

The Usefulness of Sparsifiable Inputs: How to Avoid Subexponential iO

Thomas Agrikola^{1,*}, Geoffroy Couteau^{2,*}, and Dennis Hofheinz^{3,*}

¹ Karlsruhe Institute of Technology, Karlsruhe, Germany

² IRIF, Paris-Diderot University, CNRS, France

³ ETH Zurich, Switzerland

Work done while all authors were at
Karlsruhe Institute of Technology.

Abstract. We consider the problem of removing subexponential reductions to indistinguishability obfuscation (iO) in the context of obfuscating probabilistic programs. Specifically, we show how to apply complexity absorption (Zhandry, Crypto 2016) to the recent notion of probabilistic indistinguishability obfuscation (piO, Canetti et al., TCC 2015). As a result, we obtain a variant of piO which allows to obfuscate a large class of probabilistic programs, from polynomially secure indistinguishability obfuscation and extremely lossy functions. Particularly, our piO variant is able to obfuscate circuits with specific input domains regardless of the performed computation. We then revisit several (direct or indirect) applications of piO, and obtain

- a fully homomorphic encryption scheme (without circular security assumptions),
- a multi-key fully homomorphic encryption scheme with threshold decryption,
- an encryption scheme secure under arbitrary key-dependent messages,
- a spooky encryption scheme for all circuits,
- a function secret sharing scheme with additive reconstruction for all circuits,

all from polynomially secure iO, extremely lossy functions, and, depending on the scheme, also other (but polynomial and comparatively mild) assumptions. All of these assumptions are implied by polynomially secure iO and the (non-polynomial, but very well-investigated) exponential DDH assumption. Previously, all the above applications required to assume the *subexponential* security of iO (and more standard assumptions).

Keywords: indistinguishability obfuscation, extremely lossy functions, subexponential assumptions.

1 Introduction

Obfuscation. Code obfuscation has been formalized already in the early 2000s as a cryptographic building block, by Hada [42] and Barak et al. [5], along with a

* Supported by ERC Project PREP-CRYPTO 724307.

number of early positive [23, 45, 47, 56, 61] and negative [5, 38, 61] results. However, prior to the candidate obfuscation scheme of Garg et al. [31], only relatively few positive results on obfuscation were known.

The first candidate obfuscator from [31] changed things. Their work identified indistinguishability obfuscation (iO, cf. [5, 39]) as an achievable *and* useful general notion of obfuscation: it presented a candidate indistinguishability obfuscator, along with a first highly non-trivial application. Since then, a vast number of applications have been proposed, ranging from functional [31], deniable [59], and fully homomorphic [25] encryption, over multi-party computation (e.g., [30]), to separation results (e.g., [46]). In the process, powerful techniques like “puncturing” [59] have been discovered, which have found applications even beyond obfuscation (e.g., in multi-party computation [8, 36], instantiating the Fiat-Shamir paradigm [24], and verifiable random functions [9, 40]). Besides, the notion of iO itself has been refined, and related to other notions of obfuscation [2, 10, 11, 20, 25, 50], and various different constructions of obfuscators have been presented [3, 4, 13, 53, 54, 57, 63].

Subexponential assumptions. It is currently hard to find a cryptographic primitive that can *not* be constructed from iO (in combination with another mild assumption such as the existence of one-way functions). However, some of the known iO-based constructions come only with *subexponential* reductions to iO. For instance, the only known iO-based constructions of fully homomorphic encryption [25], spooky encryption [27], and graded encoding schemes [29] suffer from reductions with a subexponential loss.

Hence, while iO has generally been recognized as an extremely powerful primitive (even to the extent being called a “central hub” for cryptography [59]), it is not at all clear if this also holds for *polynomially* secure iO. Indeed, it is conceivable that only polynomially secure iO exists, in which case much of iO’s power stands in question.

More generally, subexponential reductions (in particular to iO) are undesirable. Namely, the security of existing iO constructions is still not well-understood, and in particular current state-of-the-art constructions of iO schemes (such as [4, 53, 54]) already require subexponential computational assumptions themselves. Hence, assuming subexponential iO is a particularly risky bet. This suspicion is confirmed in part by [58], who separate polynomial and subexponential security for virtual black-box obfuscation.

Removing subexponential assumptions in general and from iO-based constructions in particular has already explicitly been considered in [35, 52] and [33, 34, 55] respectively. These works offer general techniques and ideas to turn subexponential reductions into polynomial ones. For instance, [34, 55] offer ways to replace (subexponential) iO-based constructions with (polynomial) constructions based on functional encryption. Of course, this requires a special structure of the primitive to be implemented, and is demonstrated for several primitives, including non-interactive key exchange and short signature schemes.

Our contribution. In this work, we are also concerned with substituting subexponential with polynomial reductions in iO-based constructions. Unlike [34, 55], however, we do not follow the approach of using functional encryption directly in place of iO, but instead will employ extremely lossy functions (ELFs) [62] to “absorb” subexponential complexity.⁴

We will implement a variant of probabilistic indistinguishability obfuscation (piO, introduced in [25]) using polynomially secure iO (and ELFs). piO schemes can be used to obfuscate *probabilistic* (i.e., randomized) programs, and are currently the only way to obtain, e.g., fully homomorphic encryption (FHE) schemes without circular security assumptions [25]. However, the only previous construction of piO schemes required subexponentially secure iO [25]. Hence, our construction yields the first FHE scheme from polynomially secure iO (and ELFs). Similarly, we can turn the assumption of subexponentially secure iO into polynomially secure iO (plus ELFs) in the construction of spooky encryption from [27].

Both FHE and spooky encryption are quite powerful primitives, and we obtain several “spin-off results” by revisiting their implications. For instance, when instantiating the piO-based FHE construction of [25] with our piO scheme and a suitable public-key encryption scheme, we obtain a fully key-dependent message (KDM) secure public-key encryption scheme from (polynomially secure) iO and the exponentially secure DDH assumption (and no further assumptions). Under the same assumptions, we obtain multi-key FHE with threshold decryption and function secret sharing schemes from the spooky encryption construction from [27].

On the plausibility of ELFs. One could argue that we trade one exponential assumption for another, and it is not clear that assuming polynomial iO and exponential DDH is any better than assuming only subexponential iO in the first place. Seconding Zhandry [62] here, we think that exponential DDH is a realistic assumption that is far more popular, better-investigated, and arguably more plausible than subexponential iO. Much of the currently deployed cryptography relies on (in fact a strong variant of) exponential DDH, because parameters are almost always chosen according to the best known attacks.

On the number of assumptions. Another natural observation is that iO for general circuits is already an exponential family of assumptions in itself (one for each obfuscated circuit). It might seem that this lets the challenge of relying on polynomially secure iO instead of subexponentially secure iO appear less appealing. We make two comments on that.

- First, being an exponential family of assumptions and assuming resistance against subexponential adversaries are orthogonal issues. Many cryptographic

⁴ That means that our final schemes depend on ELFs, which are currently only known to be instantiable from exponential assumptions. However, we stress that ELFs can be built from exponential variants of very standard assumptions, such as the decisional Diffie-Hellman (DDH) assumption.

assumptions have several dimensions of strengths, and relaxing the assumption in any of these dimensions is desirable.⁵ In this work, we make progress in one important dimension. By replacing subexponential iO by polynomial iO plus exponential DDH, we effectively trade an *exponential* number of subexponential hardness assumptions in exchange for a *single* (plausible, well-studied) exponential hardness assumption (plus an exponential family of polynomial hardness assumptions).

- Second, iO being an exponential family of assumptions can be considered an artificial consequence of working on the general notion of iO for *arbitrary circuits*. When using iO in concrete constructions (e.g. in all the constructions described in this paper), one almost never needs to assume iO for all circuits. It usually suffices to assume iO for a constant number of specific circuits (namely those being obfuscated in the construction and the analysis). Hence, iO is a small number of assumptions when used for building a cryptographic primitive.

