

Candidate iO From Homomorphic Encryption Schemes

Zvika Brakerski^{1*}, Nico Döttling², Sanjam Garg^{3**}, and Giulio Malavolta^{4***}

¹ Weizmann Institute of Science

² CISA Helmholtz Center for Information Security

³ UC Berkeley

⁴ UC Berkeley & Carnegie Mellon University

Abstract. We propose a new approach to construct general-purpose indistinguishability obfuscation (iO). Our construction is obtained via a new intermediate primitive that we call *split fully-homomorphic encryption* (split FHE), which we show to be sufficient for constructing iO. Specifically, split FHE is FHE where decryption takes the following two-step syntactic form: (i) A *secret* decryption step uses the secret key and produces a *hint* which is (asymptotically) shorter than the length of the encrypted message, and (ii) a *public* decryption step that only requires the ciphertext and the previously generated hint (and not the entire secret key), and recovers the encrypted message. In terms of security, the hints for a set of ciphertexts should not allow one to violate semantic security for any other ciphertexts.

Next, we show a *generic candidate* construction of split FHE based on three building blocks: (i) A standard FHE scheme with linear decrypt-and-multiply (which can be instantiated with essentially all LWE-based constructions), (ii) a linearly homomorphic encryption scheme with short decryption hints (such as the Damgård-Jurik encryption scheme, based on the DCR problem), and (iii) a cryptographic hash function (which can be based on a variety of standard assumptions). Our approach is *heuristic* in the sense that our construction is not provably secure and makes implicit assumptions about the interplay between these underlying primitives. We show evidence that this construction is secure by providing an argument in an appropriately defined oracle model.

We view our construction as a big departure from the state-of-the-art constructions, and it is in fact quite simple.

* Supported by the Israel Science Foundation (Grant No. 468/14), Binational Science Foundation (Grants No. 2016726, 2014276) and European Union Horizon 2020 Research and Innovation Program via ERC Project REACT (Grant 756482) and via Project PROMETHEUS (Grant 780701).

** Supported in part from AFOSR Award FA9550-19-1-0200, AFOSR YIP Award, NSF CNS Award 1936826, DARPA and SPAWAR under contract N66001-15-C-4065, a Hellman Award and research grants by the Okawa Foundation, Visa Inc., and Center for Long-Term Cybersecurity (CLTC, UC Berkeley). The views expressed are those of the authors and do not reflect the official policy or position of the funding agencies.

*** Part of the work done while at the Simons Institute for the Theory of Computing.

1 Introduction

The goal of program obfuscation is to transform an arbitrary circuit C into an unintelligible but functionally equivalent circuit \tilde{C} . The notion of program obfuscation was first studied by Hada [39] and Barak et al. [10]. However, these works showed that natural notions of obfuscation are impossible to realize for general functionalities. Specifically, Barak et al. [10] defined a very natural notion of security for program obfuscation called virtual black-box (VBB) security, which requires that an obfuscated program does not reveal anything beyond what could be learned from just the input-output behavior of the original program. In the same work, they showed that this notion of program obfuscation is impossible to achieve for arbitrary circuits.

In light of this impossibility result, much of the work on obfuscation focused on realizing obfuscation for special functionalities. However, this changed with the work of Garg et al. [28] that proposed the first candidate indistinguishability obfuscation (iO) construction based on multilinear maps [26]. Furthermore, Garg et al. [28] showed powerful applications of iO to tasks such as functional encryption. Loosely speaking, iO requires that the obfuscations of two circuits C_0 and C_1 that have identical input output behavior are computationally indistinguishable. Subsequently, significant work on using program obfuscation (e.g., [55, 27, 16]) has shown that most cryptographic applications of interest can be realized using iO (and one-way functions), or that iO is virtually *crypto-complete*.

Given its importance, significant effort has been poured into realizing secure obfuscation candidates. The first approach to obfuscation relied on using new candidate constructions of multilinear maps [26, 22, 33], an algebraic object that significantly expands the structure available for cryptographic construction. Unfortunately, all multilinear map construction so far have relied on ad-hoc and new computational intractability assumptions. Furthermore, attacks [21, 40] on the multilinear map candidates and attacks [51, 20] on several of the multilinear map based iO constructions [28, 18, 9] were later found. In light of these attacks, follow up works (e.g., [31]) offered constructions that defended against these attacks by giving constructions in the so-called weak multilinear map model [51]. Several of these weak multilinear map model based iO constructions are still conjectured to be secure, however, the *break-and-repair* cycle of their development has left cryptographers wary, and rightly so.

Around the time when attacks on multilinear map candidates were at an all time high, cryptographers started exploring new approaches to iO without using multilinear maps (or reducing their usage). Toward this goal, Bitansky and Vaikunthanathan [15] and Ananth and Jain [4] showed that iO could be realized assuming just functional encryption. In another approach, instead of trying to remove multilinear maps completely, Lin [42] and Lin and Vaikunthanathan [47] attempted to reduce their usage, i.e., they proposed iO constructions using only constant degree multilinear maps. With the goal of ultimately basing iO constructions on standard assumptions on bilinear maps, cryptographers started developing new ideas for realizing iO candidates from smaller constant degree multilinear maps [5, 43]. Recently, Lin and Tessaro [45] described a candidate

iO construction from degree- L multilinear maps for any $L \geq 2$ and additionally assuming PRGs with certain special locality properties. Unfortunately, it was shown the needed PRGs for the case of $L = 2$ are insecure (in fact it was proved that they cannot exist) [48, 8]. Thus, still leaving a gap between bilinear maps and iO constructions which could now be based on trilinear maps [46]. Very recent works [1, 41, 3] (and cryptanalysis [12]), develop new ideas to resolve these problems and realize constructions based on bilinear maps. However, these bilinear map based constructions, which are still conjectured to be secure, additionally rely on certain pseudorandom objects with novel security properties. Finally, we note that all the other (perhaps less popular) approaches to iO (e.g., [35]) also start from new computational hardness assumptions.

Given the prior work, it is plausible that new sources of hardness are necessary for realizing iO candidates. Thus, this break-and-repair cycle would be necessary as we understand the underlying new assumptions better. In fact, there is some evidence that iO constructions based on simpler primitives [29, 30] are hard to realize. Making progress on this dilemma is the focus of this work.

1.1 Our Results

We propose a new approach to construct general-purpose indistinguishability obfuscation. Our approach is *heuristic* but *without* using any new sources of computational hardness. In other words, our constructions use well-studied cryptographic primitives in a generic way to realize obfuscation, while still being heuristic in the sense that our constructions are not provably secure and make implicit assumptions about the interplay of the underlying primitives. The primitives we use can themselves be securely realized based on standard assumptions, namely the hardness of the learning with errors (LWE) and the decisional composite residues (DCR) problem. At a high level, our heuristics are similar in flavor to (i) the random oracle heuristic that is often used in cryptographic constructions [13] and (ii) the circular security heuristic that has been widely used in the construction of fully-homomorphic encryption schemes (FHE) [32].

Split-FHE. The starting point of our work is the fact that iO can *provably* be based on *split FHE*, a new primitive that we introduce in this work. A split FHE is an FHE scheme that allows for certain special properties of the decryption algorithm. Specifically, we consider FHE schemes for which the decryption algorithm can be split into two subroutines:

- $\rho \leftarrow \text{PDec}(\text{sk}, c)$: A *private* procedure that takes the FHE secret key and a ciphertext as input and produces a decryption hint ρ , of size much smaller than the message encrypted in c .
- $m \leftarrow \text{Rec}(\rho, c)$: A *public* procedure that takes as input the decryption hint ρ (generated by PDec) and the ciphertext c and recovers the full plaintext.

The security for a split FHE scheme requires that, for all pairs of messages (m_0, m_1) and all circuits C such that $C(m_0) = C(m_1)$, the encryption of m_0

is computationally indistinguishable from the encryption of m_1 , even given the decryption hint for the ciphertext evaluated on C .

We show that split FHE alone suffices to construct exponentially-efficient iO [44], which in turn allows us to build fully-fledged iO. Concretely, we prove the following theorem.

Theorem 1 (Informal). *Assuming sub-exponentially hard LWE and the existence of sub-exponentially secure split FHE, then there exists indistinguishability obfuscation for all circuits.*

A Generic Candidate. Next, we show a *generic candidate* construction of split FHE based on three building blocks: (i) a standard FHE scheme with linear decrypt-and-multiply (which can be instantiated with essentially all LWE-based constructions), (ii) a linearly homomorphic encryption scheme with short decryption hints (such as the Damgård-Jurik encryption scheme [23], based on the DCR problem), and (iii) a cryptographic hash functions. The security of the scheme can be based on a new conjecture on the interplay of these primitives, which we view as a natural strengthening of circular security. In this sense, it is aligned with Gentry’s heuristic step in the FHE bootstrapping theorem [32]. Additionally, our use of the cryptographic hash function has similarities to the other heuristic uses of hash functions, e.g., in the Fiat-Shamir transformation [25].

We expect that there will exist instantiations of the underlying primitives (though contrived) for which this construction is insecure. For example, if the underlying schemes are not circular secure to begin with, then the resulting split FHE would also be insecure. However, for natural instantiations of these primitives, security can be conjectured.

Evidence of security. In order to build confidence in our construction, we show evidence that the above-mentioned conjecture on the interplay between the security holds in an appropriate oracle model, inspired by the random oracle model. Thus, pushing all the heuristic aspects of the construction to an oracle. In fact, we show that security can be proved in this oracle model.

An alternate way to think of this result is that we construct split FHE based on a obfuscation for a specific program (representing the oracle), for which we can offer a relatively simple and natural heuristic implementation.

Conceptual Simplicity. Another positive feature of our construction is its conceptual simplicity, which makes it much easier to analyze and thus have confidence in. Finally, we remark that our construction is a big departure from the previously-mentioned multilinear maps based and local PRG based iO constructions and will be accessible to readers without first understanding prior iO constructions.

1.2 Technical Overview

In the following we give an informal overview of the techniques we develop in this work and we refer the reader to the technical sections for more precise statements.

