Factoring and Pairings Are Not Necessary for IO: Circular-Secure LWE Suffices

Zvika Brakerski
Weizmann Institute of Science, Rehovot, Israel

Nico Döttling
CISPA Helmholtz Center for Information Security, Saarbrücken, Germany

Sanjam Garg
University of California, Berkeley, CA, USA
NTT Research, Sunnyvale, CA, USA

Giulio Malavolta
Max Planck Institute for Security and Privacy, Bochum, Germany

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 follows the high-level outline of the recent work of Gay and Pass [STOC 2021], 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 assumption of natural homomorphic encryption schemes – specifically, they use homomorphic encryption schemes based on LWE and DCR, respectively. In this work we show how to remove the DCR assumption and remain with a scheme based on the circular security of LWE alone. Along the way we relax some of the requirements in the Gay-Pass blueprint and thus obtain a scheme that is secure under a different assumption. Specifically, we do not require security in the presence of a key-cycle, but rather only in the presence of a key-randomness cycle.

An additional contribution of our work is to point out a problem in one of the building blocks used by many iO candidates, including *all* existing provable post-quantum candidates. Namely, in the transformation from exponentially-efficient iO (XiO) from Lin, Pass, Seth and Telang [PKC 2016]. We show why their transformation inherently falls short of achieving the desired goal, and then rectify this situation by showing that *shallow* XiO (i.e. one where the obfuscator is depth-bounded) does translate to iO using LWE.

2012 ACM Subject Classification Theory of computation → Cryptographic primitives

Keywords and phrases Cryptography, Obfuscation

Digital Object Identifier 10.4230/LIPIcs.ICALP.2022.28

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://eprint.iacr.org/2020/1024`

Funding *Zvika Brakerski*: Supported by the Israel Science Foundation (Grant No. 3426/21), and by the European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).
Nico Döttling: Supported by the Helmholtz Association within the project "Trustworthy Federated Data Analytics" (TFDA) (funding number ZT-I- OO1 4).
Sanjam Garg: Supported in part by DARPA under Agreement No. HR00112020026, AFOSR Award FA9550-19-1-0200, NSF CNS Award 1936826, and research grants by the Sloan Foundation, and Visa Inc. Any opinions, findings and conclusions or recommendations ex- pressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Government or DARPA.

49th International Colloquium on Automata, Languages, and Programming (ICALP 2022).
Editors: Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff;
Article No. 28; pp. 28:1–28:20
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Giulio Malavolta: The work described in this paper has been partially supported by the German Federal Ministry of Education and Research BMBF (grant 16K15K042, project 6GEM).

1 Introduction

The goal of program obfuscation [4, 28] is to transform an arbitrary circuit Π into an unintelligible but functionally equivalent circuit $\tilde{\Pi}$. The aforementioned works showed that strong simulation-based notions of obfuscation were impossible for general purpose functionalities. However, the seemingly weaker *indistinguishability obfuscation* (iO) was not ruled out by prior work (and has in fact been shown to be the same as the best possible notion of obfuscation [27]). In broad terms, iO requires that if two circuits Π_0 and Π_1 are two implementations of the same function, then their obfuscations are computationally indistinguishable.

Garg et al. [19, 21] presented the first candidate for general purpose iO, paving the way for numerous other candidates based on a variety of mathematical structures. Although iO appears to be a weak notion of security, it has been shown to be sufficient for numerous cryptographic applications, including ones that were previously not known to exist under other assumptions (see [6, 20, 41] for examples). The first realizations of obfuscation relied an a new algebraic object called multilinear maps [15, 19, 24], which had only recently been constructed. Furthermore, the security of these objects relied on new (and poorly understood) computational intractability assumptions, or more commonly on plain heuristics. In fact, several attacks on multilinear map candidates [14, 30] and on obfuscation constructions based on multilinear maps [12, 39] were demonstrated. To defend against these attacks, several safeguards have been (e.g., [5, 13, 17, 22, 38]) proposed. Even with these heuristic safeguards, all but the schemes based on the Gentry et al. [24] 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 iO while minimizing (with the ultimate goal of removing) the use of multilinear maps [3, 33, 34, 36, 37]. These efforts culminated in replacing the use of multilinear maps with just bilinear maps [1, 2, 31], together with an additional pseudorandom generators of constant locality over the integers with polynomial stretch. Very recently this last limitation was removed by Jain, Lin and Sahai [32]. Specifically, they obtained iO based on the combined (sub-exponential) hardness of the Learning with Errors problem (LWE), a large-modulus variant of the Learning Parity with Noise problem (LPN), the existence of a pseudorandom generator in NC_0, and in addition the hardness of the external Diffie-Hellman problem in bilinear groups (SXDH). We note that the use of the pairings makes these construction insecure against quantum adversaries.

A different approach towards provably secure iO, which is more relevant to this work, was presented by Brakerski et al. [9]. They showed an iO candidate that is based on combining certain natural *homomorphic* encryption schemes. However, their construction was *heuristic* in the sense that the security argument could only be presented in the random oracle model. In a recent work, Gay and Pass [23] showed a way to remove the heuristic step and instead rely on a concrete assumption. Their construction is proved secure under the circular security of natural homomorphic encryption schemes – specifically, they use homomorphic encryption schemes based on LWE and Decisional Composite Residuosity (DCR, also known as Paillier's assumption). In terms of assumptions, their construction assumes sub-exponential security of (i) the Learning with Error (LWE) assumption, (ii) the Decisional Composite Residuosity (DCR) assumption, and (iii) a new notion of security that they call "shielded randomness leakage" (SRL). The latter essentially requires that a fully homomorphic encryption scheme (specifically the GSW encryption scheme [26]) remains secure even in the presence of a key-cycle with the Damgård-Jurik encryption scheme [16]. Moreover, the notion of security is not the standard semantic security, but rather a new notion of security with respect to leakage of ciphertext randomness. We note that this construction is insecure against quantum attackers because of the use of the Damgård-Jurik encryption scheme [16].[1] In this work, we ask:

Can we realize provably secure constructions of iO with (plausible) post-quantum security?

1.1 Our results

We obtain a general purpose iO construction based solely on the circular security of LWE-based encryption schemes. On a technical level, we achieve this by introducing a "packed" variant of the dual-Regev LWE-based encryption scheme, and showing novel ways of manipulating ciphertexts of this variant in conjunction with ciphertexts of an FHE scheme. This allows us to remove the need for DCR-based encryption from the construction of [9, 23]. Furthermore, our technique allows us to relax the SRL security property that is required, so that we no longer need to require SRL security with respect to a key-cycle, but rather only with respect to a key-randomness cycle. We put forth this potentially weaker assumption as an object for further study.

More concretely, the circular security assumption made in [23], and thus also in this work, is that a scheme (in particular a leveled FHE scheme) maintains this property even in the presence of some leakage on the randomness of the ciphertext. In [23] it is shown that standard GSW encryption [26] satisfies SRL security (under the LWE assumption), and the additional assumption is therefore that SRL security is maintained in the presence of a key-randomness cycle, connecting GSW to another encryption scheme. While this assumption falls into the category of "circular security assumptions", similarly to the ones that underlie 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.

Let us now state our results somewhat more precisely.

▶ **Theorem 1** (Informal). *Assume the (sub-exponential) hardness of the LWE problem, and the SRL security of GSW in the presence of a randomness-key cycle with a packed variant of dual-Regev, then there exists indistinguishability obfuscation for all circuits.*

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.

Shallow XiO. As an additional contribution, we identify a gap in the transformation of "exponentially efficient iO" (XiO), a notion introduced by Lin, Pass, Seth and Telang [35] that was used almost universally in prior work. We show that this transformation has an inherent problem that does not allow to recover the result as stated. This gap affects most known iO constructions and, in particular, *all post-quantum provably secure candidates.* We rectify this situation by showing that a fairly simple technical modification (i.e. constraining

[1] Concurrently, [23] updated their manuscript to also include a solution based on LWE. See Section 1.3 for additional discussion.

the compiler to be *shallow*) allows us to recover the prior results. Along the way, we develop a framework for analyzing composition of compressing encodings, which can be a useful perspective for future research in this area.

1.2 Technical Overview

We now provide a technical outline of our construction and its properties.

Obfuscation via Homomorphic Encryption. The connection between (fully) homomorphic encryption and obfuscation is fairly straightforward. Given a program Π to be obfuscated, we can provide a ciphertext c_Π which encrypts Π under an FHE scheme. This will allow to use homomorphism to derive $c_x = \mathsf{Enc}(\Pi(x))$ for all x. Now all that is needed is a way to decrypt c_x in a way that does not reveal any information on Π. Early works (e.g. [21] and followups) attempted to use this approach and provide a "defective" version of the secret key of the FHE scheme, but a different approach was suggested in [9].

