

Indistinguishability Obfuscation from Circular Security

Romain Gay*
IBM Zurich
romain.rgay@gmail.com

Rafael Pass†
Cornell Tech
rafael@cs.cornell.edu

December 8, 2020

Abstract

We show the existence of indistinguishability obfuscators ($i\mathcal{O}$) for general circuits assuming subexponential security of:

- (a) the Learning with Error (LWE) assumption (with subexponential modulus-to-noise ratio);
- (b) a *circular security conjecture* regarding the Gentry-Sahai-Water’s (GSW) encryption scheme and a Packed version of Regev’s encryption scheme.

The circular security conjecture states that a notion of leakage-resilient security, that we prove is satisfied by GSW assuming LWE, is retained in the presence of an encrypted key-cycle involving GSW and Packed Regev.

Our work thus places $i\mathcal{O}$ on qualitatively similar assumptions as unlevelled FHE, for which known constructions also rely on a circular security conjecture.

*Work done in part while at Cornell Tech

†Supported in part by NSF Award SATC-1704788, NSF Award RI-1703846, AFOSR Award FA9550-18-1-0267, DARPA SIEVE award HR00110C0086, and a JP Morgan Faculty Award. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via 2019-19-020700006. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

1 Introduction

The goal of *program obfuscation* is to “scramble” a computer program, hiding its implementation details (making it hard to “reverse-engineer”), while preserving its functionality (i.e its input/output behavior). In recent years, the notion of *indistinguishability obfuscation* ($i\mathcal{O}$) [BGI⁺01, GGH⁺13b] has emerged as the central notion of obfuscation in the cryptographic literature: roughly speaking, this notion requires that obfuscations $i\mathcal{O}(\Pi^1)$, $i\mathcal{O}(\Pi^2)$ of any two *functionally equivalent* circuits Π^1 and Π^2 (i.e. whose outputs agree on all inputs) from some class \mathcal{C} (of circuits of some bounded size) are computationally indistinguishable.

On the one hand, this notion of obfuscation is strong enough for a plethora of amazing applications (see e.g. [SW14, BCP14, BZ14, GGHR14, KNY14, KMN⁺14, BGL⁺15, CHJV14, KLW15, CLP15, BPR15, BPW16, BP15]). On the other hand, it may also plausibly exist, whereas stronger notion of obfuscations have run into strong impossibility results, even in idealized models (see e.g. [BGI⁺01, GK05, CKP15, Ps16, MMN15, LPST16]). Since the breakthrough of Garg, Gentry, Halevi, Raykova, Sahai and Waters [GGH⁺13b] that presented the first $i\mathcal{O}$ candidate, there has been an intensive effort toward obtaining a construction of $i\mathcal{O}$ based on some form of well-studied/nice assumptions. The original work [GGH⁺13b] provided a *candidate* construction based on high-degree multilinear maps (MLMs) [GGH13a, CLT13, GGH15, CLT15]; there was no proof of security based on an intractability assumption. [PST14] provided the first construction with a reduction-based proof of security, based on a strong notion of security for MLMs, similar to a sort of “Uber assumption”. [GLSW14] provided a construction based on a more concrete assumption relying on composite-order MLMs. Unfortunately, both assumptions have been broken for specific candidate constructions of MLMs [CHL⁺15, MF15].

$i\mathcal{O}$ from FE or $Xi\mathcal{O}$. Subsequently, several works have been constructing $i\mathcal{O}$ from seemingly weaker primitives, such as Functional Encryption (FE) [AJ15, BV15] or $Xi\mathcal{O}$ [LPST16], while only using standard assumptions, such as Learning with Error (LWE)¹. For both constructions, we actually need to rely on *subexponentially-secure* constructions of either FE or $Xi\mathcal{O}$, as well as subexponential security of LWE. Let us recall the notion of $Xi\mathcal{O}$ as it will be useful to us: roughly speaking, an $Xi\mathcal{O}$ is an $i\mathcal{O}$ with a very weak “exponential” efficiency requirement: the obfuscator is allowed to run in polynomial time in the size of the truth table of the function to be obfuscated, and it is only required that its outputs a program that “slightly” compresses the truth table (technically, it is sublinear in its size).

A breakthrough result by Lin [Lin16] showed how to obtain $i\mathcal{O}$ from *constant-degree* MLMs (plus standard assumptions), overcoming the black-box barriers in [Ps16, MMN15]. Her construction relies on the connection between FE and $i\mathcal{O}$. Following this result, a sequence of works (see e.g. [LV16, Lin17, LT17, AJKS18, JS18, JLMS19, AJL⁺19, GJLS20]) reduced the assumptions and the degree of the MLM — all the way down to 2-linear maps a.k.a. pairings— relying on certain types of low-degree pseudorandom generators (PRGs) to build FE. This culminated in the work of [GJLS20], whose security rely on the LWE assumption with binary errors in the presence of some PRG leakage, which despite being quite elegant, is new to their work, and as such, has not been significantly crypt-analyzed. Another line of work [Agr19, AP20] replace the use of 2-linear maps used by the aforementioned works by a noisy linear FE for inner products. While being plausibly post-quantum, these constructions are heuristic and do not provide a security reduction to a simple assumption.

A recent work by Brakerski et al [BDGM20a] presents a new type of candidate construction

¹Note that we are omitting some works that build $i\mathcal{O}$ without going through FE or $Xi\mathcal{O}$, such as [GJK18] that gives a direct heuristic construction of $i\mathcal{O}$ from tensor products, or [BIJ⁺20] that describes a candidate from Affine Determinant Programs. None of these provide a security proof.

of $Xi\mathcal{O}$ by combining a fully-homomorphic encryption (FHE) and a linear-homomorphic encryption (LHE) with certain nice properties (which can be instantiated by the Damgård-Jurik (DJ) [DJ01] encryption scheme whose security relies on the Decisional Composite Residuosity (DCR) assumption), and relying on a random oracle. More precisely, they define a new primitive called “split-FHE” and provide a candidate construction of it based on the above primitives and a random oracle, and next show how split-FHE implies $Xi\mathcal{O}$ (which by earlier work implies $i\mathcal{O}$ under standard assumptions). We highlight that [BDGM20a] does not provide any proof of security of the split-FHE construction (even in the random oracle model), but rather informally argue some intuitions, which include a) *circular security* (more on this below) of the FHE and the LHE, and b) a “*correlation conjecture*” that the FHE randomness (after FHE evaluations) does not correlate “too much” with the messages being encrypted. The correlation conjecture is not formalized, as the FHE randomness in known construction actually *does* depend on the message, so the authors simply conjecture that this correlation cannot be exploited by an attacker to break security of the $i\mathcal{O}$ (they also provide heuristic methods to weaken the correlations); as such they only get a heuristic construction.

Summarizing the above, while there have been enormous progress on realizing $i\mathcal{O}$, known constructions are either based on assumptions that are not well understood (high-degree MLMs, various low-degree PRGs assumptions and LWE with leakage type of assumptions), or the construction candidates simply do not have proofs of security.

1.1 Our Results

In this work, we provide a new $i\mathcal{O}$ construction assuming subexponential security of (a) the LWE assumption (with subexponential modulus-to-noise ratio), and (b) an (in our eyes) natural *circular security assumption* w.r.t the Gentry-Sahai-Water’s (GSW) [GSW13] FHE scheme and the DJ [DJ01] LHE scheme. Alternatively, assumption (b) can be replaced by a circular security conjecture regarding the GSW encryption scheme and a “packed” variant of Regev’s encryption scheme [Reg05, PVW08].

On a high-level, our approach follows that in [BDGM20a], but we show how to remove the heuristic arguments while instead relying on a concrete circular security assumption. We believe this constitutes strong evidence for the existence of $i\mathcal{O}$, and places $i\mathcal{O}$ on a qualitatively similar footing as *unlevelled* FHE (i.e. an FHE that support an a-priori unbounded polynomial number of operations), for which known constructions also rely on a circular security conjecture [Gen09]. We emphasize that the type of circular security conjecture that we rely on is stronger and more complex than the “plain” circular security conjecture used for unlevelled FHE. Yet on a philosophical level, we do not see any concrete evidence for why the plain circular security is more believable.²

Circular security. Circular security of encryption schemes [CL01, BRS02] considers a scenario where the attacker gets to see not only encryptions of messages, but also *encrypted key cycles*. The simplest form of circular security, referred to as *1-circular security*, requires that security holds even if the attacker gets to see not only the public key \mathbf{pk} and an encryption $\text{Enc}_{\mathbf{pk}}(\mathbf{m})$ of a message \mathbf{m} (to be secured), but also an encryption $\text{Enc}_{\mathbf{pk}}(\mathbf{sk})$ of the secret key \mathbf{sk} . A slightly more complex type of circular security, referred to as *2-circular security*, considers an encrypted key cycle of size 2 where the attacker gets to see public keys $\mathbf{pk}_1, \mathbf{pk}_2$, a length 2 encrypted secret key cycle, $\text{Enc}_{\mathbf{pk}_1}^1(\mathbf{sk}_2), \text{Enc}_{\mathbf{pk}_2}^2(\mathbf{sk}_1)$, and we require that security of $\text{Enc}_{\mathbf{pk}_1}^1$ still holds (i.e. for any m_0, m_1 , $\text{Enc}_{\mathbf{pk}_1}^1(m_0)$ is indistinguishable from $\text{Enc}_{\mathbf{pk}_1}^1(m_1)$). Encrypted key cycles commonly arise in applications of encryption scheme such as storage systems (e.g. BitLocker disk encryption utility), anonymous credentials [CL01] and most recently to construct (unlevelled) FHE [Gen09].

²See Section 1.4 for an extended comparison.

We refer to the assumption that:

If Enc^1 and Enc^2 are semantically secure, then 2-circular security holds w.r.t. $\text{Enc}^1, \text{Enc}^2$.

as the *2-circular security conjecture (2CIRC) w.r.t. $\text{Enc}^1, \text{Enc}^2$* . (We may also consider a subexponential version of this conjecture which is identically defined except that “security” is replaced by “subexponential security”.) For our purposes, we will allow the key generation procedure of Enc^1 to get the public-key pk_2 of Enc^2 as an input—for instance, this will allow Enc^1 and Enc^2 to operate over the same field.

At first sight, one may be tempted to hope that circular security holds w.r.t. *all* secure encryption schemes— $\text{Enc}^1, \text{Enc}^2$ —after all, the attacker never actually gets to see the secret key, but rather an encryption of it, which intuitively should hide it by semantic security of the encryption schemes. Yet, in recent years, counter examples to 2-circular security have been found. While for 1-circular security of *string encryption* schemes, it is easy to come up with a counter example—simply take any encryption scheme and modify it so that an encryption of \mathbf{m} outputs \mathbf{m} iff \mathbf{m} is a valid secret key, and otherwise proceeds just as before—coming up with counterexamples for 1-circular security of *bit encryption*, or 2-circular security for either string or bit encryption, is a lot harder (see e.g. [ABBC10, GH10, CGH12, Rot13, MO14, KRW15, BHW15, KW16, GKW17, WZ17]). In fact, all known counter examples are highly artificial, and require carefully embedding some trapdoor mechanism in the encryption scheme that enables decrypting the ciphertext once you see an encryption of the secret key. As far as we are aware, no “natural” counterexamples are known. Indeed, a common heuristic consists of simply assuming that 2-circular security holds for all “natural” encryption schemes that are secure; that is, 2CIRC holds for all “natural” encryption schemes—we refer to this as the 2CIRC heuristic. It is similar to the Random Oracle Heuristic [BR93]: while “contrived” counterexamples are known (see e.g., [CGH98, MRH04]), it is still commonly used for the design of practical protocols.

Leakage-resilient Circular Security. In this work, we rely on the assumption that stronger forms of security are preserved in the presence of a key cycle. More precisely, we consider a notion of \mathcal{O} -leakage resilient security where \mathcal{O} is some particular *randomness leakage oracle*; this notion enhances the standard semantic security notion by providing the attacker with access to an oracle $\mathcal{O}(\text{pk}, \mathbf{m}, \mathbf{r})$ that is parametrized by the public key pk , the message \mathbf{m} being encrypted and the randomness \mathbf{r} under which it is encrypted, while restricting the attacker to making only “valid” leakage queries (that do not trivially leak information about the message—this is formalized by letting the oracle output \perp whenever a query is invalid, and saying that the attacker fails whenever this happens). We next define the notion of *2-circular \mathcal{O} -leakage resilient security* analogously to 2-circular security, and also define a $2\text{CIRC}^{\mathcal{O}}$ conjecture (resp. a subexponential $2\text{CIRC}^{\mathcal{O}}$ conjecture) in the same way as the 2CIRC conjecture except that “security” is replaced by “ \mathcal{O} -leakage resilient security”; that is, we say that the $2\text{CIRC}^{\mathcal{O}}$ conjecture holds w.r.t. $\text{Enc}^1, \text{Enc}^2$ if the following holds:

If Enc^1 is \mathcal{O} -leakage resilient secure and Enc^2 is secure, then \mathcal{O} -leakage resilient security of Enc^1 is preserved in the presence of a length 2 key-cycle w.r.t. Enc^1 and Enc^2 .

Note that we cannot hope that $2\text{CIRC}^{\mathcal{O}}$ security holds for *all* oracles \mathcal{O} , even with respect to “natural” encryption schemes: simply consider an oracle $\mathcal{O}(\text{pk}, \mathbf{m}, \mathbf{r})$ that outputs the message \mathbf{m} iff \mathbf{m} is a valid secret key (just as in the counterexample to “plain” 1-circular security for string encryption). Thus, for $2\text{CIRC}^{\mathcal{O}}$ to be meaningful, we need to restrict not only to “natural” encryption schemes, but also to “natural” oracles \mathcal{O} .

Our first theorem shows that for a natural leakage oracle \mathcal{O}_{SRL} —which will be referred to as the “shielded randomness leakage (SRL) oracle”— $2\text{CIRC}^{\mathcal{O}_{\text{SRL}}}$ w.r.t. two standard encryption schemes (GSW and DJ) together with standard assumptions implies the existence of $i\mathcal{O}$.

Theorem 1.1 (Informally stated). *Assume the subexponential security of the LWE assumption (with subexponential modulus-to-noise ratio) and the DCR assumption, and the subexponential 2CIRC^{O_{SRL}} conjecture w.r.t. GSW and DJ. Then, $i\mathcal{O}$ exists for the class of polynomial-size circuits.*

Alternatively, we can replace the DJ encryption scheme with a “packed” variant of Regev’s encryption scheme [Reg05], which we refer to as Packed Regev. (We note that our Packed Regev is very similar to, but actually different from, the Packed Regev in [PVW08].) This construction only relies on LWE and the 2-circular security conjecture.

Theorem 1.2 (Informally stated). *Assume the subexponential security of the LWE (with subexponential modulus-to-noise ratio) assumption, and assume that the subexponential 2CIRC^{O_{SRL}} conjecture holds w.r.t. GSW and Packed Regev. Then, $i\mathcal{O}$ exists for the class of polynomial-size circuits.*

In the sequel, we refer to \mathcal{O}_{SRL} -leakage resilient security (resp. 2-circular \mathcal{O}_{SRL} -leakage resilient security) as SRL-security (resp 2-circular SRL security). We proceed to explain the notion of SRL security and how the above theorems are proven.

1.2 Shielded Randomness Leakage (SRL) Security

As mentioned above, we consider a notion of *shielded randomness leakage (SRL)* security for FHE. Roughly speaking, given two messages $\mathbf{m}^0, \mathbf{m}^1$, the attacker gets to see an FHE encryption $\mathbf{c} = \text{FHE}(\mathbf{m}^b; \mathbf{r})$ of \mathbf{m}^b for a randomly selected $b \in \{0, 1\}$, and next gets access to a “leakage oracle” $\mathcal{O}_{\text{SRL}}(\mathbf{m}^b, \mathbf{r})$ which upon every invocation sends the attacker an “extra noisy” encryption $\mathbf{c}^* = \text{FHE}(0; \mathbf{r}^*)$ of 0—we will refer to the random string \mathbf{r}^* as the “shield”. Next, the attacker can select some functions f and values α such that $f(\mathbf{m}^b) = \alpha$ —that is, we restrict the attacker to picking functions for which it *knows the output* when applying the function to the message \mathbf{m}^b ; if $f(\mathbf{m}^b) \neq \alpha$, the attacker directly fails in the game. (The reason why we add this restriction on the attacker will soon become clear). Finally, the oracle homomorphically evaluates f on the ciphertext \mathbf{c} , letting $\mathbf{c}_f = \text{FHE}(f(\mathbf{m}); \mathbf{r}_f)$ denote the evaluated ciphertext, and returns $\mathbf{r}^* - \mathbf{r}_f$. That is, the attacker gets back the randomness \mathbf{r}_f of the evaluated ciphertext masked by the “shield” \mathbf{r}^* , and as usual, the attacker’s goal is to guess the bit b . The reason why the attacker is restricted to picking functions f for which it knows the output α is that for the FHE we consider, given \mathbf{c}^* and \mathbf{c}_f , the attacker can compute $\mathbf{c}^* - \mathbf{c}_f = \text{FHE}(0 - f(\mathbf{m}^b); \mathbf{r}^* - \mathbf{r}_f)$ and thus knowing $\mathbf{r}^* - \mathbf{r}_f$ reveals $f(\mathbf{m}^b)$. So, by restricting to attackers that already know $\alpha = f(\mathbf{m}^b)$, intuitively, $\mathbf{r}^* - \mathbf{r}_f$ does not reveal anything else. Indeed, we formally prove that under the LWE assumption, the GSW encryption scheme is SRL-secure (i.e. \mathcal{O}_{SRL} -leakage resilient secure).

Theorem 1.3 (Informally stated). *Assume the LWE assumption holds (with subexponential modulus-to-noise ratio). Then, the GSW scheme is SRL-secure.*

On a very high-level, the idea behind the proof is that the encryption \mathbf{c}^* is a projection, $h_{\mathbf{A}}(\mathbf{r}^*) = \mathbf{A}\mathbf{r}^* \in \mathbb{Z}_N$, where the randomness \mathbf{r}^* used to produce \mathbf{c}^* is a vector in \mathbb{Z}_N^ℓ and \mathbf{A} is a matrix in $\mathbb{Z}_N^{n \times \ell}$ where $\ell \gg n$, that is, the map $h_{\mathbf{A}}$ that describes the encryption is compressing. Therefore, some “components” of the “shield” \mathbf{r}^* remain information-theoretically hidden. And this enables hiding the same components of \mathbf{r}_f ; furthermore, the components that are not hidden by \mathbf{r}^* are actually already revealed by $f(\mathbf{m}^b)$, which the attacker knows (as we require it to output $\alpha = f(\mathbf{m}^b)$). The formal proof of this proceeds by considering a (simplified) variant of the Micciancio-Peikert lattice trapdoor method [MP12] for generating the matrix \mathbf{A} (which is part of the public key for GSW) together with a trapdoor that enables sampling short preimages of $h_{\mathbf{A}}$ (i.e. solving the ISIS problem). Whereas traditional trapdoor preimage sampling methods require the preimage to be sampled according to some specific distribution (typically discrete Gaussian) over preimages, we will

consider a somewhat different notion: we require that given a target vector \mathbf{t} , the distribution of randomly sampled preimages of \mathbf{t} is statistically close to the distribution obtained by starting with any “short” preimage \mathbf{w} of \mathbf{t} and next adding a randomly sampled preimage of 0. Our proof relies on the fact that randomly sampled preimages can be sufficiently larger than \mathbf{w} to ensure that they “smudge” \mathbf{w} —we here rely on the fact that modulus-to-noise ratio is subexponential (which we need anyway for the security of our construction) to enable the smudging³.

2-Circular SRL Security. As mentioned, we define 2-circular SRL security as 2-circular \mathcal{O}_{SRL} -leakage resilient security; we emphasize that this security game is identically defined to the “plain” SRL security game (described above), with the only exception being that the challenge message encrypted (using $\text{Enc}_{\text{pk}_1}^1$) has the form $\text{sk}_2 || \mathbf{m}^b$ (as opposed to just being \mathbf{m}^b), and that the attacker also gets to see an encryption of sk_1 (using $\text{Enc}_{\text{pk}_2}^2$).

1.3 Overview of the $Xi\mathcal{O}$ Construction

We present a construction that makes a modular use of any LHE satisfying certain properties, and whose security relies the 2-circular SRL-security w.r.t. GSW and the LHE (i.e., that SRL security of GSW holds in the presence of a encrypted key cycle of length 2 using GSW and the LHE). To obtain a *subexponentially-secure* $Xi\mathcal{O}$ (which is required to obtain $i\mathcal{O}$ by [LPST16]), we need to strengthen the assumptions to also require subexponential security. Next, we note that the DJ LHE satisfies the desired properties. We prove that a packed version of Regev’s encryption scheme [Reg05] that is similar to, but actually different from, the packed construction from [PVW08], does so as well. We refer to our LHE simply as Packed Regev LHE.

Let us start with the construction assuming 2-circular SRL-security w.r.t. GSW and any LHE satisfying the desired properties. As mentioned, on a high-level, our construction follows similar intuitions as the BDGM construction. We combine an FHE (in our case the GSW FHE) with a (special-purpose) LHE to implement an $Xi\mathcal{O}$. In fact, in our approach, we do not directly construct an $Xi\mathcal{O}$, but rather construct an $Xi\mathcal{O}$ with *preprocessing*—this notion, which relaxes $Xi\mathcal{O}$ by allowing the obfuscator to have access to some *long* public parameter pp , was actually already considered in [LPST16] and it was noted there that subexponentially-secure $Xi\mathcal{O}$ with preprocessing also suffices to get $i\mathcal{O}$.

Towards explaining our approach, let us first recall the approach of BDGM—which relies on the DJ LHE—using a somewhat different language that will be useful for us.

The BDGM construction. The high-level idea is quite simple and very elegant. Recall that an $Xi\mathcal{O}$ is only required to work for programs Π with polynomially many inputs $n = \text{poly}(\lambda)$ where λ is the security parameter, and the obfuscators running time is allowed to be polynomial in n ; the only restriction is that the obfuscated code should be sublinear in n —we require a “slight” compression of the truth table. More precisely, the obfuscator is allowed to run in time $\text{poly}(n, \lambda)$ (i.e. polynomial time in the size of the truth table), but must output a circuit of size $\text{poly}(\lambda)n^{1-\varepsilon}$ where $\varepsilon > 0$. Assume that we have access to a special “batched” FHE which enables encrypting (and computing on) long messages of length, say m using a *short randomness* of length $\text{poly}(\lambda) \log(m)$; and furthermore that 1) given the secret key and a ciphertext \mathbf{c} , we can efficiently recover the ciphertext randomness 2) given a ciphertext \mathbf{c} and its randomness—which will also be referred to as a “hint”—one can efficiently decrypt. Given such a special FHE, it is easy to construct an $Xi\mathcal{O}$: simply cut the truth table into “chunks” of length n^ε , FHE encrypt the program Π , then, homomorphically evaluate circuits C_i for

³Another consequence of using smudging is that our lattice trapdoor mechanism and its proof become simpler than [MP12], which uses a polynomial-size modulus instead, for a better efficiency.

indices $i \in [n^{1-\varepsilon}]$ such that given the program Π as input, C_i outputs the i 'th “chunk” of the truth table, which we denote by Π_i ; finally, release the randomness \mathbf{r}_i (i.e. the “hint”) of the evaluated ciphertexts. These hints enable compressing n^ε bits into $\text{poly}(\lambda) \log(n^\varepsilon)$ bits and thus the XiO is compressing.⁴

Unfortunately, none of the known FHE constructions have short randomness. BDGM, however, observes that there are *linear* homomorphic encryptions schemes (LHE), notably the DJ LHE, that satisfy the above requirements. Moreover, many FHEs are batchable (with “long” randomness) and have “essentially” linear decryption: decryption is an inner product of the ciphertext with the secret key, then rounding. That is, the linear operations yield the plaintext with some additional small decryption noises, that are removing when rounding. So if we start off with such an FHE and additionally release an LHE encryption of the FHE secret key, we can get an FHE with the desired “batchable with short randomness” requirement: we first homomorphically evaluate the inner product of the FHE ciphertext with the encrypted FHE secret key, then simply release the randomness for the evaluated LHE ciphertext (which now is short).

But there are problems with this approach: (1) since FHE decryption requires performing both a linear operation and *rounding*, we are leaking not only Π_i but also the decryption noises, which is detrimental for the security of the FHE (2) the LHE randomness may actually leak more than just the decrypted LHE plaintext (i.e. something about how the LHE ciphertext was obtained). As BDGM shows, both of these problems can be easily overcome if we have access to many fresh LHE encryptions of some “smudging” noise (which is large enough to smudge the FHE decryption noises)⁵. Therefore, the only remaining problem is to generate these LHE encryptions of smudging noises. This is where the construction in BDGM becomes heuristic: (1) they propose to use a random oracle to generate a long sequence of randomness (2) this sequence of randomness can be interpreted as a sequence of LHE encryptions of uniformly random strings u_i for $i = 1, \dots, n^{1-\varepsilon}$, since the DJ LHE has dense ciphertext (3) they additionally provide an FHE encryption of the LHE secret key $\overline{\text{sk}}$ (note that there is now a circular security issue), on which they FHE-homomorphically evaluate a function f_i that decrypts the i 'th LHE ciphertext produced by the random oracle, and computes $\text{MSB}(u_i)$, the most significant bits of u_i (4) finally they LHE-evaluate the (partial) decryption of the evaluated FHE ciphertext (which encrypts $\text{MSB}(u_i)$); the obtained LHE ciphertext can now be subtracted from the LHE ciphertexts generated by the random oracle, to get an LHE encryption of $u_i - \text{MSB}(u_i)$, which is a noise of the appropriate size, i.e. smudging but not uniform.

One problem with this approach, however, is that while we do obtain an LHE encryption of appropriate smudging noise, it is not actually a fresh ciphertext (with fresh randomness). The issue is that the randomness \mathbf{r}_{f_i} of the evaluated FHE ciphertext of $\text{MSB}(u_i)$ may (and actually will) depend on the randomness of the original LHE ciphertext obtained by the RO. Another problem is that LHE can only compute the first step of an FHE decryption (namely, the linear operations), the LHE encryption obtained actually encrypts a message of the form: $u_i - \text{MSB}(u_i) + \text{noise}_i$. As we know, revealing the extra noise is detrimental for security (this is why we are generating LHE encryptions of smudging noises in the first place). Unfortunately, the extra noise that results from partially decrypting the FHE ciphertext depends on u_i , so the lower-order bits of the latter cannot smudge the former. BDGM here simply assumes that the attacker cannot exploit these correlations, and thus only obtain a heuristic construction.

We shall now see how to obtain the appropriate LHE encryption of smudging noises in a prov-

⁴The reason we need to cut the truth table into chunks instead of directly computing the whole output is that the size of the FHE public key and ciphertexts may grow polynomially with the length of the output of the homomorphic evaluation, i.e. the “batching capacity”. So the obfuscation is only compressing when we have a large number of chunks.

⁵They formally prove the security of their scheme in an idealized model with access to an oracle that generates fresh LHE encryptions of smudging noise.

ably secure way, relying on $\mathcal{2}$ -circular *SRL-security* of GSW and DJ—that is, \mathcal{O}_{SRL} -leakage resilient circular security of GSW and DJ.

Removing the RO. Our first task will be to remove the use of the RO. That will actually be very easy: as we have already observed, it suffices to get an \mathcal{XiO} with preprocessing to obtain $i\mathcal{O}$, so instead of using a random oracle, we will simply use a long random string as a public parameter, and interpret it as LHE encryptions of random strings.

Re-encrypting the FHE. The trickier problem will be to deal with the issue of correlations. We will here rely on the fact that we are considering a particular instantiation of the FHE: namely, using (a batched version of) the GSW encryption scheme. On a high-level, the idea for breaking the correlation is to “refresh” or re-encrypt the evaluated FHE ciphertext (which encrypts $\text{MSB}(u_i)$) to ensure that the randomness is fresh and independent of the evaluations. This way, the decryption noise itself is independent of the evaluated circuit. GSW ciphertexts can be re-randomized simply by adding a fresh extra noisy FHE encryption of 0. How do we get such encryptions? GSW ciphertexts are not dense, so we cannot put them in the public parameters, and even if they were, we still wouldn’t be able to get an encryption of 0 (we would have an encryption of a uniformly random plaintext). The public key of the GSW encryption scheme actually contains a bunch of encryptions of 0, but fewer than the amount we need (or else we wouldn’t get a compressing \mathcal{XiO}). Instead, we use the public key of the GSW encryption to generate extra noisy encryptions of 0, and we include the (many) random coins $(\mathbf{r}_i^*)_{i \in [n^{1-\epsilon}]}$ used to generate these ciphertexts as part of the public parameters of the \mathcal{XiO} (recall that the public parameters can be as long as we want). This method does indeed enable us to get a fresh FHE encryption of the most significant bits, and thus the correlation has been broken and intuitively, we should be able to get a provably secure construction. But two obstacles remain: (1) we are revealing the randomness used to re-randomize the ciphertexts, and this could hurt security, or render the re-randomization useless and (2) we still have a circular security issue (as we FHE-encrypt the LHE secret key, and LHE-encrypt the FHE secret key). Roughly speaking, the first issue will be solved by relying on SRL-security of GSW, and the second issue will be solved by our circular security conjecture.

In more detail, we note that the re-randomized evaluated FHE ciphertext of $\text{MSB}(u_i)$ and the public parameters \mathbf{r}_i^* are statistically close to freshly generated extra noisy FHE encryption of $\text{MSB}(u_i)$ using randomness \mathbf{r}_i^* , and setting the public parameter to $\mathbf{r}_i^* - \mathbf{r}_{f_i}$, where \mathbf{r}_{f_i} is the randomness of the evaluated ciphertext, before re-randomization. In other words, the re-randomization achieves a notion which we refer to as “weak circuit privacy”, where the re-randomized ciphertext is independent of the evaluated function f_i . Furthermore, noisy GSW encryptions of $\text{MSB}(u_i)$ essentially have the form of a noisy GSW encryption of 0, to which $\text{MSB}(u_i)$ is added. So, other than $\text{MSB}(u_i)$, which is truly random, $\mathbf{r}_i^* - \mathbf{r}_{f_i}$ is simply an SRL leakage on a GSW encryption of the LHE secret key $\overline{\text{sk}}$. Thus, intuitively, security should now follow from circular SRL security of GSW and the LHE.

The final construction. We summarize our final \mathcal{XiO} construction with preprocessing. The public parameter pp is a long *random* string that consists of two parts:

- The first part FHE.PubCoin will be interpreted as a sequence of rerandomization vectors \mathbf{r}^* ;
- The second part LHE.PubCoin will be interpreted as a sequence of LHE encryptions

The obfuscator, given a security parameter λ and a circuit $\Pi : \{0, 1\}^{\log n} \rightarrow \{0, 1\}$, where $n = \text{poly}(\lambda)$ proceeds as follows:

- **Output the public keys of the FHE and LHE:** The obfuscator generates a fresh key-pair $(\overline{\text{pk}}, \overline{\text{sk}})$ for the LHE, and next generate a key-pair (pk, sk) for the GSW FHE. (To make it easier for the reader to remember which key refers to which encryption scheme, we place a *line* over all keys, ciphertexts and algorithms, that correspond to the *linear* homomorphic encryption.) The modulus N of the GSW encryption is set to be the same that the modulus that defines the message space \mathbb{Z}_N of the LHE scheme. Additionally, it chooses N large enough to enable encrypting messages of size n^ε . Finally, it outputs the public keys $(\text{pk}, \overline{\text{pk}})$.
- **Output an FHE encryption of the circuit:** It outputs an FHE encryption (w.r.t. pk) of the program Π , which we denote by ct_1 .
- **Output encrypted key cycle:** It computes ct_2 , an FHE encryption of $\overline{\text{sk}}$, and $\overline{\text{ct}}$, an LHE encryption of sk . It outputs the key cycle $\text{ct}_2, \overline{\text{ct}}$.
- **Output hints:** For every $i \in [n^{1-\varepsilon}]$, it outputs a short “hint” \mathbf{r}_i computed as follows:
 - **Evaluate the circuit:** Homomorphically evaluate the circuit C_i on ct_1 and let ct_i denote the resulting evaluated FHE ciphertext — recall that ct_1 encrypts a program Π , and the circuit C_i takes a input a program Π and outputs the i 'th chunk of its truth table.
 - **Compute an FHE encryption $\text{ct}_{\text{MSB},i}$ of $\text{MSB}(u_i)$:** Consider the function $f_i(\Pi, \overline{\text{sk}})$ that ignores the input Π but uses the input $\overline{\text{sk}}$ to decrypt the i 'th LHE ciphertext from LHE.PubCoin into a plaintext u_i and outputs $\text{MSB}(u_i)$. The obfuscator homomorphically evaluates f_i on the ciphertexts ct_1, ct_2 (where, recall, ct_2 is an encryption of $\overline{\text{sk}}$). Let $\text{ct}_{\text{MSB},i} = \text{FHE}(\text{MSB}(u_i); \mathbf{r}_{f_i})$ denote the resulting evaluated FHE ciphertext.
 - **Rerandomize $\text{ct}_{\text{MSB},i}$ into $\text{ct}'_{\text{MSB},i}$:** It uses the i 'th chunk of FHE.PubCoin to get the randomness \mathbf{r}_i^* ; generates an extra noisy FHE encryption of 0 using \mathbf{r}_i^* and homomorphically adds it to $\text{ct}_{\text{MSB},i}$. Let $\text{ct}'_{\text{MSB},i} = \text{FHE}(\text{MSB}(u_i); \mathbf{r}_i^* + \mathbf{r}_{f_i})$ denote the new (re-randomized) ciphertext.
 - **Proxy re-encrypt ct_i as an LHE ciphertext $\overline{\text{ct}}_i$:** It uses $\overline{\text{ct}}$ (which, recall, is an LHE encryption of sk) to homomorphically compute the *linear* part of the FHE decryption of ct_i , which yields an LHE encryption of the value $2^\omega \cdot \Pi_i + \text{noise}_i$ where noise_i is an FHE decryption noise, and 2^ω is taken large enough so that the plaintext Π_i can be recovered by rounding.