1.1 Technical overview

The piO construction of Canetti et al. To describe our ideas, it will be helpful to briefly review the work of Canetti et al. [25]. In a nutshell, they define the notion of piO as a way to obfuscate probabilistic programs, and show how to use piO to implement the first FHE scheme without any circular security assumption. Intuitively, where the notion of iO captures that the obfuscation $iO(P)$ of a *deterministic* program P does not leak anything beyond the functionality of P , piO captures the same for *probabilistic* programs P .⁶

They also show how to implement piO with an indistinguishability obfuscator iO and a pseudorandom function (PRF) F . Namely, in order to obfuscate a probabilistic program P , Canetti et al. obfuscate the *deterministic* program P' that, on input x , runs $P(x)$ with random coins $r = F(K, x)$. Here, K is a PRF key hardcoded into P' . The security proof uses “puncturing” techniques [59] and a hybrid argument over all possible P -inputs x . More specifically, for each P -input x , separate reductions to the security of iO and F show that the execution of $P'(x)$ is secure.⁷

This proof strategy is very general and does not need to make any specific assumptions about the structure of P . (In fact, this strategy can be viewed as a specific form of “complexity leveraging”, technically similar to the conversion of selective security into adaptive security, e.g., [16].) However, the price to pay is a reduction loss which is linear in the size of the input domain (which usually is

⁵ For example, if a protocol relies on the subexponential hardness of LWE with exponential modulus-to-noise ratio, it would be desirable to achieve the same while relying only on polynomially secure LWE, even if the modulus-to-noise ratio remains exponential.

⁶ This is of course an oversimplification. Also, [25] define several types of piO security that provide a tradeoff between security and achievability.

⁷ Again, we are not very specific about the form of desired or assumed security. However, we believe that for this exposition, these specifics do not matter.

exponentially large). In particular, even after scaling security parameters suitably, Canetti et al. still require subexponentially secure iO and PRFs.

More on previous works to remove subexponentiality. There are a number of known ways to deal with subexponential reduction losses due to complexity leveraging (or related techniques). For instance, various semi-generic (pre-iO) techniques seek to achieve adaptive security (for different primitives) by establishing an algebraic or combinatorial structure on the used inputs [17, 44, 49, 60], and can sometimes be adapted to the iO setting [48]. But like the already-mentioned, somewhat more general approaches [34, 55], these works make specific assumptions about the structure of the involved computations.

A somewhat more general approach (that works for more general classes of programs) was outlined by Zhandry [62], who introduces the notion of “extremely lossy functions” (ELFs). Intuitively, an ELF is an injective function G that can be switched into an “extremely lossy mode”, in which its range is polynomially small. Such an ELF can sometimes be used to “preprocess” inputs in a cryptographic scheme, with the following benefit: a security reduction can switch the ELF to extremely lossy mode, so that only a polynomial number of (preprocessed) inputs $G(x)$ need to be considered. This simplifies a potential hybrid argument over all (preprocessed) inputs $G(x)$, and can lead to a polynomial (instead of a subexponential) reduction.

However, trying to apply this strategy to the construction and reduction of Canetti et al. (as sketched above) directly fails. Namely, in their application, inputs will be inputs x to an arbitrary (probabilistic) program P ; preprocessing them with an ELF will destroy their structure, and it is not clear how to run P on ELF-preprocessed inputs $G(x)$. Indeed, applying ELFs to realize piO requires fundamentally different techniques.

Main idea: piO with sparsifiable inputs. Instead, we will restrict ourselves to programs P that take as input an element x from a small number of (arbitrary but efficiently samplable) distributions. In other words, all possible inputs x need to be in the range of one of a small number of efficient samplers S_i . As an example, for $i \in \{0, 1\}$, sampler S_i could sample ciphertexts C that encrypt plaintext i . Moreover, we require that all inputs to a program P to be obfuscated are at some point actually sampled from some S_i according to a certain process.

Obfuscating a given probabilistic program P (that takes as inputs one or more x as above) now consists of two steps:

1. First, we *encode* all inputs x , in the sense that we compile S_i to attach a “certificate” aux to x . This certificate aux guarantees that x has really been sampled using S_i . Furthermore, the compiled sampler S_i uses preprocessed random coins of the form $G(r)$ (instead of r) for an ELF G . (When G is in injective mode, this does not affect the distribution of sampled x .) The certificate aux additionally guarantees this choice of random coins.⁸
2. Second, we produce the actual obfuscation of the probabilistic program P as follows. We use an indistinguishability obfuscator iO to obfuscate the

⁸ Looking ahead, this “certificate” will be implemented using a NIZK in our construction.

following (deterministic) variant P' of P : on inputs x_1, \dots, x_ℓ with certificates $\text{aux}_1, \dots, \text{aux}_\ell$, P' first checks the certificates aux_i and aborts if one of them is invalid. Next, P' runs $P(x_1, \dots, x_\ell)$, with random coins $F(K, (x_i)_{i=1}^\ell)$ for a PRF F and a hardcoded PRF key K . Finally, P' outputs P 's output.

Maybe the most important property of this setup is that now the sets of inputs x_i are “sparsifiable” in the following sense. If we set G to extremely lossy mode, then only a polynomial number of different random coins r can occur. Hence, each S_i will output one of only a small number of possible samples (e.g., encryptions C generated with random coins from a small set). In that sense, the set of possible inputs x_i to P has been “sparsified”, and a hybrid argument over all possible inputs as in [25] is possible with polynomial loss.

We stress that our technique of applying ELFs fundamentally differs from [62]. In [62], the constructed primitive itself ensures that G is applied on all inputs. When approaching the challenge of constructing piO, however, the input to the primitive must externally be sampled using random coins that are preprocessed with G . This process is not under the control of the primitive and therefore requires a mechanism certifying that inputs are generated according to this specific process. We implement this mechanism using the combination of compiling the sampler for the input distribution into a “certifying sampler” (step 1) and restricting correctness of the obfuscated program (step 2).

Surprisingly, our piO scheme achieves the notion of “dynamic-input piO” [25], a very strong variant of piO security. On a high level, dynamic-input piO guarantees indistinguishability between obfuscations of probabilistic programs as long as their output distributions on adversarially chosen inputs are indistinguishable. This constitutes a very strong requirement and, in fact, implies differing-inputs obfuscation [2, 5], a notion for which strong impossibility results exist [7, 32]. However, our obfuscator produces circuits which are only required to work on inputs certifiably generated according to a specific process. Hence, our piO scheme enjoys a restricted form of correctness. This enables us to circumvent the impossibility results [7, 32].

Applications. One obvious question is of course how restrictive our assumption on input domains really is. We show that our assumptions apply to two existing piO-based constructions, with a number of interesting consequences.

First, we revisit the piO-based construction of fully homomorphic encryption from [25]. Here, piO is used to obfuscate the FHE evaluation algorithm that takes two ciphertexts (say, of two bit plaintexts b_0 and b_1) as input, and outputs a ciphertext of the NAND of the two plaintexts (i.e., $b_0 \bar{\wedge} b_1$). If we set S_b to be a sampler that samples an encryption of b , this setting perfectly fits our scheme. Hence, we obtain first a leveled homomorphic encryption (LHE) scheme, and from this an FHE scheme using the high-level strategy from [25]. Hence, putting this together with our piO construction, we obtain an FHE scheme from polynomially secure iO and an ELF (and no further assumptions).

We note that the above FHE scheme is also fully key-dependent message (KDM, see [14]) secure when implemented with a suitable basic public-key encryption scheme (such as the DDH-based scheme of [18]). In that case, the

FHE is secure even when an encryption of its own secret key $C_{\text{sk}} = \text{Enc}(\text{pk}, \text{sk})$ is public. However, such an encryption C_{sk} can be transformed into an encryption $\text{Enc}(\text{pk}, f(\text{sk}))$ of an arbitrary function of sk thanks to the fully homomorphic properties of the FHE scheme. This leads to a conceptually very simple fully KDM-secure encryption scheme from polynomial assumptions (and ELFs). (We stress that we do not claim novelty for this observation. The connection between FHE and KDM security has already been observed in [6], and [27] have observed that the FHE construction of Canetti et al. preserves interesting properties of the underlying encryption scheme. However, [27] do not explicitly mention KDM security, and we find these consequences interesting enough to point out.)