Specifically, [9] considered a homomorphic evaluation that takes c_Π to c_{TT}, an encryption of the *entire truth table* of Π, i.e. to an encryption of a multi-bit value. By relying on prior generic transformations [35], they showed that one can reduce the task of constructing general-purpose obfuscation to the task of computing a "decryption" hint for c_{TT} with the following properties:

- Succinctness: The size of the decryption hint must be sublinear in the size of the truth table $|\mathsf{TT}|$.
- Simulatability: The decryption hint should not reveal any additional information besides the truth table TT.

The reason why this is helpful is that some so-called "packed-encryption" schemes have the property that a short ciphertext-dependent decryption hint suffices in order to decrypt the ciphertext, in a way that does not seem to leak the secret key of the scheme itself. While standard FHE schemes do not natively support packed encryption, it was shown in [8] that it is possible to use the so-called key-switching technique to switch from an FHE scheme into a packed-encryption scheme.

Alas, when instantiating the components of the [9] approach in its simplistic form described above, the decryption hint leaks information that renders the scheme insecure. To counter this issue, [9] proposed to inject another source of randomness: By adding freshly sampled ciphertexts of the packed-encryption scheme (which in their case was instantiated with the Damgård-Jurik scheme [16]) one can smudge the leakage of the decryption hint. However the size of these fresh ciphertext would largely exceed the size of the truth table TT. Therefore, [9] proposed to heuristically sample them from a random oracle, leveraging the fact that the ciphertexts of [16] are *dense*, i.e. a uniformly sampled string lies in the support of the encryption algorithm with all but negligible probability. This led to a candidate, but without a proof of security.

A Provably Secure Scheme. In a recent work, Gay and Pass [23] observed that for the purpose of constructing obfuscation, it suffices to consider schemes in the common random string (CRS) model where, importantly, the size of the CRS can exceed the size of the truth table. This allowed them to place the Damgård-Jurik ciphertexts in the CRS and therefore avoid relying on random-oracle-like heuristics.

They propose a new method to prove the security of this approach: Leveraging the structural property of the GSW scheme [26]. They showed that adding a GSW encryption of 0 to the evaluated FHE ciphertext (before key-switching to Damgård-Jurik) allows one to program the FHE ciphertext in the security proof. To sample these GSW encryptions of 0, they propose to draw the random coins $\mathbf{r}^*$ again from the CRS and let the evaluator recompute the correct ciphertext $\mathsf{GSW.Enc}(0; \mathbf{r}^*)$.

Taken together, these new ideas allow them to prove their construction secure against the shielded randomness leakage (SRL) security of the resulting FHE scheme. Loosely speaking, SRL security requires that semantic security of an encryption scheme is retained in the presence of an oracle that leaks the randomness $\mathbf{r}_f$ of the homomorphic evaluation of the function f over the challenge ciphertext. However the randomness $\mathbf{r}_f$ is not revealed in plain to the adversary, instead it is "shielded" by the random coins of a fresh GSW ciphertext $c = \mathsf{GSW.Enc}(0; \mathbf{r}^*)$. That is, the adversary is given $(\mathbf{r}_f - \mathbf{r}^*, c)$. In fact, the adversary can obtain polynomially-many samples from this distribution, for any function f, conditioned on the fact that the adversary knows the output of $f(m^*)$, where m^* is the hidden message.

To gain confidence in the veracity of the assumption, [23] show that the GSW encryption scheme satisfies SRL security if the (plain) LWE assumption holds. However, their obfuscation scheme requires one to publish a key cycle of GSW and Damgård-Jurik (i.e. an encryption of the GSW secrey key under Damgård-Jurik and vice versa). Thus their final assumption is that SRL security is retained in the presence of such a key cycle.

Obfuscation from Circular-Secure LWE. We wish to remove the need for the Damgård-Jurik encryption scheme from the above construction paradigm. The major obstacle to overcome consists in designing an LWE-based encryption scheme that simultaneously satisfies three properties.

- Linear Homomorphism: In order to key switch the GSW ciphertext into this form, the scheme must satisfy some weak notion of homomorphism. Specifically, it must support the homomorphic evaluation of linear functions.
- Succinct Randomness: The scheme must allow us to encrypt a long message string with a short randomness, that can then function as the decryption hint.
- Dense Ciphertexts: A uniformly sampled string must lie in the support of the encryption algorithm with all but negligible probability. This will allow us to parse the CRS as a collection of ciphertexts.[2]

Unfortunately all natural lattice-based candidates seem to fail to satisfy all of these properties. In particular, for all LWE-based schemes linear homomorphism seems to be at odds with dense ciphertexts: To ensure that the noise accumulated during the homomorphic evaluation does not impact the decryption correctness, one needs to ensure a gap between the noise bound and the modulus. More concretely, ciphertext are typically of the form $(\mathbf{a}, \mathbf{a} \cdot \mathsf{s} + e + q/2 \cdot \mathsf{m}) \in \mathbb{Z}_q^{n+1}$ where $e \ll q$, which makes them inherently sparse.

Our Solution: A Packed Variant of Dual-Regev that is also Dense-Friendly. We show that the above requirements can be relaxed. Our starting point is devising a "packed" version of the dual-Regev encryption scheme [25]. This scheme will not have dense ciphertexts so it does not fit the requirements from previous works. However, we will show how we can define, for the same scheme, a family of ciphertexts which are both "almost dense" and can inter-operate with the non-dense scheme, so as to allow to construct the obfuscator.

[2] Note that for the purpose of constructing the obfuscator, one could make do with a common reference string which can have an arbitrary distribution. However, the string needs to be parsed as a ciphertext with respect to *all* public-keys. Requiring dense ciphertexts is a simple requirement that implies this property.

Let us start with our packed dual-Regev scheme. To pack a k-bit plaintext $\mathbf{m} \in \{0,1\}^k$ in a dual-Regev ciphertext we construct the public key as a matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, which is statistically close to uniform but is sampled together with a trapdoor τ (whose role will be explained below), and another uniformly sampled matrix $\mathbf{B} \in \mathbb{Z}_q^{k \times n}$. The encryption algorithm computes a the ciphertext as

$$(\mathbf{A} \cdot \mathbf{r} + \mathbf{e}_0, \mathbf{B} \cdot \mathbf{r} + q/2 \cdot \mathbf{m} + \mathbf{e})$$

where $\mathbf{r} \leftarrow\$ \mathbb{Z}_q^n$ is the encryption randomness and the vectors $\mathbf{e}_0$ and $\mathbf{e}$ are the encryption noises, where the norm of both vectors is bounded by some $B \ll q$. The property of the trapdoor τ is that it allows to recover $\mathbf{r}$ from $\mathbf{A} \cdot \mathbf{r} + \mathbf{e}_0$. The (semantic) security of the scheme follows directly by definition of LWE. To decrypt, therefore, one can first use the trapdoor τ to recover $\mathbf{r}$ from the first m elements of the ciphertext, and then recompute the mask $\mathbf{B} \cdot \mathbf{r}$ and recover each individual bit by rounding to the closest multiple of $q/2$. Setting the parameters appropriately, we can guarantee that the decryption is always successful. One important property of this scheme is that the random coins $\mathbf{r} \in \mathbb{Z}_q^n$ are sufficient to recover the entire message and furthermore the size of $\mathbf{r}$ is succinct (in particular independent of k).

In terms of homomorphism, the scheme is straightforwardly additively homomorphic. Furthermore, it supports key switching from any scheme with almost-linear decryption as per [8].[3] In particular it is possible to take a (long) message encrypted under an FHE scheme such as GSW and convert it to an encryption of the same message under packed dual-Regev, using precomputed key-switching parameters.[4]

As explained above, this scheme does not have dense ciphertexts. At this point we make two crucial observations that will allow us to bypass this hurdle.

(1) In order to construct the obfuscator using the [9] approach, dense ciphertexts only need to enjoy a very limited form of homomorphism, they only need to support a single addition with a non-dense ciphertext.

This is essentially because the obfuscator has the following outline. It starts by considering the dense ciphertext from the CRS (or oracle in the case of the original [9]), and homomorphically bootstraps it into a non-dense FHE ciphertext by evaluating the decryption circuit. Let $\mathbf{m}$ be the (random) message that is induced by the process. Then, the FHE encryption of $\mathbf{m}$ is processed in order to create a non-dense *packed* encryption of $\mathbf{m} \oplus \mathsf{TT}$, where TT is the truth table of the program to be obfuscated (or, more accurately, a chunk of this truth table, partitioning into chunks is required in order to allow reusability of the keys). Then a single homomorphic addition between the dense and non-dense ciphertext would imply a packed encryption of the truth table. All of this can be performed by the evaluator of the obfuscated program, so all that is needed is the decryption hint for this final ciphertext, that would allow to recover TT.