Similarly, it homomorphically computes the partial FHE decryption of $\text{ct}'_{\text{MSB},i}$, which yields an LHE encryption of the value $2^{\omega'} \cdot \text{MSB}(u_i) + \text{noise}_{\text{MSB},i}$, where once again $\text{noise}_{\text{MSB},i}$ denotes an FHE decryption noise, and $2^{\omega'} = 1$ for reasons that will become clear later. We rely on the fact that GSW FHE (and many other FHE schemes) admits a flexible “scaled” evaluation algorithm, that can choose which integer 2^ω to use when performing the homomorphic evaluation (this was used also in prior works, including [BDGM20a]). The resulting LHE ciphertext is subtracted from $\text{LHE}(2^\omega \cdot \Pi_i + \text{noise}_i)$, and therefore yields $\text{LHE}(2^\omega \cdot \Pi_i + \text{noise}_i - \text{MSB}(u_i) - \text{noise}_{\text{MSB},i})$.

Finally, it homomorphically adds the LHE encryption of u_i that is part of the LHE public coins, to obtain $\overline{\text{ct}}_i = \text{LHE}(m_i)$, where $m_i = 2^\omega \cdot \Pi_i + \text{noise}_i - \text{MSB}(u_i) - \text{noise}_{\text{MSB},i} + u_i = 2^\omega \cdot \Pi_i + \text{noise}_i + \text{noise}_{\text{MSB},i} + \text{LSB}(u_i)$, where $\text{LSB}(u_i)$ denotes the least significant bits of u_i .

The integer ω' is chosen to be equal to 0 so that the smudging noise $\text{LSB}(u_i)$ is directly added to the FHE noises $\text{noise}_i - \text{noise}_{\text{MSB},i}$. As opposed to the value Π_i that we place in the higher-order bits of the plaintext, we need the smudging noise to be at the same level than the FHE noises, so they “blend” together.

- **Release hint \mathbf{r}_i for LHE ciphertext $\overline{\text{ct}}_i$:** It uses $\overline{\text{sk}}$ to recover the randomness \mathbf{r}_i of $\overline{\text{ct}}_i$ (recall that the LHE we use has a randomness recoverability property), and outputs \mathbf{r}_i .

To evaluate the obfuscated program on an input $\mathbf{x} \in \{0, 1\}^n$, that pertains to the i 'th chunk of the truth table of Π for some $i \in [n^{1-\varepsilon}]$, we compute $\overline{\text{ct}}_i$ just like the obfuscator did (note that this does not require knowing the secret key, but only information contained in the obfuscated code). Finally, we decrypt $\overline{\text{ct}}_i$ using the hint \mathbf{r}_i to recover the message m_i described above (recall that the LHE we use has the property that ciphertexts can be decrypted if you know the randomness). Finally, perform the rounding step of FHE decryption on m_i to obtain Π_i , which contains $\Pi(\mathbf{x})$.

Outline of the security proof. We provide a very brief outline of the security proof. We will rely on the fact that LHE ciphertexts (of random messages) are dense (in the set of bit strings), and additionally on the fact that both the LHE and the FHE we rely on (i.e. DJ and GSW) satisfy what we refer to as a *weak circuit privacy* notion. This notion, roughly speaking, says that *any* encryption of a message x can be rerandomized into fresh (perhaps extra noisy) encryption of $x + y$, by adding a fresh (perhaps extra noisy) encryption of y .

As usual, the proof proceeds via a hybrid argument. We start from an XiO obfuscation of a program Π^0 and transition until we get an XiO obfuscation of Π^1 , where Π^0 and Π^1 are two functionally equivalent circuits of the same size.

- **Hybrid 0: Honest $\text{XiO}(\Pi^0)$.** The first hybrid is just the honest obfuscation of the circuit Π^0 .
- **Hybrid 1: Switch to freshly encrypted $\text{ct}'_{\text{MSB},i}$.** Hybrid 1 proceeds exactly as Hybrid 0 up until the point that the ciphertexts $\text{ct}_{\text{MSB},i}$ get re-encrypted into $\text{ct}'_{\text{MSB},i}$, with the exception that FHE.PubCoin are not sampled yet. Next, instead of performing the re-encryption, we sample $\text{ct}'_{\text{MSB},i}$ as a *fresh* extra noisy encryption of $\text{MSB}(u_i)$ using randomness \mathbf{r}_i^* , and setting FHE.PubCoin to be $\mathbf{r}_i^* - \mathbf{r}_{f_i}$ (recall that \mathbf{r}_{f_i} is the randomness obtained when homomorphically evaluating f_i on the FHE encryption of sk). We finally continue the experiment in exactly the same way as in Hybrid 0.

It follows from the “weak circuit privacy” property of the FHE that Hybrid 0 and Hybrid 1 are statistically close. Note that in Hybrid 1, for each $i \in [n^{1-\varepsilon}]$, the i 'th chunk of FHE.PubCoin can be thought of as SRL leakage on the fresh encryption $\text{ct}'_{\text{MSB},i}$ computed w.r.t. function f_i , which will be useful for us later.

- **Hybrid 2: Switch LHE.PubCoin to encryptions of random strings.** Hybrid 2 proceeds exactly as Hybrid 1 except that instead of sampling LHE.PubCoin as a random string, we sample it as fresh LHE encryptions of random strings u_i , for $i = 1, \dots, n^{1-\varepsilon}$. It follows by the density property of the LHE that Hybrid 2 is statistically close to Hybrid 1.
- **Hybrid 3: Generate $\overline{\text{ct}}_i$ as a fresh encryption.** Hybrid 3 proceeds exactly as Hybrid 2 except that $\overline{\text{ct}}_i$ is generated as a fresh encryption of m_i using fresh randomness \mathbf{r}_i , and the i 'th chunk of LHE.PubCoin is instead computed homomorphically by subtracting the LHE encryption of $\text{sk}^\top(\text{ct}_i - \text{ct}_{\text{MSB},i})$ (obtained after homomorphically decrypting ct_i and $\text{ct}'_{\text{MSB},i}$ using $\overline{\text{ct}}$) from the LHE ciphertext $\overline{\text{ct}}_i$. Recall that $m_i = \text{sk}^\top(\text{ct}_i - \text{ct}_{\text{MSB},i}) + u_i$ so the above way of computing the i 'th chunk of LHE.PubCoin ensures that it is valid encryption of u_i as in Hybrid 2, but this time with non-fresh, homomorphically evaluated randomness.

It follows from the weak circuit privacy property of the LHE that Hybrid 3 and 2 are statistically close.

Note that it was possible to define this hybrid since $\text{ct}'_{\text{MSB},i}$ remains exactly the same no matter what the LHE.PubCoin are. This was not true in Hybrid 0, and we introduced Hybrid 1 to break this dependency.

Note further that in Hybrid 3, we no longer use $\overline{\text{sk}}$ (i.e. the secret key for LHE); previously it was used to recover \mathbf{r}_i .

- **Hybrid 4: Generate $\overline{\text{ct}}_i$ without FHE noises.** Hybrid 4 proceeds exactly as Hybrid 3 except that $\overline{\text{ct}}_i$ is generated as a fresh encryption of $m_i = 2^\omega \cdot \Pi_i^0 + \text{LSB}(u_i)$, whereas in Hybrid 3, it was generated as fresh encryption of $m_i = 2^\omega \cdot \Pi_i^0 + \text{LSB}(u_i) + \text{noise}_i - \text{noise}_{\text{MSB},i}$. That is, we use $\text{LSB}(u_i)$ as a smudging noise to hide the extra noise $\text{noise}_i - \text{noise}_{\text{MSB},i}$. We can do so since (1) the extra FHE noise is small and independent of $\text{LSB}(u_i)$ (2) the rest of the obfuscated code can be generated from the value $\text{LSB}(u_i) + \text{noise}_i - \text{noise}_{\text{MSB},i}$ only (in particular it does not require to know $\text{LSB}(u_i)$ itself). It follows that Hybrid 4 is statistically close to Hybrid 3.
- **Hybrid 5: Switch to encryption of Π^1 :** Hybrid 5 proceeds exactly as Hybrid 4 except that ct_1 is an encryption of Π^1 (instead of Π^0 in prior hybrids).

Note that other than the encrypted key cycle, we never use the FHE secret key, and due to Hybrid 3, we no longer use the LHE secret key. So, at first sight, Hybrid 5 ought to be indistinguishable from Hybrid 4 by circular security of the FHE and the LHE. Recall that FHE.PubCoin leaks something about the randomness used by the FHE encryption $\text{ct}'_{\text{MSB},i}$, but the leakage is exactly an SRL leakage (and note that in the experiment we do know the output α_i of the function f_i that is applied to the plaintexts encrypted in ct_1, ct_2 —namely, it is $\text{MSB}(u_i)$ where u_i is a random string selected in the experiment, see Hybrid 2). Thus, indistinguishability of Hybrid 5 and Hybrid 4 follows from 2-circular SRL-security of the FHE and the LHE.

- **Hybrids 6-10:** For $i \in [5]$, Hybrid 5 + i is defined exactly as 5 - i , except that ct_1 be an encryption of Π^1 . Statistical closeness of intermediary hybrids follows just as before.

The above sequence of hybrid allows us to conclude the following theorem.

Theorem 1.4 (Informally stated). *Assume the 2-circular SRL-security of the GSW and DJ encryption schemes. Then, there exists an $Xi\mathcal{O}$ for polynomial-size circuits taking inputs of length $\log(\lambda)$ where λ is the security parameter.*

An alternative LHE based on Packed Regev. We remark that we can obtain an alternative construction of an LHE with the desired properties by considering an *packed* version of the Regev encryption scheme. Our construction is slightly different, but similar in spirit, to the Packed Regev from [PVW08]. Recall that a (plain) Regev public key consist of a pair $\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$, where $\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{m \times n}$ with $m \geq n \log(q)$, the vector $\mathbf{s} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^n$ is the secret key, and $\mathbf{e} \in \mathbb{Z}_q^m$ is some small “noise” vector. An encryption of a message μ has the form $\mathbf{Ar}, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r} + B \cdot \mu$ where $\mathbf{r} \leftarrow_{\mathbb{R}} \{0, 1\}^m$ is the encryption randomness and B is an upper bound on the size of noise (so as to enable decryption). This scheme is linearly homomorphic, but for security, the size of the randomness $|\mathbf{r}|$ needs to be greater than $n \log(q)$, which is more than that size of the message: the randomness is too long for our purposes.

To get succinct decryption hints, we simply reuse the same randomness \mathbf{r} for many encryptions using different secret keys $\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_\ell$ and different noises $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_\ell$. The secret key is now a matrix $\mathbf{S} \in \mathbb{Z}_q^{\ell \times n}$, and the public key becomes $(\mathbf{A}, \mathbf{SA} + \mathbf{E})$ where $\mathbf{E} \in \mathbb{Z}_q^{\ell \times m}$ is a noise matrix. The encryption of a vector of messages $\mu = (\mu_1, \dots, \mu_\ell)$ is then $(\mathbf{Ar}, (\mathbf{SA} + \mathbf{E}) \mathbf{r} + B\mu)$. This is the

scheme from [PVW08]. Despite the fact that this encryption is still linearly homomorphic, and has the advantage of having rate-1 ciphertext size, its randomness is not short: to carry on the proof of security, we need to rely on the fact that \mathbf{r} contains enough bits of entropy even when the information \mathbf{Ar} (which is short) and \mathbf{Er} (that is long) is leaked. This can only be true when the dimension of \mathbf{r} , m , grows with the number of bits that are batched, ℓ .

Thus, we depart from the scheme in [PVW08] by adding a smudging noise⁶ in the ciphertext, to hide the information \mathbf{Er} . The ciphertext is of the form: $(\mathbf{Ar}, (\mathbf{SA} + \mathbf{E})\mathbf{r} + \mathbf{e}' + B \cdot \mu)$, where \mathbf{e}' is the extra smudging noise that hides the error term \mathbf{Er} , ensuring that we only have the short \mathbf{Ar} leakage and the usual proof can again be applied.

This scheme is still linearly homomorphic, but the encryption randomness is still large, as even though we reuse \mathbf{r} , the added noise terms \mathbf{e}' are large. However, we rely on the fact that knowing \mathbf{e}' is not needed for decrypting. Indeed, to decrypt, we just need to know a small vector $\tilde{\mathbf{r}} \in \mathbb{Z}_q^m$ such that $\mathbf{A}\tilde{\mathbf{r}} = \mathbf{Ar}$. That can be used to remove the term \mathbf{SAr} from the ciphertext, and recover $B \cdot \nu$ plus some small noise. To sample such vector, we use a standard trapdoor sampling mechanism as in prior works [Ajt96, GPV08, AP09, MP12]. This makes the scheme hintable with succinct hints.

We still have two (minor) obstacles, though. This scheme (as well as Regev’s original scheme or the scheme from [PVW08]) does not satisfy two of the other properties needed for our \mathcal{XiO} construction: (1) density, and (2) weak circuit privacy. But it almost does. *Extra noisy* ciphertexts, where the noise reaches the bound B are actually dense, and for extra noisy ciphertext, weak circuit privacy also holds (just as it did for GSW). So, we can directly instantiate the LHE in our \mathcal{XiO} construction with this Packed Regev construction, as long as we slightly relax the notion of an LHE to just require density when considering extra noisy ciphertexts.

Thus we can conclude:

Theorem 1.5 (Informally stated). *Assume 2-circular SRL-security of the GSW and the Packed Regev encryption schemes holds. Then, there exists an \mathcal{XiO} for polynomial-size circuits taking inputs of length $\log(\lambda)$ where λ is the security parameter.*

The proof of Theorems 1.1, 1.2 is finally concluded by upgrading Theorems 1.3, 1.4 and 1.5 to apply also in the subexponential regime, relying on the subexponential $2\text{CIRC}^{\text{O}_{\text{SRL}}}$ conjecture, and finally relying on the transformation from subexponentially-secure \mathcal{XiO} with pre-processing (and subexponential LWE) to $i\mathcal{O}$ [LPST16].

1.4 Comparing Circular SRL-security to “Plain” Circular Security

Let us make a few remarks on the 2-circular SRL-security assumption w.r.t GSW and some LHE (e.g. Damgaard Jurik or Packed-Regev). Clearly, this assumption is stronger than the 2-circular assumption w.r.t GSW and the LHE—simply consider an attacker that does not request any leakage. Additionally, we wish to highlight a few qualitative differences between “plain” circular security and circular SRL-security:

- “Plain” circular security is a simple non-interactive falsifiable assumption. Circular SRL-security is also a (relatively simple) falsifiable assumption, but the security game is now *interactive*; for the type of SRL security needed for our application, a single “parallel” SRL query suffices and such a notion of SRL security can be specified as a 5-round security game.

As we explain in more detail in Section 3.1.3, for our application, one could define a *non-interactive* falsifiable variant of SRL security—roughly speaking, where the messages and the

⁶Note that using a carefully crafted noise that needs not be of smudging size, as done in [MP12], we can “unskew” the noise \mathbf{Er} and hide the information of \mathbf{r} . We favor clarity of the exposition over efficiency and resort to using smudging noises.

leakage-selection algorithm are randomly selected—such that the subexponential hardness of this circular “random-SRL” security notion suffices⁷, but in our eyes, this non-interactive security game is less natural than the interactive one (and thus does not add much insight).

- It is also worth noting that for the notion of “plain” circular security, an alternative way of defining circular security is to require that $\text{Enc}_{\text{pk}_1}^1(\text{sk}_2), \text{Enc}_{\text{pk}_2}^2(\text{sk}_1)$ is computationally indistinguishable from $\text{Enc}_{\text{pk}_1}^1(0^{|\text{sk}_2|}), \text{Enc}_{\text{pk}_2}^2(\text{sk}_1)$ (here Enc^1 denotes GSW and Enc^2 denotes LHE); this notion (together with non-circular security) implies circular security the way we have defined it (i.e. indistinguishability of encryptions of two messages in the presence of an encrypted key cycle). However, this implication no longer holds in the context of SRL security (see Section 2.8 for more details). And this is why we are directly defining circular security as indistinguishability of encryptions of messages in the presence of an encrypted key cycle.

The above two points indicate that the circular SRL-security w.r.t. GSW and LHE is both (a-priori) stronger, and also somewhat different from a qualitative point of view than the “plain” circular security assumption. Yet, in our eyes, the main justifications for believing the latter holds true are also valid for circular SRL-security:

1. In both cases (plain and SRL), security holds in a non-circular setting, assuming LWE.
2. In both cases (plain and SRL), the security game being considered captures a simple and natural process (albeit for the case of SRL security, it is more complex).
3. Finally, just as for the notion plain circular security, it does not appear simple to even just come up with any *bit*-encryption scheme (such as GSW) that is SRL secure, but not circular SRL secure.⁸

1.5 Concurrent and Subsequent Work

A concurrent and independent breakthrough result by Jain, Lin and Sahai [JLS20] presents a construction of $i\mathcal{O}$ based on subexponential security of well-founded assumptions: (1) the SXDH assumption on asymmetric bilinear groups, (2) the LWE assumption with subexponential modulus-to-noise ratio, (3) a Boolean PRG in NC^0 , and (4) an LPN assumption over a *large field* and with a small error rate $\frac{1}{\ell^\delta}$ where $\delta > 0$ and ℓ is the dimension of the LPN secret. Assumptions (1) and (2) have widespread use and are considered standard. (3) has also been well-studied in recent years. (4) is a very natural coding problem, but the range of parameters used in (4) differs from most prior works in the cryptographic literature, a majority of which focus on a less sparse error rate (typically a constant) and/or use the field \mathbb{F}_2 .

A concurrent and independent work by Wee and Wichs [WW20] presents a new elegant heuristic instantiation of the BDGM paradigm based only on lattice-based primitives. Similarly to us, their construction proceeds by implementing $Xi\mathcal{O}$ with pre-processing. They also state a new security assumption with a circular security flavor (involving a PRF and LWE samples) under which they can prove the security of their construction: Roughly speaking, their construction proceed by reducing $Xi\mathcal{O}$ with pre-processing to the task of “oblivious LWE sampling”, and next they provide a heuristic instantiation of a protocol for performing oblivious LWE sampling. Their security assumption is

⁷This follows from a union bound as the length of both the messages $\mathbf{m}^0, \mathbf{m}^1$ and the description of the leakage-selection algorithm are “short”.

⁸Wichs and Zirdelis [WZ17] show that any public-key bit encryption scheme can modified in a way that preserves security yet violates circular security (using a special form of obfuscation that can be satisfied under LWE). The same method can be used to obtain an SRL-secure encryption scheme (by modifying GSW as in [WZ17]) that is not 1-circular SRL secure.

essentially that their protocol is a secure oblivious LWE sampler. It is worth noting that even though they also rely on the BDGM approach to implement XiO , they manage to directly construct an FHE with short randomness, relying on a “dual” variant of the GSW encryption scheme, thereby completely removing the use of any LHE (whereas we obtain short decryption hints by combining GSW with our Packed Regev).

The initial version of our paper did not contain the LWE-based instantiation of the LHE using Packed Regev (we just had the DJ-based instantiation). Following up on the initial posting of our paper, but concurrently and independently from our LWE-based construction, a preprint by Brakerski et al [BDGM20b] also provides an LWE-based way to instantiate the LHE within our framework. Differently from our construction, however, they rely on a variant of the “Dual Regev” encryption scheme, whereas we rely on regular Regev.

2 Preliminaries and Definitions

In this section, we recall some standard definitions and results. Additionally, we include a formalization of the circular security assumption that we consider.

Attackers, negligible functions and subexponential security. Below, for simplicity of exposition, we provide definitions for *polynomial security* of all the primitives we consider. As usual, we model attackers as *non-uniform probabilistic polynomial-time algorithms*, denoted *nuPPT*. We say that a function $\mu(\cdot)$ is *negligible* if for every polynomial $p(\cdot)$, there exists some $\lambda_0 \in \mathbb{N}$ such that $\mu(\lambda) \leq \frac{1}{p(\lambda)}$ for all $\lambda > \lambda_0$. The security definitions we consider will require that for every nuPPT A , there exists some negligible function μ such that for all λ , A succeeds in “breaking security” w.r.t. the security parameter λ with probability at most $\mu(\lambda)$. All the definitions that we consider can be extended to consider *subexponential security*; this is done by requiring the existence of a constant $\varepsilon > 0$, such that for every non-uniform (probabilistic) attacker \mathcal{A} with running time $\text{poly}(\lambda) \cdot 2^{\lambda^\varepsilon}$, there exists some negligible function μ such that for all λ , \mathcal{A} succeeds in “breaking security” w.r.t. the security parameter λ with probability at most $\mu(\lambda) \cdot 2^{-\lambda^\varepsilon}$ (as opposed to just $\mu(\lambda)$).

Notations. For all $n, m \in \mathbb{N}$, we write $[-n, m] = \{-n, -n+1, \dots, m\}$, $[n] = [1, n]$. For all $\alpha, \beta \in \mathbb{R}$ such that $\beta > \alpha$, we denote by $(\alpha, \beta) = \{x \in \mathbb{R}, \alpha < x < \beta\}$. For all $v_1, \dots, v_n \in \mathbb{Z}$, we denote by $\mathbf{v} = (v_1, \dots, v_n)$ the column vector in \mathbb{Z}^n . For all probabilistic polynomial time (PPT) algorithm \mathcal{A} , we denote by $y \leftarrow \mathcal{A}(x)$ the random process of running \mathcal{A} on input x and obtaining the output y . For all sets \mathcal{S} , we denote by $x \leftarrow_{\mathbb{R}} \mathcal{S}$ the process of sampling a random element x uniformly over \mathcal{S} . For any sequence $x = \{x_\lambda\}_{\lambda \in \mathbb{N}}$, we write $x \in 2^{\text{poly}(\lambda)}$ if there exists a polynomial $p(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $x_\lambda \in \mathbb{N}$ and the bit size of x_λ is less than $p(\lambda)$. For all functions $B(\cdot)$, we say an ensemble $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ of distributions that output in \mathbb{Z} is B -bounded if for all $\lambda \in \mathbb{N}$, all x in the support of \mathcal{D}_λ , we have $|x| < B(\lambda)$. We say an ensemble $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ is polynomially-bounded if there exists a polynomial $p(\cdot)$ such that $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ is p -bounded. We say two ensembles $\{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{D}_\lambda^1\}_{\lambda \in \mathbb{N}}$ are statistically close (without specifying the statistical distance) and we write $\{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{D}_\lambda^1\}_{\lambda \in \mathbb{N}}$ when for all $\lambda \in \mathbb{N}$, the distributions \mathcal{D}_λ^0 and \mathcal{D}_λ^1 have statistical distance $2^{-\Omega(\lambda)}$.

2.1 Standard Lemmas

We first recall a special case of the leftover hash lemma from [ILL89].

Lemma 2.1 (leftover hash lemma). *For all $\lambda, q, d, m \in \mathbb{N}$ such that $m \geq d\lceil \log(q) \rceil + 2\lambda$, the statistical distance between the following distributions is upper bounded by $2^{-\lambda}$:*

$$\left\{ \mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{d \times m}, \mathbf{r} \leftarrow_{\mathbb{R}} [-1, 1]^m : (\mathbf{A}, \mathbf{A}\mathbf{r}) \right\} \\ \left\{ \mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^{d \times m}, \mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_q^d : (\mathbf{A}, \mathbf{u}) \right\}.$$

We will also make use of the following standard “smudging” lemma (see e.g. [AJL⁺12] for an explicit proof of this lemma).

Lemma 2.2 (Smudging). *For all $B, B' \in \mathbb{N}$ such that $B < B'$, all $x \in [-B, B]$, the statistical distance between the following distributions is upper bounded by $\frac{B}{B'}$:*

$$\{u \leftarrow_{\mathbb{R}} [-B', B'] : u\} \\ \{u \leftarrow_{\mathbb{R}} [-B', B'] : u + x\}.$$

2.2 Indistinguishability

We start by recalling the standard definition of computational indistinguishability [GM84].

Definition 2.1 (Computational indistinguishability). *Two ensembles $\{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{D}_\lambda^1\}_{\lambda \in \mathbb{N}}$ are said to be computationally indistinguishable if for every nuPPT \mathcal{A} there exists a negligible function μ such that for every $\lambda \in \mathbb{N}$,*

$$\left| \Pr[\mathcal{A}(1^\lambda, \mathcal{D}_\lambda^0) = 1] - \Pr[\mathcal{A}(1^\lambda, \mathcal{D}_\lambda^1) = 1] \right| \leq \mu(\lambda)$$

We write $\{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{D}_\lambda^1\}_{\lambda \in \mathbb{N}}$ to denote that $\{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{D}_\lambda^1\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable.

2.3 Definition of $i\mathcal{O}$

We recall the definition of $i\mathcal{O}$ [BGI⁺01, GGH⁺13b]. Given polynomials $n(\cdot), s(\cdot), d(\cdot)$, let $\mathcal{C}_{n,s,d} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ denote the class of circuits such that for all $\lambda \in \mathbb{N}$, \mathcal{C}_λ is the set of circuits with input size $n(\lambda)$, size at most $s(\lambda)$ and depth at most $d(\lambda)$. We say that a sequence of circuits $\{\Pi_\lambda\}_{\lambda \in \mathbb{N}}$ is contained in $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ (denoted by $\{\Pi_\lambda\}_{\lambda \in \mathbb{N}} \in \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$) if for all $\lambda \in \mathbb{N}$, $\Pi_\lambda \in \mathcal{C}_\lambda$.

Definition 2.2 ($i\mathcal{O}$ for P/poly). *We say that $i\mathcal{O}$ exists for P/poly if for all polynomials $n(\cdot), s(\cdot), d(\cdot)$, there exists a tuple of PPT algorithms (Obf, Eval) such that the following holds:*

- **Correctness:** *For all $\{\Pi_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_{n,s,d}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, all $\mathbf{x} \in \{0, 1\}^{n(\lambda)}$,*

$$\Pr[\tilde{\Pi} \leftarrow \text{Obf}(1^\lambda, \Pi_\lambda) : \text{Eval}(1^\lambda, \tilde{\Pi}, \mathbf{x}) = \Pi(\mathbf{x})] \geq 1 - \mu(n)$$

- **IND-security:** *For all sequences $\{\Pi_0^\lambda\}_{\lambda \in \mathbb{N}}, \{\Pi_1^\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_{n,s,d}$ such that for all $\lambda \in \mathbb{N}$, Π_0^λ and Π_1^λ are functionally equivalent circuits, the following ensembles are computationally indistinguishable:*

$$\left\{ \tilde{\Pi} \leftarrow \text{Obf}(1^\lambda, \Pi_0^\lambda) : \tilde{\Pi} \right\}_{\lambda \in \mathbb{N}} \\ \left\{ \tilde{\Pi} \leftarrow \text{Obf}(1^\lambda, \Pi_1^\lambda) : \tilde{\Pi} \right\}_{\lambda \in \mathbb{N}}$$

2.4 Definition of XiO

We recall the definition of XiO with pre-processing [LPST16]. We restrict our attention to circuits with input length $O(\log \lambda)$: Given polynomials $n(\cdot), s(\cdot), d(\cdot)$, let $\mathcal{C}_{\log(n),s,d} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ denote the class of circuits such that for all $\lambda \in \mathbb{N}$, \mathcal{C}_λ is the set of circuits with input size $\log(n(\lambda))$, size at most $s(\lambda)$ and depth at most $d(\lambda)$.

Definition 2.3 (XiO for P^{\log}/poly). *We say XiO exists for P^{\log}/poly if there exists a polynomial $p(\cdot)$ and a constant $\varepsilon \in (0, 1)$, such that for all polynomials $n(\cdot), s(\cdot), d(\cdot)$, there exists a tuple of PPT algorithms $(\text{Gen}_{\text{Obf}}, \text{Obf}, \text{Eval})$ such that the following holds:*

- **Correctness:** *For all $\{\Pi_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_{\log(n),s,d}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, all $\mathbf{x} \in \{0, 1\}^{n(\lambda)}$:*

$$\Pr[\text{pp} \leftarrow \text{Gen}_{\text{Obf}}(1^\lambda), \tilde{\Pi} \leftarrow \text{Obf}(\text{pp}, \Pi_\lambda) : \text{Eval}(\text{pp}, \tilde{\Pi}, \mathbf{x}) = \Pi(\mathbf{x})] \geq 1 - \mu(\lambda)$$

- **Succinctness:** *For all $\{\Pi_\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_{\log(n),s,d}$, all $\lambda \in \mathbb{N}$, all pp in the support of $\text{Gen}_{\text{Obf}}(1^\lambda)$, all $\tilde{\Pi}$ in the support of $\text{Obf}(\text{pp}, \Pi_\lambda)$, we have that $|\tilde{\Pi}| \leq n(\lambda)^{1-\varepsilon} \cdot p(\lambda, s(\lambda), d(\lambda))$*
- **IND-security:** *For all sequences $\{\Pi_0^\lambda\}_{\lambda \in \mathbb{N}}, \{\Pi_1^\lambda\}_{\lambda \in \mathbb{N}} \in \mathcal{C}_{\log(n),s,d}$ such that for all $\lambda \in \mathbb{N}$, Π_0^λ and Π_1^λ are functionally equivalent circuit, the following ensembles are computationally indistinguishable:*

$$\left\{ \text{pp} \leftarrow \text{Gen}_{\text{Obf}}(1^\lambda), \tilde{\Pi} \leftarrow \text{Obf}(\text{pp}, \Pi_\lambda^0) : (\text{pp}, \tilde{\Pi}) \right\}_{\lambda \in \mathbb{N}}$$

$$\left\{ \text{pp} \leftarrow \text{Gen}_{\text{Obf}}(1^\lambda), \tilde{\Pi} \leftarrow \text{Obf}(\text{pp}, \Pi_\lambda^1) : (\text{pp}, \tilde{\Pi}) \right\}_{\lambda \in \mathbb{N}}$$

The following theorem from [LPST16] connects XiO (with pre-processing) with iO assuming the LWE assumption (we formally define the LWE assumption in Definition 3.4).

Theorem 2.4. *Assume the existence of a subexponentially-secure XiO for P^{\log}/poly , and assume subexponential security of the LWE assumption. Then there exists an (subexponentially-secure) iO for P/poly .*

2.5 Definition of Public-Key Encryption

We start by recalling the definition of public key encryption (PKE). For our purposes, we will consider PKE in a Common Reference String (CRS) model, where we first generate a CRS, and next, the key generation algorithm will take the CRS as input. This added generality will be useful to capture scenarios where multiple encryption schemes will be operating over the same ring \mathbb{Z}_N —this ring can be specified in the CRS.

Definition 2.5 (Public-Key Encryption). *A Public-Key Encryption (PKE) scheme is a tuple of PPT algorithms $(\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec})$ where:*

- $\text{CRSgen}(1^\lambda)$: *given as input the security parameter $\lambda \in \mathbb{N}$, it outputs a common reference string crs.*
- $\text{Gen}(\text{crs})$: *given as input crs, it outputs the pair (pk, sk) .*

- $\text{Enc}_{\text{pk}}(m; r)$: given as input the public key pk , a message $m \in \{0, 1\}^*$ and some randomness $r \leftarrow_{\mathbb{R}} \{0, 1\}^{\infty}$ ⁹, it outputs a ciphertext ct .
- $\text{Dec}_{\text{sk}}(\text{ct})$: given as input the secret key sk and a ciphertext ct , it deterministically outputs a plaintext.

We furthermore require these algorithms to satisfy the following correctness condition: for all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all pairs (pk, sk) in the support $\text{Gen}(\text{crs})$, all messages $m \in \{0, 1\}^*$, all ciphertexts ct in the support of $\text{Enc}_{\text{pk}}(m)$, we have:

$$\text{Dec}_{\text{sk}}(\text{ct}) = m.$$

2.6 Definition of Linearly-Homomorphic Encryption

Definition 2.6 (Linearly-Homomorphic Encryption). For any polynomial $\ell(\cdot)$, a PKE scheme $(\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec})$ is said to be a *Linearly-Homomorphic Encryption (LHE)* with plaintext size $\ell(\cdot)$, if there exists a PPT algorithm Add such that the following holds:

- For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all (pk, sk) in the support of $\text{Gen}(\text{crs})$, the public key pk contains a message space $(\mathbb{A}_{\text{pk}}, +)$, which is an Abelian group of size $|\mathbb{A}| > 2^{\ell(\lambda)}$.
- For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all (pk, sk) in the support of $\text{Gen}(\text{crs})$, all messages $m_1, m_2 \in \mathbb{A}_{\text{pk}}$, all ciphertexts ct_1, ct_2 in the support of $\text{Enc}_{\text{pk}}(m_1), \text{Enc}_{\text{pk}}(m_2)$ respectively, the algorithm $\text{Add}(\text{pk}, \text{ct}_1, \text{ct}_2)$ deterministically outputs a ciphertext in the support of $\text{Enc}_{\text{pk}}(m_1 + m_2)$, where the addition is performed in \mathbb{A}_{pk} .

