *Proc. of CRYPTO*, volume 1294 of *LNCS*, pages 198–212. Springer, 1997.
48. A. O'Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. <http://eprint.iacr.org/>.
49. A. O'Neill, C. Peikert, and B. Waters. Bi-deniable public-key encryption. In *Proc. of CRYPTO*, volume 6841 of *LNCS*, pages 525–542. Springer, 2011.
50. P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Proc. of EUROCRYPT*, volume 1592 of *LNCS*, pages 223–238. Springer, 1999.
51. J. Pollard. Kangaroos, monopoly and discrete logarithms. *Journal of Cryptology*, 13:433–447, 2000.
52. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.
53. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
54. A. Sahai and H. Seyalioglu. Worry-free encryption: Functional encryption with public keys. In *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS '10*, 2010.

55. A. Sahai and B. Waters. Fuzzy identity-based encryption. In *Proc. of EUROCRYPT*, volume 3494 of *LNCS*, pages 457–473. Springer, 2005.
56. A. Shamir. Identity-based cryptosystems and signature schemes. In *Proc. of CRYPTO*, volume 196 of *LNCS*, pages 47–53. Springer, 1984.
57. B. Waters. Functional encryption for regular languages. In *Proc. of Crypto*, volume 7417 of *LNCS*, pages 218–235. Springer, 2012.

A Definitions for functional encryption

We now recall the syntax of Functional Encryption, as defined by Boneh, Sahai and Waters [18], and their indistinguishability-based security definition.

Definition 6 ([18]). *A functionality F defined over $(\mathcal{K}, \mathcal{Y})$ is a function $F : \mathcal{K} \times \mathcal{Y} \rightarrow \Sigma \cup \{\perp\}$, where \mathcal{K} is a key space, \mathcal{Y} is a message space and Σ is an output space, which does not contain the special symbol \perp .*

Definition 7. *A functional encryption (FE) scheme for a functionality F is a tuple $\mathcal{FE} = (\text{Setup}, \text{Keygen}, \text{Encrypt}, \text{Decrypt})$ of algorithms with the following specifications:*

Setup(1^λ): *Takes as input a security parameter 1^λ and outputs a master key pair (mpk, msk) .*

Keygen(msk, K): *Given the master secret key msk and a key (i.e., a function) $K \in \mathcal{K}$, this algorithm outputs a key sk_K .*

Encrypt(mpk, Y): *On input of a message $Y \in \mathcal{Y}$ and the master public key mpk , this randomized algorithm outputs a ciphertext C .*

Decrypt($\text{mpk}, \text{sk}_K, C$): *Given the master public key mpk , a ciphertext C and a key sk_K , this algorithm outputs $v \in \Sigma \cup \{\perp\}$.*

We require that, for all $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, all keys $K \in \mathcal{K}$ and all messages $Y \in \mathcal{Y}$, if $\text{sk}_K \leftarrow \text{Keygen}(\text{msk}, K)$ and $C \leftarrow \text{Encrypt}(\text{mpk}, Y)$, with overwhelming probability, we have $\text{Decrypt}(\text{mpk}, \text{sk}_K, C) = F(K, Y)$ whenever $F(K, Y) \neq \perp$.

In some cases, we will also give a state st as input to algorithm **Keygen**, so that a stateful authority may reply to key queries in a way that depends on the queries that have been made so far. In that situation, algorithm **Keygen** may additionally update state st .

INDISTINGUISHABILITY-BASED SECURITY. From a security standpoint, what we expect from a FE scheme is that, given $C \leftarrow \text{Encrypt}(\text{mpk}, Y)$, the only thing revealed by a secret key sk_K about the underlying Y is the function evaluation $F(K, Y)$. In the natural definition of indistinguishability-based security (see, e.g., [18]), one asks that no efficient adversary be able to differentiate encryptions of Y_0 and Y_1 without obtaining secret keys sk_K such that $F(K, Y_0) \neq F(K, Y_1)$.

Definition 8 (Indistinguishability-based security). *A functional encryption scheme $\mathcal{FE} = (\text{Setup}, \text{Keygen}, \text{Encrypt}, \text{Decrypt})$ provides semantic security under chosen-plaintext attacks (or IND-CPA security) if no PPT adversary has non-negligible advantage in the following game, where $q_1 \leq q \in \text{poly}(\lambda)$:*

1. The challenger runs $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and the master public key mpk is given to the adversary \mathcal{A} .
2. The adversary adaptively makes secret key queries to the challenger. At each query, adversary \mathcal{A} chooses a key $K \in \mathcal{K}$ and obtains $\text{sk}_K \leftarrow \text{Keygen}(\text{msk}, K)$.
3. Adversary \mathcal{A} chooses distinct messages Y_0, Y_1 subject to the restriction that, if $\{K_i\}_{i=1}^{q_1}$ denotes the set of secret key queries made by \mathcal{A} at Stage 2, it holds that $F(K_i, Y_0) = F(K_i, Y_1)$ for each $i \in \{1, \dots, q_1\}$. Then, the challenger flips a fair coin $\beta \leftarrow \{0, 1\}$ and computes $C^* \leftarrow \text{Encrypt}(\text{mpk}, Y_\beta)$ which is sent as a challenge to \mathcal{A} .
4. Adversary \mathcal{A} makes further secret key queries for arbitrary keys $K \in \mathcal{K}$. However, it is required that $F(K, Y_0) = F(K, Y_1)$ at each query $K \in \{K_{q_1+1}, \dots, K_q\}$.
5. Adversary \mathcal{A} eventually outputs a bit $\beta' \leftarrow \{0, 1\}$ and wins if $\beta' = \beta$.

The adversary's advantage is defined to be $\text{Adv}_{\mathcal{A}}(\lambda) := |\Pr[\beta' = \beta] - 1/2|$, where the probability is taken over all coin tosses.

Definition 8 captures *adaptive* security in that the adversary is allowed to choose the messages Y_0, Y_1 at Stage 3. In [2], Abdalla *et al.* considered a weaker security notion, called *selective* security, where the adversary has to declare the messages Y_0, Y_1 at the very beginning of the game, before even seeing mpk (note that, in this scenario, the adversary can receive the challenge ciphertext at the same time as the public key). In this work, our goal will be to meet the strictly stronger requirements of adaptive security.

Boneh, Sahai and Waters [18] pinpointed shortcomings of indistinguishability-based definitions in the case of general functionalities, where they may fail to rule out intuitively insecure systems. Boneh *et al.* [18] proposed strong simulation-based definitions, but these have been shown to be impossible to realize in the standard model [18, 4].