We note importantly, that in prior approaches (including the [23] blueprint) the aforementioned bootstrapping creates a key cycle, since a packed ciphertext is bootstrapped into an FHE ciphertext, which is afterwards key-switched into a packed ciphertext. However, we notice that it suffices to provide an encryption of the (succinct) *randomness* of the dense ciphertext in order to apply bootstrapping, thus leading to a relaxed key-randomness circular assumption. Interestingly, this observation is not very useful for actual dense ciphertexts (since finding the randomness would require using the key), however, our relaxed notion of density described below will allow to apply it and thus relax the circularity notion as well.

[3] This is done using the by-now-standard technique of encrypting powers-of-two of the elements of the secret key of the latter scheme, so that it is possible to evaluate any inner product homomorphically.

[4] We note that the key switching parameters are quite long so it is required for our method that they are reusable.

(2) A notion of *almost-everywhere* density suffices. A ciphertext distribution is almost-everywhere dense if it is dense except for a non-dense part whose length is independent of k (the message length).

The reason that this is sufficient is that the non-dense part of the ciphertext, which we refer to as the *header*, can be generated by the obfuscator and provided to the evaluator as a part of the obfuscated program. Since the header is short, and in particular the message length k can be selected to be much longer than the header, the effect on the length of the obfuscated program will be minimal. As hinted above, since the obfuscator generates the header, it in particular also samples the randomness for the final almost-everywhere dense ciphertext. This means that the obfuscator can generate the bootstrapping parameters using this randomness without requiring a key cycle.

Dense Encryption Mode. With these observations in mind we describe an alternative encryption mode (DenseEnc) for the packed variant of dual-Regev where the bulk of the ciphertext is dense. On input a message $\mathbf{m} \in \{0,1\}^k$, the encryption algorithm in dense mode computes the following ciphertext

$$(\mathbf{A}\cdot\mathbf{r}+\mathbf{e}_0, \mathbf{B}\cdot\mathbf{r}+q/2\cdot\mathbf{m}+\mathbf{u})$$

where $\mathbf{r}$ and $\mathbf{e}_0$ are sampled as before and $\mathbf{u} \leftarrow\!\!\$\ [-q/4,+q/4]^k$. For convenience, we are going to split the ciphertexts into two blocks: The header $\mathbf{h}_0 \in \mathbb{Z}_q^m$ and the message carrier $(h_1,\ldots,h_k) \in \mathbb{Z}_q^k$. Foremost, observe that the decryption algorithm as described before still returns the correct message with probability 1, since it recovers the same $\mathbf{r}$ from $\mathbf{h}_0$. Furthermore, note that (for a fixed header) all vectors $(h_1,\ldots,h_k) \in \mathbb{Z}_q^k$ are in the support of the encryption algorithm. Since $k \gg m$, most of the elements of the ciphertext in the alternative encryption mode are dense.

One can verify that the aforementioned limited form of homomorphism indeed holds, namely that

$$\mathsf{dR.Enc}(\mathbf{m})+\mathsf{dR.DenseEnc}(\mathbf{m}')\in\mathsf{dR.DenseEnc}(\mathbf{m}\oplus\mathbf{m}').$$

This is the case since

$$\begin{aligned}&(\mathbf{A}\cdot\mathbf{r}+\mathbf{e}_0,\mathbf{B}\cdot\mathbf{r}+q/2\cdot\mathbf{m}+\mathbf{e})+(\mathbf{A}\cdot\mathbf{r}'+\mathbf{e}_0',\mathbf{B}\cdot\mathbf{r}'+q/2\cdot\mathbf{m}'+\mathbf{u})\\&=(\mathbf{A}\cdot(\mathbf{r}+\mathbf{r}')+\mathbf{e}_0+\mathbf{e}_0',\mathbf{B}\cdot(\mathbf{r}+\mathbf{r}')+q/2\cdot(\mathbf{m}\oplus\mathbf{m}')+\mathbf{e}+\mathbf{u})\\&=(\mathbf{A}\cdot\tilde{\mathbf{r}}+\tilde{\mathbf{e}}_0,\mathbf{B}\cdot\tilde{\mathbf{r}}+q/2\cdot(\mathbf{m}\oplus\mathbf{m}')+\tilde{\mathbf{u}})\end{aligned}$$

where $\tilde{\mathbf{u}}=\mathbf{e}+\mathbf{u}\in[-q/4,+q/4]^k$ with all but negligible probability over the random choice of $\mathbf{u}$, for an appropriate choice of the parameters.

Doing Away with the Header. We notice that given our two observations above, the goal of the header in the obfuscation scheme is quite minimal. The header is not needed for homomorphism, and is only needed for the purpose of extracting the randomness $\mathbf{r}$ at decryption time. We then observe that decrypting packed ciphertext is done in two contexts in the scheme. The first is when we bootstrap the almost-everywhere dense ciphertext into an FHE ciphertext, and the other is when the evaluator of the obfuscated program recovers TT from the final ciphertext. For the latter there is no need for a header since the decryption hint, i.e. the respective $\mathbf{r}$ value, is provided within the obfuscated program. For the former we do not need a header of a specific structure, but rather simply an encryption of $\mathbf{r}$ that allows bootstrapping the almost-dense ciphertext. It therefore suffices to provide $\mathsf{GSW.Enc}(\mathbf{r})$ directly, which makes the header completely redundant.

On the Assumption. Equipped with the newly developed packed version of dual-Regev we can follow the [9, 23] approach, with the aforementioned modifications, to construct the obfuscator. The resulting construction can be shown secure against the assumption that the SRL security of GSW is retained in the presence of a key cycle with the packed dual-Regev encryption scheme as presented above.

We then observe that it suffices to assume SRL security with respect to key-randomness cycles, rather than key cycles. We note that this assumption is no-stronger than key-cycle SRL since given a key-cycle it is possible to homomorphically generate a key-randomness cycle, but the converse is not known to be true.

Adding this to our observation about the redundancy of the header, the assumption we require is that SRL security is retained in the presence of a key-randomness cycle between GSW and packed dual-Regev, i.e.

$$(\mathsf{GSW.Enc}(\mathbf{r}), \mathsf{dR.Enc}(\mathsf{sk}_{\mathsf{GSW}}; \mathbf{r}))\,.$$

Since dual-Regev is randomness recoverable, this assumption is syntactically weaker than SRL security in the presence of a key-cycle: Given a GSW encryption of the dual-Regev secret key, one can homomorphically compute the randomness recovery circuit to obtain a GSW encryption of the randomness $\mathbf{r}$.

1.3 Related and Follow-up Work

Subsequently to the posting of this manuscript online (but concurrently and independently) [23] updated their manuscript to include a solution based on LWE in the place of DCR. They do not make the observations that a relaxed notion of density suffices (and is preferable) and thus they explicitly construct an encryption scheme with dense ciphertexts based on the (primal) Regev encryption scheme. The resulting scheme is more involved and in particular requires the two-key circular SRL security of GSW and (primal) Regev rather than the relaxed key-randomness circularity notion.

Wee and Wichs [42], again concurrently, presented another instantiation of the [9] approach which is arguably post-quantum secure. They rely on an indistinguishability assumption between two distributions and not directly on circular security. However, the underlying machinery developed shares many similarities with our approach. Specifically, while we essentially rely on randomness that is embedded in the CRS by interpreting it as an obliviously sampled ciphertext (which thus corresponds to one encrypted with fresh randomness), their approach is to use a pseudorandom function to transform the CRS into a randomizer for the output hint.

A follow-up work by Hopkins, Jain, and Lin [29] shows counterexamples to SRL security for general functions in the presence of a 2-key cycle, as stated in [23], and the conjecture from [42]. We stress that their findings do not imply an attack against the corresponding obfuscation scheme of [42] and [23] (as also pointed out by the authors in [29]). Rather, their results show that the veracity of SRL security depends on the concrete circuit representation of the functions under consideration. As a consequence of their findings, we updated the statement of our assumption (and adapted the analysis of our scheme) with a refined version, that further restricts the power of the adversary and more tightly characterize the security of our construction. However, the iO construction is unchanged from previous versions of this work.

A few remarks about the susceptibility of our scheme to the [29] attack are in order. In short, the attack exploits the randomness homomorphism of GSW to compute a biased leakage. The SRL function consists of a bootstrapping followed by a modular reduction (modulo 2). On the other hand, our admissible class of leakage functions consists of linear functions (modulo q) followed by a rounding (i.e. outputting the most significant bit). We are not aware of a method to establish the same correlations exploited by [29] without violating the admissibility criteria for the leakage functions. Thus, we conjecture that SRL security with respect to such leakage function holds for all natural FHE candidates (see Section 3.1 for further details).

