$i\mathcal{O}$ from LWE with Structured Preimage

Shweta Agrawal* Anuja Modi† Anshu Yadav‡ Shota Yamada§

Abstract

Constructing $i\mathcal{O}$ from simple, safe assumptions in the post-quantum regime is an important problem with wide-ranging applications. In this work, we make progress towards this goal. We formulate a new, arguably simple assumption called *LWE with Structured Preimage* that suffices to build $i\mathcal{O}$. Our assumption is instance-independent, falsifiable and non-interactive. Our assumption posits that a set of LWE (Learning With Errors) samples is indistinguishable from random given a short Gaussian matrix which is carefully chosen to satisfy the required functionality while avoiding problematic leakage.

We provide a thorough cryptanalysis of our assumption, using the extensive cryptanalytic effort against the related evasive LWE assumption (Wee, Eurocrypt 2022 and Tsabary, Crypto 2022) to gain confidence. We emphasize that our assumption is superficially similar but fundamentally different from evasive LWE. Thus, we obtain $i\mathcal{O}$ from a simple lattice assumption whose hardness is somewhat better understood as compared to brand new assumptions.

*IIT Madras, India, shweta@cse.iitm.ac.in
†IIT Madras, India, anujamodi97@gmail.com
‡IIT Kanpur, India, anshuy@cse.iitk.ac.in
§AIST Tokyo, Japan, yamada-shota@aist.go.jp

Contents

1 Introduction **3**
- 1.1 Our Results . . . 4
- 1.2 Discussion about our Assumption. . . . 5
- 1.3 Related Work . . . 6

2 Technical Overview **7**
- 2.1 $i\mathcal{O}$ via Oblivious LWE Sampling. . . . 7
- 2.2 Comparison with AKY. . . . 13
- 2.3 Cryptanalysis of LSP Assumption . . . 14

3 Preliminaries **15**
- 3.1 Indistinguishability Obfuscation . . . 16
- 3.2 Lattice Preliminaries . . . 16
- 3.3 GSW Homomorphic Encryption and Evaluation . . . 18
- 3.4 Homomorphic Evaluation Procedures . . . 19

4 LWE with Structured Preimage **19**
- 4.1 The Assumption . . . 21
- 4.2 Cryptanalysis . . . 22

5 $i\mathcal{O}$ via LWE with Structured Preimage **29**
- 5.1 Definition: Oblivious LWE Sampler . . . 29
- 5.2 Construction: Oblivious LWE Sampler . . . 30
- 5.3 Security . . . 35

1 Introduction

An important primitive in the domain of encrypted computation, which has been the subject of intense study over the last decade, is that of program obfuscation, which seeks to garble programs while preserving their input-output behaviour. Indistinguishability Obfuscation ($i\mathcal{O}$) [BGI+01] is a particular variant of program obfuscation which provides the following indistinguishability-style guarantee: given two circuits C_0, C_1 such that they have the same size and compute the same function, and an obfuscation of one of them chosen at random, a bounded adversary cannot distinguish between the two cases. A line of exciting works uncovered the power of $i\mathcal{O}$, showing that it can be used to instantiate almost every known cryptographic primitive – please see [GGH+16, JLS21] for a detailed discussion.

$i\mathcal{O}$ has been studied extensively over the last decade, over a long line of beautiful works that provide instantiations under various assumptions. The breakthrough work of Garg et al. [GGH+16] gave a first candidate construction for $i\mathcal{O}$ under a new algebraic assumption called the "Multilinear Jigsaw Puzzle". Following this, several works provided heuristic constructions under different assumptions [GGH+13, BGK+13, BR14, PST14, AGIS14, BMSZ16, CLT13, CLT15, GGH15, CHL+15, BWZ14, CGH+15, HJ15, BGH+15, Hal15, CLR15, MF15, MSZ16, DGG+18, Lin16, LV16, AS17, Lin17, LT17, GJK18, AJS18, Agr19, AP20, LM18, JLMS19, BIJ+20, GJLS21], until finally, the celebrated work of Jain, Lin and Sahai [JLS21] achieved the first provably secure construction. This construction was based on four well-studied assumptions, which were reduced to three [JLS22] and then "two and a half" [RVV24] in subsequent work. However, the crucial reliance on pairings in these constructions make them insecure in the post-quantum regime.

$i\mathcal{O}$ in the Post-Quantum Regime. In the post-quantum regime, basing security of $i\mathcal{O}$ on standard assumptions has met with much less success despite significant effort from the community. Several devastating attacks [HJL21, JLLS23, BDJ+24, HJL25, AMYY25] broke the assumptions underlying several serious candidates [WW21, DQV+21, GP21], even for weaker primitives such as *pseudorandom* obfuscation [AKY25, BDJ+24]. Even though many constructions can either be fixed or remain unbroken [MSZ16, AP20, WW21, GP21, BDGM23, BDGM20, CCMR24], their security is not well understood and the absence of reduction to a clean, well-understood assumption is dissatisfying. Thus, an important question is:

Can we build $i\mathcal{O}$ from "nice" assumptions in the post-quantum regime?

Above, by "nice" we mean an assumption that is (i) short and simple, (ii) falsifiable, (iii) non-interactive, (iv) instance-independent, (v) "safe". The term "safe" would ideally imply reduction from a standard lattice assumption, but in the absence of this, at the very least the landscape of cryptanalytic techniques should be well-understood, and the assumption should take a principled approach to avoiding known and unknown attacks. One would like to have evidence that an attack on the assumption would be surprising.

Very recently [HJL25] and [CLW25] built candidate $i\mathcal{O}$ from new, simpler assumptions from lattices. While this is encouraging progress, these assumptions require time and more cryptanalytic effort to build faith in them. Therefore, we believe it is of value to persist in identifying plausible assumptions towards this important goal.

$i\mathcal{O}$ from Lattices. The starting point of our work is the recent work of Agrawal, Kumari and Yamada [AKY25] (henceforth AKY) who constructed $i\mathcal{O}$ for pseudorandom functionalities from the (private-coin) evasive LWE assumption[1]. Upon closer examination, it is easy to see that the *construction* by AKY can already support arbitrary functions – thus, as a heuristic, this readily provides a pathway to $i\mathcal{O}$. However,

[1] A very similar construction, from the same assumption, was provided concurrently by [BDJ+24].

heuristics for $i\mathcal{O}$ abound, and the goal is to identify a nice assumption which allows to *prove* security, as discussed above.

Looking at the assumption underlying the AKY construction, it is immediate that it not only violates most of the above desiderata (it is non-falsifiable, instance-dependent and interactive) but, worse, is actually known to be *false* in general. Indeed, a recent line of works [BUW24, DJM+25, AMYY25, HJL25, BDJ+24] demonstrated multiple attacks on the private-coin evasive LWE assumption in its general form. It is instructive to study the evasive LWE assumption next.

Evasive LWE. The evasive LWE assumption roughly says that if

$$\left(\mathbf{B},\ \mathbf{P},\ \ \mathbf{s}^\intercal\mathbf{B}+\mathbf{e}_\mathbf{B}^\intercal,\ \ \mathbf{s}^\intercal\mathbf{P}+\mathbf{e}_\mathbf{P}^\intercal,\ \ \mathsf{aux}\ \right) \approx_c \left(\ \mathbf{B},\ \mathbf{P},\ \$,\ \ \$,\ \ \mathsf{aux}\right)$$

where $\$$ represents random, then

$$\left(\ \mathbf{B},\ \mathbf{P},\ \mathbf{s}^\intercal\mathbf{B}+\mathbf{e}_\mathbf{B}^\intercal,\ \ \mathbf{B}^{-1}(\mathbf{P}),\ \ \mathsf{aux}\right) \approx_c \left(\ \mathbf{B},\ \mathbf{P},\ \ \ \$,\ \ \mathbf{B}^{-1}(\mathbf{P}),\ \ \mathsf{aux}\right)$$

Above $\mathbf{B}^{-1}(\mathbf{P})$ refers to a low norm matrix, say $\mathbf{K}$, such that $\mathbf{BK} = \mathbf{P} \mod q$. Evidently, given $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$ and $\mathbf{s}^\intercal\mathbf{B}+\mathbf{e}_\mathbf{B}^\intercal$, the adversary can compute $\mathbf{s}^\intercal\mathbf{P}+\mathbf{e}_\mathbf{B}^\intercal\mathbf{K}$. The hope/conjecture is that this is *all* the adversary can do with $\mathbf{K}$.

An important distinction between variants of evasive LWE are "public-coin" and "private-coin", where the former means that the randomness used by the sampler, who outputs $\mathbf{P}$ and auxiliary information aux, is made available to the adversary, and the latter means that the sampler's random coins need to be hidden from the adversary.

While a series of works uncovered vulnerabilities in the private-coin version of evasive LWE [BUW24, DJM+25, AMYY25, HJL25, BDJ+24, AMYY26], the public coin-version remains unbroken despite significant cryptanalytic effort. However, constructing $i\mathcal{O}$ from public-coin evasive LWE seems unlikely – indeed, there are only a handful of constructions from the public-coin version and of these, most have been improved to rely on fairly weak, falsifiable lattice assumptions such as "succinct" LWE or "decomposed LWE" [Wee24, CW24, WW25, Wee25, CHW25, AMR25, LLL25]. Even much weaker objects than $i\mathcal{O}$, such as "null" $i\mathcal{O}$ are not known to be constructible from public-coin evasive LWE.

Our Approach. In this work, we take a close look at the space of attacks developed against evasive LWE and ask: can this extensive cryptanalytic effort *help* to formulate a nice assumption, whose security we can have confidence on, that can suffice for $i\mathcal{O}$? We believe that such an approach, if executed with sufficient care, can be *safer* than formulating a completely new assumption against which attacks might emerge only a few years later.

We answer the above question in the affirmative. We design a simple, falsifiable, instance-independent, non-interactive assumption to build $i\mathcal{O}$, and provide a thorough cryptanalysis of this using knowledge of attacks against evasive LWE. Our construction follows the blueprint by Wee and Wichs [WW21], who identified a simple primitive called "Oblivious LWE Sampler" (OLS) which suffices for $i\mathcal{O}$.

To build an OLS, we begin with the techniques developed by AKY for building pseudorandom functional encryption (PRFE) and enhance these to incorporate several additional safety measures. This yields a candidate construction but a construction is far from sufficient – as mentioned above, several works can give heuristics for $i\mathcal{O}$ and the core challenge is in *proving* security. The security game of OLS is fundamentally different from that of PRFE (please see Section 2.2 for a detailed discussion) and identifying a nice assumption as well as reducing security to it are both challenging and subtle.

1.1 Our Results

We formulate an assumption called "LWE with Structured Preimage" (LSP) which is simple, falsifiable, instance-independent and non-interactive. We then use this assumption to construct $i\mathcal{O}$. Our assumption

posits that a set of LWE samples is indistinguishable from random given a short Gaussian matrix which is carefully designed to satisfy the constraint required by correctness while avoiding all known attacks. A detailed description of LSP is provided in Section 4, but at a high level, it posits that the real and ideal worlds below are indistinguishable[2]. Below, $\underset{\sim}{x}$ denotes a noisy version of x.

Real World	Ideal World
$\mathbf{B}, \mathbf{D}, \mathsf{hpk}, \mathbf{P}$	$\mathbf{B}, \mathbf{D}, \mathsf{hpk}, \mathbf{P}$
$\mathbf{c}_{\mathbf{B}}^{\intercal} = \underset{\sim}{\mathbf{t}^{\intercal}\mathbf{B}}$ $\mathbf{X} = \mathsf{hpk} \cdot \mathbf{R} - \mathsf{sd} \otimes \mathbf{G}$ $\mathsf{bct} = \underset{\sim}{\mathbf{t}^{\intercal}(\mathbf{D} - (1, \mathsf{bits}(\mathbf{X})) \otimes \mathbf{G})},$	$\mathbf{c}_{\mathbf{B}}^{\intercal} : \mathsf{random}$ $\mathbf{X} : \mathsf{random}$ $\mathsf{bct} : \mathsf{random}$
$\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$	$\mathbf{K} = \begin{pmatrix} \mathbf{B} \\ \mathbf{c}_{\mathbf{B}}^{\intercal} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{P} \\ \mathbf{d}^{\intercal} + \mathsf{random}_{1/2}^{\intercal} \end{pmatrix}$

Table 1: LWE with Structured Preimage (simplified)

Above, $\mathbf{B}, \mathbf{D}$ are uniform matrices, $\mathsf{hpk} = \begin{pmatrix} \mathbf{A}_{\mathsf{fhe}} \\ \underset{\sim}{\mathbf{t}^{\intercal}\mathbf{A}_{\mathsf{fhe}}} \end{pmatrix}$ is an LWE sample which serves as the public key of the GSW FHE scheme [GSW13], while $\mathbf{X}$ is a GSW ciphertext of a random binary vector sd under the public key hpk. Next, $\mathbf{d}$ is a publicly computable vector from $\mathbf{X}$ and bct such that the relation between the ciphertexts $(\mathbf{c}_{\mathbf{B}}^{\intercal}, \mathbf{X}, \mathsf{bct})$ and the key $\mathbf{K}$ in the two worlds is preserved, and $\mathsf{random}_{1/2}$ denotes a random vector whose support is *half* the space. At a high level, choosing a random vector over the entire space would not allow to preserve the correctness of the primitive (OLS) – random over half the space captures the intuition that there is no leakage except the single output bit which is required by the functionality. Most importantly, $\mathbf{P}$ is a carefully crafted, non-random but *publicly computable* matrix, which makes it harder to misuse. Without the Gaussian matrix $\mathbf{K}$, the distributions are indistinguishable from LWE.

1.2 Discussion about our Assumption.

As seen above, our assumption is quite simple. A subtle and crucial feature of our assumption is that it is "single challenge" in flavour, namely the LWE secret is a single vector as against a matrix whose size depends on the problem instance. This makes the assumption succinct, but even more importantly, it is crucial to avoid a multi-challenge assumption since there are "incompressibility" attacks against the related evasive LWE assumption in this regime [AMYY25, BDJ+24, AMYY26]. These attacks crucially exploit the multi-challenge nature of the assumption – specifically, the pseudorandom functional encryption scheme that is shown to be impossible by [AMYY25] cannot be constructed *even* from private coin evasive LWE restricted to the single challenge setting.

Another crucial difference from evasive LWE is the extra row $\mathbf{d} + \mathsf{random}_{1/2}$ in the $\mathbf{K}$ equation in the ideal distribution. As discussed above, this is to maintain the correctness of the scheme that is constructed from the assumption while hiding everything except the most significant bit via the masking with $\mathsf{random}_{1/2}$. The informed reader may be concerned about whether we inherit the flawed heuristic from evasive LWE, namely, "correlated error is as good as i.i.d error" which has been shown to be false

[2]Technically, the preimage $\mathbf{K}$ is random and it is the image $\mathbf{BK}$ which is structured, but we use the term "structured preimage" to denote "preimage to structured matrix" for brevity.

in general [BDJ+24, HJL25, DJM+25, AMYY25]. However, via careful design of the matrix $\mathbf{P}$, we actually rely on a much weaker heuristic which posits that a suitably large PRF value can flood a small *rounding* error which is correlated with the PRF seed but is *not controllable* by the adversary and seems very hard to exploit, *even* if it chooses malicious circuit implementations as exploited by prior work [AMYY25, HJL25]. We also emphasize that this is the only heuristic aspect of our assumption, and that it has been thoroughly tested against all known attack strategies. We believe that an attack against our assumption would be very surprising and could lead to a completely new class of attacks against evasive LWE, possibly affecting the security of public-coin evasive LWE.

We emphasize that our assumption is technically incomparable to evasive LWE. Indeed, we believe it is philosophically much weaker than private coin evasive LWE since the latter implies primitives beyond $i\mathcal{O}$, such as (i) Pseudorandom Obfuscation [BDJ+24, AKY24b] and pseudorandom FE [AKY25], which are shown to be impossible [BDJ+24, AMYY25] and (ii) Succinct Witness Encryption [BDJ+24], which is not known even from $i\mathcal{O}$.

We also briefly discuss the connection with the circular evasive LWE assumption [HLL23] which was recently broken by [AMYY25]. Similarly to circular evasive LWE, the assumption includes an FHE ciphertext $\mathbf{X}$ together with an encoding bct that embeds $\mathbf{X}$. However, unlike circular evasive LWE, our FHE ciphertext does not need to encrypt its own secret but instead encodes a random string. Moreover, unlike circular evasive LWE, the choice of the matrix $\mathbf{P}$ is not free for exploitation by the adversary but fixed and carefully designed to incorporate extra safety measures that make it difficult to misuse.

Compared to public coin evasive LWE, the most crucial difference is that public coin evasive LWE is designed to avoid the "zeroizing regime" where the adversary may obtain equations in low norm secret values over the integers. Seen another way, in public coin evasive LWE, the adversary only sees conjectured pseudorandom values over the entire space upon decryption. On the other hand, in our assumption, the adversary recovers conjectured pseudorandom values over only *half* the space, as discussed above. This certainly opens up our assumption to a new class of attacks as compared to public coin evasive LWE, against which we require careful analysis. On the other hand, public coin evasive LWE is a non-falsifiable, instance-dependent, multi-stage assumption whereas ours is not. In terms of similarity, the matrix $\mathbf{P}$ is publicly computable in both assumptions.

1.3 Related Work

We provide a brief comparison of techniques with the recent works on post-quantum $i\mathcal{O}$:

- Cini, Lai and Woo [CLW25] built a lattice based $i\mathcal{O}$ from NTRU and a new indistinguishability based assumption which they call *equivocal LWE*.
- Hsieh, Jain and Lin [HJL25] built a lattice based "functional encoding" scheme, which implies $i\mathcal{O}$, from a falsifiable assumption called *circular security with random opening* (CRO). They also propose an indistinguishability based assumption, which they call IND-CRO and build OLS from this.

The assumptions in both these works, similar to ours, fall in the category of "LWE with hints". The techniques and assumption used in [CLW25] are very different from ours. The techniques and assumptions used in the Circular Random Opening (CRO) assumption of [HJL25] have more similarity but are nevertheless significantly different and technically incomparable.

In more detail, the CRO assumption involves a complex circularity between LWE samples, even when the opening is withheld. On the other hand, in LSP, there is no circularity between the LWE samples when the preimage is withheld. The motivation for introducing the complex circularity in [HJL25] is to eliminate error leakage by using "exact" openings – namely, openings that satisfy exact rather than approximate equality. This is a principled approach to avoid attacks that have emerged from exploiting noise leakage.

On the flip side, the complex circularities, being new and relatively specialized, could potentially give other types of leakage which might only surface in future work. In contrast, our assumption does allow some noise leakage but this is limited to rounding error which seems very hard to exploit. One could think of the two works taking different positions in the circularity-leakage tradeoff: both positions are heuristic and based on known attacks. We see the main advantage of our approach as deriving heavily from the cryptanalysis efforts against the evasive LWE family of assumptions, for which attacks are, by now, very well understood. This gives confidence since our assumption avoids all the dangers of a thoroughly investigated cryptanalytic landscape in a principled way. In addition, the simplicity of the assumption makes it amenable to further study and analysis.

2 Technical Overview

2.1 $i\mathcal{O}$ via Oblivious LWE Sampling.

We proceed to outline the ideas developed for constructing an oblivious LWE sampler from our new assumption.

Oblivious LWE Sampler. Wee and Wichs [WW21] defined the primitive of Oblivious LWE Sampler (OLS) and showed that it suffices to build $i\mathcal{O}$. At a high level, OLS allows for obliviously generating K LWE samples $\mathbf{b}_i^* = \mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^*$ for $i \in [K]$ with respect to a (tall) matrix $\mathbf{A}^* \in \mathbb{Z}_q^{M\times W}$ where $\mathbf{e}_i^* \in \mathbb{Z}_q^M$ is low norm, without knowing $\mathbf{s}_i^*$ and $\mathbf{e}_i^*$. More technically, OLS consists of four algorithms $(\mathsf{genCRS}, \mathsf{init}, \mathsf{Sample}, \mathsf{Sim})$ as follows:

- The genCRS algorithm on input K outputs long CRS, denoted by crs, whose length can depend on K but which must be independent of $\mathbf{A}^*$.
- The init algorithm on input $\mathbf{A}^*$ outputs succinct public parameters, denoted by pub, whose length is sublinear in $KM^{1-\epsilon}\log q$.
- The deterministic algorithm Sample on input crs, pub and index i, outputs the LWE sample $\mathbf{b}_i^*$.
- The simulator Sim on input arbitrary LWE samples $\widehat{\mathbf{b}}_i = \mathbf{A}^*\widehat{\mathbf{s}}_i + \widehat{\mathbf{e}}_i$ for all i, and $\mathbf{A}^*$, outputs crs, pub, and $\{\mathbf{s}_i^*\}_i$. We require that $(\mathsf{crs}, \mathbf{A}^*, \mathsf{pub}, \{\mathbf{s}_i^*\}_i)$ obtained by Sample, where $\{\mathbf{s}_i^*\}_i$ is recovered from $\{\mathbf{b}_i^*\}_i$ by a trapdoor for $\mathbf{A}^*$ and $(\mathsf{crs}, \mathbf{A}^*, \mathsf{pub}, \{\mathbf{s}_i^* + \widehat{\mathbf{s}}_i\}_i)$ obtained by the simulated distribution are indistinguishable.

There exists a series of transformations implying Indistinguishability Obfuscation ($i\mathcal{O}$) from OLS. OLS is known to imply Functional Encoding [WW21], which in turn implies Exponentially Efficient $i\mathcal{O}$ (XiO) [WW21], which in turn implies $i\mathcal{O}$ [LPST16]. In more detail:

Theorem 2.1 ([WW21]). Assuming that there exists a sub-exponentially secure OLS and that the sub-exponentially secure LWE assumption holds, there exists $i\mathcal{O}$.

Wee and Wichs [WW21] then provided a construction for OLS under a new conjecture, called the *homomorphic pseudorandom LWE samples* (HPLS) conjecture. While this was a principled new assumption satisfying many desiderata in the post-quantum regime – instance independence, falsifiability and non-interactiveness – it was shown by the subsequent work of [HJL21] that the assumption was sensitive to the particular implementation of the circuit implementing a specific component (PRF) used in the assumption. While this attack can be heuristically prevented by fixing a "safe" looking circuit implementation (for instance by using a universal circuit), it is dissatisfying to require such specifications in the statement of an assumption. Ideally, one would like the assumption to be independent of the choice of the circuit implementation used in its implementation. Hence we seek a different assumption to build OLS.

How far can LWE go? A natural first step is to explore the capabilities of the LWE assumption to see how much it can help us. Towards this, first observe that the component $\mathbf{A}^*\mathbf{s}_i^*$ in $\mathbf{b}_i^*$ for $i \in [K]$ can be computed using inner product FE (IPFE), which we know how to build from LWE. Recall that in an IPFE scheme, the ciphertext encrypts a vector $\mathbf{x} = (x_1, \ldots, x_n) \in \mathbb{Z}_q^n$ and the key is provided for a vector $\mathbf{y} = (y_1, \ldots, y_n) \in \mathbb{Z}_q^n$ such that the decryption algorithm outputs $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i$.

Recap of (noisy) IPFE from LWE. Here we recap a close variant[3] of the IPFE scheme of [ALS16].

- Setup samples a matrix $\mathbf{B}$ along with its trapdoor $\mathbf{B}_\gamma^{-1}$ and an another LWE matrix $\mathbf{A}_{\mathsf{IP}}$, and sets the master public and secret keys as $\mathsf{impk} = (\mathbf{B}, \mathbf{A}_{\mathsf{IP}})$ and $\mathsf{imsk} = \mathbf{B}_\gamma^{-1}$.
- $\mathsf{Enc}(\mathsf{impk}, \mathbf{x})$ samples a random $\mathbf{t}$ and computes $\mathbf{c}_{\mathsf{IP}}^\intercal = \mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal + \mathbf{x}^\intercal \otimes \mathbf{g}^\intercal$, $\mathbf{c}_{\mathbf{B}}^\intercal = \mathbf{t}^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal$ and returns $\mathsf{ict} = (\mathbf{c}_{\mathbf{B}}^\intercal, \mathbf{c}_{\mathsf{IP}}^\intercal)$.
- $\mathsf{KeyGen}(\mathsf{imsk}, \mathbf{y})$ outputs $\mathsf{isk} = \mathbf{B}^{-1}(\mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y}))$.
- $\mathsf{Dec}(\mathsf{isk}, \mathsf{ict}, \mathbf{y})$ computes

$$\begin{aligned}
\mathbf{c}_{\mathsf{IP}}^\intercal \cdot \mathbf{G}^{-1}(\mathbf{y}) - \mathbf{c}_{\mathbf{B}}^\intercal \cdot \mathsf{isk} &= (\mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal + \mathbf{x}^\intercal \otimes \mathbf{g}^\intercal)\mathbf{G}^{-1}(\mathbf{y}) - (\mathbf{t}^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal)\mathbf{B}^{-1}(\mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y})) \\
&= \mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y}) + \mathbf{e}_{\mathsf{IP}}^\intercal \mathbf{G}^{-1}(\mathbf{y}) + \mathbf{x}^\intercal \mathbf{y} - \mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y}) - \mathbf{e}_{\mathbf{B}}^\intercal \mathbf{B}^{-1}(\mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y})) \\
&= \mathbf{x}^\intercal \mathbf{y} + \mathsf{noise}_{\mathsf{IP}},
\end{aligned}$$

  where $\mathsf{noise}_{\mathsf{IP}} = \mathbf{e}_{\mathsf{IP}}^\intercal \mathbf{G}^{-1}(\mathbf{y}) - \mathbf{e}_{\mathbf{B}}^\intercal \mathbf{B}^{-1}(\mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}(\mathbf{y}))$

Encoding the mask of oblivious LWE sample: Let us begin with the modest goal of compressing $\mathbf{b}_i^* = \mathbf{A}^*\mathbf{s}_i^*$ for $i \in [K]$. We denote by $\mathbf{S}^*$ and $\mathbf{B}^*$ the concatenation of vectors $\mathbf{s}_i^*$ and $\mathbf{b}_i^*$ respectively, for $i \in [K]$. Now, it is evident that we can compute the mask $\mathbf{A}^*\mathbf{S}^*$ using IPFE described above – encode $\{\mathbf{s}_i^*\}_{i\in[K]}$ into K ciphertexts and $\{\mathbf{a}_j^*\}_{j\in[M]}$ into M keys, such that $\mathbf{A}^*\mathbf{S}^*[j,i]$, for $i \in [K], j \in [M]$ can be recovered via decryption of the i-th ciphertext with the j-th key. In more detail,

- We encode each column $\mathbf{s}_i^*$ of $\mathbf{S}^*$ into an IPFE ciphertext ict_i as above.
- Each row of $\mathbf{A}^*$ is encoded into an IPFE key isk_j as above.
- Decryption works as above to obtain $\mathbf{a}_j^*\mathbf{s}_i^* = \mathbf{A}^*\mathbf{S}^*[j,i]$, for $i \in [K], j \in [M]$.

If we could ignore the error $\mathbf{E}^* = (\mathbf{e}_1^*, \ldots, \mathbf{e}_K^*)$, then for sufficiently large K and M, we have already obtained a succinct $\mathsf{pub} = (\{\mathsf{ict}_i\}_{i\in[K]}, \{\mathsf{isk}_j\}_{j\in[M]})$, that encodes the mask $\mathbf{A}^*\mathbf{S}^*$ of $\mathbf{B}^*$. However, this is just wishful thinking since compressing the error $\mathbf{E}^*$ is the core technical challenge and the root cause of all non-standardness of lattice based assumptions underlying $i\mathcal{O}$!

Computing $\mathbf{E}^*$: Since the LWE mask $\mathbf{A}^*\mathbf{S}^*$ was compressible using an inner product functional encryption scheme, a natural question to ask is whether the error term $\mathbf{E}^*$ can be computed by a more powerful scheme. Recall that the general notion of Functional Encryption (FE) is a generalization of public key encryption where a ciphertext is associated with a vector $\mathbf{x}$, a secret key is associated with a circuit f and decryption enables recovery of $f(\mathbf{x})$ and nothing else. Since we seek to compress an LWE noise term $\mathbf{E}^*$, a matrix of dimension $O(MK)$, a natural idea is to leverage a pseudorandom generator (PRG). In particular, we choose a PRG with stretch $O(M)$ and sample K PRG seeds, $\mathsf{sd}_1, \ldots, \mathsf{sd}_K$, where sd_i is used to generate the i-th column of $\mathbf{E}^*$. If we had an FE scheme powerful enough to compute the PRG, we could encode the PRG seeds into ciphertexts and provide appropriate keys to do the expansion. However, an FE scheme powerful enough to compute a PRG with the desired succinctness properties itself implies $i\mathcal{O}$ so it would seem we are back to square one.