Chimeric FHE. Our starting point is the hybrid FHE scheme recently introduced by Brakerski et al. [17], which we recall in the following. The objective of their work is to build an FHE scheme with best possible rate (in an asymptotic sense) by leveraging the fact that most LWE-based FHE scheme admit an efficient *linear noisy decryption*. Specifically, given an FHE ciphertext c and an LWE secret key (s_1, \dots, s_n) one can rewrite the decryption operation as a linear function $L_c(\cdot)$ such that

$$L_c(s_1, \dots, s_n) = \text{ECC}(m) + e$$

where e is a B -bounded noise term and ECC is some encoding of the plaintext (in their scheme m is packed in the high-order bits so that it does not interfere with the noise term). The idea then is to encrypt the secret key (s_1, \dots, s_n) under a (high-rate) linearly homomorphic encryption (LHE) scheme, which allows one to compress evaluated FHE ciphertext by computing $L_c(\cdot)$ homomorphically.

One interesting property of this approach is that it is completely parametric in the choice of the schemes, as long as they satisfy some simple structural requirements: More concretely, one can use *any* LHE scheme as long as its plaintext domain matches the LWE modulus of the FHE scheme. As an example, one can set the LHE to be the Damgård-Jurik encryption scheme [23, 52], which we briefly recall in the following. The public key of the scheme consists of a large composite $N = pq$ and an integer ζ , and the encryption algorithm a message m computes

$$c = r^{N^\zeta} \cdot (1 + N)^m \pmod{N^{\zeta+1}}$$

for some uniform $r \leftarrow_{\$} \mathbb{Z}_N$. Note that the corresponding plaintext space is \mathbb{Z}_{N^ζ} and therefore the rate of the scheme approaches 1 as ζ grows. Furthermore, we observe that the scheme has one additional property that we refer to as *split decryption*. A scheme has split decryption if the decryption algorithm can be divided into a private and a public subroutine:

- The *private* procedure takes as input a ciphertext c and the secret key $\phi(N)$ and computes a *decryption hint*

$$\rho = c^{N^{-\zeta}} \pmod{N}$$

using the extended Euclidean algorithm. It is crucial to observe that $\rho \in \mathbb{Z}_N$ is potentially much smaller than the plaintext m .

- The *public* procedure takes as input a ciphertext c and the decryption hint ρ and recovers the plaintext by computing

$$(1 + N)^m = c/\rho^{N^\zeta} \pmod{N^{\zeta+1}}$$

and decoding m in polynomial time using the binomial theorem.

In a nutshell, the subgroup homomorphism allows one to compute a compressed version of the randomness, which can be then publicly stretched and used to

unmask the plaintext. This means that m can be fully recovered by communicating a small hint of size fixed and, in particular, independent of $|m|$. As we are going to discuss later, this property is going to be our main leverage to build general-purpose obfuscation.

Temporarily glossing over the security implications, we point out that the hybrid scheme of Brakerski et al. [17] already suffices to construct an FHE scheme with split decryption (in short, split FHE): Simply instantiate the LHE scheme with Damgård-Jurik and convert evaluated FHE ciphertexts before decryption using the algorithm described above.

Security for Split FHE. We now delve into the desired security property for a split FHE scheme. On a high level, we would like to ensure that the decryption hint does not reveal any additional information, beyond the plaintext of the corresponding ciphertext. It is instructive to observe that if we do not insist on this property, then every FHE scheme has a trivial split decryption procedure which simply outputs the secret key. We formalize this intuition as an indistinguishability definition that, roughly speaking, demands that for all plaintext pairs (m_0, m_1) and every set of circuits (C_1, \dots, C_β) such that $C_i(m_0) = C_i(m_1)$, then the encryption of m_0 and m_1 are computationally indistinguishable, even given the decryption hints ρ_i of the evaluated ciphertexts. The condition $C_i(m_0) = C_i(m_1)$ rules out trivial attacks where the distinguisher just checks the output of the evaluation. Here $\beta = \beta(\lambda)$ is an arbitrary (but a priori bounded) polynomial in the security parameter.

Unfortunately, our candidate as described above falls short in satisfying this security notion: The central problem is that our split decryption procedure reveals the complete plaintext encoded in the Damgård-Jurik ciphertext. This means that the distinguisher learns arbitrarily many relations of the form

$$L_{c_i}(s_1, \dots, s_n) = \text{ECC}(C_i(m_b)) + e_i$$

where c_i is the evaluated ciphertext and L_{c_i} is a publicly known linear function. Collecting a large enough sample allows the distinguisher to recompute the FHE secret key (s_1, \dots, s_n) via, e.g., Gaussian elimination. A standard approach to obviate this problem is to smudge the noise e_i with some mask r_i uniformly sampled from an exponentially larger domain. Thus, a natural solution would be to compute a randomizing ciphertext $d_i = \text{DJ.Enc}(\text{pk}_{\text{DJ}}, r_i)$ and output the decryption hint for

$$c_i \cdot d_i = \text{DJ.Enc}(\text{pk}_{\text{DJ}}, \text{ECC}(C_i(m_b)) + e_i + r_i) \approx \text{DJ.Enc}(\text{pk}_{\text{DJ}}, \text{ECC}(C_i(m_b)) + r_i)$$

where r_i is sampled from a domain exponentially larger than the noise bound B but small enough to allow one to decode $\text{ECC}(C_i(m_b))$. While it is possible to show that this approach indeed satisfies the security notion outlined above, it introduces an overhead in the size of the hint, which now consists of the pair (ρ_i, d_i) . Note that we cannot allow the distinguisher to recompute d_i locally as it is crucial that r_i remains hidden, so we have no other choice but append it to

the decryption hint. However the decryption hint is now of size $O(|c_i|)$, which does not satisfy our compactness requirement and makes our efforts purposeless (one can just set the decryption hint to be $C_i(m_b)$ and achieve better efficiency).

Although we appear to have encountered a roadblock, a closer look reveals that we still gained something from this approach: The ciphertext d_i encodes a (somewhat small) random value and in particular is completely independent from c_i . Furthermore, the decryption hint of $c_i \cdot d_i$ can be computed using the secret key alone. Assume for the moment that we had access to an oracle \mathcal{O} that outputs uniform Damgård-Jurik encryption of bounded random values, then our idea is to delegate the sampling of d_i to \mathcal{O} . This allows us to bypass the main obstacle: We do not need to include d_i in the decryption hint as it can be recomputed by querying \mathcal{O} . One can think of this approach as a more structured version of the Fiat-Shamir transform [25], which allows us to state the following theorem.

Theorem 2 (Informal). *Assuming the hardness of LWE and DCR, then there exists a split FHE scheme in the \mathcal{O} -hybrid model.*

Looking ahead to our end goal, another interpretation of this theorem is as a universality result: Assuming the hardness of LWE and DCR, we can bootstrap an obfuscator for a specific circuit (i.e., the one that samples a uniform Damgård-Jurik encryption of a bounded random value) to an obfuscator for all circuits.

Instantiating the Oracle. The most compelling question which arises from our main theorem is whether there exist plausible instantiations for the oracle \mathcal{O} . A first (flawed) attempt is to devise an oblivious sampling procedure for Damgård-Jurik ciphertext using a random oracle: Note that Damgård-Jurik ciphertexts live in a dense domain $\mathbb{Z}_{N^\zeta+1}$ and indeed sampling a random integer $c_i \leftarrow_s \mathbb{Z}_{N^\zeta+1}$ maps to a well-formed ciphertext with all but negligible probability. However, since c_i is uniform in the ciphertext domain, then so is the underlying plaintext $r_i \in \mathbb{Z}_{N^\zeta}$. This makes c_i unusable for our purposes since we require r_i to be bounded by some value \tilde{q} , which is exponentially smaller than N^ζ . If we were to sample r_i this way, then it would completely mask the term $\text{ECC}(C_i(m_b))$, thus making the plaintext impossible to decode.

Ideally, we would like to restrict the oblivious sampling to ciphertexts encrypting \tilde{q} -bounded messages. Unfortunately, we are not aware of the existence of any such algorithm. Instead, our idea is to still sample c_i uniformly over the complete ciphertext domain and remove the high-order bits of r_i *homomorphically*: This can be done by including an FHE encryption of the Damgård-Jurik secret key, then homomorphically evaluating the circuit that decrypts c_i and computes $-[r_i/\tilde{q}] \cdot \tilde{q}$. The evaluated ciphertext is then converted again to the Damgård-Jurik domain using the linear noisy decryption of the FHE scheme. At this point, one can obtain a well-formed encryption of a \tilde{q} -bounded value by computing

$$\begin{aligned} \text{DJ.Enc}(\text{pk}_{\text{DJ}}, -[r_i/\tilde{q}] \cdot \tilde{q} + e) \cdot c_i &= \text{DJ.Enc}(\text{pk}_{\text{DJ}}, -[r_i/\tilde{q}] \cdot \tilde{q} + e + r_i) \\ &= \text{DJ.Enc}(\text{pk}_{\text{DJ}}, (r_i \bmod \tilde{q}) + e) \end{aligned}$$

where the term $(r_i \bmod \tilde{q}) + e$ is \tilde{q} -bounded with all but negligible probability by setting $\tilde{q} \gg B$. While this approach brings us tantalizingly close to a provably secure scheme, a careful analysis highlights two lingering conjectures.

- (1) *Circular Security*: Adding and FHE encryption of the Damgård-Jurik secret key introduces a circular dependency in the security of the two schemes (recall that our construction already encodes a Damgård-Jurik encryption of the FHE secret key). While circular security falls outside of the realm of provable statements, it is widely accepted as a mild assumption and it is known to be achieved by most natural encryption schemes [11]. We stress that circular security is also inherent in the the bootstrapping theorem of Gentry [32], the only known method to construct fully (as opposed to leveled) homomorphic encryption from LWE.
- (2) *Correlations*: While the homomorphically evaluated circuit essentially ignores the low-order bits of r_i , the corresponding decryption noise e might still depend on $(r_i \bmod \tilde{q})$ in some intricate way. This might introduce some correlation and bias the distribution of the term $(r_i \bmod \tilde{q}) + e$ with respect to a uniform $u \leftarrow \mathbb{Z}_{\tilde{q}}$. However, the noise function is typically highly non-linear and therefore appears to be difficult to exploit. We also point out that the distinguisher has no control over the choice of e , which exclusively depends on an honest execution of the homomorphic evaluation algorithm. We therefore conjecture that the distribution of $(r_i \bmod \tilde{q}) + e$ is computationally indistinguishable from u .

