

Factoring and Pairings are not Necessary for iO: Circular-Secure LWE Suffices

Zvika Brakerski¹, Nico Döttling², Sanjam Garg³, and Giulio Malavolta⁴

¹Weizmann Institute of Science

²CISPA Helmholtz Center for Information Security

³UC Berkeley

⁴UC Berkeley & Carnegie Mellon University

Abstract

We construct indistinguishability obfuscation (iO) solely under circular-security properties of encryption schemes based on the Learning with Errors (LWE) problem. Circular-security assumptions were used before to construct (non-leveled) fully-homomorphic encryption (FHE), but our assumption is stronger and requires circular randomness-leakage-resilience. In contrast with prior works, this assumption can be conjectured to be post-quantum secure; yielding the first provably secure iO construction that is (plausibly) post-quantum secure.

Our work is a variant on a beautiful recent work by Gay and Pass [ePrint 2020] who showed a way to remove the heuristic step from the homomorphic-encryption based iO approach of Brakerski, Döttling, Garg, and Malavolta [EUROCRYPT 2020]. They thus obtain a construction proved secure under circular security of natural homomorphic encryption schemes — specifically, they use homomorphic encryption schemes based on LWE and DCR, respectively. In this work, we replace the DCR-based encryption with an LWE-based one and thus obtain a result solely from the circular security of LWE-based encryption schemes. Our circular security assumption is the same as in the prior work and refers to leakage on encryption randomness in the presence of key-cycles.

1 Introduction

The goal of program obfuscation [Had00, BGI⁺01] is to transform an arbitrary circuit C into an unintelligible but functionally equivalent circuit \tilde{C} . The early works on the topic casted doubts that *general purpose* obfuscation may not be cryptographically feasible. Thus, research on this topic focused on realizing obfuscation for special functions. However, somewhat surprisingly, it was shown that general purpose obfuscation is indeed possible. In particular, Garg et al. [GGH13a, GGH⁺13b] showed a cryptographic general purpose indistinguishability obfuscator (iO), which loosely speaking requires that the obfuscations of two circuits C_0 and C_1 that have identical input output behavior are computationally indistinguishable. The versatility of this seemingly weak notion iO has enabled numerous new applications in cryptography (e.g. [SW14, GGHR14, BZ14] just to name a few). Furthermore, tremendous body of work has been devoted to constructing secure realization and understanding the assumption behind them.

The first realizations of obfuscation relied on a new algebraic object called multilinear maps [GGH13a, CLT13, GGH15], which had only recently been constructed. Furthermore, the security of these objects relied on new (and poorly understood) computational intractability assumptions. In fact, several attacks on multilinear map candidates [CHL⁺15, HJ16] and on obfuscation constructions based on [MSZ16, CGH17] multilinear maps were demonstrated. To defend against these attacks, several safeguards were (e.g., [GMM⁺16, CVW18, MZ18, BGMZ18]) proposed to defend against these attacks. Even with these heuristic safeguards,

all but the schemes based on the Gentry et al. [GGH15] multilinear maps are known to be broken against quantum adversaries.

Towards the goal of avoiding heuristics and obtaining provably secure constructions, substantial effort was made towards obtaining obfuscation while minimizing (with later the hope of removing) the use of multilinear maps [Lin16, LV16, AS17, Lin17, LT17]. These efforts paid off and constructions of obfuscation replacing the use of multilinear maps with bilinear maps [Agr19, JLMS19, AJL⁺19] were recently obtained. However, these bilinear map based constructions, some of which are still conjectured to be secure, additionally relied on certain pseudorandom objects with novel security properties. In a beautiful recent work by Jain, Lin and Sahai [JLS20], this limitation was removed — specifically, they obtained iO assuming (sub-exponential security of) LWE and LPN, in addition to bilinear pairings. Here again, unfortunately, the use of the pairings makes this construction insecure against quantum adversaries.

Towards the same goal but following a completely different approach, Brakerski et al. [BDGM20] showed a construction of iO obtained by combining certain natural *homomorphic* encryption schemes. However, their construction was *heuristic* in the sense that security argument could only be presented in the random oracle model. In a beautiful recent work, Gay and Pass [GP20] showed a way to remove the heuristic step. They obtain a construction proved secure under circular security of natural homomorphic encryption schemes — specifically, they use homomorphic encryption schemes based on LWE and DCR, respectively. More specifically, their construction assumes sub-exponential security of (i) the Learning with Error (LWE) assumption, (ii) the Decisional Composite Residuosity (DCR) assumption, and (iii) the shielded leakage resilience (SRL) security of the GSW encryption scheme [GSW13] in the presence of a key-cycle with the Damgård-Jurik encryption scheme [DJ01]. This construction is also insecure against quantum attackers because of the use of the Damgård-Jurik encryption scheme [DJ01].