[3] As stated, it is insecure as an IPFE scheme, but we will modify it later, so here we present the simplified version.

The AKY *FE scheme for Pseudorandom Functionalities.* We recap the recent construction of compact FE scheme for pseudorandom functionalities by Agrawal, Kumari and Yamada [AKY25].

Preparations: Evaluation Algorithms by Boneh et al. (BGG^+*).* The seminal work of Boneh et al. [BGG+14] developed algorithms for evaluating arithmetic functions on the ciphertext as well as the public key of an ABE scheme, which form the cornerstone of several subsequent constructions. Their core technique is as follows: given an input $\mathbf{x} \in \{0,1\}^\ell$, and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, one can homomorphically evaluate a circuit $f : \{0,1\}^\ell \to \{0,1\}$ on an "input encoding" matrix of form $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ by multiplying on the right by a low norm matrix $\mathbf{H}^f_{\mathbf{A},\mathbf{x}}$ to obtain the term $\mathbf{A}_f - f(x)\mathbf{G}$. Here, $\mathbf{G}$ is a special gadget matrix defined as follows. Let $\mathbf{g} = [1, 2, 2^2, \ldots, 2^{\log q}]^\intercal$ and $\mathbf{G} = \mathbf{I} \otimes \mathbf{g}^\intercal$. In key evaluation, one can homomorphically evaluate a circuit $f : \{0,1\}^\ell \to \{0,1\}$ on $\mathbf{A}$ to obtain $\mathbf{A}_f = \mathbf{A} \cdot \mathbf{H}^f_{\mathbf{A}}$ for some low norm matrix $\mathbf{H}^f_{\mathbf{A}}$. In ciphertext evaluation, given an attribute $\mathbf{x}$ and corresponding attribute encoding of the form $\mathbf{s}^\intercal(\mathbf{A} - \mathbf{x}^\intercal \otimes \mathbf{G})$, which we refer to as BGG^+ encoding, right multiplication by $\mathbf{H}^f_{\mathbf{A},\mathbf{x}}$ yields $\mathbf{s}^\intercal(\mathbf{A}_f - f(\mathbf{x})\mathbf{G})$ without substantially blowing up the noise in the encoding since $\mathbf{H}^f_{\mathbf{A},\mathbf{x}}$ is low norm. Skipping several details, since the key generator can compute $\mathbf{A}_f = \mathbf{A} \cdot \mathbf{H}^f_{\mathbf{A}}$, it can provide a matching key which allows the decryptor to cancel out the masking term $\mathbf{s}^\intercal \mathbf{A}_f$ and proceed with decryption. We refer to $\mathbf{H}^f_{\mathbf{A}}$ and $\mathbf{H}^f_{\mathbf{A},\mathbf{x}}$ as the PK and CT evaluation matrices respectively.

Quick Recap of AKY *FE scheme.* We provide a high level description of the FE for Pseudorandom functionalities (PRFE) developed by [AKY25], which in turn is heavily based on the ideas developed in [HLL23, AKY24a].

- The setup algorithm samples matrices $\mathbf{A}_{\mathsf{att}}$ and $(\mathbf{B}, \mathbf{B}^{-1})$ of appropriate dimensions and outputs $\mathsf{mpk} := (\mathbf{A}_{\mathsf{att}}, \mathbf{B})$ and $\mathsf{msk} := \mathbf{B}^{-1}$. Here, $\mathbf{B}^{-1}$ is the trapdoor for $\mathbf{B}$ which allows to compute short preimages $\mathbf{B}^{-1}(\mathbf{U})$ for any target matrix $\mathbf{U}$.
- The encryptor on input $\mathbf{x}$ does the following:
  1. It samples a GSW secret key $\mathbf{s} = (\bar{\mathbf{s}}^\intercal, -1)^\intercal$ and a PRF seed $\mathsf{sd} \leftarrow \{0,1\}^\lambda$.
  2. It then computes a GSW ciphertext, $\mathbf{X} = \mathsf{hct}_{\mathbf{s}}(\mathbf{x}, \mathsf{sd})$, using public key $\mathbf{A}_{\mathsf{fhe}} = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{s}}^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}^\intercal_{\mathsf{fhe}} \end{pmatrix}$ and randomness $\mathbf{R}$. The notation $\mathsf{hct}_{\mathbf{s}}(\cdot)$ denotes that the ciphertext is decryptable by $\mathbf{s}$.
  3. It computes a BGG^+ encoding of $\mathbf{X}$ using randomness $\mathbf{s}$ as $\mathbf{c}^\intercal_{\mathsf{att}} := \mathbf{s}^\intercal(\mathbf{A}_{\mathsf{att}} - \mathbf{X} \otimes \mathbf{G}) + \mathbf{e}^\intercal_{\mathsf{att}}$.
  4. It computes $\mathbf{c}^\intercal_{\mathbf{B}} := \mathbf{s}^\intercal \mathbf{B} + \mathbf{e}^\intercal_{\mathbf{B}}$ and outputs the ciphertext $\mathsf{ct} = (\mathbf{c}_{\mathbf{B}}, \mathbf{c}_{\mathsf{att}}, \mathbf{X})$.
- The key generator on input $\mathsf{msk} = \mathbf{B}^{-1}$ and function f does the following.
  (a) Samples a nonce $\mathbf{r} \leftarrow \{0,1\}^\lambda$ and defines function $\mathrm{F}[f, \mathbf{r}]$, with f and $\mathbf{r}$ hardwired, as $\mathrm{F}[f, \mathbf{r}](\mathbf{x}, \mathsf{sd}) = f(\mathbf{x}) \lfloor q/2 \rceil + \mathsf{PRF}(\mathsf{sd}, \mathbf{r})$. It then computes the FHE evaluation circuit $\mathsf{VEval}_{\mathrm{F}}$ w.r.t. the function $\mathrm{F}[f, \mathbf{r}]$ (this can be computed using the knowledge of $\mathrm{F}[f, \mathbf{r}]$). Note that the circuit $\mathsf{VEval}_{\mathrm{F}}$ can be used to compute on a GSW ciphertext for an input, say $\mathbf{y}$, to recover a GSW ciphertext encoding $\mathrm{F}[f, \mathbf{r}](\mathbf{y})$.
  (b) Next, it computes the matrix $\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}}}$ for the circuit $\mathsf{VEval}_{\mathrm{F}}$ using the public matrix $\mathbf{A}_{\mathsf{att}}$. Recall that the matrix $\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}}}$ and $\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}$ (which can be computed given $\mathsf{VEval}_{\mathrm{F}}$, $\mathbf{A}_{\mathsf{att}}$ and $\mathbf{X}$) will satisfy the relation
  $$(\mathbf{A}_{\mathsf{att}} - \mathbf{X} \otimes \mathbf{G})\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} = \mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}}} - \mathsf{VEval}_{\mathrm{F}}(\mathbf{X}).$$

(c) It sets $\mathbf{A}_{\mathsf{F}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}}}$, samples $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{A}_{\mathsf{F}})$ and outputs $\mathsf{sk}_f = (\mathbf{K}, \mathbf{r})$.

– The decryption on input $\mathsf{sk}_f = (\mathbf{K}, \mathbf{r})$ and $\mathsf{ct} = (\mathbf{c}_{\mathbf{B}}, \mathbf{c}_{\mathsf{att}}, \mathbf{X})$ work as follows.

(a) It computes the matrix $\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}$ for the circuit $\mathsf{VEval}_{\mathsf{F}}$ using $\mathbf{A}_{\mathsf{att}}$ and $\mathbf{X}$.

(b) Next, it computes $\mathbf{z} := \mathbf{c}_{\mathbf{B}}^{\mathsf{T}} \cdot \mathbf{K} - \mathbf{c}_{\mathsf{att}}^{\mathsf{T}} \cdot \mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}$, rounds $\mathbf{z}$ coordinate-wise and output the most significant bits.

To see the correctness of the scheme, we note that

$$\begin{aligned} \mathbf{c}_{\mathsf{att}}^{\mathsf{T}} \cdot \mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} &= \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} + \mathbf{s}^{\mathsf{T}}(\mathsf{hct}_{\mathbf{s}}(\mathrm{F}(\mathbf{x},\mathsf{sd}))) \\ &= \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} - \mathrm{F}(\mathbf{x},\mathsf{sd}) + \mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}} \\ &\approx \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}} - \mathrm{F}(\mathbf{x},\mathsf{sd}) \end{aligned} \tag{1}$$

where $\mathbf{R}_{\mathsf{F}}$ is a matrix with small entries determined by $\mathbf{R}$ and F and the second equality follows by automatic decryption since the FHE secret key is reused as the LWE randomness. Now to remove the masking term "$\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}}$" we compute $\mathbf{c}_{\mathbf{B}}^{\mathsf{T}} \cdot \mathbf{K} \approx \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}}$ and thus $\mathbf{z} \approx \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}} - \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}} + \mathrm{F}(\mathbf{x},\mathsf{sd}) = f(\mathbf{x})\,\lfloor q/2 \rceil + \mathsf{PRF}(\mathsf{sd},\mathbf{r})$. Now, rounding gives us bits of $f(\mathbf{x})$ as long as $|\mathsf{PRF}(\mathsf{sd},\mathbf{r})| \leq q/4$. The exact decryption error is

$$\mathbf{e}_{\mathbf{B}}^{\mathsf{T}}\mathbf{K} + \mathsf{PRF}(\mathsf{sd},\mathbf{r}) - (\mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}} + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}) \tag{2}$$

The proof of security follows from private coin evasive LWE – the details are not important for this overview. However, for their proof to work it is required that f is a pseudorandom functionality.

Error Correlation Attack and Fix. Subsequent to the first online appearance of AKY, the works of [AMYY25, HJL25] showed that if we allow the adversary to choose the circuit implementation of the PRF used in the construction along with its corresponding homomorphic evaluation, then the adversary can create a correlation between the error term resulting from FHE evaluation $\mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}}$ and the PRF output $\mathsf{PRF}(\mathsf{sd},\mathbf{r})$ which is used for flooding the error. This correlation can in turn be exploited to obtain leakage that harms the security of the scheme. To address this, AKY suggested a fix that fractures the error correlation somehow, so that any remaining dependence becomes so weak that it is "dysfunctional". They achieve this via a modulus reduction operation that throws away all the accumulated error with any problematic correlations and replaces this with a benign rounding error, which seems extremely hard to exploit. They conjecture that this technique suffices to overcome error-correlation attacks, even if the circuit is chosen maliciously.

More concretely, the updated AKY scheme modifies the definition of $\mathbf{A}_{\mathsf{F}}$ so that the problematic term $\mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}}$ (and even $\mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{F}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}$) no longer appear in the decryption error (see Equation (2)). At the same time, the scheme still enables approximate recovery of $\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{F}} - \mathrm{F}(\mathbf{x},\mathsf{sd})$, as in Equation (1). To achieve this, AKY splits the output of the function into higher-order and lower-order components as

$$\mathrm{F}(\mathbf{x},\mathsf{sd}) = \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathbf{x},\mathsf{sd}) + \mathrm{F}_{\mathsf{low}}(\mathbf{x},\mathsf{sd})$$

where Δ is a parameter to be fixed later. They then define $\mathbf{H}^{\mathsf{high}}$ and $\mathbf{H}^{\mathsf{low}}$ to be the GSW homomorphic evaluation matrices corresponding to $\Delta\mathrm{F}_{\mathsf{high}}$ and $\Delta\mathrm{F}_{\mathsf{low}}$, respectively. Next, they set $\mathbf{A}_{\mathsf{high}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{high}}$ and $\mathbf{A}_{\mathsf{low}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{low}}$, and define $\mathbf{A}_{\mathsf{F}} = \Delta \cdot \left\lfloor \frac{\mathbf{A}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathsf{low}}}{\Delta} \right\rfloor$. Now, for $\mathsf{xx} \in \{\mathsf{high},\mathsf{low}\}$, we have

$$\begin{aligned} \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}^{\mathsf{T}} \cdot \mathbf{H}^{\mathsf{xx}}_{\mathbf{X}}}{\Delta} \right\rceil &= \left\lfloor \frac{\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{xx}} - \Delta\mathrm{F}_{\mathsf{xx}}(\mathbf{x},\mathsf{sd}) + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{xx}}_{\mathbf{X}} + \mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}}}{\Delta} \right\rceil \\ &= \left\lfloor \frac{\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{xx}} + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}^{\mathsf{xx}}_{\mathbf{X}} + \mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathsf{F}}}{\Delta} \right\rceil - \mathrm{F}_{\mathsf{xx}}(\mathbf{x},\mathsf{sd}) \\ &= \left\lfloor \frac{\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{xx}}}{\Delta} \right\rceil - \mathrm{F}_{\mathsf{xx}}(\mathbf{x},\mathsf{sd}), \end{aligned}$$

where the first equality follows by the same argument as in Equation (1). The third equality holds under the condition $\|\mathbf{e}_{\mathsf{fhe}}^{\mathsf{T}}\mathbf{R}_{\mathrm{F}} + \mathbf{e}_{\mathsf{att}}^{\mathsf{T}}\mathbf{H}_{\mathbf{X}}^{\mathsf{xx}}\| \ll \Delta$, unless the value $\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathsf{xx}}$ lies too close to the boundary for the rounding such that addition of the noise term (highlighted in red) causes it to cross the border. Intuitively, this event appears to occur with negligible probability. However, AKY does not address this issue formally. Looking ahead, in our construction, which uses the above machinery as one of its components, we explicitly randomize $\mathbf{A}_{\mathsf{xx}}$ to address this subtle issue, thereby ensuring that the equation holds with high probability.

Combining the cases $\mathsf{xx} = \mathsf{high}$ and $\mathsf{xx} = \mathsf{low}$, we obtain

$$\begin{aligned}
& \Delta\left\lfloor \frac{\mathbf{c}_{\mathsf{att}}\mathbf{H}_{\mathbf{X}}^{\mathsf{high}}}{\Delta} \right\rceil + \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}\mathbf{H}_{\mathbf{X}}^{\mathsf{low}}}{\Delta} \right\rceil \\
= \;& \Delta\left\lfloor \frac{\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{high}}}{\Delta} \right\rceil + \left\lfloor \frac{\mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathsf{low}}}{\Delta} \right\rceil - \Delta\mathrm{F}_{\mathsf{high}}(\mathbf{x},\mathsf{sd}) - \mathrm{F}_{\mathsf{low}}(\mathbf{x},\mathsf{sd}) \\
= \;& \Delta\Big(\mathbf{s}^{\mathsf{T}}\left\lfloor \frac{\mathbf{A}_{\mathsf{high}}}{\Delta} \right\rceil + \mathbf{e}^{\mathsf{high}}\Big) + \Big(\mathbf{s}^{\mathsf{T}}\left\lfloor \frac{\mathbf{A}_{\mathsf{low}}}{\Delta} \right\rceil + \mathbf{e}^{\mathsf{low}}\Big) - \mathrm{F}(\mathbf{x},\mathsf{sd}) \\
= \;& \mathbf{s}^{\mathsf{T}}\mathbf{A}_{\mathrm{F}} - \mathrm{F}(\mathbf{x},\mathsf{sd}) + \underbrace{\Delta\mathbf{e}^{\mathsf{high}} + \mathbf{e}^{\mathsf{low}}}_{\text{small}}
\end{aligned}$$

Here, $\mathbf{e}^{\mathsf{high}}$ and $\mathbf{e}^{\mathsf{low}}$ denote the rounding errors arising from the floor operations, which can be made sufficiently small when $\|\mathbf{s}\|$ is small. In the description of our OLS construction that follows, we incorporate this countermeasure against the attack as a design component.

Construction: Oblivious LWE Sampler. We now describe our OLS construction at a high level. The goal is to obliviously compute the LWE sample $\mathbf{b}_i^* = \mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^*$. The idea here is to have K many ciphertexts and a single secret key as part of pub such that each i^{th} ciphertext (ct_i) encodes $\mathbf{s}_i^*$, a PRG seed sd_i and $\beta \in \{0,1\}$, the secret key (sk) encodes $\mathbf{A}^* \in \mathbb{Z}_q^{M\times W}$, function PRG and crs. Decryption of ct_i using sk outputs $\mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^* = \mathbf{A}^*\mathbf{s}_i^* + \mathsf{PRG}(\mathsf{sd}_i) + \beta\cdot\mathsf{crs}_i$. Here, the masking term is generated using techniques from IPFE as described above while the PRG seed is encoded using techniques adapted from the PRFE of AKY.

- $\mathsf{genCRS}(1^\lambda, 1^K)$. It samples a $\mathsf{crs} \leftarrow \mathbb{Z}_q^{M\times K}$ and outputs it.
- $\mathsf{init}(1^K, \mathsf{crs}, \mathbf{A}^*)$.
  - Run Setup of IPFE and PRFE: Samples secret $\mathbf{S}^*$, GSW matrix $\bar{\mathbf{A}}_{\mathsf{fhe}}$, BGG+ matrix $\mathbf{A}_{\mathsf{att}}$ and IPFE matrix $\mathbf{A}_{\mathsf{IP}}$. It additionally samples $(\mathbf{B}, \mathbf{B}^{-1})$ of appropriate dimensions.
  - The ciphertext $\mathsf{ct}_i = \{\mathbf{c}_{\mathbf{B},i}, \mathsf{hct}_i, \mathsf{bct}_i, \mathsf{ict}_i\}$ for $i \in [K]$ is computed as:
    - Encrypt $\mathbf{S}^*$ using IPFE: The IPFE encoding of $\mathbf{s}_i^*$ is computed as $\mathsf{ict}_i^{\mathsf{T}} = \mathbf{t}_i^{\mathsf{T}}\mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP},i}^{\mathsf{T}} - \mathbf{s}_i^* \otimes \mathbf{g}^{\mathsf{T}}$. We also provide $\mathbf{c}_{\mathbf{B},i} = \mathbf{t}_i^{\mathsf{T}}\mathbf{B} + \mathbf{e}_{\mathbf{B},i}^{\mathsf{T}}$.
    - Encrypt sd using AKY PRFE:
      - Samples GSW secret key $\mathbf{t}_i = (\bar{\mathbf{t}}_i, -1)$ and PRG seed $\mathsf{sd}_i \leftarrow \{0,1\}^\lambda$. It then computes the GSW ciphertext of $(\mathsf{sd}_i, i, \beta = 0)$ i.e.

$$\mathsf{hct}_i = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}_i^{\mathsf{T}}\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe},i}^{\mathsf{T}} \end{pmatrix}\mathbf{R}_i - (\mathsf{sd}_i, i, 0)\otimes\mathbf{G}$$

        where $\mathbf{R}_i$ is the randomness used.
      - It then treats this as an attribute of BGG+ ABE scheme and computes the BGG+ encoding of hct_i i.e. $\mathsf{bct}_i = \mathbf{t}_i^{\mathsf{T}}(\mathbf{A}_{\mathsf{att}} - \mathsf{hct}_i \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att},i}^{\mathsf{T}}$ using randomness $\mathbf{t}_i$.

* Providing extra LWE samples: Compute $c_{\mathsf{high},i} = \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{high}} + e_{\mathsf{high},i}$ and $c_{\mathsf{low},i} = \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{low}} + e_{\mathsf{low},i}$. Set $\mathsf{ct}_i = (\mathbf{c}_{\mathbf{B},i}, \mathsf{hct}_i, \mathsf{bct}_i, \mathsf{ict}_i, c_{\mathsf{high},i}, c_{\mathsf{low},i})$

– Generate combined key for IPFE and AKY PRFE: To generate the sk, we first hardwire the crs into the function $\mathsf{LWESamp}$ that on input $(\mathsf{sd}_i, i, \beta)$, outputs $\mathsf{PRG}(\mathsf{sd}_i)^\intercal + \beta \cdot \mathsf{crs}_i$. Let $\mathrm{F} = \mathsf{LWESamp}[\mathsf{crs}]$.

* Splitting F for modulus reduction: We then define $\mathrm{F}_{\mathsf{high}}$ and $\mathrm{F}_{\mathsf{low}}$ s.t. $\mathrm{F}(\mathsf{sd}_i, i, \beta) = \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathsf{sd}_i, i, \beta) + \mathrm{F}_{\mathsf{low}}(\mathsf{sd}_i, i, \beta)$.
* Generation of matrix $\mathbf{A}_{\mathsf{err}}$: Do as follows.
  · Compute matrix $\mathbf{H}^{\mathsf{high}}$ and $\mathbf{H}^{\mathsf{low}}$ for GSW homomorphic evaluation circuit of $\Delta \cdot \mathrm{F}_{\mathsf{high}}$ and $\Delta \cdot \mathrm{F}_{\mathsf{low}}$ respectively using $\mathbf{A}_{\mathsf{att}}$[4].
  · Set $\mathbf{A}_{\mathsf{high}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{high}}$ and $\mathbf{A}_{\mathsf{low}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{low}}$.
  · Set $\mathbf{W}_{\mathsf{high}} = (1, \ldots, 1) \otimes \mathbf{w}_{\mathsf{high}}$, and $\mathbf{W}_{\mathsf{low}} = (1, \ldots, 1) \otimes \mathbf{w}_{\mathsf{low}}$.
  · Set $\mathbf{A}_{\mathsf{err}} = \Delta \cdot \left\lfloor \frac{\mathbf{A}_{\mathsf{high}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathsf{low}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor$.
* Generation of matrix $\mathbf{A}_{\mathsf{mask}}$: Set $\mathbf{A}_{\mathsf{mask}} := \mathbf{A}_{\mathsf{IP}} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$.
* Combined secret key: Compute $\mathsf{sk} := \mathbf{K} = \mathbf{B}^{-1}(\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}})$.

– Set $\mathsf{pub} = (\mathbf{A}^*, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathsf{sk}, \{\mathsf{ct}_i\}_{i \in [K]})$

- $\mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i)$. Compute matrix $\mathbf{H}_{\mathsf{hct}}^{\mathsf{high}}$ and $\mathbf{H}_{\mathsf{hct}}^{\mathsf{low}}$ and output

$$(\mathbf{b}_i^*)^\intercal = \mathbf{c}_{\mathbf{B},i}^\intercal \mathbf{K} - \mathsf{ict}_i^\intercal \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) - \left(\Delta \left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{high}} + \mathbf{c}_{\mathsf{high},i}^\intercal}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{low}} + \mathbf{c}_{\mathsf{low},i}^\intercal}{\Delta} \right\rfloor \right)$$

  where $\mathbf{c}_{\mathsf{xx},i}^\intercal = c_{\mathsf{xx},i} \cdot (1, \ldots, 1)$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$.

- $\mathsf{Sim}(\mathbf{A}^*, \{\widehat{\mathbf{b}}_i\}_i)$.

  – Set $\mathsf{crs}_i := \widehat{\mathbf{b}}_i$ for all i; sample $\mathbf{S}^* \leftarrow \mathbb{Z}_q^{W \times K}$, and compute pub as in init algorithm, where hct_i in ct_i now encrypts $(\mathsf{sd}_i, i, \beta = 1)$ for all i.
  – Output $(\mathsf{crs}, \mathsf{pub}, \{\mathbf{s}_i^*\}_i)$.

Correctness. Correctness follows broadly from the correctness of IPFE and PRFE and the algebraic structure which allowed to combine their respective keys. It is easy to see that $\mathbf{c}_{\mathbf{B},i}^\intercal \mathbf{K} = \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{err}} + \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{mask}} + \mathbf{e}_{\mathbf{B},i}^\intercal \mathbf{K}$ and $\mathsf{ict}_i^\intercal \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) = \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{mask}} + \mathbf{e}_i^\intercal \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) - (\mathbf{A}^* \mathbf{s}_i^*)^\intercal$. We note that $\Delta \cdot \left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{high}} + \mathbf{c}_{\mathsf{high},i}^\intercal}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{low}} + \mathbf{c}_{\mathsf{low},i}^\intercal}{\Delta} \right\rfloor = \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{err}} - \mathrm{F}(\mathsf{sd}_i, i, 0) + \mathsf{error}$, which follows by the similar analysis to the correctness of AKY. For more detailed analysis, we refer the readers to Section 5.

Hence, we obtain $(\mathbf{b}_i^*)^\intercal = (\mathbf{A}^* \mathbf{s}_i^*)^\intercal + \mathrm{F}(\mathsf{sd}_i, i, 0) - \mathsf{error}' = (\mathbf{A}^* \mathbf{s}_i^*)^\intercal + \mathsf{PRG}(\mathsf{sd}_i)^\intercal - \mathsf{error}' \approx (\mathbf{A}^* \mathbf{s}_i^*)^\intercal + \mathsf{PRG}(\mathsf{sd}_i)^\intercal$ since $\beta = 0$ and parameters are set such that $\mathsf{PRG}(\mathsf{sd}_i)$ drowns error'. Hence, we get $\mathbf{b}_i^* = \mathbf{A}^* \mathbf{s}_i^* + \mathbf{e}_i^*$ where $\mathbf{e}_i^* \approx \mathsf{PRG}(\mathsf{sd}_i)$.

Succinctness. We provide a high-level argument for succinctness. We refer the reader to Section 5.2 for more rigorous analysis. We observe the following:

- The ciphertext components are for inputs whose length are independent of both K and M. Hence, size of each ct_i is independent of M, K. Thus, overall size of $\{\mathsf{ct}_i\}_{i \in [K]} = O(K)$.

[4]We drop $\mathbf{A}_{\mathsf{att}}$ from subscript of $\mathbf{H}$ to ensure readability.

- Size of $\mathbf{A}_{\mathsf{att}}$ is independent of both K and M.
- Size of $\mathbf{A}^*$ depends on M but is independent of K. The size of sk depends on the number of columns in $\mathbf{B}$ and in $\mathbf{A}_{\mathsf{high/low}}$, which in turn is linear in the length of the function output, i.e., M. Specifically, note that the size of sk is independent of the size of the function F which has crs hardwired into it.
- Thus, overall we get $|\mathsf{pub}| = O(K+M) \leq (KM)^{1-\epsilon}$ for $0 < \epsilon < 1$ and sufficiently large K, M.

Security. We observe that the simulator behaves in the same way as the real world, except for only two changes: (i) it sets $\mathsf{crs}_i = \widehat{\mathbf{b}}_i$ for all $i \in [K]$, and (ii) encrypts $\beta = 1$ (instead of 0 in the real world). Therefore, when Sample is run with the simulator's output, it returns $\mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^* + \widehat{\mathbf{b}}_i = \mathbf{A}^*(\mathbf{s}_i^* + \widehat{\mathbf{s}}_i) + \mathbf{e}_i^* + \widehat{\mathbf{e}}_i$, for $\widehat{\mathbf{b}}_i = \mathbf{A}^*\widehat{\mathbf{s}}_i + \widehat{\mathbf{e}}_i$. As a basic consistency check, note that in the real world, $\mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i) = \mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^*$, which matches with the secret $\mathbf{s}_i^*$, provided to the adversary. In the simulated world, $\mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i) = \mathbf{A}^*(\mathbf{s}_i^* + \widehat{\mathbf{s}}_i) + \mathbf{e}_i^* + \widehat{\mathbf{e}}_i$, which is again consistent with the secret $\mathbf{s}_i^* + \widehat{\mathbf{s}}_i$ that the adversary receives in this world.

