

Lattice-based Secret-Key Functional Encryption for Constant-Degree Polynomials

Valerio Cini¹, Russell W. F. Lai², Akin Ünal³, and Ivy K. Y. Woo²

¹ Bocconi University, Milan, Italy

² Aalto University, Espoo, Finland

³ Institute of Science and Technology Austria, Klosterneuburg, Austria

Abstract. We present a lattice-based construction of secret-key functional encryption (FE) for low-norm polynomials of any constant degree d , hence also for NC^0 circuits. We rely on two core ingredients:

1. New trapdoor and preimage sampling algorithms for certain degree- d tensor-structured matrices, used to generate functional secret keys.
2. A new k -LWE-style assumption where short preimages of non-zero images with respect to the above tensor-structured matrix are given as hints, under which we prove that our secret-key FE scheme is selectively secure (under unbounded collusion). To gain confidence in the new assumption, we prove that the standard LWE assumption implies the degree-1 case and cryptanalyse the $d > 1$ case.

As a corollary, we obtain a new pathway to post-quantum secure indistinguishability obfuscation (iO), conditioned on the above new assumption, standard LWE, and the existence of polynomial-stretch pseudorandom generators in NC^0 . Along the way, we give a new, simple (public-key) FE scheme for linear functions with selective security under the standard LWE assumption.

1 Introduction

Functional encryption (FE) [BSW11, Wat12] allows deriving functional secret keys sk_f for a function f , which can be used to decrypt a ciphertext encrypting x to obtain $f(x)$ but nothing else. Simple constructions of public-key FE for linear functions (linear FE) exist based on standard assumptions over both groups and lattices [ABDP15, BJK15, ALS16]. While public-key FE for quadratic functions (quadratic FE) from pairings [BCFG17, Wee20] was constructed almost a decade ago, even a secret-key construction from lattices remains elusive. Beyond quadratic, FE for any higher-degree polynomials is only known from indistinguishability obfuscation (iO) or cryptographic multilinear maps.

Apart from the general research goal for lattices to achieve feature parity with pairings, constructing a lattice-based FE for beyond-linear functions is significant, since it enables new pathways to post-quantum secure iO, a quest that the state-of-the-art iO construction based on well-founded assumptions [JLS21, JLS22, RVV24] has not fulfilled. It is well-known that degree- d public-key FE, conditioned on different sizes and runtime trade-offs, is sufficient to imply iO [BV15, AJ15, AJS15], and the degree d can be down to constant [LPST16b, LV16], as low as $d = 5$ [Lin17, AS17]. The paths [KNT18, BNPW20] starting from secret-key FE are more complex, but their implications towards iO remain the same (see next subsection). There have been several attempts to construct post-quantum secure iO following an alternative path starting from “exponentially-efficient iO (XiO)” [LPST16a], where the earlier constructions are either heuristic [BDGM20] or based on subsequently cryptanalysed assumptions [WW21, GP21, DQV⁺21, BDGM22], while the most recent ones [CLW25, HJL25a] are still based on rather complex assumptions.

1.1 Our Contributions

We present a candidate lattice-based secret-key FE construction for polynomials of any constant degree d , hence also for NC^0 circuits. The core technical component of our construction is a method of sampling structured matrices together with their lattice trapdoors, which allows sampling a master secret key satisfying specific algebraic structures. Our constant-degree FE achieves compactness, with ciphertext size quasi-linear

in the message length (resp. sublinear in the description size of the supported functions) when aimed at quasi-polynomial (resp. sub-exponential) security. The selective security of our scheme is based on the $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption (Assumption 1), a secret-key variant of a degree- d generalisation of the affine $k\text{-LWE}$ assumption, the latter of which is known to be implied by the standard LWE assumption.

To gain initial confidence in our new assumption, we note that for the special case of $d = 1$, (even the public-key variant of) our assumption is implied by $k\text{-LWE}$. This implies that (the public-key variant of) our degree-1 scheme achieves the said security under standard LWE. For $d > 1$, we provide cryptanalytic evidence for the plausibility of our new assumption. While the secret-key variant survives all our cryptanalysis attempts, we describe an attack that succeeds against the public-key variant in our experiments. This is why we conjecture only the secret-key variant.

To our knowledge, no transform from secret-key FE for constant-degree polynomials to iO is explicitly stated in the literature, but one follows from straightforward adaptations of existing transforms. Briefly, [BV15, AJ15, LV16] turn a public-key FE for constant-degree polynomials with encryption time linear in the message length into a single-key public-key FE for P/poly with non-trivial runtime, i.e. sublinear in the description size of the supported functions. The same transform, applied to a secret-key FE with a sufficiently small ciphertext size rather than encryption time, turns out to yield a single-key secret-key FE for P/poly with sublinear-size ciphertexts. Feeding the latter into the transform of [BNPW16, BNPW20] from sublinear-time secret-key FE to sublinear-time XiO^4 , but with their runtime analysis applied to ciphertext size, yields XiO , hence iO under sub-exponential LWE [LPST16a] (see Appendix G for details). As a corollary, we obtain a candidate post-quantum secure iO from the sub-exponential hardness of $\text{sk-}k\text{-ALWE}^{\otimes d}$ and LWE, and sub-exponentially secure pseudorandom generators in NC^0 with polynomial stretch.

1.2 Related Works

Lattice-based constructions of FE were only known for linear functions [ABDP15, ALS16] (cf. Remark 4). Previous lattice-based attempts on FE for beyond-linear functions, while based on standard LWE, are either not compact, based on primitives without known instantiation, or only proven secure under bounded collusion [ALS16, CC17, Agr17, AR17]. Furthermore, these constructions typically compose advanced primitives such as FHE and ABE in a non-black-box manner [AR17]. Constructing a lattice-based degree- d FE for $d > 1$ without these compromises is a long-standing open problem. Tairi and Ünal [TÜ24] even ruled out the existence of lattice-based degree- d FE for $d > 1$ under a general construction framework and satisfying certain parameter constraints. In this work we explore a direct, algebraic approach of constructing constant-degree FE, hence also iO, at the cost of making a non-standard $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption. While our construction is captured by the framework [TÜ24], its parameters fall outside the impossible regime. We refer to our comprehensive cryptanalysis of $\text{sk-}k\text{-ALWE}^{\otimes d}$ in Section 7 and Appendix D for more details. Since many prior lattice-based assumptions leading to iO have been broken, we also provide in Appendix E a comparison between $\text{sk-}k\text{-ALWE}^{\otimes d}$ and those assumptions.

2 Technical overview

To understand our constant-degree FE construction, it is instructive to first overview its special case of degree 1, i.e. our new linear FE scheme, and a proof sketch of it. These will pave the way to our quadratic FE explained later in this overview, and ultimately our constant-degree FE presented in Section 6.

2.1 Our (public-key) Linear FE from $k\text{-LWE}$

We consider a message space $\mathcal{M} = (\mathbb{Z} \cap (-p/2, p/2])^k$ consisting of low-norm vectors, where p, q are two layers of moduli with $p \ll q$. Our linear FE scheme supports the class of functions $\mathcal{F} = \{f_{\mathbf{y}} : \mathbf{x} \mapsto \langle \mathbf{x}, \mathbf{y} \rangle\}_{\mathbf{y} \in \mathcal{Y}}$, with $\mathcal{Y} = \mathcal{M}$ and decryption is correct when the function value falls within $(-p/2, p/2]$. In other words, it

⁴ Called strong XiO in [BNPW16, BNPW20].

computes the inner product between a message $\mathbf{x} \in \mathcal{M}$ and a vector $\mathbf{y} \in \mathcal{Y}$ encoding the function of interest. Our scheme takes a very simple form, consisting of the following:

- The public key is a matrix $\mathbf{A} \in \mathbb{Z}_q^{(k+n) \times m}$ where $m \gg n + k$, sampled together with a trapdoor. As usual, without knowing the trapdoor, $\mathbf{A}$ is distributed close to uniform.
- The master secret key is a low-norm matrix $\mathbf{K} \in \mathbb{Z}^{m \times k}$, sampled using the trapdoor of $\mathbf{A}$, such that $\mathbf{A} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$.
- The ciphertext is an LWE sample where we encode a message $\mathbf{x} \in \mathcal{M}$ as part of the LWE secret, taking the form
$$\mathbf{c}^T = [\mathbf{x}^T \mathbf{s}^T] \cdot \mathbf{A} + p \cdot \mathbf{e}^T \bmod q. \quad (1)$$
- The functional key of a low-norm $\mathbf{y} \in \mathcal{Y}$ representing $f_{\mathbf{y}}$ is $\mathbf{k}_{\mathbf{y}} = \mathbf{K} \cdot \mathbf{y}$.

Decryption is straightforward: simply compute $\mathbf{c}^T \cdot \mathbf{k}_{\mathbf{y}} = \langle \mathbf{x}, \mathbf{y} \rangle + p \cdot \langle \mathbf{e}, \mathbf{k}_{\mathbf{y}} \rangle \bmod q$. Note that $p \cdot \langle \mathbf{e}, \mathbf{k}_{\mathbf{y}} \rangle$ is a low-norm noise term. By lifting the value to $\mathbb{Z}$ (in balanced representation) and then reducing modulo p , this noise term can be removed to recover the function value $\langle \mathbf{x}, \mathbf{y} \rangle$ if it falls in the range $(-p/2, p/2]$. We give in Appendix F the detailed linear FE scheme.

k-LWE. To argue about security, we use the k -LWE assumption which states that: Given a wide matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with $k < m$ hints $\mathbf{K} = (\mathbf{k}_1, \dots, \mathbf{k}_k)$, each being a Gaussian vector sampled over the kernel lattice $\Lambda_q^\perp(\mathbf{A})$ of $\mathbf{A}$, i.e. $\mathbf{A} \cdot \mathbf{K} = \mathbf{0}_{n \times k} \bmod q$, the following are computationally indistinguishable:

Case $b = 0$: an LWE sample $\mathbf{s}^T \cdot \mathbf{A} + \mathbf{e}^T \bmod q$, and

Case $b = 1$: a vector $\mathbf{c} = \mathbf{b} + \mathbf{e} \bmod q$, where $\mathbf{b}$ is uniformly randomly sampled subject to $\mathbf{b}^T \cdot \mathbf{K} = \mathbf{0} \bmod q$, and $\mathbf{e}$ is an error (identically distributed as that in the above LWE sample).

To interpret this, think of that given the hints $\mathbf{K}$, one can easily distinguish an LWE sample from uniform by right-multiplication: it holds $(\mathbf{s}^T \cdot \mathbf{A} + \mathbf{e}^T) \cdot \mathbf{K} \bmod q = \mathbf{e}^T \cdot \mathbf{K}$ which is low-norm. The assumption, in a sense, asserts that not much more can be done: in the second case above, similarly $\mathbf{c}^T \cdot \mathbf{K} \bmod q = \mathbf{e}^T \cdot \mathbf{K}$. In particular, the assumption states that, in presence of the hints $\mathbf{K}$, the LWE sample is still indistinguishable from a noisy random sample over a space $\{\mathbf{b} : \mathbf{b}^T \cdot \mathbf{K} = \mathbf{0} \bmod q\}$ which is much larger than the row-span of $\mathbf{A}$. Under appropriate parameters, the k -LWE assumption is known to be implied by the LWE assumption [LPSS14].

k-ALWE. To simplify our security argument, we show that k -LWE implies a variant of itself in which the hints are short preimages of a non-zero target, which we call the affine k -LWE (k -ALWE) assumption. This variant states that: Given $\mathbf{A} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$, the tuple $(\mathbf{c}, \mathbf{K})$ computed as follows are computationally indistinguishable:

- Case $b = 0$:** $\mathbf{c}^T = [\mathbf{x}^T \mathbf{s}^T] \cdot \mathbf{A} + p \cdot \mathbf{e}^T \bmod q$ is an LWE sample, with an adversarially chosen $\mathbf{x} \leftarrow \mathcal{A}(\mathbf{A})$ encoded as part of the LWE secret, and $\mathbf{K}$ is Gaussian subject to satisfying $\mathbf{A} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$.
- Case $b = 1$:** $\mathbf{c}^T = \mathbf{b}^T + p \cdot \mathbf{e}^T \bmod q$, where $\mathbf{b}$ is uniformly random, $\mathbf{e}$ is an error, and $\mathbf{K}$ is Gaussian satisfying jointly

$$\begin{bmatrix} \mathbf{A} \\ \mathbf{b}^T \end{bmatrix} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{x}^T \end{bmatrix} \bmod q. \quad (2)$$

The intuition is analogous to k -LWE – upon right-multiplying $\mathbf{K}$ to either $\mathbf{A}$ or $\mathbf{c}$ in either case, one observes the same outcomes, namely $\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}$ and a noisy $\mathbf{x}^T$.

$k\text{-LWE} \implies k\text{-ALWE}$. The proof proceeds in two steps:

1. We give a reduction from $k\text{-LWE}$ showing that, the real experiment is indistinguishable from a hybrid, where $\mathbf{K}$ is unchanged, i.e. Gaussian subject to $\mathbf{A} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$, but $\mathbf{c} = \mathbf{b} + p \cdot \mathbf{e} \bmod q$ where $\mathbf{b}$ is sampled randomly subject to $\mathbf{b}^\top \cdot \mathbf{K} = \mathbf{x}^\top \bmod q$. The reduction is straightforward: Given a $k\text{-LWE}$ instance $\mathbf{s}^\top \cdot \mathbf{A}_1 + \mathbf{e}^\top \bmod q$ and upon receiving a message $\mathbf{x}$ from an adversary, add $\mathbf{x} \cdot \mathbf{A}_0 \bmod q$ to obtain $\mathbf{x}^\top \cdot \mathbf{A}_0 + \mathbf{s}^\top \cdot \mathbf{A}_1 + \mathbf{e}^\top = [\mathbf{x}^\top \ \mathbf{s}^\top] \cdot \mathbf{A}' + \mathbf{e}^\top \bmod q$, i.e. $\mathbf{A}_0$ and $\mathbf{A}_1$ are the “top” and “bottom” parts of $\mathbf{A}'$ respectively. Picking q coprime with p and multiplying by p , we obtain $\mathbf{A} = p \cdot \mathbf{A}'$ which is again uniformly random, and the simulated LWE samples are $p \cdot ([\mathbf{x}^\top \ \mathbf{s}^\top] \cdot \mathbf{A}' + \mathbf{e}^\top) = [\mathbf{x}^\top \ \mathbf{s}^\top] \mathbf{A} + p \cdot \mathbf{e}^\top \bmod q$. That the simulation by the reduction is faithful can be shown by a statistical “reverse-sampling” argument originally from [BF11]. Roughly, a consequence of this argument is that the ordering of sampling $\mathbf{A}, \mathbf{K}, \mathbf{b}$ subject to $\begin{bmatrix} \mathbf{A} \\ \mathbf{b}^\top \end{bmatrix} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{x}^\top \end{bmatrix} \bmod q$ does not matter – they are all statistically close.
2. We apply the reverse-sampling argument of [BF11] again, swapping the ordering of sampling to first sampling $\mathbf{A}, \mathbf{b}$ at random, then sampling Gaussian $\mathbf{K}$ satisfying the constraint. This is precisely the ideal distribution in the $k\text{-ALWE}$ assumption stated above.

$k\text{-ALWE} \implies \text{Security of linear FE}$. Given the $k\text{-ALWE}$ assumption, we can prove security of our linear FE scheme as follows.

1. First swap $\mathbf{A}$ to a random matrix sampled without a trapdoor. This hop follows immediately from the properties of lattice trapdoors.
2. Next, in challenge branch $b \in \{0, 1\}$, swap the challenge ciphertext $[\mathbf{x}_b^\top \ \mathbf{s}^\top] \cdot \mathbf{A} + p \cdot \mathbf{e}^\top \bmod q$ to $\mathbf{c} = \mathbf{b} + p \cdot \mathbf{e} \bmod q$ where $\mathbf{b}$ is uniformly random, and let the functional keys be computed as $\mathbf{K} \cdot \mathbf{y}$ where $\mathbf{K}$ is Gaussian subject to the constraint Eq. (2), instantiated with $\mathbf{x} = \mathbf{x}_b$. This hop follows immediately from the $k\text{-ALWE}$ assumption stated above. At this point, the challenge ciphertext $\mathbf{c} = \mathbf{b} + p \cdot \mathbf{e} \bmod q$ is independent of the challenge bit.
3. We handle the functional keys by projecting them onto the lattice of valid queries. Let $\boldsymbol{\delta} := \mathbf{x}_1 - \mathbf{x}_0$. If $\boldsymbol{\delta} = \mathbf{0}$, the two key constraints are identical and the claim follows directly. Otherwise, let $\mathbf{Q}$ be a $\mathbb{Z}$ -basis of $H := \{\mathbf{y} \in \mathbb{Z}^k : \boldsymbol{\delta}^\top \mathbf{y} = 0\}$. Applying the Gaussian linear-transformation lemma of [GMPW20] to $\mathbf{K} \mapsto \mathbf{K}\mathbf{Q}$ shows that $\mathbf{K}\mathbf{Q}$ is statistically close to a Gaussian over the corresponding projected coset. This distribution is independent of the challenge bit because $\mathbf{x}_0^\top \mathbf{Q} = \mathbf{x}_1^\top \mathbf{Q}$. Moreover, every valid query can be written as $\mathbf{y} = \mathbf{Q}\mathbf{a}$ for some $\mathbf{a}$, and hence $\mathbf{K}\mathbf{y} = (\mathbf{K}\mathbf{Q})\mathbf{a}$. Thus, the joint distributions of all valid functional keys are statistically close.

This valid-query projection is formalized for general degree d in Lemma 15.

2.2 Generalising Linear FE to Quadratic FE

With the above linear FE and security arguments in mind, our goal is to carefully generalise the linear FE and the $k\text{-ALWE}$ assumption to an appropriate analogue in the degree- d setting, so that we can construct a degree- d FE scheme and argue about its security using a degree- d analogue of the $k\text{-ALWE}$ assumption. In this overview, we focus on $d = 2$ and defer the case with general d to later sections. Note that while the scheme described below is (functionally) correct in the public-key setting, we will only prove its security in the secret-key setting, where the LWE matrices $\mathbf{A}$ and $\tilde{\mathbf{A}}$ (to be introduced) are hidden from the adversary. The reason is an attack that succeeds against the public-key setting in our experiments, described in Section 7.2.

Our goal is to construct an FE for the functionality $f_{\mathbf{y}} : \mathbf{x} \mapsto \langle \mathbf{x}^{\otimes 2}, \mathbf{y} \rangle$ with ciphertext size sub-quadratic in the dimension of $\mathbf{x}$. Decryption is correct when the function value falls within $(-p^2/2, p^2/2]$. In a nutshell, our idea is *self-tensoring*.

Self-Tensor. Recall that the linear FE ciphertext is $\mathbf{c}^\top = [\mathbf{x}^\top \mathbf{s}^\top] \cdot \mathbf{A} + p \cdot \mathbf{e}^\top \bmod q$. Self-tensoring gives

$$\begin{aligned} (\mathbf{c}^\top)^{\otimes 2} &= [\mathbf{x}^\top \mathbf{s}^\top]^{\otimes 2} \cdot \mathbf{A}^{\otimes 2} + ([\mathbf{x}^\top \mathbf{s}^\top] \otimes (p \cdot \mathbf{e}^\top)) \cdot (\mathbf{A} \otimes \mathbf{I}_m) \\ &\quad + p^2 \cdot (\mathbf{e}^\top)^{\otimes 2} + ((p \cdot \mathbf{e}^\top) \otimes [\mathbf{x}^\top \mathbf{s}^\top]) \cdot (\mathbf{I}_m \otimes \mathbf{A}) \bmod q. \end{aligned}$$

Suppose we make the master secret key $\mathbf{K}$ taller and wider, so that $\mathbf{K} \in \mathbb{Z}^{m^2 \times k^2}$ is Gaussian satisfying

$$\mathbf{A}^{\otimes 2} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes 2} \bmod q. \quad (3)$$

Right-multiplying $\mathbf{k}_\mathbf{y} = \mathbf{K} \cdot \mathbf{y}$ to $(\mathbf{c}^\top)^{\otimes 2}$ would then yield a sum with $\langle \mathbf{x}^{\otimes 2}, \mathbf{y} \rangle$ appearing as one of the terms, as desired. However, for decryption correctness the decryptor must somehow remove the cross terms

$$([\mathbf{x}^\top \mathbf{s}^\top] \otimes (p \cdot \mathbf{e}^\top)) \cdot (\mathbf{A} \otimes \mathbf{I}_m) \cdot \mathbf{k}_\mathbf{y} + ((p \cdot \mathbf{e}^\top) \otimes [\mathbf{x}^\top \mathbf{s}^\top]) \cdot (\mathbf{I}_m \otimes \mathbf{A}) \cdot \mathbf{k}_\mathbf{y}, \quad (4)$$

which, without additional constraints on $\mathbf{K}$, are likely not low-norm since $\mathbf{A}$ and $\mathbf{s}$ are uniformly random.

Handling Cross Terms. To address this issue, an immediate idea is to restrict $\mathbf{s}$ to be low-norm (e.g. Gaussian, which can similarly be applied in the linear setting), i.e. to adopt the so-called normal-form LWE setting. This should not affect security at least heuristically since normal-form LWE is well-known to be as hard as LWE. This alone, however, is not sufficient to suppress the cross terms since $\mathbf{A}$ is uniformly random.

To move further, it is tempting to require $\mathbf{K}$ to additionally satisfy $(\mathbf{I}_m \otimes \mathbf{A}) \cdot \mathbf{K} + p \cdot \mathbf{F}_{10} = \mathbf{0} \bmod q$ and $(\mathbf{A} \otimes \mathbf{I}_m) \cdot \mathbf{K} + p \cdot \mathbf{F}_{01} = \mathbf{0} \bmod q$ for some low norm matrices $\mathbf{F}_{10}, \mathbf{F}_{01}$, which would force the cross-terms in Eq. (4) to be low-norm. Unfortunately, such a scheme is insecure, since the additional constraints on $\mathbf{K}$ imply that a collection of $\mathbf{k}_\mathbf{y}$'s is an approximate trapdoor of $\mathbf{A}$, using which it is easy to distinguish LWE samples w.r.t. $\mathbf{A}$ with low-norm secrets.⁵

Our idea of solving the cross-term problem is to introduce additional uniformly random masks. More precisely, we introduce an additional uniform matrix $\tilde{\mathbf{A}} = (\tilde{\mathbf{A}}_{00}, \tilde{\mathbf{A}}_{10}, \tilde{\mathbf{A}}_{01})$ (stacked vertically) of large enough width $\tilde{m}$ to the public key and a low norm matrix $\tilde{\mathbf{K}} \in \mathbb{Z}^{\tilde{m} \times k^2}$ to the secret key such that

$$\underbrace{\begin{bmatrix} \mathbf{A} \otimes \mathbf{A} & \tilde{\mathbf{A}}_{00} \\ \mathbf{I}_m \otimes \mathbf{A} & \tilde{\mathbf{A}}_{10} & p \cdot \mathbf{I} \\ \mathbf{A} \otimes \mathbf{I}_m & \tilde{\mathbf{A}}_{01} & p \cdot \mathbf{I} \end{bmatrix}}_{\mathcal{E}^{\otimes 2}(\mathbf{A}, \tilde{\mathbf{A}})} \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F}_{10} \\ \mathbf{F}_{01} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes 2} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bmod q. \quad (5)$$

Correspondingly, we let the ciphertext additionally include

$$\begin{aligned} \tilde{\mathbf{c}}^\top &= [\mathbf{x}^\top \mathbf{s}^\top]^{\otimes 2} \cdot \tilde{\mathbf{A}}_{00} + ([\mathbf{x}^\top \mathbf{s}^\top] \otimes (p \cdot \mathbf{e}^\top)) \cdot \tilde{\mathbf{A}}_{01} \\ &\quad + p^2 \cdot \tilde{\mathbf{e}}^\top + ((p \cdot \mathbf{e}^\top) \otimes [\mathbf{x}^\top \mathbf{s}^\top]) \cdot \tilde{\mathbf{A}}_{10} \bmod q, \end{aligned} \quad (6)$$

which intuitively serves as a linear FE ciphertext of the cross terms. Note that a ciphertext consists of $m + \tilde{m}$ elements in $\mathbb{Z}_q$. As long as $m + \tilde{m} \ll k^2$, the ciphertext size is sub-quadratic in the dimension of $\mathbf{x}$.

Now, if we extend the functional secret key to $(\mathbf{k}_\mathbf{y}, \tilde{\mathbf{k}}_\mathbf{y}) = (\mathbf{K} \cdot \mathbf{y}, \tilde{\mathbf{K}} \cdot \mathbf{y})$, then the function value $\langle \mathbf{x}^{\otimes 2}, \mathbf{y} \rangle$ can be recovered from the ciphertext as follows: Compute $(\mathbf{c}^\top)^{\otimes 2} \cdot \mathbf{k}_\mathbf{y} + \tilde{\mathbf{c}}^\top \cdot \tilde{\mathbf{k}}_\mathbf{y} = \langle \mathbf{x}^{\otimes 2}, \mathbf{y} \rangle + p^2 \cdot e \bmod q$, where the noise term

$$\begin{aligned} e &= (\mathbf{e}^\top)^{\otimes 2} \cdot \mathbf{k}_\mathbf{y} - ([\mathbf{x}^\top \mathbf{s}^\top] \otimes \mathbf{e}^\top) \cdot \mathbf{F}_{01} \cdot \mathbf{y} \\ &\quad + \tilde{\mathbf{e}}^\top \cdot \tilde{\mathbf{k}}_\mathbf{y} - (\mathbf{e}^\top \otimes [\mathbf{x}^\top \mathbf{s}^\top]) \cdot \mathbf{F}_{10} \cdot \mathbf{y} \end{aligned}$$

is low norm. Lifting the result to $\mathbb{Z}$ and reducing modulo p^2 recovers $\langle \mathbf{x}^{\otimes 2}, \mathbf{y} \rangle$.

In hindsight, our quadratic FE construction structurally resembles the pairing-based construction of [Wee20, Appendix A].

⁵ For example, from the approximation $(\mathbf{I}_m \otimes \mathbf{A}) \cdot \mathbf{k}_i \approx \mathbf{0} \bmod q$, where $\mathbf{k}_i$ denotes the i -th column of $\mathbf{K}$, one can derive $\mathbf{A} \cdot \mathbf{K}_i \approx \mathbf{0} \bmod q$, where $\mathbf{K}_i$ is the $(m \times m)$ -matrix such that $\text{vec}(\mathbf{K}_i) = \mathbf{k}_i$, which in turn corresponds to an approximate trapdoor of $\mathbf{A}$.

Further extensions. In the general degree- d case, the identity matrices at the diagonal of $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$ ($d = 2$ case given in Eq. (5)) will be multiplied with different powers of p . This is to make sure that in the decryption equation all the noise terms are multiples of p^d . For each block row, the power of p is equal to the number of times $\mathbf{A}$ appears as a Kronecker factor in the first block column. For example, for the row where the first block column is $\mathbf{A}^{\otimes i} \otimes \mathbf{I}_m^{\otimes (d-i)}$, we will have p^i times some identity matrix at the diagonal block.

In the above, we were implicitly assuming that $q \gg p^2$ so that lifting $(\mathbf{c}^\top)^{\otimes 2} \cdot \mathbf{k}_y + \tilde{\mathbf{c}}^\top \cdot \tilde{\mathbf{k}}_y \bmod q$ to $\mathbb{Z}$ followed by reduction modulo p^2 yields the correct result. To handle $p^d \gg q$, our observation is that it suffices to reduce the lifted result modulo $r = p^d \bmod q$, where r is taken as the smallest representative. If $r \ll q$, then decryption would still succeed if the value $f(\mathbf{x})$ falls into $(-r/2, r/2] \cap \mathbb{Z}$.

2.3 Generalising k -ALWE to the Quadratic Setting

By extending the k -ALWE assumption in the same way we extended the linear FE scheme to a quadratic FE scheme, we would obtain a quadratic variant of the k -ALWE assumption against which, as we will detail in Section 7, one of our attacks succeeds in experiments. We instead consider the more conservative, secret-key $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption (Assumption 1), where the LWE matrices $\mathbf{A}$ and $\tilde{\mathbf{A}}$ are not given to the adversary.

In more detail, consider a pair of experiments parametrised by $b \in \{0, 1\}$. In both, the adversary $\mathcal{A}(1^\lambda)$ specifies in advance a tuple of low-norm vectors $\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}^*$. It then receives a tuple $((\mathbf{d}_i, \tilde{\mathbf{d}}_i)_{i \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}}, \mathbf{K}, \tilde{\mathbf{K}})$ sampled in one of the following ways:

Case $b = 0$: Sample the random matrices $\mathbf{A}, \tilde{\mathbf{A}}$. Sample $(\mathbf{d}_i, \tilde{\mathbf{d}}_i)_{i \in [Q]}$ and $(\mathbf{c}, \tilde{\mathbf{c}})$ as in Eqs. (1) and (6), where $(\mathbf{d}_i, \tilde{\mathbf{d}}_i)$ encodes $\mathbf{x}_i$ and $(\mathbf{c}, \tilde{\mathbf{c}})$ encodes $\mathbf{x}^*$ as part of the LWE secrets. The matrices $(\mathbf{K}, \tilde{\mathbf{K}})$ are Gaussian subject to the constraints in Eq. (5).

Case $b = 1$: Sample the random matrices $\mathbf{B}, \tilde{\mathbf{B}} = (\tilde{\mathbf{B}}_{00}, \tilde{\mathbf{B}}_{10}, \tilde{\mathbf{B}}_{01})$. Sample $(\mathbf{d}_i, \tilde{\mathbf{d}}_i)_{i \in [Q]}$ and $(\mathbf{c}, \tilde{\mathbf{c}})$ as follows:

$$\mathbf{d}_i^\top = [\mathbf{x}_i^\top \ \mathbf{s}_i^\top \ 0] \cdot \mathbf{B} + p \cdot \mathbf{e}_i^\top \bmod q \quad (7)$$

$$\begin{aligned} \tilde{\mathbf{d}}_i^\top &= [\mathbf{x}_i^\top \ \mathbf{s}_i^\top \ 0]^{\otimes 2} \cdot \tilde{\mathbf{B}}_{00} + ([\mathbf{x}_i^\top \ \mathbf{s}_i^\top \ 0] \otimes (p \cdot \mathbf{e}_i^\top)) \cdot \tilde{\mathbf{B}}_{01} \\ &\quad + p^2 \cdot \tilde{\mathbf{e}}_i^\top + ((p \cdot \mathbf{e}_i^\top) \otimes [\mathbf{x}_i^\top \ \mathbf{s}_i^\top \ 0]) \cdot \tilde{\mathbf{B}}_{10} \bmod q, \end{aligned} \quad (8)$$

$$\mathbf{c}^\top = [\mathbf{0}^\top \ \mathbf{s}^\top \ 1] \cdot \mathbf{B} + p \cdot \mathbf{e}^\top \bmod q \quad (9)$$

$$\begin{aligned} \tilde{\mathbf{c}}^\top &= [\mathbf{0}^\top \ \mathbf{s}^\top \ 1]^{\otimes 2} \cdot \tilde{\mathbf{B}}_{00} + ([\mathbf{0}^\top \ \mathbf{s}^\top \ 1] \otimes (p \cdot \mathbf{e}^\top)) \cdot \tilde{\mathbf{B}}_{01} \\ &\quad + p^2 \cdot \tilde{\mathbf{e}}^\top + ((p \cdot \mathbf{e}^\top) \otimes [\mathbf{0}^\top \ \mathbf{s}^\top \ 1]) \cdot \tilde{\mathbf{B}}_{10} \bmod q. \end{aligned} \quad (10)$$

The matrices $(\mathbf{K}, \tilde{\mathbf{K}})$ are Gaussian subject to

$$\underbrace{\begin{bmatrix} \mathbf{B} \otimes \mathbf{B} & \tilde{\mathbf{B}}_{00} \\ \mathbf{I}_m \otimes \mathbf{B} & \tilde{\mathbf{B}}_{10} & p \cdot \mathbf{I} \\ \mathbf{B} \otimes \mathbf{I}_m & \tilde{\mathbf{B}}_{01} & p \cdot \mathbf{I} \end{bmatrix}}_{\mathcal{E}^{\otimes 2}(\mathbf{B}, \tilde{\mathbf{B}}) :=} \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G}_{10} \\ \mathbf{G}_{01} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{0}^\top \end{bmatrix}^{\otimes 2} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} + \begin{bmatrix} \begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ \mathbf{x}^{*\top} \end{bmatrix}^{\otimes 2} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \bmod q, \quad (11)$$

for some Gaussian matrices $\mathbf{G}_{10}, \mathbf{G}_{01}$.

The assumption states that the two cases are indistinguishable by any PPT $\mathcal{A}$.

We make two observations about the above experiment. First, despite the more complex expressions of $(\mathbf{d}_i, \tilde{\mathbf{d}}_i)_{i \in [Q]}$ in the $b = 1$ case, their distributions in the two cases $b \in \{0, 1\}$ are in fact identical. This can be observed by thinking of $\mathbf{A}, \tilde{\mathbf{A}}$ as appropriately indexed submatrices of $\mathbf{B}, \tilde{\mathbf{B}}$, respectively. If the matrices $\mathbf{A}, \tilde{\mathbf{A}}$ were given to $\mathcal{A}$, as in the degree-1 case, then the tuples $(\mathbf{d}_i, \tilde{\mathbf{d}}_i)_{i \in [Q]}$ could have been computed by $\mathcal{A}$ itself. Second, the tuple $(\mathbf{c}, \tilde{\mathbf{c}})$ is sampled with some randomness $\mathbf{s}$ in both $b \in \{0, 1\}$, whereas for degree $d = 1$ only in the $b = 0$ case does $\mathbf{c}$ depend on $\mathbf{s}$. On the one hand, the introduction of $\mathbf{s}$ in the $b = 1$ case makes it so that the noise term that can be recovered by the adversary (which we will discuss below) in both

cases contains contributions from $\mathbf{s}$, making attacking the assumption more difficult. On the other hand, we note that due to the simplicity of the $d = 1$ case, whether to introduce $\mathbf{s}$ or not when $b = 1$ is equivalent when $d = 1$ but not when $d > 1$.

At present, we believe that obtaining a reduction from more standard assumptions to the above is very challenging. Ultimately, this is because our assumption concerns about discrete Gaussian vectors sampled subject to non-linear constraints, whereas the majority of the lattice-based cryptography literature only handles linear transformations of discrete Gaussian (or subgaussian) distributions. For example, it seems highly non-trivial to generalise the “reverse-sampling” technique from [BF11] that we used to prove k -ALWE from k -LWE to the non-linear setting. We hope that our result motivates future studies on non-linear lattice-based cryptography. For now, we focus on understanding the assumption from a cryptanalytic perspective.

Avoiding attacks based on error term analysis. Apart from obvious strategies such as solving LWE directly, the most natural strategy for a distinguisher to violate the assumption (which would potentially translate to an attack on the FE scheme) is to examine the statistics of (some function of) the error term

$$p^2 \cdot \mathbf{f}^T = \langle \mathbf{c}^{\otimes 2}, \mathbf{K} \rangle + \langle \tilde{\mathbf{c}}, \tilde{\mathbf{K}} \rangle - ((\mathbf{x}^*)^{\otimes 2})^T \bmod q.$$

Denoting the error term for Case b by $\mathbf{f}_b$, we have

$$\begin{aligned} \mathbf{f}_0^T &= \langle \mathbf{e}^{\otimes 2}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle - \mathbf{e}^T \cdot \left(\left\langle \begin{bmatrix} \mathbf{x}^* \\ \mathbf{s} \end{bmatrix} \otimes \mathbf{I}_m, \mathbf{F}_{01} \right\rangle + \left\langle \mathbf{I}_m \otimes \begin{bmatrix} \mathbf{x}^* \\ \mathbf{s} \end{bmatrix}, \mathbf{F}_{10} \right\rangle \right), \\ \mathbf{f}_1^T &= \langle \mathbf{e}^{\otimes 2}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle - \mathbf{e}^T \cdot \left(\left\langle \begin{bmatrix} \mathbf{0} \\ \mathbf{s} \\ 1 \end{bmatrix} \otimes \mathbf{I}_m, \mathbf{G}_{01} \right\rangle + \left\langle \mathbf{I}_m \otimes \begin{bmatrix} \mathbf{0} \\ \mathbf{s} \\ 1 \end{bmatrix}, \mathbf{G}_{10} \right\rangle \right), \end{aligned}$$

where we use an extended inner product notation for matrices, i.e. $\langle \mathbf{A}, \mathbf{B} \rangle := \mathbf{A}^T \cdot \mathbf{B}$. Observe that $\mathbf{f}_0$ depends on the challenge message $\mathbf{x}^*$ chosen by the adversary while $\mathbf{f}_1$ does not. This means an adversary may attempt to distinguish by choosing an $\mathbf{x}$ such that $g(\mathbf{f}_0)$ and $g(\mathbf{f}_1)$ behave significantly differently for some function g . For example, a choice of g inspired by the attack against lattice-based FE in [TÜ24] would be $g(\cdot) = \|\cdot\|^2$. To resist this, we insist any $\mathbf{x}$ chosen by the adversary to be of ℓ_∞ -norm bounded by some β_x , and the Gaussian parameter of $\mathbf{s}$ is subexponentially larger than $\sqrt{k} \cdot \beta_x$. This way, any contribution from $\mathbf{x}$ to the error term would be masked away by that from $\mathbf{s}$.

In Section 5 we give the formal details of our new assumption for general degree $d \geq 1$, which captures the above as special case with $d = 2$. We discuss more unsuccessful attack strategies on the assumption in Section 7. We note that our assumption is falsifiable (Remark 1).