On the positive side, O'Neill [48] showed that indistinguishability-based security is equivalent to non-adaptive simulation based security (defined below) for a class of functions called preimage sampleable functions, which includes inner products. De Caro *et al.* [27] gave a general method of constructing FE schemes that achieve a meaningful definition of simulation-based security from systems that are only proved secure in the sense of indistinguishability-based definitions. Also, note that the impossibility of achieving adaptive simulation-based security for IBE, exhibited by [18] can be easily adapted to show that adaptive simulation-based security (AD-SIM) is also impossible to achieve for the inner product functionality. Thus, adaptive indistinguishability appears to be the strongest adaptive notion of security that may still be achievable, and for a wide range of practically interesting specific functionalities this notion is believed to suffice. In the following, we will aim at *full* security in the sense of Definition 8.

SIMULATION-BASED SECURITY FOR BOUNDED COLLUSIONS In this section, we define simulation based security for bounded collusions, as in [38, Defn 3.1].

Definition 9 (*q-NA-SIM- and q-AD-SIM- Security*).

Let \mathcal{F} be a functional encryption scheme for a circuit family \mathcal{C} . For every p.p.t. adversary $A = (A_1, A_2)$ and a p.p.t. simulator Sim , consider the following two experiments:

$\text{Exp}_{\mathcal{F}, A}^{\text{real}}(1^\lambda)$:	$\text{Exp}_{\mathcal{F}, \text{Sim}}^{\text{ideal}}(1^\lambda)$:
1: $(\text{MPK}, \text{MSK}) \leftarrow \text{FE.Setup}(1^\lambda)$ 2: $(x, st) \leftarrow A_1^{\text{FE.Keygen}(\text{MSK}, \cdot)}(\text{MPK})$ 3: $\text{CT} \leftarrow \text{FE.Enc}(\text{MPK}, x)$ 4: $\alpha \leftarrow A_2^{\mathcal{O}(\text{MSK}, \cdot)}(\text{MPK}, \text{CT}, st)$ 5: <i>Output</i> (x, α)	1: $\text{MPK} \leftarrow \text{FE.Setup}(1^\lambda)$ 2: $(x, st) \leftarrow A_1^{\text{FE.Keygen}(\text{MSK}, \cdot)}(\text{MPK})$ <i>Let</i> $\mathcal{V} \triangleq (C_i, C_i(x), \text{SK}_i)_{i \in [q]}$ 3: $\text{CT}, st' \leftarrow \text{Sim}(\text{MPK}, \mathcal{V}, 1^{ x })$ 4: $\alpha \leftarrow A_2^{\mathcal{O}'(\text{MSK}, st', \cdot)}(\text{MPK}, \text{CT}, st)$ 5: <i>Output</i> (x, α)

Above, C_i denote the queries made by the adversary. We distinguish between two cases of the above experiment:

1. The adaptive experiment, where:
 - the oracle $\mathcal{O}(\text{MSK}, \cdot) = \text{FE.Keygen}(\text{MSK}, \cdot)$ and
 - the oracle $\mathcal{O}'(\text{MSK}, st', \cdot)$ is the simulator, namely $\text{Sim}^{U_x(\text{MSK}, st', \cdot)}(\cdot)$ and $U_x(C) = C(x)$ for any $C \in \mathcal{C}$.

The simulator algorithm is stateful in that after each invocation, it updates the state st' which is carried over to its next invocation. We call a stateful simulator algorithm Sim admissible if, on each input C , Sim makes just a single query to its oracle $U_x(\cdot)$ on C itself.

The functional encryption scheme \mathcal{F} is then said to be q query simulation-secure for one message against adaptive adversaries (q -AD-SIM-secure, for short) if there is an admissible stateful p.p.t. simulator Sim such that for every p.p.t. adversary $A = (A_1, A_2)$ that makes at most q queries, the following two distributions are computationally indistinguishable:

$$\left\{ \text{Exp}_{\mathcal{F}, A}^{\text{real}}(1^\lambda) \right\}_{\lambda \in \mathbb{N}} \stackrel{c}{\approx} \left\{ \text{Exp}_{\mathcal{F}, \text{Sim}}^{\text{ideal}}(1^\lambda) \right\}_{\lambda \in \mathbb{N}}$$

2. The non-adaptive experiment, where the oracles $\mathcal{O}(\text{MSK}, \cdot)$ and $\mathcal{O}'(\text{MSK}, st, \cdot)$ are both the “empty oracles” that return nothing.

The functional encryption scheme \mathcal{F} is then said to be q query simulation-secure for one message against non-adaptive adversaries (q -NA-SIM-secure, for short) if there is an admissible stateful p.p.t. simulator Sim such that for every p.p.t. adversary $A = (A_1, A_2)$ that makes at most q queries, the two distributions above are computationally indistinguishable.

B Practical applications of Linear FE

On the practical front, Linear FE is already quite useful even when used directly. As pointed out by Abdalla *et al.* [2], the inner product functionality suffices for

the computation of linear functions (e.g., sums or averages) over encrypted data. As mentioned in the earlier work of Katz, Sahai and Waters [43], inner products also enable the evaluation of polynomials over encrypted data. To do this, we can simply encode a message M as a vector $\mathbf{y} = (1, M, \dots, M^d) \in \mathcal{D}^{d+1}$ and a degree- d polynomial $P[X] = \sum_{i=0}^d p_i X^i$ is encoded as a vector $\mathbf{x} = (p_0, p_1, \dots, p_d) \in \mathcal{D}^{d+1}$ for which the key $SK_{\mathbf{x}}$ is generated. Using a similar encoding, we can also evaluate multivariate polynomials of the form $P[X_1, \dots, X_d] = \prod_{i=1}^d (X_i - I_i)$ of small degree $d = O(\log \ell)$. By encoding ℓ -bit messages $M = M[1] \dots M[\ell]$ as vectors $\mathbf{y} = (M[1], \dots, M[\ell])$, the inner product functionality also allows for the computation of Hamming weights using secret keys $\mathbf{sk}_{\mathbf{x}}$ for the all-one vector $\mathbf{x} = (1, \dots, 1)$. More generally, inner products make it possible to compute the Hamming distance between an encrypted n -bit vector $\mathbf{y} = (y_1, \dots, y_n) \in \{0, 1\}^n$ and another binary vector $\mathbf{x} = (x_1, \dots, x_n) \in \{0, 1\}^n$ embedded in the key, which can be useful in biometric applications. To this end, we can simply encode \mathbf{y} and \mathbf{x} as vectors $\mathbf{Y} = (Y_1, \dots, Y_{2n}) \in \{-1, 1\}^{2n}$ and $\mathbf{X} = (X_1, \dots, X_{2n}) \in \{-1, 1\}^{2n}$ such that $X_{2i} = Y_{2i} = 1$ and $Y_{2i+1} = (-1)^{y_i}$, $X_{2i+1} = -(-1)^{x_i}$ for each $i \in \{1, \dots, n\}$. By doing so, the integer $\langle \mathbf{X}, \mathbf{Y} \rangle = \sum_{i=1}^{2n} (1 - (-1)^{x_i + y_i})$ is exactly twice the Hamming distance between \mathbf{x} and \mathbf{y} .