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 inverse 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 Turing machine M running in time $\mathsf{poly}(\lambda)$. The execution of a Turing machine M on input x and with random coins fixed to r is denoted by $M(x;r)$. We say that two distributions (D_0, D_1) are computationally (statistically, resp.) indistinguishable if for all PPT (unbounded, resp.) distinguishers, the probability to tell D_0 an D_1 apart is negligible. Matrices are denoted by $\mathbf{M}$ and vectors are denoted by $\mathbf{v}$. For convenience, we define $\mathbf{Bit}(\cdot)$ as the bit decomposition operation. We denote the infinity norm of a vector $\mathbf{v}$ by $\|\mathbf{v}\|_\infty$. We recall the smudging lemma.

▶ **Lemma 2** (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 to that of $e_2 + e_1$ as long as $B_1/B_2 = \mathsf{negl}(\lambda)$.*

2.1 Indistinguishability Obfuscation

We recall the notion of indistinguishability obfuscation (iO) from [4].

▶ **Definition 3** (Indistinguishability Obfuscation). *A PPT machine iO is an indistinguishability obfuscator for a circuit class $\{\mathfrak{C}_\lambda\}_{\lambda \in \mathbb{N}}$ if the following conditions are satisfied:*

(Functionality) For all $\lambda \in \mathbb{N}$, all circuit $\Pi \in \mathfrak{C}_\lambda$, all inputs x it holds that: $\tilde{\Pi}(x) = \Pi(x)$, where $\tilde{\Pi} \leftarrow\!\$\; \mathsf{iO}(\Pi)$.

(Indistinguishability) For all $\lambda \in \mathbb{N}$, all pairs of circuit $(\Pi_0, \Pi_1) \in \mathfrak{C}_\lambda$ such that $|\Pi_0| = |\Pi_1|$ and $\Pi_0(x) = \Pi_1(x)$ on all inputs x, it holds that the following distributions are computationally indistinguishable: $\mathsf{iO}(\Pi_0) \approx \mathsf{iO}(\Pi_1)$.

Shallow XiO. In this work we construct a weaker version of iO called (shallow) XiO, which however is sufficient (along with the LWE assumption) to construct fully-fledged iO. Loosely speaking, a shallow XiO is a indistinguishability obfuscator (with pre-processing) for $\mathsf{P}^{\mathsf{log}}/\mathsf{poly}$ with non-trivial efficiency. Here $\mathsf{P}^{\mathsf{log}}/\mathsf{poly}$ 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 $\mathsf{poly}(\lambda, |\Pi|) \cdot 2^{\eta\cdot(1-\varepsilon)}$, for some constant $\varepsilon > 0$. The runtime of the obfuscator can be any polynomial in λ, $|\Pi|$, and 2^η, except that its depth should not depend on 2^η. Furthermore, we allow the obfuscator to access a large uniform random string (the pre-processing) of size even larger than the truth table of the circuit. For a formal statement, we refer the reader to the full version [10].

2.2 The GSW Fully-Homomorphic Encryption

In the following we briefly recall the encryption scheme by Gentry, Sahai, and Waters [26] (henceforth, GSW). We denote by $n = n(\lambda)$ the lattice dimension and by $q = q(\lambda)$ the modulus (which we assume for simplicity to be even). Throughout the rest of this paper, we set $m = (n+1)(\log(q)+1)$ and $d = d(\lambda)$ as a bound on the depth of the arithmetic circuit to be evaluated.

$\mathsf{KeyGen}(1^\lambda)$: Sample a uniform matrix $\mathbf{A} \leftarrow\!\!\$\ \mathbb{Z}_q^{n\times m}$ and a vector $\mathbf{s} \leftarrow\!\!\$\ \chi^n$. Set the public key to $(\mathbf{A}, \mathbf{b} = \mathbf{s}^T\mathbf{A} + \mathbf{e}^T)$, where $\mathbf{e} \leftarrow\!\!\$\ \chi^m$. The secret key is set to $(-\mathbf{s}, 1)$.

$\mathsf{Enc}(\mathsf{pk}, \mathsf{m})$: On input a message $\mathsf{m} \in \{0,1\}$, sample a uniform $\mathbf{R} \leftarrow\!\!\$\ \{0,1\}^{m\times m}$ and compute

$$\mathbf{C} = (\mathbf{A}, \mathbf{b}) \cdot \mathbf{R} + \mathsf{m} \cdot \mathbf{G}$$

where $\mathbf{G} = (1, 2, \dots, 2^{\log(q)-1})^T \otimes I_{(n+1)}$ and $I_{(n+1)} \in \{0,1\}^{(n+1)\times(n+1)}$ denotes the identity matrix.

$\mathsf{Eval}(\mathsf{pk}, \Pi, (c_1, \dots, c_\mu))$: There exists a (deterministic) polynomial-time algorithm that allows one to compute any d-bounded depth arithmetic circuit $\Pi : \{0,1\}^n \to \{0,1\}$ homomorphically over a vector of ciphertexts $(c_1, \dots, c_\mu)$. For details about this algorithm, we refer the reader to [26]. For the purpose of this work, the only relevant information is that the evaluated ciphertext $\mathbf{c}_\Pi \in \mathbb{Z}_q^{(n+1)}$ is an $(n+1)$-dimensional vector. For multiple bits of output, the resulting ciphertext is defined to be the concatenation of the single-bit ciphertexts.

$\mathsf{Dec}(\mathsf{sk}, \mathbf{c})$: We assume without loss of generality that the input ciphertext $\mathbf{c} \in \mathbb{Z}_q^{(n+1)}$ is the output of the evaluation algorithm. Such a ciphertext defines a linear function $\boldsymbol{\ell}_\mathbf{c}$ such that

$$\boldsymbol{\ell}_\mathbf{c}(\mathsf{sk}) = q/2 \cdot \mathsf{m} + e$$

where $|e| \le \hat{B} = (m+1)^d mB$. The message m is recovered by returning the most significant bit of the output.

Note that the decryption routine of GSW consists of the application of a linear function, followed by a rounding and we refer to this property as to *almost-linear* decryption. In a slight abuse of notation, we sometimes write $\mathsf{KeyGen}(1^\lambda; q)$ to denote the above key generation algorithm with a fixed modulus q.

Alternate Encryption. For convenience we also define a modified encryption algorithm, where the output ciphertexts consists of a single column vector. An additional difference is that we sample the randomness with norm $\tilde{B} = 2^\lambda \cdot \hat{B}$.

$\mathsf{ColEnc}(\mathsf{pk}, \mathsf{m})$: On input a message m, sample a uniform $\mathbf{r} \leftarrow\!\!\$\ [-\tilde{B}, +\tilde{B}]^m$ and compute

$$\mathbf{c} = (\mathbf{A}, \mathbf{b}) \cdot \mathbf{r} + (0^n, q/2) \cdot \mathsf{m}.$$

This algorithm is going instrumental for our scheme, although ciphertexts in this form no longer support the homomorphic evaluation of arbitrary circuits. The multi-bit version of such an algorithm is defined accordingly to output the concatenation of independently sampled ciphertexts. We now recall a useful Lemma from [23].

▶ **Lemma 4** (GSW Smudging). *Let $\tilde{B} = 2^\lambda \cdot \hat{B}$. For all $\lambda \in \mathbb{N}$, for all $(\mathsf{sk}, \mathsf{pk})$ in the support of $\mathsf{KeyGen}(1^\lambda)$, for all messages $\mathbf{m} = (\mathsf{m}_1, \dots, \mathsf{m}_\mu)$, for all depth-$d$ circuit $(\Pi_1, \dots, \Pi_\tau)$, the following distributions are statistically indistinguishable*

$$\begin{pmatrix} c_1, \dots, c_\mu, \mathbf{r}_1^*, \dots, \mathbf{r}_\tau^*, \\ \mathsf{Eval}(\mathsf{pk}, \Pi_1, (c_1, \dots, c_\mu)) + \mathsf{ColEnc}(\mathsf{pk}, 0; \mathbf{r}_1^*), \dots, \\ \mathsf{Eval}(\mathsf{pk}, \Pi_\tau, (c_1, \dots, c_\mu)) + \mathsf{ColEnc}(\mathsf{pk}, 0; \mathbf{r}_\tau^*) \end{pmatrix}$$
$$\approx \begin{pmatrix} c_1, \dots, c_\mu, \mathbf{r}_1^* - \mathsf{RandEval}(\mathsf{pk}, \Pi_1, \mathbf{m}, (\mathbf{R}_1, \dots, \mathbf{R}_\mu)), \dots, \\ \mathbf{r}_\tau^* - \mathsf{RandEval}(\mathsf{pk}, \Pi_\tau, \mathbf{m}, (\mathbf{R}_1, \dots, \mathbf{R}_\mu)), \\ \mathsf{ColEnc}(\mathsf{pk}, \Pi_1(\mathsf{m}_1, \dots, \mathsf{m}_\mu); \mathbf{r}_1^*), \dots, \mathsf{ColEnc}(\mathsf{pk}, \Pi_\tau(\mathsf{m}_1, \dots, \mathsf{m}_\mu); \mathbf{r}_\tau^*) \end{pmatrix}$$

where $c_i \leftarrow\$ \mathsf{Enc}(\mathsf{pk}, \mathsf{m}_i; \mathbf{R}_i)$, $\mathbf{r}_i^* \leftarrow\$ [-\tilde{B}, +\tilde{B}]^m$, and $\mathbf{R}_i \leftarrow\$ \{0,1\}^{m \times m}$.