Our assumption $\implies$ Security of secret-key quadratic FE. We sketch a security proof of the secret-key variant of our quadratic FE from the above assumption. The proof structure is similar to that of our (public-key) linear FE. It consists of the following core steps:

1. First swap $\mathbf{A}, \tilde{\mathbf{A}}$ to random matrices sampled without a trapdoor. This hop follows immediately from trapdoor properties.
2. Next, extend the matrices $\mathbf{A}, \tilde{\mathbf{A}}$ to $\mathbf{B}, \tilde{\mathbf{B}}$, carrying one further row, so that the corresponding submatrices satisfy the constraints as defined in the first step, and simultaneously swap the challenge ciphertext $(\mathbf{c}, \tilde{\mathbf{c}})$ sampled as in Eqs. (1) and (6) to

$$\begin{aligned} \mathbf{c}^T &= [\mathbf{0}^T \ \mathbf{s}^T \ 1] \cdot \mathbf{B} + p \cdot \mathbf{e}^T \bmod q, \\ \tilde{\mathbf{c}}^T &= [\mathbf{0}^T \ \mathbf{s}^T \ 1]^{\otimes 2} \cdot \tilde{\mathbf{B}}_{00} + ([\mathbf{0}^T \ \mathbf{s}^T \ 1] \otimes p \cdot \mathbf{e}^T) \cdot \tilde{\mathbf{B}}_{01} \\ &\quad + p^2 \cdot \tilde{\mathbf{e}}^T + (p \cdot \mathbf{e}^T \otimes [\mathbf{0}^T \ \mathbf{s}^T \ 1]) \cdot \tilde{\mathbf{B}}_{10} \bmod q. \end{aligned} \tag{12}$$

For each challenge branch $b \in \{0, 1\}$, we apply the new assumption with $\mathbf{x}^* = \mathbf{x}_b^*$. The functional keys are computed as $(\mathbf{K} \cdot \mathbf{y}, \tilde{\mathbf{K}} \cdot \mathbf{y})$, where $(\mathbf{K}, \tilde{\mathbf{K}})$ are sampled subject to the constraint in Eq. (11), instantiated with $\mathbf{x}^* = \mathbf{x}_b^*$. This hop follows from the new assumption.

3. We handle the functional keys by projecting the Gaussian key matrices onto the lattice of valid queries. Let $\delta := (\mathbf{x}_1^*)^{\otimes 2} - (\mathbf{x}_0^*)^{\otimes 2}$. If $\delta = \mathbf{0}$, the two targets in Eq. (11) are identical, and the claim follows directly. Suppose therefore that $\delta \neq \mathbf{0}$. Let $\mathbf{Q}$ be a $\mathbb{Z}$ -basis of $H := \{\mathbf{y} \in \mathbb{Z}^{k^2} : \delta^T \mathbf{y} = 0\}$. We apply the Gaussian linear-transformation lemma of [GMPW20] to the map

$$[\mathbf{K}; \tilde{\mathbf{K}}; \mathbf{G}_{10}; \mathbf{G}_{01}] \mapsto [\mathbf{K}; \tilde{\mathbf{K}}; \mathbf{G}_{10}; \mathbf{G}_{01}] \mathbf{Q}$$

where $;$ denotes vertical concatenation. With the chosen parameters, the resulting distribution is statistically close to a Gaussian sampled directly subject to the projected constraint in Eq. (11). This projected constraint is independent of the challenge bit because

$$((\mathbf{x}_0^*)^{\otimes 2})^T \mathbf{Q} = ((\mathbf{x}_1^*)^{\otimes 2})^T \mathbf{Q}.$$

Every valid query $\mathbf{y}$ belongs to H and can therefore be written uniquely as $\mathbf{y} = \mathbf{Q}\mathbf{a}$. Its functional key satisfies

$$(\mathbf{K}\mathbf{y}, \tilde{\mathbf{K}}\mathbf{y}) = ((\mathbf{K}\mathbf{Q})\mathbf{a}, (\tilde{\mathbf{K}}\mathbf{Q})\mathbf{a}).$$

Thus, the joint distributions of all valid functional key answers in the two challenge branches are statistically indistinguishable.

2.4 Generating Trapdoors for Tensor Matrices

To complete our quadratic FE, the last missing piece is a way to efficiently generate the master secret key, i.e. the matrices $(\mathbf{K}, \tilde{\mathbf{K}})$ satisfying the constraints in Eq. (5). Below, we describe how to generate the matrix $\mathcal{E}^{\otimes 2}(\mathbf{A}, \tilde{\mathbf{A}})$ in Eq. (5) together with its lattice trapdoor, which allows generating $(\mathbf{K}, \tilde{\mathbf{K}})$. Its generalisation to any constant-degree $d \geq 1$ will be detailed in Section 4.

To begin, we sample $\mathbf{A} \in \mathbb{Z}_q^{(k+n) \times m}$ together with its ordinary lattice trapdoor. This is possible using standard techniques when m is sufficiently larger than $k+n$. To support degree- d polynomials, the modulus must accommodate d -fold products of preimages, so $\log q = \Omega(d \log m)$, and sampling with the trapdoor in turn needs $m = \Omega((k+n) \log q) = \Omega(d(k+n) \log m)$.

Given a trapdoor for $\mathbf{A}$, we can sample full-rank Gaussian matrices $\mathbf{U}_0 \in \mathbb{Z}^{m \times m}$ and $\mathbf{U}_1 \in \mathbb{Z}^{m \times (k+n)}$ such that

$$\mathbf{A} \cdot [\mathbf{U}_0 \ \mathbf{U}_1] = [\mathbf{0}_{(k+n) \times m} \ p \cdot \mathbf{I}_{k+n}] \bmod q.$$

Write $\mathcal{E}^{\otimes 2}(\mathbf{A})$ for the matrix obtained by removing the masking block column, i.e. the “ $\tilde{\mathbf{A}}$ column”, of $\mathcal{E}^{\otimes 2}(\mathbf{A}, \tilde{\mathbf{A}})$. Once we have obtained a trapdoor of $\mathcal{E}^{\otimes 2}(\mathbf{A})$, we can use standard trapdoor extension techniques to obtain a trapdoor for $\mathcal{E}^{\otimes 2}(\mathbf{A}, \tilde{\mathbf{A}})$. We construct an Ajtai trapdoor [Ajt96] $\mathbf{T}_2$ of $\mathcal{E}^{\otimes 2}(\mathbf{A})$ as follows:

$$\underbrace{\begin{bmatrix} \mathbf{A} \otimes \mathbf{A} \\ \mathbf{I}_m \otimes \mathbf{A} & p \cdot \mathbf{I}_{m(k+n)} \\ \mathbf{A} \otimes \mathbf{I}_m & p \cdot \mathbf{I}_{(k+n)m} \end{bmatrix}}_{\mathcal{E}^{\otimes 2}(\mathbf{A}) :=}, \quad \underbrace{\begin{bmatrix} \mathbf{U}_0 \otimes \mathbf{U}_0 & \mathbf{U}_0 \otimes \mathbf{U}_1 & \mathbf{U}_1 \otimes \mathbf{U}_0 \\ & -\mathbf{U}_0 \otimes \mathbf{I}_{k+n} & \\ & & -\mathbf{I}_{k+n} \otimes \mathbf{U}_0 \end{bmatrix}}_{\mathbf{T}_2}. \quad (13)$$

We see that $\mathbf{T}_2$ is indeed an Ajtai trapdoor of $\mathcal{E}^{\otimes 2}(\mathbf{A})$, that is:

- $\mathcal{E}^{\otimes 2}(\mathbf{A}) \cdot \mathbf{T}_2 = \mathbf{0}_{((k+n+m)^2 - m^2) \times ((k+n+m)^2 - (k+n)^2)} \bmod q$,
- $\mathbf{T}_2$ is low norm relative to q , and
- $\mathbf{T}_2$ is invertible over $\mathbb{Q}$.

The first two items are straightforward to verify, and we see that $\mathbf{T}_2$ is invertible over $\mathbb{Q}$ because it is block-triangular where the diagonal entries are Kronecker products between the invertible matrices $\mathbf{I}_{k+n}$ and $\mathbf{U}_0$. Finally, for the above to be sufficient for our FE application, we also ensure that $\mathcal{E}^{\otimes 2}(\mathbf{A})$ is primitive w.h.p., i.e. $\mathbb{Z}_q$ -linear combinations of its columns generate $\mathbb{Z}_q^{(k+n+m)^2 - m^2}$: We have that $\mathbf{A} \otimes \mathbf{A}$ is primitive since $\mathbf{A}$ is, then, due to the identity matrices on the bottom right, $\mathcal{E}^{\otimes 2}(\mathbf{A})$ is also primitive.

3 Preliminaries

Let $\lambda \in \mathbb{N}$ denote the security parameter, and $\text{poly}(\lambda)$ and $\text{negl}(\lambda)$ the set of all polynomials and negligible functions in λ , respectively. For $k, n \in \mathbb{N}$, $k \leq n$, we write $[n]$ for $\{1, \dots, n\}$ and $[k, n]$ for $\{k, \dots, n\}$. If S is a set, we write $x \leftarrow S$ for sampling a uniformly random element from S . If $\mathcal{D}$ is a distribution over S , denoted as $\mathcal{D} \sim S$, we write $x \leftarrow \mathcal{D}$ for sampling a random element from S according to the distribution $\mathcal{D}$. We use bold capital and lower-case letters, e.g. $\mathbf{A}$ and $\mathbf{b}$, to denote matrices and vectors, respectively. For a vector $\mathbf{a} = (a_1, \dots, a_n)^T \in \mathbb{Z}^n$, we write $\|\mathbf{a}\|_p := (\sum_{i \in [n]} |a_i|^p)^{1/p}$ for its ℓ_p -norm. By default $p = 2$, in which case we simply drop the subscript p . For a matrix $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_k) \in \mathbb{Z}^{n \times k}$, we write $\|\mathbf{A}\| := \max_{i \in [k]} \|\mathbf{a}_i\|$ for its maximum ℓ_2 -column-norm. For $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we say that $\mathbf{A}$ is primitive if linear combinations over $\mathbb{Z}_q$ of the columns of $\mathbf{A}$ generate $\mathbb{Z}_q^n$. For partial degree- d binomial expansion (with one argument fixed to m), we use the following notation,

$$[[h]]_a^b = \sum_{i=b}^a \binom{d}{i} h^i m^{d-i}, \quad (14)$$

which is the empty sum, hence 0, whenever $a < b$, as for $[[h]]_{d-1}^1$ at $d = 1$.

Matrix Inner Products. We generalise the notation $\langle \cdot, \cdot \rangle$ for inner products to allow matrix operands, i.e. for matrices $\mathbf{A}, \mathbf{B}$ with the same number of rows, we write $\langle \mathbf{A}, \mathbf{B} \rangle := \mathbf{A}^T \cdot \mathbf{B}$.

Matrix Tensors. Write $\otimes$ for the tensor (i.e. Kronecker) product of matrices. The tensor product of matrices satisfies the mixed product property: For all matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ and $\mathbf{D}$ of suitable dimensions, we have $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$. The (column-style) vectorisation of a matrix $\mathbf{A} \in \mathbb{Z}^{n \times m}$ is denoted by $\text{vec}(\mathbf{A}) \in \mathbb{Z}^{nm \times 1}$, the vector obtained by stacking the columns of $\mathbf{A}$ on top of one another. For matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}$ of suitable dimensions, it holds $\text{vec}(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A}) \cdot \text{vec}(\mathbf{B})$.

Degree- d Tensors. Since we will work extensively with degree- d tensor products, we introduce below a suite of shorthands. For any $d \in \mathbb{N}$ and matrix $\mathbf{A}$, we write $\mathbf{A}^{\otimes d} = \mathbf{A} \otimes \mathbf{A} \otimes \dots \otimes \mathbf{A}$ for the d -fold self-tensor of $\mathbf{A}$. Let $\{\mathbf{A}, \mathbf{B}\}^d$ denote the (ordered) set of all d -tuples of matrices where each entry of the tuple is either $\mathbf{A}$ or $\mathbf{B}$. We write $\{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = i}^d$ for the subset of tuples where $\mathbf{A}$ appears exactly i times. We adopt the convention that the tuples are sorted first by decreasing $\# \mathbf{A}$, i.e. $\{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = d}^d = \{\mathbf{A}^{\otimes d}\}$ comes first, $\{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = d-1}^d$ comes second, and so on, and then, within each group, lexicographically with respect to $\mathbf{B} < \mathbf{A}$. Extending notation e.g. $\{\mathbf{A}, \mathbf{B}\}^d$ and $\{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = i}^d$, we write e.g. $\{\mathbf{A}, \mathbf{B}\}^{\otimes d}$ for the vertical concatenation of all Kronecker products of d -tuples; and $\{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = i}^{\otimes d}$ for the vertical concatenation of all Kronecker products of d -tuples consisting of i $\mathbf{A}$'s and $d - i$ $\mathbf{B}$'s. For example,

$$\{\mathbf{A}, \mathbf{B}\}^{\otimes 3} = \left[\begin{array}{c} \{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} > 0}^{\otimes 3} \\ \mathbf{B}^{\otimes 3} \end{array} \right] = \left[\begin{array}{c} \mathbf{A}^{\otimes 3} \\ \{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} > 0}^{\otimes 3} \\ \mathbf{B}^{\otimes 3} \end{array} \right] = \left[\begin{array}{c} \mathbf{A}^{\otimes 3} \\ \{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = 2}^{\otimes 3} \\ \{\mathbf{A}, \mathbf{B}\}_{\# \mathbf{A} = 1}^{\otimes 3} \\ \mathbf{B}^{\otimes 3} \end{array} \right] = \left[\begin{array}{c} \mathbf{A} \otimes \mathbf{A} \otimes \mathbf{A} \\ \mathbf{B} \otimes \mathbf{A} \otimes \mathbf{A} \\ \mathbf{A} \otimes \mathbf{B} \otimes \mathbf{A} \\ \mathbf{A} \otimes \mathbf{A} \otimes \mathbf{B} \\ \mathbf{B} \otimes \mathbf{B} \otimes \mathbf{A} \\ \mathbf{B} \otimes \mathbf{A} \otimes \mathbf{B} \\ \mathbf{A} \otimes \mathbf{B} \otimes \mathbf{B} \\ \mathbf{B} \otimes \mathbf{B} \otimes \mathbf{B} \end{array} \right].$$

3.1 Functional Encryption

We recall the definition of functional encryption [BSW11, O'N10].

Definition 1 (Secret-key Functional Encryption). A secret-key functional encryption (FE) scheme for a class of functions $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathcal{Z}\}$ is a tuple of PPT algorithms $\text{FE} = (\text{Setup}, \text{KGen}, \text{Enc}, \text{Dec})$ defined as follows:

$\text{Setup}(1^\lambda)$: inputs a security parameter 1^λ and outputs a master secret key msk .

$\text{KGen}(\text{msk}, f)$: inputs the master secret key msk and a function $f \in \mathcal{F}$, and outputs a functional secret key sk_f .

$\text{Enc}(\text{msk}, x)$: inputs the msk and a message $x \in \mathcal{X}$, and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_f, \text{ct})$: inputs sk_f and a ciphertext ct , and returns an output $z \in \mathcal{Z}$.

$\text{INDExp}_{\text{FE}, \mathcal{A}}^b(1^\lambda)$	$\text{KGenO}(f)$
$((x_i)_{i \in [Q]}, x_0^*, x_1^*) \leftarrow \mathcal{A}(1^\lambda)$	if $f(x_0^*) \neq f(x_1^*)$ then
assert $((x_i)_{i \in [Q]}, x_0^*, x_1^*) \in \mathcal{X}^{Q+2}$	return $\perp$
$\text{msk} \leftarrow \text{Setup}(1^\lambda)$	return $\text{KGen}(\text{msk}, f)$
for $i \in [Q]$ do	
$\text{ct}_i \leftarrow \text{Enc}(\text{msk}, x_i)$	
$\text{ct}^* \leftarrow \text{Enc}(\text{msk}, x_b^*)$	
$b' \leftarrow \mathcal{A}^{\text{KGenO}}((\text{ct}_i)_{i \in [Q]}, \text{ct}^*)$	
return b'	

Fig. 1. The selective-message indistinguishability experiment for secret-key FE.

Correctness. The correctness property states that for all $\lambda \in \mathbb{N}$, all $\text{msk} \in \text{Setup}(1^\lambda)$, all functions $f \in \mathcal{F}$ and all messages $x \in \mathcal{X}$

$$\Pr \left[\text{Dec}(\text{sk}_f, \text{ct}) = f(x) \mid \begin{array}{l} \text{sk}_f \leftarrow \text{KGen}(\text{msk}, f) \\ \text{ct} \leftarrow \text{Enc}(\text{msk}, x) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Indistinguishability Security. For every stateful PPT adversary $\mathcal{A}$, define

$$\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{ind}}(\lambda) := \left| \Pr [\text{INDExp}_{\text{FE}, \mathcal{A}}^0(1^\lambda) = 1] - \Pr [\text{INDExp}_{\text{FE}, \mathcal{A}}^1(1^\lambda) = 1] \right|,$$

where the experiments $\text{INDExp}_{\text{FE}, \mathcal{A}}^b$ are defined in Fig. 1. We say that FE is selective-message IND-CPA secure if, for every stateful PPT adversary $\mathcal{A}$, the advantage $\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{ind}}(\lambda)$ is negligible.

Single-key security. If a security notion is only satisfied with the restriction that the adversary $\mathcal{A}$ only queries KGenO on a single function $f \in \mathcal{F}$, we decorate the name of the security notion with “single-key”.

It is possible to consider a multi-challenge security notion, where the adversary $\mathcal{A}$ declares multiple pair of challenge messages $(x_{0,i}^*, x_{1,i}^*)_i$ and is given the ciphertexts $(\text{ct}_i^*)_i$ of $x_{b,i}^*$ for each i . The single- and multi-challenge variants are polynomially equivalent via a standard hybrid argument.

We recall the definition of public-key FE in Appendix F.2.

3.2 Lattices and Gaussians

An m -dimensional lattice $\mathcal{L}$ is a discrete additive subgroup of $\mathbb{R}^m$. A matrix $\mathbf{B} \in \mathbb{R}^{m \times m}$ is said to be a basis of a full-rank m -dimensional lattice $\mathcal{L}$, if the columns of $\mathbf{B}$ are linearly-independent and $\mathbf{B}$ generates $\mathcal{L}$, i.e. it holds $\mathcal{L} = \mathcal{L}(\mathbf{B}) := \mathbf{B} \cdot \mathbb{Z}^m = \{\mathbf{B} \cdot \mathbf{z} : \mathbf{z} \in \mathbb{Z}^m\}$. For $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{v} \in \mathbb{Z}_q^n$, and $\mathbf{V} \in \mathbb{Z}^{n \times k}$, we define the shorthands

$$\begin{aligned} \Lambda_q^\perp(\mathbf{A}) &:= \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{u} \equiv \mathbf{0} \pmod{q}\}, \quad \Lambda_q^\mathbf{v}(\mathbf{A}) := \{\mathbf{u} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{u} \equiv \mathbf{v} \pmod{q}\}, \\ \Lambda_q^\mathbf{V}(\mathbf{A}) &:= \{\mathbf{U} \in \mathbb{Z}^{m \times k} : (\mathbf{I}_k \otimes \mathbf{A}) \cdot \text{vec}(\mathbf{U}) \equiv \text{vec}(\mathbf{V}) \pmod{q}\}. \end{aligned}$$

Gram-Schmidt orthogonalisation. For any (ordered) set $\mathbf{T} = \{\mathbf{t}_1, \dots, \mathbf{t}_m\} \subset \mathbb{R}^m$ of linearly independent vectors, we denote by $\text{GS}(\mathbf{T}) = \{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_m\}$ its Gram-Schmidt orthogonalisation, defined iteratively as follows: Set $\hat{\mathbf{t}}_1 = \mathbf{t}_1$, and for each $i \in [2 : m]$, let $\hat{\mathbf{t}}_i$ be the component of $\mathbf{t}_i$ orthogonal to the span of $\{\hat{\mathbf{t}}_1, \dots, \hat{\mathbf{t}}_{i-1}\}$. It holds $\|\hat{\mathbf{t}}_i\| \leq \|\mathbf{t}_i\|$ for all $i \in [m]$, therefore also $\|\text{GS}(\mathbf{T})\| \leq \|\mathbf{T}\|$.

Discrete Gaussians. For any possibly rectangular or non-full-rank matrix $\mathbf{B} \in \mathbb{R}^{m \times k}$, let $\Sigma := \mathbf{B}\mathbf{B}^\top$ and define $\rho_{\mathbf{B}, \mathbf{c}}(\mathbf{x}) := \exp(-\pi(\mathbf{x} - \mathbf{c})^\top \Sigma^\dagger (\mathbf{x} - \mathbf{c}))$ if $\mathbf{x} - \mathbf{c} \in \text{span}_{\mathbb{R}}(\mathbf{B})$, and $\rho_{\mathbf{B}, \mathbf{c}}(\mathbf{x}) := 0$ otherwise. This function depends on $\mathbf{B}$ only through Σ , so we also write $\rho_{\sqrt{\Sigma}, \mathbf{c}}$; when $\Sigma = \sigma^2 \mathbf{I}$ and $\mathbf{c} = \mathbf{0}$, we write ρ_σ . For any lattice coset $A \subseteq \mathbf{c} + \text{span}_{\mathbb{R}}(\mathbf{B})$, define $\mathcal{D}_{A, \mathbf{B}, \mathbf{c}}(\mathbf{x}) := \rho_{\mathbf{B}, \mathbf{c}}(\mathbf{x}) / \rho_{\mathbf{B}, \mathbf{c}}(A)$ for $\mathbf{x} \in A$, where $\rho_{\mathbf{B}, \mathbf{c}}(A) := \sum_{\mathbf{z} \in A} \rho_{\mathbf{B}, \mathbf{c}}(\mathbf{z})$; when $\mathbf{c} = \mathbf{0}$, we omit it from the notation.

Smoothing parameter. The smoothing parameter [MR04] $\eta_\epsilon(\mathcal{L})$ of a lattice $\mathcal{L}$ is the smallest value $\sigma > 0$ s.t. $\rho_{1/\sigma}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \epsilon$ where $\mathcal{L}^*$ is the dual lattice of $\mathcal{L}$. Throughout, we will consider $\epsilon = \text{negl}(\lambda)$.

Lemma 1 ([GPV08, Lem. 3.1]). *For any m -dimensional lattice $\mathcal{L}$ and $\epsilon > 0$, $\eta_\epsilon(\mathcal{L}) \leq \sqrt{\ln(2m \cdot (1 + 1/\epsilon)) / \pi} \cdot \tilde{bl}(\mathcal{L})$ where $\tilde{bl}(\mathcal{L}) := \min_{\mathbf{B}} \|\text{GS}(\mathbf{B})\|$ with the minimum taken over all bases $\mathbf{B}$ of $\mathcal{L}$. In particular for any function in $\omega(\sqrt{\log m})$ there exist a negligible function $\epsilon(m)$ such that $\eta_\epsilon(\mathcal{L}) \leq \omega(\sqrt{\log m}) \cdot \tilde{bl}(\mathcal{L})$.*

3.3 Gaussian Sampling and Lattice Trapdoors

We instantiate the Gaussian sampling algorithm of [GPV08] over q -ary lattices, which we refer to as **GPVSamp** throughout. We call an ordered set $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ of m linearly independent vectors of an m -dimensional lattice Λ a *full-rank set* of Λ . Note that a full-rank set need not generate Λ .

Lemma 2. *There is a PPT algorithm **GPVSamp** that, on input $(\mathbf{A}, \mathbf{S}, \mathbf{v}, s)$ – a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a full-rank set $\mathbf{S}$ of the m -dimensional q -ary lattice $\Lambda_q^\perp(\mathbf{A})$, a target vector $\mathbf{v} \in \mathbb{Z}_q^n$ in the image of $\mathbf{A}$, and a parameter $s \geq \|\text{GS}(\mathbf{S})\| \cdot \omega(\sqrt{\log m})$, outputs a sample from a distribution that is statistically close to $\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s}$.*

Corollary 1. *If Λ is a lattice and $\mathbf{S}$ a full-rank set of Λ , then $\tilde{bl}(\Lambda) \leq \|\text{GS}(\mathbf{S})\|$.*

The proofs of Lemma 2 and Corollary 1 are given in Appendix A.2. We also recall the syntax and properties of lattice trapdoors [GPV08, MP12].

Definition 2 (Lattice Trapdoors [GPV08]). *A lattice trapdoor scheme (over $\mathbb{Z}$) consists of PPT algorithms (**TrapGen**, **SampPre**):*

$(\mathbf{A}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^m, q)$: *Sample a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor td .*

$\mathbf{u} \leftarrow \text{SampPre}(\mathbf{A}, \text{td}, \mathbf{v}, s)$: *Given the trapdoor td of $\mathbf{A}$, a target image vector $\mathbf{v} \in \mathbb{Z}_q^n$ and a Gaussian parameter s , sample a preimage vector $\mathbf{u} \in \mathbb{Z}^m$.*

A tuple of parameters $\text{par} = (n, m, q, s)$ is said to be admissible if:

1. It holds that $\{\mathbf{A} : (\mathbf{A}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^n, 1^m, q)\} \approx_s \mathcal{U}(\mathbb{Z}_q^{n \times m})$.
2. For any $s' \geq s$ and for all but a $\text{negl}(\lambda)$ -fraction of $(\mathbf{A}, \text{td})$ in the support of $\text{TrapGen}(1^\lambda, 1^n, 1^m, q)$, the following hold:
 - For any $\mathbf{v} \in \mathbb{Z}_q^n$, it holds that $\text{SampPre}(\mathbf{A}, \text{td}, \mathbf{v}, s') \approx_s \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s'}$.
 - It holds that $\left\{ (\mathbf{u}, \mathbf{v}) \left| \begin{array}{l} \mathbf{v} \leftarrow \mathbb{Z}_q^n \\ \mathbf{u} \leftarrow \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s'} \end{array} \right. \right\} \approx_s \left\{ (\mathbf{u}, \mathbf{v}) \left| \begin{array}{l} \mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}^m, s'} \\ \mathbf{v} := \mathbf{A}\mathbf{u} \bmod q \end{array} \right. \right\}$.

When running **GPVSamp** in Lemma 2, resp. **SampPre** in Definition 2, on multiple images vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$, we simply write **GPVSamp**($\mathbf{A}, \mathbf{S}, \mathbf{V}, s$), resp. **SampPre**($\mathbf{A}, \text{td}, \mathbf{V}, s$), where $\mathbf{V} := (\mathbf{v}_1, \dots, \mathbf{v}_k)$.

4 Trapdoor for Degree- d Tensor Lattices

We describe how to generate random matrices $\mathbf{A} \in \mathbb{Z}_q^{h \times m}$ and $\tilde{\mathbf{A}} = (\tilde{\mathbf{A}}_i)_{i=d}^1 \in \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}$ together with a lattice trapdoor for

$$\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}}) := \left[\begin{array}{c|c|c} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d & \\ \hline \{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d} & \tilde{\mathbf{A}}_i & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots_{i:d-1 \rightarrow 1} & \vdots_{i:d-1 \rightarrow 1} & \ddots_{i:d-1 \rightarrow 1} \end{array} \right]. \quad (15)$$

Our method works by first generating a lattice trapdoor for the related matrix

$$\mathcal{E}^{\otimes d}(\mathbf{A}) := \left[\begin{array}{c|c} \mathbf{A}^{\otimes d} & \\ \hline \{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d} & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots_{i:d-1 \rightarrow 1} & \ddots_{i:d-1 \rightarrow 1} \end{array} \right], \quad (16)$$

i.e. taking $\tilde{\mathbf{A}}$ as the empty matrix in $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$, and then applying the standard trapdoor extension technique. The matrix $\mathcal{E}^{\otimes d}(\mathbf{A})$ in Eq. (16) can be interpreted as a $\sum_{i=1}^d \binom{d}{i} \times \sum_{i=1}^d \binom{d}{i}$ block matrix indexed by $(\mathbf{i}, \mathbf{j})$, where $\mathbf{i}, \mathbf{j}$ range over $\{0, 1\}^d \setminus \mathbf{1}_d$, and the $(\mathbf{i}, \mathbf{j})$ -block of $\mathcal{E}^{\otimes d}(\mathbf{A})$ is as follows:

- For the first column, i.e. $\mathbf{j} = \mathbf{0}_d$, the $(\mathbf{i}, \mathbf{0}_d)$ -th block is $\bigotimes_{\ell=1}^d ((1 - i_\ell) \cdot \mathbf{A} + i_\ell \cdot \mathbf{I}_m)$. That is, a tensor product whose factors are determined by the entries of $\mathbf{i} = (i_\ell)_{\ell \in [d]}$ – for each slot ℓ the factor is either a $\mathbf{A}$ or a $\mathbf{I}_m$ depending on whether i_ℓ is 0 or 1, respectively. In particular, this means the top-left block, i.e. $(\mathbf{i}, \mathbf{j}) = (\mathbf{0}_d, \mathbf{0}_d)$, is $\mathbf{A}^{\otimes d}$.
- For the remaining columns, i.e. $\mathbf{j} \neq \mathbf{0}_d$, the $(\mathbf{i}, \mathbf{j} = \mathbf{i})$ -th block is $p^{d-|\mathbf{i}|} \cdot \mathbf{I}_{h^{d-|\mathbf{i}|} \cdot m^{|\mathbf{i}|}}$. The remaining blocks are $\mathbf{0}$.

Fact 1. *The matrix $\mathcal{E}^{\otimes d}(\mathbf{A})$ defined in Eq. (16) is primitive over $\mathbb{Z}_q$ if $\mathbf{A}$ is.*

Trapdoor Construction for Tensors. In Fig. 2 we construct a PPT algorithm $\text{TrapGen}_{h,m,p,q}^{\otimes d}$ which returns $(\mathbf{A}, \mathbf{T})$, where the marginal distribution of $\mathbf{A}$ is statistically close to uniform and $\mathbf{T}$ is with overwhelming probability in λ an Ajtai trapdoor for the matrix $\mathcal{E}^{\otimes d}(\mathbf{A})$. Examples with $d \in \{2, 3\}$ are illustrated in Eq. (13) and Appendix B.1, respectively.

In words, the algorithm in Fig. 2 does the following: First, sample $\mathbf{A}$ with a trapdoor td using a lattice trapdoor scheme (Definition 2), and use td to sample preimages $\mathbf{U}_0, \mathbf{U}_1$ satisfying $\text{rank}_{\mathbb{Q}}(\mathbf{U}_0) = m$, $\mathbf{A} \cdot \mathbf{U}_0 = \mathbf{0} \bmod q$ and $\mathbf{A} \cdot \mathbf{U}_1 = p \cdot \mathbf{I}_h \bmod q$. Then, construct $\mathbf{T}$ as a $\sum_{i=1}^d \binom{d}{i} \times \sum_{i=1}^d \binom{d}{i}$ block matrix indexed by $(\mathbf{j}, \mathbf{k})$, where each of $\mathbf{j}, \mathbf{k}$ range over $\{0, 1\}^d \setminus \mathbf{1}_d$, and where $\mathbf{j} \geq \mathbf{k}$ denotes the componentwise order, i.e. $j_\ell \geq k_\ell$ for every $\ell \in [d]$, as follows:

- The top-left block, i.e. $(\mathbf{j}, \mathbf{k}) = (\mathbf{0}_d, \mathbf{0}_d)$, is $\mathbf{U}_0^{\otimes d}$. For the rest in the top row, i.e. $\mathbf{j} = \mathbf{0}_d$, the $(\mathbf{0}_d, \mathbf{k})$ -th block is $\bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0)$. That is, we construct a tensor product whose factors are determined by the entries of $\mathbf{k} = (k_\ell)_{\ell \in [d]}$ – for each slot ℓ the factor is either a $\mathbf{U}_1$ or a $\mathbf{U}_0$ depending on whether k_ℓ is 0 or 1, respectively.
- For the remaining columns, i.e. $\mathbf{k} \neq \mathbf{0}_d$, and every $\mathbf{j} \geq \mathbf{k}$ componentwise, the $(\mathbf{j}, \mathbf{k})$ -th block is $-\bigotimes_{\ell=1}^d g(j_\ell, k_\ell)$, where

$$g(0, 0) = \mathbf{I}_h, \quad g(1, 0) = \mathbf{U}_1, \quad g(1, 1) = \mathbf{U}_0. \quad (17)$$

That is, we construct the negation of a tensor product whose factors are determined slot by slot by the pair (j_ℓ, k_ℓ) , which ranges over $\{(0, 0), (1, 0), (1, 1)\}$ since $\mathbf{j} \geq \mathbf{k}$. The case $\mathbf{j} = \mathbf{k}$ recovers a tensor product whose factors are $\mathbf{I}_h$ or $\mathbf{U}_0$ depending on whether k_ℓ is 0 or 1. The remaining blocks are $\mathbf{0}$.

TrapGen $_{h,m,p,q}^{\otimes d}(1^\lambda)$	GenFullRankSet($\mathbf{A}, \text{td}, s$)
$(\mathbf{A}, \text{td}) \leftarrow \text{TrapGen}(1^\lambda, 1^h, 1^m, q)$ $\mathbf{U}_0 \leftarrow \text{GenFullRankSet}(\mathbf{A}, \text{td}, s)$ $\mathbf{U}_1 \leftarrow \text{SampPre}(\mathbf{A}, \text{td}, p \cdot \mathbf{I}_h, s)$ $\mathbf{T} := \mathbf{0}_{(m^d + [[h]]_{d-1}^1) \times (m^d + [[h]]_{d-1}^1)}$ <i>// overwrite block $(\mathbf{0}_d, \mathbf{0}_d)$ of $\mathbf{T}$:</i> $\mathbf{T}_{\mathbf{0}_d, \mathbf{0}_d} := \mathbf{U}_0^{\otimes d}$ for $\mathbf{k} = (k_\ell)_{\ell \in [d]} \in (\{0, 1\}^d \setminus \{\mathbf{0}_d, \mathbf{1}_d\})$ do <i>// overwrite block $(\mathbf{0}_d, \mathbf{k})$ and blocks $(\mathbf{j}, \mathbf{k})$ for $\mathbf{j} \geq \mathbf{k}$ of $\mathbf{T}$:</i> $\mathbf{T}_{\mathbf{0}_d, \mathbf{k}} := \bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0)$ for $\mathbf{j} = (j_\ell)_{\ell \in [d]} \in (\{0, 1\}^d \setminus \{\mathbf{1}_d\})$ with $\mathbf{j} \geq \mathbf{k}$ do $\mathbf{T}_{\mathbf{j}, \mathbf{k}} := - \bigotimes_{\ell=1}^d g(j_\ell, k_\ell)$ return $(\mathbf{A}, \mathbf{T})$	for counter $\in [\lambda]$: $\mathbf{U}_0 \leftarrow \text{SampPre}(\mathbf{A}, \text{td}, \mathbf{0}_{h \times m}, s)$ if $\text{rank}_{\mathbb{Q}}(\mathbf{U}_0) = m$ then return $\mathbf{U}_0$ return $\perp$ $\text{ExtTrapGen}_{h,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}, s^*)$ $(\mathbf{A}, \mathbf{T}) \leftarrow \text{TrapGen}_{h,m,p,q}^{\otimes d}(1^\lambda)$ $\tilde{\mathbf{A}} = (\tilde{\mathbf{A}}_i)_{i=d}^1 \leftarrow \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}; \tilde{\mathbf{K}} \leftarrow \mathcal{D}_{\mathbb{Z}, s^*}^{\tilde{m} \times k^d}$ $\begin{bmatrix} \mathbf{K} \\ \mathbf{F} \end{bmatrix} \leftarrow \text{GPVSamp}(\mathcal{E}^{\otimes d}(\mathbf{A}), \mathbf{T}, \mathbf{V} - \tilde{\mathbf{A}} \cdot \tilde{\mathbf{K}}, s^*)$ return $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$

Fig. 2. Algorithms TrapGen $_{h,m,p,q}^{\otimes d}$ (Theorem 1) and ExtTrapGen $_{h,m,\tilde{m},p,q}^{\otimes d}$ (Theorem 2). (TrapGen, SampPre) is a lattice trapdoor scheme (Definition 2). The PPT algorithm GPVSamp is given by Lemma 2. The shorthands $[[h]]_d^1, [[h]]_{d-1}^1$ are defined in Eq. (14), and the block factor g in Eq. (17).

Theorem 1. Let $d \in O(1)$ and $h, m, s, p, \log q \in \text{poly}(\lambda)$. Let $(\text{TrapGen}, \text{SampPre})$ be a (h, m, q, s) -admissible lattice trapdoor scheme. Suppose $m \geq \max(2h \log q, \lambda)$, $s \geq 3 \cdot \omega(\sqrt{\log m})$, and q is prime with $q \nmid p$. The PPT algorithm TrapGen $_{h,m,p,q}^{\otimes d}$ in Fig. 2 returns $(\mathbf{A}, \mathbf{T})$, where the marginal distribution of $\mathbf{A}$ satisfies

$$\left\{ \mathbf{A} \mid (\mathbf{A}, \mathbf{T}) \leftarrow \text{TrapGen}_{h,m,p,q}^{\otimes d}(1^\lambda) \right\} \approx_s \mathcal{U}(\mathbb{Z}_q^{h \times m}).$$

Moreover, except with probability $\text{negl}(\lambda)$, the pair $(\mathbf{A}, \mathbf{T})$ satisfies the following:

1. The columns of $\mathbf{T}$ are a full-rank set of the lattice $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}))$, where $\mathcal{E}^{\otimes d}(\mathbf{A})$ is defined in Eq. (16).
2. It holds that $\|\text{GS}(\mathbf{T})\| \leq \|\mathbf{T}\| \leq 2(s\sqrt{m})^d$.

Proof. The statistical closeness of the marginal distribution of $\mathbf{A}$ to uniform follows directly from the property of lattice trapdoor (Definition 2, Item 1). In the rest, we express $\mathcal{E}^{\otimes d}(\mathbf{A})$ and $\mathbf{T}$ as $\sum_{i=1}^d \binom{d}{i} \times \sum_{i=1}^d \binom{d}{i}$ block matrices, indexed by $(\mathbf{i}, \mathbf{j})$ and $(\mathbf{j}, \mathbf{k})$ respectively, where each of $\mathbf{i}, \mathbf{j}, \mathbf{k}$ range over $\{0, 1\}^d \setminus \mathbf{1}_d$. To recall again, the matrices $\mathcal{E}^{\otimes d}(\mathbf{A})$ and $\mathbf{T}$ take the following structure:

- For $\mathcal{E}^{\otimes d}(\mathbf{A})$: Each $(\mathbf{i}, \mathbf{0}_d)$ -th block in the first column is $\bigotimes_{\ell=1}^d ((1 - i_\ell) \cdot \mathbf{A} + i_\ell \cdot \mathbf{I}_m)$. For $\mathbf{j} \neq \mathbf{0}_d$, the $(\mathbf{i}, \mathbf{j} = \mathbf{i})$ -th block is $p^{d-|\mathbf{i}|} \cdot \mathbf{I}_{h^{d-|\mathbf{i}|} \cdot m^{|\mathbf{i}|}}$, and the rest are $\mathbf{0}$. In particular, in the first row, the top-left $(\mathbf{0}_d, \mathbf{0}_d)$ -th block is $\mathbf{A}^{\otimes d}$, and the rest are $\mathbf{0}$.
- For $\mathbf{T}$: The $(\mathbf{0}_d, \mathbf{0}_d)$ -th block is $\mathbf{U}_0^{\otimes d}$, and each remaining $(\mathbf{0}_d, \mathbf{k})$ -th block in the first row is $\bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0)$. For $\mathbf{k} \neq \mathbf{0}_d$ and every $\mathbf{j} \geq \mathbf{k}$, the $(\mathbf{j}, \mathbf{k})$ -th block is $-\bigotimes_{\ell=1}^d g(j_\ell, k_\ell)$ with g as in Eq. (17), and the rest are $\mathbf{0}$. In particular, in the first column, the top-left $(\mathbf{0}_d, \mathbf{0}_d)$ -th block is $\mathbf{U}_0^{\otimes d}$, and the rest are $\mathbf{0}$.