As our second application, we consider spooky encryption (with CRS) introduced by Dodis et al. [27]. Intuitively, a spooky encryption scheme features a particular type of homomorphism in a multi-key, multi-ciphertext setting. More precisely, given ciphertexts $\{c_i = \text{Enc}(\text{pk}_i, x_i)\}_i$, a spooky encryption scheme allows to produce ciphertexts $\{c'_i\}_i$ with $y_i = \text{Dec}(\text{sk}_i, c'_i)$ such that certain so-called “spooky” relations between the x_i ’s and the y_i ’s hold. An important subclass of spooky relations allows to ensure that the y_i ’s are random subject to $\sum_i y_i = f(x_1, \dots, x_n)$, for any polynomial-time computable function f . Dodis et al. show that spooky encryption implies (among other things) function secret sharing, and they give a piO-based instantiation of spooky encryption (without the need of a CRS). At the heart of their construction is an obfuscated public “spooky evaluation” algorithm with a hardcoded decryption key. Since this algorithm also takes ciphertexts (and a public key) as input, its input domain can be sparsified much like in the FHE case.

In contrast to the FHE application, however, the spooky encryption application contains more technical subtleties. In particular, some inputs to the “spooky evaluation” algorithm may depend on other inputs, and hence sparsifying inputs needs to proceed in a certain order. The main difficulty here is to find a suitably flexible definition of sparsification; we omit the details in this overview. We note that our results of course also yield all applications of spooky encryption, only from polynomially secure iO (and ELFs). In particular, we obtain a simple protocol for function secret sharing for all functions (with additive reconstruction) from these assumptions [21].

We believe that our new notion of obfuscation will prove useful in other applications; for example, it would likely allow to improve the recent result of [26], which constructed CCA1-secure FHE from subexponentially secure iO.

Follow-up work. In the recent work [28], Döttling and Nishimaki define the notion universal proxy re-encryption (UPRE). UPRE schemes allow a proxy to convert any ciphertext under any public key of any existing PKE scheme into a ciphertext under any public key of any possibly different existing PKE scheme. [28] instantiate UPRE based on probabilistic IO due to [25]. UPRE for all PKE schemes (including non re-randomizable ones) requires dynamic-input piO, which implies differing-inputs obfuscation. However, [28] observe that our notion of doubly-probabilistic IO suffices which yields an instantiation of UPRE for all PKE schemes based on polynomial IO and exponential DDH.

Organization. In Section 2, we introduce our notations and recall standard preliminaries. Section 3 formally introduces our new variant of piO, called dpiO. Section 4 shows how to instantiate dpiO using polynomially secure iO and ELFs. Eventually, in Section 5 and the full version [1] we revisit the construction of leveled homomorphic encryption from [25], using dpiO instead of piO. In the full version [1], we revisit the construction of spooky encryption from [27] using dpiO and analyze our new construction.

Acknowledgments

We would like to thank the anonymous reviewers for many helpful comments.

2 Preliminaries

Notations. Throughout this paper, λ denotes the security parameter. For a natural number $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, \dots, n\}$. A probabilistic polynomial time algorithm (PPT, also denoted *efficient* algorithm) runs in time polynomial in the (implicit) security parameter λ . A positive function f is *negligible* if for any polynomial p there exists a bound $B > 0$ such that, for any integer $k \geq B$, $f(k) \leq 1/p(k)$. An event depending on λ occurs with *overwhelming probability* when its probability is at least $1 - \text{negl}(\lambda)$ for a negligible function negl . Given a finite set S , the notation $x \xleftarrow{\$} S$ means a uniformly random assignment of an element of S to the variable x . The notation $\mathcal{A}^{\mathcal{O}}$ indicates that the algorithm \mathcal{A} is given oracle access to \mathcal{O} . Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \geq 0}$ be a family of sets of (possibly randomized) circuits, where \mathcal{C}_λ contains circuits of size $\text{poly}(\lambda)$. A circuit sampler for \mathcal{C} is a distribution ensemble $D = \{D_\lambda\}_{\lambda \geq 0}$, such that D_λ ranges over triples (C_0, C_1, z) with $(C_0, C_1) \in \mathcal{C}_\lambda^2$ of identical size and taking inputs of the same length, and $z \in \{0, 1\}^{\text{poly}(\lambda)}$. A class of samplers \mathbf{S} is a set of circuit samplers for \mathcal{C} .

2.1 Indistinguishability Obfuscation for General Samplers

Indistinguishability obfuscation (iO) for general samplers was introduced in [25]. This notion generalizes the classical notion of iO introduced in [5]. Informally, an iO scheme for a sampler D allows to obfuscate circuits sampled with D so that, given a sample (C_0, C_1) from D , $\text{iO}(C_0) \approx \text{iO}(C_1)$. The standard notion of iO is recovered by considering samplers over functionally equivalent deterministic circuits of the same size. Stronger notions of obfuscation, denoted piO, can be defined for samplers over *probabilistic* circuits, satisfying various indistinguishability notions. We recall below the general definition of [25] of piO for a class of samplers (using a different notion of correctness defined in [27]). The original correctness definition states that an efficient adversary given oracle access to either the original circuit or the obfuscation (with the restriction that no input can be queried twice), can not tell the difference.

Definition 1 (piO for a Class of Samplers [25, 27]). *A uniform PPT machine piO is an indistinguishability obfuscator for a class of samplers \mathbf{S} over a family $\mathcal{C} = \{C_\lambda\}_{\lambda \geq 0}$ of possibly randomized circuits if it satisfies the following conditions:*

Correctness. *For every security parameter λ , every circuit $C \in C_\lambda$, and every input x , the distributions of $C(x)$ over the random coins of C and of $\text{piO}(1^\lambda, C)(x)$ over the random coins of the obfuscator are identical.*

μ -Indistinguishability. *For every sampler $D = \{D_\lambda\}_{\lambda \geq 0} \in \mathbf{S}$, and for every non-uniform PPT machine \mathcal{A} , it holds that*

$$\begin{aligned} & \left| \Pr[(C_0, C_1, z) \stackrel{\$}{\leftarrow} D_\lambda : \mathcal{A}(C_0, C_1, \text{piO}(1^\lambda, C_0), z) = 1] \right. \\ & \left. - \Pr[(C_0, C_1, z) \stackrel{\$}{\leftarrow} D_\lambda : \mathcal{A}(C_0, C_1, \text{piO}(1^\lambda, C_1), z) = 1] \right| \leq \mu(\lambda). \end{aligned}$$

We remark that the construction of piO from [25] satisfies this notion of correctness if instantiated with a perfect puncturable PRF, see Definition 4. Note that this does not extend to multiple evaluations of the obfuscated circuit. Further, note that this notion of correctness implies that the obfuscated circuit respects the support of the original circuit.

To recover the standard notion of iO, we introduce the class \mathbf{S}^{eq} of samplers for functionally equivalent (possibly randomized) circuits, *i.e.*, samplers over triplets (C_0, C_1, z) such that $|C_0| = |C_1|$, and for any input x and random coin r , $C_0(x; r) = C_1(x; r)$. The standard iO notion is obtained by considering piO over the subclass $\mathbf{S}^{\text{det}} \subset \mathbf{S}^{\text{eq}}$ of samplers for deterministic functionally equivalent circuits. We denote by $\text{Adv}_{\text{iO}}(\mathcal{A})$ the advantage of a PPT adversary \mathcal{A} in distinguishing between the obfuscation of functionally equivalent deterministic circuits.

The work of [25] introduced four types of samplers over probabilistic circuits, which define four corresponding variants of piO: dynamic-input piO, worst-case piO, memoryless worst-case piO, and X-Ind piO. Informally, a dynamic-input sampler is required to output (possibly randomized) circuits C_0, C_1 such that the output of these circuits on a dynamically chosen input is computationally indistinguishable. The corresponding notion, dynamic-input piO, is the strongest notion defined in [25] and a randomized equivalent of the notion of differing-input obfuscation. Therefore, it inherits the implausibility results of differing-input obfuscation for general circuits [7, 32]. On the other hand, [25] shows that the weaker notion X-Ind piO can be realized from subexponentially secure iO (and subexponentially secure one-way functions). Below, we recall the notion of dynamic-input samplers and dynamic-input piO from [25].

2.2 Dynamic-Input Samplers

Definition 2 (Dynamic-Input Indistinguishable Samplers [25]).

The class $\mathbf{S}^{\text{d-Ind}}$ of dynamic-input samplers for a circuit family \mathcal{C} contains all circuits samplers $D = \{D_\lambda\}_{\lambda \in \mathbb{N}}$ for \mathcal{C} with the following properties: for every non-uniform PPT $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, the advantage $\text{Adv}_{\text{d-Ind}}(\mathcal{A}) := \Pr[\text{Exp-d-Ind}_{\mathcal{A}}(\lambda) = 1] - \frac{1}{2}$ of \mathcal{A} in the experiment Exp-d-Ind represented in Figure 1 is negligible.