In light of the above insights, we put forward the conjecture that the proposed algorithm already gives us a secure implementation of the oracle \mathcal{O} . We view this as a natural strengthening of Gentry’s heuristic for the bootstrapping theorem, which is justified by our more ambitious objective. As the conjecture pertains to standard cryptographic material (FHE and Damgård-Jurik encryption) we believe that any further insight on its veracity would substantially improve our understanding on these important and well-studied building blocks.

Finally, we mention that many heuristics can be used to weaken the correlation between the decryption noise e and the low-order bits $(r_i \bmod \tilde{q})$, such as repeated applications of FHE bootstrapping [24]. We also propose a different heuristic approach to remove correlations based on binary extractors and we refer the reader to the technical sections for further details.

From Split FHE to iO. What is left to be shown is that split FHE does indeed suffice to construct program obfuscation. With this goal in mind, we recall a surprising result by Lin et al. [44] which states that, under the assumption that the LWE problem is sub-exponentially hard, iO for all circuits is implied by an obfuscator for circuits with logarithmic-size inputs with non-trivial efficiency. Here non-trivial efficiency means that the size of the obfuscated circuit \tilde{C} with input domain $\{0, 1\}^\eta$ is at most $\text{poly}(\lambda, |\tilde{C}|) \cdot 2^{\eta \cdot (1-\epsilon)}$, for some constant $\epsilon > 0$. This means that it suffices to show that split FHE implies the existence of an obfuscator (for circuits with polynomial-size input domain) with non-trivial efficiency.

The transformation is deceptively simple (and similar to [14]): The obfuscator computes a split FHE encryption of the circuit C and partitions the input domains in $2^{\eta/2}$ disjoint sets $(P_1, \dots, P_{2^{\eta/2}})$ of equal size. Then, for each partition P_i , the algorithm homomorphically evaluates the universal circuit that evaluates C on all inputs in P_i and returns the concatenation of all outputs. Finally it returns the decryption hint ρ_i corresponding to the evaluated ciphertext. The obfuscated circuit consists of the public-key of the split FHE scheme, the encryption of C , and all of the decryption hints $(\rho_1, \dots, \rho_{2^{\eta/2}})$. Note that the obfuscated circuit can be evaluated efficiently: On input x , let P_x be the partition that contains x , then the evaluator recomputes the homomorphic evaluation (which is a deterministic operation) of C on P_x and recovers the output using the decryption hint ρ_x . As for non-trivial efficiency, since the size of each decryption hint is that of a fixed polynomial $\text{poly}(\lambda)$, the total size of the obfuscated circuit is bounded by $\text{poly}(\lambda, |C|) \cdot 2^{\eta/2}$, as desired.

Other Applications. To demonstrate that the scope of our split FHE scheme goes beyond program obfuscation, we outline two additional applications. In both cases we only rely on the hardness of the LWE and DCR problem, i.e., we do not need to introduce any new conjecture.

Two-Party Computation with Pre-Processing. We obtain a (semi-honest) two-party computation scheme for any circuit $C : \{0, 1\}^\ell \rightarrow \{0, 1\}^k$ with an input- and circuit-independent pre-processing where the communication complexity of the pre-processing phase is $\text{poly}(\lambda, k)$, whereas the communication complexity of the online phase is $\text{poly}(\lambda) + \ell$. This improves over garbled circuit-based approaches that require a pre-processing at least linear in $|C|$. The protocol works as follows: In the pre-processing phase Alice and Bob exchange their (independently sampled) public-keys for a split FHE scheme and Alice computes a randomizing ciphertext (in the scheme defined above this corresponds to a Damgård-Jurik encryption of a bounded random value), which is sent to Bob. In the online phase, Alice and Bob exchange their inputs encrypted under their own public keys (to achieve best-possible rate this can be done using hybrid encryption) and homomorphically compute the multi-key evaluation of f over both inputs. Note that multi-key evaluation is generically possible for the case of two parties by nesting the two split FHE evaluations. Then Alice consumes the randomizing ciphertext computed in the pre-processing and sends a partial decryption of the evaluated ciphertext in the form of a decryption hint. Bob can then locally complete the partial decryption using its own secret key and recover the output.

Rate-1 Reusable Garbled Circuits. The work of Goldwasser et al. [37] showed, assuming the hardness of the LWE problem, how to construct reusable garbled circuits where the size of the input encodings is $\text{poly}(\lambda, d, \ell \cdot k)$, where $C : \{0, 1\}^\ell \rightarrow \{0, 1\}^k$ and d is the depth of C . Additionally assuming the hardness of the DCR problem, we can bring down the complexity to $\text{poly}(\lambda, d, \ell) + k$.

This is done by using their scheme to garble the circuit that computes C homomorphically over the input encrypted under a split FHE scheme and returns the decryption hint of the evaluated ciphertext. This effectively removes the dependency of the underlying reusable garbled circuit on the output size k . However, we also need to include in the input encoding a randomizing Damgård-Jurik ciphertext, which reintroduces an additive overhead in k .

1.3 Related Work

In the following we discuss more in depth the relation of our approach when compared with recent candidate constructions of iO from lattices and bilinear maps [1, 41, 3]. Very informally, this line of works leverages weak pseudorandom generators (PRG) to mask the noise of the LWE decryption. However, the output domain of such a PRG is only polynomially large: This is because of the usage of bilinear groups, where the plaintext space is polynomially bounded (decryption requires one to solve a discrete logarithm). This is especially problematic because statistical/computational indistinguishability cannot hold in this regime of parameters. To circumvent this problem, all papers in this line of work assume a strict bound on the distinguisher’s success probability (e.g., 0.99) and then rely on amplification techniques. This however requires one to construct a weak PRG where the advantage of any PPT distinguisher is non-negligible but at the same time bounded by < 0.99 .

On the other hand, we rely on the Damgård-Jurik encryption scheme, where the message domain is exponential. This allows us to sample the smudging factor from a distribution that is exponentially larger than the noise bound, which is necessary in order to argue about statistical indistinguishability. Thus in our settings, conjecturing that the advantage of the distinguisher is negligible is, at least in principle, plausible.

2 Preliminaries

We denote by $\lambda \in \mathbb{N}$ the security parameter. We say that a function $\text{negl}(\cdot)$ is negligible if it vanishes faster than any polynomial. Given a set S , we denote by $s \leftarrow_s S$ the uniform sampling from S . We say that an algorithm is PPT if it can be implemented by a probabilistic machine running in time $\text{poly}(\lambda)$. We abbreviate the set $\{1, \dots, n\}$ as $[n]$. We recall the smudging lemma [6, 7].

Lemma 1 (Smudging). *Let $B_1 = B_1(\lambda)$ and $B_2 = B_2(\lambda)$ be positive integers and let $e_1 \in [B_1]$ be a fixed integer. Let $e_2 \leftarrow_s [B_2]$ 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 [28].

Definition 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 \mathcal{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

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

2.2 Learning with Errors

We recall the (decisional) learning with errors (LWE) problem as introduced by Regev [54].

Definition 2 (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 \mathcal{D} there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ it holds that

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

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 [54, 53], 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 .

3 Homomorphic Encryption

We recall the definition of homomorphic encryption in the following.

Definition 3 (Homomorphic Encryption). A homomorphic encryption scheme consists of the following efficient algorithms.

$\text{KeyGen}(1^\lambda)$: On input the security parameter 1^λ , the key generation algorithm returns a key pair (sk, pk) .

$\text{Enc}(\text{pk}, m)$: On input a public key pk and a message m , the encryption algorithm returns a ciphertext c .

$\text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell))$: On input the public key pk , an ℓ -inputs circuit C , and a vector of ciphertexts (c_1, \dots, c_ℓ) , the evaluation algorithm returns an evaluated ciphertext c .

$\text{Dec}(\text{sk}, c)$: On input the secret key sk and a ciphertext c , the decryption algorithm returns a message m .

We say that a scheme is fully-homomorphic (FHE) if it is homomorphic for all (unbounded) polynomial-size circuits. If the maximum size of the circuit that can be evaluated is bounded in the public parameters, then we call such a scheme a levelled FHE. We also consider a restricted class of homomorphism that supports linear functions and we refer to such a scheme as linearly-homomorphic encryption (LHE). We characterize correctness of a single evaluation, which suffices for our purposes. This can be extended to the more general notion of multi-hop correctness [34] if the condition specified below is required to hold for arbitrary compositions of circuits.

Definition 4 (Correctness). *A homomorphic encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{Dec})$ is correct if for all $\lambda \in \mathbb{N}$, all ℓ -inputs circuits C , all inputs (m_1, \dots, m_ℓ) , all (sk, pk) in the support of $\text{KeyGen}(1^\lambda)$, and all c_i in the support of $\text{Enc}(\text{pk}, m_i)$ it holds that*

$$\Pr[\text{Dec}(\text{sk}, \text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell))) = C(m_1, \dots, m_\ell)] = 1.$$

We require a scheme to be compact in the sense that the size of the ciphertext should not grow with the size of the evaluated circuit.

Definition 5 (Compactness). *A homomorphic encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{Dec})$ is compact if there exists a polynomial $\text{poly}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, all ℓ -inputs circuits C in the supported family, all inputs (m_1, \dots, m_ℓ) , all (sk, pk) in the support of $\text{KeyGen}(1^\lambda)$, and all c_i in the support of $\text{Enc}(\text{pk}, m_i)$ it holds that*

$$|\text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell))| = \text{poly}(\lambda) \cdot |C(m_1, \dots, m_\ell)|.$$

We define a weak notion of security (implied by the standard semantic security [38]) which is going to be more convenient to work with.