We show that, for all but a $\text{negl}(\lambda)$ -fraction of $(\mathbf{A}, \mathbf{T}) \in \text{TrapGen}_{h,m,p,q}^{\otimes d}(1^\lambda)$, it holds that (1) $\mathcal{E}^{\otimes d}(\mathbf{A}) \cdot \mathbf{T} = \mathbf{0} \bmod q$ and (2) $\mathbf{T}$ is invertible over $\mathbb{Q}$, therefore the columns of $\mathbf{T}$ are a full-rank set of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}))$.

For (1): Observe $\mathcal{E}^{\otimes d}(\mathbf{A}) \cdot \mathbf{T} \bmod q$ satisfies the following:

- For the first column, i.e. $\mathbf{k} = \mathbf{0}_d$, the only non-zero block of $\mathbf{T}$ in that column is $\mathbf{T}_{\mathbf{0}_d, \mathbf{0}_d} = \mathbf{U}_0^{\otimes d}$, so the $(\mathbf{i}, \mathbf{0}_d)$ -block is given by

$$\bigotimes_{\ell=1}^d ((1 - i_\ell) \cdot \mathbf{A} + i_\ell \cdot \mathbf{I}_m) \cdot \mathbf{U}_0^{\otimes d}.$$

Since $\mathbf{i} \neq \mathbf{1}_d$, there is at least one $\mathbf{A}$ in the first factor, and therefore at least one of the Kronecker factors of the above product is $\mathbf{A} \cdot \mathbf{U}_0 = \mathbf{0} \bmod q$.

- For the top row, i.e. $\mathbf{i} = \mathbf{0}_d$, the $(\mathbf{i}, \mathbf{k})$ -block for any $\mathbf{k} \neq \mathbf{0}_d$ is given by

$$\mathbf{A}^{\otimes d} \cdot \bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0).$$

Since $\mathbf{k} \neq \mathbf{0}_d$, there is at least one $\mathbf{U}_0$ in the second factor, and therefore at least one of the Kronecker factors of the above product is $\mathbf{A} \cdot \mathbf{U}_0 = \mathbf{0} \bmod q$. Therefore the product evaluates to zero modulo q .

- For the remaining rows, i.e. $\mathbf{i} \neq \mathbf{0}_d$, and any $\mathbf{k} \neq \mathbf{0}_d$, the only non-zero blocks of $\mathcal{E}^{\otimes d}(\mathbf{A})$ in row $\mathbf{i}$ sit in columns $\mathbf{0}_d$ and $\mathbf{i}$, so the $(\mathbf{i}, \mathbf{k})$ -block is given by

$$\begin{aligned} & \bigotimes_{\ell=1}^d ((1 - i_\ell) \cdot \mathbf{A} + i_\ell \cdot \mathbf{I}_m) \cdot \bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0) \\ & + p^{d-|\mathbf{i}|} \cdot \mathbf{I}_{h^{d-|\mathbf{i}|} \cdot m^{|\mathbf{i}|}} \cdot \mathbf{T}_{\mathbf{i}, \mathbf{k}}. \end{aligned}$$

The ℓ -th Kronecker factor of the first product is $\mathbf{A} \cdot \mathbf{U}_1 = p \cdot \mathbf{I}_h$ at $(i_\ell, k_\ell) = (0, 0)$, $\mathbf{A} \cdot \mathbf{U}_0 = \mathbf{0} \bmod q$ at $(0, 1)$, $\mathbf{U}_1$ at $(1, 0)$ and $\mathbf{U}_0$ at $(1, 1)$. We distinguish two cases. If $\mathbf{k} \not\leq \mathbf{i}$, then some slot ℓ has $(i_\ell, k_\ell) = (0, 1)$, so the first product vanishes modulo q ; and $\mathbf{T}_{\mathbf{i}, \mathbf{k}} = \mathbf{0}$ since $\mathbf{i} \not\geq \mathbf{k}$, so the second term vanishes too. If $\mathbf{k} \leq \mathbf{i}$, then no slot has $(i_\ell, k_\ell) = (0, 1)$, and collecting the factor p from each of the $d - |\mathbf{i}|$ slots with $i_\ell = 0$ gives

$$\begin{aligned} & \bigotimes_{\ell=1}^d ((1 - i_\ell) \cdot \mathbf{A} + i_\ell \cdot \mathbf{I}_m) \cdot \bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0) \\ & + p^{d-|\mathbf{i}|} \cdot \mathbf{I}_{h^{d-|\mathbf{i}|} \cdot m^{|\mathbf{i}|}} \cdot \mathbf{T}_{\mathbf{i}, \mathbf{k}} \\ & = p^{d-|\mathbf{i}|} \cdot \bigotimes_{\ell=1}^d g(i_\ell, k_\ell) - p^{d-|\mathbf{i}|} \cdot \bigotimes_{\ell=1}^d g(i_\ell, k_\ell) \\ & = \mathbf{0} \bmod q. \end{aligned}$$

We thus conclude that (all $(\mathbf{i}, \mathbf{k})$ -block of) $\mathcal{E}^{\otimes d}(\mathbf{A}) \cdot \mathbf{T} \bmod q$ equals $\mathbf{0}$.

For (2): By Lemma 12 in Appendix B.2, $\mathbf{U}_0$ obtained from `GenFullRankSet` is invertible over $\mathbb{Q}$ except with probability $\text{negl}(\lambda)$. Now observe that the first block column of $\mathbf{T}$ consists of $\mathbf{U}_0^{\otimes d}$ and zeros, so that $\det \mathbf{T} = \det(\mathbf{U}_0^{\otimes d}) \cdot \det \mathbf{L}$, where $\mathbf{L}$ is the submatrix indexed by $\mathbf{j}, \mathbf{k} \neq \mathbf{0}_d$. Since $\mathbf{T}_{\mathbf{j}, \mathbf{k}} = \mathbf{0}$ unless $\mathbf{j} \geq \mathbf{k}$, and $\mathbf{j} \geq \mathbf{k}$ with $\mathbf{j} \neq \mathbf{k}$ forces $|\mathbf{j}| > |\mathbf{k}|$, ordering those indices by increasing $|\mathbf{k}|$ makes $\mathbf{L}$ block-triangular, with diagonal blocks $-\bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{I}_h + k_\ell \cdot \mathbf{U}_0)$. All the diagonal blocks of $\mathbf{T}$ are therefore Kronecker products with factors given by $\mathbf{I}_h$ or $\mathbf{U}_0$, and since both of the latter are invertible over $\mathbb{Q}$, so is $\mathbf{T}$.

Finally, we bound the norm of $\mathbf{T}$. Recall that, by property of `SampPre`, the matrices $\mathbf{U}_0 \in \mathbb{Z}^{m \times m}$, $\mathbf{U}_1 \in \mathbb{Z}^{m \times h}$ have norm bounded by $\|\mathbf{U}_0\| \leq s\sqrt{m}$, $\|\mathbf{U}_1\| \leq s\sqrt{m}$, with overwhelming probability in $m \geq \lambda$. The rest is conditioned on this event. Below, we bound the norm of each $(\mathbf{j}, \mathbf{k})$ -block of $\mathbf{T}$:

- The $(\mathbf{0}_d, \mathbf{0}_d)$ -block clearly has norm $\|\mathbf{U}_0^{\otimes d}\| \leq (s\sqrt{m})^d$.
- For any $\mathbf{k} \in \{0, 1\}^d \setminus \{\mathbf{1}_d, \mathbf{0}_d\}$, the $(\mathbf{0}_d, \mathbf{k})$ -th block has norm

$$\left\| \bigotimes_{\ell=1}^d ((1 - k_\ell) \cdot \mathbf{U}_1 + k_\ell \cdot \mathbf{U}_0) \right\| \leq (s\sqrt{m})^d.$$

- For any $\mathbf{j}, \mathbf{k} \in \{0, 1\}^d \setminus \mathbf{1}_d$ with $\mathbf{j} \geq \mathbf{k} \neq \mathbf{0}_d$, the $(\mathbf{j}, \mathbf{k})$ -th block has norm

$$\left\| -\bigotimes_{\ell=1}^d g(j_\ell, k_\ell) \right\| \leq (s\sqrt{m})^{|\mathbf{j}|} \leq (s\sqrt{m})^{d-1},$$

since exactly the $|\mathbf{j}|$ slots with $j_\ell = 1$ contribute a factor $\mathbf{U}_1$ or $\mathbf{U}_0$ and the rest contribute $\mathbf{I}_h$, and since $\mathbf{j} \neq \mathbf{1}_d$ gives $|\mathbf{j}| \leq d - 1$.

- The remaining blocks have norm $\|\mathbf{0}\| = 0$.

Now we notice again that there are two possibilities for a column of $\mathbf{T}$: Either it lies in the first block column, i.e. it consists of a vector from the $(\mathbf{0}_d, \mathbf{0}_d)$ -block $\mathbf{U}_0^{\otimes d}$ and the rest are zeros, whose norm is upper-bounded by $(s\sqrt{m})^d$; or it lies in a block column $\mathbf{k} \neq \mathbf{0}_d$, where it meets the $(\mathbf{0}_d, \mathbf{k})$ -block together

with the $(\mathbf{j}, \mathbf{k})$ -block for each of the $2^{d-|\mathbf{k}|} - 1 \leq 2^d$ many $\mathbf{j} \in \{0, 1\}^d \setminus \{\mathbf{1}_d\}$ with $\mathbf{j} \geq \mathbf{k}$, and the rest are zeros. Summing the squared block bounds in the latter case gives

$$\|\mathbf{T}\|^2 \leq (s\sqrt{m})^{2d} + 2^d \cdot (s\sqrt{m})^{2(d-1)} = (s\sqrt{m})^{2d} \cdot \left(1 + \frac{2^d}{(s\sqrt{m})^2}\right) \leq 4 \cdot (s\sqrt{m})^{2d},$$

where the last step uses $2^d \leq 3s^2m$, which holds for λ large enough since $d \in O(1)$ while $s \geq 3 \cdot \omega(\sqrt{\log m}) \geq 3$ and $m \geq \lambda$. We conclude $\|\mathbf{T}\| \leq 2(s\sqrt{m})^d$. $\square$

Recall that $\mathcal{E}^{\otimes d}(\mathbf{A})$ (Eq. (16)) and $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$ (Eq. (15)) are almost identical, except that the latter has an additional block column $\tilde{\mathbf{A}}$. Because of Theorem 1, we can sample Gaussian preimages w.r.t. $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$ using the trapdoor $\mathbf{T}$ of $\mathcal{E}^{\otimes d}(\mathbf{A})$ by trapdoor delegation [CHKP10, CHKP12]. To ease notation in subsequent sections, we capture both the matrix- and preimage-sampling in a single algorithm $\text{ExtTrapGen}^{\otimes d}$, given in Fig. 2.

Theorem 2. *Let $d \in O(1)$ and $h, m, \tilde{m}, k, s, p, \log \tilde{s}, \log q \in \text{poly}(\lambda)$. Let $(\text{TrapGen}, \text{SampPre})$ be a (h, m, q, s) -admissible lattice trapdoor scheme, and $\text{ExtTrapGen}_{h,m,\tilde{m},p,q}^{\otimes d}$ be the PPT algorithm given in Fig. 2. Let $\mathbf{V} \in \mathbb{Z}_q^{[[h]]_d^1 \times k^d}$ be arbitrary.*

Suppose $m \geq \max(2h \log q, \lambda)$, $s \geq 3 \cdot \omega(\sqrt{\log m})$, $\tilde{s} \geq 2(s\sqrt{m})^d \cdot \omega(\sqrt{\log(m)})$, and q is prime with $q \nmid p$. Then, for any $s^* \geq \tilde{s}$, the distribution of $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}_{h,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}, s^*)$ is statistically close in λ to

$$\left\{ (\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \left| \begin{array}{l} \mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{h \times m}; \quad \tilde{\mathbf{A}} = (\tilde{\mathbf{A}}_i)_{i=d}^1 \leftarrow_{\$} \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}} \\ \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{bmatrix} \leftarrow_{\$} \mathcal{D}_{A_q^{\mathbf{V}}(\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})), s^*} \end{array} \right. \right\}. \quad (18)$$

The proof is given in Appendix B.3.

5 New Assumption

In Assumption 1 we define a family of assumptions which can be seen as a secret-key variant of a degree- d generalisation of an “affine” variant of the standard k -LWE assumption: Instead of being given $\mathbf{A}$ and hints from $\Lambda_q^\perp(\mathbf{A})$, the distinguisher is now given hints from (cosets of) $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}}))$ but not the matrices $\mathbf{A}, \tilde{\mathbf{A}}$ themselves. Following the convention that the matrix of an LWE-like assumption is public unless stated otherwise, we write $k\text{-ALWE}^{\otimes d}$ for the degree- d variant in which $\mathbf{A}, \tilde{\mathbf{A}}$ are given to the adversary, and $\text{sk-}k\text{-ALWE}^{\otimes d}$ for the variant in which they are withheld. The degree-1 case of the former is abbreviated $k\text{-ALWE}$. Assumption 1 states the $\text{sk-}k\text{-ALWE}^{\otimes d}$ variant, on which our construction relies.

Assumption 1. *Let $\text{par} = (k, n, m, \tilde{m}, p, q, \chi_s, \chi_e, \tilde{\chi}_e, \sigma, \beta_x)$ be parametrised by λ , where $k, n, m, \tilde{m}, p, q \in \mathbb{N}$, and $\chi_s, \chi_e, \tilde{\chi}_e, \sigma, \beta_x \in \mathbb{R}^+$. The $\text{sk-}k\text{-ALWE}_{\text{par}}^{\otimes d}$ problem is said to be T -hard, where $T = T(k, \lambda)$, if for any adversary $\mathcal{A}$ running in time at most T , it holds that*

$$\left| \Pr[\text{Exp}_{d,\text{par},\mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{Exp}_{d,\text{par},\mathcal{A}}^1(1^\lambda) = 1] \right| \leq \text{negl}(\lambda)$$

where the experiments $\text{Exp}_{d,\text{par},\mathcal{A}}^b$ are defined in Fig. 3. We say that the $\text{sk-}k\text{-ALWE}_{\text{par}}^{\otimes d}$ assumption holds if the problem is T -hard for every polynomial T .

Remark 1 (Falsifiability). Both experiments $\text{Exp}_{d,\text{par},\mathcal{A}}^b$ for $b \in \{0, 1\}$ can be efficiently simulated, i.e. the assumption is *falsifiable* [GW11]. Let $\mathbf{x}^*$ be given by the adversary. Define

$$\mathbf{V}_0 := \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) \bmod q,$$




Fig. 3. Experiment for $\text{sk-}k\text{-ALWE}^{\otimes d}$ (Assumption 1) with subroutines $\mathcal{E}^{\otimes d}$, $\text{LWESamp}^{\otimes d}$ and $\text{Img}^{\otimes d}$ defined in Fig. 4.

$$\mathbf{V}_1 := \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{0}^T \end{bmatrix} \right) + \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}^*)^T \end{bmatrix} \right) \bmod q.$$

To simulate $\text{Exp}_{d,\text{par},\mathcal{A}}^0$, sample $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}_{k+n,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}_0, \sigma)$. Similarly, to simulate $\text{Exp}_{d,\text{par},\mathcal{A}}^1$, run $\text{ExtTrapGen}_{k+n+1,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}_1, \sigma)$ to sample $(\mathbf{B}, \tilde{\mathbf{B}}, \mathbf{K}, \tilde{\mathbf{K}})$. In either case, compute $(\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}}$ honestly. By Theorem 2, the above simulations are faithful.

Remark 2 (Public-key variant). In the experiment defined in Fig. 3, the LWE matrices $\mathbf{A}, \tilde{\mathbf{A}}$ are not given to the adversary, different from treatments in usual LWE-like assumptions. We could consider the public-key variant $k\text{-ALWE}^{\otimes d}$ where $\mathbf{A}, \tilde{\mathbf{A}}$ are given to the adversary, in which case the adversary could compute

$\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}})$	$\parallel$	$\mathbf{M} \in \mathbb{Z}_q^{h \times m}, \tilde{\mathbf{M}} \in \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}$
parse $\tilde{\mathbf{M}} = (\tilde{\mathbf{M}}_i)_{i=d}^1 \in \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}$		
return	$\left[\begin{array}{c c c} \mathbf{M}^{\otimes d} & \tilde{\mathbf{M}}_d & \\ \hline \{\mathbf{M}, \mathbf{I}_m\}_{\# \mathbf{M}=i}^{\otimes d} & \tilde{\mathbf{M}}_i & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \hline \vdots_{i:d-1 \rightarrow 1} & \vdots_{i:d-1 \rightarrow 1} & \ddots_{i:d-1 \rightarrow 1} \end{array} \right]$ $\{\mathbf{M}, \mathbf{I}_m\}_{\# \mathbf{M} > 0}^{\otimes d}$ $\tilde{\mathbf{M}}$	
$\text{LWESamp}^{\otimes d}(\mathbf{r}, \mathbf{M}, \tilde{\mathbf{M}})$	$\parallel$	$\mathbf{r} \in \mathbb{Z}^h, \mathbf{M} \in \mathbb{Z}_q^{h \times m}, \tilde{\mathbf{M}} \in \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}$
$\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi_e}, \tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}^{\tilde{m}}, \tilde{\chi}_e}$		
$\mathbf{c}^T := \langle \mathbf{r}, \mathbf{M} \rangle + p \cdot \mathbf{e}^T \bmod q \in \mathbb{Z}_q^m$		
$\tilde{\mathbf{c}}^T := \left\langle \{\mathbf{r}, p \cdot \mathbf{e}\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \tilde{\mathbf{M}} \right\rangle + p^d \cdot \tilde{\mathbf{e}}^T \bmod q \in \mathbb{Z}_q^{\tilde{m}}$		
return $(\mathbf{c}, \tilde{\mathbf{c}})$		
$\text{Img}^{\otimes d}(\mathbf{M})$	$\parallel$	$\mathbf{M} \in \mathbb{Z}_q^{h \times k}$
return	$\{\mathbf{M}, \mathbf{0}_{m \times k}\}_{d > \# \mathbf{0}_{m \times k}}^{\otimes d} \bmod q \in \mathbb{Z}_q^{[[h]]_d^1 \times k^d}$	

Fig. 4. Subroutines for the experiments defined in Figs. 3 and 7.

$(\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}$ itself instead of by the experiment. We call the $d = 1$ case of the public-key variant the k -ALWE assumption, and we show in Appendix F.1 that this case is implied by a standard lattice assumption. As we will discuss in Section 7.2, for $d > 1$, the public-key variant is vulnerable to attacks based on continuous optimisation.

6 Constant-Degree Functional Encryption Construction

Let $d, k, n, m, \tilde{m}, p, q, r, \sigma, \chi_s, \chi_e, \tilde{\chi}_e, \beta_x, \beta_y$ be parametrised by λ and $r = p^d \bmod q$. In Fig. 5 we construct a lattice-based functional encryption scheme for a function class $\mathcal{F}$ over a message space $\mathcal{X} = \{\mathbf{x} \in \mathbb{Z}^k : \|\mathbf{x}\|_\infty \leq \beta_x\}$. We require (for correctness) the following properties to hold for any $f \in \mathcal{F}$ and $\mathbf{x} \in \mathcal{X}$:

- $f(\mathbf{x}) \in (-r/2, r/2] \cap \mathbb{Z}$, i.e. function values are bounded, and
- there exists $\mathbf{y} \in \mathbb{Z}^{k^d}$ with ℓ_1 -norm $\|\mathbf{y}\|_1 \leq \beta_y$, such that $f(\mathbf{x}) = \langle \mathbf{x}^{\otimes d}, \mathbf{y} \rangle$, i.e. f can be computed by a k -variate degree- d homogeneous integer polynomial with bounded coefficients.

Parameters. We assume the existence of a (h, m, q, s) -admissible lattice trapdoor scheme for both $h = k + n$ and $h = k + n + 1$.⁶ We require parameters as follows to satisfy the constraints from the cryptanalysis discussion of $\text{sk-}k\text{-ALWE}^{\otimes d}$ (Section 7), trapdoor sampling (Theorem 2 and Corollary 3), correctness (Theorem 3), the security proof (Theorem 4) and the efficiency analysis. The parameters are indexed by a target running time $T = T(k, \lambda)$, against which the scheme is claimed secure. Taking $T = 2^{\text{polylog}(k+\lambda)}$ gives quasi-polynomial security and $T = 2^{(k+\lambda)^{\Omega(1)}}$ the subexponential security that the path to iO requires.

The target time T trades security against compactness. Raising T widens the noise, and hence the modulus and the dimensions, so the compactness row caps how large T may be taken for a given message dimension k . **Efficiency.** We assume $\log T = \Omega(d \log(k + \lambda + \beta_x) + \log(r\beta_y))$, under which the correctness row gives $\log q = \Theta(\log T)$. The two LWE rows force the dimension n to grow with the target time: under the standard

⁶ We refer to e.g. [GPV08, MP12] for how this can be instantiated.

Setup (1^λ) $\mathbf{V} := \left\{ \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}, \mathbf{0}_{m \times k} \right\}_{d > \# \mathbf{0}_{m \times k}}^{\otimes d}$ $\parallel \mathbf{V} = \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) \in \mathbb{Z}_q^{[[k+n]]_d^1 \times k^d}$ $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}^{\otimes d}(1^\lambda, \mathbf{V}, \sigma)$ $\parallel (\mathbf{A}, \tilde{\mathbf{A}}) \in \mathbb{Z}_q^{(k+n) \times m} \times \mathbb{Z}_q^{[[k+n]]_d^1 \times \tilde{m}}$ $\parallel (\mathbf{K}, \tilde{\mathbf{K}}) \in \mathbb{Z}^{m^d \times k^d} \times \mathbb{Z}^{\tilde{m} \times k^d}$ return $\text{msk} := (\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$ KGen ($\text{msk}, \mathbf{y} \in \mathbb{Z}^{k^d}$) parse $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{msk}$ $(\mathbf{k}_y, \tilde{\mathbf{k}}_y) := (\mathbf{K} \cdot \mathbf{y}, \tilde{\mathbf{K}} \cdot \mathbf{y}) \in \mathbb{Z}^{m^d} \times \mathbb{Z}^{\tilde{m}}$ return $\text{sk}_y := (\mathbf{k}_y, \tilde{\mathbf{k}}_y)$	Enc ($\text{msk}, \mathbf{x} \in \mathbb{Z}^k$) parse $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{msk}$ $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}^m, \tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}^{\tilde{m}}$ $\mathbf{c}^T := \left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A} \right\rangle + p \cdot \mathbf{e}^T \bmod q$ $\tilde{\mathbf{c}}^T := \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e} < d}}^{\otimes d}, \tilde{\mathbf{A}} \right\rangle + p^d \cdot \tilde{\mathbf{e}}^T \bmod q$ $\parallel (\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right) \in \mathbb{Z}_q^m \times \mathbb{Z}_q^{\tilde{m}}$ return $\text{ct}_x := (\mathbf{c}, \tilde{\mathbf{c}})$ Dec (sk_y, ct_x) parse $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{ct}_x; (\mathbf{k}_y, \tilde{\mathbf{k}}_y) \leftarrow \text{sk}_y$ return $(\langle \mathbf{c}^{\otimes d}, \mathbf{k}_y \rangle + \langle \tilde{\mathbf{c}}, \tilde{\mathbf{k}}_y \rangle \bmod q) \bmod r$
--	--

Fig. 5. Secret-key FE for degree- d k -variate polynomials. $\text{ExtTrapGen}^{\otimes d}$ (with parameter $(k+n, m, \tilde{m}, p, q)$) is given in Fig. 2 and $\text{LWESamp}^{\otimes d}$ and $\text{Img}^{\otimes d}$ in Fig. 4.

Table 1. Parameters for the FE scheme in Fig. 5.

Feasibility	$d \in O(1)$
Theorem 2	q is prime, $q \nmid p$ $s \geq 3 \cdot \omega(\sqrt{\log m})$ $m \geq \Omega((k+n) \log q)$ $\sigma \geq 4(s\sqrt{m})^d \cdot \sqrt{\omega(\log(m))} \cdot ((\sqrt{k}\beta_x)^d + 1)$
Correctness	$\gamma = \sigma \cdot \left((2 \cdot \sqrt{k} + n + m)^d \cdot 2 \cdot (k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m)^{d/2} + \tilde{\chi}_e \cdot \tilde{m} \right) \cdot \beta_y$ $q \geq 2 \cdot r \cdot (\gamma + 1)$
Security against time T	$\text{LWE}_{n, m, q, \chi_s, \chi_e}$ and $\text{LWE}_{\tilde{m}, \tilde{m}, q, \chi_s, \tilde{\chi}_e}$ are hard (Strategies 1 and 2) $\chi_s, \chi_e, \tilde{\chi}_e \geq T^{1/6d} \cdot \sqrt{k} \cdot \beta_x$ (Strategy 3) $m \geq (k+n) \cdot \log T$ (Strategies 4 and 5) $\text{LWE}_{\tilde{m}, k^d, q, \mathcal{U}(\mathbb{Z}_q), p\sigma}$ is hard (Section 7.2)
Compactness	$(m + \tilde{m}) \log q \leq \text{poly}(\lambda) \cdot k^{d(1-\epsilon)}$ for a constant $\epsilon > 0$

lattice estimates they hold against time T once $n = \Omega((\log T)^2 / \log \log T)$. The $\text{LWE}_{\tilde{m}, k^d, q, \mathcal{U}(\mathbb{Z}_q), p\sigma}$ row imposes the same bound on $\tilde{m}$. Let $n = \tilde{m} := \lambda + (\log T)^2$ and abbreviate

$$\hat{n} := k + \lambda + (\log T)^2 = k + n.$$

Instantiating the remaining entries of Table 1 at the smallest admissible dimensions and modulus for the target time T , our FE scheme achieves⁷

$$\begin{aligned}
 |\text{msk}| &\leq O(\hat{n})^{2d} \cdot (\log T)^{d+1}, & \text{Time(KGen)} &\leq O(\hat{n})^{2d} \cdot (\log T)^d, \\
 |\text{ct}| &\leq O(\hat{n}) \cdot (\log T)^2, & \text{Time(Enc)} &\leq O(\hat{n})^{d+1} \cdot (\log T)^d, \\
 |\text{sk}_f| &\leq O(\hat{n})^d \cdot (\log T)^{d+1}, & \text{Time(Dec)} &\leq O(\hat{n})^d \cdot (\log T)^d
 \end{aligned}$$

where sizes are measured in bits and times are measured in number of $\mathbb{Z}_q$ operations. Counting in number of bit operations multiplies each running time by the cost $O(\log q)$ of a $\mathbb{Z}_q$ operation, where $\log q = \Theta(\log T)$.

⁷ We treat d as a variable here which is more general.

The two regimes of interest instantiate the ciphertext size as follows. Targeting quasi-polynomial security, i.e. $\log T = \text{polylog}(k + \lambda)$, we have $\hat{n} = O(k + \lambda)$ and obtain $|\text{ct}| \leq O(k + \lambda) \cdot \text{polylog}(k + \lambda)$, quasi-linear in the message length. Targeting the subexponential security that the path to iO requires (Appendix G), i.e. $\log T = (k + \lambda)^e$ for a constant $e > 0$, we have $\hat{n} = O((k + \lambda)^{\max(1, 2e)})$ and obtain $|\text{ct}| \leq O((k + \lambda)^{\max(1 + 2e, 4e)})$, polynomially superlinear in the message length. In either regime the compactness row of Table 1 caps $|\text{ct}|$ at $\text{poly}(\lambda) \cdot k^{d(1-\epsilon)}$, keeping it sublinear in the description size k^d of the supported degree- d functions, which in the subexponential regime holds as long as $\max(1 + 2e, 4e) \leq d(1 - \epsilon)$.

Theorem 3. *The FE scheme Fig. 5 with parameters as in Table 1 is correct.*

Theorem 4. *Let $\text{par} = (k, n, m, \tilde{m}, p, q, \chi_s, \chi_e, \tilde{\chi}_e, \sigma, \beta_x)$ be as in Table 1. The FE scheme for degree- d polynomials described in Fig. 5 is selectively IND-CPA secure under the $\text{sk-}k\text{-ALWE}_{\text{par}}^{\otimes d}$ assumption (Assumption 1).*

The proofs are given in Appendices C.1 and C.2, respectively.

7 Cryptanalysis

We do not know how to obtain a reduction from any standard assumption to $\text{sk-}k\text{-ALWE}^{\otimes d}$ for $d > 1$. We discuss in Appendix D.3 the hardships of such a reduction and why it would need to be quite unintuitive. In the following, we discuss different attack strategies against $\text{sk-}k\text{-ALWE}^{\otimes d}$ and the FE scheme in Fig. 5, and how they guide parameter selection and the design decision of conjecturing $\text{sk-}k\text{-ALWE}^{\otimes d}$ rather than its public-key variant $k\text{-ALWE}^{\otimes d}$ (Assumption 1). Note that for achieving iO, our FE scheme and its underlying assumption must achieve subexponential security.⁸

We first discuss in Section 7.1 strategies that fail against $k\text{-ALWE}^{\otimes d}$, the public-key variant in which $\mathbf{A}, \tilde{\mathbf{A}}$ are given to the adversary as is usual in LWE-like assumptions, and which therefore fail against $\text{sk-}k\text{-ALWE}^{\otimes d}$ as well. To rule out Strategy 3 against adversaries running in time T , we need to set $\chi_s, \chi_e, \tilde{\chi}_e \geq T^{1/6d} \cdot \sqrt{k} \cdot \beta_x$ as in Table 1, meaning that the modulus q also needs to grow with T . The correctness constraint scales q with χ_e^d , so the modulus-to-noise ratio of the LWE samples is $T^{(d-1)/6d}$ times a polynomial in $k, \lambda, \beta_x, \beta_y$, which for $d > 1$ is superpolynomial rather than polynomial. Strategy 4 considers algebraic attacks and gives an additional reason to let $n + m + \tilde{m}$ grow with $\log T$.

Strategy 5 in Section 7.2, however, is an attack on $k\text{-ALWE}^{\otimes d}$ that succeeds in our experiments, which is why we conjecture only $\text{sk-}k\text{-ALWE}^{\otimes d}$. Both strategies 4-5 are based on solving the system of polynomial equations over the integers determined by $\mathbf{A}, \tilde{\mathbf{A}}$.

Error Functions. To discuss Strategies 3, 4 and 5, we need to introduce some helper functions. With respect to $\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}$ (which fix $\mathbf{F}$) sampled as in Exp^0 , define the error function $\mathbb{Z}^{k+n+m+\tilde{m}} \rightarrow \mathbb{Z}^{k^d}$,

$$\text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}) = \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_e < d}^{\otimes d}, \mathbf{F} \right\rangle. \quad (19)$$

Similarly, with respect to $\mathbf{B}, \tilde{\mathbf{B}}, \mathbf{K}, \tilde{\mathbf{K}}$ (which fix $\mathbf{G}$) sampled as in Exp^1 , define the error functions $\mathbb{Z}^{k+n+m+\tilde{m}} \rightarrow \mathbb{Z}^{k^d}$,

$$\text{err}_1(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}) = \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \\ 0 \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_e < d}^{\otimes d}, \mathbf{G} \right\rangle, \quad (20)$$

$$\text{err}^*(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}) = \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{0}_k \\ \mathbf{s} \\ 1 \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_e < d}^{\otimes d}, \mathbf{G} \right\rangle. \quad (21)$$

⁸ In the sense that every adversary of noticeable advantage must have a time complexity of at least $2^{(k+\lambda)^e}$ for some constant $e > 0$.

Note that while err_0 and err_1 are symbolically equivalent, the distributions of $(\mathbf{K}, \tilde{\mathbf{K}})$ are different in Exp^0 and Exp^1 , since in Exp^1 they need to satisfy more constraints than in Exp^0 . Also note that in Exp^b the adversary could obtain samples of $\text{err}_b(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}})$ for $\mathbf{x}$ of choice, $\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}^m$ and $\tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}^{\tilde{m}}$ as follows: For some $j \in [Q]$, pick $\mathbf{x}_j = \mathbf{x}$ and receive from the experiment $(\mathbf{d}_j, \tilde{\mathbf{d}}_j)$ generated using random $(\mathbf{s}_j, \mathbf{e}_j, \tilde{\mathbf{e}}_j)$. Compute $\mathbf{f}_j = \text{err}_b(\mathbf{x}, \mathbf{s}_j, \mathbf{e}_j, \tilde{\mathbf{e}}_j)$ via

$$p^d \cdot \mathbf{f}_j^T := \langle \mathbf{d}_j^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{d}}_j, \tilde{\mathbf{K}} \rangle - (\mathbf{x}_j^T)^{\otimes d} \bmod q. \quad (22)$$

Note that $\mathbf{f}_j$ must be short, hence, it is indeed an integer vector. Given $(\mathbf{c}, \tilde{\mathbf{c}})$ generated in Exp^b with randomness $(\mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$, the adversary could similarly obtain $\mathbf{f}^* := \text{err}_0(\mathbf{x}^*, \mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$ in Exp^0 and $\mathbf{f}^* := \text{err}^*(\mathbf{x}^*, \mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$ in Exp^1 .

7.1 Failed Attacks against (sk-)k-ALWE $^{\otimes d}$

We first discuss Strategies 1-4, attacks that fail even against $k\text{-ALWE}^{\otimes d}$ and hence also the secret-key variant $\text{sk-}k\text{-ALWE}^{\otimes d}$.

Strategy 1: Solve LWE directly. If the matrix $\mathbf{A}$ is known, one attack would be to recover $(\mathbf{s}^*, \mathbf{e}^*)$ directly from $(\mathbf{A}, \mathbf{d})$, i.e. solving search LWE. To avoid this, the parameters should at least be such that $\text{LWE}_{n, m, q, \chi_s, \chi_e}$ is hard.

Strategy 2: Recover approximate trapdoor of $\mathbf{A}$. Let again $\mathbf{A}, \tilde{\mathbf{A}}$ be known. This strategy is to find a short $\tilde{\mathbf{K}}^*$ such that $\tilde{\mathbf{A}}_i \cdot \tilde{\mathbf{K}}^* \approx \mathbf{0} \bmod q$ for a relevant submatrix $\tilde{\mathbf{A}}_i \in \mathbb{Z}_q^{n \times \tilde{m}}$ of $\tilde{\mathbf{A}}$, so that an approximate trapdoor of $\mathbf{A}$ can be recovered, cf. discussion in Section 2.2. Setting parameters such that $\text{LWE}_{n, \tilde{m}, q, \chi_s, \tilde{\chi}_e}$ is hard avoids this attack.

Strategy 3: Analyse error terms statistically. We discuss here a statistical attack against $\text{sk-}k\text{-ALWE}^{\otimes d}$ that stems from the work of Tairi and Ünal [TÜ24] on attacking FE schemes. We first explain how their attack proceeds against our FE scheme in Fig. 5, then derive the corresponding attack against $\text{sk-}k\text{-ALWE}^{\otimes d}$. We conclude by setting parameters such that both attacks fail.

Tairi and Ünal [TÜ24] showed an impossibility result for lattice-based degree- d FE with $d > 1$. The class of lattice-based FE schemes they analysed almost captures our FE scheme. However, there are two parameter choices by which we aim to avoid their impossibility result: First, the time complexity of their attack grows polynomially in the ratio of ciphertext modulus over message modulus. Second, they prove a noticeable advantage of their attack only if the dimension of ciphertexts is linear in the dimension of the message space. We circumvent both restrictions against adversaries running in time T by letting both quantities grow with T as in Table 1: the noise widths are at least $T^{1/6d} \cdot \sqrt{k} \cdot \beta_x$, and the ciphertext dimension is at least $(k + n) \log T$, exceeding the message dimension k by a factor of at least $\log T$.

Of course, simply avoiding the formal requirements of the impossibility results of [TÜ24] does not guarantee security. Hence, we need to understand their attack procedure and consider an attack strategy that covers their approach. As the results of [TÜ24] cover a large range of lattice-based FE schemes, their attack strategy is quite complex, involving multiple transformations and reduction to an SKE scheme that is either insecure or incorrect. However, since we have a concrete FE candidate scheme, we can skip the transformation steps of [TÜ24] and directly jump to an attack on Fig. 5.

For our scheme, the attack chooses two different candidate messages $\mathbf{x}_0^* \neq \mathbf{x}_1^*$ together with a function $\mathbf{y}$ that cannot distinguish between them, i.e. $\langle \mathbf{x}_0^*, \mathbf{y} \rangle = \langle \mathbf{x}_1^*, \mathbf{y} \rangle$. The idea of [Üna20, TÜ24] is to observe the decryption noise

$$\langle \text{err}_0(\mathbf{x}_b^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle = \langle \mathbf{e}^{\otimes d}, \mathbf{K} \mathbf{y} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \mathbf{y} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#e < d}^{\otimes d}, \mathbf{F} \mathbf{y} \right\rangle.$$

If the statistical distance between $\langle \text{err}_0(\mathbf{x}_0^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$ and $\langle \text{err}_0(\mathbf{x}_1^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$ is non-negligible, and both distributions are bounded by some bound B , then [Üna20] shows that both distributions can be distinguished

by taking $O(B^3)$ samples. Most of [TÜ24] focuses on proving that a non-zero function $\mathbf{y}$ exists, which cannot distinguish between $\mathbf{x}_0^*$ and $\mathbf{x}_1^*$, but which induces a non-negligible statistical distance on $\langle \text{err}_0(\mathbf{x}_b^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$ and $\langle \text{err}_1(\mathbf{x}_b^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$. This is in fact the point where [TÜ24] needs ciphertexts to have linear size $O(k)$. Despite our ciphertexts having length $\omega(k)$, we will – for the sake of simplifying cryptanalysis – simply assume that such a function $\mathbf{y}$ does exist⁹, making the requirement of having linear ciphertexts for the result of [TÜ24] obsolete.