In this work, we ask:

Can we realize provably secure constructions of iO based solely on hard problems in lattices?

Our results. In this work, we obtain a general purpose iO construction based solely from the circular security of LWE-based encryption schemes. In other words, we remove the need for DCR-based encryption from the construction of Gay and Pass [GP20] and replace it with an LWE-based encryption satisfying similar properties.

As a result, the security of our construction relies on a circular-security assumption. Namely, the assumption that the security properties of an encryption scheme are preserved even when given an encryption of the secret key itself, or, as in this case, given a “key cycle” where the key of the first scheme is encrypted using the second scheme, and the key of the second scheme is encrypted using the first scheme. Circular security assumptions are used in the literature. Notably, they are known to imply such primitives as “unrestricted” (non-leveled) fully homomorphic encryption (FHE) schemes via Gentry’s bootstrapping paradigm [Gen09].

Concretely, the circular assumption made in [GP20], and thus also in this work is that a scheme (in particular a leveled FHE scheme) which enjoys the property that security is maintained even given some particular kind of leakage on the randomness of the ciphertext. This property is called Shielded Randomness Leakage (SLR) in [GP20]. Indeed, standard GSW encryption [GSW13] satisfies this notion, and the assumption we make is that SLR security holds even when given a key-cycle connecting GSW to another encryption scheme. The second scheme in [GP20] was based on DCR, whereas in our case it is also based on LWE. While this assumption falls into the category of “circular security assumptions”, similarly to the ones that underly bootstrapping in FHE, the concrete assumption is quite different. While in the FHE setting it was only assumed that (standard) CPA security is preserved given a key cycle, here we assume that the stronger SRL property remains intact.

We note that if we further assume that circular security also maintains post-quantum security, then our assumption becomes post-quantum secure; yielding the first provably secure iO construction that is post-quantum secure.

Let us now state the result somewhat more formally. Assuming (sub-exponential) (i) quantum hardness of the LWE problem, and (ii) the SRL security of GSW in the presence of a 2-key cycle with dual-Reggev,

we obtain the first provably secure construction of post-quantum iO from the same kind of assumption as are currently known to imply (unlevelled) fully-homomorphic encryption (FHE).

At a technical level, our construction is obtained by realizing a packed version of the dual-Regev encryption which has succinct randomness and an alternative encryption mode where the ciphertexts are “almost-everywhere” dense. These additional properties of our variant of dual-Regev allow us to replace the use of the Damgård-Jurik encryption scheme [DJ01] in Gay and Pass [GP20] with an LWE based encryption scheme.

2 Preliminaries

We denote by $\lambda \in \mathbb{N}$ the security parameter. We say that a function negl is negligible if it vanishes faster than any polynomial. Given a set S , we denote by $s \leftarrow_{\$} S$ the uniform sampling from S . We say that an algorithm is PPT if it can be implemented by a probabilistic machine running in time $\text{poly}(\lambda)$. Matrices are denoted by \mathbf{M} and vectors are denoted by \mathbf{v} . We recall the smudging lemma [AIK11, AJL⁺12].

Lemma 1 (Smudging) *Let $B_1 = B_1(\lambda)$ and $B_2 = B_2(\lambda)$ be positive integers and let $e_1 \in [-B_1, B_1]$ be a fixed integer. Let $e_2 \leftarrow_{\$} [-B_2, B_2]$ chosen uniformly at random. Then the distribution of e_2 is statistically indistinguishable from that of $e_2 + e_1$ as long as $B_1/B_2 = \text{negl}(\lambda)$.*

2.1 Indistinguishability Obfuscation

We recall the notion of indistinguishability obfuscation (iO) from [GGH⁺13b].

Definition 2.1 (Indistinguishability Obfuscation) *A PPT machine iO is an indistinguishability obfuscator for a circuit class $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ if the following conditions are satisfied:*

(Functionality) *For all $\lambda \in \mathbb{N}$, all circuit $C \in \mathcal{C}_\lambda$, all inputs x it holds that*

$$\Pr \left[\tilde{C}(x) = C(x) \mid \tilde{C} \leftarrow \text{iO}(C) \right] = 1.$$

(Indistinguishability) *For all polynomial-size distinguishers D there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all pairs of circuit $(C_0, C_1) \in \mathcal{C}_\lambda$ such that $|C_0| = |C_1|$ and $C_0(x) = C_1(x)$ on all inputs x , it holds that*

$$|\Pr [1 = D(\text{iO}(C_0))] - \Pr [1 = D(\text{iO}(C_1))]| = \text{negl}(\lambda).$$

XiO. We recall a theorem from Lin et al. [LPST16], which is going to be useful for our work.