Definition 6 (Semantic Security). *A homomorphic encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{Dec})$ is semantically secure if for all polynomial-size 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*

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

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

3.1 Linear Decrypt-and-Multiply

We consider schemes with a fine-grained correctness property. Specifically, we require that the decryption consists of the application of a linear function in the secret key, followed by some publicly computable function. Furthermore, we require that such a procedure allows us to specify an arbitrary constant ω that is multiplied to the resulting plaintext. We refer to such schemes as linear decrypt-and-multiply schemes. This property was introduced in an oral presentation by Micciancio [50] and recently formalized by Brakerski et al. [17]. We stress that all major candidate FHE constructions satisfy (or can be adapted to) such a constraint, e.g., [36, 2, 19]. We recall the definition in the following.

Definition 7 (Decrypt-and-Multiply). *We call a homomorphic encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{Dec})$ a decrypt-and-multiply scheme, if there exists bounds $B = B(\lambda)$ and $Q = Q(\lambda)$ and an algorithm Dec\&Mult such that the following holds. For every $q \geq Q$, all (sk, pk) in the support of $\text{KeyGen}(1^\lambda, q)$, every ℓ -inputs circuit C , all inputs (m_1, \dots, m_ℓ) , all c_i in the support of $\text{Enc}(\text{pk}, m_i)$ and every $\omega \in \mathbb{Z}_q$ that*

$$\text{Dec\&Mult}(\text{sk}, \text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell)), \omega) = \omega \cdot C(m_1, \dots, m_\ell) + e \pmod q$$

where Dec\&Mult is a linear function in sk over \mathbb{Z}_q and $|e| \leq B$ with all but negligible probability.

In our construction, we will need some additional structure for the modulus q . Fortunately, most LWE-based FHE schemes can be instantiated with an arbitrary q that does not depend on any secret input but only on the security parameter. Moreover, LWE-based FHE schemes can be instantiated with any (sufficiently large) modulus q without affecting the worst-case hardness of the underlying LWE problem [53]. In an abuse of notation, we often write $\text{KeyGen}(1^\lambda; q)$ to fix the modulus q in the key generation algorithm. In favor of a simpler analysis, we assume that e is always non-negative. Note that this is without loss of generality as it can be always guaranteed by adding B to the result of Dec\&Mult and setting a slightly looser bound $B = 2B$.

3.2 Split Decryption

We define the notion of homomorphic encryption with split decryption, which is going to be central in our work. Loosely speaking, a scheme has split decryption if the decryption algorithm consists of two subroutines: A private algorithm (that depends on the secret key) that on input a ciphertext c computes a *small* hint ρ , and a publicly computable algorithm that takes as input ρ and c and returns the corresponding plaintext. We henceforth refer to such schemes as *split* homomorphic encryption. We introduce the syntax in the following.

Definition 8 (Split Decryption). *A homomorphic encryption scheme $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{Dec})$ has split decryption if the decryption algorithm Dec consist of the following two subroutines.*

$\text{PDec}(\text{sk}, c)$: On input the secret key sk and a ciphertext c , the partial decryption algorithm returns a decryption hint ρ .

$\text{Rec}(\rho, c)$: On input the hint ρ and a ciphertext c , the recovery algorithm returns a message m .

The notion of correctness is extended canonically.

Definition 9 (Split Correctness). A homomorphic encryption scheme with split decryption $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{PDec}, \text{Rec})$ is correct if for all $\lambda \in \mathbb{N}$, all ℓ -inputs circuits C in the supported family, all inputs (m_1, \dots, m_ℓ) , all (sk, pk) in the support of $\text{KeyGen}(1^\lambda)$, and all c_i in the support of $\text{Enc}(\text{pk}, m_i)$ it holds that

$$\Pr[\text{Rec}(\text{PDec}(\text{sk}, c), c) = C(m_1, \dots, m_\ell)] = 1$$

where $c = \text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell))$.

Beyond the standard compactness for homomorphic encryption, a scheme with split decryption must satisfy the additional property that the size of the decryption hint ρ is independent (or, more generally, sublinear) of the size of the message. Furthermore, the size of the public key and of a fresh encryption of a message m should depend polynomially in the security parameter and otherwise be linear in the size of the output. These are the properties that make split decryption non-trivial and that are going to be our main leverage to bootstrap this primitive into more powerful machinery. We formally characterize these requirements below.

Definition 10 (Split Compactness). A homomorphic encryption scheme with split decryption $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{PDec}, \text{Rec})$ is compact if there exists a polynomial $\text{poly}(\cdot)$ and such that for all $\lambda \in \mathbb{N}$, all ℓ -inputs circuits C in the supported family, all inputs (m_1, \dots, m_ℓ) , all (sk, pk) in the support of $\text{KeyGen}(1^\lambda)$, and all c_i in the support of $\text{Enc}(\text{pk}, m_i)$ it holds that

- $|\text{pk}| \leq \text{poly}(\lambda) \cdot |C(m_1, \dots, m_\ell)|$,
- $|c_i| \leq \text{poly}(\lambda, |m_i|) \cdot |C(m_1, \dots, m_\ell)|$, and
- $|\rho| \leq \text{poly}(\lambda)$

where $\rho = \text{PDec}(\text{sk}, \text{Eval}(\text{pk}, C, (c_1, \dots, c_\ell)))$.

Finally the notion of semantic security for split schemes requires that the decryption hint ρ for a certain ciphertext does not reveal any information beyond the corresponding plaintext. Note that we define a very weak notion where the above must hold only for a bounded number of ciphertexts, and the inputs are fixed prior to the public parameters of the scheme.

Definition 11 (Split Security). A homomorphic encryption scheme with split decryption $(\text{KeyGen}, \text{Enc}, \text{Eval}, \text{PDec}, \text{Rec})$ is secure if for all polynomial-size distinguishers \mathcal{D} there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$,

all polynomials $\beta = \beta(\lambda)$, all pairs of messages (m_0, m_1) , all vectors of circuits (C_1, \dots, C_β) such that, for all $i \in [\beta]$, $C_i(m_0) = C_i(m_1)$ it holds that

$$\left| \Pr [1 = \mathcal{D}(\mathbf{pk}, c_0, \rho_{(1,0)}, \dots, \rho_{(\beta,0)})] - \Pr [1 = \mathcal{D}(\mathbf{pk}, c_1, \rho_{(1,1)}, \dots, \rho_{(\beta,1)})] \right| = \text{negl}(\lambda)$$

where $(\mathbf{sk}, \mathbf{pk}) \leftarrow \text{KeyGen}(1^\lambda)$, for all $b \in \{0, 1\}$ define $c_b \leftarrow \text{Enc}(\mathbf{pk}, m_b)$ and, for all $i \in [\beta]$ and all $b \in \{0, 1\}$, define $\rho_{(i,b)} \leftarrow \text{PDec}(\mathbf{sk}, \text{Eval}(\mathbf{pk}, C_i, c_b))$.

3.3 Damgård-Jurik Encryption

In the following we recall a variant of the Damgård-Jurik encryption linearly homomorphic encryption scheme [23]. We present a variant of the scheme that satisfies the notion of split correctness, which is going to be instrumental for our purposes. The scheme is parametrized by a non-negative integer ζ that we assume is given as input to all algorithms.

DJ.KeyGen(1^λ): On input the security parameter 1^λ , sample a uniform Blum integer $N = pq$, where p and q are λ -bits primes. Set $\mathbf{pk} = (N, \zeta)$ and $\mathbf{sk} = \varphi(N)$.

DJ.Enc(\mathbf{pk}, m): On input a message $m \in \mathbb{Z}_{N^\zeta}$, sample a random $r \leftarrow_{\$} \mathbb{Z}_N$ and compute

$$c = r^{N^\zeta} \cdot (1 + N)^m \pmod{N^{\zeta+1}}.$$

DJ.Eval($\mathbf{pk}, f, (c_1, \dots, c_\ell)$): On input a vector of ciphertexts (c_1, \dots, c_ℓ) and a linear function $f = (\alpha_1, \dots, \alpha_\ell) \in \mathbb{Z}_{N^\zeta}^\ell$, compute

$$c = \prod_{i=1}^{\ell} c_i^{\alpha_i} \pmod{N^{\zeta+1}}.$$

DJ.PDec(\mathbf{sk}, c): On input a ciphertext c , set $s = c \pmod{N}$. Then compute $N^{-\zeta}$ such that $N^\zeta \cdot N^{-\zeta} = 1 \pmod{\varphi(N)}$ using the extended Euclidean algorithm. Return

$$\rho = s^{N^{-\zeta}} \pmod{N}.$$

DJ.Rec(ρ, c): On input a hint ρ and a ciphertext c , compute

$$(1 + N)^m = c / \rho^{N^\zeta} \pmod{N^{\zeta+1}}$$

and recover m using the polynomial-time algorithm described in [23].

It is well known that the scheme satisfies (standard) semantic security assuming the intractability of the decisional composite residuosity (DCR) problem, as defined in [52]. To prove correctness, we are going to use the fact that

$$x^{N^\zeta} \pmod{N^{\zeta+1}} = (x \pmod{N})^{N^\zeta} \pmod{N^{\zeta+1}} \quad (1)$$

for all non-negative integers (x, ζ) . We refer the reader to [49] for a proof of this equality. Recall that $c = r^{N^\zeta} \cdot (1 + N)^m$ and that

$$\begin{aligned} \rho &= (c \bmod N)^{N^{-\zeta}} \bmod N \\ &= \left(r^{N^\zeta} \cdot (1 + N)^m \bmod N \right)^{N^{-\zeta}} \bmod N \\ &= \left(r^{N^\zeta} \bmod N \right)^{N^{-\zeta}} \bmod N. \end{aligned}$$

Therefore we have that

$$\begin{aligned} \rho^{N^\zeta} \bmod N^{\zeta+1} &= \left(\left(r^{N^\zeta} \bmod N \right)^{N^{-\zeta}} \bmod N \right)^{N^\zeta} \bmod N^{\zeta+1} \\ &= \left(r^{N^\zeta} \bmod N \right)^{N^{-\zeta} \cdot N^\zeta} \bmod N^{\zeta+1} \\ &= r^{N^\zeta} \bmod N^{\zeta+1} \end{aligned}$$

by an application of Equation (1). Taking the inverse on both sides of the equation above we obtain

$$\begin{aligned} c/\rho^{N^\zeta} \bmod N^{\zeta+1} &= c/r^{N^\zeta} \bmod N^{\zeta+1} \\ &= r^{N^\zeta} \cdot (1 + N)^m / r^{N^\zeta} \bmod N^{\zeta+1} \\ &= (1 + N)^m \bmod N^{\zeta+1} \end{aligned}$$

as desired for correctness. Although such a scheme does not immediately give us a *secure* split LHE, we highlight a few salient properties that we are going to leverage in our main constructions.