C Background on lattices

Let Λ be a non-zero lattice. We recall that the *smoothing parameter* of Λ is defined as $\eta_\varepsilon(\Lambda) = \min(s > 0 : \sum_{\hat{\mathbf{b}} \in \hat{\Lambda}} \text{Exp}(-\pi \|\hat{\mathbf{b}}\|^2/s^2) \leq 1 + \varepsilon)$, where $\hat{\Lambda} = \{\hat{\mathbf{b}} \in \text{span}_{\mathbb{R}}(\Lambda) : \hat{\mathbf{b}}^T \cdot \Lambda \subseteq \mathbb{Z}\}$ refers to the dual of Λ .

For a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ for some integers m, n, q , we define the lattice $\Lambda^\perp(\mathbf{A}) = \{\mathbf{x} \in \mathbb{Z}^m : \mathbf{x}^T \cdot \mathbf{A} = \mathbf{0} \pmod{q}\}$.

Lemma 6 (Adapted from [46, Le. 2.3 & 2.4]). *Let n, m, q be positive integers, and $\varepsilon, \delta > 0$. Assume that $q = p^k$ for p prime and $k \geq 1$. Assume further that*

$$m \geq \max \left(n + \frac{\log(3 + 4/(\delta\varepsilon))}{\log p}, \frac{n \log q + \log(2 + 2/(\delta\varepsilon))}{\log 2} \right).$$

Then $\eta_\varepsilon(\Lambda^\perp(\mathbf{A})) \leq 2\sqrt{\ln(2m(1 + 2/(\delta\varepsilon)))/\pi}$, except with probability $\leq \delta$ over the uniform choice of $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$.

Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice, $\sigma > 0$ and $\mathbf{c} \in \mathbb{R}^n$. We define the *lattice Gaussian distribution* of support Λ , standard deviation parameter σ and center \mathbf{c} as:

$$\forall \mathbf{b} \in \Lambda : D_{\Lambda, \sigma, \mathbf{c}}(\mathbf{b}) = \frac{\text{Exp}(-\pi \|\mathbf{b} - \mathbf{c}\|^2/\sigma^2)}{\sum_{\mathbf{x} \in \Lambda} \text{Exp}(-\pi \|\mathbf{x} - \mathbf{c}\|^2/\sigma^2)}.$$

We omit the subscript \mathbf{c} when $\mathbf{c} = \mathbf{0}$. To implement the primitives described in this work, we only need to be able to sample from 1-dimensional lattice Gaussians. Such an efficient sampler is described in [29].

We make use of the following lemmas.

Lemma 7 (Adapted from [34, Cor. 2.8]). *Let $\Lambda' \subseteq \Lambda \subseteq \mathbb{R}^n$ be two lattices with the same dimension. Let $\varepsilon \in (0, 1/2)$. Then for any $\mathbf{c} \in \mathbb{R}^n$ and any $\sigma \geq \eta_\varepsilon(\Lambda')$, the distribution $D_{\Lambda, \sigma, \mathbf{c}} \bmod \Lambda'$ is within statistical distance 2ε from the uniform distribution over Λ/Λ' .*

Lemma 8 (Adapted from [34, Le. 5.2]). *Assume the rows of $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ generate \mathbb{Z}_q^n and let $\varepsilon \in (0, 1/2)$, $\mathbf{c} \in \mathbb{Z}^m$ and $\sigma \geq \eta_\varepsilon(\Lambda^\perp(\mathbf{A}))$. Then for \mathbf{e} sampled from $D_{\mathbb{Z}^m, \sigma, \mathbf{c}}$, the distribution of the syndrome $\mathbf{e}^T \cdot \mathbf{A} \bmod q$ is within statistical distance 2ε of uniform over \mathbb{Z}_q^n .*

Note that if $q = p^k$ with p prime, then the rows of $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ generate \mathbb{Z}_q^n if and only if they generate \mathbb{Z}_p^n (once reduced modulo p). If \mathbf{A} is sampled uniformly, this occurs with probability $\geq 1 - p^{-n}$ when $m \geq 2n \log_2 p$.

Using the Lemmas 6 and 8, we obtain the following result, that we use in the proof of Theorem 2.

Lemma 9. *Let $n, m, q \geq 2$ be positive integers. Assume that $q = p^k$ for p prime and $k \geq 1$. Assume further that $m \geq 2n \log_2 q$. Let $\sigma \geq \Omega(\sqrt{n + \log m})$ and $\mathbf{c} \in \mathbb{Z}^m$. Then for $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ sampled uniformly and $\mathbf{e} \in \mathbb{Z}^m$ sampled from $D_{\mathbb{Z}^m, \sigma, \mathbf{c}}$, the distribution of the pair $(\mathbf{A}, \mathbf{e}^T \cdot \mathbf{A})$ is within statistical distance $2^{-\Omega(n)}$ of uniform over $\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^n$.*

D Missing material from Section 4

Proof of Lemma 3. The reduction from LWE to faelWE starts by sampling $\mathbf{A}' \leftarrow \mathbb{Z}_q^{t \times n}$. It aborts if it is not full-rank (modulo q): this happens with probability

$$\leq \prod_{p \text{ prime}, p|q} \left(1 - \prod_{0 \leq i < t} (1 - p^{-n+i}) \right) \leq \prod_{p \text{ prime}, p|q} (4p^{-n+t-1}) \leq 2^{-n+t+1}.$$

Else, the reduction computes $\mathbf{R} \in \mathbb{Z}_q^{n \times n}$ which is invertible and whose top $t \times n$ submatrix is \mathbf{A}' . The reduction also samples $\mathbf{s}' \leftarrow \mathbb{Z}_q^t$. The first t output samples are (\mathbf{a}'_i, s'_i) (for $i \leq t$), where \mathbf{a}'_i denote the i th row of \mathbf{A}' . The remaining samples are produced by taking a sample $(\mathbf{a}, b) \in \mathbb{Z}_q^{n-t} \times \mathbb{Z}_q$ from the given oracle, picking a fresh uniformly random $\mathbf{d} \in \mathbb{Z}_q^t$, and returning $(\mathbf{R}^T \cdot (\mathbf{d}|\mathbf{a}), b + \langle \mathbf{s}', \mathbf{d} \rangle)$.

Given uniform samples, the reduction outputs uniform samples up to statistical distance 2^{-n+t+1} . Given samples from $A_{q, \alpha, \mathbf{s}}$, the reduction outputs t samples from $A_{q, \{0\}, \mathbf{s}''}$ and the remaining samples from $A_{q, \alpha, \mathbf{s}''}$ up to statistical distance 2^{-n+t+1} , with $\mathbf{s}'' = \mathbf{R}^{-1} \cdot (\mathbf{s}'|\mathbf{s})^T \bmod q$. This proves correctness since \mathbf{R} induces a bijection on \mathbb{Z}_q^n . \square

Leftover hash lemma. We will use the following variant of the leftover hash lemma.