Now, when we fix $\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}$ and sample $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ randomly, we can consider $\langle \text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$ as a random multivariate integer polynomial of degree $d - 1$ in the unknowns $\mathbf{x} = (x_1, \dots, x_k)$. An adversary can ask for evaluations of random polynomial on points $\mathbf{x}_0^*, \mathbf{x}_1^*$, but it cannot directly learn its coefficients. A first idea is to ask for multiple samples of its evaluations on $\mathbf{x}_0^*$ and $\mathbf{x}_1^*$, taking the mean over all samples and using this to distinguish ciphertexts of $\mathbf{x}_0^*$ from ciphertexts of $\mathbf{x}_1^*$. This approach cannot work, since the coefficients of the random polynomial $\langle \text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle$ are usually centered around zero, hence $\mathbb{E}[\langle \text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle]$ is the zero polynomial. The observation of [Üna20] is that it suffices to take the means of *squares* of the random polynomial. Hence, we set $g(\mathbf{x}) := \langle \text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2$. The requirement of $\mathbf{y}$ being a useful function for the distinguisher translates to $\mathbb{E}_{\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}}[g(\mathbf{x})]$ having noticeable non-constant coefficients (also confer Lemma 18 where we compute $\mathbb{E}[g(\mathbf{x})]$ explicitly for $d = 2$). In that case, we can – by a Hoeffding bound – approximate $\mathbb{E}_{\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}}[g(\mathbf{x})]$ sufficiently if we have access to $O(N^3)$ samples where N is a bound on the magnitude of $g(\mathbf{x})$'s coefficients. At this point, [Üna20] needs that the noise distributions $\mathcal{D}_{\mathbb{Z}, \chi_s}, \mathcal{D}_{\mathbb{Z}, \chi_e}, \mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}$ of the coordinates of $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ are polynomially bounded. Hence, we can prevent this attack from succeeding by simply choosing the parameters $\chi_s, \chi_e, \tilde{\chi}_e$ of our noise vectors to be subexponentially larger than the magnitude $\sqrt{k} \cdot \beta_x$ of possible messages $\mathbf{x}$. By doing so, the adversary of [Üna20] and other statistical attackers need to query a subexponential number of ciphertexts to provably approximate the mean of $g(\mathbf{x})$ with sufficient precision.

Coming back to $\text{sk-}k\text{-ALWE}^{\otimes d}$, we expect statistical attacks against it to proceed similarly as statistical attack against our FE scheme. In fact, the statistical attack against our FE scheme in Fig. 5 where an attacker tries to approximate $\mathbb{E}[\langle \text{err}_0(\mathbf{x}_0^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2]$ and $\mathbb{E}[\langle \text{err}_0(\mathbf{x}_1^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2]$, leads to a distinguishing attack against $\text{sk-}k\text{-ALWE}^{\otimes d}$ where an adversary tries to approximate $\mathbb{E}[\langle \text{err}_0(\mathbf{x}^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2]$ and $\mathbb{E}[\langle \text{err}_1(\mathbf{x}^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2]$ over the randomness of $(\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}})$ with sufficient precision (with arbitrary vectors $\mathbf{x}^*, \mathbf{y} \neq 0$). Again, choosing parameters $\chi_s, \chi_e, \tilde{\chi}_e$ larger than $T^{1/6d} \cdot \sqrt{k} \cdot \beta_x$, these kinds of attacks against $\text{sk-}k\text{-ALWE}^{\otimes d}$ require at least T ciphertexts from the challenger, ruling out attackers with time complexity lower than T . This does not rule out Strategy 3 in general: as argued in Remark 3, the distributions of $\mathbf{f}^*$ in Exp^0 and Exp^1 plausibly remain statistically far apart even for subexponentially large randomness, so no purely statistical argument (such as noise flooding) can rule out this strategy.

Strategy 4: Recover secret randomness algebraically. A possible approach to decide $k\text{-ALWE}$ is to algebraically extract the encryption randomness from the challenge ciphertext $(\mathbf{c}, \tilde{\mathbf{c}})$. Concretely, one solves for $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ in the polynomial equation system $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$ over $\mathbb{Z}$. Due to space limitations, and since Strategy 4 is subsumed by Strategy 5, we explain the details of this attack in Appendix D.2. Here, we summarize that the equation system is of degree d and overdetermined, since it has $n + m + \tilde{m} \in o(k^d)$ variables and k^d equations. Since the number of variables $n + m + \tilde{m} \geq (k + n) \log T$ is by a factor of $\log T$ larger than k , Lemma 19 in Appendix D.2 lower bounds the degree of regularity D by $\log T$. Heuristically, one uses $D + 1$ as a proxy for the solving degree of typical algebraic solvers. This implies a time complexity of at least T for such solvers.

7.2 Attack against $k\text{-ALWE}^{\otimes d}$ but not $\text{sk-}k\text{-ALWE}^{\otimes d}$

We next discuss Strategy 5, an attack that succeeds against $k\text{-ALWE}^{\otimes d}$ in our experiments. The efficacy of this attack is what leads us to restrict our FE construction only to the secret-key setting. Given Strategy 5, an adversary that manages to extract $\mathbf{A}, \tilde{\mathbf{A}}$ from an $\text{sk-}k\text{-ALWE}^{\otimes d}$ instance would reduce $\text{sk-}k\text{-ALWE}^{\otimes d}$

⁹ The main rationale being here that we have no reason to assume that a secret key $(\mathbf{K}\mathbf{y}, \tilde{\mathbf{K}}\mathbf{y})$ for a harmless function $\mathbf{y} \neq 0$ should not induce a noticeable statistical distance on the challenge ciphertexts.

to $k\text{-ALWE}^{\otimes d}$, and so would inherit whatever Strategy 5 achieves against the latter. Therefore, we argue in Section 7.2 that it is hard to extract $\mathbf{A}, \tilde{\mathbf{A}}$ from the lattice preimages $\mathbf{K}, \tilde{\mathbf{K}}$ given in an $\text{sk-}k\text{-ALWE}^{\otimes d}$ instance.

Strategy 5: Recover secret randomness by optimization algorithms. This strategy again decides b by solving the equation system $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$. This time, however, the attacker exploits the fact that the equation system is given over the reals and uses an optimisation algorithm, for example, the gradient descent (GD) method. This is an iterative algorithm that, given a differentiable function $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$, optimizes a loss function $\|F(\mathbf{x}) - \mathbf{y}\|^2$ by starting with an initial guess $x^{(0)}$ and updating it in each step by $\mathbf{x}^{(i+1)} = \mathbf{x}^{(i)} - \eta \cdot \nabla F(\mathbf{x}^{(i)})^T (F(\mathbf{x}^{(i)}) - \mathbf{y})$. By adapting the size of update steps $\eta > 0$ appropriately, it is likely that GD converges towards a (local) minimum of $\|F(\mathbf{x}) - \mathbf{y}\|$. In the case $b = 0$, a solution $(\mathbf{s}_0, \mathbf{e}_0, \tilde{\mathbf{e}}_0)$ is guaranteed to exist for $\text{err}_0(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$. The over-determinedness of $\text{err}_0(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$ makes it more likely that GD avoids getting stuck at local minima, hence converging towards the correct solution $(\mathbf{s}_0, \mathbf{e}_0, \tilde{\mathbf{e}}_0)$.

We ran this attack at small test parameters and observed high success probabilities.¹⁰ Our experiments do not extend to the parameters the assumption requires, so we do not claim a break at those parameters. Since the attacked system becomes more over-determined as the parameters grow, which should only make the attack easier, we do not rely on $k\text{-ALWE}^{\otimes d}$. Therefore, we conjecture only $\text{sk-}k\text{-ALWE}^{\otimes d}$ for $d > 1$, in which $(\mathbf{A}, \tilde{\mathbf{A}})$ and hence $\mathbf{F}$ are withheld. This way, the adversary does not have direct access to the coefficients of the polynomial system $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$, but only an oracle that returns $\text{err}_b(\mathbf{x}^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}})$. This makes doing anything non-black-box with the polynomial system seemingly impossible.

Hardness of extracting $\mathbf{A}, \tilde{\mathbf{A}}$. As explained above, $k\text{-ALWE}^{\otimes d}$, and the public-key variant of our scheme, are jeopardized by attacks based on optimisation algorithms over the real numbers. Hence, it is imperative that an adversary cannot extract the matrices $\mathbf{F}$ resp. $\mathbf{A}, \tilde{\mathbf{A}}$ in Assumption 1, as they would allow it to get a polynomial description of $\text{err}_0(\mathbf{x}, \cdot, \cdot, \cdot)$.

We argue here that it is infeasible for an adversary to extract $\mathbf{A}, \tilde{\mathbf{A}}$ in $\text{Exp}_{d, \text{par}, \mathcal{A}}^0(1^\lambda)$. (The argument for $\text{Exp}_{d, \text{par}, \mathcal{A}}^1(1^\lambda)$ is analogous.) Recall that for $Q \in \text{poly}(\lambda)$, the adversary receives $((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}}, \mathbf{K}, \tilde{\mathbf{K}})$ in $\text{Exp}_{d, \text{par}, \mathcal{A}}^0(1^\lambda)$ where $\mathbf{A} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$, $\tilde{\mathbf{A}} \leftarrow \mathbb{Z}_q^{[[k+n]]_d^1 \times \tilde{m}}$ and

$$(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}) \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{m^d \times k^d} \times \mathcal{D}_{\mathbb{Z}, \sigma}^{\tilde{m} \times k^d} \times \mathcal{D}_{\mathbb{Z}, \sigma}^{[[k+n]]_{d-1}^1 \times k^d}$$

conditioned on having, over $\mathbb{Z}_q$,

$$\left[\begin{array}{c|c} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d \\ \hline \{\mathbf{A}, \mathbf{I}_m\}_{\# \mathbf{A}=i}^{\otimes d} & \tilde{\mathbf{A}}_i \\ \vdots & \vdots \\ \{ \underset{\# \mathbf{A} > 0}{\mathbf{A}, \mathbf{I}_m} \}_{i: d-1 \rightarrow 1}^{\otimes d} & \underbrace{\tilde{\mathbf{A}}}_{i: d-1 \rightarrow 1} \end{array} \right] p^i \cdot \mathbf{I}_{\binom{d}{i} (k+n)^i m^{d-i}} \cdot \left[\begin{array}{c} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{array} \right] = \left[\begin{array}{c} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{0}_{[[k+n]]_{d-1}^1 \times k^d} \end{array} \right]^{\otimes d}. \quad (23)$$

For $i = d$, the d -th block of Eq. (23)

$$\mathbf{A}^{\otimes d} \cdot \mathbf{K} + \tilde{\mathbf{A}}_d \cdot \tilde{\mathbf{K}} = \left[\begin{array}{c} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{array} \right]^{\otimes d} \quad (24)$$

gives an inhomogeneous system of $(n+k)^d \cdot k^d$ equations of degree d over $\mathbf{A}$ and linear over $\tilde{\mathbf{A}}_d \in \mathbb{Z}_q^{(n+k)^d \times \tilde{m}}$. For $i < d$, the i -th block of Eq. (23)

$$\{\mathbf{A}, \mathbf{I}_m\}_{\# \mathbf{A}=i}^{\otimes d} \cdot \mathbf{K} + \tilde{\mathbf{A}}_i \cdot \tilde{\mathbf{K}} + p^i \cdot \mathbf{F}_i = \mathbf{0}_{\binom{d}{i} (k+n)^i m^{d-i} \times k^d} \quad (25)$$

gives a homogeneous system of $(k+n)^i m^{d-i} \binom{d}{i} \cdot k^d$ equations which are linear in the unknown entries of $\tilde{\mathbf{A}}_i \in \mathbb{Z}_q^{\binom{d}{i} (k+n)^i m^{d-i} \times \tilde{m}}$, $\mathbf{F}_i \in \mathbb{Z}^{\binom{d}{i} (k+n)^i m^{d-i} \times k^d}$ and of degree i over the entries of $\mathbf{A} \in \mathbb{Z}_q^{(n+k) \times m}$ (where we decompose $\tilde{\mathbf{A}}$ and $\mathbf{F}$ vertically into $\tilde{\mathbf{A}}_1, \dots, \tilde{\mathbf{A}}_d$ and $\mathbf{F}_1, \dots, \mathbf{F}_d$).

¹⁰ For reproducibility, we provide a Python implementation (`Strategy5.py`) of our experiments on Strategy 5.

Below, we discuss 3 potential methods for extracting $\mathbf{A}$, $\tilde{\mathbf{A}}$ or $\mathbf{F}$ from $\mathbf{K}$, $\tilde{\mathbf{K}}$:

1. *Solve Eq. (25) for $i = 1, \dots, d-1$.* This would possibly allow to obtain $\mathbf{A}$, $\tilde{\mathbf{A}}$ or useful information on them. Note, however, that this is futile. For $i \in [d-1]$, every block equation system Eq. (25) has $(k+n)^i m^{d-i} \binom{d}{i} \cdot k^d$ unknown entries of $\mathbf{F}_i$. While the entries of $\mathbf{F}_i$ are correlated to $\mathbf{A}$, $\tilde{\mathbf{A}}$, $\mathbf{K}$, $\tilde{\mathbf{K}}$ and short, it is not clear how one would exploit those properties algebraically. As the number of entries of $\mathbf{F}_i$ equals the number of equations of Eq. (25), every equation of Eq. (25) is effectively blinded by a random coordinate of $\mathbf{F}_i$. (Also note that q and p are coprime, so we cannot reduce p^i to zero.)

From an LWE perspective, we can parse Eq. (25) as an LWE-instance where $\tilde{\mathbf{K}}$ functions as public generator matrix, $\tilde{\mathbf{A}}_i$ as LWE secret and $p^i \mathbf{F}_i$ as short noise. The LWE-sample $\tilde{\mathbf{A}}_i \cdot \tilde{\mathbf{K}} + p^i \cdot \mathbf{F}_i$ intuitively blinds $\{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}}=i}^{\otimes d} \cdot \mathbf{K}$. Of course, a real LWE-sample (+ added message) does not equal zero, so this is pure intuition and we cannot hope to prove inextractability of $\mathbf{A}$, $\tilde{\mathbf{A}}$ from Eq. (25) under LWE. Nevertheless, we require, as in Strategies 1 and 2, that $\text{LWE}_{\tilde{m}, k^d, q, \mathcal{U}(\mathbb{Z}_q), p\sigma}$ is hard, which forces $\tilde{m}$ to grow with $\log T$.

2. *Solve Eq. (24) directly.* As the degree- d equation system Eq. (24) is noise-free and strongly overdetermined, one can attempt to solve it directly. First note that we can get rid of the term $\tilde{\mathbf{A}}_d \cdot \tilde{\mathbf{K}}$ by computing a full-rank matrix $\mathbf{R} \in \mathbb{Z}_q^{k^d \times (k^d - \tilde{m})}$ with $\tilde{\mathbf{K}}\mathbf{R} = \mathbf{0}$. This yields the equation system

$$\mathbf{A}^{\otimes d} \cdot \mathbf{K}\mathbf{R} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d} \mathbf{R} \quad (26)$$

of $(k+n)^d \cdot (k^d - \tilde{m}) \approx (k \cdot (k+n))^d$ equations of degree d over the $(k+n)m$ unknown entries of $\mathbf{A}$. For a Gröbner basis approach, the number of equations is asymptotically too low. Indeed, since $m \geq (k+n) \log T$ by Table 1, we have approximately $(k \cdot (k+n))^d$ degree- d equations over $(k+n)m$ variables. By Lemma 19, this lower bounds the degree of regularity D of the homogeneous top-terms of the system by $\frac{(k+n)m}{k \cdot (k+n)} = \frac{m}{k} \geq \log T$.

The minimum runtime for an algebraic solver $\binom{(k+n)m+D}{D+1}^2$ is therefore at least T .

3. *Solve SIS for $\tilde{\mathbf{K}}$.* If one can find a short $\mathbf{v} \in \mathbb{Z}^{k^d}$ with $\tilde{\mathbf{K}}\mathbf{v} = \mathbf{0}$, then Eq. (25) gives $\{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}}=i}^{\otimes d} \cdot \mathbf{K}\mathbf{v} = -p^i \cdot \mathbf{F}_i \mathbf{v}$, so $\mathbf{K}\mathbf{v}$ is a short kernel vector of $\{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}}=i}^{\otimes d}$ up to the error term $p^i \cdot \mathbf{F}_i \mathbf{v}$, provided the latter is also short. Collecting enough such vectors would yield an approximate trapdoor, and one could hope to turn it into one for $\mathbf{A}$ and use it to distinguish the ciphertexts.

However, while it is easy to find a short trapdoor for $\tilde{\mathbf{K}} \in \mathbb{Z}^{\tilde{m} \times k^d}$ for $\tilde{m} = 1$, it becomes infeasible for large $\tilde{m}$, even given the shortness of $\tilde{\mathbf{K}}$. A usual approach to compute such a short kernel vector $\mathbf{v}$ will incur a blow-up in the coordinates of $\mathbf{v}$ that is exponential in $\tilde{m}$, so it suffices to take $\tilde{m} = \Omega(\log q)$, which is subsumed by the stronger requirement of Method 1.

We conclude that all three methods do not manage to extract useful information on $\mathbf{A}$, $\tilde{\mathbf{A}}$ in less than subexponential time. Of course, one can try to additionally take the ciphertexts $\mathbf{c}, \tilde{\mathbf{c}}, (\mathbf{d}_j, \tilde{\mathbf{d}}_j)_j$ into account. However, since every ciphertext uses fresh independent noise $\mathbf{e}, \tilde{\mathbf{e}}$, algebraically the ciphertexts will not give new additional information $\mathbf{A}$, $\tilde{\mathbf{A}}$.

Concluding, we think it is infeasible for the adversary to extract $\mathbf{A}$, $\tilde{\mathbf{A}}$ resp. $\mathbf{F}$ in the experiments of Assumption 1. Hence, while Strategy 5 succeeds against $k\text{-ALWE}^{\otimes d}$ and the public-key version of our scheme in experiments at small parameters, $\text{sk-}k\text{-ALWE}^{\otimes d}$ should resist the attack strategies in Section 7.

Summary: We conjecture that $\text{sk-}k\text{-ALWE}^{\otimes d}$ is T -hard (Assumption 1) for parameters satisfying the constraints designed to resist Strategies 1–5, summarised as follows:

- $m \geq \Omega((k+n) \cdot \log q)$ (Falsifiability).
- $n, m, \tilde{m}, q, \chi_s, \chi_e, \tilde{\chi}_e$ should be such that $\text{LWE}_{n, m, q, \chi_s, \chi_e}$ and $\text{LWE}_{n, \tilde{m}, q, \chi_s, \tilde{\chi}_e}$ are hard against time- T adversaries (Strategies 1 and 2).
- $\chi_s, \chi_e, \tilde{\chi}_e$ need to be larger than $T^{1/6d} \cdot \beta_x \cdot \sqrt{k}$ (Strategy 3).
- $m \geq (k+n) \log T$ (Strategies 4 and 5).
- For $d > 1$ we conjecture only $\text{sk-}k\text{-ALWE}^{\otimes d}$, in which $\mathbf{A}, \tilde{\mathbf{A}}$ are withheld, and not $k\text{-ALWE}^{\otimes d}$ (Strategies 4 and 5).

AI Declaration. The authors acknowledge the use of generative AI tools during the preparation of this work. Specifically, Claude Opus 5 and GPT 5.6 were used for proofreading and improving the clarity of the paper. All AI-generated content was carefully reviewed and edited by the authors, who take full responsibility for the accuracy and originality of the work.

Acknowledgments. The research of Russell W. F. Lai and Ivy K. Y. Woo are supported by Research Council of Finland grants 358951 and 358950 respectively. The research of Valerio Cini is supported by the European Research Council through an ERC Starting Grant (grant agreement No. 101077455, ObfusQation) and an MSCA Postdoctoral Fellowship (grant agreement No. 101202597, ACryL). The research of Akin Ünal is supported in whole or in part by the Austrian Science Fund (FWF) 10.55776/ESP1735525. For Open Access purposes, the author has applied a CC BY public copyright license to any author accepted manuscript version arising from this submission. We thank Erkan Tairi and Akin Ünal for pointing out a critical flaw in an earlier version of this work. Subsequently, Ünal contributed significantly in the cryptanalysis of our assumption and construction and joined as a coauthor.

References

- ABB10. Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 98–115. Springer, Berlin, Heidelberg, August 2010. 29
- ABDP15. Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *PKC 2015*, volume 9020 of *LNCS*, pages 733–751. Springer, Berlin, Heidelberg, March / April 2015. 1, 2
- ACL⁺22. Martin R. Albrecht, Valerio Cini, Russell W. F. Lai, Giulio Malavolta, and Sri Aravinda Krishnan Thyagarajan. Lattice-based SNARKs: Publicly verifiable, preprocessing, and recursively composable - (extended abstract). In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part II*, volume 13508 of *LNCS*, pages 102–132. Springer, Cham, August 2022. 52
- Agr17. Shweta Agrawal. Stronger security for reusable garbled circuits, general definitions and attacks. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 3–35. Springer, Cham, August 2017. 2
- Agr19. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 191–225. Springer, Cham, May 2019. 52, 53
- AIK06. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *computational complexity*, 15(2):115–162, 2006. 59
- AJ15. Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 308–326. Springer, Berlin, Heidelberg, August 2015. 1, 2, 58, 59
- AJL⁺19. Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 284–332. Springer, Cham, August 2019. 52, 53
- AJS15. Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Indistinguishability obfuscation from functional encryption for simple functions. *Cryptology ePrint Archive*, Report 2015/730, 2015. 1, 59
- Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. 8
- AKY24a. Shweta Agrawal, Simran Kumari, and Shota Yamada. Compact pseudorandom functional encryption from evasive LWE. *Cryptology ePrint Archive*, Report 2024/1719, 2024. 52
- AKY24b. Shweta Agrawal, Simran Kumari, and Shota Yamada. Pseudorandom multi-input functional encryption and applications. *Cryptology ePrint Archive*, Report 2024/1720, 2024. 52
- ALS16. Shweta Agrawal, Benoît Libert, and Damien Stehlé. Fully secure functional encryption for inner products, from standard assumptions. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 333–362. Springer, Berlin, Heidelberg, August 2016. 1, 2, 57, 58
- AMYY25. Shweta Agrawal, Anuja Modi, Anshu Yadav, and Shota Yamada. Zeroizing attacks against evasive and circular evasive LWE. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part II*, volume 16269 of *LNCS*, pages 259–290. Springer, Cham, December 2025. 52, 53

- AR17. Shweta Agrawal and Alon Rosen. Functional encryption for bounded collusions, revisited. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 173–205. Springer, Cham, November 2017. 2
- AS17. Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 152–181. Springer, Cham, April / May 2017. 1
- Ban93. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993. 29
- Bar04. Magali Bardet. *Étude des systèmes algébriques surdéterminés. Applications aux codes correcteurs et à la cryptographie*. Theses, Université Pierre et Marie Curie - Paris VI, December 2004. 50
- BBKK18. Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K. Kothari. Limits on low-degree pseudo-random generators (or: Sum-of-squares meets program obfuscation). In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 649–679. Springer, Cham, April / May 2018. 52, 53
- BCFG17. Carmen Elisabetta Zaira Baltico, Dario Catalano, Dario Fiore, and Romain Gay. Practical functional encryption for quadratic functions with applications to predicate encryption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 67–98. Springer, Cham, August 2017. 1
- BDGM20. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 79–109. Springer, Cham, May 2020. 1, 52
- BDGM22. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for IO: Circular-secure LWE suffices. In Mikolaj Bojanczyk, Emanuela Merelli, and David P. Woodruff, editors, *ICALP 2022*, volume 229 of *LIPICs*, pages 28:1–28:20. Schloss Dagstuhl, July 2022. 1, 60
- BDJ⁺25. Pedro Branco, Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Pseudorandom obfuscation and applications. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part V*, volume 16004 of *LNCS*, pages 663–698. Springer, Cham, August 2025. 52
- BF11. Dan Boneh and David Mandell Freeman. Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures. In Dario Catalano, Nelly Fazio, Rosario Gennaro, and Antonio Nicolosi, editors, *PKC 2011*, volume 6571 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, March 2011. 4, 7, 29, 51
- BFS03. Magali Bardet, Jean-Charles Faugère, and Bruno Salvy. Complexity of Gröbner basis computation for Semi-regular Overdetermined sequences over F_2 with solutions in F_2 . Research Report RR-5049, INRIA, 2003. 50
- BHJ⁺19. Boaz Barak, Samuel B. Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 226–250. Springer, Cham, May 2019. 52, 53
- BJK15. Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 470–491. Springer, Berlin, Heidelberg, November / December 2015. 1
- BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Berlin, Heidelberg, August 2013. 57
- BNPW16. Nir Bitansky, Ryo Nishimaki, Alain Passelègue, and Daniel Wichs. From cryptomania to obfustopia through secret-key functional encryption. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 391–418. Springer, Berlin, Heidelberg, October / November 2016. 2, 59, 60
- BNPW20. Nir Bitansky, Ryo Nishimaki, Alain Passelègue, and Daniel Wichs. From cryptomania to obfustopia through secret-key functional encryption. *Journal of Cryptology*, 33(2):357–405, April 2020. 1, 2
- BSW11. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In Yuval Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 253–273. Springer, Berlin, Heidelberg, March 2011. 1, 9
- BÜW24. Chris Brzuska, Akin Ünäl, and Ivy K. Y. Woo. Evasive LWE assumptions: Definitions, classes, and counterexamples. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 418–449. Springer, Singapore, December 2024. 52, 53

- BV15. Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *56th FOCS*, pages 171–190. IEEE Computer Society Press, October 2015. 1, 2, 58, 59
- CC17. Ran Canetti and Yilei Chen. Constraint-hiding constrained PRFs for NC^1 from LWE . In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 446–476. Springer, Cham, April / May 2017. 2
- CGH⁺15. Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrede Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 247–266. Springer, Berlin, Heidelberg, August 2015. 53
- CHKP10. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 523–552. Springer, Berlin, Heidelberg, May / June 2010. 15, 30, 32
- CHKP12. David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *Journal of Cryptology*, 25(4):601–639, October 2012. 15, 32
- CKPS00. Nicolas Courtois, Alexander Klimov, Jacques Patarin, and Adi Shamir. Efficient algorithms for solving overdefined systems of multivariate polynomial equations. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 392–407. Springer, Berlin, Heidelberg, May 2000. 50
- CLM23. Valerio Cini, Russell W. F. Lai, and Giulio Malavolta. Lattice-based succinct arguments from vanishing polynomials - (extended abstract). In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 72–105. Springer, Cham, August 2023. 52
- CLW25. Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based obfuscation from NTRU and equivocal LWE . In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 39–72. Springer, Cham, August 2025. 1, 52, 53
- CVW18. Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part II*, volume 10992 of *LNCS*, pages 577–607. Springer, Cham, August 2018. 53
- DBM⁺08. Jintai Ding, Johannes Buchmann, Mohamed Saied Emam Mohamed, Wael Said Abd Elmageed Mohamed, and Ralf-Philipp Weinmann. Mutantxl. *SCC*, (TUD-CS-2009-0142):16–22, 01 2008. 50
- DJM⁺25. Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, and Vinod Vaikuntanathan. Simple and general counterexamples for private-coin evasive LWE . In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 73–92. Springer, Cham, August 2025. 52, 53
- DQV⁺21. Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct LWE sampling, random polynomials, and obfuscation. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part II*, volume 13043 of *LNCS*, pages 256–287. Springer, Cham, November 2021. 1, 52, 53
- GGH13a. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Berlin, Heidelberg, May 2013. 52
- GGH⁺13b. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013. 52, 58
- GGH15. Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Berlin, Heidelberg, March 2015. 52
- GJLS21. Romain Gay, Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from simple-to-state hard problems: New assumptions, new techniques, and simplification. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 97–126. Springer, Cham, October 2021. 52, 53
- GMPW20. Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 623–651. Springer, Cham, May 2020. 4, 8, 29
- GP21. Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 736–749. ACM Press, June 2021. 1, 52, 53

- GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 11, 17, 29, 30
- GW11. Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108. ACM Press, June 2011. 15
- HHY26. Tzu-Hsiang Huang, Wei-Hsiang Hung, and Shota Yamada. A note on obfuscation-based attacks on private-coin evasive LWE. *CiC*, 2(4):20, 2026. 52, 53
- HJ16. Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 537–565. Springer, Berlin, Heidelberg, May 2016. 53
- HJL21. Samuel B. Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO . In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 673–700, Virtual Event, August 2021. Springer, Cham. 52, 53
- HJL25a. Yao-Ching Hsieh, Aayush Jain, and Huijia Lin. Lattice-based post-quantum iO from circular security with random opening assumption. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 3–38. Springer, Cham, August 2025. 1, 52, 53
- HJL25b. Yao-Ching Hsieh, Aayush Jain, and Huijia Lin. Lattice-based post-quantum iO from circular security with random opening assumption (part II: zeroizing attacks against private-coin evasive LWE assumptions). Cryptology ePrint Archive, Report 2025/390, 2025. 52, 53
- JKLS21. Aayush Jain, Alexis Korb, Paul Lou, and Amit Sahai. A note on the pseudorandomness of low-degree polynomials over the integers. Cryptology ePrint Archive, Report 2021/1415, 2021. 52, 53
- JLLS23. Aayush Jain, Huijia Lin, Paul Lou, and Amit Sahai. Polynomial-time cryptanalysis of the subspace flooding assumption for post-quantum iO . In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 205–235. Springer, Cham, April 2023. 52, 53
- JLMS19. Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over $\mathbb{R}$ to build iO . In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 251–281. Springer, Cham, May 2019. 52, 53
- JLS19. Aayush Jain, Huijia Lin, and Amit Sahai. Simplifying constructions and assumptions for iO . Cryptology ePrint Archive, Report 2019/1252, 2019. 52, 53
- JLS21. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73. ACM Press, June 2021. 1
- JLS22. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_p$, DLIN, and PRGs in NC^0 . In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 670–699. Springer, Cham, May / June 2022. 1
- KNT18. Fuyuki Kitagawa, Ryo Nishimaki, and Keisuke Tanaka. Obfustopia built on secret-key functional encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 603–648. Springer, Cham, April / May 2018. 1, 58
- Lin17. Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 599–629. Springer, Cham, August 2017. 1
- LPSS14. San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 315–334. Springer, Berlin, Heidelberg, August 2014. 3, 53, 57
- LPST16a. Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation with non-trivial efficiency. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 447–462. Springer, Berlin, Heidelberg, March 2016. 1, 2, 60
- LPST16b. Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Output-compressing randomized encodings and applications. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 96–124. Springer, Berlin, Heidelberg, January 2016. 1, 58
- LT17. Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 630–660. Springer, Cham, August 2017. 52
- LV16. Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In Irit Dinur, editor, *57th FOCS*, pages 11–20. IEEE Computer Society Press, October 2016. 1, 2, 58, 59

- MG02. Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a Cryptographic Perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, 2002. 30
- MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. 11, 17
- MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th FOCS*, pages 372–381. IEEE Computer Society Press, October 2004. 11
- MSZ16. Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658. Springer, Berlin, Heidelberg, August 2016. 53
- O’N10. Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010. 9
- Par10. Keith Pardue. Generic sequences of polynomials. *Journal of Algebra*, 324(4):579–590, 2010. 50
- RVV24. Seyoon Ragavan, Neekon Vafa, and Vinod Vaikuntanathan. Indistinguishability obfuscation from bilinear maps and LPN variants. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part IV*, volume 15367 of *LNCS*, pages 3–36. Springer, Cham, December 2024. 1
- Sal25. Flavio Salizzoni. An upper bound for the solving degree in terms of the degree of regularity. *Transactions on Mathematical Cryptology*, 5(1):1–7, Sep. 2025. 50
- Tsa22. Rotem Tsabary. Candidate witness encryption from lattice techniques. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part I*, volume 13507 of *LNCS*, pages 535–559. Springer, Cham, August 2022. 52, 53
- TÜ24. Erkan Tairi and Akin Ünal. Lower bounds for lattice-based compact functional encryption. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part II*, volume 14652 of *LNCS*, pages 249–279. Springer, Cham, May 2024. 2, 7, 20, 21
- Üna20. Akin Ünal. Impossibility results for lattice-based functional encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 169–199. Springer, Cham, May 2020. 20, 21
- Üna23. Akin Ünal. Worst-case subexponential attacks on PRGs of constant degree or constant locality. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part I*, volume 14004 of *LNCS*, pages 25–54. Springer, Cham, April 2023. 51
- VWW22. Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and null-IO from evasive LWE. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part I*, volume 13791 of *LNCS*, pages 195–221. Springer, Cham, December 2022. 52, 53
- Wat12. Brent Waters. Functional encryption for regular languages. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 218–235. Springer, Berlin, Heidelberg, August 2012. 1
- Wee20. Hoeteck Wee. Functional encryption for quadratic functions from k -lin, revisited. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part I*, volume 12550 of *LNCS*, pages 210–228. Springer, Cham, November 2020. 1, 5
- WW21. Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part III*, volume 12698 of *LNCS*, pages 127–156. Springer, Cham, October 2021. 1, 52, 53
- WW23. Hoeteck Wee and David J. Wu. Succinct vector, polynomial, and functional commitments from lattices. In Carmit Hazay and Martijn Stam, editors, *EUROCRYPT 2023, Part III*, volume 14006 of *LNCS*, pages 385–416. Springer, Cham, April 2023. 29, 52

A Extended Preliminaries

A.1 Additional Lemmas on Gaussians

Lemma 3 (Adapted from [Ban93, Lemma 1.5]). *For any $m \in \mathbb{N}$, $\chi > 0$,*

$$\Pr[\|\mathbf{x}\| > \chi\sqrt{m} \mid \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^m, \chi}] \leq 2^{-m}.$$

Lemma 4 (Adapted from [GPV08, Lem. 5.3]). *Let n, q be positive integers with q prime, and let $m \geq 2n \log q$. Then for all but an at most q^{-n} fraction of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and for any $\omega(\sqrt{\log m})$, there is a negligible function $\epsilon(m)$ such that $\eta_\epsilon(\Lambda_q^\perp(\mathbf{A})) \leq \omega(\sqrt{\log m})$.*

For a lattice $\Lambda \subset \mathbb{R}^n$, we call a real subspace $V \subseteq \mathbb{R}^n$ a Λ -subspace if

$$\text{span}_{\mathbb{R}}(V \cap \Lambda) = V.$$

For a full-column-rank matrix $\mathbf{S}$, we write $\sigma_{\min}(\mathbf{S})$ for its smallest singular value. Following [GMPW20, Def. 2.4, eprint], for a lattice $\Lambda \subseteq \text{span}_{\mathbb{R}}(\mathbf{S})$, we write $\eta_\epsilon(\Lambda) \leq \mathbf{S}$ if $\eta_\epsilon(\mathbf{S}^+ \Lambda) \leq 1$, where $\mathbf{S}^+$ denotes the Moore–Penrose pseudoinverse of $\mathbf{S}$.

Lemma 5 ([GMPW20, Thm. 3.1, eprint]). *For any $\epsilon \in [0, 1)$, matrix $\hat{\mathbf{S}}$ of full column rank, lattice coset $\hat{A} = \Lambda + \mathbf{a} \subset \text{span}(\hat{\mathbf{S}})$, and matrix $\mathbf{T}$ such that $\ker(\mathbf{T})$ is a Λ -subspace and $\eta_\epsilon(\Lambda \cap \ker(\mathbf{T})) \leq \hat{\mathbf{S}}$, we have*

$$\Delta(\mathbf{T} \cdot \mathcal{D}_{\hat{A}, \hat{\mathbf{S}}}, \mathcal{D}_{\mathbf{T}\hat{A}, \mathbf{T}\hat{\mathbf{S}}}) \leq \frac{\epsilon}{1 - \epsilon}.$$

Lemma 6 ([ABB10, Thm. 33, eprint]). *Let $\mathcal{L}$ be an m -dimensional lattice and $\sigma > t \cdot \eta_\epsilon(\mathcal{L})$ for some $t > 1$ and $\epsilon > 0$. The probability that m vectors sampled independently from $\mathcal{D}_{\mathcal{L}, \sigma}$ are linearly independent over $\mathbb{R}^m$ is at least $1 - \frac{1+\epsilon}{t-1}$.*

Lemma 7 (Adapted¹¹ from [BF11, Thm. 4.3, eprint]). *Suppose $m \geq 2n \log q$, $m > 2k$, $\sigma > \omega(\sqrt{\log m})$, and q is prime. Let $\mathbf{Y} \in \mathbb{Z}_q^{n \times k}$ be arbitrary. The following distributions are statistically close in n :*

$$\left\{ (\mathbf{A}, \mathbf{U}) \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{U} \leftarrow \mathcal{D}_{\mathcal{L}_q^{\mathbf{Y}}(\mathbf{A}), \sigma} \end{array} \right\} \quad \text{and} \quad \left\{ (\mathbf{A}, \mathbf{U}) \mid \begin{array}{l} \mathbf{U} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k} \\ \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} : \mathbf{A}\mathbf{U} = \mathbf{Y} \bmod q \end{array} \right\}.$$

Corollary 2. *Suppose $m \geq 2(n + \ell) \log q$, $m > 2k$, $\sigma > \omega(\sqrt{\log m})$, and q is prime. Let $\mathbf{Y} \in \mathbb{Z}_q^{n \times k}$ and $\mathbf{Z} \in \mathbb{Z}_q^{\ell \times k}$ be arbitrary. The following distributions are statistically close in λ :*

$$\left\{ (\mathbf{A}, \mathbf{B}, \mathbf{U}) \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}; \mathbf{U} \leftarrow \mathcal{D}_{\mathcal{L}_q^{\mathbf{Y}}(\mathbf{A}), \sigma} \\ \mathbf{B} \leftarrow \mathbb{Z}_q^{\ell \times m} : \mathbf{B}\mathbf{U} = \mathbf{Z} \bmod q \end{array} \right\} \quad \text{and} \quad \left\{ (\mathbf{A}, \mathbf{B}, \mathbf{U}) \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}; \mathbf{B} \leftarrow \mathbb{Z}_q^{\ell \times m} \\ \mathbf{U} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k} : \begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix} \mathbf{U} = \begin{bmatrix} \mathbf{Y} \\ \mathbf{Z} \end{bmatrix} \bmod q \end{array} \right\}.$$

Proof. Applying Lemma 7 on $(\mathbf{A}, \mathbf{U})$, the LHS distribution is statistically close in n to first sampling $\mathbf{U} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k}$ without constraint, then sample random $\mathbf{A}$ subject to $\mathbf{A} \cdot \mathbf{U} = \mathbf{Y} \bmod q$, and the rest for $\mathbf{B}$ is unchanged. Then, applying Lemma 7 again, such intermediate distribution is statistically close in $n + \ell$ to first sample random $\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix}$ then sample $\mathbf{U}$. We arrive at the RHS distribution. $\square$

The following is adapted from [WW23], where we consider a matrix $\mathbf{Y}$ instead of a column vector $\mathbf{y}$. The statement follows from repeating the original lemma over each column of $\mathbf{Y}$.