Experiment $\text{Exp-d-Ind}_{\mathcal{A}}(\lambda)$
$(C_0, C_1, z) \xleftarrow{\$} D_\lambda$
$(x, \text{st}) \xleftarrow{\$} \mathcal{A}_1(C_0, C_1, z)$ // the challenge input is chosen dynamically
$y \xleftarrow{\$} C_b(x)$ for $b \xleftarrow{\$} \{0, 1\}$
$b' \xleftarrow{\$} \mathcal{A}_2(\text{st}, C_0, C_1, z, x, y)$
return $b = b'$

Fig. 1. Experiment Exp-d-Ind for the indistinguishability property of dynamic-input samplers.

Definition 3 (dynamic-input piO). *A uniform PPT machine is a dynamic-input piO scheme if it is a piO for the class of dynamic-input samplers $\mathbf{S}^{\text{d-Ind}}$ over \mathcal{C} that includes all randomized circuits.*

Note that the class \mathbf{S}^{eq} of samplers for functionally equivalent circuits that we defined previously, is a subclass of $\mathbf{S}^{\text{d-Ind}}$: any sampler for triples (C_0, C_1, z) where C_0 and C_1 are functionally equivalent is trivially a dynamic-input sampler.

2.3 Puncturable Pseudorandom Function

A pseudorandom function (PRF) originally introduced in [37] is a tuple of PPT algorithms $F = (F.\text{KeyGen}, F.\text{Eval})$. Let \mathcal{K} denote the key space, \mathcal{X} denote the domain, and \mathcal{Y} denote the range. The key generation algorithm $F.\text{KeyGen}$ on input of 1^λ , outputs a random key from \mathcal{K} and the evaluation algorithm $F.\text{Eval}$ on input of a key K and $x \in \mathcal{X}$, evaluates the function $F: \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$. The core property of PRFs is that, on a random choice of key K , no probabilistic polynomial-time adversary should be able to distinguish $F(K, \cdot)$ from a truly random function, when given black-box access to it. Puncturable PRFs (pPRFs) have the additional property that some keys can be generated *punctured* at some point, so that they allow to evaluate the PRF at all points except for the punctured point. As observed in [19, 22, 51], it is possible to construct such punctured keys for the original construction from [37], which can be based on any one-way functions [43].

Definition 4 (Puncturable Pseudorandom Function [19, 22, 51]). *A puncturable pseudorandom function (pPRF) with punctured key space \mathcal{K}_p is a triple of PPT algorithms $(F.\text{KeyGen}, F.\text{Punct}, F.\text{Eval})$ such that*

- $F.\text{KeyGen}(1^\lambda)$ outputs a random key $K \in \mathcal{K}$,
- $F.\text{Punct}(K, x)$, on input $K \in \mathcal{K}$, $x \in \mathcal{X}$, outputs a punctured key $K\{x\} \in \mathcal{K}_p$,
- $F.\text{Eval}(K', x')$, on input a key K' (punctured or not), and a point x' , outputs an evaluation of the PRF.

We require F to meet the following conditions:

Functionality preserved under puncturing. *For all $\lambda \in \mathbb{N}$, for all $x \in \mathcal{X}$,*

$$\Pr[K \xleftarrow{\$} F.\text{KeyGen}(1^\lambda), K\{x\} \xleftarrow{\$} F.\text{Punct}(K, x): \\ \forall x' \in \mathcal{X} \setminus \{x\}: F.\text{Eval}(K, x') = F.\text{Eval}(K\{x\}, x')] = 1.$$

Pseudorandom at punctured points. For all PPT adversaries \mathcal{A} ,

$$\text{Adv}_{\text{s-cPRF}}(\mathcal{A}) := \Pr[\text{Exp-s-pPRF}_{\mathcal{A}}(\lambda) = 1] - \frac{1}{2}$$

is negligible, where Exp-s-cPRF is represented Figure 2.

We call a pPRF F perfect, if the distribution $\{F.\text{Eval}(K, x) \mid K \xleftarrow{\$} F.\text{KeyGen}(1^\lambda)\}$ is identical to the uniform distribution over \mathcal{Y} , for all inputs $x \in \mathcal{X}$.⁹

Definition 4 corresponds to a selective security notion for puncturable pseudorandom functions; adaptive security can also be considered, but will not be required in our work. For ease of notation we often write $F(\cdot, \cdot)$ instead of $F.\text{Eval}(\cdot, \cdot)$.

Experiment $\text{Exp-s-pPRF}_{\mathcal{A}}(\lambda)$

$(x^*, \text{state}) \xleftarrow{\$} \mathcal{A}(1^\lambda)$

$K \xleftarrow{\$} F.\text{KeyGen}(1^\lambda), K\{x^*\} \xleftarrow{\$} F.\text{Punct}(K, x^*)$

$b \xleftarrow{\$} \{0, 1\}, y_0 \leftarrow F.\text{Eval}(K, x^*), y_1 \xleftarrow{\$} \mathcal{Y}$

$b' \xleftarrow{\$} \mathcal{A}(\text{state}, K\{x^*\}, y_b)$

return $b = b'$

Fig. 2. Selective security game for puncturable pseudorandom functions.

2.4 Extremely Lossy Function

In this section we present extremely lossy functions (ELFs) introduced in [62]. ELFs are an extremely powerful primitive for complexity absorption allowing to replace subexponential or even exponential security assumptions with polynomial ones. Informally, an ELF is a function that can be generated in two different modes: an injective mode and an extremely lossy mode. In injective mode, the range of the ELF has exponential size whereas the range comprises only polynomially many elements in extremely lossy mode.

Definition 5 (Extremely Lossy Function [62]). An extremely lossy function ELF is an algorithm ELF.Gen which, on input (M, r) , where M is an integer and $r \in [M]$, outputs the description of a function $G: [M] \rightarrow [N]$ such that

- G can be computed in time $\text{poly}(\log M)$
- If $r = M$, G is injective with overwhelming probability (in $\log M$) over the randomness of $\text{ELF.Gen}(M, M)$;
- For any $r \in [M]$, $|G([M])| < r$ with overwhelming probability (in $\log M$) over the randomness of $\text{ELF.Gen}(M, r)$;
- **Indistinguishability:** For any large enough M , any polynomial P , and any inverse polynomial function δ , there exists a polynomial Q such that for any adversary \mathcal{A} running in time at most $P(\log M)$ and any $r \in [Q(\log M), M]$, the advantage of \mathcal{A} in distinguishing $\text{ELF.Gen}(M, M)$ from $\text{ELF.Gen}(M, r)$ is bounded by $\delta(\log M)$.

⁹ Given any pPRF F' , we can build a perfect pPRF F by sampling two keys $K_1 \xleftarrow{\$} F'.\text{KeyGen}(1^\lambda)$ and $K_2 \xleftarrow{\$} \mathcal{Y}$ in the key generation algorithm and defining the evaluation algorithm to output $F'.\text{Eval}(K_1, x) \oplus K_2$ on input of x , see [27].

In addition, we will consider extremely lossy functions satisfying *strong regularity*, as defined below.

Definition 6 (Strong regularity). *An ELF is strongly regular if for any (polynomial) r , the distribution $\{x \xleftarrow{\$} [M] : G(x)\}$ is statistically close to uniform over $G([M])$, with overwhelming probability over the choice of $G \xleftarrow{\$} \text{ELF.Gen}(M, r)$.*

We note that, if an ELF is strongly regular, it is possible to efficiently enumerate its image: the set of values obtained by evaluating an ELF on $\lambda r \log r$ random inputs, where r is a bound on the size of its image, contains the entire image of the ELF with overwhelming probability.