For the proof, we first rely on the LWE assumption to switch from a uniformly random crs in the real world to LWE samples $\{\mathbf{A}\widehat{\mathbf{s}}_i + \widehat{\mathbf{e}}_i\}_i$ in the simulated world. Next, we perform a hybrid argument over $\mathsf{ind} = 1$ to $\mathsf{ind} = K$. In the ind-th hybrid, we first replace the ciphertext components in pub corresponding to $i = \mathsf{ind}$ by uniformly random values, and then switch them back to honestly generated ciphertexts in the simulated world, where $\mathsf{hct}_{\mathsf{ind}}$ now encrypts $(\mathsf{sd}_{\mathsf{ind}}, \mathsf{ind}, 1)$ instead of $(\mathsf{sd}_{\mathsf{ind}}, \mathsf{ind}, 0)$ in the real world. To argue that the resulting distributions remain indistinguishable, we rely on our assumption of *LWE with structured preimage*, which guarantees that the IPFE ciphertext hiding $\mathbf{s}_{\mathsf{ind}}^*$, as well as the other ciphertext components hiding the PRF seed $\mathsf{sd}_{\mathsf{ind}}$ used to compute $\mathbf{e}_i^*$, remain pseudorandom even given the key sk which is a short preimage for the (structured matrix) $\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$ with respect to the LWE matrix $\mathbf{B}$. In the ideal world sk is additionally constrained to ensure that the relation between the (random) ciphertexts corresponding to ind and the key sk is preserved in both the worlds. We refer the reader to Section 5.2 for more details.

2.2 Comparison with AKY.

Below, we highlight the differences from the AKY construction and proof.

On the usage of IPFE. A crucial difference between our OLS construction and the PRFE construction by AKY is the usage of IPFE. The reason we use IPFE is subtle – recent attacks on PRFE [HJL25, AMYY25] have cleverly correlated the error term resulting from FHE evaluation to the PRF value that seeks to flood it. Yet, while these attacks break the assumption, some schemes based on these assumptions [HLL23, CW25] have remained unbroken, despite serious effort. Taking a closer look, it is because the *schemes* in these works do not allow the adversary control on how the LWE mask is generated – due to this, the pseudorandom LWE mask adds uncontrollable noise to the decryption equation which prevents the attacks from working. Our usage of IPFE creates a similar effect – the LWE mask in the generated sample is computed using a fixed procedure that cannot be manipulated by the adversary. This adds an extra layer of security to the scheme.

Randomization of Error. In our construction, we add extra random matrices, $\mathbf{W}_{\mathsf{high}}$ and $\mathbf{W}_{\mathsf{low}}$ to $\mathbf{A}_{\mathsf{high}}$ and $\mathbf{A}_{\mathsf{low}}$, respectively, in the computation of $\mathbf{A}_{\mathsf{err}}$ to randomize them. It ensures that certain error terms in the computation of $\left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{high}}}{\Delta} \right\rceil$ (resp. $\left\lfloor \frac{\mathsf{bct}_i \mathbf{H}_{\mathsf{hct}}^{\mathsf{low}}}{\Delta} \right\rceil$), introduced due to rounding, are zero with $1 - \mathsf{negl}(\lambda)$ probability.

Differences in Security. The PRFE and OLS primitives are very different especially in the security definition. A crucial difference is that in the PRFE prCT security model, ciphertexts are required to be indistinguishable from random only when the inputs $\{\mathbf{s}_i^*, \mathsf{sd}_i^*\}_i$ are hidden. In contrast, in the OLS security experiment the adversary is given $\{\mathbf{s}_i^*\}_i$ (or, $\{\mathbf{s}_i^* + \widehat{\mathbf{s}}_i\}_i$ in the ideal world). This implies that in the OLS game, the adversary can compute the mask $\mathbf{A}^*\mathbf{s}^*$ in the clear, thus legitimately obtaining a low norm error term – this places it firmly in the zeroizing regime. On the other hand, in the PRFE security game, the output of the functionality is pseudorandom and $\mathbf{s}^*$ remains hidden, so the adversary only sees pseudorandom values.

Due to the multi-challenge flavour of the evasive LWE assumption and the fact that the output of PRFE is pseudorandom, the proof follows relatively easily by plugging in the assumption to replace the ciphertexts by random. On the other hand, the LSP assumption is single challenge and there is no guarantee of pseudorandomness of functionality. Therefore, the proof must proceed via a sequence of careful hybrids to establish multi-challenge security of OLS. Please see Section 5 for details.

2.3 Cryptanalysis of LSP Assumption

In our assumption (Table 1), $\mathbf{P} = \mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$ where the latter is the IPFE matrix and the former is the PRFE matrix for computing error together with the "randomization procedure" discussed above. $\mathbf{A}_{\mathsf{err}}$ depends on the public matrix $\mathbf{A}_{\mathsf{att}}$ and the (publicly computable) matrix for the homomorphic evaluation of $\mathsf{PRG}(\mathsf{sd})$, $\mathbf{A}_{\mathsf{mask}}$ depends on $\mathbf{A}_{\mathsf{IP}}$ and $\mathbf{A}^*$. In the final assumption that we describe in Section 4, we include some extra terms to support the functionality of OLS. We now proceed to discuss why we conjecture that this assumption is secure.

To begin, we argue (Section 4) that without the preimage $\mathbf{B}^{-1}(\mathbf{P})$, provided in blue in Table 1, the remaining terms can be constructed just from the LWE assumption. So the source of non-standardness of the assumption is the preimage $\mathbf{B}^{-1}(\mathbf{P})$ since $\mathbf{P}$ is not a random matrix. We reason below why inclusion of this term is safe in our setting below:

1. *Error correlating attacks:* Our assumption inherits and refines the AKY countermeasure to the error correlating attacks of [AMYY25, HJL25]. Similarly to the updated version of AKY, we use the modulus reduction technique to get rid of any harmful correlations between error term, as detailed below.

   Let us take a step back and look at the error structure of our OLS construction if modulus reduction is not used. We drop i from the subscript for clarity of presentation. After decryption, we get $\mathsf{err} = \mathbf{e}_{\mathbf{B}}^\intercal\mathbf{K} - \mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}} - \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}_{\mathsf{hct}}^{\mathrm{F}} - \mathbf{e}_{\mathsf{IP}}^\intercal\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ where $\mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}} + \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}_{\mathsf{hct}}^{\mathrm{F}}$ error comes from $\mathsf{bct} \cdot \mathbf{H}_{\mathsf{hct}}^{\mathrm{F}}$. Since $\mathbf{K}$ is sampled randomly from an appropriate Gaussian distribution, it is not under the control of adversary and hence, $\mathbf{e}_{\mathbf{B}}^\intercal\mathbf{K}$ cannot be manipulated. The term $\mathbf{e}_{\mathsf{IP}}^\intercal\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ is independent of F_{xx} where $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$, and therefore cannot be manipulated. The contrived implementation of circuit F can induce the correlation between $\mathsf{PRG}(\mathsf{sd})$ and $\mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}}$ such that $\mathsf{PRG}(\mathsf{sd}) + \mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}} = 0 \bmod 2$.

   However, after using modulus reduction, this specific correlation is destroyed. This is because Δ is chosen such that $||\mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}} + \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}_{\mathsf{hct}}^{\mathrm{F}}|| \ll \Delta$. Hence,

$$\left\lfloor \frac{\mathsf{bct}\mathbf{H}_{\mathsf{hct}}^{\mathsf{xx}}}{\Delta} \right\rceil = \left\lfloor \frac{\mathbf{t}^\intercal\mathbf{A}_{\mathsf{err}} - \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathsf{sd}, \cdot, 0) + \mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}_{\mathsf{hct}}^{\mathsf{xx}}}{\Delta} \right\rceil$$
$$= \left\lfloor \frac{\mathbf{t}^\intercal\mathbf{A}_{\mathsf{err}}}{\Delta} \right\rceil - \mathrm{F}_{\mathsf{xx}}(\mathsf{sd}, \cdot, 0)$$

   where the last equation holds unless $\mathbf{t}^\intercal\mathbf{A}_{\mathsf{err}}$ is too close to the rounding boundary. We ensure that the marginal distribution of $\mathbf{A}_{\mathsf{err}}$ is uniformly random, by adding a random matrix $\mathbf{W}$ to it, making

the probability of the value being too close to a rounding boundary negligible. Even without this randomization, which is the case in AKY, the above equation holds up to an error term that is at most 1. However, given that the previous error correlating attack by [HJL21] uses parity information for the attack, it is prudent to completely eliminate any correlated terms.

2. *Freedom to choose* aux*:* Prior work has exploited the freedom to choose the auxiliary information aux to mount attacks – for instance, the obfuscation based attack in [VWW22], which is later formalized in [HHY25], samples $\mathbf{P}$ along with a trapdoor, and the trapdoor is hidden within an obfuscated program that is provided as part of aux. In other works such as [BUW24] and [DJM+25], an attack is mounted via a carefully designed aux which includes specific terms that are exploited to distinguish the post condition distribution. In our assumption, there is no free aux. This leaves no space for the adversary to embed any term in the distribution that could be exploited.

3. *Private* $\mathbf{P}$*:* The obfuscation based attack on private coin evasive LWE in [VWW22, HHY25] samples $\mathbf{P}$ along with a trapdoor and the trapdoor is hidden inside an obfuscated program provided to the adversary in aux. In our case, $\mathbf{P}$ is publicly computed with no private coins, which restricts the avenues that the adversary can exploit.

4. *Incompressibility arguments:* Multiple crucial factors needed to apply an incompressibility argument, used in [BDJ+24, AMYY25] to provide incompressibility attacks, are absent in our assumption. We summarize them below:

   - To begin with, it crucially requires "multi-challenge" setting, while our assumption is in a single challenge setting. In multi-challenge setting the vector $\mathbf{t}$ in our assumption would be replaced by a matrix $\mathbf{T}$. The multi-challenge setting seems quite fundamental to launching any kind of incompressibility attack. Intuitively, if there is only one challenge ciphertext, then there is enough "space" in the reduction to program all the requisite outputs and incompressibility does not apply.
   - The argument requires flexibility in choosing the inputs or aux, where WE ciphertexts, which are the key components of attacks in this regime, can be hidden. This is the primary reason why an incompressibility argument cannot be used to break vanilla public-coin evasive LWE, even in the multi-challenge setting! In our assumption both the input and the aux are fixed. Hence, there is no space to hide WE ciphertexts.

   Please see Section 4.2 for a detailed discussion.

3 Preliminaries

In this section, we present the notation and the required preliminaries, some of which are recalled from [Agr19, AMYY25].

Notations. We use bold capital letters to denote a matrix and bold small letters to represent vectors. By default, a vector $\mathbf{v} \in \mathbb{Z}_q^a$ is a column vector and $\mathbf{v}^\intercal \in \mathbb{Z}_q^{1\times a}$ is a row vector. Throughout the paper, we use λ to denote the security parameter. For any $x \in \mathbb{Z}_q$, $\lfloor x \rceil_p \in \mathbb{Z}_p = \left\lfloor x \cdot \frac{p}{q} \right\rceil$. Sometimes, we regard an element in $x \in \mathbb{Z}_q$ as an integer in $[-q/2, q/2]$. When we take an absolute value $|x|$ for $x \in \mathbb{Z}_q$, this operation means that first regarding x as an integer and then taking the absolute value. For $w \in \mathbb{Z}_q$, $\mathbf{w} = (w, \ldots, w)$ represents the vector $(1, \ldots, 1)v$. For $\mathbf{v} \in \mathbb{Z}_q^a$, $\mathbf{V} = (\mathbf{v}, \ldots, \mathbf{v})$ represents the matrix $(1, \ldots, 1) \otimes \mathbf{w}$. For any function f, we write $|f|$ to represent its size. For any matrix $\mathbf{V}$, $\mathbf{V}[i, j]$ denotes the element at i-th row and j-th column, $\mathbf{V}[i, \cdot]$ denotes the ith row of $\mathbf{V}$ and $\mathbf{V}[\cdot, j]$ denotes the jth column

of $\mathbf{V}$. We consider the infinity norm:

$$\|\mathbf{v}\| = \max_i |\mathbf{v}[i]|, \qquad \|\mathbf{V}\| = \max_i \sum_j |\mathbf{V}[i,j]|$$

With slight overload of notation, we also define $|\mathbf{v}| = \max_i |\mathbf{v}[i]|$ and $|\mathbf{V}| = \max_{i,j} |\mathbf{V}[i,j]|$. For any function or algorithm written as $\mathsf{A}[\alpha]$, α represents the hardwired value in A. For any function $\mathrm{F}[\alpha]$ and a modulus reduction parameter Δ, $\mathrm{F}_{\mathsf{high}}[\alpha]$ and $\mathrm{F}_{\mathsf{low}}[\alpha]$ are defined to be such that $\mathrm{F}[\alpha] = \Delta \mathrm{F}_{\mathsf{high}}[\alpha] + \mathrm{F}_{\mathsf{low}}[\alpha]$.

3.1 Indistinguishability Obfuscation

A uniform P.P.T machine $i\mathcal{O}$ is an indistinguishability obfuscator for a class of circuits $\{\mathcal{C}_\lambda\}_{\lambda\in\mathbb{N}}$, if the following conditions are satisfied:

1. **Correctness.** For all security parameters $\lambda \in \mathbb{N}$, for any $C \in \mathcal{C}_\lambda$ and every input $\mathbf{x}$ from the domain of C, we have that:

$$\Pr\left[C' \leftarrow i\mathcal{O}(1^\lambda, C) : C'(\mathbf{x}) = C(\mathbf{x})\right] = 1$$

   where the probability is taken over the coin-tosses of the obfuscator $i\mathcal{O}$.

2. **Indistinguishability of Equivalent Circuits.** For every ensemble of pairs of circuits $\{C_{0,\lambda}, C_{1,\lambda}\}_{\lambda\in\mathbb{N}}$ such that $C_{0,\lambda}(\mathbf{x}) = C_{1,\lambda}(\mathbf{x})$ for all $\mathbf{x}$, we have that the following ensembles of pairs of distributions are indistinguishable to any P.P.T Adv:

$$\left\{C_{0,\lambda}, C_{1,\lambda}, i\mathcal{O}(1^\lambda, C_{0,\lambda})\right\} \stackrel{c}{\approx} \left\{C_{0,\lambda}, C_{1,\lambda}, i\mathcal{O}(1^\lambda, C_{1,\lambda})\right\}$$

3.2 Lattice Preliminaries

Here, we recall some facts on lattices that are needed for the exposition of our construction. Throughout this section, n, m, and q are integers such that $n = \mathrm{poly}(\lambda)$ and $m \geq n\lceil \log q \rceil$. In the following, we define a function ρ_σ as $\rho_\sigma = \exp(-\pi\|\mathbf{x}/\sigma\|_2^2)$, where $\|\cdot\|_2$ is the Euclidean norm. We will overload the notation to denote by $\rho_\sigma(S)$ the summation of $\rho_\sigma(\mathbf{x})$ over all $\mathbf{x} \in S$ for a countable set X. For a countable set S, $\mathcal{D}_{S,\gamma}$ refers to a distribution that outputs $\mathbf{x} \in S$ with probability $\rho_\sigma(\mathbf{x})/\rho(S)$. In the following, let $\mathsf{SampZ}(\gamma)$ be a sampling algorithm for the truncated discrete Gaussian distribution over $\mathbb{Z}$ with parameter $\gamma > 0$ whose support is restricted to $z \in \mathbb{Z}$ such that $|z| \leq \sqrt{n}\gamma$. Namely, $\mathsf{SampZ}(\gamma)$ is the truncated version of $\mathcal{D}_{\mathbb{Z},\gamma}$.

Let

$$\mathbf{g} = (2^0, 2^1, \ldots, 2^{\frac{m}{n+1}-1})^\intercal, \quad \mathbf{G} = \mathbf{I}_{n+1} \otimes \mathbf{g}^\intercal$$

be the gadget vector and the gadget matrix. For $\mathbf{p} \in \mathbb{Z}_q^n$, we write $\mathbf{G}^{-1}(\mathbf{p})$ for the m-bit vector $(\mathsf{bits}(\mathbf{p}[1]), \ldots, \mathsf{bits}(\mathbf{p}[n+1]))^\intercal$, where $\mathsf{bits}(\mathbf{p}[i])$ are $m/(n+1)$ bits for each $i \in [n+1]$. The notation extends column-wise to matrices and it holds that $\mathbf{G}\mathbf{G}^{-1}(\mathbf{P}) = \mathbf{P}$.

Trapdoors. Let us consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n\times m'}$, we let $\mathbf{A}^{-1}(\mathbf{V})$ be an output distribution of $\mathsf{SampZ}(\gamma)^{m\times m'}$ conditioned on $\mathbf{A}\cdot\mathbf{A}^{-1}(\mathbf{V},\gamma) = \mathbf{V}$. A γ-trapdoor for $\mathbf{A}$ is a trapdoor that enables one to sample from the distribution $\mathbf{A}^{-1}(\mathbf{V},\gamma)$ in time $\mathrm{poly}(n,m,m',\log q)$ for any $\mathbf{V}$. We slightly overload notation and denote a γ-trapdoor for $\mathbf{A}$ by $\mathbf{A}_\gamma^{-1}$. The following properties has been established in a long sequence of works [GPV08, CHKP10, ABB10a, ABB10b, MP12, BLP+13].

Lemma 3.1 (Properties of Trapdoors). Lattice trapdoors exhibit the following properties.

1. Given $\mathbf{A}_\tau^{-1}$, one can obtain $\mathbf{A}_{\tau'}^{-1}$ for any $\tau' \geq \tau$.
2. Given $\mathbf{A}_\tau^{-1}$, one can obtain $[\mathbf{A}\|\mathbf{B}]_\tau^{-1}$ and $[\mathbf{B}\|\mathbf{A}]_\tau^{-1}$ for any $\mathbf{B}$.
3. There exists an efficient procedure $\mathsf{TrapGen}(1^n, 1^m, q)$ that outputs $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ where $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ for some $m = O(n \log q)$ and is 2^{-n}-close to uniform, where $\tau_0 = \omega(\sqrt{n \log q \log m})$.

Useful Lemmata.

Lemma 3.2 (tail and truncation of $\mathcal{D}_{\mathbb{Z},\gamma}$). There exists $B_0 \in \Theta(\sqrt{\lambda})$ such that

$$\Pr\Big[x \leftarrow \mathcal{D}_{\mathbb{Z},\gamma} : |x| > \gamma B_0(\lambda)\Big] \leq 2^{-\lambda} \quad \text{for all} \quad \gamma \geq 1 \text{ and } \lambda \in \mathbb{N}.$$

Lemma 3.3 ([[Lyu12, Lemma 4.4]]). The following hold.

1. For any $k > 0$, $\Pr[|z| > k\sigma; z \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}] \leq 2\exp(-k^2/2)$.
2. For any $k > 1$,
$$\Pr\big[\|\mathbf{z}\| > k\sigma\sqrt{m};\ \mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma}\big] < k^m \exp(\frac{m}{2}(1-k^2)).$$

Lemma 3.4 ([[ABB10a, Lemma 2.14 (Generalized Leftover Hash Lemma)]]). Let $\mathcal{H} = \{h : \mathcal{X} \to \mathcal{Y}\}$ be a 2-universal hash function family and $f : \mathcal{X} \to \mathcal{Z}$ be a function. Then for any random variable X taking value in $\mathcal{X}$, the distributions

$$(h, h(X), f(X)) \text{ and } (h, \mathcal{U}(\mathcal{Y}), f(X))$$

are within statistical distance $\sqrt{2^{-H_\infty(X)} \cdot |\mathcal{Y}||\mathcal{Z}|}$, where $\mathcal{U}(\mathcal{Y})$ represents the uniform distribution over $\mathcal{Y}$. More generally, let $X_1, \ldots X_Q$ be independent random variables taking values in $\mathcal{X}$ and let $\gamma = \min_{i\in[Q]} H_\infty(X_i)$. Then, the distributions

$$(h, h(X_1), \ldots, h(X_Q), f(X_1), \ldots, f(X_Q)) \text{ and}$$

$$(h, \mathcal{U}_1(\mathcal{Y}), \ldots, \mathcal{U}_Q(\mathcal{Y}), f(X_1), \ldots, f(X_Q))$$

are within statistical distance $Q\sqrt{2^{-\gamma} \cdot |\mathcal{Y}||\mathcal{Z}|}$.

We note that even in the special case where there is no side information of the form $f(X)$, the above lemma can be applied by setting $f(X) = \perp$ for all X and $|\mathcal{Z}| = 1$. In this case, we call the above lemma "leftover hash lemma". We also note that the family $\{\mathbf{A} \in \mathbb{Z}_q^{n\times m} : \mathbf{r} \mapsto \mathbf{A}\mathbf{r}\}_\mathbf{A}$ is 2-universal for any prime q. Furthermore, if we restrict the input space to be $\{0,1\}^m$, the same family is 2-universal for any integer q.

Lemma 3.5 (Smudging Lemma [WWW22]). Let λ be a security parameter. Take any $a \in \mathbb{Z}$ where $|a| \leq B$. Suppose $\gamma \geq B\lambda^{\omega(1)}$. Then the statistical distance between the distributions $\{z : z \leftarrow \mathcal{D}_{\mathbb{Z},\gamma}\}$ and $\{z + a : z \leftarrow \mathcal{D}_{\mathbb{Z},\gamma}\}$ is $\mathsf{negl}(\lambda)$. This also holds when z is sampled from a uniform distribution over $\mathbb{Z}_q$.

3.2.1 Hardness Assumptions

Assumption 3.6 (The LWE Assumption). Let $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda) > 2$ be integers and $\chi = \chi(\lambda)$ be a distribution over $\mathbb{Z}_q$. We say that the $\mathrm{LWE}(n, m, q, \chi)$ hardness assumption holds if for any PPT adversary $\mathcal{A}$ we have

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{s}^\intercal \mathbf{A} + \mathbf{e}^\intercal) \to 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{v}^\intercal) \to 1]| \leq \mathsf{negl}(\lambda)$$

where the probability is taken over the choice of the random coins by the adversary $\mathcal{A}$ and $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{v} \leftarrow \mathbb{Z}_q^m$. We also say that $\mathrm{LWE}(n, m, q, \chi)$ problem is subexponentially hard if the above probability is bounded by $2^{-n^\epsilon} \cdot \mathsf{negl}(\lambda)$ for some constant $0 < \epsilon < 1$ for all PPT $\mathcal{A}$.

As shown by previous works [Reg09, BLP+13], if we set $\chi = \mathsf{SampZ}(\gamma)$, the $\mathrm{LWE}(n, m, q, \chi)$ problem is as hard as solving worst case lattice problems such as gapSVP and SIVP with approximation factor $\mathrm{poly}(n) \cdot (q/\gamma)$ for some $\mathrm{poly}(n)$. Since the best known algorithms for 2^k-approximation of gapSVP and SIVP run in time $2^{\tilde{O}(n/k)}$, it follows that the above $\mathrm{LWE}(n, m, q, \chi)$ with noise-to-modulus ratio 2^{-n^ϵ} is likely to be (subexponentially) hard for some constant ϵ.

Theorem 3.7 ([Ajt96, MP12]). There exists a PPT algorithm $(\mathbf{A}, \mathbf{A}_\tau^{-1}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$ and a deterministic polynomial time algorithm $\mathsf{LWESolve}_{\mathbf{A}_\tau^{-1}}(\mathbf{b})$ such that the following holds for any $n \geq 1, q \geq 2$, and a sufficiently large $m = O(n \log q)$:

- The statistical distance between $(\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m})$ and $(\mathbf{A} : (\mathbf{A}, \mathbf{A}_\tau^{-1}) \leftarrow \mathsf{TrapGen}(1^n, 1^m, q)$ is negligible in λ.
- For any $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$ where $||\mathbf{e}|| \leq q/O(\sqrt{nm \log q})$, we have $\mathsf{LWESolve}_{\mathbf{A}_\tau^{-1}}(\mathbf{b}) = \mathbf{s}$.

3.3 GSW Homomorphic Encryption and Evaluation

We recall the format of the (leveled fully) homomorphic encryption due to [GSW13] and the correctness property. We adapt the syntax from [HLL23].

Lemma 3.8. The leveled FHE scheme works as follows:

- The keys are

$$\text{(public)} \quad \mathbf{A}_{\mathsf{fhe}} = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{s}}^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^\intercal \end{pmatrix} \in \mathbb{Z}_q^{(n+1) \times m}, \quad \text{(secret)} \quad \mathbf{s}^\intercal = (\bar{\mathbf{s}}^\intercal, -1),$$

  where $\bar{\mathbf{s}} \in \mathbb{Z}^n$,$\bar{\mathbf{A}}_{\mathsf{fhe}} \in \mathbb{Z}_q^{n \times m}$, and $\mathbf{e}_{\mathsf{fhe}}^\intercal \in \mathbb{Z}^m$.

- A ciphertext of $x \in \{0, 1\}$ is $\mathbf{X} = \mathbf{A}_{\mathsf{fhe}}\mathbf{R} - x\mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m}$, where $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$ is the encryption randomness. The decryption equation is

$$\mathbf{s}^\intercal \mathbf{X} = -\mathbf{e}_{\mathsf{fhe}}^\intercal \mathbf{R} - x\mathbf{s}^\intercal \mathbf{G} \in \mathbb{Z}_q^m,$$

  which can be used to extract x via multiplication by $\mathbf{G}^{-1}(\lfloor q/2 \rfloor \iota_{n+1})$.

Lemma 3.9. (homomorphic evaluation for vector-valued functions [HLL23]) For the scheme in Lemma 4, there is an efficient algorithm

$$\mathrm{MakeVEvalCkt}(n, m, q, C) = \mathsf{VEval}_C$$

that takes as input n, m, q and a vector-valued circuit $C : \{0, 1\}^L \to \mathbb{Z}_q^{1 \times m'}$ and outputs a circuit

$$\mathsf{VEval}_C(\mathbf{X}_1, ..., \mathbf{X}_L) = \mathbf{C},$$

taking L ciphertexts as input and outputting a new ciphertext $\mathbf{C}$ of different format.

- The depth of VEval_C is $dO(\log m \log\log q) + O(\log^2 \log q)$ for C of depth d.
- Suppose $\mathbf{X}_\ell = \mathbf{A}_{\mathsf{fhe}}\mathbf{R}_\ell - \mathbf{x}[\ell]\mathbf{G}$ for $\ell \in [L]$ with $\mathbf{x} \in \{0,1\}^L$, then

$$\mathbf{C} = \mathbf{A}_{\mathsf{fhe}}\mathbf{R}_C - \begin{pmatrix} 0^{n\times m'} \\ C(\mathbf{x}) \end{pmatrix} \in \mathbb{Z}_q^{(n+1)\times m'},$$

  where $\left\|\mathbf{R}_C^\intercal\right\| \le (m+2)^d \lceil \log q \rceil \max_{\ell\in[L]} \left\|\mathbf{R}_\ell^\intercal\right\|$. The new decryption equation is

$$\mathbf{s}^\intercal \mathbf{C} = -\mathbf{e}_{\mathsf{fhe}}^\intercal \mathbf{R}_C + C(\mathbf{x}) \in \mathbb{Z}_q^{1\times m'}.$$

3.4 Homomorphic Evaluation Procedures

In this section we describe the properties of the attribute encoding and its homomorphic evaluation. We adapt the syntax from [HLL23].

- For L-bit input, the public parameter is $\mathbf{A}_{\mathsf{att}} \in \mathbb{Z}_q^{(n+1)\times(L+1)m}$.
- The encoding of $\mathbf{x} \in \{0,1\}^L$ is

$$\mathbf{s}^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathbf{x}^\intercal) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal,$$

  where $\mathbf{s}^\intercal = (\bar{\mathbf{s}}^\intercal, -1)$ with $\bar{\mathbf{s}} \in \mathbb{Z}^n$ and $\mathbf{e}_{\mathsf{att}}^\intercal \in \mathbb{Z}^{(L+1)m}$.
- There are efficient deterministic algorithms [BTVW17]

$$\mathsf{MEvalC}(\mathbf{A}_{\mathsf{att}}, C) = \mathbf{H}^C_{\mathbf{A}_{\mathsf{att}}} \quad \text{and} \quad \mathsf{MEvalCX}(\mathbf{A}_{\mathsf{att}}, C, \mathbf{x}) = \mathbf{H}^C_{\mathbf{A}_{\mathsf{att}},x}$$

  that take as input $\mathbf{A}_{\mathsf{att}}$, a matrix-valued circuit $C : \{0,1\}^L \to \mathbb{Z}_q^{n+1\times m'}$, and (for $\mathsf{MEvalCX}$) some $\mathbf{x} \in \{0,1\}^L$, and output some matrix in $\mathbb{Z}^{(L+1)m\times m'}$.
  - Suppose C is of depth d, then $\left\|(\mathbf{H}^C_{\mathbf{A}_{\mathsf{att}}})^\intercal\right\|, \left\|(\mathbf{H}^C_{\mathbf{A}_{\mathsf{att}},\mathbf{x}})^\intercal\right\| \le (m+2)^d \lceil \log q \rceil$.
  - The matrix encoding homomorphism is $(\mathbf{A}_{\mathsf{att}} - (1,\mathbf{x}^\intercal)\otimes\mathbf{G})\mathbf{H}^C_{\mathbf{A}_{\mathsf{att}},\mathbf{x}} = \mathbf{A}_{\mathsf{att}}\mathbf{H}^C_{\mathbf{A}_{\mathsf{att}}} - C(\mathbf{x})$.