Small Hints: The scheme satisfies a weakened notion of split compactness where the decryption hint is much smaller than the message space. The hint $\rho \in \mathbb{Z}_N$ consists of $\lceil \log(N) \rceil$ bits and in particular is independent of the size of the message space \mathbb{Z}_{N^ζ} , as the integer ζ can be set to be arbitrarily large (within the range of polynomials in λ).

Simulatable Hints: Given a ciphertext c and a plaintext value m , one can efficiently compute a ciphertext \tilde{c} such that the homomorphic sum of c and \tilde{c} results in a uniform encryption of m and the corresponding decryption hint can be computed given only the random coins used to generate \tilde{c} . Concretely, let

$$\tilde{c} = \frac{r^{N^\zeta} \cdot (1 + N)^m}{c} \bmod N^{\zeta+1}$$

for some $r \leftarrow_s \mathbb{Z}_N$, then $\rho = r$.

Dense Ciphertexts: Sampling a random integer in $\mathbb{Z}_{N^{\zeta+1}}$ gives a well-formed ciphertext with all but negligible probability. This is because the group order $\varphi(N) \cdot N^\zeta$ is close to $N^{\zeta+1}$, i.e., $\frac{\varphi(N) \cdot N^\zeta}{N^{\zeta+1}} = \frac{\varphi(N)}{N} = 1 - \text{negl}(\lambda)$.

4 Split Fully-Homomorphic Encryption

In the following we present our instantiation of FHE with split decryption. We first present a scheme from standard assumptions which assumes the existence of (a structured version of) a random oracle, then we propose plausible candidates for such an oracle.

4.1 Construction in the Presence of an Oracle

Before we delve into the details of our construction we give a definition of the oracle function that we consider. The oracle is parametrized by a pair of public keys for an FHE and an LHE scheme $(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}})$ and two integers (q, \tilde{q}) . On input a bitstring $x \in \{0, 1\}^*$, the oracle returns a uniform LHE encryption of a random value in \mathbb{Z}_q and an FHE encryption of the same value rounded to the closest divisor of \tilde{q} . The oracle is deterministic and it is accessible by all parties, thus on input the same x , the oracle will always output the same pair of ciphertexts. The interface is formally defined in the following.

$\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}(x)$: On input a string $x \in \{0, 1\}^*$ return two uniformly distributed ciphertexts

$$\text{LHE.Enc}(\text{pk}_{\text{LHE}}, m) \text{ and } \text{FHE.Enc}(\text{pk}_{\text{FHE}}, -\lfloor m/\tilde{q} \rfloor \cdot \tilde{q})$$

where $m \leftarrow_s \mathbb{Z}_q$.

It is useful to observe that the oracle output, along with an LHE encryption of the FHE secret key, gives us a uniformly distributed LHE encryption of a uniform value in $\mathbb{Z}_{\tilde{q}}$. This is because we can leverage the decrypt-and-multiply algorithm $\text{Dec}\&\text{Mult}$ of the FHE scheme (matching the FHE domain with the LHE plaintext space appropriately) to compute $\text{LHE.Enc}(\text{pk}_{\text{LHE}}, -\lfloor m/\tilde{q} \rfloor \cdot \tilde{q} + \textit{noise})$, where \textit{noise} is the decryption noise of the FHE scheme. Homomorphically summing up this term with the first output of the oracle we obtain

$$\begin{aligned} \text{LHE.Enc}(\text{pk}_{\text{LHE}}, m - \lfloor m/\tilde{q} \rfloor \cdot \tilde{q} + \textit{noise}) &= \text{LHE.Enc}(\text{pk}_{\text{LHE}}, (m \bmod \tilde{q}) + \textit{noise}) \\ &\approx_s \text{LHE.Enc}(\text{pk}_{\text{LHE}}, (m \bmod \tilde{q})) \end{aligned}$$

for an appropriate choice of \tilde{q} , i.e., we obtain an ciphertext which is statistically indistinguishable from an LHE encryption of a uniform element of $\mathbb{Z}_{\tilde{q}}$.

Description. We are now in the position of giving formal description of our scheme. We assume the existence of the following primitives:

- A fully-homomorphic encryption scheme $\text{FHE} = (\text{FHE.KeyGen}, \text{FHE.Enc}, \text{FHE.Eval}, \text{FHE.Dec})$ with linear decrypt-and-multiply and with noise bound B .
- A linearly homomorphic encryption $\text{LHE} = (\text{LHE.KeyGen}, \text{LHE.Enc}, \text{LHE.Eval}, \text{LHE.PDec}, \text{LHE.Rec})$ with small and simulatable decryption hints (e.g., the Damgård-Jurik encryption scheme as described in Section 3.3).

If the underlying FHE scheme is levelled then so is going to be the resulting split FHE. Conversely, if the FHE scheme supports the evaluation of unbounded circuits, then so does the resulting split FHE construction. The scheme is formally described in the following.

KeyGen(1^λ): On input the security parameter 1^λ , sample a key pair $(\text{sk}_{\text{LHE}}, \text{pk}_{\text{LHE}}) \leftarrow \text{LHE.KeyGen}(1^\lambda)$. Let \mathbb{Z}_q be the plaintext space defined by LHE, then sample $(\text{sk}_{\text{FHE}}, \text{pk}_{\text{FHE}}) \leftarrow \text{FHE.KeyGen}(1^\lambda; q)$. Let $\text{sk}_{\text{FHE}} = (s_1, \dots, s_n) \in \mathbb{Z}_q^n$, then return

$$\text{sk} = \text{sk}_{\text{LHE}} \quad \text{and} \quad \text{pk} = (\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, c_{(\text{LHE},1)}, \dots, c_{(\text{LHE},n)})$$

where, for all $i \in [n]$, we define $c_{(\text{LHE},i)} \leftarrow \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_i)$.

Enc(pk, m): On input a message m return

$$c \leftarrow \text{FHE.Enc}(\text{pk}_{\text{FHE}}, m).$$

Eval($\text{pk}, f, (c_1, \dots, c_\ell)$): On input a circuit C with ℓ bits of input and k bits of output and a vector of ciphertexts (c_1, \dots, c_ℓ) , let, for all $j \in [k]$, C_j be the circuit that returns the j -th bit of the output of C , then compute

$$d_j \leftarrow \text{FHE.Eval}(\text{pk}_{\text{FHE}}, C_j, (c_1, \dots, c_\ell)).$$

Define the following linear function over \mathbb{Z}_q :

$$g(x_1, \dots, x_n) = \sum_{j=1}^k \text{Dec\&Mult} \left((x_1, \dots, x_n), d_j, 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \right).$$

Compute $d \leftarrow \text{LHE.Eval}(\text{pk}_{\text{LHE}}, g, (c_{(\text{LHE},1)}, \dots, c_{(\text{LHE},n)}))$, then query $(a, \tilde{a}) \leftarrow \mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}(d)$ and define the following linear function over \mathbb{Z}_q :

$$\tilde{g}(x_1, \dots, x_n, x_{n+1}, x_{n+2}) = \text{Dec\&Mult} \left((x_1, \dots, x_n), \tilde{a}, 1 \right) + x_{n+1} + x_{n+2}.$$

Return

$$c \leftarrow \text{LHE.Eval}(\text{pk}_{\text{LHE}}, \tilde{g}, (c_{(\text{LHE},1)}, \dots, c_{(\text{LHE},n)}, d, a)).$$

PDec(sk, c): On input an evaluated ciphertext c return

$$\rho \leftarrow \text{LHE.PDec}(\text{sk}_{\text{LHE}}, c).$$

Rec(ρ, c): On input an evaluated ciphertext c , compute

$$\tilde{m} \leftarrow \text{LHE.Rec}(\rho, c)$$

and return the binary representation of \tilde{m} without its $\lceil \log(\tilde{q} + (k+1)B) \rceil$ least significant bits.

Analysis. We formally analyze our scheme in the following. During the analysis, we set the parameters on demand and we show afterwards that our choices lead to a satisfiable set of constraints for which the underlying computational problems are still conjectured to be hard. The following theorem establishes correctness.

Theorem 3 (Split Correctness). *Let $q \geq 2^k + 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil}$. Let FHE be a correct fully-homomorphic encryption scheme with linear decrypt-and-multiply and let LHE be a split correct linearly-homomorphic encryption scheme. Then the scheme as described above satisfies split correctness.*

Proof. Let us rewrite

$$\tilde{m} = \text{LHE.Rec}(\rho, c) = \text{LHE.Rec}(\text{LHE.PDec}(\text{sk}_{\text{LHE}}, c), c)$$

where $c = \text{LHE.Eval}(\text{pk}_{\text{LHE}}, \tilde{g}, (c_{(\text{LHE},1)}, \dots, c_{(\text{LHE},n)}), d, a)$. We first expand the d term as

$$\begin{aligned} d &= \text{LHE.Eval}(\text{pk}_{\text{LHE}}, g, (c_{(\text{LHE},1)}, \dots, c_{(\text{LHE},n)})) \\ &= \text{LHE.Eval}(\text{pk}_{\text{LHE}}, g, (\text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_1), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_n))) \\ &= \text{LHE.Enc} \left(\text{pk}_{\text{LHE}}, \sum_{j=1}^k \text{Dec\&Mult} \left((s_1, \dots, s_n), d_j, 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \right) \right) \end{aligned}$$

by the correctness of the LHE scheme, where

$$d_j = \text{FHE.Eval}(\text{pk}_{\text{FHE}}, C_j, (c_1, \dots, c_\ell))$$

and $c_i = \text{FHE.Enc}(\text{pk}_{\text{FHE}}, m_i)$. Thus by the decrypt-and-multiply correctness of the FHE scheme we can rewrite

$$\begin{aligned} d &= \text{LHE.Enc} \left(\text{pk}_{\text{LHE}}, \sum_{j=1}^k 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \cdot C_j(m_1, \dots, m_\ell) + e_j \right) \\ &= \text{LHE.Enc} \left(\text{pk}_{\text{LHE}}, \sum_{j=1}^k 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \cdot C_j(m_1, \dots, m_\ell) + \underbrace{\sum_{j=1}^k e_j}_{\tilde{\epsilon}} \right). \end{aligned}$$