Lemma 10 (Particular case of [45, Le. 2.3]). Let $q = p^k$ for p prime and $k \geq 1$. Let $m \geq n \geq 1$. Take \mathcal{X} a distribution over \mathbb{Z}^m . Let D_0 be the uniform distribution over $\mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$ and D_1 be the distribution of $(\mathbf{A}, \mathbf{A} \cdot \mathbf{x}) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}_q^n$, where by sampling $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$ and $\mathbf{x} \leftarrow \mathcal{X}$. Then

$$\Delta(D_0, D_1) \leq \frac{1}{2} \sqrt{\sum_{i=1}^k p^{i \cdot n} \cdot \text{Pr}_i},$$

where Pr_i is the collision probability of two independent samples from $\mathcal{X} \bmod p^i$.

In particular, if the distribution $(\mathcal{X} \bmod p)$ is within statistical distance ε from the uniform distribution over \mathbb{Z}_p^m , then

$$\Delta(D_0, D_1) \leq \sqrt{(\varepsilon + p^{-m})q^n}.$$

E Proof of Theorem 1

Proof. The proof uses a sequence of games that begins with the real game and ends with a game where the adversary has no advantage at all. For each i , we denote by S_i the event that the adversary wins in Game i .

Game 0: This is the real game. In this game, the adversary \mathcal{A} is given mpk . In the challenge phase, \mathcal{A} chooses two distinct vectors $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_q^\ell$ and obtains an encryption of $\mathbf{y}_\beta = (y_{\beta,1}, \dots, y_{\beta,\ell})$, for some random $\beta \leftarrow \{0, 1\}$. At the end of the game, \mathcal{A} outputs $\beta' \in \{0, 1\}$ and we denote by S_0 the event that $\beta' = \beta$. For any vector $\mathbf{x} \in \mathbb{Z}_q^\ell$ submitted to the secret key extraction oracle, it must be the case that $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle \bmod q$.

Game 1: We modify the generation of the challenge $C_{\mathbf{y}_\beta} = (C, D, E_1, \dots, E_\ell)$. Namely, the challenger \mathcal{B} first computes

$$C = g^r \quad \text{and} \quad D = h^r, \tag{E.1}$$

for a randomly sampled $r \leftarrow \mathbb{Z}_q$. Then, it uses $\text{msk} := \{(s_i, t_i)\}_{i=1}^\ell$ to compute

$$E_i = g^{y_{\beta,i}} \cdot C^{s_i} \cdot D^{t_i}. \tag{E.2}$$

It can be observed that $C_{\mathbf{y}_\beta} = (C, D, E_1, \dots, E_\ell)$ has the same distribution as in Game 0. We hence have $\Pr[S_1] = \Pr[S_0]$.

Game 2: In this game, we modify again the generation of $C_{\mathbf{y}_\beta} = (C, D, E_1, \dots, E_\ell)$ in the challenge phase. Namely, instead of computing the pair (C, D) as in (E.1), the challenger \mathcal{B} samples $r, r' \leftarrow \mathbb{Z}_q$ and sets

$$C = g^r \quad \text{and} \quad D = h^{r+r'}.$$

The ciphertext components (E_1, \dots, E_ℓ) are still computed as per (E.2). Under the DDH assumption, this modification should not significantly affect \mathcal{A} 's view and we have $|\Pr[S_2] - \Pr[S_1]| \leq \mathbf{Adv}_B^{\text{DDH}}(\lambda)$.

In Game 2, we claim that the challenge ciphertext $C_{\mathbf{y}_\beta} = (C, D, E_1, \dots, E_\ell)$ perfectly hides $\beta \in \{0, 1\}$, so that $\Pr[S_2] = 1/2$. To see this, we first remark that, for each $i \in \{1, \dots, \ell\}$, we have

$$g^{y_{\beta,i}} \cdot C^{s_i} \cdot D^{t_i} = g^{y_{\beta,i} + \omega \cdot r' \cdot t_i} \cdot h_i^r,$$

where $\omega = \log_g(h)$, which means that an unbounded adversary can only infer

$$\begin{aligned} \mathbf{z}_\beta &= (y_{\beta,1} + \omega \cdot r' \cdot t_1, \dots, y_{\beta,\ell} + \omega \cdot r' \cdot t_\ell) \\ &= \mathbf{y}_\beta + \omega \cdot r' \cdot \mathbf{t} \in \mathbb{Z}_q^\ell \end{aligned}$$

from $C_{\mathbf{y}_\beta} = (C, D, E_1, \dots, E_\ell)$. To prove that \mathbf{z}_β does not reveal any information about $\beta \in \{0, 1\}$ to any legitimate adversary, we define $\mathbf{y} = \mathbf{y}_0 - \mathbf{y}_1 \bmod q$ and deterministically generate a \mathbb{Z}_q -basis $\mathbf{X}_{top} \in \mathbb{Z}_q^{(\ell-1) \times \ell}$ of the $(\ell-1)$ -dimensional subspace

$$\mathbf{y}^\perp = \{\mathbf{x} \in \mathbb{Z}_q^\ell \mid \langle \mathbf{x}, \mathbf{y} \rangle = 0 \bmod q\}.$$

Let $\mathbf{y}' \in \mathbb{Z}_q^\ell$ be a vector outside the subspace \mathbf{y}^\perp which we also choose in a deterministic manner. We define the invertible matrix

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{top} \\ \mathbf{y}'^T \end{bmatrix} \in \mathbb{Z}_q^{\ell \times \ell}$$

Since the rows of \mathbf{X} are deterministically generated from $\mathbf{y} \in \mathbb{Z}_q^\ell$, they are known to \mathcal{A} and it suffices to prove that $\mathbf{X} \cdot \mathbf{z}_\beta \in \mathbb{Z}_q^\ell$ is information-theoretically independent of $\beta \in \{0, 1\}$ to prove that \mathbf{z}_β does not reveal anything about β either. The first $(\ell-1)$ rows of $\mathbf{X} \cdot \mathbf{z}_\beta \in \mathbb{Z}_q^\ell$ are clearly independent of β since $\mathbf{X}_{top} \cdot \mathbf{y}_0 = \mathbf{X}_{top} \cdot \mathbf{y}_1 \bmod q$ by construction. We are thus left with the last row, which is the inner product $\langle \mathbf{y}', \mathbf{y}_\beta \rangle + \omega \cdot r' \cdot \langle \mathbf{y}', \mathbf{t} \rangle \bmod q$.