Instantiating ELFs. A construction of strongly regular extremely lossy function is given in [62]. It can be based on the exponential hardness of the decision Diffie-Hellman assumption (or any of its variants, such as the decision linear assumption), which we denote eDDH. The eDDH assumption for a group generator GroupGen (which generates a tuple (\mathbb{G}, p, g) where \mathbb{G} is a group, p is its order, and g is a generator of \mathbb{G}) states that there exists a polynomial q such that for any time bound t and probability ε , denoting $\kappa \leftarrow \log q(t, 1/\varepsilon)$, any adversary \mathcal{A} running in time at most t has advantage at most ε in distinguishing the following distributions:

$$\begin{aligned} & \{(\mathbb{G}, p, g) \xleftarrow{\$} \text{GroupGen}(1^\kappa), (a, b, c) \xleftarrow{\$} \mathbb{Z}_p^3 : (\mathbb{G}, g, g^a, g^b, g^c)\}, \\ & \{(\mathbb{G}, p, g) \xleftarrow{\$} \text{GroupGen}(1^\kappa), (a, b) \xleftarrow{\$} \mathbb{Z}_p^2 : (\mathbb{G}, g, g^a, g^b, g^{ab})\}. \end{aligned}$$

As noted in [62], groups based on elliptic curves are plausible candidates for groups where this assumption holds: in practical instantiations of DDH over elliptic curves, the size of the group is chosen assuming that the best attack takes time $O(\sqrt{p})$, hence disproving eDDH (which amounts to showing that there is an attack which takes time less than p^c for any constant c) would have considerable practical implications. Furthermore, relying on some form of exponential hardness assumption seems necessary, as a construction from polynomial hardness only would have surprising complexity-theoretic implications. More precisely, given access to only some super-logarithmic amount of non-determinism (i.e. $\omega(\log \log M)$ bits, where $[M]$ is the domain of the ELF), it is easy to distinguish between injective and lossy mode of the ELF. This is due to the fact that in lossy mode, the codomain of G has only polynomial size which means that the restriction of G to the set $D = [2^{\omega(\log \log M)}]$ (having super-polynomial cardinality) is guaranteed to have a collision (which is not the case in injective mode), and using only $\omega(\log \log M)$ bits of non-determinism this collision can be guessed.

2.5 Non-interactive Zero-Knowledge proof system

A non-interactive zero-knowledge (NIZK) proof system for a language L with witness relation R enables to prove in a non-interactive manner that some statements are in L without leaking information about corresponding witnesses. NIZK proof systems were originally introduced in [15].

Definition 7 (Non-interactive zero-knowledge proof system [41]). A non-interactive zero-knowledge (NIZK) proof system for a language $L \in \text{NP}$ (with witness relation R) is a tuple of PPT algorithms $\text{NIZK} = (\text{NIZK.Setup}, \text{NIZK.Prove}, \text{NIZK.Verify})$ such that NIZK.Setup is a common reference string generation algorithm, NIZK.Prove is a proving algorithm NIZK.Verify is a (deterministic) verification algorithm.

- $\text{NIZK.Setup}(1^\lambda)$ outputs a common reference string crs .
- $\text{NIZK.Prove}(\text{crs}, x, w)$, on input crs , a statement x and a witness w , outputs a proof π .
- $\text{NIZK.Verify}(\text{crs}, x, \pi)$, on input crs , a statement x and a proof π , outputs either 1 or 0.

We require NIZK to meet the following properties:

Perfect completeness. For every $(x, w) \in R$, we have that

$$\Pr[\text{crs} \xleftarrow{\$} \text{NIZK.Setup}(1^\lambda), \pi \xleftarrow{\$} \text{NIZK.Prove}(\text{crs}, x, w) : \text{NIZK.Verify}(\text{crs}, x, \pi) = 1] = 1.$$

Statistical soundness. For every $x \notin L$ with $|x| = \lambda$ and every (possibly unbounded) adversary \mathcal{A} , we have that

$$\Pr[\text{crs} \xleftarrow{\$} \text{NIZK.Setup}(1^\lambda), \pi \xleftarrow{\$} \mathcal{A}(\text{crs}, x) : \text{NIZK.Verify}(\text{crs}, x, \pi) = 1] < 2^{-\lambda}.$$

Computational zero-knowledge. There exists a PPT algorithm $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that for every PPT adversary \mathcal{A} ,

$$\begin{aligned} \text{Adv}_{\text{ZK}}(\mathcal{A}) := & \left| \Pr \left[\text{crs} \xleftarrow{\$} \text{NIZK.Setup}(1^\lambda) : \mathcal{A}^{\text{NIZK.Prove}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 \right] \right. \\ & \left. - \Pr \left[(\text{crs}, \tau) \xleftarrow{\$} \text{Sim}_0(1^\lambda) : \mathcal{A}^{\text{Sim}'_1(\text{crs}, \tau, \cdot, \cdot)}(\text{crs}) = 1 \right] \right| \end{aligned}$$

is negligible in λ , where $\text{Sim}'_1(\text{crs}, \tau, x, w)$ returns $\text{Sim}'_1(\text{crs}, \tau, x)$ only if $(x, w) \in R$.

For simplicity in the analysis we use a NIZK proof system that satisfies the following property: with overwhelming probability over the coins of $\text{NIZK.Setup}(1^\lambda)$, there does not exist any pair (x, π) such that $x \notin L$ and $\text{NIZK.Verify}(\text{crs}, x, \pi) = 1$. We call a NIZK that satisfies this property *almost perfectly sound*. We note that there is a simple folklore method which allows to construct an almost perfectly sound NIZK proof system starting from any statistically sound NIZK proof system. Consider a $2^{-\lambda}$ -statistically sound NIZK proof system, for statements $x \in \{0, 1\}^n$, for some polynomial $n = n(\lambda)$. Using parallel repetitions, the soundness of the proof system can be amplified to $2^{-\lambda-n}$.¹⁰ Then, it necessarily holds that for all possible crs except a $2^{-\lambda}$ fraction of them, there does not exist any pair (x, π) where $x \notin L$ and π is an accepting proof. To realize this, let E_x^{crs} denote the event that there exists a proof π such that $\text{NIZK.Verify}(\text{crs}, x, \pi) = 1$. Then, by a union bound argument, $\Pr_{\text{crs}}[\exists x \in \{0, 1\}^n \setminus L : E_x^{\text{crs}}] \leq \sum_{x \in \{0, 1\}^n \setminus L} \Pr_{\text{crs}}[E_x^{\text{crs}}] \leq 2^n \cdot 2^{-\lambda-n}$. Hence, the NIZK proof system obtained via parallel repetitions is almost perfectly sound.

¹⁰ That is, for any statement $x \notin L$, the probability $\Pr_{\text{crs}}[\exists \pi : \text{NIZK.Verify}(\text{crs}, x, \pi) = 1] \leq 2^{-\lambda-n}$.

In [12] Bitansky et al. showed that statistically sound NIZK proof systems can be obtained from polynomially secure indistinguishability obfuscation in conjunction with polynomially secure one-way functions.

3 Indistinguishability Obfuscation of Probabilistic Circuits over Distributions of Inputs

We first define the notion of a *sampler with input*. A sampler with input is a family of PPT algorithms which, on input x , sample from some distribution \mathcal{D}_x . This notion is convenient to capture the fact that, in many scenarios, the inputs to an obfuscated (probabilistic) circuit are sampled from some distribution \mathcal{D}_x , where x is some private input of a player.

Definition 8 (Sampler with Input). *We say that $ST = \{ST_\lambda\}_{\lambda \in \mathbb{N}}$ is a family of samplers with input, with input domain $\mathcal{I} = \{\mathcal{I}_\lambda\}_{\lambda \in \mathbb{N}}$, if for any $\lambda \in \mathbb{N}$, ST_λ is a set of probabilistic algorithms running in polynomial time (in 1^λ) with input domain \mathcal{I}_λ such that for any $S \in ST_\lambda$, and $x \in \mathcal{I}_\lambda$, $S(x)$ samples from $\{0, 1\}^\lambda$.*

3.1 Doubly-Probabilistic Indistinguishability Obfuscation

Below, we define a variant of indistinguishability obfuscation, that takes into account the fact that in many applications, obfuscated (probabilistic) circuits might only have to be evaluated on inputs coming from specific distributions. This is formalized by defining an encoding procedure for a sampler with input, which additionally produces auxiliary material that an obfuscated circuit can use to verify that its inputs were produced correctly, and by restricting the correctness of the obfuscated circuit to only hold for such well-formed inputs. We also refer to this auxiliary material as “certificate”.