Theorem 2.2 ([LPST16]) *Assuming sub-exponentially hard LWE, if there exists a sub-exponentially secure indistinguishability obfuscator (with pre-processing) for $\text{P}^{\text{log}}/\text{poly}$ with non-trivial efficiency, then there exists an indistinguishability obfuscator for P/poly with sub-exponential security.*

Here $\text{P}^{\text{log}}/\text{poly}$ denotes the class of polynomial-size circuits with inputs of length $\eta = O(\log(\lambda))$ and by non-trivial efficiency we mean that the size of the obfuscated circuit is bounded by $\text{poly}(\lambda, |C|) \cdot 2^{\eta \cdot (1-\varepsilon)}$, for some constant $\varepsilon > 0$. Note that the above theorem poses no restriction on the runtime of the obfuscator. Furthermore, the theorem allows the obfuscator to access a large uniform random string (the pre-processing) of size even larger than the truth table of the circuit.

2.2 Learning with Errors

Definition 2.3 (Learning with Errors) *The LWE problem is parametrized by a modulus q , positive integers n, m and an error distribution χ . The LWE problem is hard if for all polynomial-size distinguishers D there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ it holds that*

$$|\Pr [1 = D(\mathbf{A}, \mathbf{s}^\top \cdot \mathbf{A} + \mathbf{e})] - \Pr [1 = D(\mathbf{A}, \mathbf{u})]| = \text{negl}(\lambda).$$

where \mathbf{A} is chosen uniformly from $\mathbb{Z}_q^{n \times m}$, \mathbf{s} is chosen uniformly from \mathbb{Z}_q^n , \mathbf{u} is chosen uniformly from \mathbb{Z}_q^m and \mathbf{e} is chosen from χ^m .

As shown in [Reg05, PRS17], for *any* sufficiently large modulus q the LWE problem where χ is a discrete Gaussian distribution with parameter $\sigma = \alpha q \geq 2\sqrt{n}$ (i.e. the distribution over \mathbb{Z} where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\alpha)$ in *worst case* dimension n lattices. We refer to $\alpha = \sigma/q$ as the *modulus-to-noise* ratio, and by the above this quantity controls the hardness of the LWE instantiation. Hereby, LWE with polynomial α is (presumably) harder than LWE with super-polynomial or sub-exponential α . We can truncate the discrete Gaussian distribution χ to $\sigma \cdot \omega(\sqrt{\log(\lambda)})$ while only introducing a negligible error. Consequently, we omit the actual distribution χ but only use the fact that it can be bounded by a (small) value B .

Micciancio and Peikert [MP12], provide an algorithm to sample uniformly random LWE matrices together with an inversion trapdoor that allows for efficient LWE inversion. That is, there exist efficient algorithms GenTrap and Invert, such that $\text{GenTrap}(m, n, q)$ samples a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ and a *trapdoor* τ , such that

- The marginal distribution of \mathbf{A} is statistically close to uniform.
- For any $\mathbf{s} \in \mathbb{Z}_q^n$ and any $\mathbf{e} \in \mathbb{Z}_q^m$ with $\|\mathbf{e}\| < q/T$ (for some $T = \text{poly}(\lambda)$) it holds that $\text{Invert}(\tau, \mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) = \mathbf{s}$.

2.3 Public-Key Encryption

We recall the definition of public key encryption in the following.

Definition 2.4 (Public-Key Encryption) *A homomorphic encryption scheme consists of the following efficient algorithms.*

KeyGen(1^λ): *On input the security parameter 1^λ , the key generation algorithm returns a key pair (sk, pk) .*

Enc(pk, m): *On input a public key pk and a message m , the encryption algorithm returns a ciphertext c .*

Dec(sk, c): *On input the secret key sk and a ciphertext c , the decryption algorithm returns a message m .*

Definition 2.5 (Correctness) *A public-key encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is correct if for all $\lambda \in \mathbb{N}$, all messages m , all (sk, pk) in the support of $\text{KeyGen}(1^\lambda)$, and all c in the support of $\text{Enc}(\text{pk}, m)$ it holds that*

$$\text{Dec}(\text{sk}, c) = m.$$

We define a weak notion of security (implied by the standard semantic security [GM82]) which is going to be more convenient to work with.