Let $(\mathbf{s}_0, \mathbf{t}_0) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q^\ell$ denote an arbitrary pair of vectors satisfying $(h_1, \dots, h_\ell) = g^{\mathbf{s}_0} \cdot h^{\mathbf{t}_0}$ and $\text{sk}_{\mathbf{x}_i} = (\langle \mathbf{s}_0, \mathbf{x}_i \rangle, \langle \mathbf{t}_0, \mathbf{x}_i \rangle)$ for all private key queries $\{\mathbf{x}_i\}_{i=1}^{\ell-1}$. Since all queries involve vectors $\mathbf{x}_i \in \mathbf{y}^\perp$, the joint distribution of the secret vectors $(\mathbf{s}, \mathbf{t}) \in \mathbb{Z}_q^\ell \times \mathbb{Z}_q^\ell$ is

$$\{(\mathbf{s}_0 - \omega \cdot \mu \cdot \mathbf{y} \bmod q, \mathbf{t}_0 + \mu \cdot \mathbf{y} \bmod q) \mid \mu \in \mathbb{Z}_q\}$$

in the adversary's view. The conditional distribution of $\omega \cdot r' \cdot \langle \mathbf{y}', \mathbf{t} \rangle \bmod q$ is thus

$$\{\omega \cdot r' \cdot (\langle \mathbf{y}', \mathbf{t}_0 \rangle + \mu \cdot \langle \mathbf{y}', \mathbf{y} \rangle) \bmod q \mid \mu \in \mathbb{Z}_q\},$$

which is nothing but the uniform distribution over \mathbb{Z}_q since $\langle \mathbf{y}', \mathbf{y} \rangle \neq 0 \bmod q$ by construction. Since $r' \neq 0$ with overwhelming probability, this means that the term $\omega \cdot r' \cdot \langle \mathbf{y}', \mathbf{t} \rangle \bmod q$ perfectly hides $\langle \mathbf{y}', \mathbf{y}_\beta \rangle$ in the inner product $\langle \mathbf{y}', \mathbf{z}_\beta \rangle$ modulo q . \square

F Proof of Theorem 5

Proof. The proof uses a sequence of games that begins with a game where the adversary is given a real encryption of \mathbf{y}_β , for some random bit $\beta \in \{0, 1\}$, and ends with a game where $\beta \in \{0, 1\}$ is statistically independent of the adversary's view. For each i , we denote by S_i the event that the adversary wins in Game i .

Game 0: This is the actual security game. The adversary \mathcal{A} is given the master public key $\text{mpk} = (N, g, \{h_i\}_{i=1}^\ell, Y)$, where $h_i = g^{s_i} \bmod N^2$ and $\mathbf{s} = (s_1, \dots, s_\ell)^T \leftarrow D_{\mathbb{Z}, \sigma}^\ell$ is a discrete Gaussian vector. In the challenge phase, \mathcal{A} chooses two vectors $\mathbf{y}_0 = (y_{0,1}, \dots, y_{0,\ell}), \mathbf{y}_1 = (y_{1,1}, \dots, y_{1,\ell}) \in \mathbb{Z}^\ell$ of norms $\leq Y$ and obtains an encryption of \mathbf{y}_β , for some random $\beta \in \{0, 1\}$. At the end of the game, \mathcal{A} outputs $\beta' \in \{0, 1\}$ and we denote by S_0 the event that $\beta' = \beta$. For any vector $\mathbf{x} \in \mathbb{Z}^\ell$ submitted by the adversary to the secret key extraction oracle, we must have the equality $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle$ over \mathbb{Z} .

Game 1: We modify the generation of the challenge $C_{\mathbf{y}_\beta} = (C_0, C_1, \dots, C_\ell)$. Namely, the challenger \mathcal{B} first chooses $z = z_0^N \bmod N^2$, for a randomly drawn $z_0 \leftarrow \mathbb{Z}_N^*$ and computes

$$C_0 = z^2 \bmod N^2. \quad (\text{F.1})$$

Then, it uses $\text{msk} := (\{s_i\}_{i=1}^\ell, X)$ to compute

$$C_i = (1 + N)^{y_{\beta,i}} \cdot C_0^{s_i} \bmod N^2, \quad \forall i \in \{1, \dots, \ell\}.$$

The ciphertext $C_{\mathbf{y}_\beta}$ has almost the same distribution as in Game 1 as C_0 is now perfectly (instead of statistically) uniform in the subgroup of $(2N)$ th residues. We have $|\Pr[S_1] - \Pr[S_0]| \leq 2^{-\lambda}$.

Game 2: We modify again the generation of $C_{\mathbf{y}_\beta} = (C_0, C_1, \dots, C_\ell)$ in the challenge phase. Namely, instead of computing C_0 by first choosing a random N th residue z in $\mathbb{Z}_{N^2}^*$, the challenger rather samples $z \leftarrow \mathbb{Z}_{N^2}^*$ at random, computes C_0 as in as in (F.1) (so that C_0 is a square in $\mathbb{Z}_{N^2}^*$ but not a N th residue, except with negligible probability) and sets

$$C_i = (1 + N)^{y_{\beta,i}} \cdot C_0^{s_i} \bmod N^2, \quad \forall i \in \{1, \dots, \ell\}.$$

Under the DCR assumption, this modification is not noticeable to \mathcal{A} , which implies that $|\Pr[S_2] - \Pr[S_1]| \leq \text{Adv}_{\mathcal{B}}^{\text{DCR}}(\lambda)$.

Game 3: We modify the generation of C_0 in the challenge ciphertext. Instead of choosing it uniformly in \mathbb{QR}_{N^2} , the challenger picks $a_z \leftarrow \mathbb{Z}_N^*$ and $r_z \leftarrow \{1, \dots, \lfloor N/4 \rfloor\}$ and computes

$$C_0 = (1 + N)^{a_z} \cdot g^{r_z} \bmod N^2.$$

The generation of $\{C_i\}_{i=1}^\ell$ remains unchanged. The statistical distance between the distribution of C_0 in Game 3 and Game 2 is smaller than $2^{-\lambda}$. We have $|\Pr[S_3] - \Pr[S_2]| \leq 2^{-\lambda}$.

In Game 3, we have $C_0 = (1 + N)^{a_z} \cdot g^{r_z} \bmod N^2$ and we claim that

$$(C_1, \dots, C_\ell) = ((1 + N)^{y_{\beta,1} + a_z \cdot s_1 \bmod N} \cdot h_1^{r_z} \bmod N^2, \dots, (1 + N)^{y_{\beta,\ell} + a_z \cdot s_\ell \bmod N} \cdot h_\ell^{r_z} \bmod N^2) \quad (\text{F.2})$$

statistically hides $\beta \in \{0, 1\}$ so that $|\Pr[S_3] - 1/2| \leq 2^{-\lambda}$. To prove this, we use a similar argument to the one in the proof of Theorem 2.