Lemma 8 (Adapted from [WW23, Lem. 2.10]). *Let $n, m, q \in \mathbb{N}$, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be primitive over $\mathbb{Z}_q$, and $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{n \times \tilde{m}}$ be arbitrary, where $\tilde{m} \in \text{poly}(n, \log q)$. Then for all $\mathbf{Y} \in \mathbb{Z}_q^{n \times k}$ where $k \in \text{poly}(m)$, and all $\sigma \geq \eta_\epsilon(\Lambda_q^\perp(\mathbf{A}))$ where $\epsilon \leq \text{negl}(m)$, the statistical distance between the following distributions is $\text{negl}(m)$:*

$$\left[\begin{array}{c} \mathbf{K} \\ \tilde{\mathbf{K}} \end{array} \right] \leftarrow \mathcal{D}_{\mathcal{L}_q^{\mathbf{Y}}(\mathbf{A}, \tilde{\mathbf{A}}), \sigma} \quad \text{and} \quad \left\{ \left[\begin{array}{c} \mathbf{K} \\ \tilde{\mathbf{K}} \end{array} \right] \mid \begin{array}{l} \tilde{\mathbf{K}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\tilde{m} \times k}; \tilde{\mathbf{Y}} := \mathbf{Y} - \tilde{\mathbf{A}}\tilde{\mathbf{K}} \bmod q \\ \mathbf{K} \leftarrow \mathcal{D}_{\mathcal{L}_q^{\tilde{\mathbf{Y}}}(\mathbf{A}), \sigma} \end{array} \right\}.$$

¹¹ [BF11] stated the result with $\mathbf{Y} = \mathbf{0}$, although the same proof applies to arbitrary $\mathbf{Y} \in \mathbb{Z}_q^{n \times k}$.

A.2 Additional Lemmas on Gaussian Sampling

We recall the classic Gaussian-sampling algorithm from [GPV08].

Lemma 9 ([GPV08, Thm. 4.1]). *There is a probabilistic polynomial-time algorithm that, given a basis $\mathbf{B}$ of an m -dimensional lattice $\Lambda = \mathcal{L}(\mathbf{B})$, a parameter $s \geq \|\mathbf{GS}(\mathbf{B})\| \cdot \omega(\sqrt{\log m})$, and a center $\mathbf{c} \in \mathbb{R}^m$, outputs a sample from a distribution that is statistically close to $\mathcal{D}_{\Lambda, s, \mathbf{c}}$.*

Both Lemma 2 and Corollary 1 rest on the following classic conversion, which turns a full-rank set into a basis that is no worse in Gram-Schmidt norm.

Lemma 10 ([MG02, Lem. 7.1], as restated in [CHKP10, Lem. 2.1 (ePrint)]). *There is a deterministic polynomial-time algorithm $\text{ToBasis}(\mathbf{S}, \mathbf{B})$ that, on input a full-rank set $\mathbf{S} = (\mathbf{s}_1, \dots, \mathbf{s}_m)$ of the lattice $\Lambda = \mathcal{L}(\mathbf{B})$ generated by a basis $\mathbf{B}$, outputs a basis $\mathbf{R} = (\mathbf{r}_1, \dots, \mathbf{r}_m)$ of Λ whose Gram-Schmidt orthogonalisation $\mathbf{GS}(\mathbf{R}) = (\hat{\mathbf{r}}_1, \dots, \hat{\mathbf{r}}_m)$ satisfies $\|\hat{\mathbf{r}}_i\| \leq \|\hat{\mathbf{s}}_i\|$ for every $i \in [m]$, where $\mathbf{GS}(\mathbf{S}) = (\hat{\mathbf{s}}_1, \dots, \hat{\mathbf{s}}_m)$. In particular, $\|\mathbf{GS}(\mathbf{R})\| \leq \|\mathbf{GS}(\mathbf{S})\|$.*

Since a basis of $\Lambda_q^\perp(\mathbf{A})$ is computable from $\mathbf{A}$ by linear algebra, e.g. via its Hermite normal form, Lemma 10 applies to q -ary lattices without any further input.

Proof (of Corollary 1). Let $\mathbf{B}$ be any basis of Λ and put $\mathbf{R} := \text{ToBasis}(\mathbf{S}, \mathbf{B})$. Then $\mathbf{R}$ is a basis of Λ , so $\tilde{bl}(\Lambda) \leq \|\mathbf{GS}(\mathbf{R})\| \leq \|\mathbf{GS}(\mathbf{S})\|$ by Lemma 10. $\square$

Proof (of Lemma 2). We let GPVSamp be the PPT algorithm as follows: Compute a basis $\mathbf{B}$ of $\Lambda_q^\perp(\mathbf{A})$ from $\mathbf{A}$ by linear algebra, e.g. via its Hermite normal form, and put $\mathbf{B}' := \text{ToBasis}(\mathbf{S}, \mathbf{B})$, which by Lemma 10 is a basis of $\Lambda_q^\perp(\mathbf{A})$ with $\|\mathbf{GS}(\mathbf{B}')\| \leq \|\mathbf{GS}(\mathbf{S})\| \leq s/\omega(\sqrt{\log m})$. Then find arbitrary $\mathbf{c} \in \mathbb{Z}^m$ satisfying $\mathbf{A} \cdot \mathbf{c} = -\mathbf{v} \bmod q$, run the PPT algorithm given by Lemma 9 on $\mathbf{B}'$, s and $\mathbf{c}$ to obtain some $\mathbf{u} \in \mathbb{Z}^m$, and return $\mathbf{u} - \mathbf{c}$. Both added steps are deterministic and depend only on $\mathbf{A}$ and $\mathbf{S}$. Lemma 9 guarantees that $\mathbf{u}$ is distributed statistically close to $\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s, \mathbf{c}} = \mathbf{c} + \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}) - \mathbf{c}, s} = \mathbf{c} + \mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s}$, hence $\mathbf{u} - \mathbf{c}$ is close to $\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}), s}$ as wanted. $\square$

A.3 Additional Computational Assumption

In Definition 3 we recall the standard LWE assumption parametrised by Gaussian distribution of secret and error terms with parameter χ_s and χ_e respectively.

Definition 3 (LWE). *Let $\text{par} = (n, m, q, \chi_s, \chi_e)$ be parametrised by λ , where n, m, q are positive integers and χ_s, χ_e are positive real numbers. The LWE_{par} assumption states that every PPT adversary $\mathcal{A}$ satisfies*

$$\text{Adv}_{\text{par}, \mathcal{A}}^{\text{lwe}}(\lambda) := \left| \Pr \left[\mathcal{A}(\mathbf{A}, \mathbf{b}) = 1 \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n, \mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}^m \\ \mathbf{b}^\top := \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \bmod q \end{array} \right] - \Pr \left[\mathcal{A}(\mathbf{A}, \mathbf{b}) = 1 \mid \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m} \\ \mathbf{b} \leftarrow \mathbb{Z}_q^m \end{array} \right] \right| \leq \text{negl}(\lambda).$$

B Supplementary Material for Section 4

B.1 The Degree-3 Case

We illustrate the matrix $\mathcal{E}^{\otimes 3}(\mathbf{A})$ and its full-rank set $\mathbf{T}$ constructed in Theorem 1.

$$\begin{aligned}
\mathcal{E}^{\otimes 3}(\mathbf{A}) &= \begin{matrix} & \mathbf{0}_3 & (1,0,0) & (0,1,0) & (0,0,1) & (1,1,0) & (1,0,1) & (0,1,1) \\ \begin{matrix} \mathbf{0}_3 \\ (1,0,0) \\ (0,1,0) \\ (0,0,1) \\ (1,1,0) \\ (1,0,1) \\ (0,1,1) \end{matrix} & \begin{bmatrix} \mathbf{A} \otimes \mathbf{A} \otimes \mathbf{A} \\ \mathbf{I}_m \otimes \mathbf{A} \otimes \mathbf{A} \\ \mathbf{A} \otimes \mathbf{I}_m \otimes \mathbf{A} \\ \mathbf{A} \otimes \mathbf{A} \otimes \mathbf{I}_m \\ \mathbf{I}_m \otimes \mathbf{I}_m \otimes \mathbf{A} \\ \mathbf{I}_m \otimes \mathbf{A} \otimes \mathbf{I}_m \\ \mathbf{A} \otimes \mathbf{I}_m \otimes \mathbf{I}_m \end{bmatrix} \end{matrix} \\
\mathbf{T} &= \begin{matrix} & \mathbf{0}_3 & (1,0,0) & (0,1,0) & (0,0,1) & (1,1,0) & (1,0,1) & (0,1,1) \\ \begin{matrix} \mathbf{0}_3 \\ (1,0,0) \\ (0,1,0) \\ (0,0,1) \\ (1,1,0) \\ (1,0,1) \\ (0,1,1) \end{matrix} & \begin{bmatrix} \mathbf{U}_0 \otimes \mathbf{U}_0 \otimes \mathbf{U}_0 & \mathbf{U}_0 \otimes \mathbf{U}_1 \otimes \mathbf{U}_1 & \mathbf{U}_1 \otimes \mathbf{U}_0 \otimes \mathbf{U}_1 & \mathbf{U}_1 \otimes \mathbf{U}_1 \otimes \mathbf{U}_0 & \mathbf{U}_0 \otimes \mathbf{U}_0 \otimes \mathbf{U}_1 & \mathbf{U}_0 \otimes \mathbf{U}_1 \otimes \mathbf{U}_0 & \mathbf{U}_1 \otimes \mathbf{U}_0 \otimes \mathbf{U}_0 \\ -\mathbf{U}_0 \otimes \mathbf{I}_h \otimes \mathbf{I}_h & & -\mathbf{I}_h \otimes \mathbf{U}_0 \otimes \mathbf{I}_h & & & & \\ & & & -\mathbf{I}_h \otimes \mathbf{I}_h \otimes \mathbf{U}_0 & & & \\ -\mathbf{U}_0 \otimes \mathbf{U}_1 \otimes \mathbf{I}_h & -\mathbf{U}_1 \otimes \mathbf{U}_0 \otimes \mathbf{I}_h & & & -\mathbf{U}_0 \otimes \mathbf{U}_0 \otimes \mathbf{I}_h & & \\ -\mathbf{U}_0 \otimes \mathbf{I}_h \otimes \mathbf{U}_1 & & -\mathbf{U}_1 \otimes \mathbf{I}_h \otimes \mathbf{U}_0 & & & -\mathbf{U}_0 \otimes \mathbf{I}_h \otimes \mathbf{U}_0 & \\ & & -\mathbf{I}_h \otimes \mathbf{U}_0 \otimes \mathbf{U}_1 & -\mathbf{I}_h \otimes \mathbf{U}_1 \otimes \mathbf{U}_0 & & & -\mathbf{I}_h \otimes \mathbf{U}_0 \otimes \mathbf{U}_0 \end{bmatrix} \end{matrix}
\end{aligned}$$

B.2 Additional Lemmas

Lemma 11. Suppose $\mathbf{M} \in \mathbb{Z}_q^{h \times m}$ is primitive over $\mathbb{Z}_q$. Consider $\mathcal{E}^{\otimes d}(\mathbf{M})$ and $\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}})$ for some $\tilde{\mathbf{M}} \leftarrow \mathbb{Z}_q^{[[h]]_d^1 \times \tilde{m}}$, as defined in Eqs. (15) and (16). If the columns of $\mathbf{T}$ are a full-rank set of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{M}))$, then there exists a full-rank set $\mathbf{T}^*$ of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}}))$ satisfying $\|\mathbf{GS}(\mathbf{T}^*)\| = \|\mathbf{GS}(\mathbf{T})\|$.

Proof. Define the shorthands $\mathcal{E}^{\otimes d}(\mathbf{M}) = (\hat{\mathbf{M}}_0, \hat{\mathbf{M}}_1)$ and therefore $\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}}) = (\hat{\mathbf{M}}_0, \tilde{\mathbf{M}}, \hat{\mathbf{M}}_1)$, where $\hat{\mathbf{M}}_0, \hat{\mathbf{M}}_1$ are the left and right block columns of $\mathcal{E}^{\otimes d}(\mathbf{M})$, and $\tilde{\mathbf{M}}$ is the additional middle block column in $\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}})$.

Since $\mathbf{M}$ is primitive, so is $\mathcal{E}^{\otimes d}(\mathbf{M})$ by Fact 1, thus there exists $\mathbf{W} \in \mathbb{Z}^{(m^d + [[h]]_{d-1}^1) \times \tilde{m}}$ such that $\mathcal{E}^{\otimes d}(\mathbf{M}) \cdot \mathbf{W} = \tilde{\mathbf{M}} \bmod q$. Let the columns of $\mathbf{T}$ be a full-rank set of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{M}))$, i.e. $\mathcal{E}^{\otimes d}(\mathbf{M}) \cdot \mathbf{T} = \mathbf{0} \bmod q$ and $\mathbf{T}$ is invertible over $\mathbb{Q}$. Let $\mathbf{T}^* := \begin{bmatrix} \mathbf{T} & \mathbf{W} \\ -\mathbf{I}_{\tilde{m}} \end{bmatrix}$, and observe that

$$(\mathcal{E}^{\otimes d}(\mathbf{M}), \tilde{\mathbf{M}}) \cdot \mathbf{T}^* = (\mathcal{E}^{\otimes d}(\mathbf{M}), \tilde{\mathbf{M}}) \cdot \begin{bmatrix} \mathbf{T} & \mathbf{W} \\ -\mathbf{I}_{\tilde{m}} \end{bmatrix} = \mathbf{0} \bmod q,$$

and $\mathbf{T}^*$ is invertible over $\mathbb{Q}$, hence its columns are linearly independent. Since $\mathbf{T}^*$ is square of size $m^d + [[h]]_{d-1}^1 + \tilde{m}$, which is the dimension of the lattice, its columns are a full-rank set of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{M}), \tilde{\mathbf{M}}) = \Lambda_q^\perp(\hat{\mathbf{M}}_0, \tilde{\mathbf{M}}, \hat{\mathbf{M}}_1) = \Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{M}, \tilde{\mathbf{M}}))$, where the last two lattices agree up to the permutation of coordinates that moves $\tilde{\mathbf{M}}$ past $\hat{\mathbf{M}}_1$, which preserves Gram-Schmidt norms. Finally, notice that $\|\mathbf{GS}(\mathbf{T}^*)\| = \|\mathbf{GS}(\mathbf{T})\|$ by design of $\mathbf{T}^*$. $\square$

Lemma 12. Take the parameters specified by Theorem 1. The algorithm `GenFullRankSet` returns with overwhelming probability in λ a matrix $\mathbf{U}_0 \in \mathbb{Z}^{m \times m}$ that is invertible over $\mathbb{Q}$, upon being called by `TrapGen` $_{h,m,p,q}^{\otimes d}$.

Proof. We inspect `GenFullRankSet`. In each round of its for-loop, it samples a $\mathbf{U}_0$, checks if it has full rank over $\mathbb{Q}$, and returns it if so. It aborts if no $\mathbf{U}_0$ is full rank after λ trials. By properties of `SampPre` (Definition 2), each sampled $\mathbf{U}_0$ follows $\mathcal{D}_{\Lambda_q^\perp(\mathbf{A}),s}^m$. Conditioned on $\sigma \geq 3\eta_\epsilon(\Lambda_q^\perp(\mathbf{A}))$ for some negligible function ϵ , which happens with overwhelming probability by Lemma 4, we use Lemma 6 which states that each $\mathbf{U}_0$ is not full rank over $\mathbb{Q}$ with probability at most $\frac{1+\epsilon}{2}$. The overall probability that it returns $\perp$, i.e. all λ samples are not full rank over $\mathbb{Q}$, is at most $(\frac{1+\epsilon}{2})^\lambda$, negligible. $\square$

B.3 Proof of Theorem 2

Proof. We consider the following sequence of distributions of $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$:

Hyb_0 : The output $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}_{h,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}, s^*)$.

Hyb₁ : The matrix $\mathbf{A}$ is swapped to uniformly random over $\mathbb{Z}_q^{h \times m}$.

Since $\mathbf{A}$ is generated by running $\text{TrapGen}_{h,m,p,q}^{\otimes d}(1^\lambda)$, we have $\text{Hyb}_0 \approx_s \text{Hyb}_1$ by the statistical closeness of $\mathbf{A}$ to uniform in Theorem 1.

Hyb₂ : We sample $\begin{bmatrix} \mathbf{K} \\ \mathbf{F} \end{bmatrix}$ inefficiently from $\mathcal{D}_{\Lambda_q^{\mathbf{V}-\tilde{\mathbf{A}} \cdot \tilde{\mathbf{K}}}(\mathcal{E}^{\otimes d}(\mathbf{A})), s^*}$, instead of by running $\text{GPVSamp}(\mathcal{E}^{\otimes d}(\mathbf{A}), \mathbf{T}, \mathbf{V} - \tilde{\mathbf{A}}\tilde{\mathbf{K}}, s^*)$.

Theorem 1 Items 1–2 guarantee that the columns of $\mathbf{T}$ are a full-rank set of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}))$ with $\|\text{GS}(\mathbf{T})\| \leq 2(s\sqrt{m})^d$. Then, $\text{Hyb}_1 \approx_s \text{Hyb}_2$ follows from Lemma 2.

Hyb₃ : We change the preimage sampling procedure: We let $(\mathbf{K}, \tilde{\mathbf{K}})$ be the top and bottom blocks obtained from $\begin{bmatrix} \mathbf{K} \\ \mathbf{F} \\ \tilde{\mathbf{K}} \end{bmatrix} \leftarrow_s \mathcal{D}_{\Lambda_q^{\mathbf{V}}(\mathcal{E}^{\otimes d}(\mathbf{A}), \tilde{\mathbf{A}}), s^*}$, where we recall that from Equation (16) one has

$$\mathcal{E}^{\otimes d}(\mathbf{A}) = \left[\begin{array}{c|c} \mathbf{A}^{\otimes d} & \\ \hline \underbrace{\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d}}_{\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}>0}^{\otimes d}} & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \end{array} \right].$$

Notice that the matrix $\mathcal{E}^{\otimes d}(\mathbf{A})$ is primitive, i.e. $\mathcal{E}^{\otimes d}(\mathbf{A}) \cdot \mathbb{Z}_q^{m^d + \lceil [h] \rceil_{d-1}^1} = \mathbb{Z}_q^{\lceil [h] \rceil_d^1}$, with overwhelming probability in λ , because of Fact 1 together with that $\mathbf{A}$ is uniformly random. Then, $\text{Hyb}_2 \approx_s \text{Hyb}_3$ follows from trapdoor delegation [CHKP10, CHKP12], or Lemma 8. To use Lemma 8, we note that $\eta_\epsilon(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}))) \leq s^*$ is satisfied, because of Lemma 1, Theorem 1 Items 1–2, and Corollary 1, which together give $\tilde{bl}(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{A}))) \leq \|\text{GS}(\mathbf{T})\| \leq \|\mathbf{T}\|$.

Hyb₄ : We rename the uniform matrix $\begin{bmatrix} \tilde{\mathbf{A}}_d \\ \tilde{\mathbf{A}}_i \\ \vdots \\ \tilde{\mathbf{A}}_{i:d-1 \rightarrow 1} \end{bmatrix} := \tilde{\mathbf{A}}$, and we express the preimage sampling procedure as

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{bmatrix} \leftarrow_s \mathcal{D}_{\mathbb{Z}, s^*}^{(m^d + \tilde{m} + \lceil [h] \rceil_{d-1}^1) \times k^d} : \left[\begin{array}{c|c|c} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d & \\ \hline \underbrace{\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d}}_{\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}>0}^{\otimes d}} & \underbrace{\tilde{\mathbf{A}}_i}_{\tilde{\mathbf{A}}} & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots & \vdots & \ddots \\ \tilde{\mathbf{A}}_{i:d-1 \rightarrow 1} & \tilde{\mathbf{A}}_{i:d-1 \rightarrow 1} & \tilde{\mathbf{A}}_{i:d-1 \rightarrow 1} \end{array} \right] \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{bmatrix} = \mathbf{V} \bmod q.$$

This change is only syntactical, and $\text{Hyb}_3 \equiv \text{Hyb}_4$.

Notice that Hyb_4 is identical to the distribution in Eq. (18). The proof is completed. $\square$

Corollary 3. *Take the parameters stated in Theorem 2, and let n, k be such that $h = k + n$. Let*

$$\mathbf{V}_0 := \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) = \left[\begin{array}{c} \left[\begin{array}{c} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{array} \right]^{\otimes d} \\ \mathbf{0}_{\lceil [k+n] \rceil_{d-1}^1 \times k^d} \end{array} \right] \in \mathbb{Z}_q^{(\lceil [k+n] \rceil_d^d + \lceil [k+n] \rceil_{d-1}^1) \times k^d}. \quad (27)$$

With overwhelming probability in λ , the algorithm $\text{ExtTrapGen}_{h,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}_0, s^*)$ in Fig. 2 returns $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$ such that, when letting

$$\mathcal{E}^{\otimes d}(\mathbf{A}) = \left[\begin{array}{c|c} \mathbf{A}^{\otimes d} & \\ \hline \underbrace{\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d}}_{\#\mathbf{A}>0} & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots_{i:d-1 \rightarrow 1} & \ddots_{i:d-1 \rightarrow 1} \end{array} \right],$$

parsing

$$\tilde{\mathbf{A}} = \left[\begin{array}{c} \tilde{\mathbf{A}}_d \\ \hline \tilde{\mathbf{A}}_i \\ \vdots_{i:d-1 \rightarrow 1} \end{array} \right] \quad \text{and} \quad \mathbf{P} = \left[\begin{array}{c} p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots_{i:d-1 \rightarrow 1} \end{array} \right].$$

it holds

$$\begin{aligned} \left[\begin{array}{c|c} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d \\ \hline \{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d} & \tilde{\mathbf{A}}_i \\ \vdots_{i:d-1 \rightarrow 1} & \vdots_{i:d-1 \rightarrow 1} \end{array} \right] \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \end{bmatrix} &= \left[\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}>i}^{\otimes d} \middle| \tilde{\mathbf{A}} \right] \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \end{bmatrix} \\ &= \left[\begin{array}{c} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ -\mathbf{P}\mathbf{F} \end{array} \right]^{\otimes d} \bmod q \quad \text{for some } \mathbf{F} \in \mathbb{Z}^{[[k+n]]_{d-1}^1 \times k^d}, \end{aligned}$$

moreover $\|\mathbf{K}\| \leq s^* \sqrt{m^d}$, $\|\tilde{\mathbf{K}}\| \leq s^* \sqrt{\tilde{m}}$ and $\|\mathbf{F}\| \leq s^* \sqrt{[[k+n]]_{d-1}^1}$.

Proof. We invoke Theorem 2 for the stated $\mathbf{V}_0$, so that it suffices to consider the distribution in Eq. (18). Expanding $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$ according to Eq. (15), we have

$$\begin{aligned} &\underbrace{\left[\begin{array}{c|c|c} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d & \\ \hline \{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}=i}^{\otimes d} & \tilde{\mathbf{A}}_i & p^i \cdot \mathbf{I}_{\binom{d}{i} h^i m^{d-i}} \\ \vdots_{i:d-1 \rightarrow 1} & \vdots_{i:d-1 \rightarrow 1} & \ddots_{i:d-1 \rightarrow 1} \end{array} \right]}_{\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})} \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{bmatrix} = \mathbf{V}_0 \\ &= \left[\begin{array}{c} \left[\begin{array}{c} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{array} \right]^{\otimes d} \\ \mathbf{0}_{[[k+n]]_{d-1}^1 \times k^d} \end{array} \right] \bmod q. \end{aligned}$$

Rearranging the last block row, we arrive at the claimed linear relation. The norm bounds on $\mathbf{K}$, $\tilde{\mathbf{K}}$ and $\mathbf{F}$ are immediate from standard Gaussian tail bound (cf. Lemma 3). $\square$

C Proofs of Section 6

C.1 Proof of Theorem 3

Proof. The ciphertext is given by

$$(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$$

for some $\mathbf{s} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_s}^n$. This means that

$$\begin{aligned}\mathbf{c}^\top &= \left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A} \right\rangle + p \cdot \mathbf{e}^\top, \\ \tilde{\mathbf{c}}^\top &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \tilde{\mathbf{A}} \right\rangle + p^d \cdot \tilde{\mathbf{e}}^\top\end{aligned}$$

for some $\mathbf{e} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_e}^m, \tilde{\mathbf{e}} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}^{\tilde{m}}$. For the first ciphertext component, we express its degree- d self-tensor as

$$\begin{aligned}(\mathbf{c}^\top)^{\otimes d} &= \left(\left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A} \right\rangle + p \cdot \mathbf{e}^\top \right)^{\otimes d} \\ &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}^{\otimes d}, \{\mathbf{A}, \mathbf{I}_m\} \right\rangle^{\otimes d} \\ &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \{\mathbf{A}, \mathbf{I}_m\}^{\otimes d} \right\rangle \bmod q.\end{aligned}$$

Then, for honestly generated keys and ciphertexts, we observe that the following equalities hold modulo q :

$$\begin{aligned}\langle \mathbf{c}^{\otimes d}, \mathbf{k}_y \rangle &= \left\langle \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \{\mathbf{A}, \mathbf{I}_m\}^{\otimes d} \right\rangle^\top, \mathbf{K} \cdot \mathbf{y} \right\rangle \\ &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \{\mathbf{A}, \mathbf{I}_m\}^{\otimes d} \cdot \mathbf{K} \right\rangle \cdot \mathbf{y} \\ &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}} > 0}^{\otimes d} \cdot \mathbf{K} \right\rangle \cdot \mathbf{y} + p^d \cdot \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle \cdot \mathbf{y} \bmod q,\end{aligned}\tag{28}$$

and

$$\begin{aligned}\langle \tilde{\mathbf{c}}, \tilde{\mathbf{k}}_y \rangle &= \left\langle \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \tilde{\mathbf{A}} \right\rangle^\top, \tilde{\mathbf{K}} \cdot \mathbf{y} \right\rangle + p^d \cdot \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle \cdot \mathbf{y} \\ &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \tilde{\mathbf{A}} \cdot \tilde{\mathbf{K}} \right\rangle \cdot \mathbf{y} + p^d \cdot \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle \cdot \mathbf{y} \bmod q\end{aligned}\tag{29}$$

where we have used that $\langle \langle \mathbf{S}, \mathbf{A} \rangle^\top, \mathbf{K} \mathbf{Y} \rangle = \langle \mathbf{S}, \mathbf{A} \mathbf{K} \rangle \cdot \mathbf{Y}$ for any matrices $\mathbf{S}, \mathbf{A}, \mathbf{K}, \mathbf{Y}$ of suitable dimensions. Moreover, by Theorem 2:

$$\begin{aligned}\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}}) \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{F} \end{bmatrix} &= \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) \bmod q \\ \underbrace{\begin{bmatrix} \mathbf{A}^{\otimes d} & \tilde{\mathbf{A}}_d \\ \{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}}=i}^{\otimes d} & \tilde{\mathbf{A}}_i \\ \vdots & \vdots \\ i:d-1 \rightarrow 1 & i:d-1 \rightarrow 1 \end{bmatrix}}_{\{\mathbf{A}, \mathbf{I}_m\}_{\#_{\mathbf{A}} > 0}^{\otimes d}} \cdot \underbrace{\begin{bmatrix} \tilde{\mathbf{A}}_d & \tilde{\mathbf{A}}_i \\ \tilde{\mathbf{A}}_i & \tilde{\mathbf{A}}_i \\ \vdots & \vdots \\ i:d-1 \rightarrow 1 & i:d-1 \rightarrow 1 \end{bmatrix}}_{\tilde{\mathbf{A}}} &= \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) \bmod q\end{aligned}\tag{30}$$

with $\|\mathbf{F}\| \leq \sigma \sqrt{[k+n]_{d-1}^1}$ by Corollary 3 with $s^* = \sigma$. Let us introduce the following notation for the right-down block of the matrix $\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})$:

$$\mathbf{P} = \begin{bmatrix} p^i \cdot \mathbf{I}_{\binom{d}{i}(k+n)^i m^{d-i}} & & \\ & \ddots & \\ & & i:d-1 \rightarrow 1 \end{bmatrix}.$$

Again from Corollary 3, we have that

$$\left[\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}>0}^{\otimes d} | \tilde{\mathbf{A}} \right] \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \end{bmatrix} = \begin{bmatrix} \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d} \\ -\mathbf{P} \cdot \mathbf{F} \end{bmatrix} \bmod q,$$

Thus

$$\begin{aligned} \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \left[\{\mathbf{A}, \mathbf{I}_m\}_{\#\mathbf{A}>0}^{\otimes d} | \tilde{\mathbf{A}} \right] \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \end{bmatrix} \right\rangle &= \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{\#_{p \cdot \mathbf{e}} < d}^{\otimes d}, \begin{bmatrix} \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d} \\ -\mathbf{P} \cdot \mathbf{F} \end{bmatrix} \right\rangle \\ &= \left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}^{\otimes d}, \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d} \right\rangle + \hat{\mathbf{e}}^T \\ &= \left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right\rangle^{\otimes d} + \hat{\mathbf{e}}^T \\ &= (\mathbf{x}^T)^{\otimes d} + \hat{\mathbf{e}}^T \bmod q, \end{aligned}$$

where the error term is

$$\hat{\mathbf{e}}^T := \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{0 < \#_{p \cdot \mathbf{e}} < d}^{\otimes d}, -\mathbf{P} \cdot \mathbf{F} \right\rangle = -p^d \cdot \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F} \right\rangle.$$

Note that the last equality in the above holds because, expanding the definition of $\mathbf{P}$ yields

$$\begin{aligned} &\left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, p \cdot \mathbf{e} \right\}_{0 < \#_{p \cdot \mathbf{e}} < d}^{\otimes d}, - \begin{bmatrix} p^i \cdot \mathbf{I}_{\binom{d}{i}(k+n)^i m^{d-i}} & & \\ & \ddots & \\ & & \mathbf{I}_{\binom{d}{i}(k+n)^i m^{d-i}} \end{bmatrix} \cdot \mathbf{F} \right\rangle \\ &= -p^d \cdot \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \begin{bmatrix} \mathbf{I}_{\binom{d}{i}(k+n)^i m^{d-i}} & & \\ & \ddots & \\ & & \mathbf{I}_{\binom{d}{i}(k+n)^i m^{d-i}} \end{bmatrix} \cdot \mathbf{F} \right\rangle \\ &= -p^d \cdot \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F} \right\rangle. \end{aligned}$$

Combining with Eqs. (28) and (29), we arrive at

$$\begin{aligned} \langle \mathbf{c}^{\otimes d}, \mathbf{k}_y \rangle + \langle \tilde{\mathbf{c}}, \tilde{\mathbf{k}}_y \rangle &= \langle \mathbf{x}^{\otimes d}, \mathbf{y} \rangle + \left(\hat{\mathbf{e}}^T + p^d \cdot \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + p^d \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle \right) \cdot \mathbf{y} \\ &= \langle \mathbf{x}^{\otimes d}, \mathbf{y} \rangle + p^d \left(- \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F} \right\rangle + \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle \right) \cdot \mathbf{y} \\ &= f(\mathbf{x}) + r \cdot \underbrace{\left(- \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F} \right\rangle + \langle \mathbf{e}^{\otimes d}, \mathbf{K} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle \right)}_{:= \tilde{\mathbf{e}}} \cdot \mathbf{y} \bmod q, \end{aligned}$$

where we use that $p^d = r \bmod q$. Finally, since $\|\mathbf{K}\| \leq \sigma \sqrt{m^d}$, $\|\tilde{\mathbf{K}}\| \leq \sigma \sqrt{\tilde{m}}$ by Corollary 3 (taking $s^* = \sigma$), and $\|(\mathbf{x}^T, \mathbf{s}^T, \mathbf{e}^T)\| \leq \sqrt{k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m}$ using $\|\mathbf{x}\| \leq \sqrt{k} \cdot \|\mathbf{x}\|_\infty \leq \sqrt{k} \cdot \beta_x$ and Lemma 3, we have

$$|\hat{\mathbf{e}}| \leq \left(\sum_{i=1}^{d-1} \binom{d}{i} \cdot \|(\mathbf{x}^T, \mathbf{s}^T)\|^i \cdot \|\mathbf{e}\|^{d-i} \cdot \sigma \sqrt{[(k+n)]_{d-1}^1} + \|\mathbf{e}\|^d \cdot \sigma \sqrt{m^d} + \|\tilde{\mathbf{e}}\| \cdot \sigma \sqrt{\tilde{m}} \right) \cdot \|\mathbf{y}\|_1$$

$$\begin{aligned}
&\leq \left(2^d \cdot \sigma \cdot (k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m)^{d/2} \cdot \sqrt{(k+n+m)^d} + \sigma \cdot (\chi_e \sqrt{m})^d \cdot \sqrt{m^d} + \tilde{\chi}_e \cdot \sigma \cdot \tilde{m} \right) \cdot \beta_y \\
&\leq \left(2^{d+1} \cdot \sigma \cdot (k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m)^{d/2} \cdot (k+n+m)^{d/2} + \tilde{\chi}_e \cdot \sigma \cdot \tilde{m} \right) \cdot \beta_y \\
&= \sigma \cdot \left((2 \cdot \sqrt{k+n+m})^d \cdot 2 \cdot (k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m)^{d/2} + \tilde{\chi}_e \cdot \tilde{m} \right) \cdot \beta_y = \gamma,
\end{aligned}$$

where we implicitly use $\|\mathbf{x}^\top \mathbf{K} \mathbf{y}\| \leq \|\mathbf{x}^\top\| \cdot \|\mathbf{K}\| \cdot \|\mathbf{y}\|_1$ for any vectors $\mathbf{x}, \mathbf{y}$ and matrix $\mathbf{K}$ of appropriate dimensions, together with $\|\mathbf{u} \otimes \mathbf{v}\| = \|\mathbf{u}\| \cdot \|\mathbf{v}\|$. The latter is applied to each of the $\binom{d}{i}$ blocks of the vector paired with $\mathbf{F}$ indexed by a tuple with i occurrences of $(\mathbf{x}^\top, \mathbf{s}^\top)$ and $d-i$ occurrences of $\mathbf{e}$, whose ℓ_2 norm is $\|(\mathbf{x}^\top, \mathbf{s}^\top)\|^i \cdot \|\mathbf{e}\|^{d-i}$, and to the vector $\mathbf{e}^{\otimes d}$ paired with $\mathbf{K}$, whose ℓ_2 norm is $\|\mathbf{e}\|^d$. The second inequality then follows from $\max\{\|(\mathbf{x}^\top, \mathbf{s}^\top)\|, \|\mathbf{e}\|\} \leq \sqrt{k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m}$, from $\|\mathbf{e}\| \leq \chi_e \sqrt{m}$ and $\|\tilde{\mathbf{e}}\| \leq \tilde{\chi}_e \sqrt{\tilde{m}}$ by Lemma 3, and from $\sum_{i=1}^{d-1} \binom{d}{i} \leq 2^d$ and $[[k+n]]_{d-1}^1 \leq (k+n+m)^d$; the third follows from $(\chi_e \sqrt{m})^d \cdot \sqrt{m^d} \leq (k\beta_x^2 + \chi_s^2 \cdot n + \chi_e^2 \cdot m)^{d/2} \cdot (k+n+m)^{d/2}$. Together with $f(\mathbf{x}) \in (-r/2, r/2] \cap \mathbb{Z}$, we have $|f(\mathbf{x}) + r \cdot \hat{e}| < q/2$, and therefore decryption yields

$$(\langle \mathbf{c}^{\otimes d}, \mathbf{k}_y \rangle + \langle \tilde{\mathbf{c}}, \tilde{\mathbf{k}}_y \rangle \bmod q) \bmod r = f(\mathbf{x}) \bmod r = f(\mathbf{x})$$

as desired. $\square$

C.2 Proof of Theorem 4

In the following subsection, for the ease of exposition, we will introduce the following notation.

For lattice cosets A_i of lattices Λ_i of dimension n , with $i = 1, \dots, \text{height}$, and $\mathbf{S} \in \mathbb{R}^{n \cdot \text{height} \times k}$, we let

$$\mathcal{D}_{\mathcal{M}(A_1, \dots, A_{\text{height}}), \mathbf{S}} := \{\mathbf{U} : \text{vec}(\mathbf{U}) \sim \mathcal{D}_{A_1 \oplus \dots \oplus A_{\text{height}}, \mathbf{S}}\} \quad (31)$$

When $\Lambda_1 = \dots = \Lambda_{\text{height}} = \Lambda_q^\perp(\mathbf{A})$, and $A_i = \Lambda_q^{\mathbf{v}_i}(\mathbf{A})$, we will use the shorthand $\Lambda_q^\mathbf{V}(\mathbf{A})$, where $\mathbf{V} = [\mathbf{v}_1 \dots \mathbf{v}_{\text{height}}]$, to denote $\mathcal{M}(A_1, \dots, A_{\text{height}})$, so that

$$\mathcal{D}_{\mathcal{M}(\Lambda_q^{\mathbf{v}_1}(\mathbf{A}), \dots, \Lambda_q^{\mathbf{v}_{\text{height}}}(\mathbf{A})), \mathbf{S}} = \mathcal{D}_{\Lambda_q^\mathbf{V}(\mathbf{A}), \mathbf{S}} = \{\mathbf{U} : \text{vec}(\mathbf{U}) \sim \mathcal{D}_{\Lambda_q^{\mathbf{v}_1}(\mathbf{A}) \oplus \dots \oplus \Lambda_q^{\mathbf{v}_{\text{height}}}(\mathbf{A}), \mathbf{S}}\}.$$

Proof. Consider the following sequence of hybrids described in Fig. 6. Each hybrid is defined by specifying only the changes with respect to the preceding hybrid. All algorithms and oracle checks not explicitly modified remain unchanged.