Dual-Use Technique and Extension. In [BTVW17], the attribute encoded with secret $\mathbf{s}^\intercal$ is FHE ciphertexts under key $\mathbf{s}^\intercal$ (the same, "dual-use") and the circuit being $\mathsf{MEvalCX}$'ed is some HEval_C. This leads to automatic decryption. Let C be a vector-valued circuit, with co-domain $\mathbb{Z}_q^{1\times m'}$, then VEval_C is $\mathbb{Z}_q^{(n+1)\times m'}$-valued and

$$\begin{aligned} & (\mathbf{s}^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathbf{X})) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal) \cdot \mathbf{H}^{\mathsf{VEval}_C}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} \\ (\mathsf{MEvalCX}) \quad & = \mathbf{s}^\intercal \mathbf{A}_{\mathsf{att}} \mathbf{H}^{\mathsf{VEval}_C}_{\mathbf{A}_{\mathsf{att}}} - \mathbf{s}^\intercal \mathsf{VEval}_C(\mathbf{X}) + (\mathbf{e}')^\intercal \\ (\mathsf{VEval\ decryption}) \quad & = \mathbf{s}^\intercal \mathbf{A}_{\mathsf{att}} \mathbf{H}^{\mathsf{VEval}_C}_{\mathbf{A}_{\mathsf{att}}} - C(\mathbf{x}) + (\mathbf{e}'')^\intercal. \end{aligned}$$

4 LWE with Structured Preimage

In this section, we describe our main new assumption and provide a thorough analysis of its security.

Real world	**Ideal World**
params, $\bar{\mathbf{A}}_{\mathsf{fhe}}, \mathbf{B}, \mathbf{D}, \mathbf{A}^*, \widehat{\mathbf{E}}, \widehat{\mathbf{S}}$	params, $\bar{\mathbf{A}}_{\mathsf{fhe}}, \mathbf{B}, \mathbf{D}, \mathbf{A}^*, \widehat{\mathbf{E}}, \widehat{\mathbf{S}}$
$\mathbf{c}_{\mathbf{B}}^{\intercal} = \mathbf{t}^{\intercal}\mathbf{B} + \mathbf{e}_{\mathbf{B}}^{\intercal}$	$\mathbf{c}_{\mathbf{B}}^{\intercal} \leftarrow \mathbb{Z}_q^{1\times m\lceil \log q\rceil}$
$\mathbf{X} = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}^{\intercal}\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^{\intercal} \end{pmatrix} \mathbf{R} - (\mathsf{sd}, \mathsf{ind}, \beta) \otimes \mathbf{G}$	$\mathbf{X} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}$
$\mathbf{c}_{\mathbf{D}}^{\intercal} = \mathbf{t}^{\intercal}(\mathbf{D} - (((1, \mathsf{bits}(\mathbf{X})) \otimes \mathbf{G}), \mathbf{0})) + \mathbf{e}_{\mathbf{D}}^{\intercal}$	$\mathbf{c}_{\mathbf{D}}^{\intercal} \leftarrow \mathbb{Z}_q^{1\times (L_{\mathbf{X}}+1)m+W\lceil \log q\rceil+2}$
$\mathbf{K} = \mathbf{B}_{\gamma}^{-1}(\mathbf{P})$	$\mathbf{K} = \begin{pmatrix} \mathbf{B} \\ \mathbf{c}_{\mathbf{B}}^{\intercal} \end{pmatrix}_{\gamma}^{-1} \left(\begin{pmatrix} \mathbf{P} \\ \mathbf{d}^{\intercal} + \mathbf{c}'^{\intercal} \end{pmatrix} \right)$

where,
$\bar{\mathbf{A}}_{\mathsf{fhe}} \leftarrow \mathbb{Z}_q^{n\times m}$, $\mathbf{D} \leftarrow \mathbb{Z}_q^{(n+1)\times (L_{\mathbf{X}}+1)m+W\lceil \log q\rceil+2}$, $\mathbf{R} \leftarrow \{0,1\}^{m\times Lm}$,
$(\mathbf{B}, \mathbf{B}_{\gamma}^{-1}) \leftarrow \mathsf{TrapGen}(1^{n+1}, 1^{m\lceil \log q\rceil}, q)$,$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\mathsf{crs} = \mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}$,

$m = (n+1)\lceil \log q\rceil$, $L = (\lambda + \lceil \log K\rceil + 1)$, $L_{\mathbf{x}} = (n+1)mL\lceil \log q\rceil$,

$\mathbf{t}^{\intercal} = (\bar{\mathbf{t}}^{\intercal}, -1)$, $\bar{\mathbf{t}} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathbf{t}}}^{n}$, $\mathsf{sd} \leftarrow \{0,1\}^{\lambda}$, $\widehat{\mathbf{S}} = (\widehat{\mathbf{s}}_1, \ldots, \widehat{\mathbf{s}}_K) \leftarrow \mathbb{Z}_q^{W\times K}$, $\widehat{\mathbf{E}} \leftarrow \mathcal{D}_{\mathbb{Z},\widehat{\sigma}}^{M\times K}$,
$\mathbf{e}_{\mathbf{B}} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathbf{B}}}^{m\lceil \log q\rceil}$, $\mathbf{e}_{\mathsf{fhe}} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathsf{fhe}}}^{m}$,
$\mathbf{e}_{\mathbf{D}}^{\intercal} = (\mathbf{e}_{\mathsf{att}}^{\intercal}, \mathbf{e}_{\mathsf{IP}}^{\intercal}, e_{\mathsf{high}}, e_{\mathsf{low}}) \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{1\times (L_{\mathbf{X}}+1)m} \times \mathcal{D}_{\mathbb{Z},\sigma}^{1\times W\lceil \log q\rceil} \times \mathcal{D}_{\mathbb{Z},\sigma} \times \mathcal{D}_{\mathbb{Z},\sigma}$

$\mathbf{D} = (\mathbf{A}_{\mathsf{att}}, \mathbf{A}_{\mathsf{IP}}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}) \in \mathbb{Z}_q^{(n+1)\times (L_{\mathbf{X}}+1)m} \times \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil} \times \mathbb{Z}_q^{n+1} \times \mathbb{Z}_q^{n+1}$,
$\mathbf{c}_{\mathbf{D}}^{\intercal} = (\mathbf{c}_{\mathsf{att}}^{\intercal}, \mathbf{c}_{\mathsf{IP}}^{\intercal}, c_{\mathsf{high}}, c_{\mathsf{low}}) \in \mathbb{Z}_q^{1\times (L_{\mathbf{X}}+1)m} \times \mathbb{Z}_q^{1\times W\lceil \log q\rceil} \times \mathbb{Z}_q \times \mathbb{Z}_q$,

$\mathbf{c}' \leftarrow [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$,
$\mathbf{d}^{\intercal} = \Delta \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}^{\intercal}\mathbf{H}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}^{\mathsf{F}_{\mathsf{high}}} + \mathbf{c}_{\mathsf{high}}^{\intercal}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}^{\intercal}\mathbf{H}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}^{\mathsf{F}_{\mathsf{low}}} + \mathbf{c}_{\mathsf{low}}^{\intercal}}{\Delta} \right\rfloor + \mathbf{c}_{\mathsf{IP}}^{\intercal}\mathbf{G}^{-1}((\mathbf{A}^*)^{\intercal}) + \beta \cdot \mathsf{crs}_{\mathsf{ind}}$,
$\mathbf{P} = \underbrace{\Delta \left\lfloor \frac{\mathbf{A}_{\mathsf{att}}\mathbf{H}_{\mathbf{A}_{\mathsf{att}}}^{\mathsf{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathsf{att}}\mathbf{H}_{\mathbf{A}_{\mathsf{att}}}^{\mathsf{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor}_{\mathbf{A}_{\mathsf{err}}} + \underbrace{\mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}((\mathbf{A}^*)^{\intercal})}_{\mathbf{A}_{\mathsf{mask}}}$,
where for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$
$\mathbf{c}_{\mathsf{xx}}^{\intercal} = c_{\mathsf{xx}}(1, \ldots, 1)$, $\mathbf{W}_{\mathsf{xx}} = (1, \ldots, 1) \otimes \mathbf{w}_{\mathsf{xx}}$, $\mathbf{H}_{\mathbf{A}_{\mathsf{att}}}^{\mathsf{F}_{\mathsf{xx}}} = \mathsf{MEvalC}(\mathbf{A}_{\mathsf{att}}, \mathsf{VEval}_{\mathsf{F}_{\mathsf{xx}}})$,
$\mathbf{H}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}}^{\mathsf{F}_{\mathsf{xx}}} = \mathsf{MEvalCX}(\mathbf{A}_{\mathsf{att}}, \mathsf{VEval}_{\mathsf{F}_{\mathsf{xx}}}, \mathbf{X})$, $\mathsf{VEval}_{\mathsf{F}_{\mathsf{xx}}} = \mathrm{MakeVEvalCkt}(n, m, q, \Delta \cdot \mathsf{F}_{\mathsf{xx}})$,

$\mathsf{F}[\mathsf{crs}](\mathsf{sd}, \mathsf{ind}, \beta) = \mathsf{PRG}(\mathsf{sd}) + \beta \cdot \mathsf{crs}_{\mathsf{ind}}^{\intercal}$,
$\mathsf{PRG} : \{0,1\}^{\lambda} \to [-B_{\mathsf{OLS}} + \eta\ ,\ B_{\mathsf{OLS}} - \eta]^{1\times M}$; $B_{\mathsf{OLS}} \gg \widehat{\sigma}, \eta$

Table 2: LWE with structured preimage assumption

4.1 The Assumption

Assumption 4.1 (LWE with Structured Preimage). Let $\mathsf{params} = (\lambda, \Delta, q, M, K, W, n, \gamma, \sigma_{\mathbf{t}}, \sigma, \sigma_B, \sigma_{\mathsf{fhe}}, B_{\mathsf{OLS}}, \widehat{\sigma}, \eta)$ be parameters and $\mathsf{PRG} : \{0,1\}^\lambda \to [-B_{\mathsf{OLS}} + \eta, B_{\mathsf{OLS}} - \eta]^{1\times M}$ be a pseudorandom generator. We define constraints between parameters and a concrete instantiation in Section 5.2 on Page 31. We say that the (sub-exponential) LWE with Structured Preimage (LSP) holds if for all $\beta \in \{0,1\}$ and all $\mathsf{ind} \in [K]$, the real world and the ideal world in Table 2 are (sub-exponentially) indistinguishable.

Discussion: Terms in LSP Assumption. Below, we describe the role played by the various components in the assumption. We recommend the reader refresh the high-level OLS construction outlined in Section 2 since the assumption is necessarily designed to be useful for the OLS construction (detailed in Section 5).

Public Random Matrices.

- $\mathbf{A}^*$ is the public random matrix of the LWE sample $\mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$ which will be output by OLS sampler.
- $\widehat{\mathbf{S}}$ and $\widehat{\mathbf{E}}$ are parts of the OLS security definition. Hence, these are included here to support the OLS definition.
- $\bar{\mathbf{A}}_{\mathsf{fhe}}, \mathbf{D}, \mathbf{B}$ are used in the construction of our OLS as LWE matrices used in different LWE samples/encodings.

Structured Components. Now we describe the structured components, $\mathbf{P}$ and $\mathbf{d}$ of our assumption. We begin with the simpler one.

- Vector $\mathbf{d}$: We choose the encodings in the ideal world to be pseudorandom subject to the correctness of the "decryption relation" which must be satisfied in the real world, between the encodings $(\mathbf{c_B}, \mathbf{c_D}, \mathbf{X})$ and the preimage $\mathbf{K}$. Thus, the component $\mathbf{d}$ is chosen to ensure that the correctness of decryption is preserved in the ideal world.
- The matrix $\mathbf{P}$ is the most complicated and most carefully designed component of our assumption. Prior attacks on evasive LWE (described at length below) have shown that allowing complete freedom in the choice of $\mathbf{P}$ can lead to contrived counterexamples. Therefore, we craft it conscientiously to satisfy the desired structure but while incorporating several "safety" measures to prevent exploitation, even if the PRG circuit is chosen maliciously by the adversary. In detail:
  - $\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$: The first component corresponds to the computation of the error matrix $\mathbf{E}^*$ which is generated using a PRG and the second component corresponds to the computation of the LWE mask $\mathbf{A}^*\mathbf{S}^*$. Owing to the linearity of $\mathbf{A}^*\mathbf{S}^*$, it is computed using an inner-product functional encryption – as discussed in Section 2 this is an additional safety measure which restricts the control the adversary has on creating correlations in error terms. The two masks are added together so that the individual decryption of either of them is not possible.
  - *The modulus reduction* is performed by splitting the function F for computing PRG as $\mathrm{F}_{\mathsf{high}}$ and $\mathrm{F}_{\mathsf{low}}$, where the former computes the higher order bits and the latter computes the lower order bits of the PRG. By taking Δ exponentially bigger than the noise terms in the decryption, the modulus reduction ensures that *any* circuit is safe by throwing away the problematic noise terms in the decryption by rounding.

– *Addition of matrix* $\mathbf{W}$*:* To ensure that each column of $\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{W}_{\mathsf{xx}}$ is marginally random so that $(\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{W}_{\mathsf{xx}}) \bmod \Delta$ is not concentrated near the boundaries. This ensures that the noise terms are removed with overwhelming probability. Note that all the columns in $\mathbf{W}_{\mathsf{xx}} = (1, \ldots, 1) \otimes \mathbf{w}_{\mathsf{xx}}$ are the same - this is important to keep the encodings succinct, since corresponding to these terms, we need to provide additional encodings as $\mathbf{t}^\intercal\mathbf{W}_{\mathsf{xx}}$. By repeating the columns, it suffices to only provide $\mathbf{t}^\intercal\mathbf{w}_{\mathsf{xx}}$. We emphasize that repeating the columns is safe because we need only marginal, not joint randomness for our argument.

Remark 4.2 (Circuit representation of function F.). For every circuit C_{xx} computing F_{xx} (with desired depth), we define the corresponding LSP distributions using C_{xx} and conjecture their indistinguishability. Thus, the circuit may be any implementation, including a maliciously crafted one, but is fixed when defining the distributions. Our modulus reduction, IP-based computation of $\mathbf{A}^*\mathbf{S}^*$, and randomization by $\mathbf{W}$ are designed to protect against such implementations. In the security proof of our OLS construction, we only require LSP for the implementation it uses.

Encodings. The encodings $\mathbf{c_B}$, $\mathbf{X}$ and $\mathbf{c_D}$ are LWE samples which can be constructed using plain LWE, if the preimage is withheld. We describe this in detail below.

- $\mathbf{X}$ is a GSW homomorphic encryption of the PRG seed sd. Additionally, it also encodes $\mathsf{ind} \in [K]$ and β that appear because of the way OLS security is defined. At a high level, $\beta = 0$ corresponds to the real world and $\beta = 1$ corresponds to the simulated world of OLS security game. The index ind is added to support iteration over the K columns of $\mathbf{B}^*$.
- $\mathbf{c_D}$ serves multiple purposes:
  - the first part encodes $\mathbf{X}$, such that along with $\mathbf{c_B}$, it enables "automatic decryption" of homomorphically evaluated encoding of $\mathsf{PRG}(\mathsf{sd})$ from $\mathbf{X}$.
  - The second part provides a (pseudo)random mask for encoding $\mathbf{S}^*$. We remark that $\mathbf{S}^*$ does not form a part of the assumption since in IPFE encoding the secret vector is directly added to the LWE sample. The assumption only says that this mask is pseudorandom.
  - Finally, the last two components are provided to remove the extra mask introduced in the decryption equation due to the randomizing vectors $\mathbf{W}_{\mathsf{high}}$ and $\mathbf{W}_{\mathsf{low}}$.
- $\mathbf{c_B}$ serves to couple with the key $\mathbf{K}$ to enable decryption.

The Preimage. The main source of non-standardness is the preimage $\mathbf{K}$ for a non-uniform, but carefully chosen and publicly computable matrix $\mathbf{P}$ is provided with respect to $\mathbf{B}$. In the following section, we thoroughly analyze why we believe it is safe to provide this preimage.

Falsifiability of LSP. We observe that all components in both distributions are efficiently sampleable. In particular, in the real world, $\mathbf{B}$ can be sampled together with its trapdoor, enabling efficient sampling of $\mathbf{K}$. In the ideal world, $\mathbf{B}, \mathbf{c_B}$ can similarly be sampled together, along with a trapdoor that enables efficient sampling of $\mathbf{K}$. By the properties of $\mathsf{TrapGen}$, the resulting $\mathbf{c_B}$ is statistically close to uniform.

4.2 Cryptanalysis

It's just LWE if preimage is withheld. As a starting point, we observe that the pseudorandomness of the encodings follows directly from LWE in the absence of the preimage $\mathbf{K}$. We provide the proof sketch next.

H_0 : The real world.

$\mathsf{H}_1 : \mathbf{c}_\mathbf{B}^\intercal \leftarrow \mathbb{Z}_q^{1\times mw}; \quad \mathbf{X} = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \mathbf{a}_{\mathsf{fhe}}^\intercal \leftarrow \mathbb{Z}_q^{1\times m} \end{pmatrix} \mathbf{R} - (\mathsf{sd}, \mathsf{ind}, \beta) \otimes \mathbf{G};$
$\mathbf{c}_\mathbf{D}^\intercal \leftarrow \mathbb{Z}_q^{1\times (L_\mathbf{X}+1)m+W\lceil \log q\rceil +2}.$

H_2 : Same as H_1 except that $\mathbf{X} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}$. This is the ideal world.

$\mathsf{H}_0 \approx_c \mathsf{H}_1$ follows from LWE as follows. Given an LWE challenge, $\mathbf{M}$, $\mathbf{b}_\mathbf{M}^\intercal = \bar{\mathbf{t}}^\intercal \mathbf{M} + \mathbf{e}_\mathbf{M}^\intercal$ or random, where $\mathbf{M} \leftarrow \mathbb{Z}_q^{n\times mw+m+(L_\mathbf{X}+1)m+W\lceil \log q\rceil +2}$ and $\mathbf{e}_\mathbf{M} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_\mathbf{M}}^{mw+m+(L_\mathbf{X}+1)m+W\lceil \log q\rceil +2}$[5], $w = \lceil \log q \rceil$, the reduction can simulate the encodings as follows: parse $\mathbf{M}$ as $(\mathbf{M}_\mathbf{B}, \mathbf{M}_{\mathsf{fhe}}, \mathbf{M}_\mathbf{D})$ and $\mathbf{b}_\mathbf{M}$ as $(\mathbf{b}_\mathbf{B}, \mathbf{b}_{\mathsf{fhe}}, \mathbf{b}_\mathbf{D})$. Sample $\mathbf{v}_\mathbf{B} \leftarrow \mathbb{Z}_q^{mw}, \mathbf{v}_\mathbf{D} \leftarrow \mathbb{Z}_q^{(L_\mathbf{X}+1)m+W\lceil \log q\rceil +2}$. Set $\mathbf{B} = \begin{pmatrix} \mathbf{M}_\mathbf{B} \\ \mathbf{v}_\mathbf{B}^\intercal \end{pmatrix}$; $\mathbf{c}_\mathbf{B}^\intercal = \mathbf{b}_\mathbf{B}^\intercal - \mathbf{v}_\mathbf{B}^\intercal$; $\bar{\mathbf{A}}_{\mathsf{fhe}} = \mathbf{M}_{\mathsf{fhe}}$, $\mathbf{X} = \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \mathbf{b}_{\mathsf{fhe}}^\intercal \end{pmatrix} \mathbf{R} - (\mathsf{sd}, \mathsf{ind}, \beta) \otimes \mathbf{G}$; $\mathbf{D} = \begin{pmatrix} \mathbf{M}_\mathbf{D} + (((1, \mathsf{bits}(\mathbf{X})) \otimes \overline{\mathbf{G}}), \mathbf{0}) \\ \mathbf{v}_\mathbf{D}^\intercal \end{pmatrix}$ and $\mathbf{c}_\mathbf{D}^\intercal = \mathbf{b}_\mathbf{D}^\intercal - \mathbf{v}_\mathbf{D}^\intercal - (1, \mathsf{bits}(\mathbf{X}) \otimes \underline{\mathbf{G}}, \mathbf{0})$. Here $\overline{\mathbf{G}}$ represents 1st n rows of $\mathbf{G}$ and $\underline{\mathbf{G}}$ represents the last row of $\mathbf{G}$. Clearly if $\mathbf{b}_\mathbf{M}$ is an LWE sample, then the reduction simulates H_0, else H_1. Finally, $\mathsf{H}_1 \approx_s \mathsf{H}_2$ from the Leftover Hash Lemma (Lemma 3.4) which implies that $\begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \mathbf{a}_{\mathsf{fhe}}^\intercal \end{pmatrix} \mathbf{R}$, and hence $\begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \mathbf{a}_{\mathsf{fhe}}^\intercal \end{pmatrix} \mathbf{R} - (\mathsf{sd}, \mathsf{ind}, \beta) \otimes \mathbf{G}$ is statistically close to a random distribution.

Adding the Opening. To deal with the opening $\mathbf{K}$, we first perform the basic sanity check that the legitimate "semi-honest" combination of $\mathbf{K}$ with ciphertext terms $(\mathbf{c}_\mathbf{B}, \mathbf{X}, \mathbf{c}_\mathbf{D})$ is indistinguishable in the two worlds.

Firstly we note that $\mathbf{K}$ in both the worlds satisfy $\mathbf{BK} = \mathbf{P} = \mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$. Next, we analyze the relation between $(\mathbf{c}_\mathbf{B}, \mathbf{X}, \mathbf{c}_\mathbf{D})$ and $\mathbf{K}$. In the real world, $(\mathbf{c}_\mathbf{B}, \mathbf{X}, \mathbf{c}_\mathbf{D})$ and $\mathbf{K}$ are related as

$$
\begin{aligned}
& \mathbf{c}_\mathbf{B}{}^\intercal \mathbf{K} - \left(\Delta \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}{}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} + \mathbf{c}_{\mathsf{high}}^\intercal}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\mathsf{att}}{}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} + \mathbf{c}_{\mathsf{low}}^\intercal}{\Delta} \right\rfloor + \mathbf{c}_{\mathsf{IP}}{}^\intercal \mathbf{G}^{-1}(\mathbf{A}^*) \right) \\
= \; & (\mathbf{t}^\intercal \mathbf{B} + \mathbf{e}_\mathbf{B}^\intercal)\mathbf{B}^{-1}(\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}) - \left(\mathbf{t}^\intercal \left(\Delta \cdot \left\lfloor \frac{\mathbf{A}_{\mathsf{att}} \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathsf{att}} \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathbf{X}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor \right) \right. \\
& \left. - \Delta \mathrm{F}_{\mathsf{high}}(\mathsf{sd}, \mathsf{ind}, \beta) - \mathrm{F}_{\mathsf{low}}(\mathsf{sd}, \mathsf{ind}, \beta) + (\mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}})\mathbf{G}^{-1}(\mathbf{A}^*) + \Delta \mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{high}}}^\intercal + \Delta \cdot \tilde{\mathbf{e}}_{\mathsf{high}}^\intercal + \mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{low}}}^\intercal + \tilde{\mathbf{e}}_{\mathsf{low}}^\intercal \right) \\
= \; & \mathbf{t}^\intercal(\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}) - \mathbf{t}^\intercal(\mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}) + \mathrm{F}(\mathsf{sd}, \mathsf{ind}, \beta) + \widetilde{\mathbf{e}}^\intercal \\
= \; & \mathsf{PRG}(\mathsf{sd}) + \beta \cdot \mathsf{crs}_{\mathsf{ind}} + \widetilde{\mathbf{e}}^\intercal
\end{aligned} \tag{3}
$$

In the above for $\mathsf{xx} \in \{\mathsf{low}, \mathsf{high}\}$,

$$
\widetilde{\mathbf{e}}^\intercal = \mathbf{e}_\mathbf{B}^\intercal \mathbf{K} - \mathbf{e}_{\mathsf{IP}} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) - \Delta \mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{high}}}^\intercal - \Delta \cdot \tilde{\mathbf{e}}_{\mathsf{high}}^\intercal - \mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{low}}}^\intercal - \tilde{\mathbf{e}}_{\mathsf{low}}^\intercal, \tag{4}
$$

[5]We can take $\sigma_\mathbf{M}$ to be the min of all the gaussian noise parameters involved and the reduction can add additional noise of appropriate width. We omit these details here.

$$
\begin{aligned}
\tilde{\mathbf{e}}_{xx}^{\intercal} &= \left\lfloor \frac{\mathbf{t}^{\intercal}(\mathbf{A}_{att}\mathbf{H}^{F_{xx}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{W}_{xx}) - \Delta \cdot F_{xx}(sd, ind, \beta) + \mathbf{e}^{\intercal}_{F_{xx}}}{\Delta} \right\rfloor \\
&\quad - \left\lfloor \frac{\mathbf{t}^{\intercal}(\mathbf{A}_{att}\mathbf{H}^{F_{xx}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{W}_{xx}) - \Delta \cdot F_{xx}(sd, ind, \beta)}{\Delta} \right\rfloor \qquad (5) \\
\mathbf{e}^{\intercal}_{F_{xx}} &= \mathbf{e}^{\intercal}_{fhe}\mathbf{R}_{F_{xx}} + \mathbf{e}^{\intercal}_{att}\mathbf{H}^{F_{xx}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{e}^{\intercal}_{xx}, \quad \text{where } \mathbf{e}^{\intercal}_{xx} = (e_{xx}, \ldots, e_{xx}) \\
\mathbf{e}^{\intercal}_{\mathbf{t},F_{xx}} &= \left\lfloor \mathbf{t}^{\intercal} \left(\frac{\mathbf{A}_{att}\mathbf{H}^{F_{xx}}_{\mathbf{A}_{att}} + \mathbf{W}_{xx}}{\Delta} - \left\lfloor \frac{\mathbf{A}_{att}\mathbf{H}^{F_{xx}}_{\mathbf{A}_{att}} + \mathbf{W}_{xx}}{\Delta} \right\rfloor \right) \right\rfloor \\
\mathbf{R}_{F_{xx}} &= \text{randomness in the homomorphically evaluated GSW encryption of } F_{xx}(sd, ind, \beta). \qquad (6)
\end{aligned}
$$

Please refer to equations (9), (10), (11), (12) in Section 5.2 for the first equality in (3) and for the detailed computation of the error terms[6]. Thus,

$$
\begin{aligned}
\mathbf{c_B}^{\intercal}\mathbf{K} &= \left(\Delta \left\lfloor \frac{\mathbf{c}_{att}{}^{\intercal}\mathbf{H}^{F_{high}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{c}^{\intercal}_{high}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{att}{}^{\intercal}\mathbf{H}^{F_{low}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{c}^{\intercal}_{low}}{\Delta} \right\rfloor + \mathbf{c}_{IP}{}^{\intercal}\mathbf{G}^{-1}((\mathbf{A}^*)^{\intercal}) \right. \\
&\quad \left. + \beta \cdot crs_{ind} \right) + PRG(sd) + \tilde{\mathbf{e}}^{\intercal}
\end{aligned}
$$

In the ideal world, $\mathbf{K} = \begin{pmatrix} \mathbf{B} \\ \mathbf{c_B}^{\intercal} \end{pmatrix}^{-1} \left(\begin{pmatrix} \mathbf{A}_{err} + \mathbf{A}_{mask} \\ \mathbf{d}^{\intercal} + \mathbf{c}'^{\intercal} \end{pmatrix} \right)$. Thus, $(\mathbf{c_B}, \mathbf{X}, \mathbf{c_D})$ and $\mathbf{K}$ in the ideal world are related as

$$
\begin{aligned}
&\mathbf{c_B}^{\intercal}\mathbf{K} = \mathbf{d}^{\intercal} + (\mathbf{c}')^{\intercal} \\
&= \left(\Delta \left\lfloor \frac{\mathbf{c}_{att}{}^{\intercal}\mathbf{H}^{F_{high}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{c}^{\intercal}_{high}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{att}{}^{\intercal}\mathbf{H}^{F_{low}}_{\mathbf{A}_{att},\mathbf{X}} + \mathbf{c}^{\intercal}_{low}}{\Delta} \right\rfloor + \mathbf{c}_{IP}{}^{\intercal}\mathbf{G}^{-1}((\mathbf{A}^*)^{\intercal}) \right. \\
&\quad \left. + \beta \cdot crs_{ind} \right) + (\mathbf{c}')^{\intercal}
\end{aligned}
$$

we can observe that the relation between the ciphertext terms $(\mathbf{c_B}, \mathbf{X}, \mathbf{c_D})$ and $\mathbf{K}$ is the same in both the worlds except that $PRG(sd) + \tilde{\mathbf{e}}^{\intercal}$ in the real world is replaced with $(\mathbf{c}')^{\intercal}$. By setting the ranges of these terms appropriately, namely, $PRG(sd) \in [-B_{OLS} + \eta, B_{OLS} - \eta]^{1\times M}$, $\|\tilde{\mathbf{e}}\| < \eta$ for $\eta \ll B_{OLS}$ and $\mathbf{c}' \leftarrow [-B_{OLS}, B_{OLS}]^M$, one would like to conjecture that the distributions in the two worlds are indistinguishable. This requires a heuristic noise flooding conjecture, which we describe next.