Definition 2.6 (Semantic Security) *A public key encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Dec})$ is semantically secure if for all PPT distinguishers \mathcal{D} there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all pairs of message (m_0, m_1) , it holds that*

$$|\Pr[1 = \text{D}(\text{pk}, \text{Enc}(\text{pk}, m_0))] - \Pr[1 = \text{D}(\text{pk}, \text{Enc}(\text{pk}, m_1))]| = \text{negl}(\lambda)$$

where $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)$.

Circular Security. We say that two encryption schemes $(\text{KeyGen}_0, \text{Enc}_0, \text{Dec}_0)$ and $(\text{KeyGen}_1, \text{Enc}_1, \text{Dec}_1)$ form a key cycle if the distinguisher is given a cross-encryption of the secret keys $\text{Enc}(\text{pk}_1, \text{sk}_0)$ and $\text{Enc}(\text{pk}_0, \text{sk}_1)$. We say that the scheme is 2-circular secure if semantic security is retained in the presence of such a cycle.

SRL Security. Shielded randomness leakage (SRL) security says that the scheme is semantically secure even in the presence of an oracle that leaks some information about the randomness for evaluated ciphertext for adversarially chosen function for which the adversary knows the output. We refer the reader to [GP20] for a precise definition. In [GP20] it is shown that the GSW encryption scheme [GSW13] satisfies such a notion if the (plain) LWE problem is hard.

3 Packed Encryption from LWE

Here we describe the packed version of dual Regev. We denote by $n = n(\lambda)$ the lattice dimensions (which we treat as the security parameter), by $q = q(\lambda)$ the modulus (which we assume for simplicity to be even), and by $k = k(\lambda)$ the expansion factor. We set $m \geq n \log(q)$. Let `TrapGen` and `Invert` be the Trapdoor generation and inversion algorithms of [MP12].

`KeyGen`($1^n, 1^k$): Sample a uniform $n \times m$ matrix $\mathbf{A} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{n \times m}$ together with a short trapdoor τ via the trapdoor sampling algorithm $(\mathbf{A}, \tau) \leftarrow_{\mathcal{S}} \text{TrapGen}(n, m, q)$, sample uniformly random vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \leftarrow_{\mathcal{S}} \mathbb{Z}_q^m$. The public key is set to

$$(\mathbf{A}, \mathbf{b}_1, \dots, \mathbf{b}_k)$$

and the secret key is the trapdoor τ .

`Enc`(`pk`, (m_1, \dots, m_k)): To encrypt a k -bit message, sample a uniform randomness $\mathbf{r} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^m$ and a $(k+1)$ -dimensional noise vector $\mathbf{e} \leftarrow_{\mathcal{S}} \chi^{k+1}$ and return

$$c = (\mathbf{A}\mathbf{r} + e_0, \mathbf{b}_1\mathbf{r} + e_1 + m_1, \dots, \mathbf{b}_k\mathbf{r} + e_k + m_k).$$

`Dec`(`sk` = $\tau, c = (c_0, c_1, \dots, c_k)$): Use τ to recover \mathbf{r} from \mathbf{c}_0 via $\mathbf{r} = \text{Invert}(\tau, \mathbf{A}, \mathbf{c}_0)$. Compute the m_i via $m_i = \text{MSB}(c_i - \mathbf{b}_i \cdot \mathbf{r})$. Output (m_1, \dots, m_k) .

For convenience we also define an alternative encryption algorithm in the following.

`DenseEnc`(`pk`): Sample a uniform randomness $\mathbf{r} \leftarrow_{\mathcal{S}} \mathbb{Z}_q^m$ and a noise term $e_0 \leftarrow_{\mathcal{S}} \chi$ and return

$$c = (\mathbf{A}\mathbf{r} + e_0, \mathbf{b}_1\mathbf{r} + u_1, \dots, \mathbf{b}_k\mathbf{r} + u_k)$$

where $(u_1, \dots, u_k) \leftarrow_{\mathcal{S}} \mathbb{Z}_q$.

We highlight two facts about this algorithm that are going to be important for our later construction: (i) The decryption algorithm works for both `Enc` and `DenseEnc` algorithm, in fact the scheme satisfies perfect correctness in both cases. (ii) The domain of the elements (c_1, \dots, c_k) is *dense*, i.e. the support of the scheme spans the whole vector space \mathbb{Z}_q^k . Since the element \mathbf{c}_0 is small (by setting k large enough), we refer to such a property as “almost-everywhere” density.