Namely, if we consider the vector $\mathbf{y} = (y_1, \dots, y_\ell) = \frac{1}{g}(\mathbf{y}_1 - \mathbf{y}_0) \in \mathbb{Z}^\ell$, where $g = \gcd(y_{0,1} - y_{1,1}, \dots, y_{0,\ell} - y_{1,\ell})$, we know that any legal private key query $\mathbf{x} \in \mathbb{Z}^\ell$ must be in the lattice $\{\mathbf{x} \in \mathbb{Z}^\ell : \langle \mathbf{x}, \mathbf{y} \rangle = 0\}$ for which the rows of the matrix $\mathbf{X}_{top} \in \mathbb{Z}^{(\ell-1) \times \ell}$ defined in (4.1) form a basis. We also know that $\|\mathbf{y}\|^2 < N$ and we may assume that $\gcd(\|\mathbf{y}\|^2, N) = 1$ since the reduction would be able to compute a non-trivial factor of N otherwise. We may also assume that the information that the adversary can infer about $\mathbf{s} = (s_1, \dots, s_\ell) \in \mathbb{Z}^\ell$ via private key queries is completely determined by $\mathbf{X}_{top} \cdot \mathbf{s} \in \mathbb{Z}^{\ell-1}$. Let us define the matrix

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{top} \\ \mathbf{y}^T \end{bmatrix} \in \mathbb{Z}^{\ell \times \ell}.$$

Since $\ell\|\mathbf{y}\|^2 < N$, the same arguments as in the proof of Theorem 2 ensure that the matrix \mathbf{X} is invertible in \mathbb{Z}_N unless its determinant reveals a non-trivial factor of N .

Recall that the challenge ciphertext (F.2) information-theoretically reveals

$$\begin{aligned} \mathbf{z}_\beta &= (y_{\beta,1} + a_z \cdot s_1 \bmod N, \dots, y_{\beta,\ell} + a_z \cdot s_\ell \bmod N)^T \in \mathbb{Z}_N^\ell \\ &= \mathbf{y}_\beta + a_z \cdot \mathbf{s} \bmod N, \end{aligned}$$

so that we have to show that \mathbf{z}_β is statistically independent of $\beta \in \{0, 1\}$ conditionally on \mathcal{A} 's view. Since $\mathbf{X} \in \mathbb{Z}^{\ell \times \ell}$ is invertible over \mathbb{Z}_N and does not depend on $\beta \in \{0, 1\}$, it suffices to show that $\mathbf{X} \cdot \mathbf{z}_\beta \in \mathbb{Z}_N^\ell$ is statistically independent of $\beta \in \{0, 1\}$. Given that $\mathbf{X}_{top} \cdot (\mathbf{y}_0 - \mathbf{y}_1) = \mathbf{0}$ (over the integers), $\mathbf{X}_{top} \cdot \mathbf{z}_\beta \in \mathbb{Z}_N^{\ell-1}$ is clearly independent of $\beta \in \{0, 1\}$ and we only need to worry about the last row of $\mathbf{X} \cdot \mathbf{z}_\beta \in \mathbb{Z}_N^\ell$, which can be written

$$\mathbf{y}^T \cdot \mathbf{z}_\beta \bmod N = \langle \mathbf{y}_\beta, \mathbf{y} \rangle + a_z \cdot \langle \mathbf{s}, \mathbf{y} \rangle \bmod N. \quad (\text{F.3})$$

Let $\mathbf{s}_0 = (s_{0,1}, \dots, s_{0,\ell}) \in \mathbb{Z}^\ell$ denote an arbitrary vector satisfying

$$\mathbf{X}_{top} \cdot \mathbf{s}_0 = \mathbf{X}_{top} \cdot \mathbf{s} \in \mathbb{Z}^{\ell-1}, \quad \text{and} \quad h_i = g^{s_{0,i}} \bmod N^2 \quad \forall i \in [1, \ell].$$

From \mathcal{A} 's view, the distribution of the master secret key $\mathbf{s} \in \mathbb{Z}^\ell$ is $\mathbf{s}_0 + D_{\Lambda, \sigma, -\mathbf{s}_0}$, where

$$\Lambda = \{\mathbf{t} \in \mathbb{Z}^\ell \mid \mathbf{X}_{top} \cdot \mathbf{t} = \mathbf{0} \in \mathbb{Z}^{\ell-1}, \mathbf{t} = \mathbf{0} \bmod p'q'\},$$

which is the lattice $\Lambda = \mathbb{Z} \cdot \mathbf{y} \cap (p'q' \cdot \mathbb{Z})^\ell = (p'q') \cdot \mathbb{Z} \cdot \mathbf{y}$. Conditionally on $\{h_i\}_{i=1}^\ell$ and $\mathbf{X}_{top} \cdot \mathbf{s}$, the distribution of $\langle \mathbf{s}, \mathbf{y} \rangle$ is thus

$$\langle \mathbf{s}_0, \mathbf{y} \rangle + D_{(p'q')\|\mathbf{y}\|^2 \cdot \mathbb{Z}, \|\mathbf{y}\|_{\sigma, -c}},$$

where $c = \langle \mathbf{s}_0, \mathbf{y} \rangle \in \mathbb{Z}$. We consider the distribution obtained by reducing the distribution $D_{(p'q')\|\mathbf{y}\|^2 \cdot \mathbb{Z}, \|\mathbf{y}\|_{\sigma, -c}}$ over $(p'q')\|\mathbf{y}\|^2 \cdot \mathbb{Z}$ modulo the sublattice $A' = (p'q')\|\mathbf{y}\|^2 \cdot (N\mathbb{Z})$. By Lemma 7, given that $|(p'q')\|\mathbf{y}\|^2 \cdot N| < N^2\|\mathbf{y}\|^2$, $\gcd(p'q'\|\mathbf{y}\|^2, N) = 1$ and $\|\mathbf{y}\| < \sqrt{N}$, choosing $\sigma > \sqrt{\lambda} \cdot N^{5/2}$ suffices to ensure that $\langle \mathbf{s}, \mathbf{y} \rangle \bmod N$ is within distance $2^{-\lambda}$ from the uniform distribution over $A/A' \simeq \mathbb{Z}_N$ when $\{h_i\}_{i=1}^{\ell}$ and $\mathbf{X}_{top} \cdot \mathbf{s} \in \mathbb{Z}^{\ell-1}$ are given. Since a_z is invertible in \mathbb{Z}_N with all but negligible probability, the term $\langle \mathbf{y}_\beta, \mathbf{y} \rangle \bmod N$ is thus statistically hidden in the right-hand-side member of (F.3).