2.7 Definition of Fully Homomorphic Encryption

Definition 2.7 (Fully-Homomorphic Encryption). A PKE scheme $(\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec})$ is said to be a *Fully-Homomorphic Encryption (FHE)* scheme for depth $\delta(\cdot)$ circuits if there exists a PPT algorithm Eval such that for all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all pairs (pk, sk) in the support of Gen , all $n \in \mathbb{N}$, all messages $m_1, \dots, m_n \in \{0, 1\}$, all ciphertexts $\text{ct}_1, \dots, \text{ct}_n$ in the support of $\text{Enc}_{\text{pk}}(m_1), \dots, \text{Enc}_{\text{pk}}(m_n)$ respectively, all circuits $f : \{0, 1\}^n \rightarrow \{0, 1\}$ of depth at most $\delta(\lambda)$, $\text{Eval}(\text{pk}, f, \text{ct}_1, \dots, \text{ct}_n)$ deterministically outputs an evaluated ciphertext ct_f such that $\text{Dec}_{\text{sk}}(\text{ct}_f) = f(m_1, \dots, m_n)$.

Note that the depth of the circuit that are homomorphically evaluated is a priori bounded by $\delta(\lambda)$ for a polynomial δ (that is, we consider the case of *leveled* FHE). The arity of the evaluated circuits (denoted by n above), however, is unbounded. The FHE we will be using — namely, from [GSW13] — natively supports arithmetic circuits (with addition and multiplication gates), which capture Boolean circuits.

2.8 Leakage-resilient and Circular Security

We recall the standard definition of CPA-security for encryption schemes; we furthermore generalize it to a notion of \mathcal{O} -leakage resilient security, which extends the standard definition by also providing the attacker with access to a leakage oracle \mathcal{O} receiving the public key pk , the message \mathbf{m}^* being encrypted, and the randomness \mathbf{r} under which it is encrypted. Our notion of \mathcal{O} leakage-resilience restricts to attackers that only make “valid” leakage queries, where a query is said to be valid if

⁹As usual, since all algorithms are PPT we really only need to consider a finite prefix of $\{0, 1\}^\infty$ to define the uniform distribution.

the oracle does not return \perp in response to it. In more detail, to “win” in the security game, the attacker \mathcal{A} must (a) correctly guess which among two message $\mathbf{m}^0, \mathbf{m}^1$ is being encrypted, while (b) not having made any queries to \mathcal{O} on which \mathcal{O} returns \perp .

Definition 2.8 (\mathcal{O} -leakage resilient security). *We say that a public-key encryption scheme $\mathcal{PK}\mathcal{E} = (\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec})$ is \mathcal{O} -leakage resilient secure if for all stateful nuPPT adversaries \mathcal{A} , there exists some negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\Pr[\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}} = 1] \leq 1/2 + \mu(\lambda)$, where the experiment $\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}}$ is defined as follows:*

$$\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}} = \left\{ \begin{array}{l} \text{crs} \leftarrow \text{CRSgen}(1^\lambda), (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{crs}) \\ (\mathbf{m}^0, \mathbf{m}^1) \leftarrow \mathcal{A}(\text{pk}), b \leftarrow \{0, 1\} \\ \mathbf{m}^* = \mathbf{m}^b, \mathbf{r} \leftarrow_{\mathbb{R}} \{0, 1\}^\infty \\ \text{ct} = \text{Enc}_{\text{pk}}(\mathbf{m}^*; \mathbf{r}), b' \leftarrow \mathcal{A}^{\mathcal{O}(\text{pk}, \mathbf{m}^*, \mathbf{r})}(\text{ct}) \\ \text{Return } 1 \text{ if } |\mathbf{m}^0| = |\mathbf{m}^1|, b' = b \text{ and } \mathcal{O} \text{ did not return } \perp; 0 \text{ otherwise.} \end{array} \right\}$$

We say that $\mathcal{PK}\mathcal{E}$ is simply secure if the above holds when we do not give \mathcal{A} access to an oracle.

We will also consider a 2-circular secure variant of \mathcal{O} -leakage resilient security, which is similarly defined except we require indistinguishability of \mathbf{m}^0 and \mathbf{m}^1 in the presence not only of some randomness leakage, but also of an encrypted key cycle w.r.t. two public-key encryption schemes $\mathcal{PK}\mathcal{E}$ and $\overline{\mathcal{PK}\mathcal{E}}$. Note that we set the CRS of $\mathcal{PK}\mathcal{E}$ to be the public key of $\overline{\mathcal{PK}\mathcal{E}}$; this is to ensure compatibility between the schemes, i.e. for them to operate on the same ring.

Definition 2.9 (\mathcal{O} -leakage resilient 2-circular security). *We say that public-key encryption schemes $\mathcal{PK}\mathcal{E} = (\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec})$ and $\overline{\mathcal{PK}\mathcal{E}} = (\overline{\text{CRSgen}}, \overline{\text{Gen}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ are \mathcal{O} -leakage resilient 2-circular secure if for all stateful nuPPT adversaries \mathcal{A} , there exists some negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\Pr[\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}, \overline{\mathcal{PK}\mathcal{E}}} = 1] \leq 1/2 + \mu(\lambda)$, where the experiment $\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}, \overline{\mathcal{PK}\mathcal{E}}}$ is defined as follows:*

$$\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{PK}\mathcal{E}, \overline{\mathcal{PK}\mathcal{E}}} = \left\{ \begin{array}{l} \overline{\text{crs}} \leftarrow \overline{\text{CRSgen}}(1^\lambda), (\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Gen}}(\overline{\text{crs}}), (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\overline{\text{pk}}) \\ (\mathbf{m}^0, \mathbf{m}^1) \leftarrow \mathcal{A}(\overline{\text{pk}}, \text{pk}), b \leftarrow \{0, 1\}, \mathbf{m}^* = \overline{\text{sk}} \parallel \mathbf{m}^b, \mathbf{r} \leftarrow_{\mathbb{R}} \{0, 1\}^\infty, \text{ct} = \text{Enc}_{\text{pk}}(\mathbf{m}^*; \mathbf{r}) \\ \overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\overline{\text{sk}}, b' \leftarrow \mathcal{A}^{\mathcal{O}(\text{pk}, \mathbf{m}^*, \mathbf{r})}(\text{ct}, \overline{\text{ct}})) \\ \text{Return } 1 \text{ if } |\mathbf{m}^0| = |\mathbf{m}^1|, b' = b \text{ and } \mathcal{O} \text{ did not return } \perp; 0 \text{ otherwise.} \end{array} \right\}$$

We finally state the 2CIRC assumption that we will rely in our main theorem.

Definition 2.10 (2CIRC assumption). *We say that the (subexponential) 2CIRC $^{\mathcal{O}}$ assumption holds w.r.t $\mathcal{PK}\mathcal{E}$ and $\overline{\mathcal{PK}\mathcal{E}}$ if the following holds: if $\mathcal{PK}\mathcal{E}$ is (subexponentially) \mathcal{O} -leakage resilient secure and $\overline{\mathcal{PK}\mathcal{E}}$ is (subexponentially) secure, then (subexponential) \mathcal{O} -leakage resilient 2-circular security holds w.r.t $\mathcal{PK}\mathcal{E}$ and $\overline{\mathcal{PK}\mathcal{E}}$.*

A Note of the Definition of Circular Security. Let us remark that while the above definition of circular security (i.e. indistinguishability of encrypted messages in the presence of an encrypted key-cycle) is the most direct way of capturing the needs for circular security in applications (think e.g., of encrypted disk space), an alternative definition is also commonly used in the literature: it requires (a) standard security, and (b) indistinguishability of the encrypted key cycle $\text{Enc}(\overline{\text{sk}}), \overline{\text{Enc}}(\text{sk})$ from $\text{Enc}(0^{|\text{sk}|}), \overline{\text{Enc}}(\text{sk})$. We want to emphasize that whereas in the standard setting—*without a leakage oracle*—this alternative definition implies the circular security notion we gave, this is no longer true in the oracle-enhanced setting. To see this, consider an oracle that given $\text{pk}, \mathbf{m}^*, \mathbf{r}$, outputs \perp to

any query in case $\mathbf{m}^* = 0^{|\overline{\text{sk}}|}$. Leakage queries can never be useful to distinguish encryptions of $\overline{\text{sk}}$ and $0^{|\overline{\text{sk}}|}$ —as the oracle will output \perp with probability negligibly close to $1/2$ and thus the attacker’s advantage is negligible—so \mathcal{O} -leakage resilient circular security is equivalent to plain (i.e. without an oracle) circular security. Yet, leakage queries w.r.t. such an oracle may be useful when considering indistinguishability between $\overline{\text{sk}}\|\mathbf{m}_0$ and $\overline{\text{sk}}\|\mathbf{m}_1$. In fact, for the particular leakage queries that we will be relying on, this phenomena does occur: they are valid in case \mathbf{m}^* is of the form $\overline{\text{sk}}\|\mathbf{m}$ where $\overline{\text{sk}}$ is a valid secret key for $\overline{\mathcal{PKE}}$, but they are invalid if $\overline{\text{sk}} = 0^{|\overline{\text{sk}}|}$. For this reason, we directly formalize circular security as indistinguishability of encrypted messages in the presence of an encrypted key cycle.

3 Shielded Randomness Leakage Security of GSW

In this section, we define our notion of Shielded Randomness Leakage (SRL) security, which corresponds to \mathcal{O} -leakage resilience security for a particular leakage oracle \mathcal{O} . Then, we prove the GSW FHE is SRL secure under the LWE assumption.

3.1 Definition of Shielded Randomness Leakage Security

To define our notion of SRL security, we focus on FHE schemes that satisfy the following properties.

3.1.1 Batch correctness

This property states that decryption of evaluated ciphertexts solely consists of computing the inner product of the evaluated ciphertext with the secret key (both of which are vectors), then rounding. Also, a single scalar obtained by decryption can encode many output bits of the evaluated function. That is, we consider FHE scheme where the crs contains a modulus N_{crs} such that decryption of an evaluated ciphertext yields a scalar in $\mathbb{Z}_{N_{\text{crs}}}$. Our definition of FHE is flexible w.r.t. the choice of the modulus N_{crs} , which we can afford since the LWE assumption holds for essentially any (large enough) modulus. As observed in [Mic19, BDGM19, BDGM20a], most existing FHE schemes can fit this framework.

Definition 3.1 (Batch correctness). *For all polynomials δ , an FHE scheme $(\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ for depth- δ circuits satisfies batch correctness if there exist a PPT Eval' and a polynomial σ such that following holds:*

- For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$ contain a modulus $N_{\text{crs}} \in \mathbb{N}$; for all (pk, sk) in the support of $\text{Gen}(\text{crs})$, we have: pk contains $B_{\text{pk}} \in \mathbb{N}$ such that $N_{\text{crs}} \geq 2^\lambda B_{\text{pk}}$; the secret key is of the form: $\text{sk} \in \mathbb{Z}^{\sigma(\lambda)}$.
- For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all (pk, sk) in the support of $\text{Gen}(\text{crs})$, all arities $\nu \in \mathbb{N}$, all messages $m_1, \dots, m_\nu \in \{0, 1\}$, all depth- $\delta(\lambda)$ circuits f of arity ν , all ciphertexts ct_i in the support of $\text{Enc}_{\text{pk}}(m_i)$ for all $i \in [\nu]$, all scaling factors $\omega < \log(N_{\text{crs}})$, the algorithm $\text{Eval}'(\text{pk}, f, \omega, \text{ct}_1, \dots, \text{ct}_\nu)$ deterministically outputs an evaluated ciphertext $\text{ct}_f \in \mathbb{Z}_{N_{\text{crs}}}^{\sigma(\lambda)}$ such that:

$$\text{sk}^\top \text{ct}_f = 2^\omega f(\mathbf{m}) + \text{noise}_f \in \mathbb{Z}_{N_{\text{crs}}},$$

with $|\text{noise}_f| < B_{\text{pk}}$.

Note that one can recover the value $f(\mathbf{m}) \in \mathbb{Z}_{N_{\text{crs}}}$ when using the scaling factor $\omega = \lceil \log(B_{\text{pk}}) \rceil + 1$. That is, we can define $\text{Eval}(\text{pk}, f, \text{ct}_1, \dots, \text{ct}_\nu) = \text{Eval}'(\text{pk}, f, \lceil \log(B_{\text{pk}}) \rceil + 1, \text{ct}_1, \dots, \text{ct}_\nu)$.

3.1.2 Randomness homomorphism

This property states that it is possible to homomorphically evaluate a circuit f not only on the ciphertexts, but also the randomness used by the ciphertexts. The resulting evaluated randomness \mathbf{r}_f belongs to a noisy randomness space \mathcal{R}^* — typically the fresh randomness comprises noises, and the evaluated randomness consists of larger-magnitude noises. The encryption algorithm Enc^* is essentially the same as Enc except it operates on the evaluated (noisier) randomness. The ciphertext obtained by first evaluating the randomness, then using the noisy encryption algorithm Enc^* is the same as obtained by directly evaluating the original ciphertexts.

Definition 3.2 (Randomness homomorphism). *An FHE scheme $\mathcal{FHE} = (\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Dec}, \text{Eval})$ for depth- δ circuits that satisfies batch correctness (defined above) also satisfies randomness homomorphism if there exists a sequence of noisy randomness spaces $\{\mathcal{R}_\lambda^*\}_{\lambda \in \mathbb{N}}$, and the following additional PPT algorithms:*

- $\text{Eval}_{\text{rand}}(\mathbf{pk}, f, \mathbf{r}, \mathbf{m})$: given as input the public key \mathbf{pk} , a depth- $\delta(\lambda)$ circuit f of arity ν , random coins $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_\nu)$ where for all $i \in [\nu]$, $\mathbf{r}_i \in \{0, 1\}^\infty$, and messages $\mathbf{m} \in \{0, 1\}^\nu$, it deterministically outputs an evaluated randomness $\mathbf{r}_f \in \mathcal{R}_\lambda^*$.
- $\text{Enc}_{\mathbf{pk}}^*(m; \mathbf{r}^*)$: given as input the public key \mathbf{pk} , a message $m \in \mathbb{Z}_{N_{\text{crs}}}$ and the randomness $\mathbf{r}^* \in \mathcal{R}^*$, it outputs a noisy ciphertext ct^* .

We furthermore require these algorithms to satisfy the following condition: for every $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all pairs $(\mathbf{pk}, \mathbf{sk})$ in the support of $\text{Gen}(\text{crs})$, all $\nu \in \mathbb{N}$, all depth- $\delta(\lambda)$ circuits f of arity ν , all messages $m_i \in \{0, 1\}$, all randomness $\mathbf{r}_i \in \{0, 1\}^\infty$ for $i \in [\nu]$, denoting $\text{ct}_i = \text{Enc}_{\mathbf{pk}}(m_i; \mathbf{r}_i)$ and $\mathbf{r}_f = \text{Eval}_{\text{rand}}(\mathbf{pk}, f, \mathbf{r}, \mathbf{m})$, we have $\mathbf{r}_f \in \mathcal{R}_\lambda^*$ and:

$$\text{Eval}'(\mathbf{pk}, f, 0, \text{ct}_1, \dots, \text{ct}_\nu) = \text{Enc}_{\mathbf{pk}}^*(f(\mathbf{m}); \mathbf{r}_f).$$

3.1.3 Shielded Randomness-Leakage security

We proceed to formally define shielded randomness leakage (SRL) security for randomness homomorphic FHEs with batch correctness. SRL security will be defined as \mathcal{O} -leakage resilient security for a particular leakage oracle \mathcal{O}_{SRL} that given the public key \mathbf{pk} , a message \mathbf{m}^* and randomness \mathbf{r} , allows the attacker \mathcal{A} to ask to see a “shielded” version of the homomorphically evaluated randomness \mathbf{r}_f for any function f for which \mathcal{A} knows the output $f(\mathbf{m}^*)$. To make sure the attacker can only query the oracle with functions on which it knows the output, we require the attacker to also provide the output α , and the oracle outputs \perp if $f(\mathbf{m}^*) \neq \alpha$ (and thus, by the definition of \mathcal{O} -leakage resilient security, the attacker fails if it ever picks a function for which it does not know the output).

To formalize the SRL oracle, we restrict ourselves to FHE where the noisy randomness consists of integer vectors. That is, there exists a polynomial $t(\cdot)$ such that the sequence $\{\mathcal{R}_\lambda^*\}_{\lambda \in \mathbb{N}}$ is such that for all $\lambda \in \mathbb{N}$, $\mathcal{R}_\lambda^* \subseteq \mathbb{Z}^{t(\lambda)}$. Henceforth, we denote by $\mathbf{r}_1 + \mathbf{r}_2 \in \mathcal{R}_\lambda^*$ and $\mathbf{r}_1 - \mathbf{r}_2 \in \mathcal{R}_\lambda^*$ the addition and subtraction in $\mathbb{Z}^{t(\lambda)}$. We denote \mathcal{R}_λ^* by \mathcal{R}^* for simplicity.

Definition 3.3 (SRL security). *An FHE scheme \mathcal{FHE} for depth δ circuits satisfying randomness homomorphism is said to be SRL-secure if it is $\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}$ -leakage resilient secure for the following oracle $\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}$, where $\text{Eval}_{\text{rand}}$ and Enc^* are the algorithms guaranteed to exist by the definition of randomness homomorphism. Similarly, for any public-key encryption scheme $\overline{\mathcal{PKE}}$, we say 2-circular SRL security holds w.r.t. \mathcal{FHE} and $\overline{\mathcal{PKE}}$ if $\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}$ -leakage resilient 2-circular security holds w.r.t. \mathcal{FHE} and $\overline{\mathcal{PKE}}$.*

$\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}(\text{pk}, \mathbf{m}^*, \mathbf{r})$:
 $\mathbf{r}^* \leftarrow_{\mathcal{R}} \mathcal{R}^*$, $\text{ct}^* = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$
 $(f, \alpha) \leftarrow \mathcal{A}(\text{ct}^*)$
 $\mathbf{r}_f = \text{Eval}_{\text{rand}}(\text{pk}, f, \mathbf{r}, \mathbf{m}^*)$.
 If $f(\mathbf{m}^*) = \alpha$ and f is of depth at most δ , then $\text{leak} = \mathbf{r}^* - \mathbf{r}_f \in \mathcal{R}^*$.
 Otherwise, $\text{leak} = \perp$. Return leak .

Roughly speaking, given a message \mathbf{m}^* and randomness \mathbf{r} , the oracle $\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}$ samples fresh random coins \mathbf{r}^* from which it generates a noisy encryption of zero, that is sent to the adversary. The adversary next chooses a circuit f and a value $\alpha \in \mathbb{Z}$. The oracle then checks that $f(\mathbf{m}^*) = \alpha$, upon which it returns the evaluated randomness “shielded” with the randomness \mathbf{r}^* ; otherwise, it outputs \perp and in this case, the attacker fails.

In the concrete FHE we consider from [GSW13], the randomness leakage corresponds to the randomness obtained from homomorphically subtracting the evaluated challenge ciphertext from $\text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$. Revealing such leakage allows the adversary to decrypt and recover the value $0 - f(\mathbf{m}^*)$. This is why we only allow the attacker to request leakage functions f for which it knows the output $f(\mathbf{m}^*)$.

Whenever the scheme \mathcal{FHE} is clear from context, we simply write \mathcal{O}_{SRL} to denote $\mathcal{O}_{\text{SRL}}^{\mathcal{FHE}}$.

A note on the falsifiability and interactivity of (circular) SRL security. We note that both SRL (and 2-circular) SRL security of an FHE (and a PKE) is a simple and natural *interactive* falsifiable assumptions about the FHE (and the PKE): the assumption is defined as an interactive security game involving a PPT challenger \mathcal{C} , with a threshold of $1/2$ (i.e. the attacker needs to win with probability non-negligibly higher than $1/2$). Let us make some additional observations about this assumption:

- Let us first point out that for our actual application, we only need to rely on a relaxed form of SRL security where \mathcal{A} sends all its queries in parallel. Thus, such a “parallel” SRL assumption can be captured as a 6-round security game (1) \mathcal{A} first gets the public key, (2) \mathcal{A} picks the messages $\mathbf{m}^0, \mathbf{m}^1$, (3) \mathcal{A} gets the encryptions of \mathbf{m}^b and the shields, (4) \mathcal{A} selects the parallel leakage queries $\{f_i, \alpha_i\}$, (5) \mathcal{A} gets back the (shielded) randomness $\{\mathbf{r}_{f_i}\}$ and (6) \mathcal{A} finally makes a guess for b . In fact, we actually do not need CPA security for adaptively chosen messages so we can compress it to a 5-round assumption. We write down this concrete 5-round assumption that suffices for us in Appendix B.
- Next, let us note that for our application, we only need SRL security for: (1) *short* messages $\mathbf{m}^0, \mathbf{m}^1$, of length λ^ε , for some $\varepsilon \in (0, 1)$, (2) a query-selecting mechanism that has some fixed polynomial running time, and whose description size is only $\lambda^\varepsilon + O(1)$. For such applications, we can stipulate a *non-interactive* SRL-security game: the challenger picks *random* messages $\mathbf{m}^0, \mathbf{m}^1$, and a *random* query-selecting machine TM M , both of size $O(\lambda^\varepsilon)$. Next, it checks that M generates valid queries (within the a-priori fixed time-bound) w.r.t. to *both* $\mathbf{m}^0, \mathbf{m}^1$ (or in the case of circular security w.r.t. $\overline{\text{sk}}||\mathbf{m}^0$ and $\overline{\text{sk}}||\mathbf{m}^1$). If so, the challenger gives the attacker the public key, an encryption of \mathbf{m}^* , the shields and the SRL leakage. The attacker wins if it correctly guesses b . If not (i.e. some of the SRL queries were invalid), the same information *excluding the SRL leakage* is sent to the attacker.

If this non-interactive SRL-assumption is $2^{O(\lambda^{\varepsilon'})}$ -hard, where $\varepsilon' > \varepsilon$, it implies that it is still $2^{O(\lambda^{\varepsilon'})}$ -hard for *every* (short) choices of $\mathbf{m}^0, \mathbf{m}^1$ and M by a union bound over $\mathbf{m}^0, \mathbf{m}^1, M$,

and this is exactly what is needed for our $Xi\mathcal{O}$ proof, as the query-selecting machine selects leakage queries that are valid w.r.t. $\overline{\text{sk}}\|\mathbf{m}^0$ and $\overline{\text{sk}}\|\mathbf{m}^1$ (with high probability).

While the above discussion shows that we could have presented a relatively simple *non-interactive* (and falsifiable) circular SRL-assumption (which we can prove holds w.r.t. GSW based on LWE in the non-circular setting), in our opinion, doing so does not elucidate the assumption on a qualitative level. Rather, we have chosen to present the circular SRL-assumption in a more general and stronger form as we believe this interactive version is: (1) *more natural*—it captures a very natural (in our eyes) security game, and (2) it is *still secure in the non-circular setting* (based on LWE), and (3) due to the fact that the assumption is *stronger and simpler*, it becomes easier to attack (and conversely, the absence of attacks would inspire more confidence).

Let us furthermore note that for our proof, we only need to consider very specific SRL leakage queries and thus an alternative way of making the assumption “technically weaker” is to restrict to those types of queries, but we believe that would just make the assumption more ad-hoc, without adding any real reason for increased confidence in the security.

3.2 SRL Security of the GSW FHE from LWE

We now recall the FHE scheme from [GSW13], whose security relies on the LWE assumption. The variant we present uses a large modulus to permit batching many output bits in a single scalar. We prove the GSW scheme is SRL-secure (as per Definition 3.3) under the LWE assumption.

3.2.1 Learning With Error Assumption

We recall the Learning with Error (LWE) assumption with subexponential modulus-to-noise ratio. In [Reg05], Regev showed that solving the LWE problem with modulus q , dimension κ , arbitrary number of samples m , and discrete Gaussian distribution χ of standard deviation $\sigma = \alpha q \geq 2\sqrt{\kappa}$ (this is the distribution over \mathbb{Z} that follows the normal distribution of standard deviation σ , and it is such that $\Pr[e \leftarrow \chi : |e| > \sigma\sqrt{\log(\kappa)}] \in 2^{-\Omega(\kappa)}$) is at least as hard as quantumly approximating the shortest independent vector problem (SIVP) to within an approximation factor $\gamma = \tilde{\mathcal{O}}(\kappa/\alpha)$ in the *worst case* κ -dimensional lattices. His result only applied to every modulus q that is a prime power, or a product of small (poly-size) distinct primes. Later, in [PRS17], the result was generalized to any modulus q .

As typical, we choose a noise-to-modulus ratio $\alpha = 2^{-\kappa^c}$ for a constant $c \in (0, 1)$, which corresponds to the SIVP problem with an approximation factor $\gamma = \tilde{\mathcal{O}}(\kappa \cdot 2^{\kappa^c})$, which is believed to be intractable for c small enough.

Originally, the LWE assumption was defined for uniformly random secrets; however, several works [Mic01, ACPS09] showed that LWE is no easier if the secret is drawn from the noise distribution, which is the variant we will use. Finally, we use a noise distribution that is bounded with probability 1 (as opposed to all but negligible probability, as it is the case for the discrete Gaussian distribution that is commonly used). This mild strengthening is not fundamental but will make our definitions easier to work with (e.g. correctness will hold with probability 1). That is, we rely on the following LWE assumption.

Definition 3.4 (LWE assumption [Reg05, PRS17]). *For all sequences $q \in 2^{\text{poly}(\kappa)}$, all ensembles χ of efficiently sampleable distributions over \mathbb{Z} , we say that (subexponential) security of the LWE assumption holds w.r.t. the sequence q and the ensemble χ if for all polynomials $m(\cdot)$, the following*

ensembles are (subexponentially) computationally indistinguishable:

$$\left\{ \mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_{q_\kappa}^{m(\kappa) \times \kappa}, \mathbf{s} \leftarrow \chi_\kappa^\kappa, \mathbf{e} \leftarrow \chi_\kappa^{m(\kappa)}, \mathbf{z} = \mathbf{A}\mathbf{s} + \mathbf{e} \in \mathbb{Z}_{q_\kappa}^{m(\kappa)} : (\mathbf{A}, \mathbf{z}) \right\}_{\kappa \in \mathbb{N}}.$$

$$\left\{ \mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_{q_\kappa}^{m(\kappa) \times \kappa}, \mathbf{z} \leftarrow_{\mathbb{R}} \mathbb{Z}_{q_\kappa}^{m(\kappa)} : (\mathbf{A}, \mathbf{z}) \right\}_{\kappa \in \mathbb{N}}.$$

We say the (subexponential) security of the LWE assumption holds if there exists a constant $c \in (0, 1)$ such that for all sequences $q \in 2^{\text{poly}(\kappa)}$ and all polynomials B such that for all $\kappa \in \mathbb{N}$, the following holds:

- $B(\kappa) \geq 2\sqrt{\kappa \log(\kappa)}$
- $B(\kappa) \geq q_\kappa 2^{-\kappa^c}$

there exists a B -bounded ensemble χ of efficiently sampleable distributions over \mathbb{Z} , such that (subexponential) security of the LWE assumption holds w.r.t. q, χ .

3.2.2 The GSW scheme

We recall the FHE from [GSW13]. We present the leveled variant (without bootstrapping), which is parameterized by a polynomial δ that bounds the depth of the circuits that can be homomorphically evaluated. Its security relies on the LWE assumption with a subexponential modulus-to-noise ratio.

We denote the scheme by GSW_δ .

Notations. For all polynomials δ , we denote by b_δ the polynomial such that for all polynomially-bounded ensemble χ , all polynomials κ , all $\lambda \in \mathbb{N}$, the noise obtained from homomorphically evaluating circuits of depth at most $\delta(\lambda)$ on GSW ciphertexts generated with LWE noise sampled from the distribution χ_λ and LWE secret of dimension $\kappa(\lambda)$, is upper bounded by $2^{b_\delta(\lambda)}$.

• Gen(crs):

Given as input crs which contains a modulus $N \geq 2^{2\lambda + b_\delta(\lambda)}$, it chooses a sequence $\{q_n\}_{n \in \mathbb{N}}$, a B_χ -bounded ensemble $\{\chi_n\}_{n \in \mathbb{N}}$ for a polynomial B_χ and a polynomial κ such that LWE holds w.r.t. q and χ , and $q_{\kappa(\lambda)} = N$ (by the LWE assumption, given in Definition 3.4, we know such parameters exist). We abuse notations and write $\kappa = \kappa(\lambda)$, $\chi = \chi_{\kappa(\lambda)}$ and $B_\chi = B_\chi(\kappa(\lambda))$ from here on. The algorithm sets $w = (\kappa + 1)\lceil \log(N) \rceil$, $m = 2(\kappa + 1)\lceil \log(N) \rceil + 2\lambda$, $B^* = 2^\lambda(w + 1)^\delta \lceil \log(N) \rceil$ and $B = B_\chi(w + 1)^\delta \lceil \log(N) \rceil m$. Note that we have $N \geq 2^{2\lambda} B$.

It samples $\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times m}$, $\mathbf{s} \leftarrow \chi^\kappa$, $\mathbf{e} \leftarrow \chi^m$, $\mathbf{g} = (1, 2, \dots, 2^{\lceil \log(N) \rceil - 1}) \in \mathbb{Z}_N^{\lceil \log(N) \rceil}$, $\mathbf{G} = \mathbf{g}^\top \otimes \text{Id} = \begin{pmatrix} \mathbf{g}^\top & 0 & \dots \\ 0 & \mathbf{g}^\top & \\ \vdots & & \ddots \end{pmatrix} \in \mathbb{Z}_N^{(\kappa+1) \times w}$ where $\text{Id} \in \mathbb{Z}_N^{(\kappa+1) \times (\kappa+1)}$ denotes the identity matrix,

$\mathbf{U} = \begin{pmatrix} \mathbf{A} \\ \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \end{pmatrix} \in \mathbb{Z}_N^{(\kappa+1) \times m}$. It sets $\text{pk} = (B, \mathbf{U}, \mathbf{G})$, and $\text{sk} = (-\mathbf{s}, 1) \otimes \mathbf{g} \in \mathbb{Z}_N^w$. The parameters define the noisy randomness space $\mathcal{R}^* = [-B^*, B^*]^m$. It outputs (pk, sk) .

• Enc(pk, m):

Given the public pk , a message $m \in \{0, 1\}$, it samples the randomness $\mathbf{R} \leftarrow_{\mathbb{R}} [-1, 1]^{m \times w}$ and outputs the ciphertext $\text{ct} = \mathbf{U}\mathbf{R} + m\mathbf{G} \in \mathbb{Z}_N^{(\kappa+1) \times w}$. For any $\mathbf{m} \in \{0, 1\}^n$, we denote by $\text{Enc}_{\text{pk}}(\mathbf{m}; \mathbf{r})$ the concatenation of the encryptions $\text{Enc}_{\text{pk}}(m_1; \mathbf{R}_1), \dots, \text{Enc}_{\text{pk}}(m_n; \mathbf{R}_n)$.

- Eval(pk, f, ct₁, ..., ct_ν):

Given the public key \mathbf{pk} , a depth- $\delta(\lambda)$ arithmetic $f : \{0, 1\}^\nu \rightarrow \{0, 1\}$, ciphertexts $\mathbf{ct}_1, \dots, \mathbf{ct}_\nu$, it runs $\mathbf{ct}'_f \leftarrow \text{Eval}'(\mathbf{pk}, f, \omega, \mathbf{ct}_1, \dots, \mathbf{ct}_\nu)$ with scaling factor $\omega = \lceil \log(B) \rceil + 1$, where the algorithm Eval' is described below, for the batch correctness property.

We demonstrate that the GSW FHE satisfies the batch correctness property.

Proposition 1 (Batch correctness). *For all polynomials δ , the GSW_δ scheme described above satisfies batch correctness, as per Definition 3.1.*

Proof: We present the following PPT algorithm:

- Eval'(pk, f, ω, ct₁, ..., ct_ν):

Given the public \mathbf{pk} , a depth- $\delta(\lambda)$ arithmetic circuit $f : \{0, 1\}^\nu \rightarrow \mathbb{Z}_N$, a scaling factor $\omega < \log(N)$, ciphertexts $\mathbf{ct}_1, \dots, \mathbf{ct}_\nu$, it evaluates the circuit gate by gate as follows.

- Addition gate between \mathbf{ct}_i and \mathbf{ct}_j : return $\mathbf{ct}_i + \mathbf{ct}_j \in \mathbb{Z}_N^{(\kappa+1) \times w}$.
- Multiplication gate between \mathbf{ct}_i and \mathbf{ct}_j : return $\mathbf{ct}_i \cdot \text{BD}(\mathbf{ct}_j) \in \mathbb{Z}_N^{(\kappa+1) \times w}$, where $\text{BD}(\mathbf{ct}_j) \in \{0, 1\}^{w \times w}$ denotes the binary decomposition of $\mathbf{ct}_j \in \mathbb{Z}_N^{(\kappa+1) \times w}$.