3.1 Analysis

Here we argue that the scheme as described above satisfies a few properties of interest.

Semantic Security. First we argue that the scheme satisfies a strong form of semantic security, i.e. the honestly computed ciphertexts are computationally indistinguishable from uniform vectors in \mathbb{Z}_q^k .

Theorem 3.1 (Semantic Security) *If the LWE assumption holds, then the ciphertexts then for all $\lambda \in \mathbb{N}$ and all (sk, pk) in the support of `KeyGen` the following distributions are computationally indistinguishable*

$$\text{Enc}(\text{pk}, m) \approx u.$$

where $u \leftarrow_{\mathcal{S}} \mathbb{Z}_q^{k+1}$.

Proof: The security of the scheme follows routinely by an application of the Leftover-Hash Lemma [HILL99] and by k invocations of the LWE assumptions. \square

Randomness Succinctness. Here we show that our scheme satisfies the notion of randomness succinctness which, intuitively, asks that the randomness of a ciphertext is asymptotically smaller than the message space.

Theorem 3.2 (Randomness Succinctness) *There exists a polynomial $\text{poly}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all (sk, pk) in the support of KeyGen , and all elements \mathbf{r} in the corresponding randomness space, it holds that $|\mathbf{r}| \leq \text{poly}(\lambda)$.*

Proof: The randomness is a uniform vector in \mathbb{Z}_q^m and it is in particular independent of k . \square

Linear Homomorphism. The scheme is additively homomorphic over \mathbb{Z}_q^k for a bounded amount of addition. In the following we show that it can be converted to linear homomorphism (i.e. inner product with large coefficients) by encrypting all powers of 2.

Theorem 3.3 (Linear Homomorphism) *There exists a polynomial-time algorithm InnProd such that for all $\lambda \in \mathbb{N}$, all (sk, pk) in the support of KeyGen , all k -dimensional message vectors $(\mathbf{m}_1, \dots, \mathbf{m}_\ell)$, all (c_1, \dots, c_ℓ) in the support of $(\text{Enc}(\text{pk}, \mathbf{m}_1), \dots, \text{Enc}(\text{pk}, \mathbf{m}_\ell))$ and all vectors $\mathbf{y} \in \mathbb{Z}_q^\ell$ it holds that*

$$\text{Dec}(\text{sk}, \text{InnProd}(\text{pk}, (c_1, \dots, c_\ell), \mathbf{y})) = \mathbf{y}^T(\mathbf{m}_1, \dots, \mathbf{m}_\ell).$$

Proof: It is well-known that dual Regev is (bounded) additively homomorphic and so is the packed version (over \mathbb{Z}_q^k). To compute inner-products with large coefficient, one can encrypt $(m_1, \dots, m_k) \otimes \mathbf{G}$, where \mathbf{G} is the gadget matrix [MP12] of appropriate dimensions. Inner products are then computed via multiplication with the binary decomposition of the coefficients. \square

Randomness Recovery. It is well-known that dual Regev is randomness recoverable.

Theorem 3.4 (Randomness Recoverability) *There exists a polynomial-time algorithm Ext such that for all $\lambda \in \mathbb{N}$, all (sk, pk) in the support of KeyGen , all k -dimensional messages \mathbf{m} , all randomnesses $\mathbf{r} \in \mathbb{Z}_q^m$, it holds that*

$$\text{Ext}(\text{sk}, \text{Enc}(\text{pk}, \mathbf{m}; \mathbf{r})) = \mathbf{r},$$

except with negligible probability over the additional random choices made by Enc .

Proof: The algorithm Ext recovers \mathbf{r} from a ciphertext $(\mathbf{c}_0, c_1, \dots, c_k)$ in the same way as Dec , by computing $\mathbf{r} = \text{Invert}(\tau, \mathbf{A}, \mathbf{c}_0)$. The claim follows from the correctness of the inversion procedure Invert given that $\|\mathbf{e}_0\| < \delta$, which holds with overwhelming probability. \square

Decryption with Randomness. Given the randomness \mathbf{r} , one can easily decrypt a ciphertext.