Randomness Homomorphism. We recall a useful property of the GSW scheme, namely that one can alternatively evaluate functions directly over the randomness of a ciphertext to obtain the same result. More formally, we say that a homomorphic encryption scheme $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Eval}, \mathsf{Dec})$ has randomness homomorphism for the circuit class $\{\mathfrak{C}_\lambda\}_{\lambda \in \mathbb{N}}$ if there exists an efficient algorithm $\mathsf{RandEval}$ such that for all $\Pi \in \mathfrak{C}_\lambda$, all $(\mathsf{sk}, \mathsf{pk})$ in the support of KeyGen, all vectors of messages $\mathbf{m} = (\mathsf{m}_1, \dots, \mathsf{m}_\mu)$ and $\mathbf{R} = (\mathbf{R}_1, \dots, \mathbf{R}_\mu)$, all ciphertexts $(c_1, \dots, c_\mu)$ in the support of $(\mathsf{Enc}(\mathsf{pk}, \mathsf{m}_1; \mathbf{R}_1), \dots, \mathsf{Enc}(\mathsf{pk}, \mathsf{m}_\mu; \mathbf{R}_\mu))$ it holds that

$$\mathsf{Eval}(\mathsf{pk}, \Pi, (c_1, \dots, c_\mu)) = \mathsf{ColEnc}(\mathsf{pk}, \Pi(\mathbf{m}); \mathsf{RandEval}(\mathsf{pk}, \Pi, \mathbf{m}, \mathbf{R})).$$

Circuit Privacy. It is well known that the GSW encryption scheme satisfies the following notion of circuit privacy [7, 18, 40] (with a randomized evaluation algorithm).

▶ **Definition 5** (Circuit Privacy). *For all $\lambda \in \mathbb{N}$, all all $\Pi \in \mathfrak{C}_\lambda$, all $(\mathsf{sk}, \mathsf{pk})$ in the support of KeyGen, and all messages m, it holds that the following distributions are statistically indistinguishable*

$$(\mathsf{pk}, \mathsf{Enc}(\mathsf{pk}, \Pi(\mathsf{m}))) \approx (\mathsf{pk}, \mathsf{Eval}(\mathsf{pk}, \Pi, \mathsf{Enc}(\mathsf{pk}, \mathsf{m}); r)).$$

where $r \leftarrow\$ \{0,1\}^\lambda$.

3 Packed Encryption from LWE

In the following we describe a packed version of the dual-Regev encryption scheme [25]. We denote by $n = n(\lambda)$ the lattice dimension, by $q = q(\lambda)$ the modulus (which we assume for simplicity to be a power of 2), and by $k = k(\lambda)$ the expansion factor. We require the existence of a public-key encryption scheme $(\mathsf{PKE.KeyGen}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$.

$\mathsf{KeyGen}(1^\lambda, 1^k)$: Sample a uniform $k \times n$ matrix $\mathbf{B} \leftarrow\$ \mathbb{Z}_q^{k \times n}$ and a key pair of a public-key encryption scheme $(\mathsf{sk}_{\mathsf{PKE}}, \mathsf{pk}_{\mathsf{PKE}}) \leftarrow\$ \mathsf{PKE.KeyGen}(1^\lambda)$. The public key consists of $(\mathbf{B}, \mathsf{pk}_{\mathsf{PKE}})$ and the secret key is set to $\mathsf{sk}_{\mathsf{PKE}}$.

$\mathsf{Enc}(\mathsf{pk}, \mathbf{m})$: To encrypt a k-bit message $\mathbf{m} \in \{0,1\}^k$, sample a uniform randomness vector $\mathbf{r} \leftarrow\$ \mathbb{Z}_q^n$ a noise vector $\mathbf{e} \leftarrow\$ \chi^k$ and return the ciphertext

$$\mathbf{c} = (\mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{r}), \mathbf{B} \cdot \mathbf{r} + q/2 \cdot \mathbf{m} + \mathbf{e}).$$

$\mathsf{Dec}(\mathsf{sk}, \mathbf{c})$: Parse $\mathbf{c}$ as $(c_{\mathsf{PKE}}, c_1, \dots, c_k)$ and recover the random coins by decrypting $\mathbf{r} = \mathsf{PKE.Dec}(\mathsf{sk}_{\mathsf{PKE}}, c_{\mathsf{PKE}})$. Let $\mathbf{b}_i$ be the i-th row of $\mathbf{B}$. For $i = 1 \dots k$, compute $\mathsf{m}_i = \mathsf{Round}(c_i - \mathbf{b}_i \cdot \mathbf{r})$, where Round rounds to the nearest multiple of $q/2$, i.e. it returns 1 if the input is closer to $q/2$ and 0 otherwise. Output $\mathbf{m} = (\mathsf{m}_1, \dots, \mathsf{m}_k)$.

Clearly, the scheme is perfectly correct since

$$\begin{aligned}
&(\mathsf{Round}(c_1 - \mathbf{b}_1 \cdot \mathbf{r}), \ldots, \mathsf{Round}(c_k - \mathbf{b}_k \cdot \mathbf{r}))\\
&= (\mathsf{Round}(q/2 \cdot \mathsf{m}_1 + e_1), \ldots, \mathsf{Round}(q/2 \cdot \mathsf{m}_k + e_k))\\
&= (\mathsf{Round}(q/2 \cdot \mathsf{m}_1), \ldots, \mathsf{Round}(q/2 \cdot \mathsf{m}_k))\\
&= (\mathsf{m}_1, \ldots, \mathsf{m}_k)\\
&= \mathbf{m}.
\end{aligned}$$

Extended Encryption. It is not hard to see that the scheme presented above is (bounded) additively homomorphic over $\mathbb{Z}_2^k$. To lift the class of computable functions to all linear functions over $\mathbb{Z}_q^k$, we adopt the standard trick of encrypting the message multiplied by all powers of two $(1, 2, \ldots, 2^{\log(q)})$. For convenience, we define the following augmented encryption algorithm.

$\underline{\mathsf{ExtEnc}(\mathsf{pk}, \mathbf{m})}$: On input an ℓ-dimensional message $\mathbf{m} \in \mathbb{Z}_q^\ell$, let $\mathbf{g} = (1, 2, \ldots, 2^{\log(q)-1})^T$ and define

$$\mathbf{M} = \begin{bmatrix} \mathsf{m}_1 \cdot \mathbf{g} & \mathsf{m}_2 \cdot \mathbf{g} & \ldots & 0^{\log(q)} \\ 0^{\log(q)} & 0^{\log(q)} & \ldots & 0^{\log(q)} \\ \vdots & \vdots & \ddots & \vdots \\ 0^{\log(q)} & 0^{\log(q)} & \ldots & \mathsf{m}_\ell \cdot \mathbf{g} \end{bmatrix} \in \mathbb{Z}_q^{k \times \ell \cdot k \cdot \log(q)}.$$

Sample a uniform randomness matrix $\mathbf{R} \leftarrow\$ \mathbb{Z}_q^{n \times \ell \cdot k \cdot \log(q)}$ and a uniform noise matrix $\mathbf{E} \leftarrow\$ \chi^{k \times \ell \cdot k \cdot \log(q)}$. Compute

$$\mathbf{C} = \mathbf{B} \cdot \mathbf{R} + \mathbf{M} + \mathbf{E}$$

and return the ciphertext $(\mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{R}), \mathbf{C})$.

Decryption works, as before, by recovering $\mathbf{R}$ from the public-key encryption scheme and then decrypting $\mathbf{m}$ component-wise.

Almost-Everywhere Dense Encryption. For convenience, we also define an alternative encryption algorithm in the following. Note that the encryption algorithm does not take as input any message, instead it encrypts a uniform k-bit binary vector. Syntactically, this is the equivalent of a key-encapsulation mechanism.

$\underline{\mathsf{DenseEnc}(\mathsf{pk})}$: Sample a uniform randomness vector $\mathbf{r} \leftarrow\$ \mathbb{Z}_q^n$ and return the ciphertext

$$\mathbf{c} = (c_{\mathsf{PKE}}, c_1, \ldots, c_k) = (\mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{r}), \mathbf{B} \cdot \mathbf{r} + \mathbf{u}).$$

where $\mathbf{u} \leftarrow\$ \mathbb{Z}_q^k$.