Noise Flooding Conjecture. We essentially conjecture that for all PPT adversary, $PRG(sd) + \tilde{\mathbf{e}}^{\intercal} \approx (\mathbf{c}')^{\intercal}$. An optimistic (and false) hope is that one can deploy the following sequence of arguments - first, replace $PRG(sd)$ with random $\mathbf{w} \leftarrow [-B_{OLS} + \eta, B_{OLS} - \eta]^M$ due to the pseudorandomness of PRG, then flood $\tilde{\mathbf{e}}$ by $\mathbf{w}$ since $\eta \ll B_{OLS}$.

However, one cannot directly invoke the above argument because there is a subtle dependence between the "flooding noise" $PRG(sd)$ and the noise it seeks to flood $\tilde{\mathbf{e}}$. This is because the error $\tilde{\mathbf{e}}$ depends on the secret $\mathbf{t}$, GSW randomness $\mathbf{R}_{F_{high/low}}$ and errors in encodings, including $\mathbf{e}_{fhe}$, while these in turn protect the seed sd encrypted in the GSW ciphertext $\mathbf{X}$ (please see equations (4) and (5)). This leads to a subtle

[6]The error component $\tilde{\mathbf{e}}^{\intercal}_{xx}$ does not explicitly appear there, because as we will see, this is 0 w.h.p. This is analyzed in detail in Section 5.2

circularity, where on the one hand, the secrecy of sd in $\mathsf{PRG}(\mathsf{sd})$ is required to mask $\widetilde{\mathbf{e}}$, and on the other hand $\widetilde{\mathbf{e}}$ depends on secrets that hide sd.

In the context of private coin evasive LWE, a related dependence has been exploited by [AMYY25] and [HJL25] to develop an attack when the adversary is permitted freedom to choose the circuit implementing the PRG[7]. However, due to careful countermeasures, we believe that our assumption holds *even if the adversary is allowed the freedom to choose an arbitrary circuit implementing the PRG*. We will justify this claim in detail below. However, to better understand our design choices, we take a brief detour to recap the vulnerabilities in Evasive LWE.

We first recall the evasive LWE assumption.

Assumption 4.3 (Evasive LWE). [Wee22, ARYY23] Let λ be a security parameter and $n, M, m', q \in \mathbb{N}$ be parameters and. Let σ_{pre}, σ'_{pre}, σ_{post} and γ be parameters for Gaussian distributions. Let Samp be a PPT algorithm that outputs

$$\mathbf{t} \in \mathbb{Z}_q^{n+1}, \mathbf{P} \in \mathbb{Z}_q^{m' \times M}, \mathsf{aux} \in \{0,1\}^*$$

on input 1^λ. For a PPT adversary $\mathcal{A}_0, \mathcal{A}_1$, we define the following advantage functions:

$$\mathsf{Adv}^{\mathsf{pre}}_{\mathcal{A}_0}(\lambda) \stackrel{\mathsf{def}}{=} \Pr\left[\mathcal{A}_0(\mathbf{B}, \mathbf{P}, \mathbf{t}^\intercal \mathbf{B} + \mathbf{e}, \mathbf{t}^\intercal \mathbf{P} + \mathbf{e}', \mathsf{aux}) = 1\right] - \Pr\left[\mathcal{A}_0(\mathbf{B}, \mathbf{P}, \mathbf{c}_0, \mathbf{c}', \mathsf{aux}) = 1\right]$$
$$\mathsf{Adv}^{\mathsf{post}}_{\mathcal{A}_1}(\lambda) \stackrel{\mathsf{def}}{=} \Pr\left[\mathcal{A}_1(\mathbf{B}, \mathbf{P}, \mathbf{t}^\intercal \mathbf{B} + \mathbf{e}'', \mathbf{K}, \mathsf{aux}) = 1\right] - \Pr[\mathcal{A}_1(\mathbf{B}, \mathbf{P}, \mathbf{c}_0, \mathbf{K}, \mathsf{aux}) = 1]$$

$$\begin{aligned} \text{where} \quad & (\mathbf{t}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}(1^\lambda), \ \ \mathbf{B} \leftarrow \mathbb{Z}_q^{n+1 \times m'}, \ \ \mathbf{c}_0 \leftarrow \mathbb{Z}_q^{m'}, \ \ \mathbf{c}' \leftarrow \mathbb{Z}_q^{M}, \\ & \mathbf{e} \leftarrow \mathcal{D}^{m'}_{\mathbb{Z},\sigma_{\mathsf{pre}}}, \ \ \mathbf{e}' \leftarrow \mathcal{D}^{M}_{\mathbb{Z},\sigma'_{\mathsf{pre}}}, \ \ \mathbf{e}'' \leftarrow \mathcal{D}^{m'}_{\mathbb{Z},\sigma_{\mathsf{post}}}, \\ & \mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}) \text{ with standard deviation } \gamma \end{aligned}$$

We say that the *evasive* LWE (EvLWE) assumption holds if for every PPT Samp and $\mathcal{A}_1$, there exists another PPT $\mathcal{A}_0$ and a polynomial $Q(\cdot)$ such that

$$\mathsf{Adv}^{\mathsf{pre}}_{\mathcal{A}_0}(\lambda) \geq \mathsf{Adv}^{\mathsf{post}}_{\mathcal{A}_1}(\lambda)/Q(\lambda) - \mathsf{negl}(\lambda).$$

Public versus Private Coin. Evasive LWE has been studied both in the "public-coin" and "private-coin" setting, where the former means that the randomness used by the sampler is made available to the adversary, and the latter means that the sampler's random coins need to be hidden from the adversary. Here, note that it is necessary that the LWE secret $\mathbf{t}$ and error remain hidden from the adversary, whether in the public or private coin setting. In the public-coin setting, these are generated randomly from the appropriate distribution and are the *only* secrets not known to the adversary, whereas in the private coin setting, there may be additional randomness that is not known to the adversary.

Summarizing Attacks.

- *1: Error Correlating Attacks.* A common theme that runs through these attacks [AMYY25, HJL25] is that they exploit flexibility in design of $\mathbf{P}$ even though it is publicly computable. In these attacks, $\mathbf{P}$ encodes some functionality F, so that the error in the post-condition lends itself to manipulation that is not captured by the well-distributed i.i.d error in the pre-condition. It is very interesting that even though the assumption of evasive LWE originated in the desire to avoid the so-called "zeroizing regime" where attacks against iO candidates had traditionally been found [CHL+15, MSZ16, CVW18, HJL21], the new attacks show that by constructing malicious samplers, we can still end up in the zeroizing regime by cleverly manipulating the computation!

[7] In the attacks described in [AMYY25] and [HJL25], the flooding noise is generated by a PRF instead of PRG.

- *2: Private ($\mathbf{P}$, aux) Attacks.* Another set of attacks exploit the freedom to choose the matrix $\mathbf{P}$ and the auxiliary information aux to mount attacks – for instance, the obfuscation based attack in [VWW22, HHY25] samples $\mathbf{P}$ along with a trapdoor, and the trapdoor is hidden within an obfuscated program that is provided as part of aux. In other works such as [BUW24] and [DJM+25], an attack is mounted via a carefully designed aux which includes specific terms that are exploited to distinguish the post condition distribution.

- *3: Incompressibility Attacks.* An important class of attacks developed by [BDJ+24, AMYY25] are "incompressibility" attacks which show that "multi-challenge" versions of private-coin evasive LWE, where the adversary is allowed the freedom to program private coins, cannot be true. In more detail, the "multi-challenge" setting is where the secret vector $\mathbf{t}^\intercal$ in evasive LWE is instead, a matrix $\mathbf{T} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{Q_c \times (n+1)}$ and $\mathbf{K}$ is of the dimension $m' \times M$.

  In a bit more detail, the idea of the attack is to hide $Q_k \cdot Q_c$ pseudorandom values, for sufficiently large Q_k and Q_c into the decryption results obtained by combining Q_k decryption keys with Q_c ciphertexts of a functional encryption scheme, which is constructed using private coin Evasive LWE. Since the total size of the ciphertexts and decryption keys is $O(Q_k + Q_c)$ these perform the task of compressing the $Q_k \cdot Q_c$ pseudorandom values that they jointly yield. These pseudorandom values are taken to be the statement of an instance hiding witness encryption scheme (WE), and their compressed encoding via the FE ciphertexts and keys is taken to be the witness.

  Next, these WE ciphertexts are engineered into the decryption results of additional ciphertexts of the same FE scheme (with a different function key). Put together, the decryptor can perform decryption to obtain the WE ciphertexts and then use the Q_c ciphertexts and Q_k keys described previously, as the witness of these WE ciphertexts. Thus, in the real world, the WE ciphertexts can be decrypted whereas in the ideal world, where the encodings are random, they do not form a valid witness and decryption is not possible with high probability if $Q_k \cdot Q_c > O(Q_k + Q_c) + \lambda$.

- *4: Withholding-Matrix Attacks.* The work of [BUW24] presents attacks against classes of evasive LWE where either $\mathbf{B}$ or $\mathbf{P}$ are not known to the adversary. In more detail, consider the distinguisher of the pre-condition who receives the matrix $\mathbf{B}$, but not $\mathbf{P}$ – in this case, the distinguisher of the post-condition can easily recover $\mathbf{P}$ by simply computing $\mathbf{B} \cdot \mathbf{B}^{-1}(\mathbf{P})$. This may create leakage in scenarios where $\mathbf{P}$ contains secret information that was deliberately withheld from the adversary. The case where $\mathbf{B}$ is not available to the distinguisher of the pre-condition but $\mathbf{P}$ is partially available is similar. These attacks can be prevented by ensuring that the distinguisher in the pre-condition also has access to complete information about $\mathbf{B}$ and $\mathbf{P}$.

Interpreting the Attacks. Let us take a step back and try to understand the high level learning obtained from these attacks. Towards this, let us revisit the primary intuition of evasive LWE as formulated by Wee[Wee22]. One way to interpret the original formulation of Wee's evasive LWE is that it is really two assumptions rolled into one[8]:

1. **Preimages $\mathbf{B}^{-1}(\mathbf{P})$ can only be used semi-honestly**: Short Gaussian preimages $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{P})$ can only be used semi-honestly to compute samples $\mathbf{t}^\intercal\mathbf{P} + \mathbf{e}^\intercal\mathbf{K}$ from $\mathbf{t}^\intercal\mathbf{B} + \mathbf{e}^\intercal$ and cannot be exploited in any other way. Here, note that this condition implicitly demands that $\mathbf{P}$ is sufficiently "well-behaved", and in particular does not contain low norm linear dependencies that would permit "mix and match" attacks of preimages [Agr17][9].

[8] These cannot be separated in general, needless to say.

[9] As an extreme example of badly behaved $\mathbf{P} = [\mathbf{P}_1, \mathbf{P}_2]$, suppose the adversary could obtain $\mathbf{K}_1$, $\mathbf{K}_2$ such that $\mathbf{B}\mathbf{K}_1 = \mathbf{P}_1 \bmod q$ and $\mathbf{B}\mathbf{K}_2 = 2\mathbf{P}_1 \bmod q$. Then, the adversary could compute a short trapdoor to $\mathbf{B}$ simply as $2\mathbf{K}_1 - \mathbf{K}_2$, which is clearly disastrous. However, such "pre-image combination" attacks are prevented by the requirement in the pre-condition, that $\underline{\mathbf{s}\mathbf{P}_1}$, $\underline{\mathbf{s}\mathbf{P}_2}$ are jointly pseudorandom.

2. **LWE with correlated error is "no worse" than standard LWE**: By using $\mathbf{B}^{-1}(\mathbf{P})$ as above, the adversary can compute an LWE sample $\mathbf{t}^\intercal\mathbf{P} + \mathbf{e}^\intercal\mathbf{K}$ with correlated, overdefined error as against i.i.d. error. However, the intuition here is that this LWE with correlated noise can be conjectured "as good as" LWE since the attacker should never be able to remove the large mask $\mathbf{t}^\intercal\mathbf{P}$ and therefore should never obtain anything in the zeroizing regime.

We now study the two assumptions above in turn, as though they are separate. The first assumption has proven "morally, more or less" sound in that, it has only yielded attacks when the auxiliary input provided to the adversary explicitly contains contrived information, such an obfuscated program which spits out secrets when given short preimages [VWW22, BUW24]. However, for "natural" distributions of auxiliary information, we do not know even a single attack that exploits $\mathbf{K}$ to compute anything other than $\mathbf{t}^\intercal\mathbf{P} + \mathbf{e}^\intercal\mathbf{K}$, which is explicitly desired. The new problems come from the second assumption, where an i.i.d error would have sufficed for some flooding argument (allowing the pre-condition to be proven), but the correlated, overdefined error does not.

We now analyze our assumption with respect to the attacks known on evasive LWE.

1: Error Correlating Attacks. Error correlating attacks as developed by [AMYY25] and [HJL25] do not apply for our assumption for two important reasons. We begin with the most significant one - use of modulus reduction which destroys any specific correlation that could be exploited. The underlying schemes broken by [AMYY25] and [HJL25] didn't use modulus reduction and the error term equivalent to our $\widetilde{\mathbf{e}}^\intercal$ in equation (4) (let's call it $\bar{\mathbf{e}}^\intercal$) is of the form

$$\bar{\mathbf{e}}^\intercal = \mathbf{e}_{\mathbf{B}}^\intercal\mathbf{K} - \mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}} - \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}^{\mathrm{F}}_{\mathbf{A}_{\mathsf{att}},\boldsymbol{x}}.$$

The error correlation attacks by [AMYY25] and [HJL25] correlate $\mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}}$ with $\mathrm{F}(\mathbf{x})$, which was otherwise supposed to mask these error terms, in such a way that they together become 0 mod 2. After applying the modulus reduction, the error $\widetilde{\mathbf{e}}^\intercal$ still depends on $\mathbf{e}_{\mathsf{fhe}}\mathbf{R}_{\mathrm{F}_{\mathsf{xx}}}$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$, but it is now restricted to rounding errors which seems hard to control to be exploitable. Let us analyze this in more detail.

In equation (4), the first two terms, $\mathbf{e}_{\mathbf{B}}^\intercal\mathbf{K}$ and $\mathbf{e}_{\mathsf{lP}}^\intercal\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ are independent of F_{xx}[10] and hence can't be manipulated by the adversary by choosing any contrived circuit for F_{xx}. So, let us focus on $\mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{xx}}}$ and $\widetilde{\mathbf{e}}_{\mathsf{xx}}$. We have,

$$\mathbf{e}_{\mathbf{t},\mathrm{F}_{\mathsf{xx}}}^\intercal = \left\lfloor \mathbf{t}^\intercal \left(\frac{\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{W}_{\mathsf{xx}}}{\Delta} - \left\lfloor \frac{\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} + \mathbf{W}_{\mathsf{xx}}}{\Delta} \right\rfloor \right) \right\rfloor$$

This cannot cause any correlation with $\mathrm{F}[\mathsf{crs}](\mathsf{sd}, \mathsf{ind}, \beta)$, and hence with $\mathsf{PRG}(\mathsf{sd})$, for two reasons: (i) It is independent of the input $(\mathsf{sd}, \mathsf{ind}, \beta)$, (ii) It is only a rounding error that breaks any specific structure.

Lastly, we come to the key error component, $\widetilde{\mathbf{e}}_{\mathsf{xx}}$, which contains the contentious term $\mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}_{\mathsf{xx}}}$ which has been exploited before [AMYY25, HJL25]. We have,

$$\begin{aligned}
&\tilde{\mathbf{e}}_{\mathsf{xx}}^\intercal \\
&= \left\lfloor \frac{\mathbf{t}^\intercal(\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\boldsymbol{x}} + \mathbf{W}_{\mathsf{xx}}) - \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathsf{sd}, \mathsf{ind}, \beta) + \mathbf{e}_{\mathsf{fhe}}^\intercal\mathbf{R}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{e}_{\mathsf{att}}^\intercal\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\boldsymbol{x}} + \mathbf{e}_{\mathsf{xx}}^\intercal}{\Delta} \right\rfloor \\
&- \left\lfloor \frac{\mathbf{t}^\intercal(\mathbf{A}_{\mathsf{att}}\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\boldsymbol{x}} + \mathbf{W}_{\mathsf{xx}}) - \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathsf{sd}, \mathsf{ind}, \beta)}{\Delta} \right\rfloor
\end{aligned}$$

[10] Technically speaking, $\mathbf{K}$ does depend on F_{xx}, but since it is sampled randomly from the corresponding Gaussian distribution, it is out of the control of the adversary.

Above, the terms highlighted in red are those which have previously been exploited to create correlations with the flooding noise. However, by the setting of our parameters, we have $\|\mathbf{e}_{\mathsf{fhe}}^\intercal \mathbf{R}_{\mathsf{F}_{\mathsf{xx}}} + \mathbf{e}_{\mathsf{att}}^\intercal \mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\boldsymbol{X}} + \mathbf{e}_{\mathsf{xx}}^\intercal\| << \Delta$. Thus, $\widetilde{\mathbf{e}}_{\mathsf{xx}}^\intercal = 0$ with $1 - \mathsf{negl}(\lambda)$ probability – please see the detailed reasoning in Section 5.2 on Page 33.

We remark that modulus reduction is also a crucial part of CRO assumption in [HJL25] - CRO assumption defines a nice safety condition which seems to be instantiable only by using modulus reduction that removes decryption noise by rounding. Another safeguard against error correlation attacks is to fix a canonical circuit computing the function $\mathsf{F}[\mathsf{crs}](\mathsf{sd}, \mathsf{ind}, \beta) = \mathsf{PRG}(\mathsf{sd}) + \beta \cdot \mathsf{crs}_{\mathsf{ind}}$. However, we conjecture that our assumption holds even for maliciously chosen circuit representations.

2: Publicly computable matrix $\mathbf{P}$ and no aux. In our assumption there is no auxiliary information and the matrix $\mathbf{P}$ is publicly computable whose all the coins are public. This leaves no space for the adversary to embed any hint in $\mathbf{P}$ or in aux that could be exploited.

3: Incompressibility does not apply. The lower bounds based on incompressibility arguments from [AMYY25, BDJ+24] do not apply to our assumption. To see this, note that:

- The attacks crucially need the "multi-challenge" setting where the secret vector $\mathbf{t}^\intercal$ is instead, a matrix $\mathbf{T} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{Q_c \times (n+1)}$ and $\mathbf{K}$ is of the dimension $m \lceil \log q \rceil \times M$ similar to ours. The attack crucially requires a large Q_c as discussed above and does not work when $Q_c = 1$ as is the case in our assumption.
- The attacks crucially exploit the flexibility of the adversary to program the private coins (with witness encryption ciphertexts) admitted by the private coin evasive LWE assumption. Our assumption does not permit this flexibility to the attacker.

4: Matrices $\mathbf{B}, \mathbf{P}$ are public. These attacks do not apply to us since there is no pre-condition.

A Note on Presence of "LWE secret" $\widehat{\mathbf{S}}$: The reader may be concerned that the matrix $\widehat{\mathbf{S}}$, which appears like an LWE secret, is provided in the clear in our assumption. However, this apparent oddity is created by the functionality of Oblivious LWE Sampler (OLS) which our assumption will later be used to build. In detail, note that the only place where $\widehat{\mathbf{S}}$ appears in the assumption is in defining the function $\mathsf{F}[\mathsf{crs}]$, on which the matrix $\mathbf{P}$ depends, where $\mathsf{crs} = \mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}$. However, we note that we do not assume any specific distribution of crs to argue the plausibility of the assumption. Indeed, the role of crs here is simply to satisfy the functionality required by the subsequent application of the assumption. Hence, giving out $\widehat{\mathbf{S}}$ in plain does not pose any threat to the assumption. In fact, we observe that for $\beta = 0$, crs can simply be stripped off or replaced with any random value as it does not impact the function output. For $\beta = 1$, $\mathsf{crs}_{\mathsf{ind}}$ is added to the output of the function, however it's presence is just for functionality and is not really expected to provide any security.

4.2.1 Other Relevant Attacks from the Literature.

Here we examine other avenues of attack that have been studied in the literature and discuss why they do not seem to apply to our setting.

- *Attacks against Tensor Structure:* The work of Jain et al. (JLLS) [JLLS23] attacked the tensor structure which was used crucially in the construction of Devadas et al. [DQV+21]. In particular, JLLS showed that Kilian randomization on highly tensored matrices does not kill the tensor structure, and this structure can be exploited for attacks. This attack does not appear to apply to our setting since we do not rely on tensors.

- *Linearization and Sum-Of-Squares.* Linearization attacks, as the name suggests, are attacks where a system of high degree polynomial equations are linearized and then solved using standard techniques like Gaussian elimination. The sum-of-squares paradigm has been used to develop attacks [LV17, BBKK18, BHJ+19] against general expanding families of low-degree polynomials over the reals. These attacks do not seem to apply to the current setting because there aren't enough equations for the former class of attacks and the latter class of attacks requires low degree equations over the reals, whereas our PRG computation does not suffer any degree restriction, since we do not use pairings.
- *Attack against Gay-Pass Scheme.* The Gay-Pass assumption [GP21] asserts that *if* an encryption scheme is CPA secure when the adversary has access to certain leakage on the encryption randomness (called "Shielded Randomness Leakage"), then additionally publishing an encrypted key cycle does not hurt security. Counterexamples developed against this conjecture [HJL21] do not apply to our setting since the specific structure of their assumption is very different from ours.

5 $i\mathcal{O}$ via LWE with Structured Preimage

5.1 Definition: Oblivious LWE Sampler

In this section, we recall the definition of oblivious LWE sampler from [WW21, HJL25].

Syntax. Let λ be the security parameter. All other parameters implicitly depend on λ. Let $W, M, q, \sigma, B_{\mathbf{e}}$ be lattice parameters. An oblivious LWE sampler consists of the following algorithms

- $\mathsf{genCRS}(1^\lambda, 1^K) \to \mathsf{crs}$.
- $\mathsf{init}(1^K, \mathsf{crs}, \mathbf{A} \in \mathbb{Z}_q^{M\times W}) \to \mathsf{pub}$[11], where $\mathsf{bitlength}(\mathsf{pub}) \le KM^{1-\epsilon}\log q + \mathrm{poly}(\lambda)$ for some constant $\epsilon < 1$, i.e, pub is smaller than K LWE samples in $\mathbb{Z}_q^M$.
- $\mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i) = \mathbf{b}_i \in \mathbb{Z}_q^M$ is deterministic.
- $\mathsf{Sim}(1^\lambda, 1^K, \mathbf{A}, \{\widehat{\mathbf{b}}_i\}_{i\in[K]}) \to (\mathsf{crs}, \mathsf{pub}, \{\widetilde{\mathbf{s}}_i\}_{i\in[K]})$.

An oblivious LWE sampler should satisfy the following properties.

Correctness. Let $(\mathbf{A}, \mathbf{A}_\tau^{-1}) \leftarrow \mathsf{TrapGen}(1^M, 1^W, q), \mathsf{crs} \leftarrow \mathsf{genCRS}(1^\lambda, 1^K), \mathsf{pub} \leftarrow \mathsf{init}(1^K, \mathsf{crs}, \mathbf{A}), \mathbf{b} \leftarrow \mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i)$, then with overwhelming probability there exists $\widetilde{\mathbf{s}} \in \mathbb{Z}_q^W, \widetilde{\mathbf{e}} \in [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$ such that $\mathbf{b}^\intercal = \widetilde{\mathbf{s}}^\intercal\mathbf{A}^\intercal + \widetilde{\mathbf{e}}^\intercal$. In other words, the output of Sample should be LWE samples with respect to $\mathbf{A}$.

Security. The following distributions are computationally indistinguishable.

- **Real distribution:** $(\mathbf{A}, \mathbf{A}_\tau^{-1}) \leftarrow \mathsf{TrapGen}(1^M, 1^W, q), \mathsf{crs} \leftarrow \mathsf{genCRS}(1^\lambda, 1^K)$, $\mathsf{pub} \leftarrow \mathsf{init}(1^K, \mathsf{crs}, \mathbf{A})$ for all $i \in [K]$, $\mathbf{b}_i \leftarrow \mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i)$, $\widetilde{\mathbf{s}}_i = \mathsf{LWESolve}_{\mathbf{A}_\tau^{-1}}(\mathbf{b}_i)$, where $\mathsf{LWESolve}_{\mathbf{A}_\tau^{-1}}$ is the efficient LWE secret solving procedure given the trapdoor. Output $(\mathsf{crs}, \mathbf{A}, \mathsf{pub}, \{\widetilde{\mathbf{s}}_i\}_{i\in[K]})$.
- **Simulated distribution:** $(\mathbf{A}, \mathbf{A}_\tau^{-1}) \leftarrow \mathsf{TrapGen}(1^M, 1^W, q)$. For $i \in [K]$, sample $\widehat{\mathbf{s}}_i \leftarrow \mathbb{Z}_q^W, \widehat{\mathbf{e}}_i \leftarrow \mathcal{D}_{\widehat{\sigma}}^M$ and set $\widehat{\mathbf{b}}_i^\intercal = \widehat{\mathbf{s}}_i^\intercal\mathbf{A}^\intercal + \widehat{\mathbf{e}}_i^\intercal$. Simulate $(\mathsf{crs}, \mathsf{pub}, \{\widetilde{\mathbf{s}}_i\}_{i\in[K]}) \leftarrow \mathsf{Sim}(1^\lambda, 1^K, \mathbf{A}, \{\widehat{\mathbf{b}}_i\}_{i\in[K]})$ and output $(\mathsf{crs}, \mathbf{A}, \mathsf{pub}, \{\widetilde{\mathbf{s}}_i + \widehat{\mathbf{s}}_i\}_{i\in[K]})$.

[11] In [WW21, HJL25], init does not take crs as input. However, the compiler to build functional encoding (FEd) from an OLS works even when init takes crs as an input. In particular, the CRS of FEd (let's call it FEd.crs) is the same as the CRS of underlying OLS (call it OLS.crs) and FEd.Enc that calls OLS.init does take FEd.crs = OLS.crs as an input. Hence, this can be provided as an input to the OLS.init algorithm.

In other words, the output distribution of oblivious LWE sampling is indistinguishable to the simulated distribution, which hides an additional LWE shift.

5.2 Construction: Oblivious LWE Sampler

In this section, for any matrix $\mathbf{M} \in \mathbb{Z}_q^{a\times b}$, $\mathbf{m}_i$ represents the i-th column of $\mathbf{M}$, for all $i \in [b]$. We will need a pseudorandom generator $\mathsf{PRG} : \{0,1\}^\lambda \to [-B_{\mathsf{OLS}}+\eta, B_{\mathsf{OLS}}-\eta]^{1\times M}$ that can be evaluated by a circuit of depth at most $\mathsf{dep}(\lambda) = \mathrm{poly}(\lambda)$.