For the a variable we have that $a = \text{LHE.Enc}(\text{pk}_{\text{LHE}}, r)$, for some uniform $r \leftarrow_s \mathbb{Z}_q$, by definition of the oracle $\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}$. Recall that

$$\tilde{g}(x_1, \dots, x_n, x_{n+1}, x_{n+2}) = \text{Dec\&Mult}((x_1, \dots, x_n), \tilde{a}, 1) + x_{n+1} + x_{n+2}.$$

where $\tilde{a} = \text{FHE.Enc}(\text{pk}_{\text{FHE}}, -\lfloor r/\tilde{q} \rfloor \cdot \tilde{q})$. Thus $c = \text{LHE.Enc}(\text{pk}_{\text{LHE}}, \tilde{m})$ where

$$\begin{aligned} \tilde{m} &= \text{Dec\&Mult}((s_1, \dots, s_n), \tilde{a}, 1) + \sum_{j=1}^k 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \cdot C_j(m_1, \dots, m_\ell) + \tilde{e} + r \\ &= -\lfloor r/\tilde{q} \rfloor \cdot \tilde{q} + e + \sum_{j=1}^k 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \cdot C_j(m_1, \dots, m_\ell) + \tilde{e} + r \\ &= \sum_{j=1}^k 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil + j} \cdot C_j(m_1, \dots, m_\ell) + \tilde{e} + e + \underbrace{(r \bmod \tilde{q})}_{\tilde{r}} \end{aligned}$$

by the correctness of the FHE scheme. Note that the sum $\tilde{e} + e$ is bounded from above by $(k+1) \cdot B$, whereas the term \tilde{r} is trivially bounded from above by \tilde{q} . This implies that the output of the circuit is encoded in the higher order bits of \tilde{m} with probability 1, for a large enough q .

We then argue about the split security of the scheme. We remark that we analyze security in the presence of an oracle and we refer the reader to Section 4.2 for concrete instantiations.

Theorem 4 (Split Security). *Let $\tilde{q} \geq 2^\lambda \cdot (k+1) \cdot B$ and let $q \geq 2^\lambda \cdot \tilde{q}$. Let FHE be a semantically secure fully-homomorphic encryption scheme and let LHE be a semantically secure linearly homomorphic encryption scheme with simulatable decryption hints. Then the scheme as described above satisfies split security in the $\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}$ -hybrid model.*

Proof. Let $(m_0, m_1, C_1, \dots, C_\beta)$ be the inputs specified by the adversary at the beginning of the experiments. Consider the following series of hybrids.

Hybrid \mathcal{H}_0 : Is defined as the original experiment. Denote the distribution induced by the random coins of the challenger by

$$(\text{pk}, c = \text{FHE.Enc}(\text{pk}_{\text{FHE}}, m_b), \rho_1, \dots, \rho_\beta)$$

where

$$\text{pk} = (\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_1), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_n))$$

and ρ_i is computed as $\text{PDec}(\text{sk}, \text{Eval}(\text{pk}, C_i, c))$.

Hybrids $\mathcal{H}_1 \dots \mathcal{H}_\beta$: Let $d^{(i)}$ be the variable d defined during the execution of $\text{Eval}(\text{pk}, C_i, c)$. The i -th hybrid \mathcal{H}_i is defined to be identical to \mathcal{H}_{i-1} , except that the oracle $\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}$ on input $d^{(i)}$ is programmed to output some a (along with a well-formed \tilde{a}) such that the resulting c is of the form

$$c = \text{LHE.Enc}(\text{pk}_{\text{LHE}}, \text{ECC}(C_i(m_b)) + \tilde{e} + e + r - \lfloor r/\tilde{q} \rfloor \cdot \tilde{q})$$

where ECC is the high-order bits encoding defined in the evaluation algorithm, $\tilde{e} + e$ is the sum of the decryption noises of the ciphertexts $(d^{(1)}, \dots, d^{(k)}, \tilde{a})$, as defined in the evaluation algorithm, and $r \leftarrow_{\mathcal{S}} \mathbb{Z}_q$. Then $\tilde{\rho}_i$ is defined to be the decryption hint of c computed using the random coins of a .

First observe that $\tilde{e} + e$ is efficiently computable given the secret key of the FHE scheme and therefore $\tilde{\rho}_i$ is also computable in polynomial time. It is important to observe that the distribution of c is identical to the previous hybrid and the difference lies only in the way $\tilde{\rho}_i$ is computed. Since the LHE scheme has simulatable hints, it follows that the distribution of \mathcal{H}_i is identical to that of \mathcal{H}_{i-1} and the change described here is only syntactical. That is,

$$\begin{aligned} & (\text{pk}, \text{FHE.Enc}(\text{pk}_{\text{FHE}}, m_b), \tilde{\rho}_1, \dots, \tilde{\rho}_{i-1}, \rho_i, \rho_{i+1}, \dots, \rho_\beta) \\ &= (\text{pk}, \text{FHE.Enc}(\text{pk}_{\text{FHE}}, m_b), \tilde{\rho}_1, \dots, \tilde{\rho}_{i-1}, \tilde{\rho}_i, \rho_{i+1}, \dots, \rho_\beta). \end{aligned}$$

Hybrids $\mathcal{H}_{\beta+1} \dots \mathcal{H}_{2\beta}$: The $(\beta + i)$ -th hybrid differs from the previous one in the sense that a is programmed such that

$$\begin{aligned} c &= \text{LHE.Enc}(\text{pk}_{\text{LHE}}, \text{ECC}(C_i(m_b))) + \tilde{e} + e + \lfloor r/\tilde{q} \rfloor \cdot \tilde{q} + \tilde{r} - \lfloor r/\tilde{q} \rfloor \cdot \tilde{q} \\ &= \text{LHE.Enc}(\text{pk}_{\text{LHE}}, \text{ECC}(C_i(m_b))) + \tilde{e} + e + \tilde{r} \end{aligned}$$

where $\tilde{r} \leftarrow_{\mathcal{S}} \mathbb{Z}_{\tilde{q}}$. Note that the distributions induced by the two hybrids differ only in case where $r \in R$, where $R = \{q - (q \bmod \tilde{q}), \dots, q\}$. Since $\tilde{q}/q \leq 2^{-\lambda}$ we have that the two distributions are statistically close.

Hybrids $\mathcal{H}_{2\beta+1} \dots \mathcal{H}_{3\beta}$: The $(\beta + i)$ -th hybrid is defined to be identical to the previous ones except that a is programmed such that

$$c = \text{LHE.Enc}(\text{pk}_{\text{LHE}}, \text{ECC}(C_i(m_b)) + \tilde{r}).$$

I.e., the noise term \tilde{e} is omitted from the computation. Thus the only difference with respect to the previous hybrid is whether the noise term $\tilde{e} + e$ is included in the ciphertext or not. Since $\tilde{e} + e$ is bounded from above by $(k + 1) \cdot B$ and $\tilde{q} \geq 2^\lambda \cdot (k + 1) \cdot B$, by Lemma 1 the distribution induced by this hybrid is statistically indistinguishable from that of the previous one.

Hybrids $\mathcal{H}_{3\beta+1} \dots \mathcal{H}_{3\beta+n}$: The $(3\beta + i)$ -th hybrid is defined as the previous one, except that the ciphertext $c_{(\text{LHE}, i)}$ in the public parameters is computed as the encryption of 0. Note that the secret key of the LHE scheme is no longer used in the computation of $(\tilde{\rho}_1, \dots, \tilde{\rho}_\beta)$ and therefore indistinguishability follows from an invocation of the semantic security of the LHE scheme. Specifically, the following distributions are computationally indistinguishable

$$\begin{aligned} & \left(\text{LHE.Enc}(\text{pk}_{\text{LHE}}, 0), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, 0), \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_i), \right. \\ & \left. \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_{i+1}), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_n) \right) \\ & \approx_c \left(\text{LHE.Enc}(\text{pk}_{\text{LHE}}, 0), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, 0), \text{LHE.Enc}(\text{pk}_{\text{LHE}}, 0), \right. \\ & \left. \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_{i+1}), \dots, \text{LHE.Enc}(\text{pk}_{\text{LHE}}, s_n) \right). \end{aligned}$$

Hybrid $\mathcal{H}_{3\beta+n}^{(b)}$: We define the hybrid $\mathcal{H}_{3\beta+n}^{(b)}$ as $\mathcal{H}_{3\beta+n}$ with the challenger bit fixed to b . Note that the distribution induced by these hybrids is

$$(\mathbf{pk}, c = \text{FHE.Enc}(\mathbf{pk}_{\text{FHE}}, m_b), \tilde{\rho}_1, \dots, \tilde{\rho}_\beta)$$

where

$$\mathbf{pk} = (\mathbf{pk}_{\text{FHE}}, \mathbf{pk}_{\text{LHE}}, \text{LHE.Enc}(\mathbf{pk}_{\text{LHE}}, 0), \dots, \text{LHE.Enc}(\mathbf{pk}_{\text{LHE}}, 0)).$$

Observe that the secret key of the FHE scheme is no longer encoded in the public parameters and is not needed to compute $(\tilde{\rho}_1, \dots, \tilde{\rho}_\beta)$ either. It follows that any advantage that the adversary has in distinguishing $\mathcal{H}_{3\beta+n}^{(0)}$ from $\mathcal{H}_{3\beta+n}^{(1)}$ cannot be greater than the advantage in distinguishing $\text{FHE.Enc}(\mathbf{pk}_{\text{FHE}}, m_0)$ from $\text{FHE.Enc}(\mathbf{pk}_{\text{FHE}}, m_1)$. Thus, computational indistinguishability follows from an invocation of the semantic security of the FHE scheme. This concludes our proof.