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 $\mathsf{DenseEnc}$ algorithms, where the plaintext of $\mathsf{DenseEnc}$ corresponds to $(\mathsf{Round}(u_1), \ldots, \mathsf{Round}(u_k))$. In fact, the scheme satisfies perfect correctness in both cases. (ii) The domain of the elements $(c_1, \ldots, c_k)$ is *dense*, i.e. the support of the scheme spans the entire vector space $\mathbb{Z}_q^k$. Since the element c_{PKE} is small (i.e. independent of k) for an appropriate choice of the public-key encryption scheme, we refer to such a property as *almost-everywhere* density.

Semantic Security. 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 $\mathbb{Z}_q^k$. Semantic security for the extended encryption ExtEnc and the dense encryption $\mathsf{DenseEnc}$ follows along the same lines.

▶ **Theorem 6** (Semantic Security). *If* $(\mathsf{PKE.KeyGen}, \mathsf{PKE.Enc}, \mathsf{PKE.Dec})$ *is semantically secure and the LWE assumption holds, then for all* $\lambda \in \mathbb{N}$ *and all messages* m *it holds that the following distributions are computationally indistinguishable*

$$(\mathsf{pk}, \mathsf{Enc}(\mathsf{pk}, \mathsf{m})) \approx (\mathsf{pk}, \mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{z}), \mathbf{u}).$$

where $(\mathsf{sk}, \mathsf{pk}) \leftarrow\$ \mathsf{KeyGen}(1^\lambda, 1^k)$, $\mathbf{z} \leftarrow\$ \mathbb{Z}_q^n$, *and* $\mathbf{u} \leftarrow\$ \mathbb{Z}_q^k$.

Proof. The security of the scheme follows routinely by an invocation of semantic security of the public-key encryption scheme and an invocation of the LWE assumption. ◀

3.1 Key-Randomness SRL Security

We state a version of SRL security [23] tailored for our specific instance and adapted to the randomness-key circularity assumption (rather than the 2-key circularity, as stated in [23]).

▶ **Definition 7** (Key-Randomness SRL Security). *Let* $(\mathsf{GSW.KeyGen}, \mathsf{GSW.Enc}, \mathsf{GSW.Eval}, \mathsf{GSW.Dec})$ *be the GSW encryption scheme and* $(\mathsf{dR.KeyGen}, \mathsf{dR.Enc}, \mathsf{dR.Eval}, \mathsf{dR.Dec})$ *be the packed dual-Regev encryption scheme. Fix messages* $(\mathsf{m}_0, \mathsf{m}_1)$, *polynomials* $\tau = \tau(\lambda)$ *and* $k = k(\lambda)$ *and an adversary* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$. *Consider the following experiment.*

$\mathsf{Exp}_{\mathsf{SRL}}^{(b)}(\mathcal{A})$*:*

- *Sample* $(\bar{\mathsf{sk}}, (\bar{\mathsf{pk}}, \mathbf{B})) \leftarrow\$ \mathsf{dR.KeyGen}(1^\lambda, 1^k)$ *and* $(\mathsf{sk}, \mathsf{pk}) \leftarrow\$ \mathsf{GSW.KeyGen}(1^\lambda)$
- *Compute* $\mathbf{c} \leftarrow\$ \mathsf{GSW.Enc}(\mathsf{pk}, \mathsf{m}_b)$
- $\{\mathbf{P}_i, \mathbf{p}_i\}_{i=1\ldots\tau} = \mathcal{A}_1(\mathsf{pk}, \mathbf{B}, \mathbf{c})$
- *Compute* $\left(\bar{c}_{\mathsf{sk}}, \bar{\mathbf{C}}_{\mathsf{sk}}\right) = \mathsf{dR.ExtEnc}(\bar{\mathsf{pk}}, \mathsf{sk}; \mathbf{S})$ *and* $\mathbf{c}_{\mathbf{S}} \leftarrow\$ \mathsf{GSW.Enc}(\mathsf{pk}, \mathbf{S}; \mathbf{R}_{\mathbf{S}})$
- *For all* $i = 1 \ldots \tau$*:*
    - *Sample* $\mathbf{C}_i^* = \mathsf{ColEnc}(0; \mathbf{r}_i^*)$ *where* $\mathbf{r}_i^* \leftarrow\$ [-\tilde{B}, +\tilde{B}]^{m\cdot k}$
    - *Sample* $\mathbf{t}_i \leftarrow\$ \mathbb{Z}_q^n$
    - *Sample* $\mathbf{c}_{i,\mathbf{r}} = \mathsf{GSW.Enc}(\mathsf{pk}, \mathbf{S} \cdot \mathsf{Bit}(\ell_i) + \mathbf{t}_i; \mathbf{R}_i) \leftarrow\$ \mathsf{GSW.Eval}(\mathsf{pk}, \cdot\mathsf{Bit}(\ell_i) + \mathbf{t}_i, \mathbf{c}_{\mathbf{S}})$
- *Output* $\mathcal{A}_2(\bar{\mathbf{C}}_{\mathsf{sk}}, \{\mathbf{c}_{i,\mathbf{r}}, \mathbf{C}_i^*, \mathbf{u}_i, \mathbf{t}_i, \mathbf{r}_{\psi,i} - \mathbf{r}_i^*\}_{i=1\ldots\tau})$

Here, letting ℓ_i *be the linear function associated with* $\mathbf{C}_i^* + \mathbf{P}_i + q/2 \cdot \mathbf{v}_i$, *we set*

$$\mathbf{r}_{\psi,i} = \mathsf{RandEval}(\mathsf{pk}, \psi_i, \mathbf{S} \cdot \mathsf{Bit}(\ell_i) + \mathbf{t}_i, \mathbf{R}_i) \text{ and}$$
$$\psi_i(\mathbf{Z}) = \mathsf{Round}\left(\mathbf{B} \cdot \mathbf{t}_i + q/2 \cdot \mathbf{p}_i + \mathbf{w}_i - \bar{\mathbf{C}}_{\mathsf{sk}} \cdot \mathsf{Bit}(\ell_i) - \mathbf{B} \cdot \mathbf{Z}\right)$$

where $\mathbf{w}_i \leftarrow\$ [-q/4, q/4]^k$, $\mathbf{v}_i \leftarrow\$ \{0,1\}^k$, *and* $\mathbf{u}_i = \mathbf{w}_i + q/2 \cdot \mathbf{v}_i$, *for all* $i = 1 \ldots \tau$. *We say that an adversary* $\mathcal{A}$ *is admissible if for all* $i = 1 \ldots \tau$ *it holds that* $\mathbf{P}_i \in \mathsf{GSW.Enc}(\mathsf{pk}, \mathbf{p}_i)$. *The KR-SRL assumption conjectures that it holds for all admissible PPT adversaries* $\mathcal{A}$, *all messages* $(\mathsf{m}_0, \mathsf{m}_1)$ *and all polynomials* $\tau = \tau(\lambda)$ *and* $k = k(\lambda)$ *that*

$$|\Pr[\mathsf{Exp}_{\mathsf{SRL}}^{(1)}(\mathcal{A}) = 1] - \Pr[\mathsf{Exp}_{\mathsf{SRL}}^{(0)}(\mathcal{A}) = 1]| \leq \mathsf{negl}(\lambda).$$

Since the SRL leakage depends on the specific circuit representation of the functions $(\psi_1 \ldots \psi_\tau)$, we propose a natural implementation for a class of functions that suffices to capture all possible leakage functions. Specifically, observe that ψ_i consist of a linear function (computed over $\mathbb{Z}_q$) followed by a rounding to the nearest multiple of $q/2$. Since all inputs are bit-wise encrypted the computation of modular additions (and multiplication by constants) is done via a canonical boolean circuit (see [11] for a concrete example) and the rounding is obtained by simply returning the ciphertext containing the most significant bit of the output.

4 Constructing (Shallow) XiO

In the following we present the construction of shallow XiO from the GSW scheme $(\mathsf{GSW.KeyGen}, \mathsf{GSW.Enc}, \mathsf{GSW.Eval}, \mathsf{GSW.Dec})$ and the packed version of the dual-Regev encryption $(\mathsf{dR.KeyGen}, \mathsf{dR.Enc}, \mathsf{dR.Eval}, \mathsf{dR.Dec})$ as described in Section 3.

4.1 Construction

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

- A sequence of randomization vectors $(\mathbf{r}_1^*, \dots, \mathbf{r}_{2^{\eta-\log(k)}}^*)$ for the GSW scheme $\mathsf{GSW.PubCoin}$, where each $\mathbf{r}_i^* = (\mathbf{r}_{i,1}^*, \dots, \mathbf{r}_{i,k}^*) \in [-\tilde{B}, +\tilde{B}]^{m \cdot k}$.
- A sequence of dense ciphertexts $(h_1, \dots, h_{2^{\eta-\log(k)}})$ for packed dual-Regev scheme $\mathsf{dR.PubCoin}$, where each $h_i = (h_{i,1}, \dots, h_{i,k}) \in \mathbb{Z}_q^k$.

On input the security parameter 1^λ and the circuit $\Pi : \{0,1\}^\eta \to \{0,1\}$, the obfuscator proceeds as follows.