When counting probabilities, we find that \mathcal{A} 's advantage in the real game can be bounded as

$$|\Pr[S_0] - 1/2| \leq \mathbf{Adv}_B^{\text{DCR}}(\lambda) + 2^{-\lambda+1},$$

which is thus negligible if the DCR assumption holds. \square

G Proof of Theorem 6

Proof. The proof proceeds similarly to the proof of Theorem 5 and uses the same sequence of games. The only modification is in the argument to argue that the adversary's view is independent of $\beta \in \{0, 1\}$ in the final game.

For each i , we denote by S_i the event that the adversary wins in Game i . We first recall the final game of the sequence:

Game 3': At the beginning of the game, the challenger sets $g = g'^{2N} \bmod N^2$, where $g' \leftarrow \mathbb{Z}_{N^2}^*$, and picks $\mathbf{s} = (s_1, \dots, s_\ell)^T \leftarrow D_{\mathbb{Z}^\ell, \sigma}$ before defining the public parameters $h_i = g^{s_i} \bmod N^2$ for each $i \in \{1, \dots, \ell\}$. The master public key $\text{mpk} = (N, g, \{h_i\}_{i=1}^{\ell})$ is given to the adversary \mathcal{A} . In the challenge phase, \mathcal{A} chooses two distinct vectors $\mathbf{y}_0, \mathbf{y}_1 \in \mathbb{Z}_N^\ell$. The challenger flips a random coin $\beta \leftarrow \{0, 1\}$ and encrypts \mathbf{y}_β by choosing $a_z \leftarrow \mathbb{Z}_N$, $r_z \leftarrow \{0, \dots, \lfloor N/4 \rfloor\}$ and setting

$$\begin{aligned} C_0 &= (1 + N)^{a_z} \cdot g^{r_z} \bmod N^2, \\ C_i &= (1 + N)^{y_{\beta, i} + a_z \cdot s_i} \cdot h_i^{r_z} \bmod N^2 \quad \forall i \in \{1, \dots, \ell\}. \end{aligned}$$

Note that any vector $\mathbf{x} \in \mathbb{Z}_N^\ell$ queried by \mathcal{A} to the private key generation oracle must satisfy $\langle \mathbf{x}, \mathbf{y}_0 \rangle = \langle \mathbf{x}, \mathbf{y}_1 \rangle \bmod N$. At each query, the adversary \mathcal{A} is given a secret key $(\bar{\mathbf{x}}, \mathbf{z}_\mathbf{x})$ exactly as in the real scheme. When \mathcal{A} halts, it outputs $\beta' \in \{0, 1\}$ and wins in the event that $\beta' = \beta$.

Let us define $\mathbf{y} = \mathbf{y}_1 - \mathbf{y}_0 \in \mathbb{Z}_N^\ell$ and let $\mathbf{x}_i \in \mathbb{Z}_p^\ell$ be the vectors corresponding to the secret key queries made by \mathcal{A} . Since \mathcal{A} is a legitimate adversary, we know that $\langle \mathbf{x}_i, \mathbf{y} \rangle = 0 \bmod N$ for each secret key query $\mathbf{x}_i \in \mathbb{Z}_N^\ell$. Note that $\langle \mathbf{x}_i, \mathbf{y} \rangle = 0 \bmod p$ and $\langle \mathbf{x}_i, \mathbf{y} \rangle = 0 \bmod q$. We may further assume that \mathbf{y} is non-zero both modulo p and modulo q (otherwise, the reduction has found a non-trivial factor of N).

We have to show that \mathcal{A} 's view is statistically independent of $\beta \in \{0, 1\}$ in Game 3'. To this end, we use a similar analysis to that of Theorem 3. In

particular, we cannot use the same matrix \mathbf{X}_{top} as in (4.1) since, if we define \mathbf{X} to be the matrix formed by that matrix \mathbf{X}_{top} and \mathbf{y}^T , we cannot guarantee that the determinant of $\mathbf{X}\mathbf{X}^T$ is not a multiple of N . One option would be to define \mathbf{X} as the $\ell \times \ell$ matrix formed by the \mathbb{Z}_N -independent private key queries and \mathbf{y}^T . One difficulty is that such an \mathbf{X} may not be invertible (we may have $\langle \mathbf{y}, \mathbf{y} \rangle = 0 \pmod N$). Further, after the challenge phase, we cannot immediately rule out that the queried vectors \mathbf{x}_i somehow depend on $\beta \in \{0, 1\}$. To avoid a circularity, we use an inductive argument as in the proof of Theorem 3.

For each j from 0 to $\ell - 1$, we consider the adversary's view after exactly j private key queries for linearly independent vectors over \mathbb{Z}_N . Since \mathcal{A} 's view is trivially independent of $\beta \in \{0, 1\}$ before the challenge phase, we only need to worry about post-challenge queries and we assume that the challenge ciphertext has been generated when the j th query occurs. By induction, we will show that, for any $j \in \{1, \dots, \ell - 1\}$, the view of the adversary remains statistically independent of β after the first j queries. In particular, the $(j + 1)$ th linearly independent private key query will be statistically independent of β as well. It will also ensure that, for $j = \ell - 1$, \mathcal{A} 's view will remain statistically independent of β in Game 3', which is exactly what we have to show.

Our induction hypothesis thus considers a counter value $j \in \{0, \dots, \ell - 1\}$ and assumes that, at this point st, is independent of β . At this stage, st contains j tuples $(\mathbf{x}_i, \bar{\mathbf{x}}_i, \mathbf{z}_i)$. The $\bar{\mathbf{x}}_i$'s form a basis of both

$$\mathbf{y}^{\perp p} := \{\mathbf{x} \in \mathbb{Z}_p^\ell : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod p\} \quad \text{and} \quad \mathbf{y}^{\perp q} := \{\mathbf{x} \in \mathbb{Z}_q^\ell : \langle \mathbf{x}, \mathbf{y} \rangle = 0 \pmod q\}.$$

We deterministically extend the $\{\bar{\mathbf{x}}_i\}_{i=1}^j$ into $\{\bar{\mathbf{x}}_i\}_{i=1}^{\ell-1}$ that is a basis of both \mathbf{y}_p^\perp and \mathbf{y}_q^\perp which is statistically independent of β . This is achieved by linear algebra modulo N , and, if any Gaussian elimination step fails because $\mathbb{Z}/N\mathbb{Z}$ is not a field, the reduction succeeds as it has found a factor of N . This extension of $\{\bar{\mathbf{x}}_i\}_{i=1}^j$ into $\{\bar{\mathbf{x}}_i\}_{i=1}^{\ell-1}$ can be seen as having the challenger making dummy private key queries for itself and creating the corresponding $\bar{\mathbf{x}}_i$'s in \mathbb{Z}^ℓ so as to get a full basis of both $\mathbf{y}^{\perp p}$ and $\mathbf{y}^{\perp q}$. Note that we may assume that the challenger knows \mathbf{y} as we only need to apply the inductive argument to post-challenge queries.