However, this approach faces two issues. First, the inputs to an obfuscated circuit might not be sampled “all at once” from a single distribution; rather, they can come from different and independent sources. We capture this behavior by defining *ℓ -source obfuscation*, to account for the fact that different inputs might have been sampled independently. Second, when inputs are sampled by different parties, there might still be interdependencies which must be accounted for. For example, a party might sample an input (e.g. a public key of an encryption scheme), pass it to a second party, who then samples a second input from a distribution that is parametrized by the first input (e.g. a ciphertext under that public key). We handle this possibility by ordering the ℓ inputs to the obfuscated circuit, and by considering a *stateful* sampler with input S : when S is used to generate the i 'th sample y_i , it receives in addition to its input a state $\text{stf}(y_1, \dots, y_{i-1})$, where stf is some fixed efficiently computable *state function* (which depends on the particular application), and the y_j are outputs sampled by the first $i - 1$ sources. The state function captures the fact that a particular application might define an arbitrary communication pattern, and specifies which samples a party should have access to when generating his sample.

Additionally, we admit the possibility that a sampler produces some additional correlated output, that will not serve as input to an obfuscated circuit. Hence,

there is no need to “certify” this input using the auxiliary information, and we call this output unauthenticated output. Continuing the use case from above, given a sampler producing some public key, the unauthenticated part of that sampler’s output could be a corresponding secret key.

Definition 9 (Doubly-Probabilistic Indistinguishability Obfuscation (dpiO)). *Let ℓ be an integer. Let $\{\text{stf}_\lambda : (\{0, 1\}^\lambda \cup \{\perp\})^{\ell-1} \rightarrow \mathcal{T}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of efficiently computable functions. Let $\mathcal{SI} = \{\mathcal{SI}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of samplers with inputs, with input domain $\{\mathcal{T}_\lambda \times \mathcal{I}\}_{\lambda \in \mathbb{N}}$. Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of (probabilistic) circuits, and let \mathbf{CS} be a class of circuit samplers over \mathcal{C} . An ℓ -source dpiO scheme for $(\text{stf}, \mathcal{SI}, \mathcal{C}, \mathbf{CS})$ is a triple of PPT algorithms $(\text{Setup}, \text{Encode}, \text{Obfuscate})$ such that*

- $\text{Setup}(1^\lambda)$, on input the security parameter (in unary), outputs public parameters pp ;
- $\text{Encode}(\text{pp}, S)$, on input the public parameters pp , and a sampler with input $S \in \mathcal{SI}_\lambda$, outputs an encoded sampler S' ;
- $\text{Obfuscate}(\text{pp}, S, C)$, on input public parameters pp , a sampler with input $S \in \mathcal{SI}_\lambda$, and a circuit $C \in \mathcal{C}_{\ell, \lambda}$, outputs a circuit C' of size $\text{poly}(\lambda, |C|)$. We call C' an obfuscation of C with respect to S .

We further assume that the outputs of S on any input (state, x) is of the form $(y; y')$ (looking ahead, we will call y the authenticated output, and y' the unauthenticated output). The scheme should satisfy the three properties given below.

Informally, the first security requirement ensures that, on any (adversarially chosen) input x , state state , and sampler with input S , the sampler S' obtained by encoding S outputs samples of the form $(y, \text{aux}; y')$ where $(y; y')$ is distributed as an output of $S(\text{state}, x)$, and aux does not leak any non-trivial information about the inputs. This is formalized by requiring the existence of a simulator that can simulate aux given only y .

Definition 10 (Simulatability of Encodings). *An ℓ -source dpiO scheme for $(\text{stf}, \mathcal{SI}, \mathcal{C}, \mathbf{CS})$ satisfies simulatability of encodings if for any large enough λ and any (stateful) PPT adversary \mathcal{A} , there exists a PPT simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that the advantage of \mathcal{A} in distinguishing the experiments $\text{Exp}^{0\text{-enc}}$ and $\text{Exp}^{1\text{-enc}}$ represented on Figure 3 is negligible. We denote by $\text{Adv}_{\text{enc}}(\mathcal{A})$ the advantage of \mathcal{A} in this experiment.*

We now introduce the restricted correctness requirement. Intuitively, it states the following: in an honest scenario, the inputs (y_1, \dots, y_ℓ) should be constructed using the sampler with input S . The restricted correctness property guarantees that if the inputs have indeed been constructed “according to S ”, then the obfuscated circuit will behave correctly, and its output distribution (taken over the coins of the obfuscator) will be (statistically) indistinguishable from the output distribution of the circuit C (taken over its internal random coins). Note that this statistical indistinguishability does not extend to multiple evaluations.

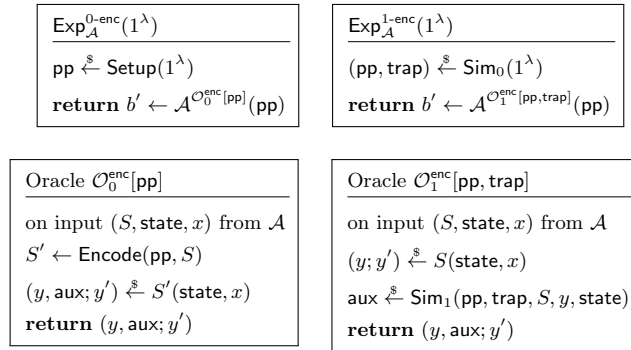


Fig. 3. Experiments $\text{Exp}_{\mathcal{A}}^{0\text{-enc}}(1^\lambda)$ and $\text{Exp}_{\mathcal{A}}^{1\text{-enc}}(1^\lambda)$ for the simulatability of encodings in an ℓ -source dpiO. The PPT algorithm \mathcal{A} can interact polynomially many times with either $\mathcal{O}_0^{\text{enc}}[\text{pp}]$ or $\mathcal{O}_1^{\text{enc}}[\text{pp}, \text{trap}]$. \mathcal{A} wins the experiment when it outputs $b' = b$ in $\text{Exp}_{\mathcal{A}}^{b\text{-enc}}(1^\lambda)$

Additionally, when evaluated on such inputs, the obfuscated circuit respects the support of the original circuit.

To make this definition meaningful, we need a way to let the obfuscated circuit verify that the inputs are well-formed. Note that we do not want to ensure that they were generated through S with uniformly random coins, but only that they were generated through S with *some* random coins (and some input). To make this verification possible, we let the parties generate their input using the encoded sampler S' instead. This encoded sampler should correctly sample as S , but it will in addition produce auxiliary information which can be used by the obfuscated program to verify that the inputs were honestly constructed (more formally, for a given y , that there exists an input x , coins r , and an unauthenticated part y' such that $(y; y') = S(x; r)$).

A small technicality is that we must allow the sampler with input to depend on state information, to capture the possible interdependencies between the inputs. This means that the auxiliary information will have to certify that an input was generated correctly, with respect to some state that the obfuscated circuit might not have access too (which would prevent it from verifying the certificate). However, this issue disappears by restricting the interdependencies to only involve a state computed from the previous *samples* (as opposed to more complex interdependencies which would involve, for example, the coins used to produce these samples). In this case, the obfuscated circuit can check the certificates in an incremental way: it first checks that y_1 was correctly constructed with respect to the state $\text{st}_\lambda(\perp, \dots, \perp)$, then it checks that y_2 was correctly constructed with respect to the state $\text{st}_\lambda(y_1, \perp, \dots, \perp)$, and so on.

Definition 11 (Statistical Restricted Correctness). *An ℓ -source dpiO scheme for $(\text{stf}, \mathcal{SI}, \mathcal{C}, \mathbf{CS})$ satisfies restricted correctness if for any large enough $\lambda \in \mathbb{N}$, any $S \in \mathcal{SI}_\lambda$, $(x_1, \dots, x_\ell) \in \mathcal{I}_\lambda^\ell$, and $C \in \mathcal{C}_{\ell\lambda}$, the advantage of any (possibly unbounded) adversary \mathcal{A} in distinguishing the experiments $\text{Exp}^{0\text{-rcorr}}$ and $\text{Exp}^{1\text{-rcorr}}$ represented on Figure 4 is negligible. We denote by $\text{Adv}_{\text{rcorr}}(\mathcal{A})$ the advantage of \mathcal{A} in this experiment. Additionally, we require that the encoded*