Parameters. When instantiating the LHE scheme with the Damgård-Jurik encryption scheme (as described in Section 3.3) and the FHE scheme with any LWE-based scheme with linear decrypt-and-multiply (e.g., the scheme proposed in [36]) we obtain a split FHE which satisfies the notion of split compactness: The hint ρ is of size $N = \text{poly}(\lambda)$ and in particular is arbitrarily smaller than the size of the plaintext space $q = N^\zeta$. For essentially any choice of the LWE-based FHE scheme with modulus q , the size of the public key and fresh ciphertexts depends polynomially in λ and linearly in $\log(q) = \log(N^\zeta)$, which gives us the desired bound. The analysis above sets the following additional constraints:

- $q \geq 2^k + 2^{\lceil \log(\tilde{q} + (k+1)B) \rceil}$,
- $q \geq 2^\lambda \cdot \tilde{q}$, and
- $\tilde{q} \geq 2^\lambda \cdot (k+1) \cdot B$

which are always satisfied for $q = N^\zeta$, by setting the integer ζ to be large enough. Note that this choice of parameters fixes the modulus of the FHE with linear decrypt-and-multiply to \mathbb{Z}_{N^ζ} , which is super-polynomially larger than the noise bound B . Finally, the LWE parameter n is free and can be set to any value for which the corresponding problem (with super-polynomial modulus-to-noise ratio) is conjectured to be hard.

4.2 Instantiating the Oracle

To complete the description of our scheme, we discuss a few candidate instantiations for the oracle $\mathcal{O}_{(\mathbf{pk}_{\text{FHE}}, \mathbf{pk}_{\text{LHE}}, q, \tilde{q})}$. We require the underlying LHE scheme to have a dense ciphertext domain (which is the case for the Damgård-Jurik encryption scheme). Both of our proposal introduce new circularity assumptions between the FHE and the LHE schemes.

An alternate way to think of the oracle in Theorem 4 is to see it as an obfuscation for a special program, which is sufficient for realizing split FHE. The candidate constructions that we provide below can be seen as a very natural and simple obfuscation of this special program.

A Simple Candidate. Let \mathfrak{C} be the ciphertext domain of LHE. Our first instantiation hardwires an FHE encryption of the LHE secret key $c_{\text{FHE}} \leftarrow \text{FHE.Enc}(\text{pk}_{\text{FHE}}, \text{sk}_{\text{LHE}})$. We fix the random coins of the algorithm (whenever needed) by drawing them from the evaluation of a cryptographic hash function Hash over the input. The intuition for our candidate is very simple: The LHE ciphertext is obviously sampled without the knowledge of the underlying plaintext (which is the reason why we need dense ciphertexts) whereas the FHE term is computed by evaluating the decryption circuit homomorphically and rounding the resulting message to the closest multiple of \tilde{q} .

$\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}(x)$: On input a string $x \in \{0, 1\}^*$ sample $y \leftarrow_{\text{s}} \mathfrak{C}$, using $\text{Hash}(x)$ as the random coins, then compute

$$\tilde{y} \leftarrow \text{FHE.Eval}(\text{pk}_{\text{FHE}}, - \lfloor \text{LHE.Dec}(\cdot, y) / \tilde{q} \rfloor \cdot \tilde{q}, c_{\text{FHE}})$$

and return (y, \tilde{y}) .

Observe that y is an element in the ciphertext domain of LHE and it is of the form $y = \text{LHE.Enc}(\text{pk}_{\text{LHE}}, m)$, for some $m \in \mathbb{Z}_q$, since LHE has a dense ciphertext domain. Furthermore, by the correctness of the FHE and the LHE scheme, we have that

$$\begin{aligned} \tilde{y} &= \text{FHE.Eval}(\text{pk}_{\text{FHE}}, - \lfloor \text{LHE.Dec}(\cdot, y) / \tilde{q} \rfloor \cdot \tilde{q}, c_{\text{FHE}}) \\ &= \text{FHE.Eval}(\text{pk}_{\text{FHE}}, - \lfloor \text{LHE.Dec}(\cdot, y) / \tilde{q} \rfloor \cdot \tilde{q}, \text{FHE.Enc}(\text{pk}_{\text{FHE}}, \text{sk}_{\text{LHE}})) \\ &= \text{FHE.Enc}(\text{pk}_{\text{FHE}}, - \lfloor \text{LHE.Dec}(\text{sk}_{\text{LHE}}, y) / \tilde{q} \rfloor \cdot \tilde{q}) \\ &= \text{FHE.Enc}(\text{pk}_{\text{FHE}}, - \lfloor m / \tilde{q} \rfloor \cdot \tilde{q}). \end{aligned}$$

It follows that the pair (y, \tilde{y}) is syntactically well formed. However, a closer look to the oracle instantiation reveals two lingering assumptions.

- (1) *Circular Security*: The addition of $c_{\text{FHE}} = \text{FHE.Enc}(\text{pk}_{\text{FHE}}, \text{sk}_{\text{LHE}})$ introduces a circular dependency in the security of the LHE and FHE schemes (recall that our split FHE construction includes in the public key an encryption of sk_{FHE} under pk_{LHE}). Circular security is however widely considered to be a very mild assumption and currently is the only known approach to construct plain (as opposed to levelled) FHE from LWE via the bootstrapping theorem [32].
- (2) *Correlations*: Although \tilde{y} is an FHE encryption of the correct value, it is not necessarily uniformly distributed, conditioned on y . In particular the randomness of \tilde{y} may depend in some intricate way on the low-order bits of m . For the specific case of LWE-based schemes, the *noise* term might carry some information about $m \bmod \tilde{q}$, which could introduce some harmful correlation. However, the noise function is typically highly non-linear and therefore appears to be difficult to exploit. We also stress that we only consider honest executions of the FHE.Eval algorithm.

While (1) can be regarded as a standard assumption, we view (2) as a natural conjecture which we believe holds true for any natural/known candidate

instantiation of the FHE and LHE schemes. In light of these considerations, we conjecture that the implementation as describe above already leads to a secure split FHE scheme.

Towards Removing Correlations. A natural approach towards removing the correlation of the LHE and FHE ciphertexts is that of ciphertext sanitization [24]: One could expect that repeatedly bootstrapping the FHE ciphertext would decorrelate the noise from the companion LHE ciphertext. Unfortunately our settings are different than those typically considered in the literature, in the sense that the sanitization procedure must be carried out by the distinguisher and therefore cannot use private random coins. Although it appears hard to formally analyze the effectiveness of these methods in our settings, we expect that these techniques might (at least heuristically) help to obliterate harmful correlations. In this work we take a different route and we suggest a simple heuristic method to prevent correlations. In a nutshell, the idea is to sample a set of random plaintexts and define the random string as the sum of a uniformly sampled subset S of these plaintext. The key observation is that subset sum is a linear operation and therefore can be performed directly in the LHE scheme, which implies that the *leakage* of the FHE scheme cannot depend on S . As for the previous construction, our instantiation contains and a ciphertext $c_{\text{FHE}} = \text{FHE.Enc}(\text{pk}_{\text{FHE}}, \text{sk}_{\text{LHE}})$. The scheme is parametrized by some $\sigma \in \text{poly}(\lambda)$, which defines the size of the set S . In the following description we present the algorithm as randomized, although this simplification can be easily bypassed with standard techniques (e.g., computing the random coins using a cryptographic hash $\text{Hash}(x)$).

$\mathcal{O}_{(\text{pk}_{\text{FHE}}, \text{pk}_{\text{LHE}}, q, \tilde{q})}(x)$: On input a string $x \in \{0, 1\}^*$ sample a random set $S \leftarrow_{\$} \{0, 1\}^\sigma$. Then, for all $i \in [\sigma]$, do the following:

- If $S_i = 1$, sample a uniform $y_i \leftarrow_{\$} \mathcal{C}$.
- If $S_i = 0$, sample a uniform encryption $y_i \leftarrow_{\$} \text{LHE.Enc}(\text{pk}_{\text{LHE}}, m_i)$, for a random known m_i .

Then compute

$$\tilde{y} \leftarrow \text{FHE.Eval} \left(\text{pk}_{\text{FHE}}, - \sum_{i=1}^{\sigma} [\text{LHE.Dec}(\cdot, y_i) / \tilde{q}] \cdot \tilde{q}, c_{\text{FHE}} \right).$$

Let f be the following linear function

$$f(x_1, \dots, x_{|S|}) = \sum_{i \in S} x_i + \sum_{i \notin S} [m_i / \tilde{q}] \cdot \tilde{q}$$

then compute $y \leftarrow \text{LHE.Eval}(\text{pk}_{\text{LHE}}, f, \{y_i\}_{i \in S})$ and return (y, \tilde{y}) .

To see why the implementation is syntactically correct, observe that

$$\begin{aligned}
\tilde{y} &= \text{FHE.Eval} \left(\text{pk}_{\text{FHE}}, - \sum_{i=1}^{\sigma} \lfloor \text{LHE.Dec}(\cdot, y_i) / \tilde{q} \rfloor \cdot \tilde{q}, c_{\text{FHE}} \right) \\
&= \text{FHE.Enc} \left(\text{pk}_{\text{FHE}}, - \sum_{i=1}^{\sigma} \lfloor \text{LHE.Dec}(\text{sk}_{\text{LHE}}, y_i) / \tilde{q} \rfloor \cdot \tilde{q} \right) \\
&= \text{FHE.Enc} \left(\text{pk}_{\text{FHE}}, - \sum_{i=1}^{\sigma} \lfloor m_i / \tilde{q} \rfloor \cdot \tilde{q} \right)
\end{aligned}$$

by the evaluation correctness of the FHE scheme. Invoking to the correctness of the LHE scheme we have that

$$\begin{aligned}
y &= \text{LHE.Eval} (\text{pk}_{\text{LHE}}, f, \{y_i\}_{i \in S}) \\
&= \text{LHE.Eval} (\text{pk}_{\text{LHE}}, f, \{\text{LHE.Enc}(\text{pk}_{\text{LHE}}, m_i)\}_{i \in S}) \\
&= \text{LHE.Enc} \left(\text{pk}_{\text{LHE}}, \sum_{i \in S} m_i + \sum_{i \notin S} \lfloor m_i / \tilde{q} \rfloor \cdot \tilde{q} \right) \\
&= \text{LHE.Enc} \left(\text{pk}_{\text{LHE}}, \underbrace{\sum_{i \in S} (m_i \bmod \tilde{q})}_{\tilde{m}} + \sum_{i=1}^{\sigma} \lfloor m_i / \tilde{q} \rfloor \cdot \tilde{q} \right)
\end{aligned}$$

which is exactly what we want, except that \tilde{m} is slightly larger than \tilde{q} , by a factor of at most σ . This can still be used in our main construction by adjusting the error correcting code accordingly. The intuition why we believe that this variant is secure is that the leakage in the FHE randomness cannot depend on the set S , since the distributions of all y_i are statistically close (recall that LHE has dense ciphertexts). Thus, S (which is chosen uniformly) resembles the behavior of a binary extractor on $(m_i \bmod \tilde{q})$. Nevertheless, proving a formal statement remains an interesting open question.