We define $\mathbf{X}_{top} \in \mathbb{Z}^{(\ell-1) \times \ell}$ as the matrix whose i th row is $\bar{\mathbf{x}}_i$, for all $i \in \{1, \dots, \ell - 1\}$ (including the genuine and dummy private key queries). Via private key queries, the information obtained by the adversary amounts to a subset of the coordinates of $\mathbf{X}_{top} \cdot \mathbf{s} \in \mathbb{Z}^{(\ell-1)}$.

Let $\mathbf{x}' \in \mathbb{Z}_N^\ell$ be a vector outside $\mathbf{y}^{\perp p}$ and $\mathbf{y}^{\perp q}$ which is chosen as a deterministic function of $\{\bar{\mathbf{x}}_i\}_{i=1}^j$ and \mathbf{y} and let $\mathbf{X}_{bot} \in \mathbb{Z}^{1 \times \ell}$ be the canonical lift of $(\mathbf{x}')^T$ over the integers. By construction of $\mathbf{X}_{bot} \in \mathbb{Z}^{1 \times \ell}$, we have $\mathbf{X}_{bot} \cdot \mathbf{y} \neq 0 \pmod N$. Now, let us define the matrix

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{top} \\ \mathbf{X}_{bot} \end{bmatrix} \in \mathbb{Z}^{\ell \times \ell}.$$

By construction, this matrix \mathbf{X} is invertible modulo N . Moreover, by induction and construction, it is statistically independent of $\beta \in \{0, 1\}$.

In Game 3', the challenge ciphertext information-theoretically reveals

$$\begin{aligned} \mathbf{z}_\beta &= (y_{\beta,1} + a_z \cdot s_1 \bmod N, \dots, y_{\beta,\ell} + a_z \cdot s_\ell \bmod N)^T \in \mathbb{Z}_N^\ell \\ &= \mathbf{y}_\beta + a_z \cdot \mathbf{s} \bmod N \end{aligned}$$

and it thus suffices to prove that \mathbf{z}_β is almost independent of $\beta \in \{0,1\}$ conditionally on $\{h_i = g^{s_i} \bmod N^2\}_{i=1}^\ell$ and $\mathbf{X}_{top} \cdot \mathbf{s} \in \mathbb{Z}^{\ell-1}$. Since the rows of \mathbf{X} are known to \mathcal{A} in the information theoretic sense, it is sufficient to prove that $\mathbf{X} \cdot \mathbf{z}_\beta \in \mathbb{Z}_N^\ell$ is statistically independent of β .

Since $\mathbf{X}_{top} \cdot (\mathbf{y}_0 - \mathbf{y}_1) = \mathbf{0} \bmod N$, we remark that $\mathbf{X}_{top} \cdot \mathbf{z}_\beta \in \mathbb{Z}_N^{\ell-1}$ is necessarily independent of $\beta \in \{0,1\}$ and it suffices to consider the distribution of $\mathbf{X}_{bot} \cdot \mathbf{z}_\beta \bmod N$. Let $\mathbf{s}_0 = (s_{0,1}, \dots, s_{0,\ell})^T \in \mathbb{Z}^\ell$ be an arbitrary vector satisfying the equalities $h_i = g^{s_{0,i}} \bmod N^2$ and $\mathbf{X}_{top} \cdot \mathbf{s}_0 = \mathbf{X}_{top} \cdot \mathbf{s}$ (over \mathbb{Z}). Conditionally on the adversary's view, the distribution of $\mathbf{s} \in \mathbb{Z}^\ell$ is $\mathbf{s}_0 + D_{\Lambda, \sigma, -\mathbf{s}_0}$, where

$$\Lambda = \{\mathbf{t} \in \mathbb{Z}^\ell \mid \mathbf{X}_{top} \cdot \mathbf{t} = \mathbf{0} \in \mathbb{Z}^{\ell-1}, \mathbf{t} = \mathbf{0} \bmod p'q'\},$$

which is a one-dimensional lattice in \mathbb{Z}^ℓ . We have

$$\Lambda = \mathbf{y}' \cdot \mathbb{Z} \cap (p'q' \cdot \mathbb{Z})^\ell = (p'q') \cdot \mathbf{y}' \cdot \mathbb{Z}$$

for some $\mathbf{y}' \in \mathbb{Z}^\ell$. Note that, by the Chinese Remainder Theorem, there exists an invertible $\alpha \in \mathbb{Z}_N^*$ such that $\mathbf{y}' = \alpha \cdot \mathbf{y} \bmod N$ as, if we had $\alpha = 0 \bmod p$ or $\alpha = 0 \bmod q$, it would contradict the definition of \mathbf{y}' .

If we define Λ' to be the lattice generated by the rows of $\mathbf{X}_{top} \in \mathbb{Z}^{(\ell-1) \times \ell}$, we have $\|\mathbf{y}'\| = \det(\Lambda/(p'q')) \leq \det \Lambda'$ and Hadamard's bound implies $\|\mathbf{y}'\| \leq (\sqrt{\ell}N)^{\ell-1}$. Since $(p'q') \cdot \|\mathbf{y}'\| \leq (\sqrt{\ell}N)^\ell$, the smoothing parameter of the sublattice $(N \cdot \Lambda)$ must be bounded by $(\sqrt{\ell}N)^{\ell+1}$. By choosing $\sigma > \sqrt{\lambda}(\sqrt{\ell}N)^{\ell+1}$, we obtain that the distribution $(D_{\Lambda, \sigma, -\mathbf{s}_0} \bmod (N \cdot \Lambda))$ is within $2^{-\lambda}$ distance from the uniform distribution over $\Lambda/(N \cdot \Lambda)$ which is isomorphic to $\mathbf{y} \cdot \mathbb{Z}_N$ because $\gcd(p'q', N) = 1$ and $\mathbf{y}' = \alpha \cdot \mathbf{y} \bmod N$ for some $\alpha \in \mathbb{Z}_N^*$. Since $\mathbf{X}_{bot} \cdot \mathbf{y} \neq 0 \bmod p$ and $\mathbf{X}_{bot} \cdot \mathbf{y} \neq 0 \bmod q$, we find that $\mathbf{X}_{bot} \cdot \mathbf{s} \bmod p$ and $\mathbf{X}_{bot} \cdot \mathbf{s} \bmod q$ are statistically close to the uniform distribution over \mathbb{Z}_p and \mathbb{Z}_q , respectively.

Since $\gcd(a_z, N) = 1$ with overwhelming probability, this implies that the product $\mathbf{X}_{bot} \cdot \mathbf{z}_\beta \bmod N$ does not carry any significant information about β , as claimed.

Putting the above altogether, the adversary's advantage in the initial game can be bounded as

$$|\Pr[S_0] - 1/2| \leq \mathbf{Adv}_B^{\text{DCR}}(\lambda) + 2^{-\lambda+1}.$$

□