Setting the Public Keys: Sample a dual-Regev key pair $(\bar{\mathsf{sk}}, \bar{\mathsf{pk}}) \leftarrow\!\!\$\ \mathsf{dR.KeyGen}(1^\lambda, 1^k)$ and GSW key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\!\$\ \mathsf{GSW.KeyGen}(1^\lambda; q)$, where q is the modulus defined by the dual-Regev scheme. Compute a bit-by-bit GSW encryption $\mathbf{c}_\Pi \leftarrow\!\!\$\ \mathsf{GSWEnc}(\mathsf{pk}, \Pi)$ of the binary representation of the circuit Π.

Compute a Key Encryption: Compute a dual-Regev extended encryption of the GSW secret key $\left(\bar{c}_{\mathsf{sk}}, \bar{\mathbf{C}}_{\mathsf{sk}}\right) = \mathsf{dR.ExtEnc}(\bar{\mathsf{pk}}, \mathsf{sk}; \mathbf{S})$. where $\mathsf{sk} \in \mathbb{Z}_q^{n+1}$ and $\mathbf{S} \leftarrow\!\!\$\ \mathbb{Z}_q^{n \times \log(q) \cdot k \cdot (n+1)}$.

Decryption Hints: For all indices $i \in \{0,1\}^{\eta - \log(k)}$, do the following.

Evaluate the Circuit: Let $\Phi_{i,j} : \{0,1\}^{|\Pi|} \to \{0,1\}$ be the universal circuit that, on input a circuit description Π, returns the j-th bit of the i-th block (where each block consists of k bits) of the corresponding truth table. Compute

$$\mathbf{C}_i = \begin{bmatrix} \mathsf{GSW.Eval}(\mathsf{pk}, \Phi_{i,1}, \mathbf{c}_\Pi) \\ \dots \\ \mathsf{GSW.Eval}(\mathsf{pk}, \Phi_{i,k}, \mathbf{c}_\Pi) \end{bmatrix} \in \mathbb{Z}_q^{k \times (n+1)}.$$

Compute the Low-Order Bits: Sample $\mathbf{r}_i \leftarrow\!\!\$\ \mathbb{Z}_q^n$ and compute $\mathbf{c}_{i,\mathbf{r}} \leftarrow\!\!\$\ \mathsf{GSW.Enc}(\mathsf{pk}, \mathbf{r}_i)$. Parse the i-th block of $\mathsf{dR.PubCoin}$ as

$$(h_{i,1}, \dots, h_{i,k}) = \mathbf{B} \cdot \mathbf{r}_i + (u_{i,1}, \dots, u_{i,k}) \in \mathbb{Z}_q^k$$

for some $(u_{i,1}, \dots, u_{i,k}) \in \mathbb{Z}_q^k$. Let $\Psi_{i,j} : \{0,1\}^\lambda \to \{0,1\}$ be the circuit that, on input $\mathbf{r}_i$, computes the decryption of the j-th bit encrypted in $(h_{i,1}, \dots, h_{i,k})$. I.e. it computes $\mathsf{Round}((h_{i,1}, \dots, h_{i,k}) - \mathbf{B} \cdot \mathbf{r}_i)$. Compute homomorphically the matrix of ciphertexts

$$\mathbf{C}_{i,\mathsf{Round}} = \begin{bmatrix} \mathsf{GSW.Eval}(\mathsf{pk}, \Psi_{i,1}, \mathbf{c}_{i,\mathbf{r}}) \\ \dots \\ \mathsf{GSW.Eval}(\mathsf{pk}, \Psi_{i,k}, \mathbf{c}_{i,\mathbf{r}}) \end{bmatrix} \in \mathbb{Z}_q^{k \times (n+1)}.$$

Rerandomize the Ciphertext: Parse the i-th block of $\mathsf{GSW.PubCoin}$ as $(\mathbf{r}_{i,1}^*, \dots, \mathbf{r}_{i,k}^*) \in [-\tilde{B}, +\tilde{B}]^{m \cdot k}$ and compute

$$\mathbf{C}'_{i,\mathsf{Round}} = \mathbf{C}_{i,\mathsf{Round}} + \begin{bmatrix} \mathsf{GSW.ColEnc}(\mathsf{pk}, 0; \mathbf{r}_{i,1}^*) \\ \dots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, 0; \mathbf{r}_{i,k}^*) \end{bmatrix} \in \mathbb{Z}_q^{k \times (n+1)}.$$

Proxy Re-Encrypt: Define $\mathbf{D}_i$ as the vector of GSW ciphertexts resulting from the homomorphic sum of $\mathbf{C}'_{i,\mathsf{Round}}$ and $\mathbf{C}_i$, i.e. $\mathbf{D}_i = \mathbf{C}'_{i,\mathsf{Round}} + \mathbf{C}_i$. Observe that $\mathbf{D}_i$ consists of k GSW ciphertexts and let $\boldsymbol{\ell}_{i,j} \in \mathbb{Z}_q^{(n+1)}$ be the linear function associated with the decryption of the j-th ciphertext. Define $\boldsymbol{\ell}_i = (\boldsymbol{\ell}_{i,1}, \ldots, \boldsymbol{\ell}_{i,k})$ and compute

$$\bar{\mathbf{c}}_i = \bar{\mathbf{C}}_{\mathsf{sk}} \cdot \mathsf{Bit}(\boldsymbol{\ell}_i) + (h_{i,1}, \ldots, h_{i,k}) \in \mathbb{Z}_q^k$$

where the function $\mathsf{Bit} : \mathbb{Z}_q^{k\cdot(n+1)} \to \{0,1\}^{\log(q)\cdot k\cdot(n+1)}$ is the bit decomposition operator.

Release Hint: Compute the i-th decryption hint as

$$\boldsymbol{\rho}_i = \mathbf{S} \cdot \mathsf{Bit}(\boldsymbol{\ell}_i) + \mathbf{r}_i \in \mathbb{Z}_q^n.$$

Output: The obfuscated circuit consists of the public keys $(\mathsf{pk}, \bar{\mathsf{pk}})$, the matrix $\bar{\mathbf{C}}_{\mathsf{sk}}$, the GSW encryption of the circuit $\mathbf{c}_\Pi$, the encryption headers $(\mathbf{c}_{1,\mathbf{r}}, \ldots, \mathbf{c}_{2^{\eta-\log(k)},\mathbf{r}})$, and the decryption hints $(\boldsymbol{\rho}_1, \ldots, \boldsymbol{\rho}_{2^{\eta-\log(k)}})$.

To evaluate the obfuscated circuit on input x, let i be the index of the block of the truth table of Π that contains $\Pi(x)$. The evaluator computes $\bar{\mathbf{c}}_i$ as specified above (note that all the operations are public, given the information included in the obfuscated circuit) and recovers $\Pi^{(i)}$ (the i-th block of the truth table of Π) by computing

$$\Pi^{(i)} = \mathsf{Round}(\bar{\mathbf{c}}_i - \mathbf{B} \cdot \boldsymbol{\rho}_i)$$

where $\mathsf{Round} : \mathbb{Z}_q^k \to \{0,1\}^k$ rounds the input to the nearest multiple of $q/2$.

Correctness. To see why the evaluation algorithm is correct, recall that

$$\bar{\mathbf{c}}_i = \bar{\mathbf{C}}_{\mathsf{sk}} \cdot \mathsf{Bit}(\boldsymbol{\ell}_i) + (h_{i,1}, \ldots, h_{i,k}).$$

First observe that $(\mathbf{r}_i, h_{i,1}, \ldots, h_{i,k})$ define a ciphertext in the support of the algorithm $\mathsf{dR.DenseEnc}(\bar{\mathsf{pk}})$, which we rewrite as

$$\begin{aligned}\mathsf{dR.DenseEnc}(\bar{\mathsf{pk}}) &= (\mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{r}_i), \mathbf{B} \cdot \mathbf{r}_i + (u_{i,1}, \ldots, u_{i,k})) \\ &= (\mathsf{PKE.Enc}(\mathsf{pk}_{\mathsf{PKE}}, \mathbf{r}_i), \mathbf{B} \cdot \mathbf{r}_i + \mathbf{u}_i)\,.\end{aligned}$$