- $\text{Hyb}_{b,0}$: The security experiment $\text{INDExp}_{\text{FE}, \mathcal{A}}^b$. The adversary declares messages $(\mathbf{x}_1, \dots, \mathbf{x}_Q)$ and challenge messages $(\mathbf{x}_0^*, \mathbf{x}_1^*)$. The experiment samples $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}_{k+n, m, \tilde{m}, p, q}^{\otimes d}(\mathbf{1}^\lambda, \mathbf{V}_0, \sigma)$, for

$$\begin{aligned}
\mathbf{V}_0 &= \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \right) = \left[\left\{ \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d}, \mathbf{0}_{m \times k} \right\}_{d > \# \mathbf{0}_{m \times k} > 0} \right] = \begin{bmatrix} \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}^{\otimes d} \\ \mathbf{0}_{[[k+n]]_{d-1}^1 \times k^d} \end{bmatrix} \\
\mathbf{S}_0 &:= \mathbf{I}_{k^d} \otimes \sigma \cdot \mathbf{I}_{m^d + \tilde{m} + [[k+n]]_{d-1}^1}
\end{aligned} \quad (32)$$

as in the construction. The adversary is then given ciphertexts $(\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}$ encrypting $(\mathbf{x}_j)_{j \in [Q]}$ and computed as

$$(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$$

where $\mathbf{s}_j \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, and a challenge ciphertext $(\mathbf{c}, \tilde{\mathbf{c}})$ encrypting $\mathbf{x}_b^*$ and computed as

$$(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_b^* \\ \mathbf{s}^* \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$$

$\text{Hyb}_{b,0}(1^\lambda)$ $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_0^*, \mathbf{x}_1^*) \leftarrow \mathcal{A}(1^\lambda)$ $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}}) \leftarrow \text{ExtTrapGen}_{k+n,m,\tilde{m},p,q}^{\otimes d}(1^\lambda, \mathbf{V}_0, \sigma)$ for $j \in [Q]$ do : $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$ $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_b^* \\ \mathbf{s}^* \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$ $b' \leftarrow \mathcal{A}^{\text{KGen}((\mathbf{K}, \tilde{\mathbf{K}}), \cdot)}((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}})$ $\text{Hyb}_{b,2}(1^\lambda)$ $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_0^*, \mathbf{x}_1^*) \leftarrow \mathcal{A}(1^\lambda)$ $\mathbf{B} \leftarrow \mathbb{Z}_q^{(k+n+1) \times m}, \tilde{\mathbf{B}} \leftarrow \mathbb{Z}_q^{[[k+n+1]]_d^1 \times \tilde{m}}$ $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}) \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_1}$ for $j \in [Q]$ do : $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_k \\ \mathbf{s}^* \\ 1 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ $b' \leftarrow \mathcal{A}^{\text{KGen}((\mathbf{K}, \tilde{\mathbf{K}}), \cdot)}((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}})$	$\text{Hyb}_{b,1}(1^\lambda)$ $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_0^*, \mathbf{x}_1^*) \leftarrow \mathcal{A}(1^\lambda)$ $\mathbf{A} \leftarrow \mathbb{Z}_q^{(k+n) \times m}, \tilde{\mathbf{A}} \leftarrow \mathbb{Z}_q^{[[k+n]]_d^1 \times \tilde{m}}$ $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}) \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_0}(\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})), \mathbf{S}_0}$ for $j \in [Q]$ do : $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$ $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_b^* \\ \mathbf{s}^* \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right)$ $b' \leftarrow \mathcal{A}^{\text{KGen}((\mathbf{K}, \tilde{\mathbf{K}}), \cdot)}((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}})$ $\text{Hyb}_{b,3}(1^\lambda)$ $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_0^*, \mathbf{x}_1^*) \leftarrow \mathcal{A}(1^\lambda)$ $\mathbf{B} \leftarrow \mathbb{Z}_q^{(k+n+1) \times m}, \tilde{\mathbf{B}} \leftarrow \mathbb{Z}_q^{[[k+n+1]]_d^1 \times \tilde{m}}$ $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}) \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_\mathbf{Q}}$ for $j \in [Q]$ do : $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$ $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_k \\ \mathbf{s}^* \\ 1 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ $b' \leftarrow \mathcal{A}^{\text{KGen}((\mathbf{K}, \tilde{\mathbf{K}}), \mathbf{Q}^{-1}(\cdot))}((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}})$
--	---

Fig. 6. Hybrids for the proof of Theorem 4. The shorthands $\mathbf{V}_0, \mathbf{V}_1^{(b)}, \mathbf{Q}, \mathbf{S}_0, \mathbf{S}_1, \mathbf{S}_\mathbf{Q}$ are defined in Eqs. (32) to (34). In every box, the KGen oracle first performs the validity check defined in $\text{Hyb}_{b,0}$ on the original query; the displayed expressions specify its answer when the query is valid. The $\text{Hyb}_{b,3}$ box depicts the case $\delta \neq 0$; when $\delta = 0$, we define $\text{Hyb}_{b,3} := \text{Hyb}_{b,2}$. Differences between consecutive hybrids are highlighted in **red**.

where $\mathbf{s}^* \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_s}^n$. The key-generation oracle, on input $\mathbf{y} \in \mathbb{Z}^{k^d}$, first checks whether

$$((\mathbf{x}_0^*)^{\otimes d})^T \mathbf{y} = ((\mathbf{x}_1^*)^{\otimes d})^T \mathbf{y}.$$

If this condition does not hold, it returns $\perp$. Otherwise, it returns

$$\text{sk}_{\mathbf{y}} := (\mathbf{K}\mathbf{y}, \tilde{\mathbf{K}}\mathbf{y}).$$

The experiment returns whatever the adversary returns.

- $\text{Hyb}_{b,1}$: This is exactly $\text{Hyb}_{b,0}$, except for the following changes to how $\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}$, and $\tilde{\mathbf{K}}$ are computed:
 - sample $\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{(k+n) \times m}$, $\tilde{\mathbf{A}} \leftarrow \$ \mathbb{Z}_q^{[[k+n]]_d^1 \times \tilde{m}}$
 - sample $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}) \leftarrow \$ \mathcal{D}_{\Lambda_q^{\mathbf{V}_0}(\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})), \mathbf{S}_0}$,
 - set $\text{msk} = (\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$.
- $\text{Hyb}_{b,2}$: This is exactly $\text{Hyb}_{b,1}$, except for the following changes to how $\mathbf{K}, \tilde{\mathbf{K}}, (\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}$, and $(\mathbf{c}, \tilde{\mathbf{c}})$ are computed:
 - sample $\mathbf{B} \leftarrow \$ \mathbb{Z}_q^{(k+n+1) \times m}$, $\tilde{\mathbf{B}} \leftarrow \$ \mathbb{Z}_q^{[[k+n+1]]_d^1 \times \tilde{m}}$,
 - for $j \in [Q]$, let

$$(\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$$

with $\mathbf{s}_j \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_s}^n$,

- set

$$(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_k \\ \mathbf{s}^* \\ 1 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$$

with $\mathbf{s}^* \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \chi_s}^n$,

- define

$$\begin{aligned} \mathbf{V}_1^{(b)} &= \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{0}^T \end{bmatrix} \right) + \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}_b^*)^T \end{bmatrix} \right) \\ \mathbf{S}_1 &:= \mathbf{I}_{k^d} \otimes \sigma \cdot \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1}. \end{aligned} \tag{33}$$

- sample $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}) \leftarrow \$ \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_1}$.
- $\text{Hyb}_{b,3}$: This is exactly $\text{Hyb}_{b,2}$ except for how $\mathbf{K}, \tilde{\mathbf{K}}$ are generated and the functional key queries are computed:
 - let $\boldsymbol{\delta} := (\mathbf{x}_1^*)^{\otimes d} - (\mathbf{x}_0^*)^{\otimes d}$. If $\boldsymbol{\delta} = \mathbf{0}$, define $\text{Hyb}_{b,3}$ to be identical to $\text{Hyb}_{b,2}$. Otherwise, define

$$H := \left\{ \mathbf{y} \in \mathbb{Z}^{k^d} : \boldsymbol{\delta}^T \mathbf{y} = 0 \right\},$$

and let $\mathbf{Q}$ be a deterministically chosen matrix whose columns form a $\mathbb{Z}$ -basis of H .

- define

$$\mathbf{Z}_b := \left[\begin{bmatrix} \mathbf{I}_{k^d} \\ ((\mathbf{x}_b^*)^{\otimes d})^T \end{bmatrix} \right], \quad \mathbf{S}_{\mathbf{Q}} := \sigma \left(((\mathbf{Q}^T \mathbf{Q})^{1/2} \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right). \tag{34}$$

- sample $(\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}) \leftarrow \$ \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}} \mathbf{Q}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_{\mathbf{Q}}}$.

- The validity check performed by the key-generation oracle remains unchanged. For every valid query $\mathbf{y}$, the oracle computes the unique¹² $\mathbf{a} = \mathbf{Q}^{-1}(\mathbf{y}) \in \mathbb{Z}^{k^d-1}$ satisfying

$$\mathbf{y} = \mathbf{Q}\mathbf{a},$$

and returns

$$\text{sk}_{\mathbf{y}} := (\mathbf{K}\mathbf{a}, \tilde{\mathbf{K}}\mathbf{a}).$$

We now analyze each adjacent pair of hybrids.

Lemma 13. *For parameters par as in Table 1 and every $b \in \{0, 1\}$, $\text{Hyb}_{b,0} \approx_s \text{Hyb}_{b,1}$.*

Proof. By the properties of the $\text{ExtTrapGen}^{\otimes d}$ (Figure 2) algorithm (Theorem 2), the distribution of $(\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{K}, \tilde{\mathbf{K}})$ are statistically indistinguishable between $\text{Hyb}_{b,0}$ and $\text{Hyb}_{b,1}$. Therefore $\text{Hyb}_{b,0} \approx_s \text{Hyb}_{b,1}$. $\square$

Lemma 14. *For parameters par as in Table 1 and every $b \in \{0, 1\}$, under the $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption, $\text{Hyb}_{b,1} \approx_c \text{Hyb}_{b,2}$.*

Proof. The idea is to show that the view of the adversary in both hybrids can be efficiently simulated given the instance

$$((\mathbf{d}_{\text{Inst},j}, \tilde{\mathbf{d}}_{\text{Inst},j})_{j \in [Q]}, (\mathbf{c}_{\text{Inst}}, \tilde{\mathbf{c}}_{\text{Inst}}), (\mathbf{K}_{\text{Inst}}, \tilde{\mathbf{K}}_{\text{Inst}}))$$

from the $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption. The reduction $\mathcal{B}_b$ works as follows:

- it receives $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_0^*, \mathbf{x}_1^*)$ from the adversary $\mathcal{A}$,
- it forwards $(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}_b^*)$ to the $\text{sk-}k\text{-ALWE}^{\otimes d}$ challenger, and receives back

$$((\mathbf{d}_{\text{Inst},j}, \tilde{\mathbf{d}}_{\text{Inst},j})_{j \in [Q]}, (\mathbf{c}_{\text{Inst}}, \tilde{\mathbf{c}}_{\text{Inst}}), (\mathbf{K}_{\text{Inst}}, \tilde{\mathbf{K}}_{\text{Inst}})),$$

- it sets

$$\begin{aligned} (\mathbf{d}_j, \tilde{\mathbf{d}}_j) &:= (\mathbf{d}_{\text{Inst},j}, \tilde{\mathbf{d}}_{\text{Inst},j}) && \text{for } j \in [Q], \\ (\mathbf{c}, \tilde{\mathbf{c}}) &:= (\mathbf{c}_{\text{Inst}}, \tilde{\mathbf{c}}_{\text{Inst}}), \\ (\mathbf{K}, \tilde{\mathbf{K}}) &:= (\mathbf{K}_{\text{Inst}}, \tilde{\mathbf{K}}_{\text{Inst}}), \end{aligned}$$

- it returns $(\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}$, and $(\mathbf{c}, \tilde{\mathbf{c}})$ and answers its functional key queries as follows. On input $\mathbf{y} \in \mathbb{Z}^{k^d}$, it first checks whether

$$((\mathbf{x}_0^*)^{\otimes d})^T \mathbf{y} = ((\mathbf{x}_1^*)^{\otimes d})^T \mathbf{y}.$$

If the equality does not hold, it returns $\perp$. Otherwise, it returns

$$(\mathbf{K}_{\text{Inst}}\mathbf{y}, \tilde{\mathbf{K}}_{\text{Inst}}\mathbf{y}).$$

Finally, when $\mathcal{A}$ returns a bit b' , the reduction $\mathcal{B}_b$ returns the same bit b' .

We claim that

- if the reduction is in the $\text{Exp}_{d,\text{par},\mathcal{B}_b}^0(1^\lambda)$ experiment from Fig. 3, the view of the adversary $\mathcal{A}$ is identical to $\text{Hyb}_{b,1}$,
- if the reduction is in the $\text{Exp}_{d,\text{par},\mathcal{B}_b}^1(1^\lambda)$ experiment from Fig. 3, the view of the adversary $\mathcal{A}$ is identical to $\text{Hyb}_{b,2}$.

¹² Since the columns of $\mathbf{Q}$ form a $\mathbb{Z}$ -basis of H , the map $\mathbf{Q} : \mathbb{Z}^{k^d-1} \rightarrow H$ sending $\mathbf{a}$ to $\mathbf{Q}\mathbf{a}$ is bijective. We write $\mathbf{Q}^{-1} : H \rightarrow \mathbb{Z}^{k^d-1}$ for its inverse, so that $\mathbf{Q}^{-1}(\mathbf{y})$ is the unique integer vector $\mathbf{a}$ satisfying $\mathbf{y} = \mathbf{Q}\mathbf{a}$.

It suffices to study the joint distribution of

$$((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, (\mathbf{c}, \tilde{\mathbf{c}}), (\mathbf{K}, \tilde{\mathbf{K}})).$$

If $\mathcal{B}_b$ is interacting with the $\text{sk-}k\text{-ALWE}^{\otimes d}$ challenger in the $\text{Exp}_{d, \text{par}, \mathcal{B}_b}^0(1^\lambda)$ experiment, then

$$\begin{aligned} (\mathbf{d}_j, \tilde{\mathbf{d}}_j) &\leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right) \quad \text{for } j \in [Q], \\ (\mathbf{c}, \tilde{\mathbf{c}}) &\leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_b^* \\ \mathbf{s}_b^* \end{bmatrix}, \mathbf{A}, \tilde{\mathbf{A}} \right) \\ (\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}) &\leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_0}(\mathcal{E}^{\otimes d}(\mathbf{A}, \tilde{\mathbf{A}})), \mathbf{S}_0} \end{aligned}$$

for $\mathbf{A} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$, $\tilde{\mathbf{A}} \leftarrow \mathbb{Z}_q^{[[k+n]]_d^1 \times \tilde{m}}$, $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, which generates a view of the adversary $\mathcal{A}$ identical to that in $\text{Hyb}_{b,1}$.

If $\mathcal{B}_b$ is interacting with the $\text{sk-}k\text{-ALWE}^{\otimes d}$ challenger in the $\text{Exp}_{d, \text{par}, \mathcal{B}_b}^1(1^\lambda)$ experiment, then

$$\begin{aligned} (\mathbf{d}_j, \tilde{\mathbf{d}}_j) &\leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right) \quad \text{for } j \in [Q], \\ (\mathbf{c}, \tilde{\mathbf{c}}) &\leftarrow \text{LWESamp}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_k \\ \mathbf{s}^* \\ 1 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right) \\ (\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}) &\leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_1} \end{aligned}$$

for $\mathbf{B} \leftarrow \mathbb{Z}_q^{(k+n+1) \times m}$, $\tilde{\mathbf{B}} \leftarrow \mathbb{Z}_q^{[[k+n+1]]_d^1 \times \tilde{m}}$, $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, $\mathbf{s}^* \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_s}^n$, which generates a view of the adversary $\mathcal{A}$ identical to that in $\text{Hyb}_{b,2}$.

We conclude that, under the $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption, one has $\text{Hyb}_{b,1} \approx_c \text{Hyb}_{b,2}$ $\square$

Lemma 15. *For parameters par as in Table 1 and every $b \in \{0, 1\}$,*

$$\text{Hyb}_{b,2} \approx_s \text{Hyb}_{b,3}.$$

Proof. If

$$\boldsymbol{\delta} = (\mathbf{x}_1^*)^{\otimes d} - (\mathbf{x}_0^*)^{\otimes d} = \mathbf{0},$$

then $\text{Hyb}_{b,3}$ is defined to be identical to $\text{Hyb}_{b,2}$, and the claim follows. Hence, for the remainder of the proof, we assume that $\boldsymbol{\delta} \neq \mathbf{0}$.

Notice that the only change between the two hybrids is how $\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G}$ are sampled and how the KGen oracle uses these matrices. It therefore suffices to show that the distribution of

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} \quad \text{with} \quad \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_1}$$

as in $\text{Hyb}_{b,2}$ is statistically close to the distribution of

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \leftarrow \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)}} \mathbf{Q}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_\mathbf{Q}}$$

as in $\text{Hyb}_{b,3}$.

By column-wise vectorisation,

$$\text{vec} \left(\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} \right) = \left(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \text{vec} \left(\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \right). \quad (35)$$

Moreover,

$$\begin{aligned} \left(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \mathbf{S}_1 \mathbf{S}_1^T \left(\mathbf{Q} \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) &= \sigma^2 \left(\mathbf{Q}^T \mathbf{Q} \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \\ &= \mathbf{S}_Q \mathbf{S}_Q^T. \end{aligned} \quad (36)$$

Thus, by Eq. (36) and the definition of discrete Gaussians, the parameters $(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) \mathbf{S}_1$ and $\mathbf{S}_Q$ define the same Gaussian distribution.

By Lemma 16, with overwhelming probability over the choice of $\mathbf{B}, \tilde{\mathbf{B}}$, the image coset under the transformation in Eq. (35) is

$$\Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right),$$

the kernel of the transformation is a subspace of the lattice underlying the input coset, and

$$\eta_\epsilon \left(\ker \left(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \cap \left(\Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \right)^{k^d} \right) \leq \sigma = \sigma_{\min}(\mathbf{S}_1).$$

Since $\mathbf{S}_1 = \sigma \mathbf{I}$, the preceding inequality is equivalent to

$$\eta_\epsilon \left(\ker \left(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}} \right) \cap \left(\Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \right)^{k^d} \right) \leq \mathbf{S}_1$$

in the matrix-smoothing notation of Lemma 5. Since $\mathbf{S}_1$ is full rank, all the hypotheses of Lemma 5 are satisfied. Together with Eq. (36), we obtain

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} \approx_s \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}), \mathbf{S}_Q)}. \quad (37)$$

It remains to compare the answers to functional key queries. Every valid query $\mathbf{y}$ belongs to H , and hence there is a unique $\mathbf{a} \in \mathbb{Z}^{k^d - 1}$ such that

$$\mathbf{y} = \mathbf{Q} \mathbf{a}.$$

The answer in $\text{Hyb}_{b,2}$ can therefore be written as

$$\begin{aligned} \text{sk}_{\mathbf{y}} &= (\mathbf{K} \mathbf{y}, \tilde{\mathbf{K}} \mathbf{y}) \\ &= (\mathbf{K} \mathbf{Q} \mathbf{a}, \tilde{\mathbf{K}} \mathbf{Q} \mathbf{a}). \end{aligned}$$

This is exactly how $\text{Hyb}_{b,3}$ answers the query, with $\mathbf{K} \mathbf{Q}, \tilde{\mathbf{K}} \mathbf{Q}$ replaced by the matrices sampled in that hybrid. The ciphertexts are identically distributed in the two hybrids.

We conclude that $\text{Hyb}_{b,2} \approx_s \text{Hyb}_{b,3}$, as claimed. $\square$

The proof above invokes the following technical lemma to verify the image-coset, kernel-subspace, and smoothing conditions required by Lemma 5 for the valid-query projection.

Lemma 16. *Adopt the notation in $\text{Hyb}_{b,3}$ and suppose that*

$$\boldsymbol{\delta} = (\mathbf{x}_1^*)^{\otimes d} - (\mathbf{x}_0^*)^{\otimes d} \neq \mathbf{0}.$$

Write $\boldsymbol{\delta} = (\delta_1, \dots, \delta_{k^d})^T$ and define

$$\bar{\delta} := \frac{\delta}{\gcd(|\delta_1|, \dots, |\delta_{k^d}|)}.$$

For parameters par as in Table 1, the following hold with overwhelming probability over the choice of $\mathbf{B}, \tilde{\mathbf{B}}$:

1.

$$\Lambda_q^{\mathbf{V}_1^{(b)}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \mathbf{Q} = \Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right); \quad (38)$$

2.

$$\ker \left(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \cap \left(\Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \right)^{k^d} = \left\{ \bar{\boldsymbol{\delta}} \otimes \mathbf{w} : \mathbf{w} \in \Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \right\}; \quad (39)$$

3. The real subspace $\ker(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1})$ is a lattice subspace with respect to $(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})))^{k^d}$.

4. There exists a negligible function ϵ such that

$$\eta_\epsilon \left(\ker \left(\mathbf{Q}^T \otimes \mathbf{I}_{m^d + \tilde{m} + [[k+n+1]]_{d-1}^1} \right) \cap \left(\Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right) \right)^{k^d} \right) \leq \sigma. \quad (40)$$

Proof. In order to ease notation, let

$$\text{rank} := m^d + \tilde{m} + [[k+n+1]]_{d-1}^1.$$

1. We prove Eq. (38) by double inclusion.

For the first inclusion, let

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \in \Lambda_q^{\mathbf{V}_1^{(b)}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right).$$

Then

$$\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} = \mathbf{V}_1^{(b)} \mathbf{Q} \pmod{q}.$$

Therefore,

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} \in \Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right),$$

which proves the first inclusion.

For the opposite inclusion, since $\bar{\boldsymbol{\delta}}$ is primitive and the columns of $\mathbf{Q}$ form a $\mathbb{Z}$ -basis of $\{\mathbf{y} \in \mathbb{Z}^{k^d} : \bar{\boldsymbol{\delta}}^T \mathbf{y} = 0\}$, there exists $\mathbf{v} \in \mathbb{Z}^{k^d}$ such that

$$\bar{\boldsymbol{\delta}}^T \mathbf{v} = 1 \quad \text{and} \quad [\mathbf{Q} \mathbf{v}] \in \text{GL}_{k^d}(\mathbb{Z}).$$

In particular,

$$[\mathbf{Q} \mathbf{v}]^{-1} \in \mathbb{Z}^{k^d \times k^d}.$$

By Theorem 1 and Fact 1, and since adjoining columns preserves primitivity, $\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})$ is primitive except with negligible probability. We condition on this event. Let

$$\begin{bmatrix} \widehat{\mathbf{K}} \\ \widehat{\tilde{\mathbf{K}}} \\ \widehat{\mathbf{G}} \end{bmatrix} \in \Lambda_q^{\mathbf{V}_1^{(b)} \mathbf{Q}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right).$$

By primitivity, there exists an integral vector $\mathbf{w}$ satisfying $\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \mathbf{w} = \mathbf{V}_1^{(b)} \mathbf{v} \pmod{q}$. Define

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} := \begin{bmatrix} \widehat{\mathbf{K}} \\ \widehat{\tilde{\mathbf{K}}} \\ \widehat{\mathbf{G}} \end{bmatrix} \mathbf{w} [\mathbf{Q} \mathbf{v}]^{-1}.$$

This matrix is integral. Moreover,

$$\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} = \left[\mathbf{V}_1^{(b)} \mathbf{Q} \mathbf{V}_1^{(b)} \mathbf{v} \right] [\mathbf{Q} \mathbf{v}]^{-1} = \mathbf{V}_1^{(b)} \pmod{q}.$$

Hence

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \in \Lambda_q^{\mathbf{V}_1^{(b)}} \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right).$$

Finally,

$$[\mathbf{Q} \mathbf{v}]^{-1} \mathbf{Q} = \begin{bmatrix} \mathbf{I}_{k^d-1} \\ \mathbf{0}_{1 \times (k^d-1)} \end{bmatrix},$$

and therefore

$$\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \mathbf{Q} = \left[\begin{bmatrix} \widehat{\mathbf{K}} \\ \widehat{\tilde{\mathbf{K}}} \\ \widehat{\mathbf{G}} \end{bmatrix} \middle| \mathbf{w} \right] \begin{bmatrix} \mathbf{I}_{k^d-1} \\ \mathbf{0}_{1 \times (k^d-1)} \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{K}} \\ \widehat{\tilde{\mathbf{K}}} \\ \widehat{\mathbf{G}} \end{bmatrix}.$$

This proves the opposite inclusion and hence Item 1.

2. We prove Eq. (39) by double inclusion.

For the first inclusion, let

$$\mathbf{w} \in \Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right).$$

Since $\mathbf{Q}^T \bar{\boldsymbol{\delta}} = \mathbf{0}$, we have

$$(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) (\bar{\boldsymbol{\delta}} \otimes \mathbf{w}) = (\mathbf{Q}^T \bar{\boldsymbol{\delta}}) \otimes \mathbf{w} = \mathbf{0}.$$

Thus, $\bar{\boldsymbol{\delta}} \otimes \mathbf{w}$ belongs to the lattice on the left-hand side of Eq. (39).

For the opposite inclusion, let

$$\mathbf{w}_1, \dots, \mathbf{w}_{k^d} \in \Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right)$$

and suppose that

$$(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) \begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_{k^d} \end{bmatrix} = \mathbf{0}.$$

Since the real kernel of $\mathbf{Q}^T$ is generated by $\bar{\boldsymbol{\delta}}$, there exists $\mathbf{w} \in \mathbb{R}^{\text{rank}}$ such that

$$\begin{bmatrix} \mathbf{w}_1 \\ \vdots \\ \mathbf{w}_{k^d} \end{bmatrix} = \bar{\boldsymbol{\delta}} \otimes \mathbf{w}.$$

Because $\bar{\boldsymbol{\delta}}$ is primitive, there exist $c_1, \dots, c_{k^d} \in \mathbb{Z}$ such that

$$\sum_{i=1}^{k^d} c_i \bar{\boldsymbol{\delta}}_i = \mathbf{1}.$$

It follows that

$$\mathbf{w} = \sum_{i=1}^{k^d} c_i \mathbf{w}_i \in \Lambda_q^\perp \left(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \right).$$

This proves the opposite inclusion and hence Item 2.

3. Since the columns of $\mathbf{Q}$ form a basis of H , the matrix $\mathbf{Q}$ has rank $k^d - 1$. Therefore, the real kernel of $\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}$ has dimension rank .

By Eq. (39), the intersection of this kernel with the product lattice is isomorphic to $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}))$, and hence has rank rank . It follows that

$$\text{span}_{\mathbb{R}} \left(\ker(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) \cap \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right)^{k^d} \right) = \ker(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}).$$

Thus, by definition, $\ker(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}})$ is a lattice subspace with respect to $\left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right)^{k^d}$.

4. The map

$$\mathbf{w} \mapsto \bar{\boldsymbol{\delta}} \otimes \mathbf{w}$$

scales Euclidean norms by $\|\bar{\boldsymbol{\delta}}\|$. Consequently,

$$\eta_\epsilon \left(\ker(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) \cap \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right)^{k^d} \right) = \|\bar{\boldsymbol{\delta}}\| \cdot \eta_\epsilon \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right). \quad (41)$$

The hypotheses of Theorem 1 hold with $h = k + n + 1$ by Table 1. Moreover, a uniform $\mathbf{B} \in \mathbb{Z}_q^{(k+n+1) \times m}$ is primitive except with negligible probability. By Theorem 1, the matrix output by $\text{TrapGen}_{k+n+1, m, p, q}^{\otimes d}(1^\lambda)$ is statistically close to uniform. It therefore suffices to consider $(\mathbf{B}, \mathbf{T}) \leftarrow \text{TrapGen}_{k+n+1, m, p, q}^{\otimes d}(1^\lambda)$.

By Lemma 11, there exists a full-rank set $\mathbf{T}^*$ of $\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}))$ satisfying $\|\text{GS}(\mathbf{T}^*)\| = \|\text{GS}(\mathbf{T})\|$. Hence, by Corollary 1, with overwhelming probability,

$$\tilde{bl} \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right) \leq \|\text{GS}(\mathbf{T}^*)\| = \|\text{GS}(\mathbf{T})\| \leq \|\mathbf{T}\| \leq 2(s\sqrt{m})^d, \quad (42)$$

where the last inequality follows from Theorem 1, whose hypotheses hold by Table 1, choosing the implicit constant in $m = \Omega((k+n) \log q)$ sufficiently large and noting that $n \geq \lambda$.

By Lemma 1, there exists a negligible function ϵ such that

$$\eta_\epsilon \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right) \leq \omega \left(\sqrt{\log \text{rank}} \right) \tilde{bl} \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right) \leq 2(s\sqrt{m})^d \omega \left(\sqrt{\log \text{rank}} \right).$$

Moreover,

$$\begin{aligned} \|\bar{\boldsymbol{\delta}}\| &\leq \|\boldsymbol{\delta}\| \\ &\leq \|(\mathbf{x}_1^*)^{\otimes d}\| + \|(\mathbf{x}_0^*)^{\otimes d}\| \\ &= \|\mathbf{x}_1^*\|^d + \|\mathbf{x}_0^*\|^d \\ &\leq 2(\sqrt{k}\beta_x)^d. \end{aligned}$$

Combining these bounds with Eq. (41), we obtain

$$\eta_\epsilon \left(\ker(\mathbf{Q}^T \otimes \mathbf{I}_{\text{rank}}) \cap \left(\Lambda_q^\perp(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})) \right)^{k^d} \right) \leq 4(s\sqrt{m})^d (\sqrt{k}\beta_x)^d \omega \left(\sqrt{\log \text{rank}} \right) \leq \sigma.$$

For the last inequality, we use $\log \text{rank} = O(\log m)$ and $\sqrt{\omega(\log m)} = \omega(\sqrt{\log m})$, together with the choice of σ in Table 1. This proves Item 4. $\square$

Lemma 17. *For parameters par as in Table 1, we have $\text{Hyb}_{0,3} \equiv \text{Hyb}_{1,3}$.*

Proof. First suppose that $\delta = (\mathbf{x}_1^*)^{\otimes d} - (\mathbf{x}_0^*)^{\otimes d} = \mathbf{0}$. By definition, it follows that $\mathbf{V}_1^{(0)} = \mathbf{V}_1^{(1)}$. Therefore, $\text{Hyb}_{0,2} \equiv \text{Hyb}_{1,2}$. By the definition of $\text{Hyb}_{b,3}$ in the case $\delta = \mathbf{0}$, it follows that

$$\text{Hyb}_{0,3} \equiv \text{Hyb}_{1,3}.$$

Suppose now that $\delta \neq \mathbf{0}$. Since the columns of $\mathbf{Q}$ form a $\mathbb{Z}$ -basis of

$$H = \left\{ \mathbf{y} \in \mathbb{Z}^{k^d} : \delta^T \mathbf{y} = 0 \right\},$$

we have

$$\delta^T \mathbf{Q} = \mathbf{0}.$$

By the definition of $\text{Img}^{\otimes d}$, for every $b \in \{0, 1\}$,

$$\text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}_b^*)^T \end{bmatrix} \right) = \begin{bmatrix} \mathbf{0}_{((k+n+1)^d - 1) \times k^d} \\ ((\mathbf{x}_b^*)^{\otimes d})^T \\ \mathbf{0}_{[[k+n+1]]_{d-1}^1 \times k^d} \end{bmatrix}.$$

Indeed, the only nonzero row of

$$\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}_b^*)^T \end{bmatrix}^{\otimes d}$$

is the last row of its leading $(k+n+1)^d \times k^d$ block, and this row is $((\mathbf{x}_b^*)^{\otimes d})^T$. All the subsequent blocks in the definition of $\text{Img}^{\otimes d}$ are zero.

The first summand in the definition of $\mathbf{V}_1^{(b)}$ is independent of b . Therefore,

$$\begin{aligned} \mathbf{V}_1^{(1)} \mathbf{Q} - \mathbf{V}_1^{(0)} \mathbf{Q} &= \left(\text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}_1^*)^T \end{bmatrix} \right) - \text{Img}^{\otimes d} \left(\begin{bmatrix} \mathbf{0}_{k \times k} \\ \mathbf{0}_{n \times k} \\ (\mathbf{x}_0^*)^T \end{bmatrix} \right) \right) \mathbf{Q} \\ &= \left(\begin{bmatrix} \mathbf{0}_{((k+n+1)^d - 1) \times k^d} \\ ((\mathbf{x}_1^*)^{\otimes d})^T \\ \mathbf{0}_{[[k+n+1]]_{d-1}^1 \times k^d} \end{bmatrix} - \begin{bmatrix} \mathbf{0}_{((k+n+1)^d - 1) \times k^d} \\ ((\mathbf{x}_0^*)^{\otimes d})^T \\ \mathbf{0}_{[[k+n+1]]_{d-1}^1 \times k^d} \end{bmatrix} \right) \mathbf{Q} \\ &= \begin{bmatrix} \mathbf{0}_{((k+n+1)^d - 1) \times (k^d - 1)} \\ \delta^T \mathbf{Q} \\ \mathbf{0}_{[[k+n+1]]_{d-1}^1 \times (k^d - 1)} \end{bmatrix} \\ &= \mathbf{0}_{[[k+n+1]]_d^1 \times (k^d - 1)}. \end{aligned}$$

We have thus shown that

$$\mathbf{V}_1^{(0)} \mathbf{Q} = \mathbf{V}_1^{(1)} \mathbf{Q}. \quad (43)$$

The matrices $\mathbf{B}$ and $\tilde{\mathbf{B}}$ have the same distribution in $\text{Hyb}_{0,3}$ and $\text{Hyb}_{1,3}$. Conditioned on any fixed choice of these matrices, Eq. (43) implies

$$\mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(0)} \mathbf{Q}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_\mathbf{Q}} = \mathcal{D}_{\Lambda_q^{\mathbf{V}_1^{(1)} \mathbf{Q}}(\mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}})), \mathbf{S}_\mathbf{Q}}.$$

Thus, the joint distribution of

$$(\mathbf{B}, \tilde{\mathbf{B}}, \mathbf{K}, \tilde{\mathbf{K}}, \mathbf{G})$$

is identical in the two hybrids.

The ciphertexts are also identically distributed: their generation in $\text{Hyb}_{b,3}$ is exactly the same as in $\text{Hyb}_{b,2}$ and does not depend on b .

It remains to compare the key-generation oracles. Both hybrids use the same validity condition

$$((\mathbf{x}_0^*)^{\otimes d})^T \mathbf{y} = ((\mathbf{x}_1^*)^{\otimes d})^T \mathbf{y}$$

and the same matrix $\mathbf{Q}$. Therefore, for every valid query $\mathbf{y}$, both hybrids compute the same unique $\mathbf{a} \in \mathbb{Z}^{k^d-1}$ satisfying $\mathbf{y} = \mathbf{Q}\mathbf{a}$ and return

$$(\mathbf{K}\mathbf{a}, \tilde{\mathbf{K}}\mathbf{a}).$$

Hence, the two experiments are identical, i.e., $\text{Hyb}_{0,3} \equiv \text{Hyb}_{1,3}$ as claimed. $\square$

Combining Lemmas 13 to 15 and 17, we obtain

$$\text{Hyb}_{0,0} \approx_s \text{Hyb}_{0,1} \approx_c \text{Hyb}_{0,2} \approx_s \text{Hyb}_{0,3} \equiv \text{Hyb}_{1,3} \approx_s \text{Hyb}_{1,2} \approx_c \text{Hyb}_{1,1} \approx_s \text{Hyb}_{1,0}.$$

Therefore, under the $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption,

$$\text{INDEX}_{\text{FE}, \mathcal{A}}^0 = \text{Hyb}_{0,0} \approx_c \text{Hyb}_{1,0} = \text{INDEX}_{\text{FE}, \mathcal{A}}^1,$$

which proves the theorem. $\square$

D Supplementary Material for the Cryptanalysis

This appendix collects the material supporting Section 7.

D.1 Supplementary Material for Strategy 3

Recall that Strategy 3 (Section 7.1) distinguishes by approximating the mean of $g(\mathbf{x}) = \langle \text{err}_0(\mathbf{x}, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}), \mathbf{y} \rangle^2$ over the randomness $(\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}})$ for a fixed function $\mathbf{y}$: against the FE scheme by comparing $\mathbb{E}[g(\mathbf{x}_0^*)]$ with $\mathbb{E}[g(\mathbf{x}_1^*)]$ for two messages, and against $\text{sk-}k\text{-ALWE}^{\otimes d}$ by comparing $\mathbb{E}[g(\mathbf{x}^*)]$ with the corresponding mean for err^* at the same $\mathbf{x}^*$. The following Lemma 18 provides the explicit form of $\mathbb{E}[g(\mathbf{x})]$ in the $d = 2$ case.

Lemma 18. *Let $d = 2$, and fix $\mathbf{K}, \tilde{\mathbf{K}}, \mathbf{F}$ and $\mathbf{y} \in \mathbb{Z}^{k^2}$. For $\mathbf{s} \leftarrow \$\mathcal{D}_{\mathbb{Z}, \chi_s}^n$, $\mathbf{e} \leftarrow \$\mathcal{D}_{\mathbb{Z}, \chi_e}^m$, $\tilde{\mathbf{e}} \leftarrow \$\mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}^{\tilde{m}}$, we have*

$$\begin{aligned} & \mathbb{E} \left[\left(\langle \mathbf{e}^{\otimes d}, \mathbf{K}\mathbf{y} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F}\mathbf{y} \right\rangle \right)^2 \right] \\ &= \left(\mu_e^{(4)} - 3\sigma_e^4 \right) \cdot \sum_{i=1}^m \mathbf{m}_{i,i}^2 + \sigma_e^4 \cdot ((\text{tr}(\mathbf{M}))^2 + \text{tr}(\mathbf{M}\mathbf{M}^T) + \text{tr}(\mathbf{M}^2)) \\ & \quad + \sigma_e^2 \cdot \|(\mathbf{F}_1 + \mathbf{F}_3)\mathbf{x}\|^2 + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}}\mathbf{y}\|^2 + \sigma_e^2 \cdot \sigma_s^2 \cdot \langle \mathbf{F}_2 + \mathbf{F}_4, \mathbf{F}_2 + \mathbf{F}_4 \rangle_F \end{aligned}$$

where

$$\begin{aligned} \sigma_e^2 &= \mathbb{E}_{e \leftarrow \$\mathcal{D}_{\mathbb{Z}, \chi_e}} [e^2] \approx \chi_e^2/2\pi, & \tilde{\sigma}_e^2 &= \mathbb{E}_{\tilde{e} \leftarrow \$\mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}} [\tilde{e}^2] \approx \tilde{\chi}_e^2/2\pi, & \sigma_s^2 &= \mathbb{E}_{s \leftarrow \$\mathcal{D}_{\mathbb{Z}, \chi_s}} [s^2] \approx \chi_s^2/2\pi, \\ \mu_e^{(4)} &= \mathbb{E}_{e \leftarrow \$\mathcal{D}_{\mathbb{Z}, \chi_e}} [e^4] \approx 3\sigma_e^4, \end{aligned}$$

the matrix $\mathbf{M} = (\mathbf{m}_{i,j})_{i,j}$ is a reordering of $\mathbf{K}\mathbf{y}$ such that $\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle = \mathbf{e}^T \mathbf{M} \mathbf{e}$ (the value above does not depend on which such $\mathbf{M}$ is taken: the identity determines the symmetric part $\mathbf{S} = (\mathbf{M} + \mathbf{M}^T)/2$, and the value depends on $\mathbf{M}$ only through $\mathbf{S}$, since $\text{tr}(\mathbf{M}) = \text{tr}(\mathbf{S})$, $\mathbf{m}_{i,i} = \mathbf{s}_{i,i}$ and $\text{tr}(\mathbf{M}\mathbf{M}^T) + \text{tr}(\mathbf{M}^2) = 2\text{tr}(\mathbf{S}^2)$), and the

matrices $\mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3, \mathbf{F}_4$ stem by reordering and partitioning the coordinates of $\mathbf{F}\mathbf{y} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \mathbf{f}_3 \\ \mathbf{f}_4 \end{bmatrix}$ such that

$$\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle = \mathbf{e}^T \mathbf{F}_1 \mathbf{x},$$

$$\begin{aligned}
\langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle &= \mathbf{e}^\top \mathbf{F}_2 \mathbf{s}, \\
\langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle &= \mathbf{e}^\top \mathbf{F}_3 \mathbf{x}, \\
\langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle &= \mathbf{e}^\top \mathbf{F}_4 \mathbf{s},
\end{aligned}$$

and $\langle (a_{i,j})_{i,j}, (b_{i,j})_{i,j} \rangle_F = \sum_{i,j} a_{i,j} b_{i,j}$ denotes the entry-wise scalar product for matrices.