By recursively applying the above operations, one can turn the ciphertexts $\mathbf{ct}_1, \dots, \mathbf{ct}_\nu$ into $\mathbf{C}_f^i = \mathbf{UR}_f^i + f_i(\mathbf{m})\mathbf{G} \in \mathbb{Z}_N^{(\kappa+1) \times w}$, where $f_i(\mathbf{m}) \in \{0, 1\}$ denotes the i 'th bit of the binary decomposition of $f(\mathbf{m}) \in \mathbb{Z}_N$, that is, $f(\mathbf{m}) = \sum_{i=0}^{\lceil \log(N) \rceil - 1} 2^i f_i(\mathbf{m})$. For all $i \in [0, \lceil \log(N) \rceil - 1]$, we have $\|\mathbf{R}_f^i\|_\infty \leq (w+1)^\delta$. By definition of the matrix \mathbf{G} , choosing the $\kappa \cdot \lceil \log(N) \rceil + i + \omega + 1$ 'th column of \mathbf{C}_f^i yields:

$$\mathbf{c}_f^i = \left(\mathbf{Ar}_f^i, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}_f^i + 2^{\omega+i} f_i(\mathbf{m}) \right) \in \mathbb{Z}_N^{\kappa+1}.$$

Summing up for all $i \in [0, \lceil \log(N) \rceil - 1]$, we get: $\mathbf{ct}'_f = (\mathbf{Ar}_f, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}_f + 2^\omega f(\mathbf{m})) \in \mathbb{Z}_N^{\kappa+1}$, where $\mathbf{r}_f = \sum_{i=0}^{\lceil \log(N) \rceil - 1} \mathbf{r}_f^i$ of norm $\|\mathbf{r}_f\|_\infty \leq (w+1)^\delta \lceil \log(N) \rceil$. It outputs the evaluated ciphertext $\mathbf{ct}_f = \text{BD}(\mathbf{ct}'_f) \in \{0, 1\}^w$.

The evaluated ciphertext $\mathbf{ct}_f \in \{0, 1\}^w$ is such that:

$$\mathbf{sk}^\top \mathbf{ct}_f = -\mathbf{s}^\top \mathbf{Ar}_f + (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}_f + 2^\omega f(\mathbf{m}) = 2^\omega f(\mathbf{m}) + \text{noise}_f \in \mathbb{Z}_N,$$

where $\text{noise}_f = \mathbf{e}^\top \mathbf{r}_f$. We have $|\text{noise}_f| \leq (w+1)^\delta \lceil \log(N) \rceil B_\chi m = B$. □

We turn to proving that it also satisfies the randomness homomorphism property.

Proposition 2 (Randomness homomorphism). *For all polynomials δ , the GSW_δ scheme satisfies the randomness homomorphism property as per Definition 3.2.*

Proof: We present the following PPT algorithms:

- Enc*(pk, m; r*):

Given the public \mathbf{pk} , a message $m \in \mathbb{Z}$, the randomness $\mathbf{r}^* \in [-B^*, B^*]^m$, it computes $\mathbf{ct}' = (\mathbf{Ar}^*, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}^* + m) \in \mathbb{Z}_N^{\kappa+1}$, and outputs $\mathbf{ct} = \text{BD}(\mathbf{ct}') \in \{0, 1\}^w$.

- Eval_{rand}(pk, f, (R_i)_{i∈[ν]}, (m_i)_{i∈[ν]}):

This algorithm is similar to the ciphertext evaluation algorithm. Namely, given the public \mathbf{pk} , a depth- $\delta(\lambda)$ arithmetic circuit $f : \{0, 1\}^\nu \rightarrow \mathbb{Z}_N$, randomness $\mathbf{R}_1, \dots, \mathbf{R}_\nu \in [-1, 1]^{m \times w}$, it evaluates the circuit gate by gate as follows.

- Addition gate between \mathbf{R}_i and \mathbf{R}_j : return $\mathbf{R}_i + \mathbf{R}_j \in \mathbb{Z}_N^{m \times w}$.
- Multiplication gate between \mathbf{R}_i and \mathbf{R}_j : compute $\mathbf{ct}_j = \text{Enc}_{\mathbf{pk}}(m_j; \mathbf{R}_j)$, return $\mathbf{R}_i \text{BD}(\mathbf{ct}_j) + m_i \mathbf{R}_j \in \mathbb{Z}_N^{m \times w}$, where $\text{BD}(\mathbf{ct}_j) \in \{0, 1\}^{w \times w}$ denotes the binary decomposition of $\mathbf{ct}_j \in \mathbb{Z}_N^{(\kappa+1) \times w}$.

By recursively applying the above operations, one can turn the randomness $\mathbf{R}_1, \dots, \mathbf{R}_n$ into $\mathbf{R}_f^i \in \mathbb{Z}_N^{m \times w}$ such that: $\mathbf{C}_f^i = \mathbf{U} \mathbf{R}_f^i + f_i(\mathbf{m}) \mathbf{G} \in \mathbb{Z}_N^{(\kappa+1) \times w}$, for all $i \in [0, \lceil \log(N) \rceil - 1]$; and $\|\mathbf{R}_f^i\|_\infty \leq (w+1)^\delta$. By definition of the matrix \mathbf{G} , choosing the $\kappa \lceil \log(N) \rceil + i + 1$ 'th column of \mathbf{R}_f^i yields $\mathbf{r}_f^i \in \mathbb{Z}_N^m$ such that: $\mathbf{c}_f^i = \left(\mathbf{A} \mathbf{r}_f^i, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}_f^i + 2^i f_i(\mathbf{m}) \right) \in \mathbb{Z}_N^{\kappa+1}$, and $\|\mathbf{r}_f^i\|_\infty \leq (w+1)^\delta$. Summing up for all $i \in [0, \lceil \log(N) \rceil - 1]$, we get: $\mathbf{r}_f = \sum_{i=0}^{\lceil \log(N) \rceil - 1} \mathbf{r}_f^i \in \mathcal{R}^*$ such that $\mathbf{ct}'_f = \left(\mathbf{A} \mathbf{r}_f, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \mathbf{r}_f + f(\mathbf{m}) \right) \in \mathbb{Z}_N^{\kappa+1}$. It outputs the evaluated randomness \mathbf{r}_f . \square

3.2.3 SRL Security

Before proving the SRL security of GSW under the LWE assumption, we describe new trapdoor generation and pre-image sampling algorithms that are inspired by those from [MP12]. As in prior works, the trapdoor generation algorithm generates a matrix $\mathbf{U} \in \mathbb{Z}_N^{d \times m}$ that is statistically close to uniformly random over $\mathbb{Z}_N^{d \times m}$, together with an associated trapdoor $T_{\mathbf{U}}$. The pre-image sampling algorithm, given a target vector $\mathbf{t} \in \mathbb{Z}_N^d$, produces a short pre-image, that is, a short vector $\mathbf{r} \in \mathbb{Z}_N^m$ such that $\mathbf{U} \mathbf{r} = \mathbf{t}$. In these works, the distribution of these short pre-images is independent of the trapdoor — typically they follow a discrete (spherical) Gaussian distribution. Our requirements are slightly different: a pre-image produced by our sampling algorithm when given as input a target vector $\mathbf{t} \in \mathbb{Z}_N^d$ should be statistically close to a pre-image produced by our sampling algorithm when given as input the vector $\mathbf{0} \in \mathbb{Z}_N^d$, shifted by a much smaller pre-image of \mathbf{t} . That is, if a very short pre-image is given, adding a somewhat short pre-image of $\mathbf{0}$ (produced by the sampling algorithm) to it will produce a pre-image that looks like a fresh output of the sampling algorithm on input \mathbf{t} . This inherently requires smudging size noises, which implies the use of an exponential-size modulus q . In fact this property is not known to hold for existing trapdoor generation and pre-image sampling algorithms using polynomial-size modulus.

We prove this property for the concrete algorithms provided in [MP12], which we simplify since we can afford to use smudging-size noises. We provide a self-contained description of the scheme and its proofs here.

Lattice trapdoors.

- TrapGen($1^\lambda, N, d$):

Given as input the security parameter $\lambda \in \mathbb{N}$, a modulus $N \in \mathbb{N}$, a dimension $d \in \mathbb{N}$, it sets $\tilde{m} = d \lceil \log(N) \rceil + 2\lambda$, $w = d \lceil \log(N) \rceil$, $m = \tilde{m} + w$, computes the gadget matrix $\mathbf{G} = \mathbf{g}^\top \otimes \text{Id} = \begin{pmatrix} \mathbf{g}^\top & 0 & \cdots \\ 0 & \mathbf{g}^\top & \\ \vdots & & \ddots \end{pmatrix} \in \mathbb{Z}_N^{d \times w}$ where $\text{Id} \in \mathbb{Z}_N^{d \times d}$ denotes the identity matrix and $\mathbf{g} = (1, 2, \dots, 2^{\lceil \log(N) \rceil - 1}) \in \mathbb{Z}_N^{\lceil \log(N) \rceil}$, $\tilde{\mathbf{U}} \leftarrow_{\mathcal{R}} \mathbb{Z}_N^{d \times \tilde{m}}$, $\mathbf{R} \leftarrow_{\mathcal{R}} [-1, 1]^{\tilde{m} \times w}$, $\mathbf{U} = (\tilde{\mathbf{U}} \parallel -\tilde{\mathbf{U}} \mathbf{R} + \mathbf{G}) \in \mathbb{Z}_N^{d \times m}$, $T_{\mathbf{U}} = \mathbf{R}$. It outputs $(\mathbf{U}, T_{\mathbf{U}})$.

• PrelmSamp($\mathbf{U}, T_{\mathbf{U}}, \mathbf{t}, B$):

Given as input the matrix \mathbf{U} , the trapdoor $T_{\mathbf{U}}$, a target vector $\mathbf{t} \in \mathbb{Z}_N^d$ and a bound $B \in \mathbb{N}$, it samples $\mathbf{v} \leftarrow_{\mathbb{R}} [-B, B]^m$, sets $\mathbf{b} = \text{BD}(\mathbf{U}\mathbf{v} + \mathbf{t}) \in \{0, 1\}^{w \times w}$ which denotes the binary decomposition of $\mathbf{U}\mathbf{v} + \mathbf{t} \in \mathbb{Z}_N^d$. It outputs $\begin{pmatrix} \mathbf{R}\mathbf{b} \\ \mathbf{b} \end{pmatrix} - \mathbf{v} \in \mathbb{Z}_N^m$.

We show the following properties hold.

Proposition 3 (Correctness of TrapGen). *For all λ, N, d , writing $m = 2d\lceil\log(N)\rceil + 2\lambda$, the following distributions have statistical distance at most $2^{-\lambda}$:*

$$\left\{ \mathbf{U} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{d \times m} : \mathbf{U} \right\} \\ \left\{ (\mathbf{U}, T_{\mathbf{U}}) \leftarrow \text{TrapGen}(1^\lambda, N, d) : \mathbf{U} \right\}.$$

Proof: The proposition follows readily from Lemma 2.1 (leftover hash lemma). \square

Proposition 4 (Correctness of PrelmSamp). *For all $\lambda, q, d, B \in \mathbb{N}$, all $(\mathbf{U}, T_{\mathbf{U}})$ in the support of $\text{TrapGen}(1^\lambda, N, d)$, all $\mathbf{t} \in \mathbb{Z}_N^d$, all $\mathbf{r} \in \mathbb{Z}_N^m$ in the support of $\text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{t}, B)$, are such $\mathbf{U}\mathbf{r} = \mathbf{t}$ and $\|\mathbf{r}\|_\infty < B + w$.*

Proof: Straightforward. \square

Proposition 5 (Security). *For all $\lambda, N, d, B \in \mathbb{N}$, writing $m = 2d\lceil\log(N)\rceil + 2\lambda$, for all $\mathbf{w} \in \mathbb{Z}_N^m$ such that $\|\mathbf{w}\|_\infty < B'$, the statistical distance of the two following distributions is upper-bounded by B'/B :*

$$\{(\mathbf{U}, T_{\mathbf{U}}) \leftarrow \text{TrapGen}(1^\lambda, N, d), \tilde{\mathbf{r}}_0 \leftarrow_{\mathbb{R}} \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B) : \tilde{\mathbf{r}}_0 + \mathbf{w} \in \mathbb{Z}_N^m\} \\ \{(\mathbf{U}, T_{\mathbf{U}}) \leftarrow \text{TrapGen}(1^\lambda, N, d), \tilde{\mathbf{r}} \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{U}\mathbf{w}, B) : \tilde{\mathbf{r}}\}$$

Proof: By definition of PrelmSamp we have: $\tilde{\mathbf{r}}_0 = \begin{pmatrix} \mathbf{R}\mathbf{b} \\ \mathbf{b} \end{pmatrix} - \mathbf{v}$ where $\mathbf{v} \leftarrow_{\mathbb{R}} [-B, B]^m$ and $\mathbf{b} = \text{BD}(\mathbf{U}\mathbf{v}) \in \{0, 1\}^w$. For all $\mathbf{w} \in \mathbb{Z}_N^m$ such that $\|\mathbf{w}\|_\infty < B'$, by Lemma 2.2 (smudging), the following distributions have statistical distance B'/B : $\{\mathbf{v} \leftarrow_{\mathbb{R}} [-B, B]^m : \mathbf{v}\}$ and $\{\mathbf{v} \leftarrow_{\mathbb{R}} [-B, B]^m : \mathbf{v} + \mathbf{w}\}$. This implies that $\tilde{\mathbf{r}}_0 + \mathbf{w} \approx \begin{pmatrix} \mathbf{R}\mathbf{b}' \\ \mathbf{b}' \end{pmatrix} - \mathbf{v}$, where $\mathbf{b}' = \text{BD}(\mathbf{U}\mathbf{v} + \mathbf{U}\mathbf{w})$. The latter is identically distributed to $\text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{U}\mathbf{w}, B)$. \square

Theorem 3.5 (SRL security). *Assume the (subexponential) security of the LWE assumption holds. Then, for all polynomials δ , GSW_δ is (subexponentially) SRL secure.*

Proof: For all nuPPT adversaries \mathcal{A} , all $\lambda \in \mathbb{N}$, we use the following hybrid experiments.

• $\mathcal{H}_{\lambda, \mathcal{A}}^0$: is the experiment $\text{Exp}_{\lambda, \mathcal{A}}^{\text{GSW}_\delta}$ from Definition 2.8.

• $\mathcal{H}_{\lambda, \mathcal{A}}^1$: is as $\mathcal{H}_{\lambda, \mathcal{A}}^0$ except the LWE sample $\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top$ from the public key is switched to a uniformly random vector using the LWE assumption. That is, the public key is computed as follows: $\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times m}$, $\mathbf{v} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^m$, $\mathbf{U} = \begin{pmatrix} \mathbf{A} \\ \mathbf{v}^\top \end{pmatrix}$; the gadget matrix \mathbf{G} is computed as in $\mathcal{H}_{\lambda, \mathcal{A}}^0$, and $\text{pk} = (B, \mathbf{U}, \mathbf{G})$. The secret key is also computed as in $\mathcal{H}_{\lambda, \mathcal{A}}^0$ (but now it is uncorrelated with pk), namely: $\mathbf{s} \leftarrow_{\mathbb{R}} \chi^\kappa$, $\text{sk} = (-\mathbf{s}, 1) \otimes \mathbf{g} \in \mathbb{Z}_N^w$. The challenge ciphertext is computed as $\text{ct} = \text{Enc}_{\text{pk}}(\mathbf{m}^*; \mathbf{r})$

and the oracle $\mathcal{O}_{\text{SRL}}(\text{pk}, \mathbf{m}^*, \mathbf{r})$ behaves as in $\mathcal{H}_{\lambda, \mathcal{A}}^0$. The LWE assumption implies that for all nuPPT \mathcal{A} , there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, $|\Pr[\mathcal{H}_{\lambda, \mathcal{A}}^0 = 1] - \Pr[\mathcal{H}_{\lambda, \mathcal{A}}^1 = 1]| \leq \mu(\lambda)$, since these experiments can be efficiently simulated from \mathbf{U} and the winning condition can be efficiently checked.

- $\mathcal{H}_{\lambda, \mathcal{A}}^2$: is as $\mathcal{H}_{\lambda, \mathcal{A}}^1$ except the matrix \mathbf{U} from the public key is sampled from $(\mathbf{U}, T_{\mathbf{U}}) \leftarrow \text{TrapGen}(1^\lambda, N, \kappa + 1)$. By Property 3, this is statistically close (within statistical distance $2^{-\Omega(\lambda)}$) to generating a uniformly random $\mathbf{U} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{(\kappa+1) \times m}$ as done in $\mathcal{H}_{\lambda, \mathcal{A}}^1$. The experiments can be simulated from \mathbf{U} , thus, for all nuPPT \mathcal{A} , we have:

$$\{\mathcal{H}_{\lambda, \mathcal{A}}^1\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^2\}_{\lambda \in \mathbb{N}}.$$

- $\mathcal{H}_{\lambda, \mathcal{A}}^3$: is as $\mathcal{H}_{\lambda, \mathcal{A}}^2$ except we use the oracle $\tilde{\mathcal{O}}_{\text{SRL}}$ instead of \mathcal{O}_{SRL} :

$\tilde{\mathcal{O}}_{\text{SRL}}(\text{pk}, \mathbf{m}^*, \text{ct})$:
 $\mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{R}^*$, $\text{ct}^* = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$
 $(f, \alpha) \leftarrow \mathcal{A}(\text{ct}^*)$
 $\text{BD}(\text{ct}'_f) = \text{Eval}'(\text{pk}, f, 0, \text{ct})$. Parse $\text{ct}'_f = (\mathbf{A}\mathbf{r}_f, \mathbf{v}^\top \mathbf{r}_f + f(\mathbf{m}^*)) \in \mathbb{Z}_N^{\kappa+1}$.
 Compute $\mathbf{t}_f = (\mathbf{A}\mathbf{r}_f, \mathbf{v}^\top \mathbf{r}_f) \in \mathbb{Z}_N^{\kappa+1}$, and $\tilde{\mathbf{r}}_f \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{t}_f, B^* 2^{-\lambda/2})$.
 If $f(\mathbf{m}^*) = \alpha$, and f is of depth δ , then $\text{leak} = \mathbf{r}^* - \tilde{\mathbf{r}}_f \in \mathcal{R}^*$.
 Otherwise, $\text{leak} = \perp$. Return leak .

Note that the oracle $\tilde{\mathcal{O}}_{\text{SRL}}$ only takes as input the public key pk , the message $\mathbf{m}^* \in \{0, 1\}^*$ and the challenge ciphertext $\text{ct} = \text{Enc}_{\text{pk}}(\mathbf{m}^*; \mathbf{r})$, but not the randomness \mathbf{r} itself. Instead of computing the evaluated randomness $\mathbf{r}_f = \text{Eval}_{\text{rand}}(\text{pk}, f, \mathbf{m}^*, \mathbf{r})$, it computes a small $\tilde{\mathbf{r}}_f$ that is consistent with the evaluated ciphertext $\text{ct}'_f = \text{BD}(\text{ct}'_f)$, that is, such that $\text{ct}'_f = (\mathbf{A}\tilde{\mathbf{r}}_f, \mathbf{v}^\top \tilde{\mathbf{r}}_f + f(\mathbf{m}))$. Clearly, the distributions: $(\text{ct}, \mathbf{r}_f)$, which corresponds to $\mathcal{H}_{\lambda, \mathcal{A}}^2$ and $(\text{ct}, \tilde{\mathbf{r}}_f)$, which corresponds to $\mathcal{H}_{\lambda, \mathcal{A}}^3$ are distinct — for one thing, the first distribution has less entropy than the second distribution where $\tilde{\mathbf{r}}_f$ is sampled freshly. However, the value $\tilde{\mathbf{r}}_f$ is shielded by the noisy randomness $\mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{R}^*$. Because it is of much larger magnitude than \mathbf{r}_f and $\tilde{\mathbf{r}}_f$, the latter can smudge the difference $\delta_f = \mathbf{r}_f - \tilde{\mathbf{r}}_f$, which would successfully transition from $\mathcal{H}_{\lambda, \mathcal{A}}^2$ to $\mathcal{H}_{\lambda, \mathcal{A}}^3$. To effectively hide δ_f , we need to make sure $\mathbf{r}^* \in \mathbb{Z}^m$ itself is hidden. Partial information is revealed in $\text{ct}^* = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$, of the form $\mathbf{U}\mathbf{r}^* \in \mathbb{Z}_N^{\kappa+1}$. Intuitively, the component of \mathbf{r}^* along \mathbf{U} is revealed by ct^* , but the remaining entropy of \mathbf{r}^* is hidden; in particular, its component along \mathbf{U}^\perp , the orthogonal space of \mathbf{U} , is hidden. Because $\tilde{\mathbf{r}}_f$ is consistent with ct'_f , we have $\mathbf{U}\delta_f = \mathbf{0}$; that is, δ_f is orthogonal to \mathbf{U} . The orthogonal component of \mathbf{r}^* can simply smudge δ_f . This argument is formalized in Lemma 3.1. Overall, for all nuPPT \mathcal{A} , we have:

$$\{\mathcal{H}_{\lambda, \mathcal{A}}^2\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^3\}_{\lambda \in \mathbb{N}}.$$

To complete the proof of Theorem 3.5, we now show that for all nuPPT \mathcal{A} , all $\lambda \in \mathbb{N}$, we have: $\Pr[\mathcal{H}_{\lambda, \mathcal{A}}^3 = 1] \leq 1/2$. To do so, we consider the event fail (and the complementary event $\overline{\text{fail}}$), which happens when \mathcal{A} chooses a pair of messages $(\mathbf{m}^0, \mathbf{m}^1)$ and makes a query of the form (f^*, α^*) to \mathcal{O}_{SRL} such that $f^*(\mathbf{m}^0) \neq f^*(\mathbf{m}^1)$.

First, we show that $\Pr[\mathcal{H}_{\lambda, \mathcal{A}}^3 = 1 | \text{fail}] \leq 1/2$. This follows from the fact that conditioning on fail , the query (f^*, α^*) makes \mathcal{O}_{SRL} output \perp with probability 1/2 over the choice of $b \leftarrow_{\mathbf{R}} \{0, 1\}$, in which case the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^3$ outputs 0.

Finally, we show that $\Pr[\mathcal{H}_{\lambda, \mathcal{A}}^3 = 1 | \overline{\text{fail}}] \leq 1/2$. This follows from the fact that in the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^3$, the only information revealed about the random bit b is $f(\mathbf{m}^*)$ where $\mathbf{m}^* = \mathbf{m}^b$. Since we

condition on $\overline{\text{fail}}$, we know that $f(\mathbf{m}^0) = f(\mathbf{m}^1)$. Thus, there is no information revealed about b and $\Pr[\mathcal{H}_{\lambda, \mathcal{A}}^3 = 1 | \overline{\text{fail}}] = 1/2$. \square

Lemma 3.1. *For all nuPPT \mathcal{A} , we have: $\{\mathcal{H}_{\lambda, \mathcal{A}}^2\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^3\}_{\lambda \in \mathbb{N}}$.*

Proof: We introduce intermediate hybrids $\mathcal{H}_{\lambda, \mathcal{A}}^{2,i}$ for $i = 1, 2$ which are defined for all nuPPT adversaries \mathcal{A} and all $\lambda \in \mathbb{N}$ as follows.

The experiment $\mathcal{H}_{\lambda, \mathcal{A}}^{2,1}$ is as $\mathcal{H}_{\lambda, \mathcal{A}}^2$ except it uses the following oracle $\mathcal{O}_{\text{SRL}}^1$ instead of \mathcal{O}_{SRL} .

$\mathcal{O}_{\text{SRL}}^1(\text{pk}, \mathbf{m}^*, \mathbf{r})$:

$\mathbf{r}^* \leftarrow_{\mathbb{R}} \mathcal{R}^*$, $\tilde{\mathbf{r}}_0 \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B^* 2^{-\lambda/2})$, $\text{ct}^* = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^* - \tilde{\mathbf{r}}_0)$

$(f, \alpha) \leftarrow \mathcal{A}(\text{ct}^*)$

$\mathbf{r}_f = \text{Eval}_{\text{rand}}(\text{pk}, f, \mathbf{r}, \mathbf{m}^*)$.

If $f(\mathbf{m}^*) = \alpha$ and f is of depth δ , then $\text{leak} = \mathbf{r}^* - \tilde{\mathbf{r}}_0 - \mathbf{r}_f \in \mathcal{R}^*$.

Otherwise, $\text{leak} = \perp$. Return leak .

We first prove that for all nuPPT \mathcal{A} , we have:

$$\{\mathcal{H}_{\lambda, \mathcal{A}}^2\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^{2,1}\}_{\lambda \in \mathbb{N}}.$$

The only difference between these experiments is that $\mathcal{O}_{\text{SRL}}^1$ subtract a pre-image of $\mathbf{0} \in \mathbb{Z}_N^m$ from the shield, that is, it uses $\mathbf{r}^* - \tilde{\mathbf{r}}_0$ with $\mathbf{r}^* \leftarrow_{\mathbb{R}} \mathcal{R}^*$ and $\tilde{\mathbf{r}}_0 \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B^* 2^{-\lambda/2})$ instead of \mathbf{r}^* .

By Property 4, $\tilde{\mathbf{r}}_0 \in \mathbb{Z}_N^m$ is such that $\|\tilde{\mathbf{r}}_0\|_{\infty} < B^* 2^{-\lambda/2}$. Thus, by Lemma 2.2 (smudging), the following distributions have statistical distance at most $2^{-\lambda/2}$:

$$\{\mathbf{r}^* \leftarrow_{\mathbb{R}} [-B^*, B^*]^m : \mathbf{r}^*\} \quad \text{and} \quad \{\mathbf{r}^* \leftarrow_{\mathbb{R}} [-B^*, B^*]^m : \mathbf{r}^* - \tilde{\mathbf{r}}_0\}.$$

The leftmost distribution corresponds to the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^2$ (with post-processing), whereas the rightmost distribution corresponds to the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^{2,1}$ (with the same post-processing). This completes the proof that $\{\mathcal{H}_{\lambda, \mathcal{A}}^2\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^{2,1}\}_{\lambda \in \mathbb{N}}$.

Now, we introduce another intermediate hybrid, $\mathcal{H}_{\lambda, \mathcal{A}}^{2,2}$, which uses the oracle $\mathcal{O}_{\text{SRL}}^2$ instead of $\mathcal{O}_{\text{SRL}}^1$, where $\mathcal{O}_{\text{SRL}}^2$ behaves just as $\mathcal{O}_{\text{SRL}}^1$ with the only exception that ct^* is encrypted using the randomness \mathbf{r}^* (as opposed to randomness $\mathbf{r}^* - \tilde{\mathbf{r}}_0$):

$\mathcal{O}_{\text{SRL}}^2(\text{pk}, \mathbf{m}^*, \mathbf{r})$:

$\mathbf{r}^* \leftarrow_{\mathbb{R}} \mathcal{R}^*$, $\text{ct}^* = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$

$(f, \alpha) \leftarrow \mathcal{A}(\text{ct}^*)$

$\mathbf{r}_f = \text{Eval}_{\text{rand}}(\text{pk}, f, \mathbf{r}, \mathbf{m}^*)$, $\tilde{\mathbf{r}}_0 \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B^* 2^{-\lambda/2})$.

If $f(\mathbf{m}^*) = \alpha$ and f is of depth δ , then $\text{leak} = \mathbf{r}^* - \tilde{\mathbf{r}}_0 - \mathbf{r}_f \in \mathcal{R}^*$.

Otherwise, $\text{leak} = \perp$. Return leak .

By Property 4, we have, $\mathbf{U}\tilde{\mathbf{r}}_0 = \mathbf{0}$. This implies $\text{Enc}_{\text{pk}}^*(0; \mathbf{r}^* - \tilde{\mathbf{r}}_0) = \text{Enc}_{\text{pk}}^*(0; \mathbf{r}^*)$. Thus, for all nuPPT \mathcal{A} , we have:

$$\{\mathcal{H}_{\lambda, \mathcal{A}}^{2,1}\}_{\lambda \in \mathbb{N}} = \{\mathcal{H}_{\lambda, \mathcal{A}}^{2,2}\}_{\lambda \in \mathbb{N}}.$$

To conclude the proof of this lemma, we now prove that for all nuPPT \mathcal{A} , we have:

$$\{\mathcal{H}_{\lambda, \mathcal{A}}^{2,2}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_{\lambda, \mathcal{A}}^3\}_{\lambda \in \mathbb{N}}.$$

To do so, we note that $\mathbf{r}_f \in \mathbb{Z}_N^m$ is such that $\|\mathbf{r}_f\|_\infty < B^*2^{-\lambda}$. Moreover, it is independent of the vector $\tilde{\mathbf{r}}_0 \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B^*2^{-\lambda/2})$. Therefore, we can use Proposition 5, which states that for all vectors $\mathbf{r}_f \in \mathbb{Z}_N^m$ such that $\|\mathbf{r}_f\|_\infty < B^*2^{-\lambda}$, the following distributions have statistical distance at most $2^{-\lambda/2}$:

$$\{\tilde{\mathbf{r}}_0 \leftarrow_{\mathbb{R}} \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{0}, B^*2^{-\lambda/2}) : \tilde{\mathbf{r}}_0 + \mathbf{r}_f \in \mathbb{Z}_N^m\} \text{ and } \{\tilde{\mathbf{r}}_f \leftarrow \text{PrelmSamp}(\mathbf{U}, T_{\mathbf{U}}, \mathbf{U}\mathbf{r}_f, B^*2^{-\lambda/2}) : \tilde{\mathbf{r}}_f\}.$$

The leftmost distribution corresponds to the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^{2,2}$ (with pre and post-processing), whereas the rightmost distribution corresponds to the experiment $\mathcal{H}_{\lambda, \mathcal{A}}^3$ (with the same pre and post-processing). \square

4 Hintable Linearly Homomorphic Encryption

BDGM [BDGM20a] introduced the notion of “hintable” Linearly Homomorphic Encryption (LHE). Roughly speaking, an LHE scheme is said to be hintable if there is a secret-key algorithm that given a ciphertext, produces a “short” decryption hint. The latter can be used to decrypt the ciphertext is was generated from, without the secret key. It is also possible to generate a hint from a ciphertext only knowing the random coins used to produce that ciphertext (but without knowledge of the secret key), and the hints generated in these two ways should be statistically close. For our purposes (and as explained in the introduction), we will need to consider a notion of a hintable LHE satisfying a “weak circuit privacy” notion.

Additionally, we here generalize the notion of a hintable LHE to also consider “packed” LHE, where we can encrypt a vector of messages. Additionally, (just as we did for FHE), we will consider LHE with two encryption modes: a “normal” and an “extra noisy” mode. Linear functions can be evaluated on normal encryptions; furthermore *one* addition with a noisy encryption can be performed. More additions with noisy encryptions would lead to ill-formed ciphertexts that cannot be decrypted properly. (We introduce these extra generalizations to be able to obtain an instantiation based on LWE; these extra generalizations are not needed to capture DJ).

More precisely, we consider the notion of an $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE which enables operating over a plaintext space of length $\ell_2(\lambda)$ vectors over \mathbb{Z}_N for some modulus $|N| \geq \ell_1(\lambda)$, and release hints of size $h(\lambda)$. We will be interested in schemes where either ℓ_1 or ℓ_2 can be made arbitrarily big, while keeping h the same (i.e., the hint will become significantly shorter than a single group element, or it will be significantly smaller than the packing capacity). We consider LHE schemes where decryption only recovers the encrypted message approximately, with some extra small noise. The parameter α quantifies the noise magnitude.

We will present two constructions satisfying the notion of a hintable packed LHE. The first one is the Damgård-Jurik [DJ01] encryption scheme which is proven secure under the DCR assumption: this construction considers the setting where $\ell_2(\lambda) = 1$ (i.e., there is no packing, and instead the group elements are directly much larger than the size of the hint), and $\alpha(\lambda) = 0$, i.e. the decryption recovers the encrypted message perfectly. The second construction will instead be a tweaked version of Packed-Regev [Reg05, PVW08] which is secure under the LWE assumption; in this construction, $\ell_1(\lambda)$ is small (comparable to the hint size), but instead $\ell_2(\lambda)$ can be made arbitrarily large (i.e., we can pack a large number of elements into a ciphertext and still keep the hint size small). The approximation parameter $\alpha(\lambda) = 2^{\lambda+1}$ (note that we will use a modulus N of exponential size in λ for this scheme).

4.1 Definition of Hintable LHE

We proceed to the formal definition.