5 Split Fully-Homomorphic Encryption \implies Obfuscation

In order to construct fully-fledged iO from split FHE, we rely on a theorem from Lin et al. [44], which we recall in the following. Roughly speaking, the theorem states that, under the assumption that the LWE problem is sub-exponentially hard, it suffices to consider circuits with a polynomial-size input domain and obfuscators that output obfuscated circuits of size slightly sublinear in size of the truth table of the circuit.

Theorem 5 ([44]). *Assuming sub-exponentially hard LWE, if there exists a sub-exponentially secure indistinguishability obfuscator for $\text{P}^{\log}/\text{poly}$ with non-trivial efficiency, then there exists an indistinguishability obfuscator for P/poly with sub-exponential security.*

Here $\text{P}^{\log}/\text{poly}$ 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, which can be as large as $\text{poly}(\lambda, |C|) \cdot 2^\eta$.

In the following we show how to construct an obfuscator for $\text{P}^{\log}/\text{poly}$ with non-trivial efficiency. We assume only the existence of a (levelled) split FHE scheme $\text{sFHE} = (\text{KeyGen}, \text{Enc}, \text{Eval}, \text{PDec}, \text{Rec})$.

$\text{iO}(C)$: On input the description of a circuit C , sample a fresh key pair $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(1^\lambda)$ and compute $c \leftarrow \text{Enc}(\text{pk}, C)$. For all $i \in [2^{\eta/2}]$ define the universal circuit \mathfrak{U}_i as

$$\mathfrak{U}_i(C) = C\left((i-1) \cdot 2^{\eta/2}\right) \parallel \dots \parallel C\left(i \cdot 2^{\eta/2} - 1\right).$$

Then compute $c_i \leftarrow \text{Eval}(\text{pk}, \mathfrak{U}_i, c)$ and $\rho_i \leftarrow \text{PDec}(\text{sk}, c_i)$. The obfuscated circuit is defined to be $(\text{pk}, c, \rho_1, \dots, \rho_{2^{\eta/2}})$.

First we discuss how to evaluate an obfuscated circuit: On input some $x \in \{0, 1\}^\eta$, parse it as an integer and round it to the nearest multiple of $2^{\eta/2}$ (let such integer be \bar{x}) such that $\bar{x} \leq x$. Then compute $c_{\bar{x}} \leftarrow \text{Eval}(\text{pk}, \mathfrak{U}_{\bar{x}}, c)$ and $m \leftarrow \text{Rec}(\rho_{\bar{x}}, c_{\bar{x}})$. Read the output as the $(x - \bar{x})$ -th bit of m .

Analysis. Note that the runtime of the obfuscator is dominated by $2^{\eta/2}$ evaluations of the split FHE ciphertext, where each subroutine homomorphically evaluates the circuit C $2^{\eta/2}$ -many times. Thus the total runtime of the obfuscator is in the order of $\text{poly}(\lambda, |C|) \cdot 2^\eta$. We now argue that our obfuscator has non-trivial efficiency in terms of output size. We analyze the size of each component of the obfuscated circuit:

- By the compactness of the split FHE scheme, the public key pk grows linearly with the size of the output domain, i.e., $2^{\eta/2}$, and polynomially in the security parameter.
- The ciphertext c grows linearly with the size of the encrypted message and therefore, by the compactness of the split FHE scheme, bounded by $\text{poly}(\lambda, |C|) \cdot 2^{\eta/2}$.
- Each decryption hint ρ_i is of size $\text{poly}(\lambda)$, since the underlying split FHE is compact. As an obfuscated circuit consists of $2^{\eta/2}$ -many decryption hints, the size of the vector $(\rho_1, \dots, \rho_{2^{\eta/2}})$ is $\text{poly}(\lambda) \cdot 2^{\eta/2}$.

It follows that the total size of the obfuscated circuit is bounded from above by $\text{poly}(\lambda, |C|) \cdot 2^{\eta/2}$. What is left to be shown is that our obfuscator satisfies the notion of indistinguishability obfuscation.

Theorem 6 (Indistinguishability Obfuscation). *Let sFHE be a sub-exponentially secure levelled split FHE scheme. Then the scheme as described above is a sub-exponentially secure indistinguishability obfuscator.*

Proof. By the perfect correctness of the split FHE scheme it follows that the obfuscated circuit is functionally equivalent to the plain circuit. Indistinguishability follows immediately from the split security of sFHE: If the split FHE is secure against a distinguisher running in sub-exponential time, then so is iO.

References

1. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 191–225. Springer, Heidelberg, May 2019.
2. Jacob Alperin-Sheriff and Chris Peikert. Faster bootstrapping with polynomial error. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 297–314. Springer, Heidelberg, August 2014.
3. Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 284–332. Springer, Heidelberg, August 2019.
4. Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 308–326. Springer, Heidelberg, August 2015.
5. Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 152–181. Springer, Heidelberg, April / May 2017.
6. Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. How to garble arithmetic circuits. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 120–129. IEEE Computer Society Press, October 2011.
7. 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, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Heidelberg, April 2012.
8. Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Praves K. Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). Cryptology ePrint Archive, Report 2017/312, 2017. <http://eprint.iacr.org/2017/312>.
9. Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 221–238. Springer, Heidelberg, May 2014.
10. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001.
11. Boaz Barak, Iftach Haitner, Dennis Hofheinz, and Yuval Ishai. Bounded key-dependent message security. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 423–444. Springer, Heidelberg, May / June 2010.

12. Boaz Barak, Samuel B. Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 226–250. Springer, Heidelberg, May 2019.
13. Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.
14. Nir Bitansky, Ryo Nishimaki, Alain Passelègue, and Daniel Wichs. From cryptomania to obfustopia through secret-key functional encryption. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 391–418. Springer, Heidelberg, October / November 2016.
15. Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *56th FOCS*, pages 171–190. IEEE Computer Society Press, October 2015.
16. Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 480–499. Springer, Heidelberg, August 2014.
17. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. Cryptology ePrint Archive, Report 2019/720, 2019.
18. Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 1–25. Springer, Heidelberg, February 2014.
19. Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014*, pages 1–12. ACM, January 2014.
20. Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 278–307. Springer, Heidelberg, April / May 2017.
21. Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12. Springer, Heidelberg, April 2015.
22. Jean-Sébastien Coron, Tancreède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 476–493. Springer, Heidelberg, August 2013.
23. 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*, volume 1992 of *LNCS*, pages 119–136. Springer, Heidelberg, February 2001.
24. Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 294–310. Springer, Heidelberg, May 2016.
25. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

26. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Heidelberg, May 2013.
27. Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure MPC from indistinguishability obfuscation. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 74–94. Springer, Heidelberg, February 2014.
28. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
29. Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. Lower bounds on obfuscation from all-or-nothing encryption primitives. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 661–695. Springer, Heidelberg, August 2017.
30. Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. When does functional encryption imply obfuscation? In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 82–115. Springer, Heidelberg, November 2017.
31. 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, Part II*, volume 9986 of *LNCS*, pages 241–268. Springer, Heidelberg, October / November 2016.
32. Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.
33. Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 498–527. Springer, Heidelberg, March 2015.
34. Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. i-Hop homomorphic encryption and rerandomizable Yao circuits. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 155–172. Springer, Heidelberg, August 2010.
35. Craig Gentry, Charanjit S. Jutla, and Daniel Kane. Obfuscation using tensor products. Cryptology ePrint Archive, Report 2018/756, 2018.
36. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013.
37. Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 555–564. ACM Press, June 2013.
38. Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th ACM STOC*, pages 365–377. ACM Press, May 1982.
39. Satoshi Hada. Zero-knowledge and code obfuscation. In Tatsuaki Okamoto, editor, *ASIACRYPT 2000*, volume 1976 of *LNCS*, pages 443–457. Springer, Heidelberg, December 2000.
40. Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 537–565. Springer, Heidelberg, May 2016.

41. Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over \mathbb{R} to build iO . In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 251–281. Springer, Heidelberg, May 2019.
42. Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 28–57. Springer, Heidelberg, May 2016.
43. Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 599–629. Springer, Heidelberg, August 2017.
44. Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation with non-trivial efficiency. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 447–462. Springer, Heidelberg, March 2016.
45. Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from bilinear maps and block-wise local prgs. Cryptology ePrint Archive, Report 2017/250, Version 20170320:142653, 2017.
46. Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 630–660. Springer, Heidelberg, August 2017.
47. Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In Irit Dinur, editor, *57th FOCS*, pages 11–20. IEEE Computer Society Press, October 2016.
48. Alex Lombardi and Vinod Vaikuntanathan. Limits on the locality of pseudorandom generators and applications to indistinguishability obfuscation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 119–137. Springer, Heidelberg, November 2017.
49. Giulio Malavolta and Sri Aravinda Krishnan Thyagarajan. Homomorphic time-lock puzzles and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 620–649. Springer, Heidelberg, August 2019.
50. Daniele Micciancio. From linear functions to fully homomorphic encryption. Technical report, 2019. <https://bacrypto.github.io/presentations/2018.11.30-Micciancio-FHE.pdf>.
51. Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658. Springer, Heidelberg, August 2016.
52. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 223–238. Springer, Heidelberg, May 1999.
53. 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 ACM STOC*, pages 461–473. ACM Press, June 2017.
54. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.
55. Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.