Theorem 3.5 (Decryption with Randomness) *There exists a polynomial-time algorithm Rec such that for all $\lambda \in \mathbb{N}$, all (sk, pk) in the support of KeyGen , all k -dimensional messages \mathbf{m} , all randomnesses $\mathbf{r} \in \mathbb{Z}_q^m$, it holds that*

$$\text{Rec}(\text{pk}, \mathbf{r}, \text{Enc}(\text{pk}, \mathbf{m}; \mathbf{r})) = \mathbf{m}.$$

Proof: For all $i = 1 \dots k$ compute $\mathbf{b}_i \mathbf{r}$ and round to the nearest multiple of $q/2$ to recover m_i . \square

4 Constructing XiO

In the following we outline the construction of XiO using and FHE scheme and the LHE scheme presented in this work. Since the scheme and the analysis is largely unchanged from [GP20], we only provide a high-level overview highlighting the differences. The notation is taken from [GP20] in favor of clarity of exposition.

4.1 Construction

The scheme assumes a long uniform string that is, for convenience, split in two chunks:

1. A sequence of randomization vectors for the GSW FHE scheme `FHE.PubCoin`.
2. A sequence of simulated LHE encryptions `LHE.PubCoin`.

On input the security parameter 1^λ and the circuit Π , the obfuscator proceeds as follows.

Setting the Public Keys: Sample an FHE key pair (sk, pk) and an LHE $(\bar{\text{sk}}, \bar{\text{pk}}) \leftarrow \text{KeyGen}(1^n, 1^k)$ with matching modulus q . Compute an FHE encryption $c_1 \leftarrow \text{GSWEnc}(\text{pk}, C_\Pi)$ where C_Π is the circuit that on input some index i computes the i -th block of the truth table of Π .

Compute a Key Cycle: Compute an FHE encryption of the LHE secret key $c_2 \leftarrow \text{GSWEnc}(\text{pk}, \bar{\text{sk}})$ and an LHE encryption of the FHE secret key $\bar{c} \leftarrow \text{Enc}(\bar{\text{pk}}, \text{sk})$.

Decryption Hints: For all indices $i \in \{0, 1\}^{\log(n^{1-\varepsilon})}$, for some constant ε , do the following.

Evaluate the Circuit: Homomorphically evaluate C_Π on i and let $c_{1,i}$ be the resulting ciphertext.

Compute the Encryption Header: Sample a uniform $\mathbf{r} \leftarrow_{\$} \mathbb{Z}_q^m$ and a noise term $e \leftarrow_{\$} \chi$ and return $h = \mathbf{A}\mathbf{r} + e$.

Compute the Low-Order Bits: Let the i -th block of `LHE.PubCoin` be

$$(h_1, \dots, h_k) = (\mathbf{b}_1 \mathbf{r} + u_1, \dots, \mathbf{b}_k \mathbf{r} + u_k)$$

for some $(u_1, \dots, u_k) \in \mathbb{Z}_q^k$. Compute homomorphically over c_2 the function f , which takes as input a dual Regev ciphertext (h, h_1, \dots, h_k) , computes the decryption algorithm and returns

$$(-\text{MSB}(u_1), \dots, -\text{MSB}(u_k)).$$

Note that (h, h_1, \dots, h_k) is a ciphertext in the support of the alternative encryption algorithm `DenseEnc`. Denote the resulting ciphertext by c_{MSB} .

Rerandomize the Ciphertext: Use the i -th block \mathbf{r}^* of the `FHE.PubCoin` to compute an FHE encryption of 0 and compute

$$c'_{\text{MSB}} = c_{\text{MSB}} + \text{GSWEnc}(\text{pk}, 0; \mathbf{r}^*).$$

Proxy Re-Encrypt: Combine $c_{1,i}$ and c'_{MSB} into a single FHE encryption d (by staggering the plaintexts in different bits) and compute

$$\bar{c}_i \leftarrow \text{InnerProd}(\bar{\text{pk}}, \bar{c}, d) + (h, h_1, \dots, h_k).$$

Release Hint: Release the randomness of the resulting LHE ciphertext by computing $\text{Ext}(\bar{\text{sk}}, \bar{c}_i)$.

Output: The obfuscated circuit consists of the public keys, the decryption hints, and the headers.