$\mathsf{genCRS}(1^\lambda, 1^K)$**.** The genCRS algorithm does the following.

1. Sample $\mathsf{crs} \leftarrow \mathbb{Z}_q^{M\times K}$.
2. Define a function $\mathsf{LWESamp}[\mathsf{crs}] : \{0,1\}^{1\times L} \to \mathbb{Z}_q^{1\times M}$ with crs being hardwired as follows: it takes as input $(\mathsf{sd}, \mathsf{bits}(i), \beta)$, where $\mathsf{sd} \in \{0,1\}^\lambda, \beta \in \{0,1\}$ and $i \in [K]$, and outputs $\mathsf{PRG}(\mathsf{sd}) + \beta \cdot \mathsf{crs}_i^\intercal \in \mathbb{Z}_q^{1\times M}$. Here, $L = \lambda + \lceil \log K \rceil + 1$ and $\Delta \in \mathbb{Z}$ is some fixed constant such that Δ divides q. Let the depth of circuit be $\mathsf{dep}(\lambda) = \mathrm{poly}(\lambda)$.
3. Output $\mathbf{CRS} = (\mathsf{LWESamp}, \mathsf{crs}, \Delta)$.

$\mathsf{init}(1^K, \mathbf{CRS}, \mathbf{A}^* \in \mathbb{Z}_q^{M\times W})$**.** The init algorithm does the following.

1. Sample $\mathbf{S}^* \leftarrow \mathbb{Z}_q^{W\times K}, \bar{\mathbf{A}}_{\mathsf{fhe}} \leftarrow \mathbb{Z}_q^{n\times m}, \mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}, \mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{x}}+1)m}$, $\mathbf{w}_{\mathsf{high}} \leftarrow \mathbb{Z}_q^{n+1}, \mathbf{w}_{\mathsf{low}} \leftarrow \mathbb{Z}_q^{n+1}$ where $L_{\mathbf{x}} = (n+1)mL\lceil \log q\rceil$ and $L = \lambda + \lceil \log K\rceil + 1$.
2. Sample $(\mathbf{B}, \mathbf{B}_\gamma^{-1}) \leftarrow \mathsf{TrapGen}(1^{n+1}, 1^{mw}, q)$ where $w \in O(\log q)$.
3. For all $i \in [K]$,
   (a) Sample $\bar{\mathbf{t}}_i \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_t}^n$ and set $\mathbf{t}_i = (\bar{\mathbf{t}}_i^\intercal, -1)^\intercal$.
   (b) Sample $\mathbf{e}_{\mathbf{B},i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathbf{B}}}^{mw}$ and compute $\mathbf{c}_{\mathbf{B},i}^\intercal = \mathbf{t}_i^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B},i}^\intercal \in \mathbb{Z}_q^{1\times mw}$.
   (c) Compute FHE ciphertext: sample PRG seed $\mathsf{sd}_i \leftarrow \{0,1\}^\lambda$, $\mathbf{e}_{\mathsf{fhe},i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathsf{fhe}}}^m, \mathbf{R}_i \leftarrow \{0,1\}^{m\times mL}$ and compute GSW encryption as follows.

$$\mathbf{A}_{\mathsf{fhe},i} := \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}_i^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe},i}^\intercal \end{pmatrix}, \mathsf{hct}_i := \mathbf{A}_{\mathsf{fhe},i}\mathbf{R}_i - (\mathbf{x}_i^\intercal) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1)\times mL}$$

   where $\mathbf{x}_i^\intercal = (\mathsf{sd}_i, \mathsf{bits}(i), 0) \in \mathbb{Z}_q^{1\times L}$. Let $L_{\mathbf{x}} = (n+1)mL\lceil \log q\rceil$ be the bit length of hct_i.
   (d) Compute BGG+ encoding of hct_i as follows: $\mathbf{e}_{\mathsf{att},i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathsf{att}}}^{(L_{\mathbf{x}}+1)m}$ and

$$\mathsf{bct}_i^\intercal := \mathbf{t}_i^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}_i)) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att},i}^\intercal \in \mathbb{Z}_q^{1\times(L_{\mathbf{x}}+1)m}$$

   (e) Compute IPFE encoding of $\mathbf{s}_i^*$ as: $\mathbf{e}_{\mathsf{IP},i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\mathsf{IP}}}^{W\lfloor \log q\rceil}$, $\mathsf{ict}_i^\intercal = \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP},i}^\intercal - (\mathbf{s}_i^*)^\intercal \otimes \mathbf{g}^\intercal$.
   (f) Compute $c_{\mathsf{xx},i} = \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{xx}} + e_{\mathsf{xx},i}$ where $e_{\mathsf{xx},i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$.
   (g) Set $\mathsf{ct}_i = (\mathbf{c}_{\mathbf{B},i}, \mathsf{hct}_i, \mathsf{bct}_i, \mathsf{ict}_i, c_{\mathsf{high},i}, c_{\mathsf{low},i})$.
4. Parse $\mathbf{CRS}$ as $(\mathsf{LWESamp}[\mathsf{crs}], \mathsf{crs})$ and set $\mathrm{F} = \mathsf{LWESamp}[\mathsf{crs}]$. Define functions $\mathrm{F}_{\mathsf{high}}$ and $\mathrm{F}_{\mathsf{low}}$ such that $\mathrm{F}(\mathsf{sd}, \mathsf{bits}(i), \beta) = \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathsf{sd}, \mathsf{bits}(i), \beta) + \mathrm{F}_{\mathsf{low}}(\mathsf{sd}, \mathsf{bits}(i), \beta)$. Note that both $\Delta \cdot \mathrm{F}_{\mathsf{high}}$ and $\mathrm{F}_{\mathsf{low}}$ can also be computed by a circuit of depth $d = \mathrm{poly}(\mathsf{dep})$.
5. Define the homomorphic evaluation circuit $\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}} = \mathsf{MakeVEvalCkt}(n, m, q, \Delta \cdot \mathrm{F}_{\mathsf{xx}})$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$. From Lemma 3.9, the depth of $\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}$ is $dO(\log m \log\log q) + O(\log^2 \log q)$.

6. Compute $\mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} = \mathsf{MEvalC}(\mathbf{A}_{\mathsf{att}}, \mathsf{VEval}_{\mathsf{F}_{\mathsf{xx}}}) \in \mathbb{Z}_q^{(L_{\mathbf{X}}+1)m \times M}$ and $\mathbf{A}_{\mathsf{F}_{\mathsf{xx}}} = \mathbf{A}_{\mathsf{att}} \cdot \mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}}$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$.
7. Set $\mathbf{W}_{\mathsf{xx}} = (1, \ldots, 1) \otimes \mathbf{w}_{\mathsf{xx}} \in \mathbb{Z}_q^{(n+1)\times M}$ for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$.
8. Compute $\mathbf{A}_{\mathsf{err}} = \Delta \cdot \left\lfloor \frac{\mathbf{A}_{\mathsf{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathsf{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor$.
9. Compute $\mathbf{A}_{\mathsf{mask}} = \mathbf{A}_{\mathsf{IP}}\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ and $\mathbf{A}_{\mathsf{F}} = \mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$.
10. Set $\mathsf{sk} = \mathbf{K} \leftarrow \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathsf{F}})$.
11. Output $\mathsf{pub} = (\mathbf{A}^*, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathsf{sk}, \{\mathsf{ct}_i\}_{i\in[K]})$

$\mathsf{Sample}(\mathbf{CRS}, \mathsf{pub}, i)$**.** The sample algorithm does the following.

1. Parse $\mathsf{pub} = (\mathbf{A}^*, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathsf{sk}, \{\mathsf{ct}_i\}_{i\in[K]})$, $\mathbf{CRS} = (\mathsf{LWESamp}[\mathsf{crs}], \mathsf{crs}, \Delta)$.
2. Define $\mathsf{F}_{\mathsf{high}}$ and $\mathsf{F}_{\mathsf{low}}$ as in the init algorithm, and compute $\mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} = \mathsf{MEvalCX}(\mathbf{A}_{\mathsf{att}}, \mathsf{VEval}_{\mathsf{F}_{\mathsf{xx}}}, \mathsf{hct}_i)$ and set $\mathbf{c}^\intercal_{\mathsf{xx},i} = c_{\mathsf{xx},i} \cdot (1, \ldots, 1) \in \mathbb{Z}_q^{1\times M}$, for $\mathsf{xx} \in \{\mathsf{high}, \mathsf{low}\}$.
3. Compute
$$(\mathbf{b}_i^*)^\intercal = \mathbf{c}^\intercal_{\mathbf{B},i} \cdot \mathbf{K} - \left(\Delta \cdot \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathsf{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{high},i}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathsf{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{low},i}}{\Delta} \right\rfloor \right) - \mathsf{ict}_i^\intercal \cdot \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal).$$

$\mathsf{Sim}(1^\lambda, 1^K, \mathbf{A}^*, \{\widehat{\mathbf{b}}_i\}_{i\in[K]})$**.** The Simulator does the following.

1. Sample $\mathbf{S}^* \leftarrow \mathbb{Z}_q^{W\times K}$.
2. Set $\mathsf{crs} = (\widehat{\mathbf{b}}_1, \ldots, \widehat{\mathbf{b}}_K)$ and $\mathbf{CRS} = (\mathsf{LWESamp}[\mathsf{crs}], \mathsf{crs}, \Delta)$.
3. For all $i \in [K]$, sample $\mathsf{sd}_i \leftarrow \{0,1\}^\lambda$ and set $\mathbf{x}_i^\intercal = (\mathsf{sd}_i, \mathsf{bits}(i), 1) \in \mathbb{Z}_q^{1\times L}$.
4. For all $i \in [K]$, set $\mathsf{ct}_i = (\mathbf{c}_{\mathbf{B},i}, \mathsf{hct}_i, \mathsf{bct}_i, \mathsf{ict}_i, c_{\mathsf{high},i}, c_{\mathsf{low},i})$ where $\mathbf{c}_{\mathbf{B},i}$, hct_i, bct_i, ict_i, $c_{\mathsf{high},i}, c_{\mathsf{low},i}$ are computed as in Step 3 of init algorithm.
5. Compute sk as in Step 10 of init algorithm.
6. Set $\mathsf{pub} = (\mathbf{A}^*, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathsf{sk}, \{\mathsf{ct}_i\}_{i\in[K]})$, where $\mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}$ are sampled as in the init algorithm.
7. Return $(\mathbf{CRS}, \mathsf{pub}, \{\widetilde{\mathbf{s}}_i = \mathbf{s}_i^*\}_{i\in[K]})$.

Parameters Constraints. We list here the set of constraints needed for correctness and security.

- Δ divides q.
- $(m+2)^d \lceil \log q \rceil m < \widehat{\eta}, \quad 2^{dO(\log\lambda)} < \widehat{\eta}$ where d is as defined in Step 4 of init algorithm.
- $\|\mathbf{e}^\intercal_{\mathsf{fhe},i}\mathbf{R}_{\mathsf{F}_{\mathsf{xx}},i} + \mathbf{e}^\intercal_{\mathsf{att},i}\mathbf{H}^{\mathsf{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{e}^\intercal_{\mathsf{high},i}\|$ is exponentially smaller Δ for $\mathsf{xx} = \{\mathsf{high}, \mathsf{low}\}$.
- $O(\sigma_{\mathbf{B}}\gamma + \Delta\sigma_{\mathbf{t}} + W\sigma_{\mathsf{IP}})\mathrm{poly}(\lambda) \le \eta$
- $\sqrt{\lambda}\widehat{\sigma}$ is exponentially less than B_{OLS}
- η is exponentially less than B_{OLS}
- $q \ge O(\sqrt{MW\log q})B_{\mathsf{OLS}}$

Parameters Setting. We set the parameters as below.

$$\widehat{\eta} = 2^{O(\mathsf{dep}\cdot\log^3\lambda)},\ n = \mathrm{poly}(\lambda, \mathsf{dep}),\ m = O(n\log q),\ \gamma = \tau = O\left(\sqrt{(n+1)\log q}\right),$$
$$\sigma_{\mathbf{t}} = \sigma_{\mathsf{fhe}} = \sigma_{\mathsf{att}} = \sigma_{\mathsf{IP}} = \sigma = 2^{2\lambda},\ \sigma_{\mathbf{B}} = 2^{3\lambda},\ \eta \geq O(\Delta\sigma_{\mathbf{t}} + W\sigma_{\mathsf{IP}} + \sigma_{\mathbf{B}}\gamma)\mathrm{poly}(\lambda),$$
$$\Delta = 2^{4\lambda}\widehat{\eta}, B_{\mathsf{OLS}} = 2^{\lambda}\eta, \widehat{\sigma} = \frac{B_{\mathsf{OLS}}}{2^{2\lambda}}, q \geq O(\sqrt{MW\log q})B_{\mathsf{OLS}}.$$

Succinctness. We now prove the succinctness.

Claim 5.1. Let $(\mathbf{A}^*, (\mathbf{A}^*_\tau)^{-1}) \leftarrow \mathsf{TrapGen}(1^W, 1^M, q)$. For $\mathbf{CRS} \leftarrow \mathsf{genCRS}(1^\lambda, 1^K)$, $\mathsf{pub} \leftarrow \mathsf{init}(1^K, \mathbf{CRS}, \mathbf{A}^*)$ as defined in the construction, we have $\mathsf{bitlength}(\mathsf{pub}) \leq KM^{1-\epsilon}\log q + \mathrm{poly}(\lambda)$ for $\epsilon < 1$.

Proof. To prove the claim, we start with $|\mathsf{ct}_i|$.

$$\begin{aligned}
\mathsf{bitlength}(\mathsf{ct}_i) &= \mathsf{bitlength}(\mathbf{c}_{\mathbf{B},i}, \mathsf{hct}_i, \mathsf{bct}_i, \mathsf{ict}_i, c_{\mathsf{high},i}, c_{\mathsf{low},i}) \\
&= mw\log q + (n+1)mL\log q + (L_{\mathbf{X}}+1)m\log q + W\log^2 q + 2\log q \\
&\leq \mathrm{poly}_1(\lambda) + \mathrm{poly}_2(\lambda) + \mathrm{poly}_3(\lambda) + \mathrm{poly}_4(\lambda) + \mathrm{poly}_5(\lambda) \\
&\leq \mathrm{poly}_{\mathsf{ct}}(\lambda)
\end{aligned}$$

In the above, W is some fixed polynomial in λ. Now,

$$\begin{aligned}
\mathsf{bitlength}(\mathsf{pub}) &= \mathsf{bitlength}(\mathbf{A}^*, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathsf{sk}) + K\cdot\mathsf{bitlength}(\mathsf{ct}_i) \\
&\leq MW\log q + (n+1)W\log^2 q + (n+1)^2m^2L\log^2 q + mwM\log q + K\mathrm{poly}_{\mathsf{ct}}(\lambda) \\
&\leq M\mathrm{poly}_6(\lambda) + \mathrm{poly}_7(\lambda) + K\mathrm{poly}_{\mathsf{ct}}(\lambda) \\
&\leq (M+K)\mathrm{poly}(\lambda)
\end{aligned}$$

where the third inequality follows as $n, m, w, L, W, \log q, L_{\mathbf{X}}$ are bounded by some fixed polynomial in λ and by taking $\mathrm{poly}(\lambda)$ as the maximum of $\mathrm{poly}_6(\lambda)$, $\mathrm{poly}_7(\lambda)$, $\mathrm{poly}_{\mathsf{ct}}(\lambda)$.

We can easily see that for sufficiently large K and M, i.e., $K, M > \mathrm{poly}(\lambda)$, we have $(M+K)\mathrm{poly}(\lambda) \leq KM^{1-\epsilon}\log q + \mathrm{poly}'(\lambda)$.

We divide the analysis in two cases: (i) $K \geq M$, and (ii) $K < M$, and show that the above inequality holds.

Case 1: $K \geq M$**.** Here, $LHS = (M+K)\mathrm{poly}(\lambda) \leq 2K\mathrm{poly}(\lambda)$. Note that $RHS = KM^{1-\epsilon}\log q + \mathrm{poly}'(\lambda) > 2KM^{1-\epsilon}$. Therefore, it suffices to prove that $2K\mathrm{poly}(\lambda) < 2KM^{1-\epsilon}$, which is implied by showing $\mathrm{poly}(\lambda) < M^{1-\epsilon}$. This is satisfied, for $\epsilon < 1 - \frac{\log(\mathrm{poly}(\lambda))}{\log M}$. Since $M > \mathrm{poly}(\lambda)$, $0 < \epsilon < 1$.

Case 1: $K < M$**.** Here, $LHS = (M+K)\mathrm{poly}(\lambda) \leq 2M\mathrm{poly}(\lambda)$. Hence, it suffices to prove that $2M\mathrm{poly}(\lambda) < KM^{1-\epsilon}\log q + \mathrm{poly}'(\lambda)$, which is implied by showing $M\mathrm{poly}(\lambda) < KM^{1-\epsilon}$, i.e., $M^\epsilon < K/\mathrm{poly}(\lambda)$. That is, $\epsilon < \frac{\log K - \log(\mathrm{poly}(\lambda))}{\log M}$. Since $M > K > \mathrm{poly}(\lambda)$, $0 < \epsilon < 1$.

For a fixed constant $0 < \epsilon < 0$ we can choose q, K, M such that $q > 2^{2\mathrm{poly}(\lambda)}$, $M > \mathrm{poly}(\lambda)^{1/(1-\epsilon)}$ and $K > M^\epsilon$. □

Correctness. We now prove the correctness.

Claim 5.2. Let $(\mathbf{A}^*, (\mathbf{A}^*_\tau)^{-1}) \leftarrow \mathsf{TrapGen}(1^M, 1^W, q)$. For $\mathbf{CRS} \leftarrow \mathsf{genCRS}(1^\lambda, 1^K)$, $\mathsf{pub} \leftarrow \mathsf{init}(1^K, \mathbf{CRS}, \mathbf{A}^*)$, $\mathbf{b}^*_i = \mathsf{Sample}(\mathbf{CRS}, \mathsf{pub}, i)$ for all $i \in [K]$ as defined in the construction, the probability that the output of Sample is an LWE sample with respect to $\mathbf{A}^*$ i.e., $\mathbf{b}^*_i = \mathbf{A}^*\mathbf{s}^*_i + \mathbf{e}^*_i$ where $\mathbf{e}^*_i \in [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$ is $1 - \mathsf{negl}(\lambda)$.

Proof. To prove the claim, we start analyzing the individual term as below.

- We first start by analyzing $\mathsf{ict}^\intercal_i \cdot \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ term.

$$\begin{aligned}\mathsf{ict}^\intercal_i \cdot \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) &= (\mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{IP}} + \mathbf{e}^\intercal_{\mathsf{IP},i} - (\mathbf{s}^*_i)^\intercal \otimes \mathbf{g}^\intercal) \cdot \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) \\ &= \mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{IP}} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) + \mathbf{e}^\intercal_{\mathsf{IP},i} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) - (\mathbf{s}^*_i)^\intercal (\mathbf{A}^*)^\intercal \\ &= \mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{mask}} + \mathbf{e}^\intercal_{\mathsf{IP},i} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) - (\mathbf{s}^*_i)^\intercal (\mathbf{A}^*)^\intercal\end{aligned} \tag{7}$$

  where the last equality follows as $\mathbf{A}_{\mathsf{mask}} = \mathbf{A}_{\mathsf{IP}} \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ from construction.

- We then analyze the term $\mathbf{c}^\intercal_{\mathbf{B},i} \cdot \mathbf{K}$.

$$\mathbf{c}^\intercal_{\mathbf{B},i} \cdot \mathbf{K} = (\mathbf{t}^\intercal_i \mathbf{B} + \mathbf{e}^\intercal_{\mathbf{B},i}) \cdot \mathbf{K} = \mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}} + \mathbf{e}^\intercal_{\mathbf{B},i} \mathbf{K} = \mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{err}} + \mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{mask}} + \mathbf{e}^\intercal_{\mathbf{B},i} \mathbf{K} \tag{8}$$

  where the last equality follows as $\mathbf{A}_{\mathrm{F}} = \mathbf{A}_{\mathsf{err}} + \mathbf{A}_{\mathsf{mask}}$ from the construction.

- Next, we analyze the term $\mathsf{bct}^\intercal_i \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i}$ for all $i \in [K]$, $\mathsf{xx} \in \{\mathsf{low}, \mathsf{high}\}$.

$$\begin{aligned}\mathsf{bct}^\intercal_i \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i} &= (\mathbf{t}^\intercal_i (\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}_i)) \otimes \mathbf{G}) + \mathbf{e}^\intercal_{\mathsf{att},i}) \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i} \\ &= \mathbf{t}^\intercal_i \mathbf{A}_{\mathsf{att}} \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}} - \mathbf{t}^\intercal_i \mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}(\mathsf{hct}_i) + \mathbf{e}^\intercal_{\mathsf{att},i} \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i} \\ &= \mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} - \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathbf{x}^\intercal_i) + \mathbf{e}^\intercal_{\mathsf{fhe},i} \mathbf{R}_{\mathrm{F}_{\mathsf{xx}},i} + \mathbf{e}^\intercal_{\mathsf{att},i} \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i} \\ &= \mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} - \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathbf{x}^\intercal_i) + \mathbf{e}^\intercal_{\mathrm{F}_{\mathsf{xx}},i}\end{aligned} \tag{9}$$

  where $\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}(\mathsf{hct}_i) = \mathbf{A}_{\mathsf{fhe},i} \mathbf{R}_{\mathrm{F}_{\mathsf{xx}},i} - \begin{pmatrix} 0^{(n+1)\times M} \\ \Delta \cdot \mathrm{F}_{\mathsf{xx}}(\mathbf{x}^\intercal_i) \end{pmatrix}$, and
  $\mathbf{e}^\intercal_{\mathrm{F}_{\mathsf{xx}},i} = \mathbf{e}^\intercal_{\mathsf{fhe},i} \mathbf{R}_{\mathrm{F}_{\mathsf{xx}},i} + \mathbf{e}^\intercal_{\mathsf{att},i} \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i}$.

- By Lemma 3.9, we have, for $\mathsf{xx} \in \{\mathsf{low}, \mathsf{high}\}$,

$$||\mathbf{R}^\intercal_{\mathrm{F}_{\mathsf{xx}},i}|| \leq (m+2)^d \lceil \log q \rceil \max_{k \in [L]} ||\mathbf{R}^\intercal_{i,k}|| \leq (m+2)^d \lceil \log q \rceil m \leq \widehat{\eta},$$

  and using depth bound from Section 3.4, we have

$$||(\mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i})^\intercal|| \leq (m+2)^{d_{\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}}} \lceil \log q \rceil \leq 2^{dO(\log \lambda)} \leq \widehat{\eta}$$

  where $\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}$ is the depth of the circuit $\mathsf{VEval}_{\mathrm{F}_{\mathsf{xx}}}$. Next, we have $||\mathbf{e}^\intercal_{\mathrm{F}_{\mathsf{xx}},i}|| \leq ||\mathbf{e}^\intercal_{\mathsf{fhe},i} \mathbf{R}_{\mathrm{F}_{\mathsf{xx}},i} + \mathbf{e}^\intercal_{\mathsf{att},i} \mathbf{H}^{\mathrm{F}_{\mathsf{xx}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i}|| \leq \sqrt{\lambda}\sigma_{\mathsf{fhe}} \widehat{\eta} m + \sqrt{\lambda}\sigma_{\mathsf{att}} \widehat{\eta}(L_{\mathbf{X}} + 1)m$.

- Now, for all $i \in [K]$, and $\mathsf{xx} \in \{\mathsf{low}, \mathsf{high}\}$

$$\begin{aligned}\left\lfloor \frac{\mathsf{bct}^\intercal_i \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}}, \mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{high},i}}{\Delta} \right\rfloor &= \left\lfloor \frac{\mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}_{\mathsf{high}}} - \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathbf{x}^\intercal_i) + \mathbf{e}^\intercal_{\mathrm{F}_{\mathsf{high}},i} + \mathbf{t}^\intercal_i \mathbf{W}_{\mathsf{high}} + \mathbf{e}^\intercal_{\mathsf{high},i}}{\Delta} \right\rfloor \\ &= \left\lfloor \frac{\mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}_{\mathsf{high}}} - \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathbf{x}^\intercal_i) + \mathbf{t}^\intercal_i \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor \\ &= \left\lfloor \frac{\mathbf{t}^\intercal_i \mathbf{A}_{\mathrm{F}_{\mathsf{high}}} + \mathbf{t}^\intercal_i \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor - \mathrm{F}_{\mathsf{high}}(\mathbf{x}^\intercal_i)\end{aligned} \tag{10}$$

The first equality follows because $\mathbf{c}_{\mathsf{xx},i}^\intercal = (c_{\mathsf{xx},i},\ldots,c_{\mathsf{xx},i}) = \mathbf{t}_i^\intercal(\mathbf{w}_{\mathsf{xx}},\ldots,\mathbf{w}_{\mathsf{xx}}) + (e_{\mathsf{xx},i},\ldots,e_{\mathsf{xx},i}) = \mathbf{t}_i^\intercal\mathbf{W}_{\mathsf{xx}} + \mathbf{e}_{\mathsf{xx},i}^\intercal$. The second equality holds with probability $1-\mathsf{negl}(\lambda)$. This is because

- For all $k \in [M]$, the marginal distribution of $\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}}[\cdot,k] + \mathbf{W}_{\mathsf{xx}}[\cdot,k]$ is uniformly random as $\mathbf{W}_{\mathsf{xx}}[\cdot,k] = \mathbf{w}_{\mathsf{xx}} \leftarrow \mathbb{Z}_q$.
- This implies that for each $k \in [M]$, the marginal probability that $\mathbf{t}_i(\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{W}_{\mathsf{xx}})[k] \bmod \Delta \geq \Delta - \|(\mathbf{e}_{\mathrm{F}_{\mathsf{xx}},i}^\intercal + \mathbf{e}_{\mathsf{xx},i}^\intercal)[k]\|$ is at most $\frac{\|(\mathbf{e}_{\mathrm{F}_{\mathsf{xx}},i}^\intercal + \mathbf{e}_{\mathsf{xx},i}^\intercal)[k]\|}{\Delta} \leq \frac{2^{3\lambda}\widehat{\eta}}{\Delta} \leq \mathsf{negl}(\lambda)$, since $2^{3\lambda}\widehat{\eta} \ll \Delta$.
- Hence, the probability that the second equality does not hold is $M \cdot \mathsf{negl}(\lambda)$ by union bound. This implies, the second equality holds with probability $1-\mathsf{negl}(\lambda)$.