Definition 4.1 (hintable LHE). For any polynomial $\ell_1, \ell_2, h, \alpha$, an $(\ell_1, \ell_2, h, \alpha)$ -Hintable Packed LHE comprises the following PPT algorithms:

- $\text{CRSgen}(1^\lambda)$: given as input the security parameter $\lambda \in \mathbb{N}$, it outputs crs .
- $\text{Gen}(\text{crs})$: given as input crs , it outputs the tuple $(\text{pk}, \text{sk}, \text{td})$, where td a trapdoor that will be used to compute decryption hints. This tuple defines the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$, where $N \geq 2^{\ell_1(\lambda)}$.
- $\text{Enc}_{\text{pk}}(\mathbf{x})$: given as input the public key pk and a vector in $\mathbf{x} \in \mathbb{Z}_N^\nu$, it outputs a ciphertext ct .
- $\text{Enc}_{\text{pk}}^*(\mathbf{m})$: given as input the public key pk and a message in $\mathbf{m} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, it outputs a noisy ciphertext ct^* .
- $\text{Eval}(\text{pk}, \text{ct}, \text{ct}^*, \mathbf{y})$: given as input the public key pk , ciphertexts ct , a noisy ciphertext ct^* and a function $\mathbf{y} \in [-1, 1]^{\nu \ell_2}$, it outputs an evaluated ciphertext $\text{ct}_{\mathbf{y}}$.
- $\text{Dec}(\text{sk}, \text{ct}^*)$: given as input the secret key and a (noisy or evaluated) ciphertext ct^* , it outputs a plaintext.
- $\text{SecHint}(\text{td}, \text{ct}^*)$: given as input the secret trapdoor td and a (noisy or evaluated) ciphertext ct^* , it outputs a decryption hint ρ .
- $\text{PubHint}(\text{pk}, r)$: given as input the public key and some random coins $r \in \mathcal{R}^*$, where \mathcal{R}^* denotes the randomness space of Enc^* , it outputs a hint ρ .
- $\text{Rec}(\text{pk}, \text{ct}^*, \rho)$: given as input a (noisy or evaluated) ciphertext and a decryption hint ρ , it outputs a plaintext.

These PPT algorithms additionally need to satisfy the properties listed below.

Property 4.1 (α -approximate correctness). For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all tuples $(\text{pk}, \text{sk}, \text{td})$ in the support of $\text{Gen}(\text{crs})$ that define the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$ where $N \geq 2^{\ell_1(\lambda)}$, all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, we have: $\Pr \left[\text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}), \mathbf{m}' = \text{Dec}_{\text{sk}}(\text{ct}^*) : \|\mathbf{m}' - \mathbf{m}\|_\infty < 2^{\alpha(\lambda)} \right] \in 1 - 2^{-\Omega(\lambda)}$.

Property 4.2 (Linear Homomorphism). For all polynomials $\nu(\cdot)$, all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all tuples $(\text{pk}, \text{sk}, \text{td})$ in the support of $\text{Gen}(\text{crs})$ that define the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$, all vectors $\mathbf{x} \in \mathbb{Z}_N^{\nu(\lambda)}$, all ciphertexts ct in the support of $\text{Enc}_{\text{pk}}(\mathbf{x})$, all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, all ciphertexts ct^* in the support of $\text{Enc}_{\text{pk}}^*(\mathbf{m})$, all vectors $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_{\ell_2}) \in \{0, 1\}^{\nu(\lambda)\ell_2}$, we have: $\text{Eval}(\text{pk}, \text{ct}, \text{ct}^*, \mathbf{y})$ deterministically outputs a ciphertext in the support of $\text{Enc}_{\text{pk}}^*(m_1 + \mathbf{x}^\top \mathbf{y}_1, \dots, m_{\ell_2} + \mathbf{x}^\top \mathbf{y}_{\ell_2})$.

Property 4.3 (α -approximate correctness of the secret hints). For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all tuples $(\text{pk}, \text{sk}, \text{td})$ in the support of $\text{Gen}(\text{crs})$ that define the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$, for all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, we have: $\Pr \left[\text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}), \rho \leftarrow \text{SecHint}(\text{sk}, \text{ct}^*), \mathbf{m}' = \text{Rec}(\text{pk}, \text{ct}^*, \rho) : \|\mathbf{m}' - \mathbf{m}\|_\infty \leq 2^{\alpha(\lambda)} \right] \in 1 - 2^{-\Omega(\lambda)}$.

Property 4.4 (Equivalence between public and secret hints). For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \{(\text{pk}, \text{sk}, \text{td}) \leftarrow \text{Gen}(1^\lambda), \text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}), \rho \leftarrow \text{SecHint}(\text{sk}, \text{ct}^*) : (\text{pk}, \text{ct}^*, \rho)\} \\ & \{(\text{pk}, \text{sk}, \text{td}) \leftarrow \text{Gen}(1^\lambda), r \leftarrow_{\mathcal{R}^*} \mathcal{R}^*, \rho \leftarrow \text{PubHint}(\text{pk}, r), \text{ct}^* = \text{Enc}_{\text{pk}}^*(\mathbf{m}; r) : (\text{pk}, \text{ct}^*, \rho)\} \end{aligned}$$

Property 4.5 (*h-succinctness of hints*). For all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all tuples $(\text{pk}, \text{sk}, \text{td})$ in the support of $\text{Gen}(\text{crs})$ that define the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$, all messages $\mathbf{x} \in \mathbb{Z}_N^{\ell_2(\lambda)}$, all ciphertexts ct^* in the support of $\text{Enc}_{\text{pk}}^*(\mathbf{x})$, all hints ρ in the support of $\text{SecHint}(\text{sk}, \text{ct}^*)$ are of size at most $h(\lambda)$.

Property 4.6 (*Weak circuit privacy*). For all polynomials $\nu(\cdot)$, all $\lambda \in \mathbb{N}$, all crs in the support of $\text{CRSgen}(1^\lambda)$, all tuples $(\text{pk}, \text{sk}, \text{td})$ in the support of $\text{Gen}(\text{crs})$ that define the message space $\mathbb{Z}_N^{\ell_2(\lambda)}$, all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, all vectors $\mathbf{x} \in \mathbb{Z}_N^{\nu(\lambda)}$, all vectors $\mathbf{y} \in [-1, 1]^{\nu(\lambda)\ell_2(\lambda)}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \{\text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mathbf{x}), \text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}), \text{ct}_\mathbf{y} = \text{Eval}(\text{pk}, \text{ct}, \text{ct}^*, \mathbf{y}) : (\text{pk}, \text{crs}, \text{ct}, \text{ct}^*, \text{ct}_\mathbf{y})\} \\ & \{\text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mathbf{x}), \text{ct}_\mathbf{y} \leftarrow \text{Enc}_{\text{pk}}^*(\mu), \text{ct}^* = \text{Eval}(\text{pk}, \text{ct}, \text{ct}_\mathbf{y}, -\mathbf{y}) : (\text{crs}, \text{pk}, \text{ct}, \text{ct}^*, \text{ct}_\mathbf{y})\}, \end{aligned}$$

where $\mu = (m_1 + \mathbf{x}^\top \mathbf{y}_1, \dots, m_{\ell_2} + \mathbf{x}^\top \mathbf{y}_{\ell_2}) \in \mathbb{Z}_N^{\ell_2}$.

Property 4.7 (*Density of the noisy ciphertexts*). There exists a polynomial $s(\cdot)$ and a poly-time deterministic function CTsample such that the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \left\{ \text{crs} \leftarrow \text{CRSgen}(1^\lambda), (\text{pk}, \text{sk}, \text{td}) \leftarrow \text{Gen}(\text{crs}), \mathbf{r} \leftarrow_{\mathbb{R}} \{0, 1\}^{s(\lambda)} : \text{CTsample}(\mathbf{r}) \right\}. \\ & \left\{ \text{crs} \leftarrow \text{CRSgen}(1^\lambda), (\text{pk}, \text{sk}, \text{td}) \leftarrow \text{Gen}(\text{crs}), \mathbf{m} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\ell_2(\lambda)}, \text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}) : \text{ct}^* \right\}. \end{aligned}$$

In Appendix A.1, we present the DJ LHE from the DCR assumption, and show that it is a hintable LHE. Now we present a variant of the packed Regev encryption scheme from the LWE assumption.

4.2 Packed-Regev, a Hintable LHE from LWE

We present a packed version of Regev encryption scheme [Reg05], and demonstrate that it can be used to satisfy our notion of a $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE, where ℓ_1 and ℓ_2 can be any polynomial and $h(\lambda)$ is a polynomial that is larger than ℓ_1 but is independent of ℓ_2 . That is, the hint is large in comparison to individual group elements, but we can pack an arbitrary polynomial number of elements and still use the same hint size. We have $\alpha(\lambda) = \lambda + 1$.

As explained in the introduction, our construction is slightly different, but similar in spirit, to the Packed-Regev from [PVW08]. Just as in [PVW08], the idea is to individually encrypt the ℓ_2 different components of the message vector but reusing the *same* randomness \mathbf{r} (but different parts of the secret key) for the components. In contrast to [PVW08], as we do not want the length of the randomness to grow with ℓ_2 , to prove security of the scheme, we add an extra smudging noise term \mathbf{e}' to each encrypted component.

The hint for an encrypted message is a short pre-image of the ciphertext hear \mathbf{Ar} , where \mathbf{r} is the randomness of the encryption. To enable efficiently recovering this hint, we will generate the lattice given in the public key together with a standard lattice trapdoor that enables sampling random short pre-images as in [Ajt96, GPV08, AP09, MP12].

To satisfy the properties of density and weak circuit privacy, we will rely on a extra noisy Packed-Regev encryption which proceeds just like the normal one but uses a much larger amount of randomness (so that it covers the whole set of strings for density, and so that it smudges the noises of evaluated ciphertexts for weak circuit privacy).

We will show how the Packed-Regev encryption scheme satisfies our notion of a hintable packed LHE.

Theorem 4.2. *Assume (subexponential) security of the LWE assumption holds. Then, for all polynomials ℓ_1 , there exists some polynomial h such that for all polynomials ℓ_2 , there exists a (subexponentially) secure $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE, with $\alpha(\lambda) = \lambda + 1$.*

4.2.1 Trapdoor Sampling

The Packed-Regev scheme makes use of the following lattice trapdoor mechanism: prior works [Ajt96, GPV08, AP09, MP12] show that there exist PPT algorithms `TrapGen` and `PrelmSamp`, an ensemble $\{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ of efficiently sampleable distributions over \mathbb{Z} such that the following holds.

- `TrapGen` $(1^\lambda, N, d)$:

Given as input the security parameter $\lambda \in \mathbb{N}$, a modulus $N \in \mathbb{N}$, a dimension $d \in \mathbb{N}$, it outputs $(\mathbf{A}, T_{\mathbf{A}})$, where $\mathbf{A} \in \mathbb{Z}_N^{d \times m}$, $m \in \Theta(d \log(N))$ and $T_{\mathbf{A}}$ is a trapdoor. The matrix \mathbf{A} is statistically close to uniform, that is, for all $\lambda, N, d \in \mathbb{N}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$: $\{(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, d) : \mathbf{A}\}$ and $\{\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times m} : \mathbf{A}\}$.

- `PrelmSamp` $(\mathbf{A}, T_{\mathbf{A}}, \mathbf{t})$:

Given as input the matrix \mathbf{A} , the trapdoor $T_{\mathbf{A}}$, a target vector $\mathbf{t} \in \mathbb{Z}_N^d$, it outputs $\mathbf{r} \in \mathbb{Z}_N^m$. For all $\lambda, d, N \in \mathbb{N}$, all $(\mathbf{A}, T_{\mathbf{A}})$ in the support of `TrapGen` $(1^\lambda, N, d)$, all $\mathbf{t} \in \mathbb{Z}_N^d$, `PrelmSamp` $(\mathbf{A}, T_{\mathbf{A}}, \mathbf{t})$ outputs $\mathbf{r} \in \mathbb{Z}_N^m$ such that $\mathbf{A}\mathbf{r} = \mathbf{t} \in \mathbb{Z}_N^d$ and $\|\mathbf{r}\|_\infty < 2^{\lambda/2}$ with probability $1 - 2^{-\Omega(\lambda)}$ over its random coin.

We require the output \mathbf{r} to follow some distribution that does not depend on the actual trapdoor $T_{\mathbf{A}}$, namely, for all $\lambda, d, N \in \mathbb{N}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \{(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, d), \mathbf{r} \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda^m, \mathbf{r}' \leftarrow_{\mathbb{R}} \text{PrelmSamp}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{A}\mathbf{r}) : (\mathbf{r}', \mathbf{A}\mathbf{r})\} \\ & \{(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, d), \mathbf{r} \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda^m : (\mathbf{r}, \mathbf{A}\mathbf{r})\}. \end{aligned}$$

For our purposes, we want the distributions \mathcal{D}_λ to be of smudging size. That is, for all polynomials $p(\cdot)$, all $\lambda, q \in \mathbb{N}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \{r \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda : r \in \mathbb{Z}_N\} \\ & \{r \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda : r + p(\lambda) \in \mathbb{Z}_N\}. \end{aligned}$$

Finally, by the leftover hash lemma, since m is large enough and \mathcal{D}_λ has enough entropy, for all $\lambda, N, q \in \mathbb{N}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$: $\{\mathbf{r} \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda^m, (\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, d) : (\mathbf{A}, \mathbf{A}\mathbf{r})\}$ and $\{\mathbf{r} \leftarrow_{\mathbb{R}} \mathcal{D}_\lambda^m, (\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, d), \mathbf{u} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^d : (\mathbf{A}, \mathbf{u})\}$.

4.2.2 The Construction

We now proceed to describing the Packed-Regev scheme, which is parameterized by polynomials ℓ_1 and ℓ_2 . We denote the scheme by $\text{P-Regev}_{\ell_1, \ell_2}$.

- `CRSgen` (1^λ) :

It simply outputs $\text{crs} = 1^\lambda$, i.e. there is no proper crs for that scheme.

- `Gen` (crs) :

Given as input $\text{crs} = 1^\lambda$, it chooses $q = \{q_\kappa\}_{\kappa \in \mathbb{N}}$ with $q_\kappa = 2^{\kappa^c}$, a B_χ -bounded ensemble $\chi = \{\chi_\kappa\}_{\kappa \in \mathbb{N}}$

of efficiently sampleable distributions over \mathbb{Z} with $B_\chi(\kappa) = \kappa$, and a polynomial $\kappa(\lambda) = \ell_1(\lambda)^{1/c}$, where $c \in (0, 1)$ is the constant from Definition 3.4, and $N = q_{\kappa(\lambda)}$. This choice of parameters ensures that the LWE assumption implies LWE holds for q, χ (i.e. the sequence q and the polynomial B_χ satisfy the requirement from Definition 3.4), and $N = 2^{\ell_1(\lambda)}$. We abuse notations and write $\kappa = \kappa(\lambda)$, $\chi = \chi_{\kappa(\lambda)}$ and $B_\chi = B_\chi(\kappa(\lambda))$ from here on.

Then, the algorithm samples $(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^\lambda, N, \kappa)$, $\mathbf{S} \leftarrow \chi^{\ell_2 \times \kappa}$, $\mathbf{E} \leftarrow \chi^{\ell_2 \times m}$, and sets $\text{pk} = (N, \mathbf{A}, \mathbf{SA} + \mathbf{E}) \in \mathbb{N} \times \mathbb{Z}_N^{\kappa \times m} \times \mathbb{Z}_N^{\ell_2 \times m}$, $\text{sk} = \mathbf{S}$ and $\text{td} = T_{\mathbf{A}}$. It outputs $(\text{pk}, \text{sk}, \text{td})$.

• Enc_{pk}($\mathbf{x} \in \mathbb{Z}_N^\nu$):

Given the public pk , a vector $\mathbf{x} \in \mathbb{Z}_N^\nu$, it samples $\mathbf{R} \leftarrow [-1, 1]^{m \times \nu \ell_2}$, $\mathbf{E}' \leftarrow_{\mathbf{R}} [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \nu \ell_2}$ and outputs the ciphertext $\text{ct} = (\mathbf{AR}, (\mathbf{SA} + \mathbf{E})\mathbf{R} + \mathbf{E}' + \mathbf{x}^\top \otimes \text{Id}_{\ell_2}) \in \mathbb{Z}_N^{(\kappa + \ell_2) \times \nu \ell_2}$, where $\text{Id}_{\ell_2} \in \mathbb{Z}_N^{\ell_2 \times \ell_2}$

denotes the identity matrix, and $\mathbf{x}^\top \otimes \text{Id}_{\ell_2} = \begin{pmatrix} \mathbf{x}^\top & 0 & \cdots \\ 0 & \mathbf{x}^\top & \\ \vdots & & \ddots \end{pmatrix} \in \mathbb{Z}_N^{\ell_2 \times \nu \ell_2}$.

• Enc_{pk}^{*}($\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$):

Given the public pk , a message $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, it samples $\mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\lambda^m$, where \mathcal{D}_λ is the efficiently sampleable distribution over \mathbb{Z} related to TrapGen and PrelmSamp; $\mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}$, and outputs the noisy ciphertext $\text{ct}^* = (\mathbf{Ar}^*, (\mathbf{SA} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}) \in \mathbb{Z}_N^{\kappa + \ell_2}$.

• Eval($\text{pk}, \text{ct}^*, \mathbf{y}$):

Given the public pk , ciphertext $\text{ct} \in \mathbb{Z}_N^{(\kappa + \ell_2) \times \nu \ell_2}$, noisy ciphertext $\text{ct}^* \in \mathbb{Z}_N^{\kappa + \ell_2}$ and a vector $\mathbf{y} \in [-1, 1]^{\nu \ell_2}$, it outputs the evaluated ciphertext $\text{ct} \cdot \mathbf{y} + \text{ct}^* \in \mathbb{Z}_N^{\kappa + \ell_2}$.

• SecHint(td, ct^*):

Given as input the secret key sk and a (noisy or evaluated) ciphertext $\text{ct}^* \in \mathbb{Z}_N^{\kappa + \ell_2}$ of the form (\mathbf{t}, \mathbf{z}) where $\mathbf{t} \in \mathbb{Z}_N^\kappa$ and $\mathbf{z} \in \mathbb{Z}_N^{\ell_2}$, it samples $\rho \leftarrow \text{PrelmSamp}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{t})$ and outputs the hint $\rho \in \mathbb{Z}_N^m$.

• PubHint(pk, r):

Given as input the public key pk and the random coins $r = (\mathbf{r}^*, \mathbf{e}^*)$ where $\mathbf{r}^* \in \mathcal{D}_\lambda^m$, $\mathbf{e}^* \in [-2^\lambda, 2^\lambda]^{\ell_2}$ used to produce a noisy ciphertext, it outputs $\mathbf{r}^* \in \mathbb{Z}_N^m$.

• Rec($\text{pk}, \text{ct}^*, \rho$):

Given as input the public key pk , a (noisy or evaluated) ciphertext $\text{ct}^* \in \mathbb{Z}_N^{\kappa + \ell_2}$ of the form $\text{ct} = (\mathbf{t}, \mathbf{z})$ where $\mathbf{t} \in \mathbb{Z}_N^\kappa$, $\mathbf{z} \in \mathbb{Z}_N^{\ell_2}$ and a hint $\rho \in \mathbb{Z}_N^m$, it outputs $\mathbf{d} = \mathbf{z} - (\mathbf{SA} + \mathbf{E})\rho \in \mathbb{Z}_N^{\ell_2}$.

• Dec_{sk}(ct^*):

Given as input the secret key sk and a (noisy or evaluated) ciphertext $\text{ct}^* = (\mathbf{t}, \mathbf{z})$ with $\mathbf{t} \in \mathbb{Z}_N^\kappa$, $\mathbf{z} \in \mathbb{Z}_N^{\ell_2}$, it outputs $\mathbf{d} = \mathbf{z} - \mathbf{St} \in \mathbb{Z}_N^{\ell_2}$.

The proof of Theorem 4.2 follows from the propositions and theorem below (which demonstrate that Packed-Regev Scheme satisfies the desired properties of a hintable packed LHE, as well as security).

Proposition 6 ($\lambda+1$ -approximate correctness). *The LHE presented above satisfies $\lambda+1$ -approximate correctness, as defined in Property 4.1.*

Proof: A ciphertext ct^* in the support of $\text{Enc}_{\text{pk}}^*(\mathbf{m})$ is of the form $\text{ct}^* = (\mathbf{t}, \mathbf{z})$ with $\mathbf{t} = \mathbf{A}\mathbf{r}^* \in \mathbb{Z}_q^\kappa$ and $\mathbf{z} = (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{m} \in \mathbb{Z}_N^{\ell_2}$. The vector $\mathbf{d} \in \mathbb{Z}_N^{\ell_2}$ output by the decryption is of the form $\mathbf{m} + \text{noise}$ where $\text{noise} = \mathbf{E}\mathbf{r}^* + \mathbf{e}^*$. For all $\mathbf{E} \in \mathbb{Z}_N^{\ell_2 \times m}$ such that $\|\mathbf{E}\|_\infty \leq B_\chi$, with probability $1 - 2^{-\Omega(\lambda)}$ over the choices of $\mathbf{e}^* \leftarrow [-2^\lambda, 2^\lambda]^{\ell_2}$ and $\mathbf{r}^* \leftarrow \mathcal{D}_\chi^m$, we have $\|\text{noise}\|_\infty \leq 2^\lambda + B_\chi m 2^{\lambda/2} \leq 2^{\lambda+1}$. \square

Proposition 7 (Linear Homomorphism). *The LHE presented above satisfies Property 4.2.*

Proof: The ciphertext produced by $\text{Enc}_{\text{pk}}(\mathbf{x})$ has the form $\text{ct} = (\mathbf{A}\mathbf{R}, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{R} + \mathbf{E}' + \mathbf{x}^\top \otimes \text{Id}_{\ell_2}) \in \mathbb{Z}_N^{(\kappa+\ell_2) \times \nu\ell_2}$, and $\text{ct}^* = (\mathbf{A}\mathbf{r}^*, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{x}^*)$. For all $\mathbf{y} \in [-1, 1]^{\nu\ell_2}$, $\text{Eval}(\text{pk}, \text{ct}, \text{ct}^*, \mathbf{y})$ outputs the evaluated ciphertext $\text{ct}_\mathbf{y} = (\mathbf{A}(\mathbf{R}\mathbf{y} + \mathbf{r}^*), (\mathbf{S}\mathbf{A} + \mathbf{E})(\mathbf{R}\mathbf{y} + \mathbf{r}^*) + \mathbf{E}'\mathbf{y} + \mathbf{e}^* + \mu) \in \mathbb{Z}_N^{\kappa+\ell_2}$ where $\mu = (m_1 + \mathbf{x}^\top \mathbf{y}_1, \dots, m_{\ell_2} + \mathbf{x}^\top \mathbf{y}_{\ell_2}) \in \mathbb{Z}_N^{\ell_2}$ which is in the support of $\text{Enc}_{\text{pk}}^*(\mu)$. \square

Proposition 8 ($\lambda+1$ -approximate correctness of the secret hints). *The LHE presented above satisfies $\lambda+1$ -approximate correctness of the secret hints, as defined in Property 4.3.*

Proof: For all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, $\text{Enc}_{\text{pk}}^*(\mathbf{m})$ is of the form $\text{ct}^* = (\mathbf{t}, \mathbf{z}) \in \mathbb{Z}_N^{\kappa \times \ell_2}$, where $\mathbf{t} = \mathbf{A}\mathbf{r}^*$ and $\mathbf{z} = (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{x}^*$. The algorithm SecHint computes $\rho \leftarrow \text{PrelmSamp}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{t})$, which is such that $\mathbf{A}\rho = \mathbf{t}$. Next, the algorithm Rec computes $\mathbf{d} = \mathbf{z} - (\mathbf{S}\mathbf{A} + \mathbf{E})\rho = \mathbf{m} + \text{noise}$, where $\text{noise} = \mathbf{e}^* + \mathbf{E}(\mathbf{r}^* - \rho)$. For all $\mathbf{E} \in \mathbb{Z}_N^{\ell_2 \times m}$ such that $\|\mathbf{E}\|_\infty \leq B_\chi$, with probability $1 - 2^{-\Omega(\lambda)}$ over the choices of $\mathbf{e}^* \leftarrow [-2^\lambda, 2^\lambda]^{\ell_2}$, $\mathbf{r}^* \leftarrow \mathcal{D}_\chi^m$ and the random coins used to produce ρ , we have $\|\text{noise}\|_\infty \leq 2^\lambda + 2mB_\chi 2^{\lambda/2} \leq 2^{\lambda+1}$. \square

Proposition 9 (Equivalence between public and secret hints). *The LHE presented above satisfies Property 4.4.*

Proof: We aim at proving that for all $\lambda \in \mathbb{N}$, all messages $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\mathcal{D}_0 = \left\{ \begin{array}{l} (\text{pk} = (\mathbf{A}, \mathbf{S}\mathbf{A} + \mathbf{E}), \text{sk} = \mathbf{S}, \text{td} = T_{\mathbf{A}}) \leftarrow \text{Gen}(1^\lambda), \mathbf{r}^* \leftarrow \mathcal{D}_\chi^m, \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2} \\ \rho \leftarrow \text{PrelmSamp}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{A}\mathbf{r}^*) : (\mathbf{A}\mathbf{r}^*, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}, \rho) \end{array} \right\}$$

$$\mathcal{D}_1 = \left\{ \begin{array}{l} (\text{pk} = (\mathbf{A}, \mathbf{S}\mathbf{A} + \mathbf{E}), \text{sk} = \mathbf{S}, \text{td} = T_{\mathbf{A}}) \leftarrow \text{Gen}(1^\lambda), \mathbf{r}^* \leftarrow \mathcal{D}_\chi^m \\ \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2} : (\mathbf{A}\mathbf{r}^*, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}, \mathbf{r}^*) \end{array} \right\}$$

By Lemma 2.2 (smudging), distribution \mathcal{D}_0 has statistical distance at most $2^{-\Omega(\lambda)}$ with $\mathcal{D}'_0 = \{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \mathbf{r}^* \leftarrow \mathcal{D}_\chi^m, \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}, \rho \leftarrow \text{PrelmSamp}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{A}\mathbf{r}^*) : (\mathbf{A}\mathbf{r}^*, \mathbf{S}\mathbf{A}\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}, \rho)\}$. By the property of PrelmSamp , \mathcal{D}'_0 has statistical distance at most $2^{-\Omega(\lambda)}$ with $\mathcal{D}'_1 = \{(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \mathbf{r}^* \leftarrow \mathcal{D}_\chi^m, \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2} : (\mathbf{A}\mathbf{r}^*, \mathbf{S}\mathbf{A}\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}, \mathbf{r}^*)\}$. Finally, using smudging again, we have $\mathcal{D}'_1 \approx_s \mathcal{D}_1$. \square

Proposition 10 (Succinctness of hints). *The LHE presented above satisfies $h(\lambda) = (\lambda/2 + 1)m$ succinctness, where $m = 2\kappa \log(q) + 2\lambda$.*

Proof: With probability $1 - 2^{-\Omega(\lambda)}$ over the random coins of SecHint , the hint output by SecHint belong to $[-2^{\lambda/2}, 2^{\lambda/2}]^m$. \square

Proposition 11 (Density of the noisy ciphertexts). *The LHE presented above satisfies Property 4.7.*

Proof: For all $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, a noisy ciphertext $\text{ct}^* \leftarrow_{\mathbf{R}} \text{Enc}_{\text{pk}}^*(\mathbf{m})$ is of the form $\text{ct}^* = (\mathbf{A}\mathbf{r}^*, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{m}) \in \mathbb{Z}_N^{\kappa+\ell_2}$. When $\mathbf{m} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{\ell_2}$, the second part of the ciphertext is uniformly random over $\mathbb{Z}_N^{\ell_2}$. That is, for all $\lambda \in \mathbb{N}$, for all pairs (pk, sk) in the support of $\text{Gen}(1^\lambda)$, the following distributions are identical: $\{\mathbf{m} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{\ell_2}, \text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}) : \text{ct}^*\}$ and $\{\mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\chi^m, \mathbf{w} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{\ell_2} : (\mathbf{A}\mathbf{r}^*, \mathbf{w})\}$. We conclude the proof using the properties of PrelmSamp , which imply that the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$: $\{\mathbf{A} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{\kappa \times m}, \mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\chi^m : (\mathbf{A}, \mathbf{A}\mathbf{r}^*)\}$ and $\{\mathbf{A} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^{\kappa \times m}, \mathbf{u} \leftarrow_{\mathbf{R}} \mathbb{Z}_N^\kappa : (\mathbf{A}, \mathbf{u})\}$. \square

Proposition 12 (Weak circuit privacy). *The LHE presented above satisfies Property 4.6.*

Proof: For all $\mathbf{x} \in \mathbb{Z}_N^\nu$, $\mathbf{y} \in [-1, 1]^{\nu\ell_2}$, we aim at proving the following distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\begin{aligned} \mathcal{D}_0 &= \{ (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mathbf{x}), \text{ct}^* \leftarrow \text{Enc}_{\text{pk}}^*(\mathbf{m}), \text{ct}_{\mathbf{y}} = \text{Eval}(\text{pk}, \text{ct}, \text{ct}^*, \mathbf{y}) : (\text{ct}, \text{ct}^*, \text{ct}_{\mathbf{y}}) \} \\ \mathcal{D}_1 &= \left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda), \text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mathbf{x}) \\ \text{ct}_{\mathbf{y}} \leftarrow \text{Enc}_{\text{pk}}^*(\mu), \text{ct}^* = \text{Eval}(\text{pk}, \text{ct}, \text{ct}_{\mathbf{y}}, -\mathbf{y}) : (\text{ct}, \text{ct}^*, \text{ct}_{\mathbf{y}}) \end{array} \right\}, \end{aligned}$$

where $\mu = (m_1 + \mathbf{x}^\top \mathbf{y}_1, \dots, m_{\ell_2} + \mathbf{x}^\top \mathbf{y}_{\ell_2}) \in \mathbb{Z}_N^{\ell_2}$.

In distribution \mathcal{D}_0 , we have:

- $\text{ct} = (\mathbf{A}\mathbf{R}, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{R} + \mathbf{E}' + \mathbf{x}^\top \otimes \text{Id}_{\ell_2}),$
- $\text{ct}^* = (\mathbf{A}\mathbf{r}^*, (\mathbf{S}\mathbf{A} + \mathbf{E})\mathbf{r}^* + \mathbf{e}^* + \mathbf{x}^*),$
- $\text{ct}_{\mathbf{y}} = (\mathbf{A}(\mathbf{R}\mathbf{y} + \mathbf{r}^*), (\mathbf{S}\mathbf{A} + \mathbf{E})(\mathbf{R}\mathbf{y} + \mathbf{r}^*) + \mathbf{E}'\mathbf{y} + \mathbf{e}^* + \mu).$

We show that it has statistical distance $2^{-\Omega(\lambda)}$ from \mathcal{D}_1 by a series of claims.

Claim 1. For all $\mathbf{y} \in [-1, 1]^{\nu\ell_2}$, the following distributions have statistical distance $2^{-\Omega(\lambda)}$:

$$\begin{aligned} &\left\{ \mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\lambda^m, \mathbf{R} \leftarrow_{\mathbf{R}} [-1, 1]^{m \times \nu\ell_2} : (\mathbf{R}, \mathbf{r}^*, \mathbf{r}^* + \mathbf{R}\mathbf{y}) \right\} \\ &\left\{ \mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\lambda^m, \mathbf{R} \leftarrow_{\mathbf{R}} [-1, 1]^{m \times \nu\ell_2} : (\mathbf{R}, \mathbf{r}^* - \mathbf{R}\mathbf{y}, \mathbf{r}^*) \right\}, \end{aligned}$$

by the properties of `PrelmSamp`.

Claim 2. For all $\mathbf{y} \in [-1, 1]^{\nu\ell_2}$, the following distributions have statistical distance $2^{-\Omega(\lambda)}$:

$$\begin{aligned} &\left\{ \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}, \mathbf{E}' \leftarrow [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \nu\ell_2} : (\mathbf{E}', \mathbf{e}^*, \mathbf{e}^* + \mathbf{E}'\mathbf{y}) \right\} \\ &\left\{ \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}, \mathbf{E}' \leftarrow [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \nu\ell_2} : (\mathbf{E}', \mathbf{e}^* - \mathbf{E}'\mathbf{y}, \mathbf{e}^*) \right\}. \end{aligned}$$

by Lemma 2.2 (smudging).

Claim 1 and Claim 2 imply the following claim.

Claim 3. For all $\mathbf{y} \in [-1, 1]^{\nu\ell_2}$, the following distributions have statistical distance $2^{-\Omega(\lambda)}$:

$$\begin{aligned} \mathcal{D}'_0 &= \left\{ \begin{array}{l} \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}, \mathbf{E}' \leftarrow [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \nu\ell_2}, \mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\lambda^m \\ \mathbf{R} \leftarrow_{\mathbf{R}} [-1, 1]^{m \times \nu\ell_2} : (\mathbf{R}, \mathbf{r}^*, \mathbf{r}^* + \mathbf{R}\mathbf{y}, \mathbf{E}', \mathbf{e}^*, \mathbf{e}^* + \mathbf{E}'\mathbf{y}) \end{array} \right\} \\ \mathcal{D}'_1 &= \left\{ \begin{array}{l} \mathbf{e}^* \leftarrow_{\mathbf{R}} [-2^\lambda, 2^\lambda]^{\ell_2}, \mathbf{E}' \leftarrow [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \nu\ell_2}, \mathbf{r}^* \leftarrow_{\mathbf{R}} \mathcal{D}_\lambda^m \\ \mathbf{R} \leftarrow_{\mathbf{R}} [-1, 1]^{m \times \nu\ell_2} : (\mathbf{R}, \mathbf{r}^* - \mathbf{R}\mathbf{y}, \mathbf{r}^*, \mathbf{E}', \mathbf{e}^* - \mathbf{E}'\mathbf{y}, \mathbf{e}^*) \end{array} \right\}. \end{aligned}$$

Now we prove the following claim, that will conclude the proof of Proposition 12.