The obfuscated circuit is evaluated block-wise by the evaluator, who recomputes \bar{c}_i as specified above and uses the corresponding decryption hint to recover the plaintext via the `Rec` algorithm of dual Regev. Note that the decryption returns the correct output since

$$\begin{aligned}
\bar{c}_i &= \text{InnerProd}(\bar{\mathbf{pk}}, \bar{c}, d) + (h, h_1, \dots, h_k) \\
&= \text{Enc}(\bar{\mathbf{pk}}, (m_1 - \text{MSB}(u_1), \dots, m_k - \text{MSB}(u_k))) + (h, h_1, \dots, h_k) \\
&= \text{Enc}(\bar{\mathbf{pk}}, (m_1 - \text{MSB}(u_1), \dots, m_k - \text{MSB}(u_k))) + (\mathbf{Ar} + e, \mathbf{b}_1\mathbf{r} + u_1, \dots, \mathbf{b}_k\mathbf{r} + u_k) \\
&= (\mathbf{Ar}' + e, \mathbf{b}_1\mathbf{r}' + m_1 - \text{MSB}(u_1) + u_1, \dots, \mathbf{b}_k\mathbf{r}' + m_k - \text{MSB}(u_k) + u_k) \\
&= (\mathbf{Ar}' + e, \mathbf{b}_1\mathbf{r}' + m_1 + \nu_1, \dots, \mathbf{b}_k\mathbf{r}' + m_k + \nu_k)
\end{aligned}$$

where (ν_1, \dots, ν_k) are small, which is a well-formed ciphertext.

Compression is obtained by setting k to be large enough polynomial overhead dictated by the FHE encryption. Note that the size of the LHE encryption of \mathbf{sk} potentially grows with k , so its size has to be amortized by setting ε appropriately. We refer the reader to [GP20] for a concrete choice of parameters.

4.2 Analysis

In the following we analyze the security of our scheme.

Theorem 4.1 (XiO Security) *If the FHE and LHE schemes are 2-circular SRL secure, then the XiO scheme as described above is secure.*

Proof: We provide a high-level overview of the security analysis. This is mostly unchanged from [GP20], except for a few steps that we highlight.

Hybrid 0: This is the original obfuscation of the circuit Π_0 .

Hybrid 1: Here we sample c'_{MSB} as a fresh encryption of $(-\text{MSB}(u_1), \dots, -\text{MSB}(u_k))$ using randomness \mathbf{r}^* and setting the corresponding block of `FHE.PubCoin` to $\mathbf{r}^* - \mathbf{r}_c$, where \mathbf{r}_c is the randomness of c_{MSB} .

This hybrid is statistically close by the weak circuit privacy of the FHE scheme (same as in [GP20]).

Hybrid 2: Here the i -th block of `LHE.PubCoin` is computed by

$$\bar{h} = \text{DenseEnc}(\bar{\mathbf{pk}}) = (h_0, h_1, \dots, h_k),$$

that is, it holds that $h_i = \mathbf{b}_i \cdot \mathbf{r} + u_i$ is uniform in \mathbb{Z}_q . Thus the distribution is identical to that of the previous hybrid.

Hybrid 3: Here we generate \bar{c}_i as a fresh encryption of $(m_1 + u'_1 - \text{MSB}(u'_1), \dots, m_k + u'_k - \text{MSB}(u'_k))$ for uniformly random u'_j using fresh randomness and compute the encryption header together with the corresponding block of `LHE.PubCoin` as

$$\bar{c}_i = \text{InnProd}(\bar{\mathbf{pk}}, \bar{c}, d).$$

Since the $u_j - \text{MSB}(u_j)$ are uniformly random in $[-q/4, q/4]$ and the modulus-to-noise ratio for the e_i is super-polynomial, it holds that $u_j - \text{MSB}(u_j)$ and $u'_j - \text{MSB}(u'_j) + e_j$ are statistically close. It follows that hybrid 3 and hybrid 4 are statistically close.

Hybrid 4: Here we compute \bar{c}_i using fresh noise. Statistical indistinguishability follows from Lemma 1 (same as in [GP20]).

Hybrid 5: Here we switch to encrypting Π_1 instead. The computational indistinguishability follows from a reduction to the circular SRL security of the FHE and the LHE scheme (same as in [GP20]).

Hybrid 6-10: Undo all the changes except that now we encrypt Π_1 instead of Π_0 .

□

4.3 On the Assumption

When instantiating LHE with the packed version of dual Regev and the FHE with GSW, our assumption states that SRL security of GSW (which can be shown to hold in the stand alone settings) is retained in the presence of a 2-key cycle with (packed) dual Regev.

We observe that we can further modify the XiO scheme described in Section 4.1 to reduce against a *weaker* assumption, although somewhat more cumbersome to state. More specifically, instead of an encryption a trapdoor τ under the GSW key, we can simply provide the evaluator with an encryption of each randomness vector \mathbf{r} (as defined in the computation of the header). Note that this modification does not affect correctness, since the trapdoor was only used to recompute \mathbf{r} , nor succinctness, since the vectors $\mathbf{r} \in \mathbb{Z}_q^m$ are small. This modification removes the 2-key cycle although we still have a randomness-key circularity in the dependency of the two schemes.