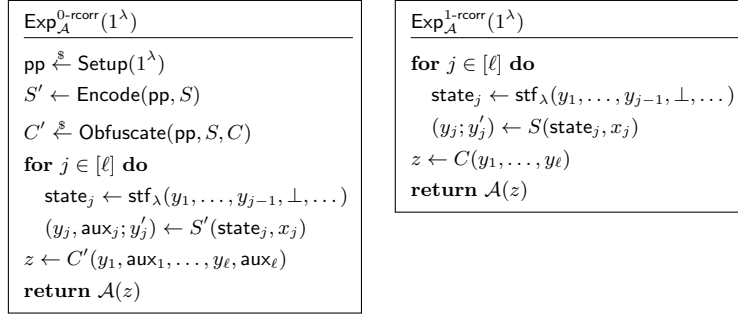


Fig. 4. Experiments $\text{Exp}_{\mathcal{A}}^{0\text{-rcorr}}(1^\lambda)$ and $\text{Exp}_{\mathcal{A}}^{1\text{-rcorr}}(1^\lambda)$ for the restricted correctness property an ℓ -source dpiO. \mathcal{A} wins the experiment when it outputs $b' = b$ in $\text{Exp}_{\mathcal{A}}^{b\text{-rcorr}}(1^\lambda)$ when $b \stackrel{\$}{\leftarrow} \{0, 1\}$.

sampler and the obfuscated circuit respect the support of the original sampler and the original circuit, respectively. That is for all $\text{pp} \leftarrow \text{Setup}(1^\lambda)$ and all $S' \leftarrow \text{Encode}(\text{pp}, S)$ and all $C' \leftarrow \text{Obfuscate}(\text{pp}, S, C)$, we have that for all inputs (state, x) , $S'(\text{state}, x) \in \text{Supp}(S(\text{state}, x))$ and for all $(y_1, \text{aux}_1, \dots, y_\ell, \text{aux}_\ell)$ produced as in $\text{Exp}^{0\text{-rcorr}}$, $C'(y_1, \text{aux}_1, \dots, y_\ell, \text{aux}_\ell) \in \text{Supp}(C(y_1, \dots, y_\ell))$.

We now introduce the indistinguishability notion. It is close in spirit to the standard indistinguishability notion for obfuscation of probabilistic circuits of [25]. However, in our scenario, the security notion must account for the fact that a set of public parameters pp is generated in a setup phase; the indistinguishability property of obfuscated circuits must therefore hold when (polynomially) many circuits are obfuscated with respect to a single string of public parameters. This suggests an oracle-based security notion.

Definition 12 (Indistinguishability with Respect to CS). *An ℓ -source dpiO scheme for $(\text{stf}, \mathcal{S}, \mathcal{C}, \mathbf{CS})$ satisfies indistinguishability with respect to \mathbf{CS} if for every circuit sampler $D = \{D_\lambda\}_{\lambda \in \mathbb{N}} \in \mathbf{CS}$, for any large enough λ , the advantage of any PPT adversary \mathcal{A} in distinguishing the experiments $\text{Exp}^{0\text{-ind}}$ and $\text{Exp}^{1\text{-ind}}$ represented on Figure 5 is negligible. We denote by $\text{Adv}_{\text{ind}}(\mathcal{A})$ the advantage of \mathcal{A} in this experiment.*




Fig. 5. Experiment $\text{Exp}_{\mathcal{A}}^{b\text{-ind}}(1^\lambda)$ for the indistinguishability with respect to \mathbf{CS} in an ℓ -source dpiO. The PPT algorithm \mathcal{A} can interact polynomially many times with $\mathcal{O}_b^{\text{ind}}[\text{pp}, D_\lambda]$. The oracle $\mathcal{O}_b^{\text{ind}}[\text{pp}, D_\lambda]$ is stateful and has (pp, D_λ) hardcoded in its description. \mathcal{A} wins the experiment when it outputs $b' = b$ in $\text{Exp}_{\mathcal{A}}^{b\text{-ind}}(1^\lambda)$ when $b \stackrel{\$}{\leftarrow} \{0, 1\}$.

4 Construction

In this section, we will construct an ℓ -source dpiO scheme (for any constant ℓ), for samplers with input over an input domain \mathcal{I} of polynomial size¹¹, and dynamic-input indistinguishable circuit-samplers. Our construction relies on polynomially-secure indistinguishability obfuscation, a perfect puncturable pseudorandom function, an almost perfectly sound non-interactive zero-knowledge proof system, and an extremely lossy function.

4.1 Overview

We start by providing a high-level overview of our construction. The Setup procedure generates parameters for the ELF and for the NIZK proof system. To encode a sampler with input S , we define the encoded sampler S' as follows: on input $(\text{state}, x; r)$, S' computes $(y; y') \stackrel{\$}{\leftarrow} S(\text{state}, x; G(r))$ and $\text{aux} \stackrel{\$}{\leftarrow} \text{NIZK.Prove}(y, L_{\text{state}}^{G,S}, (y', x, r))$, and outputs $(y, \text{aux}; y')$. Here, G is the ELF defined by the public parameters, and the language $L_{\text{state}}^{G,S}$ contains all values y for which there exists (y', x, r) such that $(y; y') = S(\text{state}, x, G(r))$. We call *valid input* a value $y \in L_{\text{state}}^{G,S}$. Note that when G is in injective mode, $L_{\text{state}}^{G,S}$ will in general be a trivial language. The simulatability of the encodings directly follows from the injectivity of G , and the zero-knowledge property of the proof system.

We construct the **Obfuscate** algorithm for a circuit C as follows (we assume a single source in this overview for simplicity). It first samples a pPRF key K for the pPRF F . Then, it returns an obfuscation of the following circuit: on input (y, aux) , run NIZK.Verify on aux to check that y is a valid input (and output \perp otherwise). Set $r \leftarrow F(K, y)$, and output $C(y; r)$. Restricted correctness follows from the correctness of the NIZK scheme. For indistinguishability between obfuscations of two dynamic-input indistinguishable circuits (C_0, C_1) , we follow the standard puncturing strategy of [25]: we proceed through a sequence of hybrids, with successive modifications of the obfuscated circuit. For every possible input y , we construct a sequence of hybrids where the outputs $C_0(y; r)$ are gradually replaced by $C_1(y; r)$. Each replacement relies on the security of the iO scheme, the PRF security, and the dynamic-input indistinguishability of C_0 and C_1 .

The main issue of this approach is that the number of possible inputs y (hence the number of hybrids) is exponential – indeed, this is the reason why the piO scheme of [25] requires subexponentially secure primitives (iO and PRF). To get around this issue, we first switch G to an appropriate extremely lossy mode, that the adversary cannot distinguish from the injective mode. Now, the soundness of the NIZK proof system ensures that all valid inputs y are of the form $S(\text{state}, x; G(r))$ for some (x, r) (omitting y' for simplicity). For a given **state**, the quantity of such values is bounded by the size of the range of G (which is polynomial), times the size of the input domain \mathcal{I} . Therefore, in all applications where the inputs to the obfuscated circuit are sampled using private inputs from a small domain, we can base security on polynomially secure iO.

¹¹ We note that the output domain of such samplers can be of exponential size.

$\text{Setup}(1^\lambda)$ $\text{crs} \xleftarrow{\$} \text{NIZK.Setup}(1^\lambda)$ $G \xleftarrow{\$} \text{ELF.Gen}(M, M)$ return pp $\leftarrow (\text{crs}, G)$	$\text{Encode}(\text{pp}, S)$ define S'^{pp} as follows :
$\text{Obfuscate}(\text{pp}, S, C)$ $K \xleftarrow{\$} \text{F.KeyGen}(1^\lambda)$ define \bar{C} as follows:	<div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> Circuit $S'^{\text{pp}}(\text{state}, x; r_1, r_2)$ $(y; y') \leftarrow S(\text{state}, x; G(r_1))$ $\pi \xleftarrow{\\$} \text{NIZK.Prove}(\text{crs},$ $\quad st = (G, S, \text{state}, y), w = (y', x, r_1); r_2)$ return $(y, \pi; y')$ </div> return S'^{pp}
<div style="border: 1px solid black; padding: 5px; margin-bottom: 5px;"> Circuit $\bar{C}[\text{stf}, (\text{crs}, G), S, C, K](x)$ $\text{parse } x =: ((y_1, \text{aux}_1), \dots, (y_\ell, \text{aux}_\ell))$ $\text{state}_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ if $\neg(\forall j \in [\ell]: \text{NIZK.Verify}(\text{crs}, (G, S, \text{state}_j, y_j), \text{aux}_j) = 1)$ then \quad return \perp $r := \text{F.Eval}(K, (y_1, \dots, y_\ell))$ $y := C((y_1, \dots, y_\ell); r)$ return y </div> $\Lambda \xleftarrow{\$} \text{iO}(\bar{C})$ return Λ	

Fig. 6. Construction of ℓ -source dpIO scheme $\text{dpiO} = (\text{Setup}, \text{Encode}, \text{Obfuscate})$.