- From Equation (10), we have, for all $i \in [K]$,

$$
\begin{aligned}
&\Delta \cdot \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{high},i}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{low},i}}{\Delta} \right\rfloor \qquad (11)\\
&= \Delta \left\lfloor \frac{\mathbf{t}_i^\intercal(\mathbf{A}_{\mathrm{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}})}{\Delta} \right\rfloor - \Delta \mathrm{F}_{\mathsf{high}}(\mathbf{x}_i^\intercal) + \left\lfloor \frac{\mathbf{t}_i^\intercal(\mathbf{A}_{\mathrm{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}})}{\Delta} \right\rfloor - \mathrm{F}_{\mathsf{low}}(\mathbf{x}_i^\intercal)\\
&= \Delta \mathbf{t}_i^\intercal \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \Delta \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{high}}} - \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathbf{x}_i^\intercal) + \mathbf{t}_i^\intercal \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor\\
&\quad + \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{low}}} - \mathrm{F}_{\mathsf{low}}(\mathbf{x}_i^\intercal)\\
&= \mathbf{t}_i^\intercal \left(\Delta \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor \right) - \mathrm{F}(\mathbf{x}_i^\intercal) + \Delta \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{high}}} + \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{low}}}\\
&= \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{err}} - \mathrm{F}(\mathbf{x}_i^\intercal) + \Delta \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{high}}} + \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{low}}}
\end{aligned}
$$

where $\mathbf{A}_{\mathsf{err}} = \Delta \cdot \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{high}}} + \mathbf{W}_{\mathsf{high}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{low}}} + \mathbf{W}_{\mathsf{low}}}{\Delta} \right\rfloor$ from the construction and $\mathrm{F}(\mathbf{x}_i^\intercal) = \Delta \cdot \mathrm{F}_{\mathsf{high}}(\mathbf{x}_i^\intercal) + \mathrm{F}_{\mathsf{low}}(\mathbf{x}_i^\intercal)$ by definition of F. We will now analyze the error $\mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{xx}}}$ for $\mathsf{xx} = \mathsf{high}$ and $\mathsf{xx} = \mathsf{low}$

$$
\begin{aligned}
\mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{xx}}} &= \left\lfloor \frac{\mathbf{t}_i^\intercal(\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{W}_{\mathsf{xx}})}{\Delta} \right\rfloor - \mathbf{t}_i^\intercal \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{W}_{\mathsf{xx}}}{\Delta} \right\rfloor \qquad (12)\\
&= \left\lfloor \mathbf{t}_i^\intercal \underbrace{\left(\frac{\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{W}_{\mathsf{xx}}}{\Delta} - \left\lfloor \frac{\mathbf{A}_{\mathrm{F}_{\mathsf{xx}}} + \mathbf{W}_{\mathsf{xx}}}{\Delta} \right\rfloor \right)}_{<1} \right\rfloor\\
&\implies \|\mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{xx}}}\| \leq n\|\mathbf{t}_i\|
\end{aligned}
$$

- Hence from Equation (8), (7) and (11), we have for all $i \in [K]$,

$$
\begin{aligned}
(\mathbf{b}_i^*)^\intercal &= \mathbf{c}^\intercal_{\mathbf{B},i} \cdot \mathbf{K} - \left(\Delta \cdot \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{high},i}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}_i^\intercal \cdot \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}_i} + \mathbf{c}^\intercal_{\mathsf{low},i}}{\Delta} \right\rfloor \right) - \mathsf{ict}_i^\intercal \cdot \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)\\
&= \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{err}} + \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{mask}} + \mathbf{e}^\intercal_{\mathbf{B},i}\mathbf{K} - \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{err}} + \mathrm{F}(\mathsf{sd}_i, \mathsf{bits}(i), 0) - \Delta \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{high}}}\\
&\quad - \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{low}}} - \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{mask}} - \mathbf{e}^\intercal_{\mathsf{IP},i}\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) + (\mathbf{s}_i^*)^\intercal(\mathbf{A}^*)^\intercal\\
&= (\mathbf{s}_i^*)^\intercal(\mathbf{A}^*)^\intercal + \underbrace{\mathsf{PRG}(\mathsf{sd}_i) + (\mathbf{e}_i'')^\intercal}_{=(\mathbf{e}_i^*)^\intercal}\\
&= (\mathbf{A}^*\mathbf{s}_i^* + \mathbf{e}_i^*)^\intercal
\end{aligned}
$$

where $\mathbf{x}_i^\intercal = (\mathsf{sd}_i, \mathsf{bits}(i), 0)$, $(\mathbf{e}_i'')^\intercal = \mathbf{e}^\intercal_{\mathbf{B},i}\mathbf{K} - \Delta \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{high}}} - \mathbf{e}^\intercal_{\mathbf{t}_i,\mathrm{F}_{\mathsf{low}}} - \mathbf{e}^\intercal_{\mathsf{IP},i}\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$ and $\mathrm{F}(\mathsf{sd}_i, \mathsf{bits}(i), 0) = \mathsf{PRG}(\mathsf{sd}_i) + \beta \cdot \mathsf{crs}_i = \mathsf{PRG}(\mathsf{sd}_i)$ as $\beta = 0$.

- Now we have $||(\mathbf{e}_i'')^\intercal|| \leq \lambda\sigma_{\mathbf{B}}\gamma mw + \sqrt{\lambda}\Delta\sigma_t + \sqrt{\lambda}\sigma_{\mathbf{t}} + \sqrt{\lambda}\sigma_{\mathsf{IP}} W \log q \leq O(\sigma_{\mathbf{B}}\gamma + \Delta\sigma_{\mathbf{t}} + W\sigma_{\mathsf{IP}})\mathrm{poly}(\lambda) \leq \eta$.
- This implies with probability $1 - \mathsf{negl}(\lambda)$, there exists $\mathbf{s}_i^*$ and $\mathbf{e}_i^* \in [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$, for which $\mathbf{b}_i^*$ is an LWE sample with respect to $\mathbf{A}^*$.

□

5.3 Security

Next, we prove the security of OLS.

Theorem 5.3. Assuming the sub-exponential hardness of LWE and that sub-exponential LSP holds (Assumption 4.1), the construction of OLS as in Section 5.2 is sub-exponentially secure as per the definition in Section 5.1.

Proof. We prove the theorem via a sequence of hybrids $\mathsf{Hyb}_0, \mathsf{Hyb}_1, \ldots, \mathsf{Hyb}_5$ where Hyb_0 is the same as the real world of OLS and Hyb_5 is the same as the simulated world of OLS. In particular, we prove that $\mathsf{Hyb}_0 \approx \mathsf{Hyb}_1 \approx \ldots \approx \mathsf{Hyb}_5$. Let $\mathcal{A}$ be any PPT adversary distinguishing the real and the simulated world of OLS.

The following tables describe the changes across the hybrids highlighted with red and contain only the term provided to $\mathcal{A}$. Additionally, when we say xx in the following hybrids, we say it holds for both $\mathsf{xx} = \mathsf{high}$ and $\mathsf{xx} = \mathsf{low}$. Also, in the GSW ciphertext hct_i, we only write i instead of $\mathsf{bits}(i)$ to keep the presentation simple.

	Hyb_0 (OLS real world)	Hyb_1	Hyb_2	Hyb_3
$\mathbf{A}^*$	with trapdoor $(\mathbf{A}_\tau^*)^{-1}$	– do –	$\leftarrow \mathbb{Z}_q^{M\times W}$	– do –
crs	$\leftarrow \mathbb{Z}_q^{M\times K}$	– do –	– do –	$\mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}$
$\mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}$	$\leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil\log q\rceil}, \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$	– do –	– do –	– do –
$\mathsf{ct}_i, \forall\, i \in [K] := \begin{cases} \mathbf{c}_{\mathbf{B},i}^\intercal \\ \mathsf{hct}_i \\ \mathsf{bct}_i^\intercal \\ \mathsf{ict}_i^\intercal \\ c_{\mathsf{xx},i} \end{cases}$	$\mathbf{t}_i^\intercal\mathbf{B} + \mathbf{e}_{\mathbf{B},i}^\intercal$, $\begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}_i^\intercal\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe},i}^\intercal \end{pmatrix}\mathbf{R}_i - (\mathsf{sd}_i, i, 0)\otimes\mathbf{G}$, $\mathbf{t}_i^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}_i))\otimes\mathbf{G}) + \mathbf{e}_{\mathsf{att},i}^\intercal$, $\mathbf{t}_i^\intercal\mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP},i}^\intercal - \mathbf{s}_i^{*\intercal}\otimes\mathbf{g}^\intercal$ $\mathbf{t}_i^\intercal\mathbf{w}_{\mathsf{high}} + e_{\mathsf{high},i}$, $\mathbf{t}_i^\intercal\mathbf{w}_{\mathsf{low}} + e_{\mathsf{low},i}$	– do –	– do –	– do –
sk	$\mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$	– do –	– do –	– do –
LWE secrets	$\mathbf{s}_1^*, \ldots, \mathbf{s}_i^*, \ldots, \mathbf{s}_K^*$, where $\forall\, i \in [K]$, $\mathbf{s}_i^* = \mathsf{LWESolve}_{(\mathbf{A}_\tau^*)^{-1}}(\mathsf{Sample}(\mathsf{crs}, \mathsf{pub}, i))$	directly w/o $(\mathbf{A}_\tau^*)^{-1}$	– do –	– do –
			In Hyb_3, $\widehat{\mathbf{S}} \leftarrow \mathbb{Z}_q^{W\times K}, \widehat{\mathbf{E}} \leftarrow \mathcal{D}_{\mathbb{Z},\widehat{\sigma}}^{M\times K}$	

	Hyb_4	Hyb_5 (OLS Sim world)
$\mathbf{A}^*$	$\leftarrow \mathbb{Z}_q^{M\times W}$	(with trapdoor $(\mathbf{A}_\tau^*)^{-1}$)
crs	$\mathbf{A}^*\widehat{\mathbf{S}}+\widehat{\mathbf{E}}$	– do –
$\mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}$	$\leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil\log q\rceil}, \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{x}}+1)m}$	– do –
$\mathsf{ct}_i, \forall i\in[K] := \begin{cases} \mathbf{c}_{\mathbf{B},i}^\intercal \\ \mathsf{hct}_i \\ \mathsf{bct}_i^\intercal \\ \mathsf{ict}_i^\intercal \\ c_{\mathsf{xx},i} \end{cases}$	$\mathbf{t}_i^\intercal\mathbf{B}+\mathbf{e}_{\mathbf{B},i}^\intercal$, $\begin{pmatrix}\bar{\mathbf{A}}_{\mathsf{fhe}}\\ \bar{\mathbf{t}}_i^\intercal\bar{\mathbf{A}}_{\mathsf{fhe}}+\mathbf{e}_{\mathsf{fhe},i}^\intercal\end{pmatrix}\mathbf{R}_i-(\mathsf{sd}_i,i,1)\otimes\mathbf{G}$, $\mathbf{t}_i^\intercal(\mathbf{A}_{\mathsf{att}}-(1,\mathsf{bits}(\mathsf{hct}_i))\otimes\mathbf{G})+\mathbf{e}_{\mathsf{att},i}^\intercal$, $\mathbf{t}_i^\intercal\mathbf{A}_{\mathsf{IP}}+\mathbf{e}_{\mathsf{IP},i}^\intercal-\mathbf{s}_i^{*\intercal}\otimes\mathbf{g}^\intercal$ $\mathbf{t}_i^\intercal\mathbf{w}_{\mathsf{high}}+e_{\mathsf{high},i}$, $\mathbf{t}_i^\intercal\mathbf{w}_{\mathsf{low}}+e_{\mathsf{low},i}$	– do –
sk	$\mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$	– do –
LWE secrets	$(\mathbf{s}_1^*+\widehat{\mathbf{s}}_1),\ldots,(\mathbf{s}_i^*+\widehat{\mathbf{s}}_i),\ldots,(\mathbf{s}_K^*+\widehat{\mathbf{s}}_K)$	– do –

Indistinguishability between the hybrids

Claim 5.4. $\mathsf{Hyb}_0 \equiv \mathsf{Hyb}_1$

Proof. Hyb_0 and Hyb_1 differ only in the way LWE secrets $\mathbf{s}_i^*$ for all $i\in[K]$, are computed - recomputed using a trapdoor for $\mathbf{A}^*$ in Hyb_0 and directly returned in Hyb_1. The two hybrids are equivalent in the adversary's view due to Theorem 3.7. □

Claim 5.5. $\mathsf{Hyb}_1 \approx_s \mathsf{Hyb}_2$.

Proof. Hyb_1 and Hyb_2 differ only in the way matrix $\mathbf{A}^*$ is sampled - sampled along with a trapdoor as $(\mathbf{A}^*,(\mathbf{A}_\tau^*)^{-1})\leftarrow\mathsf{TrapGen}(1^M,1^W,q)$ in Hyb_1 and randomly as $\mathbf{A}^*\leftarrow\mathbb{Z}_q^{M\times W}$ in Hyb_2. The two hybrids are statistically indistinguishable due to Theorem 3.7. Note that $(\mathbf{A}_\tau^*)^{-1}$ is not used in Hyb_2, hence the hybrid is well defined. □

Claim 5.6. $\mathsf{Hyb}_2 \approx_c \mathsf{Hyb}_3$.

Proof. Hyb_2 and Hyb_3 differ only in the way crs is generated - uniformly at random in Hyb_2 and computed as LWE matrix $\mathbf{A}^*\widehat{\mathbf{S}}+\widehat{\mathbf{E}}$ in Hyb_3 where $\widehat{\mathbf{S}}\leftarrow\mathbb{Z}_q^{W\times K}, \widehat{\mathbf{E}}\leftarrow\mathcal{D}_{\mathbb{Z},\widehat{\sigma}}^{M\times K}$. The indistinguishability between the two hybrids follows from the hardness of LWE, since $\widehat{\mathbf{S}}$ and $\widehat{\mathbf{E}}$ are not used anywhere else. □

We defer the proof of indistinguishability between Hyb_3 and Hyb_4 to Claim 5.8.

Claim 5.7. $\mathsf{Hyb}_4 \approx_s \mathsf{Hyb}_5$

Proof. In Hyb_5, $\mathbf{A}^*$ is sampled with a trapdoor i.e. $(\mathbf{A}^*,(\mathbf{A}_\tau^*)^{-1})\leftarrow\mathsf{TrapGen}(1^M,1^W,q)$ instead of sampling it uniformly at random. The statistical indistinguishability argument between the two hybrids again follows from Theorem 3.7. □

Claim 5.8. Assuming Assumption 4.1, $\mathsf{Hyb}_3 \approx_c \mathsf{Hyb}_4$.

Proof. To prove the above claim, we introduce sub-hybrids $\mathsf{G}_{\mathsf{ind}}$ for $\mathsf{ind} \in [0, K]$ between Hyb_3 and Hyb_4, where $\mathsf{G}_0 = \mathsf{Hyb}_3$ and $\mathsf{G}_K = \mathsf{Hyb}_4$. Then in Claim 5.9, we show that $\mathsf{G}_{\mathsf{ind}-1} \approx \mathsf{G}_{\mathsf{ind}}$ for all $\mathsf{ind} \in [K]$, which in turn implies $\mathsf{Hyb}_3 \approx \mathsf{Hyb}_4$ since K can be atmost $\mathrm{poly}(\lambda)$, for some polynomial poly.

$\mathsf{G}_{\mathsf{ind}}$ for $\mathsf{ind} \in [0, K]$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M \times W}$, $\mathsf{crs} = \mathbf{A}^* \widehat{\mathbf{S}} + \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1) \times W \lceil \log q \rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L_{\mathbf{X}}+1)m}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B},i}^\intercal, \mathsf{hct}_i, \\ \mathsf{bct}_i^\intercal, \\ \mathsf{ict}_i^\intercal, \\ c_{\mathsf{high},i}, c_{\mathsf{low},i} \end{pmatrix} = \begin{cases} \mathbf{t}_i^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B},i}^\intercal, \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}_i^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe},i}^\intercal \end{pmatrix} \mathbf{R}_i - (\mathsf{sd}_i, i, 1) \otimes \mathbf{G}, \\ \mathbf{t}_i^\intercal (\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}_i)) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att},i}^\intercal, \\ \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP},i}^\intercal - \mathbf{s}_i^{*\intercal} \otimes \mathbf{g}^\intercal, & i \le \mathsf{ind} \\ \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{high}} + e_{\mathsf{high},i}, \quad \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{low}} + e_{\mathsf{low},i} \\ \\ \mathbf{t}_i^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B},i}^\intercal, \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}_i^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe},i}^\intercal \end{pmatrix} \mathbf{R}_i - (\mathsf{sd}_i, i, 0) \otimes \mathbf{G}, \\ \mathbf{t}_i^\intercal (\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}_i)) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att},i}^\intercal, \\ \mathbf{t}_i^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP},i}^\intercal - \mathbf{s}_i^{*\intercal} \otimes \mathbf{g}^\intercal, & i > \mathsf{ind} \\ \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{high}} + e_{\mathsf{high},i}, \quad \mathbf{t}_i^\intercal \mathbf{w}_{\mathsf{low}} + e_{\mathsf{low},i} \end{cases}$$

$\mathsf{sk} = \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$

LWE secrets: $(\mathbf{s}_1^* + \widehat{\mathbf{s}}_1), \ldots, (\mathbf{s}_{\mathsf{ind}}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}), \mathbf{s}_{\mathsf{ind}+1}^*, \ldots, \mathbf{s}_K^*$

Claim 5.9. $\mathsf{G}_{\mathsf{ind}-1} \approx_c \mathsf{G}_{\mathsf{ind}}$ for all $\mathsf{ind} \in [K]$.

Proof. To prove the claim, we first define the following two simplified distributions, which focus only on $i = \mathsf{ind}$ based terms since the other terms corresponding to $i \neq \mathsf{ind}$ remain unchanged between the two hybrids.

$D_{\mathsf{ind},0}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$

$\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{t}^\intercal\mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal, & \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}^\intercal\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^\intercal \end{pmatrix}\mathbf{R} - (\mathsf{sd}, \mathsf{ind}, 0)\otimes\mathbf{G}, \\ \mathbf{t}^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}))\otimes\mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal, & \mathbf{t}^\intercal\mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal - \mathbf{s}^{*\intercal}\otimes\mathbf{g}^\intercal, \\ \mathbf{t}^\intercal\mathbf{w}_{\mathsf{high}} + e_{\mathsf{high}}, & \mathbf{t}^\intercal\mathbf{w}_{\mathsf{low}} + e_{\mathsf{low}} \end{cases}$$

$\mathsf{sk} = \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$

LWE secret: $\mathbf{s}^*$

$D_{\mathsf{ind},1}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$

$\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{t}^\intercal\mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal, & \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}^\intercal\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^\intercal \end{pmatrix}\mathbf{R} - (\mathsf{sd}, \mathsf{ind}, 1)\otimes\mathbf{G}, \\ \mathbf{t}^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}))\otimes\mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal, & \mathbf{t}^\intercal\mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal - \mathbf{s}^{*\intercal}\otimes\mathbf{g}^\intercal, \\ \mathbf{t}^\intercal\mathbf{w}_{\mathsf{high}} + e_{\mathsf{high}}, & \mathbf{t}^\intercal\mathbf{w}_{\mathsf{low}} + e_{\mathsf{low}} \end{cases}$$

$\mathsf{sk} = \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$

LWE secret: $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}$

In the above distributions, F is as defined in Assumption 4.1, with $\mathsf{crs} = \mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}$ and $\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$ are computed as in init algorithm.

We then show that, for all $\mathsf{ind} \in [K]$:

1. if $\mathcal{A}$ can distinguish between $\mathsf{G}_{\mathsf{ind}-1}$ and $\mathsf{G}_{\mathsf{ind}}$ with non-negligible probability, then there exists a PPT adversary $\mathcal{B}$ who distinguishes $D_{\mathsf{ind},0}$ and $D_{\mathsf{ind},1}$ with non-negligible probability.
2. $D_{\mathsf{ind},0} \approx_c D_{\mathsf{ind},1}$.

Proof of Step 1. $\mathcal{B}$ is defined as follows. It receives ind from $\mathcal{A}$, which it forwards to its challenger, $\mathcal{C}$ and receives $D_{\mathsf{ind},\theta} = (\mathbf{A}^*, \widehat{\mathbf{S}}, \widehat{\mathbf{E}}, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \mathbf{B}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}, \overline{\mathbf{A}}_{\mathsf{fhe}}, \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, c_{\mathsf{high}}, c_{\mathsf{low}}, \mathsf{sk}, \mathbf{s}^* + \theta\cdot\widehat{\mathbf{s}}_{\mathsf{ind}})$, where θ is the challenge bit of $\mathcal{C}$. It then computes $\mathsf{crs} = \mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}$. It samples $\mathbf{t}_i, \mathbf{s}_i^*$ for all $i \in [K]\setminus\{\mathsf{ind}\}$ on its own.It then computes all the terms corresponding to $i \neq \mathsf{ind}$, $\{\mathbf{c}_{\mathbf{B},i}^\intercal, \mathsf{hct}_i, \mathsf{bct}_i^\intercal, \mathsf{ict}_i^\intercal, c_{\mathsf{high},i}, c_{\mathsf{low},i}\}_{i\in[K]\setminus\{\mathsf{ind}\}}$ by itself, and substitutes the terms received from $\mathcal{C}$ for $i = \mathsf{ind}$ as $(\mathbf{c}_{\mathbf{B},\mathsf{ind}}^\intercal = \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}_{\mathsf{ind}} = \mathsf{hct}, \mathsf{bct}_{\mathsf{ind}}^\intercal = \mathsf{bct}^\intercal, \mathsf{ict}_{\mathsf{ind}}^\intercal = \mathsf{ict}^\intercal, c_{\mathsf{high},\mathsf{ind}} = c_{\mathsf{high}}, c_{\mathsf{low},\mathsf{ind}} = c_{\mathsf{low}})$. $\mathcal{B}$ then forwards $(\mathbf{A}^*, \mathsf{crs}, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}, \{\mathbf{c}_{\mathbf{B},i}^\intercal, \mathsf{hct}_i, \mathsf{bct}_i^\intercal, \mathsf{ict}_i^\intercal, c_{\mathsf{high},i}, c_{\mathsf{low},i}\}_{i\in[K]}, \mathsf{sk}, \{\mathbf{s}_i^* + \widehat{\mathbf{s}}_i\}_{i<\mathsf{ind}}, \mathbf{s}^* + \theta\cdot\widehat{\mathbf{s}}_{\mathsf{ind}}, \{\mathbf{s}_i^*\}_{i>\mathsf{ind}})$ to $\mathcal{A}$. In the end, $\mathcal{A}$ outputs θ', which $\mathcal{B}$ returns to $\mathcal{C}$. It is easy to see that if $\theta = 0$, then $\mathcal{B}$ simulates $\mathsf{G}_{\mathsf{ind}-1}$, else $\mathsf{G}_{\mathsf{ind}}$. Hence, the advantage of $\mathcal{B}$ is the same as the advantage of $\mathcal{A}$.

Proof of Step 2. To prove Step 2, we define the following intermediate hybrid distributions $\overline{\mathsf{G}}_{0,0}, \overline{\mathsf{G}}_{0,1}, \ldots, \overline{\mathsf{G}}_{0,5}$ between $D_{\mathsf{ind},0}$ and $D_{\mathsf{ind},1}$, with $D_{\mathsf{ind},0} = \overline{\mathsf{G}}_{0,0}$ and $D_{\mathsf{ind},1} = \overline{\mathsf{G}}_{0,5}$, and prove that $\overline{\mathsf{G}}_{0,0} \approx \overline{\mathsf{G}}_{0,1} \approx \ldots \approx \overline{\mathsf{G}}_{0,5}$.

$\overline{\mathsf{G}}_{0,0} = D_{\mathsf{ind},0}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$

$\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{t}^\intercal\mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal, & \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}^\intercal\bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^\intercal \end{pmatrix}\mathbf{R} - (\mathsf{sd}, \mathsf{ind}, 0)\otimes\mathbf{G}, \\ \mathbf{t}^\intercal(\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct}))\otimes\mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal, & \mathbf{t}^\intercal\mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal - \mathbf{s}^{*\intercal}\otimes\mathbf{g}^\intercal, \\ \mathbf{t}^\intercal\mathbf{w}_{\mathsf{high}} + e_{\mathsf{high}}, & \mathbf{t}^\intercal\mathbf{w}_{\mathsf{low}} + e_{\mathsf{low}} \end{cases}$$

$\mathsf{sk} = \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$

LWE secrets: $\mathbf{s}^*$

$\overline{\mathsf{G}}_{0,1}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$

$\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{w}_{\mathbf{B}}^\intercal \leftarrow \mathbb{Z}_q^{1\times mw}, \mathbf{W}_{\mathbf{X}} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}, \\ \mathbf{w}_{\mathsf{att}}^\intercal \leftarrow \mathbb{Z}_q^{1\times(L_{\mathbf{X}}+1)m}, \mathbf{w}_{\mathsf{IP}}^\intercal - (\mathbf{s}^*)^\intercal\otimes\mathbf{g}^\intercal; \ (\mathbf{w}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{W\lceil \log q\rceil}), \\ w_{\mathsf{high}} \leftarrow \mathbb{Z}_q, \ w_{\mathsf{low}} \leftarrow \mathbb{Z}_q \end{cases}$$

$$\mathsf{sk} = \begin{pmatrix} \mathbf{B} \\ \mathbf{w}_{\mathbf{B}}^\intercal \end{pmatrix}^{-1}\left(\begin{pmatrix} \mathbf{A}_{\mathrm{F}} \\ \mathbf{d}^\intercal + \mathbf{c}'^\intercal \end{pmatrix}\right), \text{for } \mathbf{c}' \leftarrow [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$$

$$\mathbf{d}^\intercal = \Delta\left\lfloor \frac{\mathsf{bct}^\intercal\mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^\intercal}{\Delta}\right\rfloor + \left\lfloor \frac{\mathsf{bct}^\intercal\mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^\intercal}{\Delta}\right\rfloor + \mathbf{w}_{\mathsf{IP}}^\intercal\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal)$$

LWE secret: $\mathbf{s}^*$

$\overline{\mathsf{G}}_{0,2}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times(L_{\mathbf{X}}+1)m}$

$\mathbf{B}$, $\mathbf{w}_{\mathsf{high}}$, $\mathbf{w}_{\mathsf{low}}$, $\overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{w}_{\mathbf{B}}^\intercal \leftarrow \mathbb{Z}_q^{1\times mw}, \mathbf{W}_{\mathbf{X}} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}, \\ \mathbf{w}_{\mathsf{att}}^\intercal \leftarrow \mathbb{Z}_q^{1\times(L_{\mathbf{X}}+1)m}, \mathbf{w}_{\mathsf{IP}}^\intercal \leftarrow \mathbb{Z}_q^{1\times W\lceil \log q\rceil}) \\ w_{\mathsf{high}} \leftarrow \mathbb{Z}_q, \ w_{\mathsf{low}} \leftarrow \mathbb{Z}_q \end{cases}$$

$$\mathsf{sk} = \begin{pmatrix} \mathbf{B} \\ \mathbf{w}_{\mathbf{B}}^\intercal \end{pmatrix}^{-1}\left(\begin{pmatrix} \mathbf{A}_{\mathrm{F}} \\ \mathbf{d}^\intercal + \mathbf{c}'^\intercal \end{pmatrix}\right), \text{for } \mathbf{c}' \leftarrow [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$$

$$\mathbf{d}^\intercal = \Delta\left\lfloor \frac{\mathsf{bct}^\intercal\mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^\intercal}{\Delta}\right\rfloor + \left\lfloor \frac{\mathsf{bct}^\intercal\mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^\intercal}{\Delta}\right\rfloor + \mathbf{w}_{\mathsf{IP}}^\intercal\mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) + (\mathbf{A}^*\mathbf{s}^*)^\intercal$$

LWE secret: $\mathbf{s}^*$

$\overline{\mathsf{G}}_{0,3}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times (L_{\mathbf{X}}+1)m}$

$\mathbf{B}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}, \overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{w}_{\mathbf{B}}^\intercal \leftarrow \mathbb{Z}_q^{1\times mw}, \mathbf{W}_{\mathbf{X}} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}, \\ \mathbf{w}_{\mathsf{att}}^\intercal \leftarrow \mathbb{Z}_q^{1\times (L_{\mathbf{X}}+1)m}, \mathbf{w}_{\mathsf{IP}}^\intercal \leftarrow \mathbb{Z}_q^{1\times W\lceil \log q\rceil}) \\ w_{\mathsf{high}} \leftarrow \mathbb{Z}_q, \quad w_{\mathsf{low}} \leftarrow \mathbb{Z}_q \end{cases}$$

$$\mathsf{sk} = \begin{pmatrix} \mathbf{B} \\ \mathbf{w}_{\mathbf{B}}^\intercal \end{pmatrix}^{-1} \left(\begin{pmatrix} \mathbf{A}_{\mathrm{F}} \\ \mathbf{d}^\intercal + \mathbf{c}'^\intercal + \widehat{\mathbf{e}}_{\mathsf{ind}}^\intercal \end{pmatrix} \right), \text{ for } \mathbf{c}' \leftarrow [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$$

$$\mathbf{d}^\intercal = \Delta \left\lfloor \frac{\mathsf{bct}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^\intercal}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^\intercal}{\Delta} \right\rfloor + \mathbf{w}_{\mathsf{IP}}^\intercal \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) + (\mathbf{A}^*(\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}^\intercal))^\intercal$$

LWE secret: $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}^\intercal$

$\overline{\mathsf{G}}_{0,4}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times (L_{\mathbf{X}}+1)m}$

$\mathbf{B}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}, \overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{w}_{\mathbf{B}}^\intercal \leftarrow \mathbb{Z}_q^{1\times mw}, \mathbf{W}_{\mathbf{X}} \leftarrow \mathbb{Z}_q^{(n+1)\times mL}, \\ \mathbf{w}_{\mathsf{att}}^\intercal \leftarrow \mathbb{Z}_q^{1\times (L_{\mathbf{X}}+1)m}, \mathbf{w}_{\mathsf{IP}}^\intercal - (\mathbf{s}^*)^\intercal \otimes \mathbf{g}^\intercal; (\mathbf{w}_{\mathsf{IP}}^\intercal \leftarrow \mathbb{Z}_q^{W\lceil \log q\rceil}) \\ w_{\mathsf{high}} \leftarrow \mathbb{Z}_q, \quad w_{\mathsf{low}} \leftarrow \mathbb{Z}_q \end{cases}$$

$$\mathsf{sk} = \begin{pmatrix} \mathbf{B} \\ \mathbf{w}_{\mathbf{B}}^\intercal \end{pmatrix}^{-1} \left(\begin{pmatrix} \mathbf{A}_{\mathrm{F}} \\ \mathbf{d}^\intercal + \mathbf{c}'^\intercal \end{pmatrix} \right), \text{ for } \mathbf{c}' \leftarrow [-B_{\mathsf{OLS}}, B_{\mathsf{OLS}}]^M$$

$$\mathbf{d}^\intercal = \Delta \left\lfloor \frac{\mathsf{bct}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^\intercal}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}^\intercal \mathbf{H}^{\mathrm{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^\intercal}{\Delta} \right\rfloor + \mathbf{w}_{\mathsf{IP}}^\intercal \mathbf{G}^{-1}((\mathbf{A}^*)^\intercal) + \underbrace{(\mathbf{A}^* \widehat{\mathbf{s}}_{\mathsf{ind}} + \widehat{\mathbf{e}}_{\mathsf{ind}})^\intercal}_{\mathsf{crs}_{\mathsf{ind}}^\intercal \text{ in } \mathrm{F}}$$

LWE secret: $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}^\intercal$

$\overline{\mathsf{G}}_{0,5} = D_{\mathsf{ind},1}$

$\mathbf{A}^* \leftarrow \mathbb{Z}_q^{M\times W}$, $\widehat{\mathbf{S}}, \widehat{\mathbf{E}}$, $\mathbf{A}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{(n+1)\times W\lceil \log q\rceil}$, $\mathbf{A}_{\mathsf{att}} \leftarrow \mathbb{Z}_q^{(n+1)\times (L_{\mathbf{X}}+1)m}$

$\mathbf{B}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}, \overline{\mathbf{A}}_{\mathsf{fhe}}$

$$\begin{pmatrix} \mathbf{c}_{\mathbf{B}}^\intercal, \mathsf{hct}, \\ \mathsf{bct}^\intercal, \mathsf{ict}^\intercal, \\ c_{\mathsf{high}}, c_{\mathsf{low}} \end{pmatrix} = \begin{cases} \mathbf{t}^\intercal \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\intercal, \quad \begin{pmatrix} \bar{\mathbf{A}}_{\mathsf{fhe}} \\ \bar{\mathbf{t}}^\intercal \bar{\mathbf{A}}_{\mathsf{fhe}} + \mathbf{e}_{\mathsf{fhe}}^\intercal \end{pmatrix} \mathbf{R} - (\mathsf{sd}, \mathsf{ind}, 1) \otimes \mathbf{G}, \\ \mathbf{t}^\intercal (\mathbf{A}_{\mathsf{att}} - (1, \mathsf{bits}(\mathsf{hct})) \otimes \mathbf{G}) + \mathbf{e}_{\mathsf{att}}^\intercal, \quad \mathbf{t}^\intercal \mathbf{A}_{\mathsf{IP}} + \mathbf{e}_{\mathsf{IP}}^\intercal - \mathbf{s}^{*\intercal} \otimes \mathbf{g}^\intercal, \\ \mathbf{t}^\intercal \mathbf{w}_{\mathsf{high}} + e_{\mathsf{high}}, \quad \mathbf{t}^\intercal \mathbf{w}_{\mathsf{low}} + e_{\mathsf{low}} \end{cases}$$

$\mathsf{sk} = \mathbf{B}_\gamma^{-1}(\mathbf{A}_{\mathrm{F}})$

LWE secret: $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}^\intercal$

Indistinguishability between the intermediate hybrids.