References

- [Agr19] Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 191–225, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [AIK11] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. In Rafail Ostrovsky, editor, *52nd Annual Symposium on Foundations of Computer Science*, pages 120–129, Palm Springs, CA, USA, October 22–25, 2011. IEEE Computer Society Press.
- [AJL⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 483–501, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [AJL⁺19] Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 284–332, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.
- [AS17] Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 152–181, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 79–109, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.
- [BGI⁺01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.

- [BGMZ18] James Bartusek, Jiaxin Guan, Fermi Ma, and Mark Zhandry. Return of GGH15: Provable security against zeroizing attacks. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018: 16th Theory of Cryptography Conference, Part II*, volume 11240 of *Lecture Notes in Computer Science*, pages 544–574, Panaji, India, November 11–14, 2018. Springer, Heidelberg, Germany.
- [BZ14] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 480–499, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [CGH17] Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part III*, volume 10212 of *Lecture Notes in Computer Science*, pages 278–307, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- [CHL⁺15] Jung Hee Cheon, KyooHyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 3–12, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 476–493, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 577–607, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [DJ01] Ivan Damgård and Mats Jurik. A generalisation, a simplification and some applications of Paillier’s probabilistic public-key system. In Kwangjo Kim, editor, *PKC 2001: 4th International Workshop on Theory and Practice in Public Key Cryptography*, volume 1992 of *Lecture Notes in Computer Science*, pages 119–136, Cheju Island, South Korea, February 13–15, 2001. Springer, Heidelberg, Germany.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st Annual ACM Symposium on Theory of Computing*, pages 169–178, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17, Athens, Greece, May 26–30, 2013. Springer, Heidelberg, Germany.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual Symposium on Foundations of Computer Science*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 498–527, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany.

- [GGHR14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure MPC from indistinguishability obfuscation. In Yehuda Lindell, editor, *TCC 2014: 11th Theory of Cryptography Conference*, volume 8349 of *Lecture Notes in Computer Science*, pages 74–94, San Diego, CA, USA, February 24–26, 2014. Springer, Heidelberg, Germany.
- [GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th Annual ACM Symposium on Theory of Computing*, pages 365–377, San Francisco, CA, USA, May 5–7, 1982. ACM Press.
- [GMM⁺16] Sanjam Garg, Eric Miles, Pratyay Mukherjee, Amit Sahai, Akshayaram Srinivasan, and Mark Zhandry. Secure obfuscation in a weak multilinear map model. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B: 14th Theory of Cryptography Conference, Part II*, volume 9986 of *Lecture Notes in Computer Science*, pages 241–268, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.
- [GP20] Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. Cryptology ePrint Archive, Report 2020/1010, 2020.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [Had00] Satoshi Hada. Zero-knowledge and code obfuscation. In Tatsuaki Okamoto, editor, *Advances in Cryptology – ASIACRYPT 2000*, volume 1976 of *Lecture Notes in Computer Science*, pages 443–457, Kyoto, Japan, December 3–7, 2000. Springer, Heidelberg, Germany.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.
- [HJ16] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 537–565, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [JLMS19] Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over \mathbb{R} to build $i\mathcal{O}$. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 251–281, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [JLS20] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. Cryptology ePrint Archive, Report 2020/1003, 2020.
- [Lin16] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 28–57, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [Lin17] Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 599–629, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany.

- [LPST16] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation with non-trivial efficiency. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016: 19th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 9615 of *Lecture Notes in Computer Science*, pages 447–462, Taipei, Taiwan, March 6–9, 2016. Springer, Heidelberg, Germany.
- [LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 630–660, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany.
- [LV16] Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In Irit Dinur, editor, *57th Annual Symposium on Foundations of Computer Science*, pages 11–20, New Brunswick, NJ, USA, October 9–11, 2016. IEEE Computer Society Press.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EURO-CRYPTO 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 700–718, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.
- [MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 629–658, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- [MZ18] Fermi Ma and Mark Zhandry. The MMap strikes back: Obfuscation and new multilinear maps immune to CLT13 zeroizing attacks. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018: 16th Theory of Cryptography Conference, Part II*, volume 11240 of *Lecture Notes in Computer Science*, pages 513–543, Panaji, India, November 11–14, 2018. Springer, Heidelberg, Germany.
- [PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th Annual ACM Symposium on Theory of Computing*, pages 461–473, Montreal, QC, Canada, June 19–23, 2017. ACM Press.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th Annual ACM Symposium on Theory of Computing*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press.