Claim 4. There exists a (possibly inefficient) simulator \mathcal{S} that given as input $\mathbf{v} \in (\mathbb{Z}_N^{m \times (\nu \ell_2 + 2)} \times \mathbb{Z}_N^{\ell_2 \times (\nu \ell_2 + 2)})$, outputs a tuple $(\text{ct}, \text{ct}^*, \text{ct}_y) \in \mathbb{Z}_N^{(\kappa + \ell_2) \times \nu \ell_2} \times (\mathbb{Z}_N^{\kappa + \ell_2})^2$, such that when \mathcal{S} is fed with an input from distribution \mathcal{D}'_0 , it produces an output following the distribution \mathcal{D}_0 , whereas when fed with an input coming from distribution \mathcal{D}'_1 , it produces an output following the distribution \mathcal{D}_1 .

Given as input $(\mathbf{R}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{E}', \mathbf{e}_1, \mathbf{e}_2)$, \mathcal{S} computes:

- $\text{ct} = (\mathbf{AR}, (\mathbf{SA} + \mathbf{E})\mathbf{R} + \mathbf{E}' + \mathbf{x}^\top \otimes \text{Id}_{\ell_2})$,
- $\text{ct}^* = (\mathbf{Ar}_1, (\mathbf{SA} + \mathbf{E})\mathbf{r}_1 + \mathbf{e}_1 + \mathbf{m})$,
- $\text{ct}_y = (\mathbf{Ar}_2, (\mathbf{SA} + \mathbf{E})\mathbf{r}_2 + \mathbf{e}_2 + \mu)$.

□

Theorem 4.3 (Security). *Assume (subexponential) security of the LWE assumption holds. Then, for all polynomials ℓ_1 and ℓ_2 , P-Regev $_{\ell_1, \ell_2}$ is (subexponentially) secure.*

Proof: In a nutshell, Regev encryption uses the LWE assumption to switch the LWE samples from the public key $\mathbf{SA} + \mathbf{E}$ to uniformly random, then use the leftover hash lemma to extract the entropy from the randomness used to encrypt the messages. The problem here is that for succinctness, we use small randomness, and large messages: the randomness does not hold enough entropy to hide the messages. Instead, we use the randomness and extra smudging noise to generate fresh LWE samples, and then hide the messages. We now provide the formal proof, which uses the hybrids experiment listed below.

- \mathcal{H}_λ^0 : as per Definition 2.8. Namely, the experiment generates $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\lambda)$ where $\text{pk} = (N, \mathbf{A}, \mathbf{SA} + \mathbf{E})$, sends pk to the adversary who chooses a pair $(\mathbf{x}^0, \mathbf{x}^1) \in \mathbb{Z}_N^{\nu_0} \times \mathbb{Z}_N^{\nu_1}$ (wlog. we can assume $\nu_0 = \nu_1 = \nu$). The experiment samples $b \leftarrow_{\mathbb{R}} \{0, 1\}$ and computes the challenge ciphertext as $\text{ct} = (\mathbf{AR}, (\mathbf{SA} + \mathbf{E})\mathbf{R} + \mathbf{E}' + \mathbf{x}^b \otimes \text{Id}_{\ell_2})$, which is sent to the adversary, who wins if it guesses b successfully.

- \mathcal{H}_λ^1 : this experiment is the same as \mathcal{H}_λ^0 except the challenge ciphertext is now of the form $\text{ct} = (\mathbf{AR}, \mathbf{SAR} + \mathbf{E}' + \mathbf{x}^b \otimes \text{Id}_{\ell_2})$. Recall that $\mathbf{E} \leftarrow_{\mathbb{R}} \chi^{\ell_2 \times m}$ where χ is B_χ -bounded, for a polynomial B_χ , $\mathbf{R} \leftarrow_{\mathbb{R}} [-1, 1]^{m \times \ell_2}$ for a polynomial m and $\mathbf{E}' \leftarrow_{\mathbb{R}} [-2^{\lambda/2}, 2^{\lambda/2}]^{\ell_2 \times \ell_2}$. Thus, we can use Lemma 2.2 (smudging) to argue that $\mathbf{E}' + \mathbf{ER} \approx_s \mathbf{E}'$ with statistical distance $2^{-\Omega(\lambda)}$. The first distribution corresponds to \mathcal{H}_λ^0 (with pre and post-processing), whereas the second distribution corresponds to \mathcal{H}_λ^1 (with pre and post-processing).

- \mathcal{H}_λ^2 : this experiment is the same as \mathcal{H}_λ^1 , except the challenge ciphertext is now of the form $\text{ct} = (\mathbf{U}, \mathbf{SU} + \mathbf{E}' + \mathbf{x}^b \otimes \text{Id}_{\ell_2})$, where $\mathbf{U} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times \ell_2}$. The fact that $\{\mathcal{H}_\lambda^1\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^2\}_{\lambda \in \mathbb{N}}$ follows readily from the leftover hash lemma, which states that the following distributions have statistical distance $2^{-\Omega(\lambda)}$:

$$\begin{aligned} & \{\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times m}, \mathbf{R} \leftarrow_{\mathbb{R}} [-1, 1]^{m \times \ell_2} : (\mathbf{A}, \mathbf{AR})\} \\ & \{\mathbf{A} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times m}, \mathbf{U} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times \ell_2} : (\mathbf{A}, \mathbf{U})\}. \end{aligned}$$

The first distribution corresponds to \mathcal{H}_λ^1 (with post-processing), whereas the second distribution corresponds to \mathcal{H}_λ^2 (with the same post-processing).

- \mathcal{H}_λ^3 : this experiment is the same as \mathcal{H}_λ^2 , except the challenge ciphertext is now of the form

$\text{ct} = (\mathbf{U}, \mathbf{W})$, where $\mathbf{W} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times \nu \ell_2}$, and the public key is now of the form $\text{pk} = (N, \mathbf{A}, \mathbf{V})$ where $\mathbf{V} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\ell_2 \times m}$. We show that $\{\mathcal{H}_\lambda^2\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{H}_\lambda^3\}_{\lambda \in \mathbb{N}}$. This holds by the properties of `TrapGen`, which state that the matrix \mathbf{A} is statistically close to uniform over $\mathbb{Z}_N^{\kappa \times m}$. Then, we rely on the LWE assumption, which implies that $(\mathbf{A}, \mathbf{S}\mathbf{A} + \mathbf{E}, \mathbf{U}, \mathbf{S}\mathbf{U} + \mathbf{E}' + \mathbf{x}^b \otimes \text{Id}_{\ell_2}) \approx_c (\mathbf{A}, \mathbf{V}, \mathbf{U}, \mathbf{W} + \mathbf{x}^b \otimes \text{Id}_{\ell_2}) \equiv (\mathbf{A}, \mathbf{V}, \mathbf{U}, \mathbf{W})$, where $\mathbf{W} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\kappa \times \nu \ell_2}$ and $\mathbf{V} \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\ell_2 \times m}$. We conclude by noting that in the experiment \mathcal{H}_λ^3 , the adversary's view does not depend on the random bit b . \square

5 Constructing XiO for $\mathbb{P}^{\text{log/poly}}$

We present a modular construction of XiO $\mathbb{P}^{\text{log/poly}}$ from the GSW FHE scheme for circuits of depth δ , denoted by GSW_δ , for sufficiently large δ , and any $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE for sufficiently small h, α and sufficiently large ℓ_1 and ℓ_2 —recall that h measures the LHE succinctness, α quantifies the approximate correctness, while ℓ_1, ℓ_2 measure the plaintext size, or "batching capacity" of the scheme; ℓ_1 intuitively represents how many bits can be packed in a scalar, i.e. how large is the modulus in use, whereas ℓ_2 measures how many such scalars are recovered when decrypting one ciphertext. We prove the IND-security of our XiO construction from 2-circular SRL security of the GSW scheme and the hintable LHE.

Outline. In the rest of this section, we abstract out 2 additional properties of the GSW FHE that we will rely on to enable a modular proof. This is done in Section 5.1. Then, we present our XiO construction in Section 5.2. Afterwards, in Section 5.3 we prove the IND-security of our XiO assuming the 2-circular SRL security of the GSW scheme and the hintable LHE. Finally, in Section 5.4 we instantiate our modular construction with Packed-Regev LHE, which yields our main result, namely the existence of XiO (and then iO) from the subexponential LWE assumption and the subexponential 2CIRC conjecture w.r.t. GSW and Packed-Regev LHE. In Appendix A.2 we provide an instantiation with Damgård Jurik LHE encryption.

5.1 Additional properties for GSW

Weak Circuit Privacy of GSW_δ . As mentioned in the introduction, we will rely on the fact that the GSW encryption scheme also satisfies a notion of "weak circuit privacy" similar to the one defined for LHE. More precisely, we show that GSW satisfies a property that involves a public-key algorithm that re-randomizes evaluated ciphertext so that they look like fresh ciphertexts from the support of the noise encryption algorithm Enc^* . Namely, we show that there exists a PPT algorithm ReRand that takes as input the public key pk , an evaluated ciphertext ct , some random coins $\mathbf{r}^* \in \mathcal{R}^*$, and outputs an evaluated ciphertext ct computed as described below.

- $\text{ReRand}(\text{pk}, \text{ct}; \mathbf{r}^*)$:

Given $\text{pk} = (B, \mathbf{U}, \mathbf{G})$, $\text{ct} \in \{0, 1\}^w$ and $\mathbf{r}^* \leftarrow_{\mathbb{R}} [-B^*, B^*]^m$, it computes $\text{ct}' \in \mathbb{Z}_N^{\kappa+1}$ whose binary decomposition is ct , computes $\tilde{\text{ct}} = \text{ct}' + \mathbf{U}\mathbf{r}^* \in \mathbb{Z}_N^{\kappa+1}$, and outputs the re-randomized ciphertext $\text{BD}(\tilde{\text{ct}}) \in \{0, 1\}^w$.

Theorem 5.1 (weak circuit privacy). *For all polynomials ν, ℓ, δ , all $\lambda \in \mathbb{N}$, all crs containing a modulus $N \in \mathbb{N}$ such that $N \geq 2^{2\lambda}B$, where B is an upper-bound on the noise obtained from homomorphically evaluating circuits of depth at most $\delta(\lambda)$, all pairs (pk, sk) in the support of $\text{Gen}(\text{crs})$, all messages $\mu \in \{0, 1\}^{\nu(\lambda)}$, all depth- $\delta(\lambda)$ circuits $f_1, \dots, f_{\ell(\lambda)} : \{0, 1\}^{\nu(\lambda)} \rightarrow \mathbb{Z}_N$, the following*

distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\mathcal{D}_0 : \left\{ \begin{array}{l} \text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mu), \forall j \in [\ell(\lambda)], \text{ct}_{f_j} = \text{Eval}'(\text{pk}, f_j, 0, \text{ct}), \mathbf{r}_j^* \leftarrow_{\mathcal{R}} \mathcal{R}^* \\ \text{ct}_{f_j}^* = \text{ReRand}(\text{pk}, \text{ct}_{f_j}; \mathbf{r}^*) : \left(\text{ct}, (\mathbf{r}_j^*, \text{ct}_{f_j}^*)_{j \in [\ell(\lambda)]} \right) \end{array} \right\}$$

$$\mathcal{D}_1 : \left\{ \begin{array}{l} r \leftarrow_{\mathcal{R}} \mathcal{R}^{\nu(\lambda)}, \text{ct} = \text{Enc}_{\text{pk}}(\mu; r), \forall j \in [\ell(\lambda)], \mathbf{r}_j^* \leftarrow_{\mathcal{R}} \mathcal{R}^*, \text{ct}_{f_j}^* = \text{Enc}_{\text{pk}}^*(f_j(\mu); \mathbf{r}_j^*) \\ \mathbf{r}_{f_j} = \text{Eval}_{\text{rand}}(\text{pk}, f_j, r, \mu) : \left(\text{ct}, (\mathbf{r}_j^* - \mathbf{r}_{f_j}, \text{ct}_{f_j}^*)_{j \in [\ell(\lambda)]} \right) \end{array} \right\}$$

Proof: By batch correctness of the scheme, for all $j \in [\ell(\lambda)]$, the evaluated ciphertext is of the form $\text{ct}_{f_j} = (\mathbf{A}\mathbf{r}_{f_j}, (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)\mathbf{r}_{f_j} + f_j(\mu)) \in \mathbb{Z}_N^{\kappa+1}$, and the re-randomized ciphertext is of the form $\text{ct}_{f_j}^* = (\mathbf{A}(\mathbf{r}_{f_j} + \mathbf{r}_j^*), (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)(\mathbf{r}_{f_j} + \mathbf{r}_j^*) + f_j(\mu)) \in \mathbb{Z}_N^{\kappa+1}$, where $\|\mathbf{r}_{f_j}\|_\infty < (w+1)^\delta \lceil \log(N) \rceil = 2^{-\lambda} B^*$ and $\mathbf{r}_j^* \leftarrow_{\mathcal{R}} [-B^*, B^*]^m$. By Lemma 2.2 (smudging), the following distributions have statistical distance at most $2^{-\lambda}$:

$$(\mathbf{r}_j^*)_{j \in [\ell]} \approx_s (\mathbf{r}_j^* - \mathbf{r}_{f_j})_{j \in [\ell]}.$$

The leftmost distribution corresponds to \mathcal{D}_0 (with pre and post-processing), whereas the rightmost distribution corresponds to \mathcal{D}_1 (with pre and post-processing). \square

Proposition 13 ($B(2^\lambda+1)$ -approximate correctness of refreshed evaluated ciphertexts). *For all polynomials ν, δ , all $\lambda \in \mathbb{N}$, all crs containing a large enough modulus $N \in \mathbb{N}$, all messages $\mu \in \{0, 1\}^{\nu(\lambda)}$, all circuits $f : \{0, 1\}^{\nu(\lambda)} \rightarrow \mathbb{Z}_N$ of depth at most $\delta(\lambda)$, we have: $\Pr \left[(\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{crs}), \text{ct} \leftarrow \text{Enc}_{\text{pk}}(\mu), \text{ct}_f = \text{Eval}(\text{pk}, f, 0, \text{ct}), \text{ct}'_f \leftarrow \text{ReRand}(\text{pk}, \text{ct}_f) : |\text{sk}^\top \text{ct}'_f - f(\mu)| \leq B(2^\lambda + 1) \right] \in 1 - 2^{-\Omega(\lambda)}$.*

Proof: The ciphertext ct_f is the binary decomposition of $(\mathbf{A}(\mathbf{r}^* + \mathbf{r}_f), (\mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top)(\mathbf{r}^* + \mathbf{r}_f) + f(\mu)) \in \mathbb{Z}_N^{\kappa+1}$, where $|\mathbf{e}^\top \mathbf{r}_f| < B$ by batch correctness and $|\mathbf{e}^\top \mathbf{r}^*| < B_\chi B^* m = 2^\lambda B$ with probability $1 - 2^{-\Omega(\lambda)}$ over the choices of $\mathbf{e} \leftarrow \chi^m$. \square

5.2 XiO Construction

We directly dive into the formal description of the construction, see the introduction for a detailed overview.

We present a modular construction of XiO for the class of circuits $\mathcal{C}_{\log(n), s, d}$ for polynomials n, s, d , from the following building blocks:

- the GSW FHE scheme for depth δ circuits, denoted by $\text{GSW}_\delta = (\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Enc}^*, \text{Dec}, \text{Eval}, \text{Eval}', \text{Eval}_{\text{rand}}, \text{ReRand})$, presented in Section 3.2.2. The depth δ is chosen sufficiently large to handle the homomorphic evaluations of the circuits described below.
- an $(\ell_1, \ell_2, h, \alpha)$ -Hintable Packed LHE, denoted by $\mathcal{LHE}_{b, n, \varepsilon} = (\overline{\text{Gen}}, \overline{\text{Enc}}, \overline{\text{Enc}^*}, \overline{\text{Dec}}, \overline{\text{Eval}}, \overline{\text{SecHint}}, \overline{\text{PubHint}}, \overline{\text{Rec}}, \overline{\text{VerKey}})$ where h is independent of n to ensure succinctness; $\ell_1(\lambda) \geq b(\lambda) + 2\lambda$, where 2^b is a bound on the noise obtained when FHE evaluating circuits of depth at most δ^{10} ; moreover $(\ell_1(\lambda) - b(\lambda) - 2\lambda) \cdot \ell_2(\lambda) \geq n^\varepsilon(\lambda)$, where $\varepsilon \in (0, 1)$ is defined below (see the paragraph about succinctness); finally $\alpha(\lambda) \leq \lambda + 1$.

¹⁰To make sure these parameters are instantiable, we require that LHE decryption is of poly-logarithmic depth, which ensures that δ and therefore B only depend poly-logarithmically on ℓ_1 .

Notations. For every program Π with $\log(n)$ bits inputs, every $\varepsilon \in (0, 1)$, the truth table can be written as $(\Pi_i)_{i \in [n^{1-\varepsilon}]}$, where each chunk Π_i contains n^ε bits. The chunks Π_i themselves can be subdivided into sub-chunks $\Pi_i = (\Pi_{i,j})_{j \in [\ell_2]}$, where each sub-chunk $\Pi_{i,j}$ contains n^ε/ℓ_2 bits. For all $i \in [n^{1-\varepsilon}]$ and $j \in [\ell_2]$, we denote by $C_{i,j}$ the circuit that takes as input a program Π of size s and outputs $\Pi_{i,j}$.

The construction: We proceed to the construction.

• Gen_{Obf}(1^λ):

Set the parameters:

- Choose a constant $0 < \varepsilon < 1$ that is small enough so as to ensure succinctness of the scheme (see paragraph succinctness below).
- Let $|\overline{\text{ct}}|(\cdot)$, $|\overline{\mathbf{r}^*}|(\cdot)$ be polynomials such that for every $\lambda \in \mathbb{N}$, every $(\overline{\text{pk}}, \overline{\text{sk}})$ in the support of $\overline{\text{Gen}}(1^\lambda)$ that defines the message space $\mathbb{Z}_N^{\ell_2}$ and the noisy randomness space \mathcal{R}^* , every message $\mathbf{m} \in \mathbb{Z}_N^{\ell_2}$, every ciphertext in the support of $\overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{m})$ has a bit size at most $|\overline{\text{ct}}|(\lambda)$ and every $\mathbf{r}^* \in \mathcal{R}^*$ has bit size at most $|\overline{\mathbf{r}^*}|(\lambda)$.
- $\text{FHE.PubCoin} \leftarrow_{\mathbb{R}} \{0, 1\}^{n^{1-\varepsilon} \cdot \ell_2 \cdot |\overline{\mathbf{r}^*}|}$, $\text{LHE.PubCoin} \leftarrow_{\mathbb{R}} \{0, 1\}^{n^{1-\varepsilon} \cdot |\overline{\text{ct}}|}$.

Return $\text{pp} = (\text{FHE.PubCoin}, \text{LHE.PubCoin})$.

• Obf(pp, 1^n , Π):

Sample the following parameters:

- $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Gen}}(1^\lambda)$ that defines the message space $\mathbb{Z}_N^{\ell_2}$.
- $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(\overline{\text{pk}})$ that defines the noisy randomness space \mathcal{R}^* , where $\text{sk} \in \mathbb{Z}_N^w$, and pk contains the noise bound B ; we write $b = \lceil \log(B) \rceil$.

Compute the following ciphertexts:

- $\text{ct}_1 \leftarrow \text{Enc}_{\text{pk}}(\Pi)$.
- $\text{ct}_2 \leftarrow \text{Enc}_{\text{pk}}(\overline{\text{sk}})$.
- $\overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk})$.

For all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, compute the following:

- $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1) \in \{0, 1\}^w$, where the circuit $C_{i,j}$ is defined above. The homomorphic evaluation is performed with scaling factor $b + 2\lambda$.
- $\text{ct}_{\text{MSB},i,j} = \text{Eval}'(\text{pk}, f_{i,j}, 0, \text{ct}_2) \in \{0, 1\}^w$, where the circuit $f_{i,j}$ takes as input a bit string $\mathbf{a} \in \{0, 1\}^{|\overline{\text{sk}}|}$. It checks that \mathbf{a} is the secret key associated with $\overline{\text{pk}}$, that is, it runs $\overline{\text{VerKey}}(\overline{\text{pk}}, \mathbf{a})$. If the latter outputs 0, $f_{i,j}$ outputs 0. Otherwise, it uses \mathbf{a} as an LHE secret key to compute $\mathbf{v}_i = \overline{\text{Dec}}_{\mathbf{a}}(\text{LHE.PubCoin}_i)$, where LHE.PubCoin_i is interpreted as an LHE ciphertext $\overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{u}_i)$, with $\mathbf{u}_i \in \mathbb{Z}_N^{\ell_2}$, by density of the LHE ciphertext space. Then it computes $v_{i,j} \in \mathbb{Z}_N$, the j 'th coordinate of $\mathbf{v}_i \in \mathbb{Z}_N^{\ell_2}$ and outputs the most significant bits of $v_{i,j}$, of the form: $\text{MSB}(v_{i,j}) = v_{i,j} - \text{LSB}(v_{i,j}) \in \mathbb{Z}_N^{\ell_2}$, where the (shifted) least significant bits are of the form: $\text{LSB}(v_{i,j}) = v_{i,j} \bmod B2^{2\lambda} - B2^{2\lambda}/2 \in \mathbb{Z}_N$. The homomorphic evaluation is performed with scaling factor 1.

- Parse $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}^* \in \mathcal{R}^*$ and compute $\text{ct}'_{\text{MSB},i,j} = \text{ReRand}(\text{pk}, \text{ct}_{\text{MSB},i,j}; \mathbf{r}_{i,j}^*) \in \mathbb{Z}_N^{\kappa+1}$.
- Compute $\text{ct}_i = (\text{ct}_{i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$, $\text{ct}'_{\text{MSB},i} = (\text{ct}'_{\text{MSB},i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$.
- Compute $\overline{\text{ct}}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \text{LHE.PubCoin}_i, \text{ct}_i - \text{ct}'_{\text{MSB},i})$.
- Compute $\rho_i \leftarrow \overline{\text{SecHint}}(\overline{\text{sk}}, \overline{\text{ct}}_i)$.

Return $\tilde{\Pi} = (\text{pk}, \overline{\text{pk}}, \text{ct}_1, \text{ct}_2, \overline{\text{ct}}, \{\rho_i\}_{i \in [n^{1-\varepsilon}]})$.

• Eval(pp, $\tilde{\Pi}$, \mathbf{x}):

- Let $i \in [n^{1-\varepsilon}]$ such that $\Pi(\mathbf{x})$ belongs to the i 'th chunk of the truth table of Π . Compute $\overline{\text{ct}}_i$ as described above.
- Recover $m_i \leftarrow \overline{\text{Rec}}(\overline{\text{pk}}, \overline{\text{ct}}_i, \rho_i)$.
- Compute $m'_i = \lfloor 2^{-2\lambda}/B \cdot m_i \rfloor$, which contains $\Pi(\mathbf{x})$.

We now proceed to prove Theorem 5.2.

Succinctness. By h -succinctness of $\mathcal{LHE}_{b,n,\varepsilon}$, for all $i \in [n^{1-\varepsilon}]$, we have $|\rho_i| \leq h(\lambda)$ for a polynomial h that is independent of n . The rest of the obfuscated circuit $\tilde{\Pi}$ is of size $p(\lambda, n^\varepsilon, d)$, for a polynomial p that is independent of n and ε . Thus, there exists a constant $c \in \mathbb{N}$ (independent of ε) and a polynomial q (independent of n) such that $|\tilde{\Pi}| \in n^{c\varepsilon}(\lambda) \cdot q(\lambda, d) + n^{1-\varepsilon}(\lambda) \cdot h(\lambda)$. For succinctness, we pick an appropriately small $0 < \varepsilon < 1/c$.

Correctness.

- By the batch correctness of the GSW scheme (Proposition 1), for all $i \in [n^{1-\varepsilon}]$ and $j \in [\ell_2]$, we have:

$$\text{sk}^\top \text{ct}_{i,j} = 2^{2\lambda} B \cdot \Pi_{i,j} + \text{noise}_{i,j} \in \mathbb{Z}_N,$$

where $|\text{noise}_{i,j}| < B$.

- By the density of the noisy ciphertexts of $\mathcal{LHE}_{b,n,\varepsilon}$, for all $i \in [n^{1-\varepsilon}]$, we have $\text{LHE.PubCoin}_i = \overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{u}_i)$ with $\mathbf{u}_i \in \mathbb{Z}_N^{\ell_2}$.
- By the $(2^\lambda + 1)B$ -approximate correctness of refreshed evaluated ciphertexts of the GSW scheme (Proposition 13), for all $i \in [n^{1-\varepsilon}]$ and $j \in [\ell_2]$, we have:

$$\text{sk}^\top \text{ct}'_{\text{MSB},i,j} = \text{MSB}(u_{i,j}) + \text{noise}_{\text{MSB},i,j} \in \mathbb{Z}_N,$$

where $|\text{noise}_{\text{MSB},i,j}| < (2^\lambda + 1)B$.

- By linear homomorphism of $\mathcal{LHE}_{b,n,\varepsilon}$ (Property 4.2), the ciphertext $\overline{\text{ct}}_i$ is in the support of $\overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{m}_i)$, $\mathbf{m}_i = (m_{i,j})_{j \in [\ell_2]}$ of the form:

$$m_{i,j} = B2^{2\lambda} \cdot \Pi_{i,j} + \text{LSB}(u_{i,j}) + \text{noise}_{i,j} + \text{noise}_{\text{MSB},i,j} \in \mathbb{Z}_N.$$

- By α -correctness of the secret hints of $\mathcal{LHE}_{b,n,\varepsilon}$ (Property 4.3), the evaluator of the obfuscated circuit recovers the message $\overline{\mathbf{m}}_i \in \mathbb{Z}_N^{\ell_2}$ such that $\|\overline{\mathbf{m}}_i - \mathbf{m}_i\|_\infty < 2^\alpha(\lambda)$. That is, for all $j \in [\ell_2]$, $\overline{m}_{i,j} = m_{i,j} + \overline{\text{noise}}_{i,j} \in \mathbb{Z}_N$, where $|\overline{\text{noise}}_{i,j}| < 2^\alpha(\lambda)$.

- With probability $1 - 2^{-\Omega(\lambda)}$ over the choice of $\mathbf{u}_i \leftarrow_{\mathcal{R}} \mathbb{Z}_N^{\ell_2}$, we have for all $j \in [\ell_2]$, $|\text{LSB}(u_{i,j}) + \text{noise}_{i,j} + \text{noise}_{\text{MSB},i} + \text{noise}_{i,j}| < B2^{2\lambda}/2$. Thus, $m'_{i,j} = \lfloor \overline{m}_{i,j} \cdot 2^{-2\lambda}/B \rfloor = \Pi_{i,j}$ for all $j \in [\ell_2]$, and the evaluator outputs $\Pi(\mathbf{x})$.

5.3 IND Security from 2CIRC

We now state our theorem.

Theorem 5.2. *Assume that for all polynomials δ, b , all constants $\varepsilon \in (0, 1)$, there exists a polynomial h s.t. for all polynomials n , there exist polynomials ℓ_1, ℓ_2, α and an $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE denoted by $\mathcal{LHE}_{b,\varepsilon,n}$ s.t. for all $\lambda \in \mathbb{N}$:*

- $\alpha(\lambda) \leq \lambda + 1$
- $\ell_1(\lambda) \geq b(\lambda)$
- $(\ell_1(\lambda) - b(\lambda)) \cdot \ell_2(\lambda) > n(\lambda)^\varepsilon$
- 2-circular SRL security (resp. subexponential 2-circular SRL security) holds w.r.t. $\mathcal{LHE}_{b,\varepsilon,n}$ and GSW_δ .

Then XiO (resp. subexponentially secure XiO) for $\text{P}^{\log}/\text{poly}$ exists.

Proof: We now prove that the XiO scheme presented in Section 5 is IND-secure, provided 2-circular SRL security (as per Definition 3.3) holds w.r.t. GSW_δ and $\mathcal{LHE}_{b,\varepsilon,n}$.

We proceed via a hybrid argument using the experiments described below for all $b \in \{0, 1\}$ and all $\lambda \in \mathbb{N}$.

- $\mathcal{H}_\lambda^{b,0}$: this is the experiment from Definition 2.3. For completeness, we describe it here.
 - Generation of pp : for all $i \in [n^{1-\varepsilon}]$, $\text{LHE.PubCoin}_i \leftarrow_{\mathcal{R}} \{0, 1\}^{|\overline{\text{ct}}|}$, for all $j \in [\ell_2]$, $\text{FHE.PubCoin}_{i,j} \leftarrow_{\mathcal{R}} \mathcal{R}^*$, where \mathcal{R}^* denotes the noisy randomness space of \mathcal{FHE} . Return $\text{pp} = \left((\text{LHE.PubCoin}_i)_{i \in [n^{1-\varepsilon}]}, (\text{FHE.PubCoin}_{i,j})_{i \in [n^{1-\varepsilon}], j \in [\ell_2]} \right)$.
 - Generation of $\widetilde{\Pi}_b$: $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Gen}}(1^\lambda)$, $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(\overline{\text{pk}})$, $\text{ct}_1 \leftarrow \text{Enc}_{\text{pk}}(\Pi)$, $\text{ct}_2 \leftarrow \text{Enc}_{\text{pk}}(\overline{\text{sk}})$, $\overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk})$. For all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, compute the following:
 - $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1) \in \{0, 1\}^w$;
 - $\text{ct}_{\text{MSB},i,j} = \text{Eval}'(\text{pk}, f_{i,j}, 0, \text{ct}_2) \in \{0, 1\}^w$;
 - Parse $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}^* \in \mathcal{R}^*$ and compute $\text{ct}'_{\text{MSB},i,j} = \text{ReRand}(\text{pk}, \text{ct}_{\text{MSB},i,j}; \mathbf{r}_{i,j}^*) \in \{0, 1\}^w$.
 - Compute $\text{ct}_i = (\text{ct}_{i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$, $\text{ct}'_{\text{MSB},i} = (\text{ct}'_{\text{MSB},i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$.
 - $\overline{\text{ct}}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \text{LHE.PubCoin}_i, \text{ct}_i - \text{ct}'_{\text{MSB},i})$;
 - $\rho_i \leftarrow \overline{\text{SecHint}}(\overline{\text{sk}}, \overline{\text{ct}}_i)$.

Return $\widetilde{\Pi}_b = (\text{pk}, \overline{\text{pk}}, \text{ct}_1, \text{ct}_2, \overline{\text{ct}}, (\rho_i)_{i \in [n^{1-\varepsilon}]})$.

• $\mathcal{H}_\lambda^{b,1}$: the experiment samples LHE.PubCoin as in $\mathcal{H}_\lambda^{b,0}$, but does not sample FHE.PubCoin just yet; it then generates $\widetilde{\Pi}_b$ as in $\mathcal{H}_\lambda^{b,0}$ up until the point that the $\text{ct}_{\text{MSB},i,j}$ get re-randomized into $\text{ct}'_{\text{MSB},i,j}$ via ReRand. Next, instead of performing the re-randomization, it samples $\text{ct}'_{\text{MSB},i,j}$ as a *fresh* extra noisy encryption of $\text{MSB}(u_{i,j})$ using randomness $\mathbf{r}_{i,j}^* \leftarrow_{\mathcal{R}} \mathcal{R}^*$, and setting $\text{FHE.PubCoin}_{i,j}$ to be $\mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}}$, where $\mathbf{r}_{f_{i,j}}$ denotes the evaluated randomness computed via $\text{Eval}_{\text{rand}}$. Afterwards, the experiment continues exactly the same way as in the experiment $\mathcal{H}_\lambda^{b,0}$.

We show that for all $b \in \{0, 1\}$, we have:

$$\{\mathcal{H}_\lambda^{b,0}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,1}\}_{\lambda \in \mathbb{N}},$$

using the weak circuit privacy of \mathcal{FHE} (Theorem 5.1).

The latter states that for all $\lambda \in \mathbb{N}$, all $(\overline{\text{pk}}, \overline{\text{sk}})$ in the support of $\overline{\text{Gen}}(1^\lambda)$, all (pk, sk) in the support of $\text{Gen}(\overline{\text{pk}})$, for all depth d -circuits and in particular the functions $f_{i,j}$ defined previously, these two distributions have statistical distance at most $2^{-\Omega(\lambda)}$:

$$\mathcal{D}_\lambda^0 : \left\{ \begin{array}{l} r \leftarrow_{\mathcal{R}} ([-1, 1]^{m \times w})^{|\overline{\text{sk}}|}, \text{ct} = \text{Enc}_{\text{pk}}(\overline{\text{sk}}; r), \forall i \in [n^{1-\varepsilon}], j \in [\ell_2], \text{ct}_{f_{i,j}} = \text{Eval}'(\text{pk}, f_{i,j}, 0, \text{ct}) \\ \mathbf{r}_{i,j}^* \leftarrow_{\mathcal{R}} \mathcal{R}^*, \text{ct}_{f_{i,j}}^* = \text{ReRand}(\text{pk}, \text{ct}_{f_{i,j}}, \mathbf{r}_{i,j}^*) : \left(\overline{\text{pk}}, \text{pk}, \text{ct}, \left(\mathbf{r}_{i,j}^*, \text{ct}_{f_{i,j}}^* \right)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]} \right) \end{array} \right\}$$

$$\mathcal{D}_\lambda^1 : \left\{ \begin{array}{l} r \leftarrow_{\mathcal{R}} ([-1, 1]^{m \times w})^{|\overline{\text{sk}}|}, \text{ct} = \text{Enc}_{\text{pk}}(\overline{\text{sk}}; r), \forall i \in [n^{1-\varepsilon}], j \in [\ell_2], \mathbf{r}_{i,j}^* \leftarrow_{\mathcal{R}} \mathcal{R}^* \\ \text{ct}_{f_{i,j}}^* = \text{Enc}_{\text{pk}}^*(\text{MSB}(u_{i,j}); \mathbf{r}_{i,j}^*) \\ \mathbf{r}_{f_{i,j}} = \text{Eval}_{\text{rand}}(\text{pk}, f_{i,j}, r, \overline{\text{sk}}) : \left(\overline{\text{pk}}, \text{pk}, \text{ct}, \left(\mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}}, \text{ct}_{f_{i,j}}^* \right)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]} \right) \end{array} \right\}.$$

We design an inefficient simulator \mathcal{S} that given a tuple $(\overline{\text{pk}}, \text{pk}, \text{ct}, (\mathbf{r}_{i,j}, \text{ct}_{i,j})_{i \in [n^{1-\varepsilon}], j \in [\ell_2]})$, simulates the adversary view in the XiO security experiment. That is, we show that when fed with an input distributed according to \mathcal{D}_λ^0 , \mathcal{S} simulates the experiment $\mathcal{H}_\lambda^{b,0}$, whereas it simulates the experiment $\mathcal{H}_\lambda^{b,1}$ when fed with an input distributed according to \mathcal{D}_λ^1 .