Claim 5.10. Assuming Assumption 4.1, $\overline{\mathsf{G}}_{0,0} \approx_c \overline{\mathsf{G}}_{0,1}$.

Proof. Let $\mathcal{A}$ be a PPT adversary distinguishing $\overline{\mathsf{G}}_{0,0}$ and $\overline{\mathsf{G}}_{0,1}$. We construct a PPT adversary $\mathcal{B}$ that breaks Assumption 4.1. $\mathcal{B}$ is defined as follows.

1. $\mathcal{B}$ sends $\mathsf{index} = \mathsf{ind}, \beta = 0$ to its own challenger $\mathcal{C}$ and receives the following,

$$\left(\bar{\mathbf{A}}_{\mathsf{fhe}}, \mathbf{B}, \mathbf{D} = (\mathbf{A}_{\mathsf{att}}, \mathbf{A}_{\mathsf{IP}}, \mathbf{w}_{\mathsf{high}}, \mathbf{w}_{\mathsf{low}}), \mathbf{A}^*, \widehat{\mathbf{E}}, \widehat{\mathbf{S}}, \mathbf{c}_{\mathbf{B}}^{\mathsf{T}}, \mathbf{X}, \mathbf{c}_{\mathbf{D}}^{\mathsf{T}} = (\mathbf{c}_{\mathsf{att}}^{\mathsf{T}}, \mathbf{c}_{\mathsf{IP}}^{\mathsf{T}}, c_{\mathsf{high}}, c_{\mathsf{low}}), \mathbf{K}\right)$$

2. Samples $\mathbf{s}^* \leftarrow \mathbb{Z}_q^W$, and compute $\mathsf{ict}^{\mathsf{T}} = \mathbf{c}_{\mathsf{IP}}^{\mathsf{T}} - (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}$.
3. Sets $\mathsf{sk} = \mathbf{K}$, $\mathsf{hct} = \mathbf{X}$, $\mathsf{bct} = \mathbf{c}_{\mathsf{att}}$.
4. Output $(\mathbf{A}^*, \widehat{\mathbf{S}}, \widehat{\mathbf{E}}, \mathbf{A}_{\mathsf{IP}}, \mathbf{A}_{\mathsf{att}}), (\mathbf{c}_{\mathbf{B}}^{\mathsf{T}}, \mathsf{hct}, \mathsf{bct}, \mathsf{ict}, c_{\mathsf{high}}, c_{\mathsf{low}}, \mathsf{sk}), \mathbf{s}^*)$ to $\mathcal{A}$.

It is easy to see that if $\mathcal{C}$ sent the real-world distribution of Assumption 4.1, then $\mathcal{B}$ simulated the distribution $\overline{\mathsf{G}}_{0,0}$, else it simulated the distribution $\overline{\mathsf{G}}_{0,1}$. Hence, the proof of Claim 5.10. □

Claim 5.11. $\overline{\mathsf{G}}_{0,1} \equiv \mathsf{G}_{0,2}$.

Proof. The only difference between $\overline{\mathsf{G}}_{0,1}$ and $\overline{\mathsf{G}}_{0,2}$ are:

- in $\overline{\mathsf{G}}_{0,1}$, we have $\mathsf{ict}^{\mathsf{T}} = \mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} - (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}$, while in $\overline{\mathsf{G}}_{0,2}$, we have $\mathsf{ict}^{\mathsf{T}} = \mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} = (\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} + (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}) - (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}$, where $\mathbf{w}_{\mathsf{IP}} \leftarrow \mathbb{Z}_q^{W\lceil \log q \rceil}$.
- in $\overline{\mathsf{G}}_{0,1}$,

$$\mathbf{d}^{\mathsf{T}} = \Delta \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^{\mathsf{T}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^{\mathsf{T}}}{\Delta} \right\rfloor + \mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} \mathbf{G}^{-1}((\mathbf{A}^*)^{\mathsf{T}}),$$

  in $\overline{\mathsf{G}}_{0,2}$,

$$\begin{aligned}\mathbf{d}^{\mathsf{T}} &= \Delta \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^{\mathsf{T}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^{\mathsf{T}}}{\Delta} \right\rfloor \\ &\quad + \mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} \mathbf{G}^{-1}((\mathbf{A}^*)^{\mathsf{T}}) + (\mathbf{A}^* \mathbf{s}^*)^{\mathsf{T}} \\ &= \Delta \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{high}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{high}}^{\mathsf{T}}}{\Delta} \right\rfloor + \left\lfloor \frac{\mathsf{bct}^{\mathsf{T}} \mathbf{H}^{\mathsf{F}_{\mathsf{low}}}_{\mathbf{A}_{\mathsf{att}},\mathsf{hct}} + \mathbf{c}_{\mathsf{low}}^{\mathsf{T}}}{\Delta} \right\rfloor \\ &\quad + (\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} + (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}) \mathbf{G}^{-1}((\mathbf{A}^*)^{\mathsf{T}})\end{aligned}$$

  .

Thus, effectively, the only change between the two hybrids is that $\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}}$ in the former is changed to $\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} + (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}$ in the latter. Hence, the two hybrids are equivalent because for $\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} \leftarrow \mathbb{Z}_q^{1\times W\lceil \log q \rceil}$, $\mathbf{w}_{\mathsf{IP}}^{\mathsf{T}} + (\mathbf{s}^*)^{\mathsf{T}} \otimes \mathbf{g}^{\mathsf{T}}$ is also uniformly distributed in $\mathbb{Z}_q^{1\times W\lceil \log q \rceil}$. □

Claim 5.12. $\overline{\mathsf{G}}_{0,2} \approx_s \overline{\mathsf{G}}_{0,3}$.

Proof. We can observe that the only difference between the two hybrids is the following:

1. $\mathbf{s}^*$ in the former changed to $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}$ in the latter, where $\mathbf{s}^* \leftarrow \mathbb{Z}_q^W$.
2. Addition of $\widehat{\mathbf{e}}_{\mathsf{ind}}$ to $\mathbf{d}^{\mathsf{T}} + \mathbf{c}'^{\mathsf{T}}$ in the computation of key sk in the latter hybrid.

Hence, the two hybrids are statistically indistinguishable because

1. For $\mathbf{s}^* \leftarrow \mathbb{Z}_q^W$, $\mathbf{s}^* + \widehat{\mathbf{s}}_{\mathsf{ind}}$ is also uniformly distributed in $\mathbb{Z}_q^W$.

2. $\mathbf{c}'^{\intercal} \approx_s \mathbf{c}'^{\intercal} + \widehat{\mathbf{e}}_{\mathsf{ind}}^{\intercal}$ by noise flooding (Lemma 3.5), since $\sqrt{\lambda}\widehat{\sigma}$ is exponentially less than B_{OLS} from the parameter setting.

□

Claim 5.13. $\overline{\mathsf{G}}_{0,3} \equiv \overline{\mathsf{G}}_{0,4}$.

Proof. The two hybrids differ only as follows:

- Replacing $\mathsf{ict}^{\intercal} = \mathbf{w}_{\mathsf{IP}}^{\intercal}$ in $\overline{\mathsf{G}}_{0,3}$ with $\mathbf{w}_{\mathsf{IP}}^{\intercal} - (\mathbf{s}^*)^{\intercal} \otimes \mathbf{g}^{\intercal}$ in $\overline{\mathsf{G}}_{0,4}$ by the similar arguments as in the proof of Claim 5.11.
- Renaming $\mathbf{d} + \widehat{\mathbf{e}}_{\mathsf{ind}}$ as $\mathbf{d}$. This is just a conceptual change and does not introduce any real change between the two hybrids.

Hence, the two hybrids are equivalent, since for $\mathbf{w}_{\mathsf{IP}}^{\intercal} \leftarrow \mathbb{Z}_q^{1\times W\lceil \log q\rceil}$, $\mathbf{w}_{\mathsf{IP}}^{\intercal} - (\mathbf{s}^*)^{\intercal} \otimes \mathbf{g}^{\intercal}$ is also uniformly distributed in $\mathbb{Z}_q^{1\times W\lceil \log q\rceil}$. □

Claim 5.14. Assuming Assumption 4.1, $\overline{\mathsf{G}}_{0,4} \approx_c \overline{\mathsf{G}}_{0,5}$.

Proof. The indistinguishability argument is the same as that in Claim 5.10 except that the reduction sets and sends $\beta = 1$ instead of $\beta = 0$ to $\mathcal{C}$.

□□□□

Acknowledgments. The first and the second authors are supported by the Arvind Raghunathan Center for the Theory of Computing (aRtCS), the Mphasis F1 Foundation, and the National Quantum Mission initiative of the Department of Science and Technology, India. Shota Yamada's work was partially supported by JST K Program JPMJKP24U3 and JST CREST JPMJCR22M1.

References

[ABB10a] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 553–572. Springer, Heidelberg, May / June 2010. (Cited on page 16, 17.)

[ABB10b] 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, Heidelberg, August 2010. (Cited on page 16.)

[AGIS14] Prabhanjan Ananth, Divya Gupta, Yuval Ishai, and Amit Sahai. Optimizing obfuscation: Avoiding barrington's theorem. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, CCS '14, page 646–658, New York, NY, USA, 2014. Association for Computing Machinery. (Cited on page 3.)

[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, Heidelberg, August 2017. (Cited on page 26.)

[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, Heidelberg, May 2019. (Cited on page 3, 15.)

[AJS18] Prabhanjan Ananth, Aayush Jain, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: iO from LWE, bilinear maps, and weak pseudorandomness. Cryptology ePrint Archive, Report 2018/615, 2018. https://eprint.iacr.org/2018/615. (Cited on page 3.)

[Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 99–108, New York, NY, USA, 1996. Association for Computing Machinery. (Cited on page 18.)

[AKY24a] Shweta Agrawal, Simran Kumari, and Shota Yamada. Attribute Based Encryption for Turing Machines from Lattices. In *Crypto*, 2024. (Cited on page 9.)

[AKY24b] Shweta Agrawal, Simran Kumari, and Shota Yamada. Pseudorandom Multi-Input Functional Encryption and Applications, 2024. (Cited on page 6.)

[AKY25] Shweta Agrawal, Simran Kumari, and Shota Yamada. Pseudorandom fe, io and applications. In *Theory of Cryptography: 23rd International Conference, TCC 2025, Aarhus, Denmark, December 1–5, 2025, Proceedings, Part II*, page 222–258, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 3, 6, 9.)

[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, Heidelberg, August 2016. (Cited on page 8.)

[AMR25] Damiano Abram, Giulio Malavolta, and Lawrence Roy. Key-homomorphic computations for ram: Fully succinct randomised encodings and more. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 236–268, Cham, 2025. Springer Nature Switzerland. (Cited on page 4.)

[AMYY25] Shweta Agrawal, Anuja Modi, Anshu Yadav, and Shota Yamada. Zeroizing attacks against evasive & circular evasive lwe. In *Theory of Cryptography: 23rd International Conference, TCC 2025, Aarhus, Denmark, December 1–5, 2025, Proceedings, Part II*, page 259–290, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 3, 4, 5, 6, 10, 13, 14, 15, 25, 26, 27, 28.)

[AMYY26] Shweta Agrawal, Anuja Modi, Anshu Yadav, and Shota Yamada. Evasive inner product functional encryption is broken too. *IACR Communications in Cryptology*, 3(3), 2026. (Cited on page 4, 5.)

[AP20] Shweta Agrawal and Alice Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 110–140. Springer, Heidelberg, May 2020. (Cited on page 3.)

[ARYY23] Shweta Agrawal, Mélissa Rossi, Anshu Yadav, and Shota Yamada. Constant input attribute based (and predicate) encryption from evasive and tensor LWE. In *CRYPTO 2023, Part IV*, LNCS, pages 532–564. Springer, Heidelberg, August 2023. (Cited on page 25.)

[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, *Advances in Cryptology – EUROCRYPT 2017*, pages 152–181, Cham, 2017. Springer International Publishing. (Cited on page 3.)

[BBKK18] Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). In *Advances in Cryptology–EUROCRYPT 2018: 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29-May 3, 2018 Proceedings, Part II 37*, pages 649–679. Springer, 2018. (Cited on page 29.)

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for iO: Circular-secure LWE suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. https://eprint.iacr.org/2020/1024. (Cited on page 3.)

[BDGM23] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate io from homomorphic encryption schemes. *J. Cryptol.*, 36(3), June 2023. (Cited on page 3.)

[BDJ+24] Pedro Branco, Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Pseudorandom obfuscation and applications. Cryptology ePrint Archive, Paper 2024/1742, 2024. (Cited on page 3, 4, 5, 6, 15, 26, 28.)

[BGG+14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit abe and compact garbled circuits. In *EUROCRYPT*. Springer, 2014. (Cited on page 9.)

[BGH+15] Zvika Brakerski, Craig Gentry, Shai Halevi, Tancrède Lepoint, Amit Sahai, and Mehdi Tibouchi. Cryptanalysis of the quadratic zero-testing of GGH. Cryptology ePrint Archive, Report 2015/845, 2015. https://eprint.iacr.org/2015/845. (Cited on page 3.)

[BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001. (Cited on page 3.)

[BGK+13] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. Cryptology ePrint Archive, Report 2013/631, 2013. https://eprint.iacr.org/2013/631. (Cited on page 3.)

[BHJ+19] Boaz Barak, Samuel B Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In *Advances in Cryptology–EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I 38*, pages 226–250. Springer, 2019. (Cited on page 29.)

[BIJ+20] James Bartusek, Yuval Ishai, Aayush Jain, Fermi Ma, Amit Sahai, and Mark Zhandry. Affine determinant programs: A framework for obfuscation and witness encryption. In Thomas Vidick, editor, *ITCS 2020*, volume 151, pages 82:1–82:39. LIPIcs, January 2020. (Cited on page 3.)

[BLP+13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013. (Cited on page 16, 18.)

[BMSZ16] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. Post-zeroizing obfuscation: New mathematical tools, and the case of evasive circuits. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 764–791. Springer, Heidelberg, May 2016. (Cited on page 3.)

[BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 1–25. Springer, Heidelberg, February 2014. (Cited on page 3.)

[BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 264–302. Springer, Heidelberg, November 2017. (Cited on page 19.)

[BUW24] Chris Brzuska, Akin Ünal, and Ivy K. Y. Woo. Evasive lwe assumptions: Definitions, classes, and counterexamples. In *Advances in Cryptology – ASIACRYPT 2024: 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9–13, 2024, Proceedings. Part IV*, page 418–449, Berlin, Heidelberg, 2024. Springer-Verlag. (Cited on page 4, 15, 26, 27.)

[BWZ14] Dan Boneh, David J. Wu, and Joe Zimmerman. Immunizing multilinear maps against zeroizing attacks. Cryptology ePrint Archive, Report 2014/930, 2014. https://eprint.iacr.org/2014/930. (Cited on page 3.)

[CCMR24] Ran Canetti, Claudio Chamon, Eduardo R. Mucciolo, and Andrei E. Ruckenstein. Towards general-purpose program obfuscation via local mixing. In *Theory of Cryptography: 22nd International Conference, TCC 2024, Milan, Italy, December 2–6, 2024, Proceedings, Part IV*, page 37–70, Berlin, Heidelberg, 2024. Springer-Verlag. (Cited on page 3.)

[CGH+15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrède 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, Heidelberg, August 2015. (Cited on page 3.)

[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, Heidelberg, May / June 2010. (Cited on page 16.)

[CHL+15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, Heidelberg, April 2015. (Cited on page 3, 25.)

[CHW25] Jeffrey Champion, Yao-Ching Hsieh, and David J. Wu. Registered abe and adaptively-secure broadcast encryption from succinct lwe. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology – CRYPTO 2025*, pages 3–34, Cham, 2025. Springer Nature Switzerland. (Cited on page 4.)

[CLR15] Jung Hee Cheon, Changmin Lee, and Hansol Ryu. Cryptanalysis of the new CLT multilinear maps. Cryptology ePrint Archive, Report 2015/934, 2015. https://eprint.iacr.org/2015/934. (Cited on page 3.)

[CLT13] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013*, pages 476–493, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. (Cited on page 3.)

[CLT15] Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 267–286. Springer, Heidelberg, August 2015. (Cited on page 3.)

[CLW25] Valerio Cini, Russell W. F. Lai, and Ivy K. Y. Woo. Lattice-based obfuscation from ntru and equivocal lwe. In *Advances in Cryptology – CRYPTO 2025: 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2025, Proceedings, Part VII*, page 39–72, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 3, 6.)

[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, Heidelberg, August 2018. (Cited on page 25.)

[CW24] Jeffrey Champion and David J. Wu. Distributed broadcast encryption from lattices. In *Theory of Cryptography: 22nd International Conference, TCC 2024, Milan, Italy, December 2–6, 2024, Proceedings, Part III*, page 156–189, Berlin, Heidelberg, 2024. Springer-Verlag. (Cited on page 4.)

[CW25] Valerio Cini and Hoeteck Wee. Faster abe for turing machines from circular evasive lwe. In Serge Fehr and Pierre-Alain Fouque, editors, *Advances in Cryptology – EUROCRYPT 2025*, pages 94–125, Cham, 2025. Springer Nature Switzerland. (Cited on page 13.)

[DGG+18] Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Pratyay Mukherjee. Obfuscation from low noise multilinear maps. In Debrup Chakraborty and Tetsu Iwata, editors, *INDOCRYPT 2018*, volume 11356 of *LNCS*, pages 329–352. Springer, Heidelberg, December 2018. (Cited on page 3.)

[DJM+25] Nico Döttling, Abhishek Jain, Giulio Malavolta, Surya Mathialagan, and Vinod Vaikuntanathan. Simple and general counterexamples for private-coin evasive lwe. In *CRYPTO*. Springer-Verlag, 2025. (Cited on page 4, 6, 15, 26.)

[DQV+21] Lalita Devadas, Willy Quach, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Succinct lwe sampling, random polynomials, and obfuscation. In *Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II 19*, pages 256–287. Springer, 2021. (Cited on page 3, 28.)

[GGH+13] 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. (Cited on page 3.)

[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, Heidelberg, March 2015. (Cited on page 3.)

[GGH+16] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing*, 45(3):882–929, 2016. (Cited on page 3.)

[GJK18] Craig Gentry, Charanjit S. Jutla, and Daniel Kane. Obfuscation using tensor products. Cryptology ePrint Archive, Report 2018/756, 2018. https://eprint.iacr.org/2018/756. (Cited on page 3.)

[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 *Advances in Cryptology – EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part III*, page 97–126, Berlin, Heidelberg, 2021. Springer-Verlag. (Cited on page 3.)

[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. (Cited on page 3, 29.)

[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. (Cited on page 16.)

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013. (Cited on page 5, 18.)

[Hal15] Shai Halevi. Graded encoding, variations on a scheme. Cryptology ePrint Archive, Report 2015/866, 2015. https://eprint.iacr.org/2015/866. (Cited on page 3.)

[HHY25] Tzu-Hsiang Huang, Wei-Hsiang Hung, and Shota Yamada. A note on obfuscation-based attacks on private-coin evasive LWE. *IACR Cryptol. ePrint Arch.*, page 421, 2025. (Cited on page 15, 26.)

[HJ15] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. Cryptology ePrint Archive, Report 2015/301, 2015. https://eprint.iacr.org/2015/301. (Cited on page 3.)

[HJL21] Sam Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying io. In *Advances in Cryptology–CRYPTO 2021: 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16–20, 2021, Proceedings, Part II 41*, pages 673–700. Springer, 2021. (Cited on page 3, 7, 15, 25, 29.)

[HJL25] 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). In *Advances in Cryptology – CRYPTO 2025: 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2025, Proceedings, Part VII*, page 3–38, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 3, 4, 6, 10, 13, 14, 25, 27, 28, 29.)

[HLL23] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 415–434. IEEE, 2023. (Cited on page 6, 9, 13, 18, 19.)

[JLLS23] Aayush Jain, Huijia Lin, Paul Lou, and Amit Sahai. Polynomial-time cryptanalysis of subspace flooding assumption for post-quantum io. In *Advances in Cryptology – EUROCRYPT 2023: 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part I*, page 205–235, Berlin, Heidelberg, 2023. Springer-Verlag. (Cited on page 3, 28.)

[JLMS19] Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over r to build io. In *Advances in Cryptology – EUROCRYPT 2019: 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19–23, 2019, Proceedings, Part I*, page 251–281, Berlin, Heidelberg, 2019. Springer-Verlag. (Cited on page 3.)

[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. (Cited on page 3.)

[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, Heidelberg, May / June 2022. (Cited on page 3.)

[Lin16] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 28–57. Springer, Heidelberg, May 2016. (Cited on page 3.)

[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, Heidelberg, August 2017. (Cited on page 3.)

[LLL25] Hanjun Li, Huijia Lin, and George Lu. Succinct garbled circuits with low-depth garbling algorithms. Cryptology ePrint Archive, Paper 2025/2308, 2025. (Cited on page 4.)

[LM18] Huijia Lin and Christian Matt. Pseudo flawed-smudging generators and their application to indistinguishability obfuscation. Cryptology ePrint Archive, Report 2018/646, 2018. https://eprint.iacr.org/2018/646. (Cited on page 3.)

[LPST16] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation with non-trivial efficiency. In *Proceedings, Part II, of the 19th IACR International Conference on Public-Key Cryptography — PKC 2016 - Volume 9615*, page 447–462, Berlin, Heidelberg, 2016. Springer-Verlag. (Cited on page 7.)

[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, Heidelberg, August 2017. (Cited on page 3.)

[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. (Cited on page 3.)

[LV17] Alex Lombardi and Vinod Vaikuntanathan. Minimizing the complexity of goldreich's pseudorandom generator. *Cryptology ePrint Archive*, 2017. (Cited on page 29.)

[Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012. (Cited on page 17.)

[MF15] Brice Minaud and Pierre-Alain Fouque. Cryptanalysis of the new multilinear map over the integers. Cryptology ePrint Archive, Report 2015/941, 2015. https://eprint.iacr.org/2015/941. (Cited on page 3.)

[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, Heidelberg, April 2012. (Cited on page 16, 18.)

[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, Heidelberg, August 2016. (Cited on page 3, 25.)

[PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 500–517. Springer, Heidelberg, August 2014. (Cited on page 3.)

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):34:1–34:40, 2009. (Cited on page 18.)

[RVV24] Seyoon Ragavan, Neekon Vafa, and Vinod Vaikuntanathan. Indistinguishability obfuscation from bilinear maps and lpn variants. In *Theory of Cryptography: 22nd International Conference, TCC 2024, Milan, Italy, December 2–6, 2024, Proceedings, Part IV*, page 3–36, Berlin, Heidelberg, 2024. Springer-Verlag. (Cited on page 3.)

[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, Heidelberg, December 2022. (Cited on page 15, 26, 27.)

[Wee22] Hoeteck Wee. Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 217–241. Springer, Heidelberg, May / June 2022. (Cited on page 25, 26.)

[Wee24] Hoeteck Wee. Circuit abe with poly(depth,λ)-sized ciphertexts and keys from lattices. In *Advances in Cryptology – CRYPTO 2024: 44th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2024, Proceedings, Part III*, page 178–209, Berlin, Heidelberg, 2024. Springer-Verlag. (Cited on page 4.)

[Wee25] Hoeteck Wee. Almost optimal kp and cp-abe for circuits from succinct lwe. In *Advances in Cryptology – EUROCRYPT 2025: 44th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Madrid, Spain, May 4–8, 2025, Proceedings, Part III*, page 34–62, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 4.)

[WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious lwe sampling. In *Advances in Cryptology – EUROCRYPT 2021: 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17–21, 2021, Proceedings, Part III*, page 127–156, Berlin, Heidelberg, 2021. Springer-Verlag. (Cited on page 3, 4, 7, 29.)

[WW25] Hoeteck Wee and David J. Wu. Unbounded distributed broadcast encryption and registered abe from succinct lwe. In *Advances in Cryptology – CRYPTO 2025: 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17–21, 2025, Proceedings, Part III*, page 204–235, Berlin, Heidelberg, 2025. Springer-Verlag. (Cited on page 4.)

[WWW22] Brent Waters, Hoeteck Wee, and David J. Wu. Multi-authority ABE from lattices without random oracles. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 651–679. Springer, Heidelberg, November 2022. (Cited on page 17.)