Proof. Let us first compute for $d = 2$

$$\begin{aligned}
& \left(\langle \mathbf{e}^{\otimes d}, \mathbf{K}\mathbf{y} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}, \mathbf{F}\mathbf{y} \right\rangle \right)^2 \\
&= \left(\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle - \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle \right)^2 \\
&= \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle^2 + 2 \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle - 2 \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle \\
&\quad + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle^2 - 2 \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle \cdot \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle + \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle^2.
\end{aligned}$$

Since the variables $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ are independent of each other and zero-centered, the means of $\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle$

and $\langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle \cdot \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle$ are zero. Let $\mathbf{M} = (\mathbf{m}_{i,j})_{i,j} \in \mathbb{Z}^{m \times m}$ be a reordering of $\mathbf{K}\mathbf{y} \in \mathbb{Z}^{m^2}$ as a matrix

such that $\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle = \mathbf{e}^\top \mathbf{M} \mathbf{e}$. The mean of $\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle^2$ is given by

$$\mathbb{E} [\langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle^2] = \mathbb{E} [(\mathbf{e}^\top \mathbf{M} \mathbf{e})^2] = \sum_{i,j,u,v=1}^m \mathbf{m}_{i,j} \cdot \mathbf{m}_{u,v} \cdot \mathbb{E}_{e_1, \dots, e_m \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}} [e_i e_j e_u e_v].$$

Since the coordinates of $\mathbf{e}$ are independent, identically distributed and zero-centered, $\mathbb{E} [e_i e_j e_u e_v]$ vanishes unless the four indices split into two coinciding pairs. It equals $\mu_e^{(4)}$ if $i = j = u = v$, and σ_e^4 in each of the three remaining cases $i = j \neq u = v$, $i = u \neq j = v$ and $i = v \neq j = u$. Collecting the four cases, we get

$$\begin{aligned}
\mathbb{E} [(\mathbf{e}^\top \mathbf{M} \mathbf{e})^2] &= \mu_e^{(4)} \sum_{i=1}^m \mathbf{m}_{i,i}^2 + \sigma_e^4 \left(\sum_{i \neq u} \mathbf{m}_{i,i} \mathbf{m}_{u,u} + \sum_{i \neq j} \mathbf{m}_{i,j}^2 + \sum_{i \neq j} \mathbf{m}_{i,j} \mathbf{m}_{j,i} \right) \\
&= \left(\mu_e^{(4)} - 3\sigma_e^4 \right) \sum_{i=1}^m \mathbf{m}_{i,i}^2 + \sigma_e^4 ((\text{tr}(\mathbf{M}))^2 + \text{tr}(\mathbf{M}\mathbf{M}^\top) + \text{tr}(\mathbf{M}^2)) =: \theta_{\mathbf{M}},
\end{aligned}$$

where the second step subtracts the diagonal from each of the three full sums, using $\sum_{i,u} \mathbf{m}_{i,i} \mathbf{m}_{u,u} = (\text{tr}(\mathbf{M}))^2$, $\sum_{i,j} \mathbf{m}_{i,j}^2 = \text{tr}(\mathbf{M}\mathbf{M}^\top)$ and $\sum_{i,j} \mathbf{m}_{i,j} \mathbf{m}_{j,i} = \text{tr}(\mathbf{M}^2)$. We continue with the terms that are linear in $\mathbf{e}$. For any fixed $\mathbf{a} = (a_1, \dots, a_m)$, we have

$$\mathbb{E} [\langle \mathbf{e}, \mathbf{a} \rangle^2] = \sum_{i,j=1}^m a_i a_j \cdot \mathbb{E}_{e_1, \dots, e_m \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}} [e_i e_j] = \sum_{i=1}^m a_i^2 \cdot \mathbb{E}_{e_i \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}} [e_i^2] = \sigma_e^2 \cdot \|\mathbf{a}\|^2,$$

as $\mathbb{E}_{e_i \leftarrow \mathcal{D}_{\mathbb{Z}, \chi_e}} [e_i^2] = \sigma_e^2$ and the means of the cross-terms $e_i \cdot e_j$ vanish if $i \neq j$. The same computation, with $\tilde{\mathbf{e}} \leftarrow \mathcal{D}_{\mathbb{Z}, \tilde{\chi}_e}^{\tilde{m}}$ in place of $\mathbf{e}$ and $\mathbf{a} = \tilde{\mathbf{K}}\mathbf{y}$, gives $\mathbb{E} \left[\langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}}\mathbf{y} \rangle^2 \right] = \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}}\mathbf{y}\|^2$.

Now, decompose $\mathbf{F}\mathbf{y} = \begin{bmatrix} \mathbf{f}_1 \\ \mathbf{f}_2 \\ \mathbf{f}_3 \\ \mathbf{f}_4 \end{bmatrix}$ such that

$$\left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle = \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle + \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle + \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle.$$

The mean of

$$\begin{aligned} & \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle \\ &= \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle + \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle + \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + \langle \mathbf{e}^{\otimes 2}, \mathbf{K}\mathbf{y} \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle \end{aligned}$$

does vanish as $\mathbf{e}$ appears in degree 3 in each summand and is zero-centered (also note that $\mathbf{e}, \mathbf{s}$ are independent and $\mathbf{x}$ is fixed).

We are left with determining the mean of

$$\begin{aligned} & \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F}\mathbf{y} \right\rangle^2 \\ &= (\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle + \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle + \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle)^2 \\ &= \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle^2 + 2 \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle + 2 \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + 2 \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle \\ &\quad + \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle^2 + 2 \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + 2 \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle \\ &\quad + \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle^2 + 2 \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle + \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle^2. \end{aligned}$$

By usual arguments, the means of the cross-terms $\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle$, $\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle$, $\langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle$, and $\langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle$ vanish. For the remaining terms, organize $\mathbf{f}_1, \dots, \mathbf{f}_4$ into matrices $\mathbf{F}_1, \dots, \mathbf{F}_4$ such that $\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle = \mathbf{e}^T \mathbf{F}_1 \mathbf{x}$, $\langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle = \mathbf{e}^T \mathbf{F}_2 \mathbf{s}$, $\langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle = \mathbf{e}^T \mathbf{F}_3 \mathbf{x}$ and $\langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle = \mathbf{e}^T \mathbf{F}_4 \mathbf{s}$. For $\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle^2$, the identity above with $\mathbf{a} = \mathbf{F}_1 \mathbf{x}$ gives

$$\mathbb{E} \left[\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle^2 \right] = \mathbb{E} \left[(\mathbf{e}^T \mathbf{F}_1 \mathbf{x})^2 \right] = \sigma_e^2 \cdot \|\mathbf{F}_1 \mathbf{x}\|^2,$$

and analogously we get $\mathbb{E} \left[\langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle^2 \right] = \sigma_e^2 \cdot \|\mathbf{F}_3 \mathbf{x}\|^2$. For the cross-term $\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle = (\mathbf{e}^T \mathbf{F}_1 \mathbf{x}) \cdot (\mathbf{e}^T \mathbf{F}_3 \mathbf{x})$, set $\mathbf{a} = \mathbf{F}_1 \mathbf{x}$ and $\mathbf{b} = \mathbf{F}_3 \mathbf{x}$. We have

$$\begin{aligned} & \mathbb{E} \left[(\mathbf{e}^T \mathbf{F}_1 \mathbf{x}) \cdot (\mathbf{e}^T \mathbf{F}_3 \mathbf{x}) \right] = \mathbb{E} \left[\mathbf{e}^T \mathbf{a} \cdot \mathbf{e}^T \mathbf{b} \right] = \mathbb{E} \left[\sum_{i,j} e_i a_i \cdot e_j b_j \right] = \sum_{i,j} \mathbb{E} [e_i a_i \cdot e_j b_j] = \sum_i \mathbb{E} [e_i^2 \cdot a_i b_i] \\ &= \sigma_e^2 \cdot \langle \mathbf{a}, \mathbf{b} \rangle = \sigma_e^2 \cdot \langle \mathbf{F}_1 \mathbf{x}, \mathbf{F}_3 \mathbf{x} \rangle. \end{aligned}$$

For $\langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle = (\mathbf{e}^T \mathbf{F}_2 \mathbf{s})(\mathbf{e}^T \mathbf{F}_4 \mathbf{s})$, we set $(f_{2,i,j}) = \mathbf{F}_2$, $(f_{4,i,j}) = \mathbf{F}_4$ and get

$$\begin{aligned} \mathbb{E}[(\mathbf{e}^T \mathbf{F}_2 \mathbf{s})(\mathbf{e}^T \mathbf{F}_4 \mathbf{s})] &= \mathbb{E} \left[\left(\sum_{i,j} e_i f_{2,i,j} s_j \right) \left(\sum_{u,v} e_u f_{4,u,v} s_v \right) \right] = \sum_{i,j} \mathbb{E}[e_i f_{2,i,j} s_j \cdot e_i f_{4,i,j} s_j] \\ &= \sigma_e^2 \sigma_s^2 \sum_{i,j} f_{2,i,j} \cdot f_{4,i,j} = \sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_2, \mathbf{F}_4 \rangle_F, \end{aligned}$$

since $\langle \mathbf{F}_2, \mathbf{F}_4 \rangle_F = \sum_{i,j} f_{2,i,j} \cdot f_{4,i,j}$. Analogously, we get $\mathbb{E}[\langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle^2] = \sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_2, \mathbf{F}_2 \rangle_F = \sigma_e^2 \sigma_s^2 \cdot \|\mathbf{f}_2\|^2$ and $\mathbb{E}[\langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle^2] = \sigma_e^2 \sigma_s^2 \cdot \|\mathbf{f}_4\|^2$. In conclusion, we get

$$\begin{aligned} &\mathbb{E} \left[\left(\langle \mathbf{e}^{\otimes d}, \mathbf{K} \mathbf{y} \rangle + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \mathbf{y} \rangle - \left\langle \left\{ \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{e} \right\}_{0 < \#_{\mathbf{e}} < d}^{\otimes d}, \mathbf{F} \mathbf{y} \right\rangle \right)^2 \right] \\ &= \mathbb{E} \left[\langle \mathbf{e}^{\otimes 2}, \mathbf{K} \mathbf{y} \rangle^2 + \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \mathbf{y} \rangle^2 + \left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F} \mathbf{y} \right\rangle^2 \right] \\ &= \theta_M + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}} \mathbf{y}\|^2 + \mathbb{E} \left[\left\langle \begin{bmatrix} \mathbf{x} \otimes \mathbf{e} \\ \mathbf{s} \otimes \mathbf{e} \\ \mathbf{e} \otimes \mathbf{x} \\ \mathbf{e} \otimes \mathbf{s} \end{bmatrix}, \mathbf{F} \mathbf{y} \right\rangle^2 \right] \\ &= \theta_M + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}} \mathbf{y}\|^2 + \mathbb{E} \left[(\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle + \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle + \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle)^2 \right] \\ &= \theta_M + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}} \mathbf{y}\|^2 + \mathbb{E} \left[\langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle^2 + 2 \langle \mathbf{x} \otimes \mathbf{e}, \mathbf{f}_1 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle + \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle^2 \right. \\ &\quad \left. + 2 \langle \mathbf{s} \otimes \mathbf{e}, \mathbf{f}_2 \rangle \cdot \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle + \langle \mathbf{e} \otimes \mathbf{x}, \mathbf{f}_3 \rangle^2 + \langle \mathbf{e} \otimes \mathbf{s}, \mathbf{f}_4 \rangle^2 \right] \\ &= \theta_M + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}} \mathbf{y}\|^2 + \sigma_e^2 \cdot \|\mathbf{F}_1 \mathbf{x}\|^2 + 2\sigma_e^2 \cdot \langle \mathbf{F}_1 \mathbf{x}, \mathbf{F}_3 \mathbf{x} \rangle + \sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_2, \mathbf{F}_2 \rangle_F \\ &\quad + 2\sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_2, \mathbf{F}_4 \rangle_F + \sigma_e^2 \cdot \|\mathbf{F}_3 \mathbf{x}\|^2 + \sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_4, \mathbf{F}_4 \rangle_F \\ &= \theta_M + \sigma_e^2 \cdot \|(\mathbf{F}_1 + \mathbf{F}_3) \mathbf{x}\|^2 + \tilde{\sigma}_e^2 \cdot \|\tilde{\mathbf{K}} \mathbf{y}\|^2 + \sigma_e^2 \sigma_s^2 \cdot \langle \mathbf{F}_2 + \mathbf{F}_4, \mathbf{F}_2 + \mathbf{F}_4 \rangle_F. \end{aligned}$$

Hence, the claim follows. $\square$

Remark 3 (Why not rule out Strategy 3 by statistical arguments?). We observe that, for any fixed $\mathbf{x}^*$, the image spaces $\{\text{err}_0(\mathbf{x}^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}) \mid \mathbf{s} \in \mathbb{Z}^n, \mathbf{e} \in \mathbb{Z}^m, \tilde{\mathbf{e}} \in \mathbb{Z}^{\tilde{m}}\}$ and $\{\text{err}^*(\mathbf{x}^*, \mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}) \mid \mathbf{s} \in \mathbb{Z}^n, \mathbf{e} \in \mathbb{Z}^m, \tilde{\mathbf{e}} \in \mathbb{Z}^{\tilde{m}}\}$ are (discrete subsets of) $(n+m+\tilde{m})$ -dimensional manifolds in the space $\mathbb{R}^{k^d}$. These spaces overlap, whenever there exist $\mathbf{s}_0, \mathbf{s}^*, \mathbf{e}_0, \mathbf{e}^*, \tilde{\mathbf{e}}_0, \tilde{\mathbf{e}}^*$ such that $\text{err}_0(\mathbf{x}^*, \mathbf{s}_0, \mathbf{e}_0, \tilde{\mathbf{e}}_0) = \text{err}^*(\mathbf{x}^*, \mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$. Now, a minimal $\tilde{m}$ -dimensional intersection space is given when $\mathbf{e}_0 = \mathbf{e}^* = 0$, as we have in this case $\text{err}_0(\mathbf{x}^*, \mathbf{s}_0, \mathbf{e}_0, \tilde{\mathbf{e}}_0) = \langle \tilde{\mathbf{e}}, \tilde{\mathbf{K}} \rangle = \text{err}^*(\mathbf{x}^*, \mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$. However, the probability that $\mathbf{e} = 0$ in the actual experiments is negligible. If we restrict ourselves to $\mathbf{e} \neq 0$, it seems plausible to assume that both image spaces do not intersect. Indeed, since $\mathbf{K}$, $\tilde{\mathbf{K}}$ and $\mathbf{F}$ are tall and have full rank, we would (intuitively) need $\mathbf{e}_0 = \mathbf{e}^*$, $\tilde{\mathbf{e}}_0 = \tilde{\mathbf{e}}^*$ and $\mathbf{s}_0 = \mathbf{s}^*$ for an intersection. This implies a large statistical distance between the distributions of $\mathbf{f}^*$ in Exp^0 and Exp^1 and that we cannot hope to rule out this attack strategy purely based on statistical arguments (like noise-flooding), even if the randomness $(\mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$ are subexponentially large.

D.2 Strategy 4: Algebraic Methods

The idea of Strategy 4 is to check if the challenge ciphertext $(\mathbf{c}, \tilde{\mathbf{c}})$ has been correctly generated in Exp^b . Note that if $b = 0$, the distribution of $(\mathbf{c}, \tilde{\mathbf{c}})$ follows the same distribution as the auxiliary ciphertexts

$(\mathbf{d}_j, \tilde{\mathbf{d}}_j)$, namely $\text{LWESamp}^{\otimes d}((\mathbf{x}, \mathbf{s}), \mathbf{A}, \tilde{\mathbf{A}})$. However, if $b = 1$, the candidate ciphertext $(\mathbf{c}, \tilde{\mathbf{c}})$ is sampled by $\text{LWESamp}^{\otimes d}((\mathbf{0}_k, \mathbf{s}, 1), \mathbf{B}, \tilde{\mathbf{B}})$. Therefore, this strategy tries to decide b in the $k\text{-ALWE}^{\otimes d}$ experiment by searching for $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ such that

$$(\mathbf{c}, \tilde{\mathbf{c}})^T = (\langle (\mathbf{x}^*, \mathbf{s}), \mathbf{A} \rangle + p \cdot \mathbf{e}^T, \langle \{(\mathbf{x}^*, \mathbf{s}), p \cdot \mathbf{e}\}_{\#_{p \cdot \mathbf{e}} < d}, \tilde{\mathbf{A}} \rangle + p^d \cdot \tilde{\mathbf{e}}^T) \pmod{q}.$$

If it can find such randomness $\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}}$ that correctly opens $(\mathbf{c}, \tilde{\mathbf{c}})^T$ to $\mathbf{x}^*$, then it outputs $b = 0$. Otherwise, it claims $b = 1$. However, instead of considering the relevant equation system induced by $(\mathbf{c}, \tilde{\mathbf{c}})$ over $\mathbb{Z}_q$, it considers the integral equation system $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$ induced by the left-over noise. Note, for doing so, the attacker needs to know the Gaussian matrix $\mathbf{F}$, which it can deduce from $\mathbf{K}, \tilde{\mathbf{K}}$ and the matrices $\mathbf{A}, \tilde{\mathbf{A}}$. This strategy therefore targets $k\text{-ALWE}^{\otimes d}$, in which they are given to the adversary. Now, the equation system $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$ has a multitude of interesting properties:

- It is over-determined, with k^d equations but only $n + m + \tilde{m} = o(k^d)$ variables.
- The coefficients of err_b are sampled from some (correlated) discrete Gaussian distributions with parameter σ hence, implying sufficient genericity of the system for large enough σ .
- Each polynomial in err_b has $\binom{n+m+d}{d} - \binom{n+d}{d} + \tilde{m}$ distinct monomials.
- When $b = 0$, the system admits at least one solution $(\mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$. Assuming that the precautions against Strategy 3 are taken, the solution has a subexponentially large norm. Since the system is highly over-determined, when $b = 0$, the system likely has exactly one solution $(\mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$ but no solution if $b = 1$.

Since the number of monomials is greater than the number of equations whenever $n, m, \tilde{m}$ meet the bound $n + m + \tilde{m} \geq m \geq (k + n) \log T$ of Table 1, the system cannot be solved by linearisation. Due to the subexponentially large norm of the solution $(\mathbf{s}^*, \mathbf{e}^*, \tilde{\mathbf{e}}^*)$, lattice-based methods are also unlikely to work.

D.2.1 Macaulay Matrix Methods. Algebraic solving algorithms, however, can solve polynomial equation systems of this kind. For example, Macaulay matrix-based methods like eXtended Linearisation (XL) [CKPS00] or Mutant-XL [DBM⁺08] compute a (non-trivial) Groebner basis of the equation system if it is satisfiable (and the trivial basis $\{1\}$ if it is unsatisfiable). To estimate the time complexity of the Mutant-XL algorithm, we can use the fact that its solving degree is upper-bounded by $1 +$ the degree of regularity of the equation system [Sal25]. By assuming the equation system to be *semi-regular* [Par10, BFS03, Bar04], we can bound the degree of regularity concretely by computing the Hilbert-series of the top terms of the involved polynomials. As far as it concerns this discussion, it suffices to know that the number of equations is k^d and the degree of the system is d . Since the number of variables is large, $n + m + \tilde{m} \geq m \geq (k + n) \log T$ by Table 1, the degree of regularity¹³ D of a semi-regular system in this parameter regime is at least $\log T$ by Lemma 19. Assuming that the time complexity of the Mutant-XL algorithm is at least $\binom{n+m+\tilde{m}+D+1}{D+1}^2$ [Sal25], where D is the degree of regularity, this implies a time complexity for Macaulay matrix-based solving algorithms of at least $((n + m + \tilde{m})/(D + 1))^{2(D+1)} \geq T^2$.

Lemma 19. *Let $k, N, D, d \geq 1$ be such that $N > D \cdot k$ and $D \geq d$. An equation system of k^d homogeneous equations of degree d over N variables cannot have a degree of regularity equal or less than D .*

Proof. Assume that the opposite holds. Then, the space of all degree- $(D - d)$ monomials multiplied with each of the k^d homogeneous polynomials must contain all monomials of degree D . This implies

$$k^d \cdot \binom{N + D - d - 1}{D - d} \geq \binom{N + D - 1}{D}.$$

This inequality however implies

$$k^d \geq \frac{N + D - 1}{D} \cdots \frac{N + D - d}{D - d + 1}.$$

¹³ The *degree of regularity* D of polynomials $f_1, \dots, f_m$ is defined as the infimum D such that the graded ideal $(f_1^{\text{top}}, \dots, f_m^{\text{top}})$ contains all monomials of degree D , where f_i^{top} denotes the highest-degree terms of f_i .

As $\frac{N+D-1}{D}$ is the smallest factor in the above product, we get

$$k \geq \frac{N+D-1}{D} \geq \frac{N}{D}.$$

This directly contradicts the required inequality of the claim. $\square$

However, there is a gap in the above argument. Concretely, while the system at hand has a lot of random coefficients it is missing some monomials. Hence, it is in general *not* semi-regular. In this case, it is hard to bound its degree of regularity, and the time complexity for Macaulay matrix-based solving algorithms might be arbitrarily high (or low). To counteract, we present an alternative algebraic attack algorithm in Section D.2.2 based on polynomial tests resp. annihilators.

D.2.2 Polynomial Tests on k -ALWE. We discuss here an alternative to Strategy 4 above. While Strategy 4 tries to solve for $(\mathbf{s}, \mathbf{e}, \tilde{\mathbf{e}})$ such that $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot) = \mathbf{f}^*$, this alternative strategy aims lower and simply tries to discern the distribution $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ from $\text{err}^*(\mathbf{x}^*, \cdot, \cdot, \cdot)$. For this end, we interpret $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ as a polynomial PRG mapping seeds of length $n + m + \tilde{m}$ to somewhat pseudorandom outputs of length k^d . Following [Üna23], we can guarantee the existence of a non-zero test polynomial $h \in \mathbb{Z}[y_1, \dots, y_{k^d}]$ that vanishes on all images of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ (hence, it distinguishes the outputs of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ from random vectors with high advantage). Because of compactness, we must have $n + m + \tilde{m} \in o(k^d)$. Therefore, [Üna23] upper bounds the degree of h by $O((n + m + \tilde{m})^d / k^d)^{1/(d-1)}$. We assume that this upper bound is asymptotically tight. Under this assumption, the degree of h lies in $\Theta((\log T)^{d/(d-1)})$, since $n + m + \tilde{m} \geq m \geq (n + k) \log(T)$ by the bound of Table 1. Making the additional assumption that h is a dense polynomial, we get that evaluating h takes time at least $\binom{k^d}{\deg h} = 2^{\Omega(\deg h \cdot d \log k)} \geq T$.

There is again a caveat here, as h needs to discern images of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ from images of $\text{err}^*(\mathbf{x}^*, \cdot, \cdot, \cdot)$. While h is guaranteed to vanish on the images of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$, we cannot exclude the case that h also vanishes on images of $\text{err}^*(\mathbf{x}^*, \cdot, \cdot, \cdot)$ with overwhelming probability. However, if the image spaces of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ and $\text{err}^*(\mathbf{x}^*, \cdot, \cdot, \cdot)$ significantly differ, it seems unlikely that a randomly chosen annihilator h of $\text{err}_b(\mathbf{x}^*, \cdot, \cdot, \cdot)$ does with overwhelming probability vanish on $\text{err}^*(\mathbf{x}^*, \cdot, \cdot, \cdot)$.

D.3 Difficulty of Reductions from Standard Assumptions

A natural strategy to obtain a reduction from a “more standard” assumption to $\text{sk-}k\text{-ALWE}^{\otimes d}$ is to generalise the reduction from $k\text{-LWE}$ to $k\text{-ALWE}$ in Theorem 5. A core ingredient of that reduction is a statistical “reverse-sampling” lemma (Lemma 7) originated from [BF11], which roughly states that, for any image $\mathbf{Y}$, the ordering of sampling uniform $\mathbf{A}$ and Gaussian $\mathbf{U}$ subject to $\mathbf{A} \cdot \mathbf{U} = \mathbf{Y} \bmod q$ does not matter – they are statistically close to each other. Corollary 2 states an even more fine-grained result – one can first sample part of $\mathbf{A}$, then $\mathbf{U}$, then the remaining part of $\mathbf{A}$. Generalising the proof of Theorem 5 to $d > 1$ likely requires a degree- d analogue of the reverse-sampling lemma, where the matrix $\mathbf{A}$ has a degree- d tensor structure (hence not uniformly random), which does not exist in the literature.

We showcase a simple example highlighting the difficulties which arise due to the lack of such an analogue. For concreteness, we consider in Fig. 7 a hybrid instance generator $\text{InstSamp}_{0.5}^{\otimes d}$ for the experiment $\text{Exp}_{d, \text{par}, \mathcal{A}}^{0.5}$ defined in Fig. 3. The only difference between $\text{Exp}_{d, \text{par}, \mathcal{A}}^0$ and $\text{Exp}_{d, \text{par}, \mathcal{A}}^{0.5}$ is that, in the latter, $(\mathbf{K}, \tilde{\mathbf{K}})$ satisfy additional random non-linear constraints whose coefficients are not given to the adversary. Intuitively, one would expect that $\text{Exp}_{d, \text{par}, \mathcal{A}}^0$ and $\text{Exp}_{d, \text{par}, \mathcal{A}}^{0.5}$ are statistically close. Indeed, if the constraints were linear then the reverse-sampling lemma would have implied the statistical closeness of these experiments immediately. However, without a degree- d analogue of the reverse-sampling lemma, we do not know how to formally prove this claim.

Another difficulty is that, even if we were bestowed a degree- d reverse-sampling lemma, the hypothetical analogue of Theorem 5 would only give a reduction from a “homogeneous” variant of $\text{sk-}k\text{-ALWE}^{\otimes d}$ which, although simpler in the sense that the assumption may no longer be interactive, is still a degree- d and

```

InstSamp $_{0.5}^{\otimes d}(\mathbf{x}_1, \dots, \mathbf{x}_Q, \mathbf{x}^*)$ 


---


 $\mathbf{s}^* \leftarrow \chi_s^n, \mathbf{B} \leftarrow \mathbb{Z}_q^{(k+n+1) \times m}, \tilde{\mathbf{B}} \leftarrow \mathbb{Z}_q^{[[k+n+1]]_d^1 \times \tilde{m}}$ 
for  $j \in [Q]$  do  $\mathbf{s}_j \leftarrow \chi_s^n, (\mathbf{d}_j, \tilde{\mathbf{d}}_j) \leftarrow \text{LWESamp}^{\otimes d} \left( \begin{bmatrix} \mathbf{x}_j \\ \mathbf{s}_j \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ 
 $(\mathbf{c}, \tilde{\mathbf{c}}) \leftarrow \text{LWESamp}^{\otimes d} \left( \begin{bmatrix} \mathbf{x}^* \\ \mathbf{s}^* \\ 0 \end{bmatrix}, \mathbf{B}, \tilde{\mathbf{B}} \right)$ 
 $\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} \leftarrow \mathbb{Z}_{\sigma}^{\mathcal{D}_{\mathbb{Z}, \sigma}^{m^d \times k^d} \times \mathcal{D}_{\mathbb{Z}, \sigma}^{\tilde{m} \times k^d} \times \mathcal{D}_{\mathbb{Z}, \sigma}^{[[k+n+1]]_{d-1}^1 \times k^d}} : \mathcal{E}^{\otimes d}(\mathbf{B}, \tilde{\mathbf{B}}) \cdot \begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \\ \mathbf{G} \end{bmatrix} = \text{Img}^{\otimes d} \left( \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{0}_k^{\top} \end{bmatrix} \right) \bmod q$ 
return  $((\mathbf{d}_j, \tilde{\mathbf{d}}_j)_{j \in [Q]}, \mathbf{c}, \tilde{\mathbf{c}}, \mathbf{K}, \tilde{\mathbf{K}})$ 

```

Fig. 7. Hybrid instance generator $\text{InstSamp}_{0.5}^{\otimes d}$ for $\text{Exp}_{d, \text{par}, \mathcal{A}}^{0.5}$ defined in Fig. 3.

thus non-standard variant of LWE. Fundamentally, we lack tools for analysing any non-linear variants of lattice assumptions, such as those introduced in [ACL⁺22, WW23, CLM23, CLW25, HJL25a], but not only $\text{sk-}k\text{-ALWE}^{\otimes d}$ and its hypothetical homogeneous variant.

From another point of view, we can see Strategy 3 as an intuitive separation between $\text{sk-}k\text{-ALWE}^{\otimes d}$ and standard lattice assumptions like LWE. Note that Strategy 3 has a noticeable advantage if the involved parameters of $\text{sk-}k\text{-ALWE}^{\otimes d}$ are polynomial. If there would be a “simple” reduction from LWE to $\text{sk-}k\text{-ALWE}^{\otimes d}$ for superpolynomial parameters that preserves the modulus q , then it should also work for polynomial parameters, leading to a contradiction as $\text{sk-}k\text{-ALWE}^{\otimes d}$ is plausibly broken for polynomial parameters, while LWE is not. Of course, a reduction from standard assumptions can involve techniques like noise flooding that cannot work over polynomial parameters. However, as pointed out in Remark 3, statistical smudging techniques can most likely not be applied to $\text{sk-}k\text{-ALWE}^{\otimes d}$.

E Comparison with Related Assumptions

We give here a brief overview of similar experimental lattice-based assumptions from which iO has been built, then compare them to $\text{sk-}k\text{-ALWE}^{\otimes d}$ and highlight the structural differences.

Overview. A number of works build iO from integer-valued [Agr19, AJL⁺19, JLMS19, JLS19, GJLS21] or block-local [LT17] PRGs. These are broken by statistical or sum-of-squares methods [BBKK18, BHJ⁺19, JKLS21]. Gay and Pass [GP21] base the security of their iO scheme on the conjectured 2-circular shielded randomness leakage resilience of GSW, and Wee and Wichs [WW21] on the homomorphic pseudorandom LWE samples conjecture for dual-GSW. Both assumptions are broken by [HJL21], which embeds malicious circuits into the respective challengers. Devadas et al. [DQV⁺21] propose an LWE sampling assumption whose noise matrices can be directly extracted via linear algebra [JLLS23]. A line of works [BDJ⁺25, AKY24b, AKY24a, VWW22] construct FE and (null-)iO from private-coin evasive LWE [Tsa22, VWW22]. Attacks on this assumption [BÜW24, HHY26, DJM⁺25, HJL25b, AMYY25] exploit malicious sampling algorithms that leak important data when supplied with short preimages. [HJL25a, CLW25] formalise the heuristic of [BDGM20] by introducing the circular random opening and equivocal LWE assumptions. Both state that FHE ciphertexts in the presence of HE ciphertexts encrypting the FHE secret key remain uncompromised, even when short unstructured preimages are given as openings of derived HE ciphertexts, where the former assumption additionally involves circular FHE and HE ciphertexts. To the best of our knowledge, neither assumption is broken. There are further paths to iO, for instance via multilinear maps [GGH13a, GGH⁺13b, GGH15] with

their own barrage of zeroizing and annihilation attacks [CGH⁺15, HJ16, MSZ16, CVW18], which are out of scope for this section.

Comparison and distinction. The $\text{sk-}k\text{-ALWE}^{\otimes d}$ assumption can be interpreted as a kind of polynomial integer PRG, where one has to distinguish the functions err_0 in Eq. (19) and err^* in Eq. (21). In this regard it is comparable to the assumptions of [Agr19, AJL⁺19, JLMS19, JLS19, GJLS21], with the important distinction that our “PRG” consists of dense (i.e. non-sparse) polynomials whose coefficients are superpolynomially large. Indeed, these coefficients come from short lattice points which do not expose additional structure. Hence the statistical and sum-of-squares attacks of [BBKK18, BHJ⁺19, JKLS21] are not applicable in our case. The assumptions underlying [GP21, WW21] are highly interactive, allowing an adversary to specify malicious circuits. In private-coin evasive LWE [Tsa22, VWW22] the interactivity is limited, yet attacks can still specify malicious sampling algorithms. Such attacks are not possible in our case, as we restrict interaction to fixed inputs chosen upfront, and we thereby avoid attacks in the style of [HJL21, BÜW24, HHY26, DJM⁺25, HJL25b, AMYY25]. The assumption of [DQV⁺21] has a significant flaw with respect to the underlying tensor structure, which enables a linear extraction of the noise terms [JLLS23]. It is plausible that $\text{sk-}k\text{-ALWE}^{\otimes d}$ avoids such linear flaws, as it provides only unstructured short lattice points to an adversary and, for $d = 1$, is implied by $k\text{-LWE}$ (Theorem 5). We consider $\text{sk-}k\text{-ALWE}^{\otimes d}$ to be somewhat comparable to the unbroken assumptions of [HJL25a, CLW25], in that all three are of limited interaction and hand out unstructured short lattice points. Relative to [HJL25a, CLW25], ours is fully algebraic and not parametrised by another encryption scheme, which makes it potentially easier to analyse further. Overall, the assumptions surveyed above grant more adversarial freedom and reveal significantly more complex information in the form of polynomial-size circuits, whereas $\text{sk-}k\text{-ALWE}^{\otimes d}$ exposes only constant-degree polynomials and ciphertexts. A further difference is that several of them are used to construct iO schemes directly, whereas our construction takes the full path via compact FE.

F Public-key FE for Linear Functions from LWE

We formalise the linear FE scheme overviewed in Section 2.1.

F.1 LWE to $k\text{-ALWE}$

In Definition 4 we recall the (standard, homogeneous) $k\text{-LWE}$ assumption parametrised by the Gaussian parameters χ_s and χ_e for the secret and error distributions. For suitable parameters, the $k\text{-LWE}$ assumption is known to be implied by the LWE assumption [LPSS14]¹⁴.

Definition 4 ($k\text{-LWE}$ [LPSS14]). *Parametrise $\text{par} = (k, n, m, q, \chi_s, \chi_e, \sigma)$ by λ , where k, n, m, q are positive integers with $k \leq m$, and $\chi_s, \chi_e, \sigma \in \mathbb{R}^+$. The kLWE_{par} assumption states that, for any PPT $\mathcal{A}$,*

$$\text{Adv}_{\text{par}, \mathcal{A}}^{\text{kLWE}}(\lambda) := |\Pr[\text{kLWE}_{\text{par}, \mathcal{A}}^0(1^\lambda) = 1] - \Pr[\text{kLWE}_{\text{par}, \mathcal{A}}^1(1^\lambda) = 1]| \leq \text{negl}(\lambda)$$

where the experiments $\text{kLWE}_{\text{par}, \mathcal{A}}^b$ are defined in Fig. 8.

Consider $k\text{-ALWE}^{\otimes d}$, the public-key variant of Assumption 1, with $d = 1$. We observe that the “ $\tilde{\mathbf{A}}$ track” of the experiments in Fig. 3 is simply an extension of the “ $\mathbf{A}$ track”. More precisely, when $\chi_e = \tilde{\chi}_e$, so that the noise is identically distributed across both tracks, the assumption with parameter $(m, \tilde{m})$ is equivalent to the assumption with parameter $(m + \tilde{m}, 0)$. By renaming $m + \tilde{m}$ to m , we obtain the assumption formulated below, which we call the $k\text{-ALWE}$ assumption.

¹⁴ [LPSS14] considered uniformly distributed $\mathbf{s}$; their proof applies to $\mathbf{s}$ of arbitrary distributions, i.e. $k\text{-LWE}$ with any distribution for $\mathbf{s}$ is implied by LWE with the same distribution for $\mathbf{s}$.

$\text{kLWE}_{\text{par},\mathcal{A}}^0(1^\lambda)$	$\text{kLWE}_{\text{par},\mathcal{A}}^1(1^\lambda)$
$\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{n \times m}; \mathbf{e} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m, \chi_e}$	$\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{n \times m}; \mathbf{e} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m, \chi_e}$
$\mathbf{U} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m \times k, \sigma} : \mathbf{AU} = \mathbf{0}_{n \times k} \bmod q$	$\mathbf{U} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m \times k, \sigma} : \mathbf{AU} = \mathbf{0}_{n \times k} \bmod q$
$\mathbf{s} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^n, \chi_s}$	$\mathbf{d} \leftarrow \$ \mathbb{Z}_q^m : \mathbf{d}^\top \mathbf{U} = \mathbf{0} \bmod q$
$\mathbf{b}^\top := \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \bmod q$	$\mathbf{b} := \mathbf{d} + \mathbf{e} \bmod q$
return $\mathcal{A}(\mathbf{A}, \mathbf{U}, \mathbf{b})$	return $\mathcal{A}(\mathbf{A}, \mathbf{U}, \mathbf{b})$

Fig. 8. Experiments for k -LWE assumption.

Definition 5 (k -ALWE). Parametrise $\text{par} = (k, n, m, p, q, \chi_s, \chi_e, \sigma)$ by λ , where k, n, m, p, q are positive integers with $k \leq m$, and $\chi_s, \chi_e, \sigma \in \mathbb{R}^+$. The k -ALWE $_{\text{par}}$ assumption states that, for any PPT $\mathcal{A}$,

$$\text{Adv}_{\text{par},\mathcal{A}}^{\text{kAlwe}}(\lambda) := |\Pr[k\text{-ALWE}_{\text{par},\mathcal{A}}^0(1^\lambda) = 1] - \Pr[k\text{-ALWE}_{\text{par},\mathcal{A}}^1(1^\lambda) = 1]| \leq \text{negl}(\lambda)$$

where the experiments $k\text{-ALWE}_{\text{par},\mathcal{A}}^b$ are defined in Fig. 9.