Given $(\overline{\text{pk}}, \text{pk}, \text{ct}, (\mathbf{r}_{i,j}, \text{ct}_{f_{i,j}}^*)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]})$, \mathcal{S} (inefficiently) recovers $\overline{\text{sk}}$ from $\overline{\text{pk}}$, sk from pk , and the randomness r from ct (more precisely \mathcal{S} samples some uniformly random $\overline{\text{sk}}$, sk , r among those that match $\overline{\text{pk}}$, pk and ct). It samples $\text{LHE.PubCoin} \leftarrow_{\mathcal{R}} \{0, 1\}^{n^{1-\varepsilon} \cdot |\overline{\text{ct}}|}$, and for all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, sets $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}$, and $\text{pp} = (\text{LHE.PubCoin}, (\text{FHE.PubCoin}_{i,j})_{i,j})$. It computes $\text{ct}_1 \leftarrow \text{Enc}_{\text{pk}}(\Pi_b)$, $\overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk})$.

For all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, \mathcal{S} computes the following:

- $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1) \in \{0, 1\}^w$;
- $\text{ct}'_{\text{MSB},i,j} = \text{ct}_{i,j} \in \{0, 1\}^w$.
- Compute $\text{ct}_i = (\text{ct}_{i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$, $\text{ct}'_{\text{MSB},i} = (\text{ct}'_{\text{MSB},i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$.
- $\overline{\text{ct}}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \text{LHE.PubCoin}_i, \text{ct}_i - \text{ct}'_{\text{MSB},i})$;
- $\rho_i \leftarrow \overline{\text{SecHint}}(\overline{\text{sk}}, \overline{\text{ct}}_i)$.

The simulator sets $\widetilde{\Pi}_b = (\text{pk}, \overline{\text{pk}}, \text{ct}_1, \text{ct}_2, \overline{\text{ct}}, (\rho_i)_{i \in [n^{1-\varepsilon}]})$, and returns $(\text{pp}, \widetilde{\Pi}_b)$. It is clear from the description of the simulator \mathcal{S} that when the latter is fed with an input distributed according to \mathcal{D}_λ^0 , it simulates the experiment $\mathcal{H}_\lambda^{b,0}$, whereas it simulates the experiment $\mathcal{H}_\lambda^{b,1}$ when fed with an input distributed according to \mathcal{D}_λ^1 .

- $\mathcal{H}_\lambda^{b,2}$: this experiment is the same as $\mathcal{H}_\lambda^{b,1}$, except that instead of sampling LHE.PubCoin_i as random strings, they are sampled as fresh LHE ciphertexts of random plaintexts, that is, of the form $\text{LHE.PubCoin}_i \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{u}_i)$ for $\mathbf{u}_i \leftarrow_{\mathbb{R}} \mathbb{Z}_N^{\ell_2}$. By density of the ciphertexts of \mathcal{LHE} Property 4.7, we have:

$$\{\mathcal{H}_\lambda^{b,1}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,2}\}_{\lambda \in \mathbb{N}}.$$

- $\mathcal{H}_\lambda^{b,3}$: this experiment is the same as $\mathcal{H}_\lambda^{b,2}$, except the ciphertexts $\overline{\text{ct}}_i$ are generated as fresh noisy LHE encryptions of the messages $\mathbf{m}_i = \text{sk}^\top(\text{ct}_i - \text{ct}'_{\text{MSB},i}) + \mathbf{u}_i \in \mathbb{Z}_N^{\ell_2}$, and the LHE.PubCoin_i are instead computed homomorphically by subtracting the LHE encryption of the message $\widetilde{\mathbf{m}}_i = \text{sk}^\top(\text{ct}_i + \text{ct}'_{\text{MSB},i}) \in \mathbb{Z}_N^{\ell_2}$ from the fresh noisy encryption of \mathbf{m}_i . That is, for all $i \in [n^{1-\varepsilon}]$, $\overline{\text{ct}}_i \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}^*(\mathbf{m}_i)$, $\text{LHE.PubCoin}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \overline{\text{ct}}_i, -\text{ct}_i + \text{ct}'_{\text{MSB},i})$.

Note that it is possible to define this hybrid since $\text{ct}'_{\text{MSB},i}$ remains exactly the same no matter what LHE.PubCoin_i is. This was not true in $\mathcal{H}_\lambda^{b,0}$, and we introduced $\mathcal{H}_\lambda^{b,1}$ to break this dependency.

We show that for all $b \in \{0, 1\}$, we have:

$$\{\mathcal{H}_\lambda^{b,2}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,3}\}_{\lambda \in \mathbb{N}},$$

using the weak circuit privacy property of $\mathcal{LHE}_{b,n,\varepsilon}$ (Property 4.6).

The latter states that for all $\lambda \in \mathbb{N}$, all $(\overline{\text{pk}}, \overline{\text{sk}})$ in the support of $\overline{\text{Gen}}(1^\lambda)$, all vectors $\mathbf{x} \in \mathbb{Z}_N^w$ and in particular $\mathbf{x} = \text{sk} \in \mathbb{Z}_N^w$, all $\mathbf{u}_i \in \mathbb{Z}_N^{\ell_2}$, all functions $\mathbf{y}^i = (\mathbf{y}_1^i, \dots, \mathbf{y}_{\ell_2}^i) \in [-1, 1]^{w\ell_2}$ and in particular the vector $\text{ct}_i - \text{ct}'_{\text{MSB},i} \in [-1, 1]^{w\ell_2}$ defined previously for all $i \in [n^{1-\varepsilon}]$, the following distributions have statistical distance $2^{-\Omega(\lambda)}$:

$$\mathcal{D}_\lambda^0 = \left\{ \begin{array}{l} \overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk}), \forall i \in [n^{1-\varepsilon}], \widetilde{\text{ct}}_i \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}^*(\mathbf{u}_i) \\ \overline{\text{ct}}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \widetilde{\text{ct}}_i, \text{ct}_i - \text{ct}'_{\text{MSB},i}) : (\overline{\text{pk}}, \overline{\text{ct}}, (\widetilde{\text{ct}}_i, \overline{\text{ct}}_i)_{i \in [n^{1-\varepsilon}]}) \end{array} \right\}$$

$$\mathcal{D}_\lambda^1 = \left\{ \begin{array}{l} \overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk}), \forall i \in [n^{1-\varepsilon}], \overline{\text{ct}}_i \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}^*(\mathbf{m}_i) \\ \widetilde{\text{ct}}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \overline{\text{ct}}_i, -\text{ct}_i + \text{ct}'_{\text{MSB},i}) : (\overline{\text{pk}}, \overline{\text{ct}}, (\widetilde{\text{ct}}_i, \overline{\text{ct}}_i)_{i \in [n^{1-\varepsilon}]}) \end{array} \right\},$$

where for all $i \in [n^{1-\varepsilon}]$, $\mathbf{m}_i = \text{sk}^\top(\text{ct}_i - \text{ct}'_{\text{MSB},i}) + \mathbf{u}_i \in \mathbb{Z}_N^{\ell_2}$.

We design an inefficient simulator \mathcal{S} that given a tuple $(\overline{\text{pk}}, \overline{\text{ct}}, (\widetilde{\text{ct}}_i, \overline{\text{ct}}_i)_{i \in [n^{1-\varepsilon}]})$, simulates the adversary view in the XiO security experiment. That is, we show that when fed with an input distributed according to \mathcal{D}_λ^0 , \mathcal{S} simulates the experiment $\mathcal{H}_\lambda^{b,2}$, whereas it simulates the experiment $\mathcal{H}_\lambda^{b,3}$ when fed with an input distributed according to \mathcal{D}_λ^1 .

Given $(\overline{\text{pk}}, \overline{\text{ct}}, (\widetilde{\text{ct}}_i, \overline{\text{ct}}_i)_{i \in [n^{1-\varepsilon}]})$, \mathcal{S} (inefficiently) recovers $\overline{\text{sk}}$ from $\overline{\text{pk}}$, sk from $\overline{\text{ct}}$, pk from sk and \mathbf{u}_i from $\widetilde{\text{ct}}_i$ for all $i \in [n^{1-\varepsilon}]$. It generates $\text{ct}_1 \leftarrow \text{Enc}_{\text{pk}}(\Pi_b)$, $r \leftarrow_{\mathbb{R}} ([-1, 1]^{m \times w})^{|\overline{\text{sk}}|}$, $\text{ct}_2 = \text{Enc}_{\text{pk}}(\overline{\text{sk}}; r)$, for all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1)$, $\mathbf{r}_{i,j}^* \leftarrow_{\mathbb{R}} \mathcal{R}^*$, where \mathcal{R}^* denotes the noisy randomness space of \mathcal{FHE}_d , $\text{ct}'_{\text{MSB},i,j} = \text{Enc}_{\text{pk}}^*(\text{MSB}(u_{i,j}); \mathbf{r}_{i,j}^*)$, $\mathbf{r}_{f_{i,j}} = \text{Eval}_{\text{rand}}(\text{pk}, f_{i,j}, r, \overline{\text{sk}})$, $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}}$, $\rho_i \leftarrow \overline{\text{SecHint}}(\text{sk}, \overline{\text{ct}}_i)$, $\text{LHE.PubCoin}_i = \widetilde{\text{ct}}_i$.

It returns $\text{pp} = ((\text{LHE.PubCoin}_i)_i, (\text{FHE.PubCoin}_{i,j})_{i,j})$ and $\widetilde{\Pi}_b = (\text{pk}, \overline{\text{pk}}, \text{ct}_1, \text{ct}_2, \overline{\text{ct}}, (\rho_i)_{i \in [n^{1-\varepsilon}]})$. It is clear from the description of the simulator \mathcal{S} that when the latter is fed with an input distributed according to \mathcal{D}_λ^0 , it simulates $\mathcal{H}_\lambda^{b,2}$, whereas it simulates $\mathcal{H}_\lambda^{b,3}$ when fed with an input distributed according to \mathcal{D}_λ^1 .

- $\mathcal{H}_\lambda^{b,4}$: it is the same experiment as $\mathcal{H}_\lambda^{b,3}$, except the hints ρ_i for $i \in [n^{1-\varepsilon}]$ are computed using

$\overline{\text{PubHint}}(\overline{\text{pk}}, r_i)$, where r_i denotes the randomness used to produce the ciphertext $\overline{\text{ct}}_i$; instead of $\overline{\text{SecHint}}(\overline{\text{sk}}, \overline{\text{ct}}_i)$. By Property 4.4 of the LHE, we have $\{\mathcal{H}_\lambda^{b,3}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,4}\}_{\lambda \in \mathbb{N}}$. Note that in the experiment $\mathcal{H}_\lambda^{b,4}$, we no longer use the LHE secret key $\overline{\text{sk}}$.

- $\mathcal{H}_\lambda^{b,5}$: it is the same experiment as $\mathcal{H}_\lambda^{b,4}$, except that the ciphertexts $\overline{\text{ct}}_i$ are generated as a fresh encryption of a message $\mathbf{m}_i \in \mathbb{Z}_N^{\ell_2}$ of the form $(m_{i,1}, \dots, m_{i,\ell_2})$ where for all $j \in [\ell_2]$, we have $m_{i,j} = B2^{2\lambda} \cdot \Pi_{i,j}^b + \text{LSB}(u_{i,j}) \in \mathbb{Z}_N$. Recall that $\text{LSB}(u_{i,j}) = u_{i,j} \bmod B2^{2\lambda} - B2^{2\lambda}/2 \in \mathbb{Z}_N$. This is instead of having $m_{i,j} = \text{sk}^\top(\text{ct}_{i,j} + \text{ct}'_{\text{MSB},i,j}) + u_{i,j} = B2^{2\lambda} \cdot \Pi_{i,j}^b + \text{LSB}(u_{i,j}) + \text{noise}_{i,j} + \text{noise}_{\text{MSB},i} \in \mathbb{Z}_N$, where $\text{noise}_{i,j} = \mathbf{e}^\top \mathbf{r}_{C_{i,j}} \in \mathbb{Z}_N$ and $\text{noise}_{\text{MSB},i} = \mathbf{e}^\top \mathbf{r}_{i,j}^* \in \mathbb{Z}_N$, $\mathbf{r}_{C_{i,j}}$ is the randomness obtained when evaluating the circuit $C_{i,j}$ on the FHE ciphertext ct_1 , $\mathbf{r}_{i,j}^* \leftarrow_{\mathbb{R}} \mathcal{R}^*$, and $\mathbf{e} \leftarrow \chi^m$ is used to generate pk .

Note that $\text{noise}_{i,j}$ and $\text{noise}_{\text{MSB},i}$ are deterministic functions of pk , $\mathbf{r}_{i,j}^*$ and the randomness $r \in [-1, 1]^{m \times w}$ used to produce ct_1 . In particular, they are independent of $\text{LSB}(u_{i,j})$.

We show that $\{\mathcal{H}_\lambda^{b,3}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,4}\}_{\lambda \in \mathbb{N}}$. To do so, for all $\lambda \in \mathbb{N}$, we exhibit two distributions \mathcal{D}_λ^0 and \mathcal{D}_λ^1 , together with a (possibly inefficient) simulator \mathcal{S} , such that (1) \mathcal{D}_λ^0 and \mathcal{D}_λ^1 have statistical distance $2^{-\Omega(\lambda)}$, and (2) for all $\beta \in \{0, 1\}$, when fed with an input from distribution $\mathcal{D}_\lambda^\beta$, \mathcal{S} produces the adversary view as in the experiment $\mathcal{H}_\lambda^{b,4+\beta}$.

The distributions are defined as follows (the differences are highlighted in red):

$$\mathcal{D}_\lambda^0 = \left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{crs}), r \leftarrow_{\mathbb{R}} ([-1, 1]^{m \times w})^s, \forall i \in [n^{1-\varepsilon}], j \in [\ell_2], \mathbf{r}_{i,j}^* \leftarrow_{\mathbb{R}} \mathcal{R}^* \\ \gamma_{i,j} \leftarrow_{\mathbb{R}} [-B2^{2\lambda}/2 + 1, B2^{2\lambda}/2] : \left(\text{pk}, r, (\gamma_{i,j}, \mathbf{r}_{i,j}^*)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]} \right) \end{array} \right\}$$

$$\mathcal{D}_\lambda^1 = \left\{ \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\text{crs}), r \leftarrow_{\mathbb{R}} ([-1, 1]^{m \times w})^s, \forall i \in [n^{1-\varepsilon}], j \in [\ell_2], \mathbf{r}_{i,j}^* \leftarrow_{\mathbb{R}} \mathcal{R}^* \\ \gamma_{i,j} \leftarrow_{\mathbb{R}} [-B2^{2\lambda}/2 + 1, B2^{2\lambda}/2] : \left(\text{pk}, r, (\gamma_{i,j} + \text{noise}_{i,j} + \text{noise}_{\text{MSB},i,j}, \mathbf{r}_{i,j}^*)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]} \right) \end{array} \right\},$$

where $\text{noise}_{i,j}$ $\text{noise}_{\text{MSB},i,j}$ are functions of pk , r and $\mathbf{r}_{i,j}^*$ defined as below, namely, $\text{noise}_{i,j} = \mathbf{e}^\top \mathbf{r}_{C_{i,j}}$ and $\text{noise}_{\text{MSB},i,j} = \mathbf{e}^\top \mathbf{r}_{i,j}^*$.

We show that these distributions have statistical distance $2^{-\Omega(\lambda)}$. The only difference is that in \mathcal{D}_λ^1 , an extra noise $\text{noise}_{i,j} + \text{noise}_{\text{MSB},i,j}$ is added to the random value $s_{i,j}$. This noise is small, indeed $|\text{noise}_{i,j} + \text{noise}_{\text{MSB},i,j}| \leq B(2^\lambda + 1)$ (see the correctness section for more details). Moreover, $\gamma_{i,j}$ is sampled uniformly at random over $[-B2^{2\lambda}/2 + 1, B2^{2\lambda}/2]$, independently of the other values output by the distributions. Thus, we can use the value $\gamma_{i,j}$ to smudge the noise $\text{noise}_{i,j} + \text{noise}_{\text{MSB},i,j}$. That is, by Lemma 2.2 (smudging), the statistical distance of the two distributions is $2^{-\Omega(\lambda)}$.

Now, we proceed to describe the simulator \mathcal{S} . Given as input the tuple $(\text{pk}, r, (v_{i,j}, \mathbf{r}_{i,j}^*)_{i \in [n^{1-\varepsilon}], j \in [\ell_2]})$, the simulator (inefficiently) recovers sk from pk , samples $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Gen}}(1^\lambda)$, generates $\overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk})$, $\text{ct}_1 = \text{Enc}_{\text{pk}}(\Pi_b; r)$. It samples $r' \leftarrow_{\mathbb{R}} ([-1, 1]^{m \times w})^{|\overline{\text{sk}}|}$, computes $\text{ct}_2 = \text{Enc}_{\text{pk}}(\overline{\text{sk}}; r')$.

For all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, samples $\omega_{i,j} \leftarrow_{\mathbb{R}} \mathbb{Z}_N / (2^{2\lambda} B)$, and sets $m_{i,j} = B2^{2\lambda} \cdot \Pi_{i,j}^b + v_{i,j} + \omega_{i,j} \in \mathbb{Z}_N$. Note that $(\gamma_{i,j}, \omega_{i,j})$ is identically distributed to $(\text{LSB}(u_{i,j}), \text{MSB}(u_{i,j}))$ for $u_{i,j} \leftarrow_{\mathbb{R}} \mathbb{Z}_N$.

It sets $\mathbf{m}_i = (m_{i,1}, \dots, m_{i,\ell_2}) \in \mathbb{Z}_N^{\ell_2}$, samples $r_i \leftarrow_{\mathbb{R}} \overline{\mathcal{R}}$, where $\overline{\mathcal{R}}$ denotes the randomness space of $\overline{\text{Enc}}$, computes $\overline{\text{ct}}_i = \overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{m}_i; r_i)$ and $\rho_i \leftarrow \overline{\text{PubHint}}(\overline{\text{pk}}, r_i)$. It computes $\text{ct}'_{\text{MSB},i,j} = \text{Enc}_{\text{pk}}^*(\omega_{i,j}; \mathbf{r}_{i,j}^*)$, $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1)$, $\text{ct}_i = (\text{ct}_{i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$, $\text{ct}'_{\text{MSB},i} = (\text{ct}'_{\text{MSB},i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$, and $\text{LHE.PubCoin}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \overline{\text{ct}}_i, -\text{ct}_i + \text{ct}'_{\text{MSB},i})$. It computes $\mathbf{r}_{f_{i,j}} = \text{Eval}_{\text{rand}}(\text{pk}, f_{i,j}, \text{ct}_2)$ where the functions $f_{i,j}$ are defined as before, and sets $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}} \in \mathbb{Z}_N^m$.

It returns $\text{pp} = (\text{LHE.PubCoin}_i, \text{FHE.PubCoin}_{i,j})_{i \in [n^{1-\varepsilon}]}$, and $\tilde{\Pi}_b = (\text{pk}, \overline{\text{pk}}, \text{ct}_1, \text{ct}_2, \overline{\text{ct}}, (\rho_i)_{i \in [n^{1-\varepsilon}]})$. When \mathcal{S} is fed with an input distributed according to \mathcal{D}_λ^0 , it simulates the experiment $\mathcal{H}_\lambda^{b,4}$, whereas it simulates the experiment $\mathcal{H}_\lambda^{b,5}$ when fed with distribution \mathcal{D}_λ^1 . Thus, we have:

$$\{\mathcal{H}_\lambda^{b,4}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{b,5}\}_{\lambda \in \mathbb{N}}.$$

- $\{\mathcal{H}_\lambda^{0,5}\}_{\lambda \in \mathbb{N}} \approx_c \{\mathcal{H}_\lambda^{1,5}\}_{\lambda \in \mathbb{N}}$: To complete the proof, we show that $\{\mathcal{H}_\lambda^{0,5}\}_{\lambda \in \mathbb{N}}$ is computationally indistinguishable from $\{\mathcal{H}_\lambda^{1,5}\}_{\lambda \in \mathbb{N}}$. These two ensembles are the same except the former obfuscates the program Π^0 , whereas the latter obfuscates Π^1 . Note that other than the encrypted key cycle, we never use the FHE secret key, and due to hybrid $\{\mathcal{H}_\lambda^{b,4}\}_{\lambda \in \mathbb{N}}$, we no longer use the LHE secret key. The coins FHE.PubCoin exactly correspond to an SRL leakage on the FHE ciphertext ct_2 (and note that in the experiment we do know the output $\alpha_{i,j}$ of the function $f_{i,j}$ that is applied to the plaintexts encrypted in ct_1, ct_2 —namely, it is $\text{MSB}(u_{i,j})$ where $u_{i,j}$ is a random element of \mathbb{Z}_N selected in the experiment, see Hybrid $\mathcal{H}_\lambda^{b,2}$). Thus, (subexponential) indistinguishability of $\{\mathcal{H}_\lambda^{0,5}\}_{\lambda \in \mathbb{N}}$ and $\{\mathcal{H}_\lambda^{1,5}\}_{\lambda \in \mathbb{N}}$ follows from (subexponential) 2-circular SRL-security of \mathcal{FHE}_d and $\mathcal{LHE}_{b,n,\varepsilon}$.

Namely, for any nuPPT (resp. subexponential) distinguisher \mathcal{A} , we show there exists a nuPPT (resp. subexponential) reduction \mathcal{B} such that for all $\lambda \in \mathbb{N}$, we have:

$$\Pr \left[b \leftarrow_{\mathbb{R}} \{0, 1\}, (\text{pp}, \tilde{\Pi}^b) \leftarrow \mathcal{H}_\lambda^{b,5} : \mathcal{A}(\text{pp}, \tilde{\Pi}^b) = b \right] \leq \Pr \left[\text{Exp}_{\lambda, \mathcal{B}}^{\mathcal{FHE}_d, \mathcal{LHE}_{b,n,\varepsilon}} = 1 \right] + 2^{-\Omega(\lambda)},$$

where the experiment $\text{Exp}_{\lambda, \mathcal{B}}^{\mathcal{FHE}_d, \mathcal{LHE}_{b,n,\varepsilon}}$ is described in Definition 2.9.

We now proceed to describe the reduction \mathcal{B} . Given the public keys $\text{pk}, \overline{\text{pk}}$, \mathcal{B} sends the pair (Π^0, Π^1) to the 2-circular SRL security experiment, upon which it receives the ciphertexts $\text{ct} = (\text{ct}_1 \| \text{ct}_2)$ and $\overline{\text{ct}}$, where ct_1 encrypts Π^0 or Π^1 , ct_2 encrypts $\overline{\text{sk}}$, and $\overline{\text{ct}}$ encrypts sk . It samples $u_{i,j} \leftarrow_{\mathbb{R}} \mathbb{Z}_N$ for all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, and generates the following:

- Generation of pp :
 - To generate LHE.PubCoin_i :
 - * It samples $r_i \leftarrow_{\mathbb{R}} \overline{\mathcal{R}}$, where $\overline{\mathcal{R}}$ denotes the randomness space of $\overline{\text{Enc}}$, and computes $\overline{\text{ct}}_i = \overline{\text{Enc}}_{\overline{\text{pk}}}(\mathbf{m}_i; r_i)$, where for all $j \in [\ell_2]$, the j 'th coordinate of \mathbf{m}_i is of the form: $m_{i,j} = B2^{2\lambda} \cdot \Pi_{i,j}^b + \text{LSB}(u_{i,j}) \in \mathbb{Z}_N$. Note that this does not require to know the bit b , since $\Pi_{i,j}^0 = \Pi_{i,j}^1$ for all $i \in [n^{1-\varepsilon}]$, $j \in [\ell_2]$, because the programs Π^0 and Π^1 are functionally equivalent.
 - * It computes $\rho_i \leftarrow \overline{\text{PubHint}}(\overline{\text{pk}}, r_i)$.
 - * It computes $\text{ct}_{i,j} = \text{Eval}'(\text{pk}, C_{i,j}, b + 2\lambda, \text{ct}_1)$.
 - * Then, it queries its \mathcal{O}_{SRL} oracle, to obtain a fresh, extra noisy encryption $\text{Enc}_{\text{pk}}^*(0; \mathbf{r}_{i,j}^*)$. It leaves the oracle \mathcal{O}_{SRL} pending.
 - * It adds the vector $(\mathbf{0}, \text{MSB}(u_{i,j})) \in \mathbb{Z}_N^{\kappa+1}$ to the vector whose binary decomposition is $\text{Enc}_{\text{pk}}^*(0; \mathbf{r}_{i,j}^*)$, which yields a vector $\text{ct}_{i,j}^* \in \mathbb{Z}_N^{\kappa+1}$ whose binary decomposition is $\text{Enc}_{\text{pk}}^*(\text{MSB}(u_{i,j}); \mathbf{r}_{i,j}^*)$. Then, it computes $\text{ct}'_{\text{MSB},i,j} = \text{BD}(\text{ct}_{i,j}^*) \in \{0, 1\}^w$.
 - * It computes $\text{ct}_i = (\text{ct}_{i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$ and $\text{ct}'_{\text{MSB},i} = (\text{ct}'_{\text{MSB},i,j})_{j \in [\ell_2]} \in \{0, 1\}^{w\ell_2}$.
 - * Finally, it computes $\text{LHE.PubCoin}_i = \overline{\text{Eval}}(\overline{\text{pk}}, \overline{\text{ct}}, \overline{\text{ct}}_i, -\text{ct}_i - \text{ct}'_{\text{MSB},i})$.
 - To generate $\text{FHE.PubCoin}_{i,j}$: it answers the pending oracle \mathcal{O}_{SRL} with the function $f_{i,j}$ and the value $\alpha = \text{MSB}(u_{i,j})$. With probability $1 - 2^{-\Omega(\lambda)}$ over the random coins used to produce The oracle \mathcal{O}_{SRL} returns the leakage $\mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}} \in \mathcal{R}^*$. The reduction sets $\text{FHE.PubCoin}_{i,j} = \mathbf{r}_{i,j}^* - \mathbf{r}_{f_{i,j}}$.

It sets $\text{pp} = ((\text{LHE.PubCoin}_i)_{i \in [n^{1-\varepsilon}]}, (\text{FHE.PubCoin}_{i,j})_{i \in [n^{1-\varepsilon}], j \in [\ell_2]})$.

- Generation of $\tilde{\Pi}$: it sets $(\text{pk}, \overline{\text{pk}}, \overline{\text{ct}}, \text{ct}, (\rho_i)_{i \in [n^{1-\varepsilon}]})$ computed as described above.

The reduction \mathcal{B} sends $(\text{pp}, \tilde{\Pi})$ to the distinguisher \mathcal{A} , which outputs a bit b' . Finally, \mathcal{B} outputs the bit b' . When ct_1 encrypts Π^0 , the reduction simulates the experiment $\mathcal{H}_\lambda^{0.5}$ to \mathcal{A} , whereas it simulates the experiment $\mathcal{H}_\lambda^{1.5}$ when ct_1 encrypts Π^1 . Thus, we have $\Pr \left[\text{Exp}_{\lambda, \mathcal{B}}^{\mathcal{FHE}_d, \mathcal{LHE}_{b,n,\varepsilon}} = 1 \right] \geq \Pr \left[b \leftarrow_{\text{R}} \{0, 1\}, (\text{pp}, \tilde{\Pi}^b) \leftarrow \mathcal{H}_\lambda^{b.5} : \mathcal{A}(\text{pp}, \tilde{\Pi}^b) = b \right]$. Overall, we have shown that:

$$\begin{aligned} \{\mathcal{H}_\lambda^{0.1}\}_{\lambda \in \mathbb{N}} &\approx_s \{\mathcal{H}_\lambda^{0.2}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{0.3}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{0.4}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{0.5}\}_{\lambda \in \mathbb{N}} \approx_c \\ \{\mathcal{H}_\lambda^{1.5}\}_{\lambda \in \mathbb{N}} &\approx_s \{\mathcal{H}_\lambda^{1.4}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{1.3}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{1.2}\}_{\lambda \in \mathbb{N}} \approx_s \{\mathcal{H}_\lambda^{1.1}\}_{\lambda \in \mathbb{N}}. \end{aligned}$$

□

5.4 Instantiation with Packed-Regev LHE

We now state the results when instantiating our modular XiO construction with the Packed-Regev LHE, presented in Section 4.2. This construction is parameterized by polynomials ℓ_1 and ℓ_2 , and it is denoted by $\text{P-Regev}_{\ell_1, \ell_2}$.

By combining Theorem 4.3 (i.e. security of $\text{P-Regev}_{\ell_1, \ell_2}$ for all polynomials ℓ_1 and ℓ_2 under LWE) and Theorem 3.5 (i.e. SRL-security of the GSW_δ for all polynomials δ under LWE), we get:

Lemma 5.1. *Assume the (subexponential) the LWE assumption holds. Assume further that for all polynomials δ, ℓ_1 and ℓ_2 , $2\text{CIRC}^{\text{SRL}}$ holds w.r.t. GSW_δ and $\text{P-Regev}_{\ell_1, \ell_2}$. Then, for all polynomials δ, ℓ_1 and ℓ_2 , (subexponential) 2-circular SRL security holds w.r.t. GSW_δ and $\text{P-Regev}_{\ell_1, \ell_2}$.*

By combining Lemma 5.1 above with Theorem 4.2, which states that for all polynomials ℓ_1 , there exists a polynomial h such that for all polynomials ℓ_2 , $\text{P-Regev}_{\ell_1, \ell_2}$ is an $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE with $\alpha(\lambda) = \lambda + 1$, together with Theorem 5.2, we get:

Theorem 5.3. *Assume the (subexponential) LWE assumption holds. Assume further that for all polynomials δ, ℓ_1 and ℓ_2 , the (subexponential) $2\text{CIRC}^{\text{SRL}}$ conjecture holds w.r.t. GSW_δ and $\text{P-Regev}_{\ell_1, \ell_2}$. Then (subexponentially-secure) XiO for $\text{P}^{\log}/\text{poly}$ exists.*

Finally, combining Theorem 5.3 with Theorem 2.4 (i.e., $i\mathcal{O}$ from XiO and LWE) we get our main theorem:

Theorem 5.4. *Assume the subexponential LWE assumption holds. Assume further that for all polynomials δ, ℓ_1 and ℓ_2 , the subexponential $2\text{CIRC}^{\text{SRL}}$ conjecture holds w.r.t. GSW_δ and $\text{P-Regev}_{\ell_1, \ell_2}$. Then subexponentially-secure $i\mathcal{O}$ for P/poly exists.*

Acknowledgments

We wish to thank Hoeteck Wee and Daniel Wichs for their insightful feedback on a previous eprint version of this paper. In particular, they encouraged us to present a game-based definition of SRL security (as opposed to the indistinguishability based definition we presented in our earlier draft), and to clarify differences between circular SRL-security and “plain” circular security.

References

- [ABBC10] T. Acar, M. Belenkiy, M. Bellare, and D. Cash. Cryptographic agility and its relation to circular encryption. In *EUROCRYPT 2010, LNCS 6110*, pages 403–422. Springer, Heidelberg, May / June 2010.
- [ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO 2009, LNCS 5677*, pages 595–618. Springer, Heidelberg, August 2009.
- [Agr19] S. Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In *EUROCRYPT 2019, Part I, LNCS 11476*, pages 191–225. Springer, Heidelberg, May 2019.
- [AJ15] P. Ananth and A. Jain. Indistinguishability obfuscation from compact functional encryption. In *CRYPTO 2015, Part I, LNCS 9215*, pages 308–326. Springer, Heidelberg, August 2015.
- [AJKS18] P. Ananth, A. Jain, D. Khurana, and A. 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>.
- [AJL⁺12] G. Asharov, A. Jain, A. López-Alt, E. Tromer, V. Vaikuntanathan, and D. Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In *EUROCRYPT 2012, LNCS 7237*, pages 483–501. Springer, Heidelberg, April 2012.
- [AJL⁺19] P. Ananth, A. Jain, H. Lin, C. Matt, and A. Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. Cryptology ePrint Archive, Report 2019/643, 2019. <https://eprint.iacr.org/2019/643>.
- [Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996.
- [AP09] J. Alwen and C. Peikert. Generating Shorter Bases for Hard Random Lattices. In *26th International Symposium on Theoretical Aspects of Computer Science STACS 2009*, Proceedings of the 26th Annual Symposium on the Theoretical Aspects of Computer Science, pages 75–86, Freiburg, Germany, February 2009. IBFI Schloss Dagstuhl.
- [AP20] S. Agrawal and A. Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In *EUROCRYPT 2020, Part I, LNCS*, pages 110–140. Springer, Heidelberg, May 2020.
- [BCP14] E. Boyle, K.-M. Chung, and R. Pass. On extractability obfuscation. In *TCC*, pages 52–73, 2014.
- [BDGM19] Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In *TCC 2019, Part II, LNCS*, pages 407–437. Springer, Heidelberg, March 2019.
- [BDGM20a] Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Candidate iO from homomorphic encryption schemes. In *EUROCRYPT 2020, Part I, LNCS*, pages 79–109. Springer, Heidelberg, May 2020.