Thus $\mathbf{C}'_{i,\mathsf{Round}}$ and $\mathbf{C}_i$ are in the support of

$$\begin{bmatrix}\mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_1)) \\ \ldots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_k))\end{bmatrix} \text{ and } \begin{bmatrix}\mathsf{GSW.ColEnc}(\mathsf{pk}, \Pi_1^{(i)}) \\ \ldots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, \Pi_k^{(i)})\end{bmatrix}$$

respectively, by the evaluation correctness of the GSW scheme and by Lemma 2. Furthermore, recall that $\mathbf{D}_i = \mathbf{C}'_{i,\mathsf{Round}} + \mathbf{C}_i$. By an invocation of Lemma 2, we have that $\mathbf{D}_i$ is in the support of

$$\begin{aligned}&\begin{bmatrix}\mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_1)) \\ \ldots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_k))\end{bmatrix} + \begin{bmatrix}\mathsf{GSW.ColEnc}(\mathsf{pk}, \Pi_1^{(i)}) \\ \ldots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, \Pi_k^{(i)})\end{bmatrix} \\ \approx &\begin{bmatrix}\mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_1) \oplus \Pi_1^{(i)}) \\ \ldots \\ \mathsf{GSW.ColEnc}(\mathsf{pk}, \mathsf{Round}(u_k) \oplus \Pi_k^{(i)})\end{bmatrix}\end{aligned}$$

with all but negligible probability. By the almost-linear decryption of GSW, it follows that

$$\bar{\mathbf{C}}_{\mathsf{sk}} \cdot \mathsf{Bit}(\boldsymbol{\ell}_i) = \mathbf{B} \cdot \tilde{\mathbf{s}}_i + \boldsymbol{\xi}_i + \boldsymbol{\zeta}_i + q/2 \cdot \left(\mathsf{Round}(u_1) \oplus \Pi_1^{(i)}, \dots, \mathsf{Round}(u_k) \oplus \Pi_k^{(i)}\right)$$

where $\boldsymbol{\xi}_i$ is the decryption noise of the packed dual-Regev scheme (i.e. the subset sum of the noise terms of $\bar{\mathbf{C}}_{\mathsf{sk}}$) and $\boldsymbol{\zeta}_i$ is the decryption noise of the GSW ciphertext. It follows that $\|\boldsymbol{\xi}_i\|_\infty \le B \cdot \log(q) \cdot k \cdot (n+1)$ and, by Lemma 2, $\|\boldsymbol{\zeta}_i\|_\infty \le \tilde{B}$ with all but negligible probability. Note that, by linearity we have that $\tilde{\mathbf{s}}_i = \mathbf{S} \cdot \mathsf{Bit}(\boldsymbol{\ell}_i)$. Consequently, it holds that

$$\begin{aligned}\bar{\mathbf{c}}_i &= \mathbf{B} \cdot \tilde{\mathbf{s}}_i + \boldsymbol{\xi}_i + \boldsymbol{\zeta}_i + q/2 \cdot \left(\mathsf{Round}(u_1) \oplus \Pi_1^{(i)}, \dots, \mathsf{Round}(u_k) \oplus \Pi_k^{(i)}\right) + \mathbf{B} \cdot \mathbf{r}_i + \mathbf{u}_i \\ &= \mathbf{B} \cdot (\tilde{\mathbf{s}}_i + \mathbf{r}_i) + \boldsymbol{\xi}_i + \boldsymbol{\zeta}_i + q/2 \cdot \left(\mathsf{Round}(u_1) \oplus \Pi_1^{(i)}, \dots, \mathsf{Round}(u_k) \oplus \Pi_k^{(i)}\right) + \mathbf{u}_i \\ &= \mathbf{B} \cdot (\tilde{\mathbf{s}}_i + \mathbf{r}_i) + q/2 \cdot \Pi^{(i)} + \mathbf{v}_i \\ &= \mathbf{B} \cdot \boldsymbol{\rho}_i + q/2 \cdot \Pi^{(i)} + \mathbf{v}_i\end{aligned}$$

where $\mathbf{v}_i = \mathbf{u}_i + q/2 \cdot \mathsf{Round}(\mathbf{u}_i) + \boldsymbol{\xi}_i + \boldsymbol{\zeta}_i$ and $\|\mathbf{v}_i\|_\infty < q/4$ with all but negligible probability, over the random choice of $\mathbf{u}_i$. This is because $\mathbf{D}_i$ is statistically close to a fresh GSW encryption of $(\mathsf{Round}(u_1), \dots, \mathsf{Round}(u_k)) \oplus \Pi^{(i)}$, by Lemma 4. Therefore we have that

$$\begin{aligned}\mathsf{Round}\left(\mathbf{c}_i - \mathbf{B} \cdot \boldsymbol{\rho}_i\right) &= \mathsf{Round}\left(\mathbf{B} \cdot \boldsymbol{\rho}_i + q/2 \cdot \Pi^{(i)} + \mathbf{v}_i - \mathbf{B} \cdot \boldsymbol{\rho}_i\right) \\ &= \mathsf{Round}\left(q/2 \cdot \Pi^{(i)} + \mathbf{v}_i\right) \\ &= \mathsf{Round}\left(q/2 \cdot \Pi^{(i)}\right) \\ &= \Pi^{(i)}\end{aligned}$$

with the same probability. Due to space constraints, we defer the analysis of our scheme to the full version [10].

References

1 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. doi:10.1007/978-3-030-17653-2_7.

2 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. doi:10.1007/978-3-030-26954-8_10.

3 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. doi:10.1007/978-3-319-56620-7_6.

4 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. doi:10.1007/3-540-44647-8_1.

5 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. doi:10.1007/978-3-030-03810-6_20.

6 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. doi:10.1007/978-3-662-44371-2_27.

7 Florian Bourse, Rafaël del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *Advances in Cryptology – CRYPTO 2016, Part II*, volume 9815 of *Lecture Notes in Computer Science*, pages 62–89, Santa Barbara, CA, USA, August 14–18 2016. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-53008-5_3.

8 Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 407–437, Nuremberg, Germany, December 1–5 2019. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-36033-7_16.

9 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. doi:10.1007/978-3-030-45721-1_4.

10 Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for io: Circular-secure lwe suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. URL: https://ia.cr/2020/1024.

11 Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd Annual Symposium on Foundations of Computer Science*, pages 97–106, Palm Springs, CA, USA, October 22–25 2011. IEEE Computer Society Press. doi:10.1109/FOCS.2011.12.

12 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. doi:10.1007/978-3-319-56617-7_10.

13 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. doi:10.1007/978-3-319-96881-0_20.

14 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. doi:10.1007/978-3-662-46800-5_1.

15 Jean-Sébastien Coron, Tancrède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 476–493, Santa Barbara, CA, USA, August 18–22 2013. Springer, Heidelberg, Germany. doi:10.1007/978-3-642-40041-4_26.

16 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. doi:10.1007/3-540-44586-2_9.

17 Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Pratyay Mukherjee. Obfuscation from low noise multilinear maps. In Debrup Chakraborty and Tetsu Iwata, editors, *Progress in Cryptology - INDOCRYPT 2018: 19th International Conference in Cryptology in India*, volume 11356 of *Lecture Notes in Computer Science*, pages 329–352, New Delhi, India, December 9–12 2018. Springer, Heidelberg, Germany. doi:10.1007/978-3-030-05378-9_18.

18 Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016, Part I*, volume 9665 of *Lecture Notes in Computer Science*, pages 294–310, Vienna, Austria, May 8–12 2016. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-49890-3_12.

19 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. doi:10.1007/978-3-642-38348-9_1.

20 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. doi:10.1007/978-3-642-54242-8_4.

21 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. doi:10.1109/FOCS.2013.13.

22 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. doi:10.1007/978-3-662-53644-5_10.

23 Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. Cryptology ePrint Archive, Report 2020/1010, 2020.

24 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. doi:10.1007/978-3-662-46497-7_20.

25 Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, May 17–20 2008. ACM Press. doi:10.1145/1374376.1374407.

26 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. doi:10.1007/978-3-642-40041-4_5.

27 Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In Salil P. Vadhan, editor, *TCC 2007: 4th Theory of Cryptography Conference*, volume 4392 of *Lecture Notes in Computer Science*, pages 194–213, Amsterdam, The Netherlands, February 21–24 2007. Springer, Heidelberg, Germany. doi:10.1007/978-3-540-70936-7_11.

28 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. doi:10.1007/3-540-44448-3_34.

29 Sam Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying io, 2021.

30 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. doi:10.1007/978-3-662-49890-3_21.

31 Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials overa $\mathbb{R}$ 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. doi:10.1007/978-3-030-17653-2_9.

32 Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. Cryptology ePrint Archive, Report 2020/1003, 2020.

33 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. doi:10.1007/978-3-662-49890-3_2.

34 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. doi:10.1007/978-3-319-63688-7_20.

35 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. doi:10.1007/978-3-662-49387-8_17.

36 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. doi:10.1007/978-3-319-63688-7_21.

37 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. doi:10.1109/FOCS.2016.11.

38 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. doi:10.1007/978-3-030-03810-6_19.

39 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. doi:10.1007/978-3-662-53008-5_22.

40 Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 536–553, Santa Barbara, CA, USA, August 17–21 2014. Springer, Heidelberg, Germany. doi:10.1007/978-3-662-44371-2_30.

41 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. doi:10.1145/2591796.2591825.

42 Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious lwe sampling. Cryptology ePrint Archive, Report 2020/1042, 2020.