$k\text{-ALWE}_{1,\text{par},\mathcal{A}}^0(1^\lambda)$	$k\text{-ALWE}_{1,\text{par},\mathcal{A}}^1(1^\lambda)$
$\mathbf{x} \leftarrow \$ \mathcal{A}(1^\lambda) \quad // \quad \mathbf{x} \in \mathcal{X} \subseteq \mathbb{Z}^k$	$\mathbf{x} \leftarrow \$ \mathcal{A}(1^\lambda) \quad // \quad \mathbf{x} \in \mathcal{X} \subseteq \mathbb{Z}^k$
$\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{(k+n) \times m}, \mathbf{e} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m, \chi_e}$	$\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{(k+n) \times m}, \mathbf{e} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^m, \chi_e}$
$\mathbf{s} \leftarrow \$ \mathcal{D}_{\mathbb{Z}^n, \chi_s}$	$\mathbf{b} \leftarrow \$ \mathbb{Z}_q^m$
$\mathbf{c}^\top := \left\langle \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix}, \mathbf{A} \right\rangle + p \cdot \mathbf{e}^\top \bmod q \in \mathbb{Z}_q^m$	$\mathbf{c}^\top := \mathbf{b}^\top + p \cdot \mathbf{e}^\top \bmod q$
$\mathbf{K} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k} : \mathbf{A} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$	$\mathbf{K} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k} : \begin{bmatrix} \mathbf{A} \\ \mathbf{b}^\top \end{bmatrix} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{x}^\top \end{bmatrix} \bmod q$
return $\mathcal{A}(\mathbf{A}, \mathbf{c}, \mathbf{K})$	return $\mathcal{A}(\mathbf{A}, \mathbf{c}, \mathbf{K})$

Fig. 9. Experiments for k -ALWE $_{\text{par}}$. (Public-key variant of Fig. 3 with $d = 1$.)

We show that k -ALWE is implied by a standard assumption: It is well-known that there exists a reduction from LWE to k -LWE (Definition 4) for $k < m$, the following theorem provides a reduction from k -LWE to k -ALWE.

Theorem 5. Let $k, n, m, p, q \in \mathbb{N}$ such that $m \geq 2(k + n + 1) \log q$, $\sigma > \omega(\sqrt{\log m})$ and q is prime. Denote $\text{par}_0 = (k, n, m, q, \chi_s, \chi_e, \sigma)$ and $\text{par}_1 = (k, n, m, p, q, \chi_s, \chi_e, \sigma)$. If the $\text{kLWE}_{\text{par}_0}$ assumption holds, then the $k\text{-ALWE}_{\text{par}_1}$ assumption holds.

Proof. We consider a sequence of hybrids:

Hyb₀ : The $k\text{-ALWE}_{\text{par}_1}$ assumption experiment for $b = 0$, detailed in Fig. 9. In brief, the experiment samples random $\mathbf{A} \leftarrow \$ \mathbb{Z}_q^{(k+n) \times m}$, the challenge consists of $(\mathbf{K}, \mathbf{c})$ where the preimages $\mathbf{K} \leftarrow \$ \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k}$ are subject to $\mathbf{A} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$, and $\mathbf{c}^\top := (\mathbf{x}^\top, \mathbf{s}^\top) \mathbf{A} + p \cdot \mathbf{e}^\top \bmod q$ where $\mathbf{x}$ is that chosen by $\mathcal{A}$ in its first stage.

Hyb₁ : Same as **Hyb₀**, except that the vector $\mathbf{c}$ is computed as $\mathbf{c} = \mathbf{b} + p \cdot \mathbf{e} \bmod q$, where $\mathbf{b}$ is sampled uniformly random over $\mathbb{Z}_q^m$ subject to $\mathbf{b}^\top \mathbf{K} = \mathbf{x}^\top \bmod q$.

By Lemma 20 below, it holds that $\text{Hyb}_0 \approx_c \text{Hyb}_1$.

Hyb_2 : Same as Hyb_1 , except that we change the sampling procedure of matrices and vector $\mathbf{A}, \mathbf{K}, \mathbf{b}$: First sample random $\mathbf{A} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$ and $\mathbf{b} \leftarrow \mathbb{Z}_q^m$, then sample preimages $\mathbf{K} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{m \times k}$ subject to

$\begin{bmatrix} \mathbf{A} \\ \mathbf{b}^\top \end{bmatrix} \cdot \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \\ \mathbf{x}^\top \end{bmatrix} \pmod{q}$. Note that this is identical to the $k\text{-ALWE}_{\text{par}_1}$ assumption experiment for $b = 1$, detailed in Fig. 9.

By Corollary 2, we have $\text{Hyb}_1 \approx_s \text{Hyb}_2$. This completes the proof.

Lemma 20. *Under the parameters stated in Theorem 5 and the $k\text{LWE}_{\text{par}_0}$ assumption, $\text{Hyb}_0 \approx_c \text{Hyb}_1$.*

Proof. Given a PPT distinguisher $\mathcal{A}$ against the $k\text{-ALWE}_{\text{par}_1}$ problem, we construct a distinguisher $\mathcal{B}$ against Hyb_0 and Hyb_1 . Upon receiving a $k\text{LWE}_{\text{par}_0}$ problem instance $(\mathbf{A}'_1, \mathbf{K}, \mathbf{c}_1) \in \mathbb{Z}_q^{n \times m} \times \mathbb{Z}^{m \times k} \times \mathbb{Z}_q^m$, where $\mathbf{A}'_1$ is uniformly random and the preimages $\mathbf{K}$ satisfy $\mathbf{A}'_1 \mathbf{K} = \mathbf{0}_{n \times k} \pmod{q}$, let it proceed as follows:

- Sample uniform $\mathbf{A}'_0 \leftarrow \mathbb{Z}_q^{k \times m}$ subject to $\mathbf{A}'_0 \mathbf{K} = p^{-1} \cdot \mathbf{I}_k \pmod{q}$. Let $\mathbf{A} := p \cdot \begin{bmatrix} \mathbf{A}'_0 \\ \mathbf{A}'_1 \end{bmatrix}$.
- Send $\mathbf{A}$ to $\mathcal{A}$, and receive $\mathbf{x}$ from $\mathcal{A}$.
- Set $\mathbf{c}^\top := p \cdot (\mathbf{x}^\top \mathbf{A}'_0 + \mathbf{c}_1^\top) \pmod{q}$.
- Send $(\mathbf{c}, \mathbf{K})$ to $\mathcal{A}$, and return whatever $\mathcal{A}$ returns.

By Corollary 2, the matrices $(\begin{bmatrix} \mathbf{A}'_0 \\ \mathbf{A}'_1 \end{bmatrix}, \mathbf{K})$ are jointly statistically close to those from the following: First sample uniform $\begin{bmatrix} \mathbf{A}'_0 \\ \mathbf{A}'_1 \end{bmatrix} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$ then sample preimages $\mathbf{K}$ such that $\begin{bmatrix} \mathbf{A}'_0 \\ \mathbf{A}'_1 \end{bmatrix} \mathbf{K} = \begin{bmatrix} p^{-1} \cdot \mathbf{I}_k \\ \mathbf{0} \end{bmatrix} \pmod{q}$. This is equivalent to first sample uniform $\mathbf{A} = \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{A}_1 \end{bmatrix} = p \cdot \begin{bmatrix} \mathbf{A}'_0 \\ \mathbf{A}'_1 \end{bmatrix} \leftarrow \mathbb{Z}_q^{(k+n) \times m}$ then sample preimages $\mathbf{K}$ such that $\mathbf{A} \mathbf{K} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0} \end{bmatrix}$ because q is prime, i.e. exactly as in both Hyb_0 and Hyb_1 . Below, we analyse the conditional distribution of $\mathbf{c}$.

Suppose $\mathcal{B}$ interacts with the $k\text{LWE}$ experiment with bit $b = 0$, so that $\mathbf{c}_1^\top = \mathbf{s}^\top \mathbf{A}'_1 + \mathbf{e}^\top \pmod{q}$ for some error $\mathbf{e}$. Then

$$\mathbf{c}^\top = p \cdot (\mathbf{x}^\top \mathbf{A}'_0 + \mathbf{s}^\top \mathbf{A}'_1 + \mathbf{e}^\top) = (\mathbf{x}^\top, \mathbf{s}^\top) \mathbf{A} + p \cdot \mathbf{e}^\top \pmod{q}.$$

Thus, overall the simulation is statistically close to that in Hyb_0 .

Else, suppose $\mathcal{B}$ interacts with the $k\text{LWE}$ experiment with bit $b = 1$, so that $\mathbf{c}_1 = \mathbf{d}_1 + \mathbf{e} \pmod{q}$ for some random $\mathbf{d}_1 \leftarrow \mathbb{Z}_q^m$ subject to $\mathbf{d}_1^\top \mathbf{K} = \mathbf{0}_k^\top \pmod{q}$ and some error $\mathbf{e}$. Observe that $\mathbf{x}^\top \mathbf{A}_0 \pmod{q}$ satisfies $\mathbf{x}^\top \mathbf{A}_0 \cdot \mathbf{K} = \mathbf{x}^\top \pmod{q}$ by construction. We have

$$\begin{aligned} \mathbf{c}^\top &= p \cdot (\mathbf{x}^\top \mathbf{A}'_0 + \mathbf{d}_1^\top + \mathbf{e}^\top) \pmod{q} \\ &\equiv \mathbf{x}^\top \mathbf{A}_0 + \mathbf{d}_1^\top + p \cdot \mathbf{e}^\top \pmod{q} \\ &\equiv \mathbf{d}^\top + p \cdot \mathbf{e}^\top \pmod{q} \end{aligned}$$

where $\mathbf{d}^\top \leftarrow \mathbb{Z}_q^m : \mathbf{d}^\top \mathbf{K} = \mathbf{x}^\top \pmod{q}$. The first $\equiv$ follows from $p \cdot \mathbf{d}_1 \equiv \mathbf{d}_1$ since q is prime, and the second $\equiv$ is by Lemma 21 below. Overall, we conclude that $\mathcal{B}$ statistically simulates Hyb_1 to $\mathcal{A}$ in this case. $\square$

Lemma 21. *Let $m, q, k \in \mathbb{N}$. Fix any $\mathbf{z} \in \mathbb{Z}_q^k$ and $\mathbf{K} \in \mathbb{Z}_q^{m \times k}$. Suppose there exists $\hat{\mathbf{d}} \in \mathbb{Z}_q^m$ such that $\hat{\mathbf{d}}^\top \mathbf{K} = \mathbf{z}^\top \pmod{q}$. Then it holds that*

$$\left\{ \mathbf{d} \mid \mathbf{d} \leftarrow \mathbb{Z}_q^m : \mathbf{d}^\top \mathbf{K} = \mathbf{z}^\top \pmod{q} \right\} \equiv \left\{ \mathbf{d} \mid \begin{array}{l} \mathbf{d}_1 \leftarrow \mathbb{Z}_q^m : \mathbf{d}_1^\top \mathbf{K} = \mathbf{0}^\top \pmod{q} \\ \mathbf{d} := \mathbf{d}_1 + \hat{\mathbf{d}} \pmod{q} \end{array} \right\}.$$

$\text{INDExp}_{\text{FE}, \mathcal{A}}^b(1^\lambda)$	$\text{KGenO}(f)$
$(x_0^*, x_1^*) \leftarrow \mathcal{A}(1^\lambda)$	if $f(x_0^*) \neq f(x_1^*)$ then
assert $(x_0^*, x_1^*) \in \mathcal{X}^2$	return $\perp$
$(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$	return $\text{KGen}(\text{msk}, f)$
$\text{ct}^* \leftarrow \text{Enc}(\text{mpk}, x_b^*)$	
$b' \leftarrow \mathcal{A}^{\text{KGenO}}(\text{mpk}, \text{ct}^*)$	
return b'	

Fig. 10. The selective-message indistinguishability experiment for public-key functional encryption.

Proof. Denote $S_0 := \{\mathbf{y} \in \mathbb{Z}_q^m : \mathbf{y}^T \mathbf{K} = \mathbf{0}^T \bmod q\}$ and $S_z := \{\mathbf{y} \in \mathbb{Z}_q^m : \mathbf{y}^T \mathbf{K} = \mathbf{z}^T \bmod q\}$. The LHS distribution in the lemma statement is $\mathbf{d} \leftarrow \$ S_z$, and the RHS distribution is $\mathbf{d} \leftarrow \$ (S_0 + \hat{\mathbf{d}})$.

We have $S_0 + \hat{\mathbf{d}} \subseteq S_z$, since $(S_0 + \hat{\mathbf{d}}) \cdot \mathbf{K} = \hat{\mathbf{d}}^T \mathbf{K} = \mathbf{z}^T \bmod q$. Therefore $|S_0 + \hat{\mathbf{d}}| = |S_0| \leq |S_z|$. We also have $\hat{\mathbf{d}} - S_z \subseteq S_0$, since $(\hat{\mathbf{d}} - S_z) \cdot \mathbf{K} = \hat{\mathbf{d}}^T \mathbf{K} - \mathbf{z}^T = \mathbf{0}^T \bmod q$. Therefore $|\hat{\mathbf{d}} - S_z| = |S_z| \leq |S_0|$. Together, we get $|S_0| = |S_0 + \hat{\mathbf{d}}| = |S_z|$, thus $S_0 + \hat{\mathbf{d}} = S_z$ which implies the statement. $\square$

F.2 Public-key FE for Linear Functions

We recall definitions for public-key functional encryption.

Definition 6 (Public-key Functional Encryption). A public-key functional encryption (FE) scheme for a class of functions $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathcal{Z}\}$ is a tuple of PPT algorithms $\text{FE} = (\text{Setup}, \text{KGen}, \text{Enc}, \text{Dec})$ defined as follows:

$\text{Setup}(1^\lambda)$: takes as input a security parameter 1^λ , outputs a master public and secret key pair (mpk, msk) .
 $\text{KGen}(\text{msk}, f)$: takes as input the master secret key msk and a function $f \in \mathcal{F}$, and outputs a functional secret key sk_f .
 $\text{Enc}(\text{mpk}, x)$: takes as input the master public key mpk and a message $x \in \mathcal{X}$, and outputs a ciphertext ct .
 $\text{Dec}(\text{sk}_f, \text{ct})$: takes as input a functional secret key sk_f and a ciphertext ct , and returns an output $z \in \mathcal{Z}$.

Correctness. FE is correct if for all $\lambda \in \mathbb{N}$, all $(\text{mpk}, \text{msk}) \in \text{Setup}(1^\lambda)$, all function $f \in \mathcal{F}$ and all message $x \in \mathcal{X}$

$$\Pr \left[\text{Dec}(\text{sk}_f, \text{ct}) = f(x) \mid \begin{array}{l} \text{sk}_f \leftarrow \text{KGen}(\text{msk}, f) \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, x) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Indistinguishability Security. For every stateful PPT adversary $\mathcal{A}$, define

$$\text{Adv}_{\text{FE}, \mathcal{A}}^{\text{ind}}(\lambda) := \left| \Pr [\text{INDExp}_{\text{FE}, \mathcal{A}}^0(1^\lambda) = 1] - \Pr [\text{INDExp}_{\text{FE}, \mathcal{A}}^1(1^\lambda) = 1] \right|,$$

where the experiments $\text{INDExp}_{\text{FE}, \mathcal{A}}^b$ are defined in Fig. 10.

Construction. We state our public-key linear FE scheme, i.e. the public-key variant of the $d = 1$ specialisation of the scheme in Fig. 5. This FE supports bounded-norm linear functions on bounded-norm evaluation points over $\mathbb{Z}$.

Let $k, n, m, p, q, r, \sigma, \beta_x, \beta_y$ be parametrised by λ , where $r \in (-q/2, q/2] \cap \mathbb{Z}$. The linear FE scheme in Fig. 11 supports a function class $\mathcal{F}$ over a message space $\mathcal{X} \subseteq \{\mathbf{x} \in \mathbb{Z}^k : \|\mathbf{x}\|_\infty \leq \beta_x\}$. We require (for correctness) the following properties to hold for any $f \in \mathcal{F}$ and $\mathbf{x} \in \mathcal{X}$:

- $f(\mathbf{x}) \in (-r/2, r/2] \cap \mathbb{Z}$, i.e. function values are bounded, and
- there exists $\mathbf{y} \in \mathbb{Z}^k$ with ℓ_1 -norm $\|\mathbf{y}\|_1 \leq \beta_y$, such that $f(\mathbf{x}) = \langle \mathbf{x}, \mathbf{y} \rangle$, i.e. f can be computed by a k -variate degree-1 homogeneous integer polynomial with bounded coefficients.

Relative to Fig. 5, in Fig. 11 we have simplified $\text{ExtTrapGen}^{\otimes 1}$ by omitting GenFullRankSet which is redundant for $d = 1$, we let $p = r$, and we simply take $\tilde{m} = 0$. (Alternatively, we can view $(\mathbf{A}, \tilde{\mathbf{A}})$ and $\begin{bmatrix} \mathbf{K} \\ \tilde{\mathbf{K}} \end{bmatrix}$ as $\mathbf{A}$ and $\mathbf{K}$ respectively, and rename $m + \tilde{m}$ as m). It is possible to take the LWE secret $\mathbf{s} \leftarrow \mathbb{Z}_q^n$ as uniform; we keep it as Gaussian in Fig. 11 for consistency with our degree- d scheme.

Setup(1^λ)	KGen(mpk, msk, $\mathbf{y} \in \mathbb{Z}^k$)
$\mathbf{V} := \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix}$	$\mathbf{k}_y := \mathbf{K} \cdot \mathbf{y}$
$(\mathbf{A}, \mathbf{K}) \leftarrow \text{ExtTrapGen}_{k+n, m, 0, r, q}^{\otimes 1}(1^\lambda, \mathbf{V}, \sigma)$	return $\text{sk}_y := \mathbf{k}_y$
mpk := $\mathbf{A}$; msk := $\mathbf{K}$	Enc(mpk, $\mathbf{x} \in \mathcal{X}$)
return (mpk, msk)	$\mathbf{s} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^n$; $\mathbf{e} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m$
$\text{ExtTrapGen}_{k+n, m, 0, r, q}^{\otimes 1}(1^\lambda, \mathbf{V}, \sigma)$	$\mathbf{c}^\top := [\mathbf{x}^\top \ \mathbf{s}^\top] \cdot \mathbf{A} + p \cdot \mathbf{e}^\top \bmod q$
$(\mathbf{A}, \text{td}_\mathbf{A}) \leftarrow \text{TrapGen}(1^\lambda, 1^{k+n}, 1^m, q)$	return $\text{ct}_\mathbf{x} := \mathbf{c}$
$\mathbf{K} \leftarrow \text{SampPre}(\mathbf{A}, \text{td}_\mathbf{A}, \mathbf{V}, \sigma)$	Dec($\text{ct}_\mathbf{x}, \text{sk}_y$)
return $(\mathbf{A}, \mathbf{K}) \in \mathbb{Z}_q^{(k+n) \times m} \times \mathbb{Z}^{m \times k}$	return $(\text{ct}_\mathbf{x}^\top \cdot \text{sk}_y \bmod q) \bmod r$

Fig. 11. Linear functional encryption scheme (public-key version of Fig. 5 with degree $d = 1$).

Parameters. We assume a $(k + n, m, q, \sigma)$ -admissible lattice trapdoor scheme ($\text{TrapGen}, \text{SampPre}$), and set parameters as follows for correctness and security:

$$\begin{aligned}
m &\geq \max(2(k + n + 1) \log q, \lambda), & \gamma &= \sigma \cdot \left((2 \cdot \sqrt{k + n + m}) \cdot \left(\sqrt{k} \beta_x + 2 \cdot \sigma \cdot \sqrt{n + m} \right) \right) \cdot \beta_y, \\
q &\geq 2 \cdot (r \cdot \gamma + r/2) \text{ is prime.}
\end{aligned} \tag{44}$$

We have the following corollary, which follows directly from Theorem 5 and a public-key adaptation of Theorems 3 and 4.

Corollary 4. *The public-key FE scheme for linear functions in Fig. 11 with parameters as in Eq. (44) is correct and selectively IND-CPA secure under the kLWE_{par} assumption (Definition 4) with $\text{par} = (k, n, m, q, \sigma, \sigma, \sigma)$.*

The corollary below follows since the kLWE assumption (Definition 4) is implied by the LWE assumption [LPSS14]. LWE with the (optional) Gaussian LWE secret is as hard as LWE with uniform secret [BLMR13].

Corollary 5. *The public-key FE scheme for linear functions in Fig. 11 is correct and selectively IND-CPA secure under the LWE assumption.*

Remark 4 (Comparison with [ALS16]). In the linear FE scheme of [ALS16], the master public- and secret keys are $(\mathbf{A} \in \mathbb{Z}_q^{n \times m}, \mathbf{U} \in \mathbb{Z}_q^{n \times k})$ and $\mathbf{Z} \in \mathbb{Z}^{m \times k}$ respectively satisfying $\mathbf{AZ} = \mathbf{U} \bmod q$, a ciphertext encoding $\mathbf{x} \in \mathcal{M}$ is $(\mathbf{s}^\top \cdot \mathbf{A} + \mathbf{e}_0^\top, \mathbf{s}^\top \cdot \mathbf{U} + \lfloor q/K \rfloor \cdot \mathbf{x}^\top + \mathbf{e}_1^\top) \bmod q$, and the functional secret key of a low-norm $\mathbf{y} \in \mathcal{Y}$ is $(\mathbf{K} \cdot \mathbf{y}, \mathbf{y})$. Packing the terms into matrix notation, we obtain the relations

$$\begin{bmatrix} \mathbf{I}_k & \mathbf{0}_{k \times m} \\ \mathbf{U} & \mathbf{A} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{I}_k \\ -\mathbf{Z} \end{bmatrix} = \begin{bmatrix} \mathbf{I}_k \\ \mathbf{0}_{n \times k} \end{bmatrix} \bmod q$$

$$\mathbf{c}^T = [\mathbf{c}_1^T \ \mathbf{c}_0^T] = [\lfloor q/K \rfloor \cdot \mathbf{x}^T \ \mathbf{s}^T] \cdot \begin{bmatrix} \mathbf{I}_k & \mathbf{0}_{k \times m} \\ \mathbf{U} & \mathbf{A} \end{bmatrix} + [\mathbf{e}_1^T \ \mathbf{e}_0^T] \bmod q,$$

which is almost identical to the corresponding relations in our linear FE scheme up to the choice of “lower-bit” vs “higher-bit” encodings. The main difference between our scheme and that of [ALS16] is in key generation. In our view, sampling the keys with a trapdoor allows a more natural degree- d generalisation.

Remark 5 (Insecurity without $\mathbf{K}$). With the same construction paradigm, one could alternatively consider a scheme where $\text{msk} = \mathbf{K}$ is replaced by simply the trapdoor $\text{td}_{\mathbf{A}}$ of $\mathbf{A}$, and a functional key for $\mathbf{y}$ is sampled as $\text{sk}_{\mathbf{y}} = \mathbf{A}^{-1} \begin{pmatrix} \mathbf{0} \\ \mathbf{y} \end{pmatrix}$. Correctness follows in the same way. Unfortunately, such modified scheme is insecure. Indeed, an adversary can query functional keys for $\mathbf{y}$ ’s that are linearly dependent over $\mathbb{Z}$, for example, $\mathbf{y}_i^T = (i, \dots, i)$ for many distinct i ’s. Upon obtaining $\text{sk}_{\mathbf{y}_i}$ for $i = 1, 2, \dots, m$, the adversary can compute $\text{sk}_{\mathbf{y}_i} - i \cdot \text{sk}_{\mathbf{y}_1}$ to obtain m short preimages of $\mathbf{0}$, the differences being likely linearly independent over $\mathbb{Q}$, since each $\text{sk}_{\mathbf{y}_i}$ is Gaussian distributed. As such, the adversary would likely obtain an Ajtai trapdoor of $\mathbf{A}$. In contrast, similar attack strategy would not apply to the scheme in Fig. 11, since here each functional key $\text{sk}_{\mathbf{y}}$ is deterministic conditioned on msk .

G From secret-key FE for NC^0 to iO

We consider the following efficiency notions of an FE scheme:

Definition 7. Let $\delta, \epsilon > 0$ be constants. A secret-key FE scheme for a function class $\mathcal{C} = \mathcal{C}_\lambda$ with maximum circuit size $S = S(\lambda) := \max_{f \in \mathcal{C}} |f|$ and a message space $\mathcal{M} = \{0, 1\}^N$, where $N = N(\lambda)$, is said to be

- ϵ -sublinear-size, if for all $\lambda \in \mathbb{N}$, all $\mathbf{x} \in \mathcal{M}$ and all $\text{msk} \in \text{Setup}(1^\lambda)$, all ciphertexts $\text{ct} \in \text{FE.Enc}(\text{msk}, \mathbf{x})$ have size $|\text{ct}| \leq \text{poly}(\lambda, N) \cdot S^{1-\epsilon}$,
- δ -efficient, if for all $\lambda \in \mathbb{N}$, all $\mathbf{x} \in \mathcal{M}$ and all $\text{msk} \in \text{Setup}(1^\lambda)$, all ciphertexts $\text{ct} \in \text{FE.Enc}(\text{msk}, \mathbf{x})$ have size $|\text{ct}| \leq \text{poly}(\lambda) \cdot N^{1+\delta}$.

We say an FE is sublinear-size (resp. efficient) to mean that it is ϵ -sublinear-size for some constant $\epsilon > 0$ (resp. δ -efficient for some constant $\delta > 0$).

An ϵ -sublinear-size FE is sometimes also called weakly sublinearly compact [LPST16b] in the literature. We recall the definition of a pseudorandom generator (PRG) and its stretch.

Definition 8 (PRG). A secure pseudorandom generator (PRG) for input length $\nu = \nu(\lambda)$ with $s(\nu)$ -stretch is a deterministic function $f : \{0, 1\}^\nu \rightarrow \{0, 1\}^{s(\nu)}$ satisfying $\{f(r) \mid r \leftarrow \mathcal{R} \{0, 1\}^\nu\} \approx_c \mathcal{U}(\{0, 1\}^{s(\nu)})$.

We recall the definition of indistinguishability obfuscation.

Definition 9 (Indistinguishability Obfuscation (iO) [GGH⁺13b]). An indistinguishability obfuscation for a circuit family $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a pair of PPT algorithms $\text{iO} = (\text{Obf}, \text{Eval})$ satisfying the following properties:

- *Correctness:* For all $\lambda \in \mathbb{N}$, all circuits $\Gamma \in \mathcal{C}_\lambda$, all inputs x it holds that $\text{Eval}(\text{Obf}(\Gamma), x) = \Gamma(x)$.
- *Security:* For any PPT adversary $\mathcal{A}$, any pair of circuits $(\Gamma_0, \Gamma_1) \in \mathcal{C}_\lambda$ such that $|\Gamma_0| = |\Gamma_1|$ and $\Gamma_0(x) = \Gamma_1(x)$ on all inputs x , it holds that

$$|\Pr[\mathcal{A}(\text{Obf}(\Gamma_0)) = 1] - \Pr[\mathcal{A}(\text{Obf}(\Gamma_1)) = 1]| \leq \text{negl}(\lambda).$$

Lemma 22 below follows from a transform which (or whose variants) appeared in multiple previous works, e.g. [AJ15, BV15, LV16, KNT18]. The only change is in the parameters of the underlying and resulting FE schemes, which are adapted to suit our FE scheme.

Lemma 22 (Adapted from [LV16, Proof of Theorem 5]). *Let $\alpha > \delta > 0$ be constants. Assuming there exist*

- *a (resp. sub-exponential) secure PRG in NC^0 with $\nu^{1+\alpha}$ -stretch, and*
- *a (resp. sub-exponential) selectively secure δ -efficient secret-key FE for NC^0 ,*

then there exists a (resp. sub-exponential) selectively secure ϵ -sublinear-size single-key secret-key FE for P/poly , where $\epsilon = \frac{\alpha-\delta}{1+\alpha} > 0$.

Proof. We use the transform given by [LV16, Theorem 5], which is a special case of those of [BV15, Theorem 4.3], [AJ15, Theorem 3] and [AJS15], except that we are in the secret-key setting, and we recall it below.

Let PRG be a PRG in NC^0 with $\nu^{1+\alpha}$ -stretch, FE be a δ -efficient secret-key FE for NC^0 where $\delta < \alpha$, and $\text{RE} = (\text{Decomp}, \text{Enc}, \text{Eval})$ be the decomposable randomised encoding scheme of [AIK06, Theorem 4.14] (see e.g. [LV16] for a detailed definition), which supports the class of P/poly circuits and which is (computationally) secure assuming the existence of a PRG in NC^1 (which is implied by a PRG in NC^0). The ϵ -sublinear-size secret-key FE for P/poly is as follows:

- The master secret-key msk is that of the FE.
- A ciphertext encrypting a message $\mathbf{x}$ is $\text{ct} \leftarrow \text{FE.Enc}(\text{msk}, (\mathbf{x}, s_1, s_2, 0))$, where $s_1, s_2 \leftarrow \{0, 1\}^{\ell^{1/(1+\alpha)}}$ are random bits, and $\ell = \ell(\lambda, \mathbf{x}, |f|) \in \mathbb{N}$ is a parameter of RE specifying its number of components.
- To generate a function key for a function f , do the following:
 - Decompose f into $(F_i, S_i, I_i)_{i \in [\ell]} \leftarrow \text{RE.Decomp}(f)$, where each F_i is a circuit representing part of f , $S_i \subseteq [N]$ and $I_i \subseteq [\ell]$ are sets of indices.
 - For each $i \in [\ell]$, sample a random bit $b_i \leftarrow \{0, 1\}$.
 - The function key consists of $\text{sk}_{f_i} \leftarrow \text{FE.KGen}(\text{msk}, f_i)$ for each $i \in [\ell]$, for the function

$$f_i(\mathbf{x}, s_1, s_2, b) = \begin{cases} \text{RE.Enc}(F_i, \mathbf{x}|_{S_i}, \text{PRG}(s_1)|_{I_i}) & \text{if } b = 0, \\ b_i \oplus \text{PRG}(s_2)|_{I_i} & \text{if } b = 1, \end{cases}$$

where the notation $\cdot|_S$ means taking the subset of indices specified by the set S .

- To decrypt, run $\text{FE.Dec}(\text{sk}_{f_i}, \text{ct})$ for each $i \in [\ell]$ to get the randomised encoding, then run RE.Eval on the randomised encoding to obtain $f(\mathbf{x})$.

The only difference between the above and [LV16, Theorem 5] is that (1) we have instantiated RE with that of [AIK06, Theorem 4.14] to support the class of P/poly circuits, and (2) we are in the secret-key setting. This RE satisfies the efficiency requirements of [LV16, Theorem 5], namely, RE.Enc is in NC^0 , an encoding of $f \in \text{poly}(\lambda)$ has number of components $\ell \leq \text{poly}(\lambda) \cdot |f|$, and RE.Enc takes ℓ bits of randomness.

Correctness follows from that of FE and RE, where we note that f_i is supported by FE since both RE.Enc and PRG are in NC^0 . One can directly verify that the security proof is unchanged, in particular, being in the secret-key setting does not alter the proof (which is also noted in [AJS15, P.2]). We refer to [LV16, Theorem 5], [BV15, Theorem 4.3], or [AJ15, Theorem 3] for the details.

The above FE is ϵ -sublinear-size, since the resulting ciphertext has size

$$\begin{aligned} |\text{ct}| &\leq \text{poly}(\lambda) \cdot |(\mathbf{x}, s_1, s_2, 0)|^{1+\delta} \\ &= \text{poly}(\lambda) \cdot (N + 2 \cdot \ell^{1/(1+\alpha)} + 1)^{1+\delta} \\ &= \text{poly}(\lambda, N) \cdot \ell^{(1+\delta)/(1+\alpha)} \\ &\leq \text{poly}(\lambda, N) \cdot |f|^{(1+\delta)/(1+\alpha)} = \text{poly}(\lambda, N) \cdot |f|^{1-\epsilon}, \end{aligned}$$

where the first $\leq$ is by the δ -efficient property of FE, the second $\leq$ by the property $\ell \leq \text{poly}(\lambda) \cdot |f|$ of RE stated above. Set $\epsilon = \frac{\alpha-\delta}{1+\alpha} > 0$ to get the last equality. $\square$

The following lemma shows that secret-key single-key FE for P/poly is sufficient to imply a shallow sub-exponentially secure iO for $\text{P}^{\log}/\text{poly}$ with non-trivial efficiency (a.k.a. XiO). It is a direct adaptation of the result in [BNPW16].

Lemma 23 (Adapted from [BNPW16, Theorem 3.1 and Remark 3.1]). *If there exists a ϵ -sublinear-size (resp. sub-exponential) single-key selectively-secure secret-key FE for P/poly , then there exists a (resp. sub-exponential) secure indistinguishability obfuscator for $\mathsf{P}^{\log}/\text{poly}$ where the size of the obfuscated circuit is bounded by $2^{\gamma N} \text{poly}(\lambda, S)$, where S is the description size of the circuit, N is its input length, $0 < \gamma = \frac{c}{c+\epsilon} < 1$ and $c \in O(1)$ is a constant.*

Proof. The following XiO construction is from [BNPW16, Theorem 3.1 and Remark 3.1], originally used to show that a single-key single-input FE with non-trivial efficiency can be transformed into a strong XiO¹⁵. We show that the same construction can transform a single-key single-input FE *without* non-trivial efficiency to a (standard, non-strong) XiO.

Let FE be a single-key single-input secret-key FE for P/poly with ϵ -sublinear-size ciphertext. Let $c \geq 1$ be a constant such that the time to derive a key for f is bounded by $|f|^c \cdot \text{poly}(\lambda)$. Denote $t := c + \epsilon - 1 > 0$.

Consider a circuit Γ with input length N , hence its truth table size is $M := 2^N$. Denote its truth table entries as $\Gamma(i, j)$, where $i \in [M^{1/(t+1)}], j \in [M^{t/(t+1)}]$. We consider the following XiO:

- Generate $\text{msk} \leftarrow \text{FE.Setup}(1^\lambda)$.
- For each $j \in [M^{t/(t+1)}]$ do the following:
 - Let Γ_j be the (1-input) circuit that, on input $i \in [M^{1/(t+1)}]$, returns $\Gamma(i, j)$. Compute FE ciphertext ct_j encrypting Γ_j .
- Compute FE secret key $\text{sk}_f \leftarrow \text{FE.KGen}(\text{msk}, f)$, where f is the circuit that, on input Γ_j , outputs the evaluations $(\Gamma_j(1), \dots, \Gamma_j(M^{1/(t+1)}))$.
- The XiO consists of all $M^{t/(t+1)}$ ciphertexts ct_j and the secret key sk_f .

The above construction is identical to that in [BNPW16, Theorem 3.1 and Remark 3.1], so the security proof is unchanged. In particular, selective security of FE implies that of the XiO, see [BNPW16] for the details.

The size compression of the XiO is analogous to the runtime analysis in [BNPW16] for strong XiO.

The function f evaluates Γ_j to $M^{1/(t+1)}$ points and $|\Gamma_j| \leq |\Gamma|$, hence $|f| \leq M^{1/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda)$. The size of sk_f is

$$|\text{sk}_f| \leq |f|^c \leq M^{c/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda).$$

Since FE is ϵ -sublinear-size, each ciphertext ct_j has size bounded by

$$|\text{ct}| \leq |f|^{(1-\epsilon)} \cdot \text{poly}(|\Gamma|, \lambda) \leq M^{(1-\epsilon)/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda).$$

Overall, the above XiO has size

$$\begin{aligned} M^{t/(t+1)} \cdot |\text{ct}| + |\text{sk}_f| &\leq M^{(t+1-\epsilon)/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda) + M^{c/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda) \\ &= M^{c/(t+1)} \cdot \text{poly}(|\Gamma|, \lambda) \\ &= 2^{c/(t+1) \cdot N} \cdot \text{poly}(|\Gamma|, \lambda) \\ &= 2^{c/(c+\epsilon) \cdot N} \cdot \text{poly}(|\Gamma|, \lambda). \end{aligned}$$

The compression factor is $0 < \gamma = c/(c+\epsilon) < 1$. □

The following lemma from [LPST16a] is well-known. It states that, under the sub-exponential hardness of LWE, to achieve iO, it suffices to achieve XiO.

Lemma 24 ([LPST16a, BDGM22]). *Assuming sub-exponentially hard LWE, if there exists a sub-exponentially secure indistinguishability obfuscator for $\mathsf{P}^{\log}/\text{poly}$ with depth at most $\text{poly}(\lambda, N)$ and with non-trivial efficiency, i.e. the size of the obfuscated circuit is bounded by $2^{\gamma N} \text{poly}(\lambda, S)$ for some $\gamma < 1$, where S is the description size of the circuit and N its input length, then there exists an indistinguishability obfuscator for P/poly with sub-exponential security.*

¹⁵ A strong XiO is an XiO but additionally with non-trivial runtime efficiency guarantee.

In the above, $P^{\log}/\text{poly}$ denotes the class of circuits with description size $S = \text{poly}(\lambda)$ and input length $N = \text{polylog}(\lambda)$. The lemma poses no restriction on the runtime of the obfuscator, which can be as large as $2^N \text{poly}(\lambda, S)$, except that its depth should be at most $\text{poly}(\lambda, N)$.

Finally, in Lemma 22 it suffices for the δ -efficient FE to support degree- d polynomials for some fixed large enough d . Targeting time $T = 2^{(k+\lambda)^e}$, the FE of Section 6 has $|\text{ct}| \leq O((k + \lambda)^{\max(1+2e, 4e)})$, i.e. it is δ -efficient with $\delta = \max(2e, 4e - 1)$. Combining Theorem 4 with Lemmas 22 to 24 gives Corollary 6.

Corollary 6. *Assuming*

- *the (resp. sub-exponential) LWE assumption holds,*
- *the (resp. sub-exponential) $\text{sk-k-ALWE}^{\otimes d}$ assumption holds for some sufficiently large $d \in O(1)$, in the sub-exponential case against adversaries running in time $2^{(k+\lambda)^e}$ for a constant $e > 0$,*
- *the existence of a (resp. sub-exponential) secure PRG in NC^0 with $\nu^{1+\alpha}$ -stretch for some constant $\alpha > 0$, and*
- *in the sub-exponential case, $\max(2e, 4e - 1) < \alpha$,*

then there exists a (resp. sub-exponential) secure indistinguishability obfuscator for P/poly .