- [BDGM20b] Z. Brakerski, N. Döttling, S. Garg, and G. 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>.
- [BGI⁺01] B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (im)possibility of obfuscating programs. In *CRYPTO 2001, LNCS 2139*, pages 1–18. Springer, Heidelberg, August 2001.
- [BGL⁺15] N. Bitansky, S. Garg, H. Lin, R. Pass, and S. Telang. Succinct randomized encodings and their applications. *IACR Cryptology ePrint Archive*, 2015:356, 2015.
- [BHW15] A. Bishop, S. Hohenberger, and B. Waters. New circular security counterexamples from decision linear and learning with errors. In *ASIACRYPT 2015, Part II, LNCS 9453*, pages 776–800. Springer, Heidelberg, November / December 2015.
- [BIJ⁺20] J. Bartusek, Y. Ishai, A. Jain, F. Ma, A. Sahai, and M. Zhandry. Affine determinant programs: A framework for obfuscation and witness encryption. In *ITCS 2020*, pages 82:1–82:39. LIPIcs, January 2020.
- [BP15] N. Bitansky and O. Paneth. ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In *TCC 2015, Part II, LNCS 9015*, pages 401–427. Springer, Heidelberg, March 2015.
- [BPR15] N. Bitansky, O. Paneth, and A. Rosen. On the cryptographic hardness of finding a Nash equilibrium. In *56th FOCS*, pages 1480–1498. IEEE Computer Society Press, October 2015.
- [BPW16] N. Bitansky, O. Paneth, and D. Wichs. Perfect structure on the edge of chaos - trapdoor permutations from indistinguishability obfuscation. In *TCC 2016-A, Part I, LNCS 9562*, pages 474–502. Springer, Heidelberg, January 2016.
- [BR93] M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
- [BRS02] J. Black, P. Rogaway, and T. Shrimpton. Encryption-scheme security in the presence of key-dependent messages. Cryptology ePrint Archive, Report 2002/100, 2002. <http://eprint.iacr.org/2002/100>.
- [BV15] N. Bitansky and V. Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In *56th FOCS*, pages 171–190. IEEE Computer Society Press, October 2015.
- [BZ14] D. Boneh and M. Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 480–499, 2014.
- [CGH98] R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited (preliminary version). In *30th ACM STOC*, pages 209–218. ACM Press, May 1998.
- [CGH12] D. Cash, M. Green, and S. Hohenberger. New definitions and separations for circular security. In *PKC 2012, LNCS 7293*, pages 540–557. Springer, Heidelberg, May 2012.

- [CHJV14] R. Canetti, J. Holmgren, A. Jain, and V. Vaikuntanathan. Indistinguishability obfuscation of iterated circuits and RAM programs. Cryptology ePrint Archive, Report 2014/769, 2014. <http://eprint.iacr.org/2014/769>.
- [CHL⁺15] J. H. Cheon, K. Han, C. Lee, H. Ryu, and D. Stehlé. Cryptanalysis of the multilinear map over the integers. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 3–12, 2015.
- [CKP15] R. Canetti, Y. T. Kalai, and O. Paneth. On obfuscation with random oracles. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, pages 456–467, 2015.
- [CL01] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT 2001, LNCS 2045*, pages 93–118. Springer, Heidelberg, May 2001.
- [CLP15] K.-M. Chung, H. Lin, and R. Pass. Constant-round concurrent zero-knowledge from indistinguishability obfuscation. In *CRYPTO 2015, Part I, LNCS 9215*, pages 287–307. Springer, Heidelberg, August 2015.
- [CLT13] J.-S. Coron, T. Lepoint, and M. Tibouchi. Practical multilinear maps over the integers. In *CRYPTO 2013, Part I, LNCS 8042*, pages 476–493. Springer, Heidelberg, August 2013.
- [CLT15] J.-S. Coron, T. Lepoint, and M. Tibouchi. New multilinear maps over the integers. In *CRYPTO 2015, Part I, LNCS 9215*, pages 267–286. Springer, Heidelberg, August 2015.
- [DJ01] I. Damgård and M. Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In *PKC 2001, LNCS 1992*, pages 119–136. Springer, Heidelberg, February 2001.
- [Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
- [GGH13a] S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT 2013, LNCS 7881*, pages 1–17. Springer, Heidelberg, May 2013.
- [GGH⁺13b] S. Garg, C. Gentry, S. Halevi, M. Raykova, A. Sahai, and B. Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- [GGH15] C. Gentry, S. Gorbunov, and S. Halevi. Graph-induced multilinear maps from lattices. In *TCC 2015, Part II, LNCS 9015*, pages 498–527. Springer, Heidelberg, March 2015.
- [GGHR14] S. Garg, C. Gentry, S. Halevi, and M. Raykova. Two-round secure MPC from indistinguishability obfuscation. In *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, pages 74–94, 2014.
- [GH10] M. Green and S. Hohenberger. Cpa and cca-secure encryption systems that are not 2-circular secure, 2010. matthewdgreen@gmail.com 14686 received 16 Mar 2010, last revised 18 Mar 2010.

- [GJK18] C. Gentry, C. S. Jutla, and D. Kane. Obfuscation using tensor products. In *Electronic Colloquium on Computational Complexity (ECCC)*, page 149, 2018.
- [GJLS20] R. Gay, A. Jain, H. Lin, and A. Sahai. Indistinguishability obfuscation from simple-to-state hard problems: New assumptions, new techniques, and simplification. Technical report, Cryptology ePrint Archive, Report 2020/764, 2020. <https://eprint.iacr.org/2020/764>, 2020.
- [GK05] S. Goldwasser and Y. T. Kalai. On the impossibility of obfuscation with auxiliary input. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 553–562, 2005.
- [GKW17] R. Goyal, V. Koppula, and B. Waters. Separating semantic and circular security for symmetric-key bit encryption from the learning with errors assumption. In *EUROCRYPT 2017, Part II, LNCS 10211*, pages 528–557. Springer, Heidelberg, April / May 2017.
- [GLSW14] C. Gentry, A. Lewko, A. Sahai, and B. Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. Cryptology ePrint Archive, Report 2014/309, 2014.
- [GM84] S. Goldwasser and S. Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *40th ACM STOC*, pages 197–206. ACM Press, May 2008.
- [GSW13] C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO 2013, Part I, LNCS 8042*, pages 75–92. Springer, Heidelberg, August 2013.
- [ILL89] R. Impagliazzo, L. A. Levin, and M. Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989.
- [JLMS19] A. Jain, H. Lin, C. Matt, and A. Sahai. How to leverage hardness of constant-degree expanding polynomials over \mathbb{R} to build $i\mathcal{O}$. In *EUROCRYPT 2019, Part I, LNCS 11476*, pages 251–281. Springer, Heidelberg, May 2019.
- [JLS20] A. Jain, H. Lin, and A. Sahai. Indistinguishability obfuscation from well-founded assumptions. Cryptology ePrint Archive, Report 2020/1003, 2020. <https://eprint.iacr.org/2020/1003>.
- [JS18] A. Jain and A. Sahai. How to leverage hardness of constant-degree expanding polynomials over \mathbb{R} to build $i\mathcal{O}$. Cryptology ePrint Archive, Report 2018/973, 2018. <https://eprint.iacr.org/2018/973>.
- [KLW15] V. Koppula, A. B. Lewko, and B. Waters. Indistinguishability obfuscation for turing machines with unbounded memory. In *47th ACM STOC*, pages 419–428. ACM Press, June 2015.

- [KMN⁺14] I. Komargodski, T. Moran, M. Naor, R. Pass, A. Rosen, and E. Yogev. One-way functions and (im)perfect obfuscation. In *55th FOCS*, pages 374–383. IEEE Computer Society Press, October 2014.
- [KNY14] I. Komargodski, M. Naor, and E. Yogev. Secret-sharing for NP. In *ASIACRYPT 2014, Part II, LNCS 8874*, pages 254–273. Springer, Heidelberg, December 2014.
- [KRW15] V. Koppula, K. Ramchen, and B. Waters. Separations in circular security for arbitrary length key cycles. In *TCC 2015, Part II, LNCS 9015*, pages 378–400. Springer, Heidelberg, March 2015.
- [KW16] V. Koppula and B. Waters. Circular security separations for arbitrary length cycles from LWE. In *CRYPTO 2016, Part II, LNCS 9815*, pages 681–700. Springer, Heidelberg, August 2016.
- [Lin16] H. Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In *EUROCRYPT 2016, Part I, LNCS 9665*, pages 28–57. Springer, Heidelberg, May 2016.
- [Lin17] H. Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In *CRYPTO 2017, Part I, LNCS 10401*, pages 599–629. Springer, Heidelberg, August 2017.
- [LPST16] H. Lin, R. Pass, K. Seth, and S. Telang. Indistinguishability obfuscation with non-trivial efficiency. In *PKC 2016, Part II, LNCS 9615*, pages 447–462. Springer, Heidelberg, March 2016.
- [LT17] H. Lin and S. Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In *CRYPTO 2017, Part I, LNCS 10401*, pages 630–660. Springer, Heidelberg, August 2017.
- [LV16] H. Lin and V. Vaikuntanathan. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In *57th FOCS*, pages 11–20. IEEE Computer Society Press, October 2016.
- [MF15] B. Minaud and P.-A. Fouque. Cryptanalysis of the new multilinear map over the integers. Cryptology ePrint Archive, Report 2015/941, 2015. <http://eprint.iacr.org/>.
- [Mic01] D. Micciancio. Improving lattice based cryptosystems using the hermite normal form. In *Cryptography and Lattices Conference — CaLC 2001, Lecture Notes in Computer Science 2146*, pages 126–145, Providence, Rhode Island, 29–30 March 2001. Springer-Verlag.
- [Mic19] D. Micciancio. From linear functions to fully homomorphic encryption. <https://bacrypto.github.io/presentations/2018.11.30-micciancio-fhe.pdf>. Technical report, 2019.
- [MMN15] M. Mahmoody, A. Mohammed, and S. Nematihaji. More on impossibility of virtual black-box obfuscation in idealized models. *IACR Cryptology ePrint Archive*, 2015:632, 2015.

- [MO14] A. Marcedone and C. Orlandi. Obfuscation \Rightarrow (IND-CPA security $\not\Rightarrow$ circular security). In *SCN 14, LNCS 8642*, pages 77–90. Springer, Heidelberg, September 2014.
- [MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT 2012, LNCS 7237*, pages 700–718. Springer, Heidelberg, April 2012.
- [MRH04] U. M. Maurer, R. Renner, and C. Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In *TCC 2004, LNCS 2951*, pages 21–39. Springer, Heidelberg, February 2004.
- [Pai99] P. Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT’99, LNCS 1592*, pages 223–238. Springer, Heidelberg, May 1999.
- [PRS17] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *49th ACM STOC*, pages 461–473. ACM Press, June 2017.
- [Ps16] R. Pass and a. shelat. Impossibility of VBB obfuscation with ideal constant-degree graded encodings. In *TCC 2016-A, Part I, LNCS 9562*, pages 3–17. Springer, Heidelberg, January 2016.
- [PST14] R. Pass, K. Seth, and S. Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In *CRYPTO 2014, Part I, LNCS 8616*, pages 500–517. Springer, Heidelberg, August 2014.
- [PVW08] C. Peikert, V. Vaikuntanathan, and B. Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO 2008, LNCS 5157*, pages 554–571. Springer, Heidelberg, August 2008.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
- [Rot13] R. Rothblum. On the circular security of bit-encryption. In *TCC 2013, LNCS 7785*, pages 579–598. Springer, Heidelberg, March 2013.
- [SW14] A. Sahai and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
- [WW20] H. Wee and D. Wichs. Candidate obfuscation via oblivious lwe sampling. Cryptology ePrint Archive, Report 2020/1042, 2020. <https://eprint.iacr.org/2020/1042>.
- [WZ17] D. Wichs and G. Zirdelis. Obfuscating compute-and-compare programs under LWE. In *58th FOCS*, pages 600–611. IEEE Computer Society Press, October 2017.

A XiO from the 2-circular SRL Security of the GSW FHE and the DJ LHE

A.1 Hintable LHE from DCR

We consider the Damgård Jurik (DJ) Linearly Homomorphic Encryption scheme from [DJ01], which generalizes Paillier’s encryption scheme [Pai99] to larger message spaces, whose security relies on the Decisional Composite Residuosity (DCR) assumption. This scheme is not needed for our main result,

but serves as a good warm-up for understanding the notion of a hintable packed LHE (and leads to our simplest construction of an XiO). As mentioned in the introduction, BDGM already showed that the DJ scheme is a hintable LHE; since our notion of a hintable LHE is somewhat different, for completeness, we here include the standard proof. We start by recalling the DCR assumption under which the DJ scheme is proven secure.

Definition A.1 (Decisional Composite Residuosity (DCR) assumption [Pai99]). *We say that security of the DCR assumption holds if there exists a PPT algorithm RSAsample that on input a security parameter λ , outputs a pair $(N, \phi(N))$ where N is a 2λ -bits integer, ϕ denotes Euler's totient function; such that $\gcd(\phi(N), N) = 1$ and such that for all polynomial $\zeta(\cdot)$, the following ensembles are computationally indistinguishable:*

$$\begin{aligned} \{\mathcal{D}_\lambda^0\}_{\lambda \in \mathbb{N}} &= \left\{ (N, \phi(N)) \leftarrow \text{RSAsample}(1^\lambda); r \leftarrow_{\mathbb{R}} \mathbb{Z}_M : r^{N^{\zeta(\lambda)}} \in \mathbb{Z}_{N^{\zeta(\lambda)+1}} \right\}_{\lambda \in \mathbb{N}} \\ \{\mathcal{D}'_\lambda\}_{\lambda \in \mathbb{N}} &= \left\{ (N, \phi(N)) \leftarrow \text{RSAsample}(1^\lambda); u \leftarrow_{\mathbb{R}} \mathbb{Z}_{N^{\zeta(\lambda)+1}} : u \in \mathbb{Z}_{N^{\zeta(\lambda)+1}} \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

We say the DCR assumption holds if 1-security of the LWE assumption holds, and that the subexponential DCR assumption holds if there exists $\varepsilon > 0$ such that 2^{λ^ε} -security of the DCR assumptions holds.

As explained in [DJ01] in further details, the algorithm $\text{RSAsample}(1^\lambda)$ samples two safe primes p, q of λ bits each, and compute the RSA modulus $N = pq$.

We will show how the DJ encryption scheme satisfies our notion of a hintable packed LHE (actually even without any packing):

Theorem A.2. *Assume (subexponential) security of the DCR assumption. Then, for every polynomial ℓ_1 , there exists a (subexponentially) secure $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE, where $h(\lambda) = 2\lambda$, $\ell_2(\lambda) = 1$, and $\alpha(\lambda) = 0$.*

A.1.1 The DJ scheme

Given a polynomial $\ell_1(\cdot)$, we recall the LHE from [DJ01] when operating on plaintext of size ℓ_1 (recall that $\ell_2 = 1$ so there is no packing). That is, the scheme is parameterized by a polynomial ℓ_1 , and we call it DJ_{ℓ_1} . For this scheme, the hint is simply the randomness used to encrypt a message, and noisy and normal encryptions behave in the same way:

- $\text{CRSgen}(1^\lambda)$:

It simply outputs $\text{crs} = 1^\lambda$, i.e. there is no proper crs for that scheme.

- $\text{Gen}(\text{crs})$:

Given $\text{crs} = 1^\lambda$, it uses the sampling algorithm from Definition A.1, $(M, \phi(M)) \leftarrow \text{RSAsample}(1^\lambda)$, where $M \in \mathbb{N}$ is a 2λ -bit modulus, ϕ denotes Euler's totient function, and we have $\gcd(\phi(M), M) = 1$. Then it chooses a polynomial $\zeta(\cdot)$ such that $2^{\ell_1(\lambda)+2\lambda} > N \geq 2^{\ell_1(\lambda)}$, where $N = M^{\zeta(\lambda)}$. For simplicity of the notations, we write $\zeta = \zeta(\lambda)$. It sets $\text{pk} = (N, \zeta)$, $\text{sk} = \text{td} = \phi(M)$ and outputs $(\text{pk}, \text{sk}, \text{td})$. The plaintext space is $\mathbb{Z}_N = \mathbb{Z}_{M^\zeta}$. The randomness space for Enc is \mathbb{Z}_M^* , the ciphertext space is $\mathbb{Z}_{M^\zeta}^*$, and the function space is \mathbb{Z}_N , that is $t = N$.

- $\text{Enc}_{\text{pk}}(\mathbf{x})$:

Given the public pk , a vector $\mathbf{x} \in \mathbb{Z}_N^\nu$, for all $i \in [\nu]$, it samples $r_i \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*$ and compute $\text{ct}_i = r_i^{M^\zeta} \cdot (1 + M)^{x_i} \in \mathbb{Z}_{M^{\zeta+1}}^*$. It outputs the ciphertext $\text{ct} = (\text{ct}_1, \dots, \text{ct}_\nu)$.

- Enc_{pk}^{*}(m):

Given the public pk , a message $m \in \mathbb{Z}_N$, it samples $r \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*$ and outputs the noisy ciphertext $\text{ct}^* = r^{M^\zeta} \cdot (1 + M)^m \in \mathbb{Z}_{M^{\zeta+1}}^*$.

- Eval(pk, ct, ct^{*}, y):

Given as input the public key pk , ciphertext $\text{ct} \in \mathbb{Z}_{M^{\zeta+1}}^{*\nu}$, noisy ciphertext $\text{ct}^* \in \mathbb{Z}_{M^{\zeta+1}}^*$, and a vector $\mathbf{y} \in [-1, 1]^\nu$, it outputs the evaluated ciphertext $\text{ct}^* \cdot \prod_{i \in [\nu]} \text{ct}_i^{y_i} \in \mathbb{Z}_{M^{\zeta+1}}^*$, where \cdot denotes the integer multiplication in $\mathbb{Z}_{M^{\zeta+1}}^*$.

- SecHint(td, ct^{*}):

Given the secret key td and a noisy ciphertext $\text{ct}^* \in \mathbb{Z}_{M^{\zeta+1}}^*$, it computes $d = \text{ct} \bmod M$. Since $\gcd(M^\zeta, \phi(M)) = 1$, it can compute $M^{-\zeta} \in \mathbb{Z}$ such that $M^\zeta \cdot M^{-\zeta} = 1 \bmod \phi(M)$. It outputs the hint $d^{M^{-\zeta}} \in \mathbb{Z}_M^*$.

- PubHint(pk, r):

Given the public key pk and some randomness $r \in \mathbb{Z}_M^*$, it outputs the hint $\rho = r$.

- Rec(pk, ct^{*}, ρ):

Given the public key pk , a noisy ciphertext ct^* , and a hint $\rho \in \mathbb{Z}_M^*$, it computes $d = \text{ct} \cdot r^{-M^\zeta} \in \mathbb{Z}_{M^{\zeta+1}}^*$, where ρ^{-M^ζ} is the inverse of ρ^{M^ζ} in $\mathbb{Z}_{M^{\zeta+1}}^*$. Then, it applies Paillier's decryption recursively to obtain $x \in \mathbb{Z}_{M^\zeta}$. It outputs $x \in \mathbb{Z}_N$.

- Dec_{sk}(ct^{*}):

Given the secret key sk , a noisy ciphertext ct^* , it runs $\text{Ext}(\text{sk}, \text{ct}^*)$ to recover the randomness $r \in \mathbb{Z}_M^*$, then outputs $\text{Rec}(\text{pk}, \text{ct}^*, r)$.

The proof of Theorem A.2 follows from the propositions and theorem below (which demonstrate that DJ satisfies the desired properties of a hintable packed LHE, as well as security).

Proposition 14 (Linear Homomorphism). *The LHE presented above satisfies Property 4.2.*

Proof: For all $i \in [\nu]$, let ct_i be a ciphertext in the support of $\text{Enc}_{\text{pk}}(x_i)$, that is, of the form $\text{ct}_i = (r_i)^{M^\zeta} \cdot (1 + M)^{x_i} \in \mathbb{Z}_{M^{\zeta+1}}^*$, and let ct^* be a ciphertext in the support of $\text{Enc}_{\text{pk}}^*(x^*)$, that is, of the form $\text{ct}^* = r^{M^\zeta} \cdot (1 + M)^{x^*} \in \mathbb{Z}_{M^{\zeta+1}}^*$. For all $\mathbf{y} \in \{0, 1\}^\nu$, the evaluated ciphertext $\text{ct}_{\mathbf{y}}$ is of the form:

$$\left(r \prod_{i \in [\nu]} r_i^{y_i} \right)^{M^\zeta} \cdot (1 + M)^{x^* + \sum_{i \in [\nu]} x_i y_i} \in \mathbb{Z}_{M^{\zeta+1}}^*$$
, which is in the support of $\text{Enc}_{\text{pk}}^*(x^* + \sum_{i \in [\nu]} x_i y_i)$. \square

Proposition 15 (0-approximate correctness of secret hints). *The LHE presented above satisfies 0-approximate correctness, as defined in Property 4.3.*

Proof: Let $\text{ct}^* = r^{M^\zeta} \cdot (1 + M)^{x^*} \in \mathbb{Z}_{M^{\zeta+1}}^*$, where $x^* \in \mathbb{Z}_{M^\zeta}$ is the message and $r \in \mathbb{Z}_M^*$ is the randomness used to produce the ciphertext. We have $\text{ct}^* \bmod M = r^{M^\zeta \bmod \phi(M)} \in \mathbb{Z}_M^*$. Since $\gcd(\phi(M), M^\zeta) = 1$, there is $M^{-\zeta} \in \mathbb{Z}$ such that $M^\zeta \cdot M^{-\zeta} = 1 \bmod \phi(M)$. Thus, the algorithm $\text{SecHint}(\text{td}, \text{ct})$ outputs $r \in \mathbb{Z}_M^*$. That is, the hint output by SecHint is the randomness used to produce ct^* .

The algorithm $\text{Rec}(\text{pk}, \text{ct}^*, r)$ computes $d = \text{ct}^* \cdot r^{-M^\zeta} = (1 + M)^{x^*} \in \mathbb{Z}_{M^{\zeta+1}}^*$ where r^{-M^ζ} is the inverse of r^{M^ζ} in $\mathbb{Z}_{M^{\zeta+1}}^*$. Then it applies Paillier's decryption recursively to obtain $x^* \in \mathbb{Z}_{M^\zeta}$. It outputs x^* . \square

Proposition 16 (Equivalence between public hints and secret hints). *The LHE presented above satisfies Property 4.4.*

Proof: As seen in the proof of Proposition 15, the hint recovered by SecHint given the trapdoor td and a ciphertext ct^* is the randomness r used by Enc^* to produce ct^* , which is also what $\text{PubHint}(\text{pk}, r)$ outputs. \square

Proposition 17 (0-approximate correctness). *The LHE presented above satisfies 0-approximate correctness, as defined in Property 4.1.*

Proof: This directly follows from the 0-approximate correctness of the secret hints (Property 4.3), since decryption first computes the hint using the secret key, then recovers the message using the hint. \square

Proposition 18 (h -succinctness of hints). *The LHE presented above satisfies $h(\lambda) = 2\lambda$ -succinctness.*

Proof: For all $\lambda \in \mathbb{N}$, for all pairs (pk, sk) in the support of $\text{Gen}(1^\lambda)$ that define the message space \mathbb{Z}_N , for all $x^* \in \mathbb{Z}_N$, all ciphertext ct^* in the support of $\text{Enc}_{\text{pk}}^*(x^*)$, we have: all hints ρ in the support of $\text{SecHint}(\text{sk}, \text{ct}^*)$ are in \mathbb{Z}_M^* , where M is an RSA modulus of size at most 2λ bits. \square

Proposition 19 (Weak circuit privacy). *The LHE presented above satisfies Property 4.6 (weak circuit privacy).*

Proof: For any message $\mathbf{x} = (x_1, \dots, x_\nu) \in \mathbb{Z}_N^\nu$, $x^* \in \mathbb{Z}_N$, $\mathbf{y} \in \{0, 1\}^\nu$, we aim at proving that following distributions are identical:

$$\mathcal{D}_0 : \left\{ \begin{array}{l} \forall i \in [\nu], r_i \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*, \text{ct}_i = r_i^{M^\zeta} \cdot (1 + M)^{x_i}, r \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*, \text{ct}^* = r^{M^\zeta} \cdot (1 + M)^{x^*} \\ \text{ct}_{\mathbf{y}} = (r \prod_i r_i^{y_i})^{M^\zeta} \cdot (1 + M)^{x^* + \mathbf{x}^\top \mathbf{y}} : ((\text{ct}_i)_i, \text{ct}^*, \text{ct}_{\mathbf{y}}) \end{array} \right\}$$

$$\mathcal{D}_1 : \left\{ \begin{array}{l} \forall i \in [\nu], r_i \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*, \text{ct}_i = r_i^{M^\zeta} \cdot (1 + M)^{x_i}, r \leftarrow_{\mathbb{R}} \mathbb{Z}_M^*, \text{ct}^* = (r / \prod_i r_i^{y_i})^{M^\zeta} \cdot (1 + M)^{x^*} \\ \text{ct}_{\mathbf{y}} = r^{M^\zeta} \cdot (1 + M)^{x^* + \mathbf{x}^\top \mathbf{y}} : ((\text{ct}_i)_i, \text{ct}^*, \text{ct}_{\mathbf{y}}) \end{array} \right\}.$$

This relies on the fact that for all $i \in [\nu]$, all $r_i \in \mathbb{Z}_M^*$, all $y_i \in \mathbb{Z}_{M^\zeta}$, the following distributions are identical: $\mathcal{D}'_0 = \{r \leftarrow_{\mathbb{R}} \mathbb{Z}_M^* : ((r_i)_i, r, r \cdot \prod_i r_i^{y_i})\}$ and $\mathcal{D}'_1 = \{r \leftarrow_{\mathbb{R}} \mathbb{Z}_M^* : ((r_i)_i, r / (\prod_i r_i^{y_i}), r)\}$. The distribution \mathcal{D}'_0 corresponds to \mathcal{D}_0 , whereas \mathcal{D}'_1 corresponds to \mathcal{D}_1 . \square

Proposition 20 (Density of the noisy ciphertexts). *The LHE presented above satisfies Property 4.7 (density of the noisy ciphertexts).*

Proof: One can sample a uniform random value $u \leftarrow_{\mathbb{R}} \mathbb{Z}_{M^{\zeta+1}}$ from $\lceil \log(M^{\zeta+1}) \rceil$ random bits. The random value $u \in \mathbb{Z}_{M^{\zeta+1}}$ can be written $u = r^{M^\zeta} \cdot (1 + M)^x$ where $x \in \mathbb{Z}_{M^\zeta}$ and $r \in \mathbb{Z}_M^*$ with probability $1 - \frac{\phi(N)}{N} > 1 - \frac{3}{2^\lambda}$ over the choice of $u \leftarrow_{\mathbb{R}} \mathbb{Z}_{M^{\zeta+1}}$. \square

Theorem A.3 (Security [DJ01]). *Assuming the (subexponential) DCR assumption, the DJ scheme is (subexponentially) secure.*

A.2 Instantiation with DJ LHE

Now we instantiate the modular XiO construction described in Section 5.2 with the DJ LHE presented in Section A.1 that is parameterized by a polynomial ℓ_1 ; we denote it by DJ_{ℓ_1} . For all polynomials δ , we denote by GSW_δ the GSW FHE scheme for depth δ circuits.

By combining Theorem A.3 (i.e. security of DJ_{ℓ_1} for all polynomials ℓ_1 under DCR) and Theorem 3.5 (i.e. SRL-security of the GSW_δ for all polynomials δ under LWE), we get:

Lemma A.1. *Assume the (subexponential) DCR and the (subexponential) LWE assumptions hold. Assume further that for all polynomials δ and ℓ_1 , the $2\text{CIRC}^{\text{O}_{\text{SRL}}}$ assumption holds w.r.t. GSW_δ and DJ_{ℓ_1} . Then, for all polynomials δ and ℓ_1 , (subexponential) 2-circular SRL security holds w.r.t. GSW_δ and DJ_{ℓ_1} .*

By combining Lemma A.1 above with Theorem A.2, which states that for all polynomials ℓ_1 , DJ_{ℓ_1} is an $(\ell_1, \ell_2, h, \alpha)$ -hintable packed LHE with $\ell_2(\lambda) = 1$, $h(\lambda) = 2\lambda$ and $\alpha(\lambda) = 0$, together with Theorem 5.2, we get:

Theorem A.4. *Assume the (subexponential) DCR and (subexponential) LWE assumptions hold. Assume further that for all polynomials δ , ℓ_1 , the $2\text{CIRC}^{\text{O}_{\text{SRL}}}$ assumption holds w.r.t. GSW_δ and DJ_{ℓ_1} . Then (subexponentially-secure) XiO for $\text{P}^{\text{log}}/\text{poly}$ exists.*

Finally, combining Theorem A.4 with Theorem 2.4 (i.e., $i\text{O}$ from XiO and LWE) yields:

Theorem A.5. *Assume the subexponential DCR and subexponential LWE assumptions hold. Assume further that for all polynomials δ and ℓ_1 the $2\text{CIRC}^{\text{O}_{\text{SRL}}}$ holds w.r.t. GSW_δ and DJ_{ℓ_1} . Then subexponentially-secure $i\text{O}$ for P/poly exists.*

B Concrete Assumption

In Theorem 5.2, we show how to build XiO whose security relies on the 2-circular SRL security between the GSW FHE and an LHE with special properties. As mentioned in Section 3.1.3, it actually suffices to assume the following variant of 2-circular SRL security (that is weaker than that of definition Definition 3.3).

Definition B.1 (One-shot SRL circular security). *We say that one-shot SRL 2-circular security holds w.r.t. an FHE scheme $\mathcal{FHE} = (\text{CRSgen}, \text{Gen}, \text{Enc}, \text{Enc}^*, \text{Eval}', \text{Eval}_{\text{rand}}, \text{Dec})$ that satisfies batch correctness (as per Definition 3.1) and randomness homomorphism (as per Definition 3.2) and a PKE scheme $\overline{\text{PKE}}$ if for all stateful nuPPT adversaries \mathcal{A} , there exists some negligible function $\mu(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\Pr[\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{FHE}} = 1] \leq 1/2 + \mu(\lambda)$, where the experiment $\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{FHE}}$ is defined as follows:*

$$\text{Exp}_{\lambda, \mathcal{A}}^{\mathcal{FHE}} = \left\{ \begin{array}{l} (\mathbf{m}^0, \mathbf{m}^1, q) \leftarrow \mathcal{A}(1^\lambda), \overline{\text{crs}} \leftarrow \overline{\text{CRSgen}}(1^\lambda) \\ (\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \text{Gen}(\overline{\text{crs}}), (\text{pk}, \text{sk}) \leftarrow \text{Gen}(\overline{\text{pk}}), b \leftarrow \{0, 1\} \\ \mathbf{r} \leftarrow_{\mathcal{R}} \mathcal{R}^{|\overline{\text{sk}}||\mathbf{m}^b|}, \text{ct} = \text{Enc}_{\text{pk}}(\overline{\text{sk}}||\mathbf{m}^b; \mathbf{r}), \overline{\text{ct}} \leftarrow \overline{\text{Enc}}_{\overline{\text{pk}}}(\text{sk}) \\ \forall i \in [q(\lambda)] : r_i^* \leftarrow_{\mathcal{R}} \mathcal{R}^*, \text{ct}_i = \text{Enc}_{\text{pk}}^*(\mathbf{0}; r_i^*) \\ (f_i, \alpha_i)_{i \in [q(\lambda)]} \leftarrow \mathcal{A}(\text{pk}, \overline{\text{pk}}, \text{ct}, \overline{\text{ct}}, (\text{ct}_i)_{i \in [q(\lambda)]}) \\ \forall i \in [q(\lambda)], \mathbf{r}_{f_i} = \text{Eval}_{\text{rand}}(\text{pk}, f_i, \mathbf{r}, \overline{\text{sk}}||\mathbf{m}^b), \text{leak}_i = r_i^* - \mathbf{r}_{f_i} \\ b' \leftarrow \mathcal{A}((\text{leak}_i)_{i \in [q(\lambda)]}) \\ \text{Return } 1 \text{ if } |\mathbf{m}^0| = |\mathbf{m}^1|, b' = b \text{ and for all } i \in [q(\lambda)], f_i(\overline{\text{sk}}||\mathbf{m}^b) = \alpha_i; 0 \text{ otherwise.} \end{array} \right\}$$