4.2 Construction

For our construction, we employ a perfectly sound NIZK proof system for the following (parametrized) language

$$L_{\text{state}}^{G,S} := \{y \mid \exists(y', x, r) : (y; y') = S(\text{state}, x; G(r))\}.$$

Let $\ell \in \mathbb{N}$ be a constant, let $\{\text{stf}_\lambda : (\{0, 1\}^\lambda \cup \{\perp\})^{\ell-1} \rightarrow \mathcal{T}_\lambda\}_\lambda$ be a family of efficiently computable state functions, and let $\mathcal{C} = \{C_\lambda\}_\lambda$ be a family of (randomized) circuits with random space $\{0, 1\}^M$ (where $M = M(\lambda)$ is polynomial). Let \mathcal{SI} be a family of samplers with input domain \mathcal{I} of polynomial size. Further, let $\mathbf{S}^{\text{d-Ind}}$ be the class of dynamic-input indistinguishable samplers (over \mathcal{C}).

Theorem 13. *If ELF is a strongly regular extremely lossy function, iO is a perfectly correct polynomially secure IO scheme, F is a polynomially secure perfect puncturable PRF, and NIZK is a perfectly sound polynomially zero-knowledge NIZK proof system for the family of languages $\{L_{\text{state}}^{G,S}\}_{\text{state}, G, S}$, then $\text{dpiO} = (\text{Setup}, \text{Encode}, \text{Obfuscate})$ defined in Figure 6 is an ℓ -source dpIO scheme for $(\text{stf}, \mathcal{SI}, \mathcal{C}, \mathbf{S}^{\text{d-Ind}})$.*

As noted in Section 2.5, almost perfectly correct NIZKs can be constructed from polynomially-secure indistinguishability obfuscation and extremely lossy functions. ELFs also imply the existence of one-way functions, hence of perfect puncturable PRFs [37, 43]. Therefore, we get as corollary:

Corollary 14. *Assuming polynomially-secure indistinguishability obfuscation and extremely lossy functions, there exists (for any constant ℓ) an ℓ -source doubly-probabilistic indistinguishability obfuscation scheme for the class of dynamic-input circuit-samplers, and input-samplers with a polynomial size input domain.*

Proof (of Theorem 13). We prove that dpiO as defined in Figure 6 satisfies simulatability of encodings (cf. Definition 10), statistical restricted correctness (cf. Definition 11), and indistinguishability (cf. Definition 12).

Simulatability of encodings. We prove that there exists a PPT simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ such that for every PPT adversary \mathcal{A} , the advantage $\text{Adv}_{\text{enc}}(\mathcal{A})$ is negligible. By the zero-knowledge property of NIZK, there exists a simulator $(\text{NIZK.Sim}_0, \text{NIZK.Sim}_1)$. We construct a simulator $\text{Sim} = (\text{Sim}_0, \text{Sim}_1)$ as follows:

- Sim_0 produces the CRS using $(\text{crs}, \tau) \stackrel{\$}{\leftarrow} \text{NIZK.Sim}_0(1^\lambda)$, samples the parameters of the ELF G in injective mode, and outputs $\text{pp} := (\text{crs}, G)$ together with $\text{trap} := \tau$.
- Sim_1 on input (pp, trap) , a sampler S , a state state , and a value y sampled via $(y; y') \stackrel{\$}{\leftarrow} S(\text{state}, x)$, Sim_1 produces a simulated proof via $\pi \stackrel{\$}{\leftarrow} \text{NIZK.Sim}_1(\text{crs}, \tau, (G, S, \text{state}, y))$ and outputs $\text{aux} := \pi$.

Let \mathcal{A} be a PPT adversary on the simulatability property of dpiO . We prove indistinguishability between the real and the simulated distribution via a series of hybrids starting from the simulated game $\text{Exp}_{\mathcal{A}}^{1\text{-enc}}(1^\lambda)$.

Game \mathbf{G}_0 : This game is identical to $\text{Exp}_{\mathcal{A}}^{1\text{-enc}}(1^\lambda)$. We remark that in this game, the tuple $(y; y')$ is produced using the adversarially chosen sampler S on input of the adversarially chosen state state and input x supplied with true randomness.

Game \mathbf{G}_1 : This game is identical to \mathbf{G}_0 except for the fact that for each query (S, state, x) , the sampler S is supplied with randomness $G(r)$ for uniform r (instead of true randomness). Due to the strong regularity of G and by a standard hybrid argument over all queries, the statistical distance between \mathbf{G}_0 and \mathbf{G}_1 is negligible.

Game \mathbf{G}_2 : This game is the same as \mathbf{G}_1 with the difference that crs is produced honestly using $\text{NIZK.Setup}(1^\lambda)$. Additionally, for each adversarial query (S, state, x) , the proof π is produced honestly by $\text{NIZK.Prove}(\text{crs}, (G, S, \text{state}, y), (y', x, r))$, where $G(r)$ are the random coins supplied to the sampler S . The view of \mathcal{A} in game \mathbf{G}_2 is distributed exactly as in the real game $\text{Exp}_{\mathcal{A}}^{0\text{-enc}}(1^\lambda)$.

We construct a PPT adversary \mathcal{B} on the zero-knowledge property of NIZK. Given a CRS crs , \mathcal{B} samples an ELF G in injective mode and invokes \mathcal{A} on input of $\text{pp} := (\text{crs}, G)$. Each time \mathcal{A} queries its oracle on (S, state, x) , \mathcal{B} draws random coins r and invokes the sampler S on input of (state, x) with random coins $G(r)$ to obtain $(y; y')$. In order to produce π , \mathcal{B} calls its prove oracle on input (G, S, state, y) with witness (y', x, r) . Therefore, if \mathcal{B} is supplied with an honest CRS and honestly generated proofs, \mathcal{B} perfectly simulates \mathbf{G}_2 for \mathcal{A} , else \mathcal{B} perfectly simulates \mathbf{G}_1 . Hence, $|\Pr[\text{out}_2 = 1] - \Pr[\text{out}_3 = 1]| \leq \text{Adv}_{\text{zk}}(\mathcal{B})$. This concludes the proof.

Restricted Correctness. Let $S \in \mathbf{SI}_\lambda$ be an arbitrary sampler with input, let y_1, \dots, y_ℓ be arbitrary values from the input domain \mathcal{I}_λ , and let C be a circuit from the family $\mathcal{C}_{\ell\lambda}$. To prove the correctness of dpiO , we proceed over a series of hybrids.

Game \mathbf{G}_0 : This game is the ideal game $\text{Exp}_{\mathcal{A}}^{1\text{-rcorr}}(1^\lambda)$. As the sampler S is called using true randomness whereas in $\text{Exp}_{\mathcal{A}}^{0\text{-rcorr}}(1^\lambda)$ samples are generated using $G(r)$, where r is truly random, we need an intermediate hybrid.

Game \mathbf{G}_1 : This game is identical to \mathbf{G}_0 with the difference that each call of the sampler S is supplied with $G(r)$ as randomness (where r is sampled uniformly for each call). Due to the strong regularity of G , and by a hybrid argument over all calls of S , the statistical distance between \mathbf{G}_0 and \mathbf{G}_1 is negligible.

Game \mathbf{G}_2 : This game is the real game $\text{Exp}_{\mathcal{A}}^{0\text{-rcorr}}(1^\lambda)$.

We now argue that the view of \mathcal{A} in game \mathbf{G}_1 is distributed identically to its view in \mathbf{G}_2 . \mathbf{G}_2 samples public parameters pp via $\text{Setup}(1^\lambda)$ and S' an encoded sampler via $S' \leftarrow \text{Encode}(\text{pp}, S)$. Further, (y_j, aux_j) are sampled as $\text{state}_j \leftarrow \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ and $(y_j, \text{aux}_j, y'_j) \stackrel{\$}{\leftarrow} S'(\text{state}_j, x_j)$, for $j \in [\ell]$. Let A be the obfuscation $A \stackrel{\$}{\leftarrow} \text{Obfuscate}(\text{pp}, S, C)$ of the circuit C with respect to sampler S . Due to the perfect correctness of iO, A has the same functionality as $\bar{C}[\text{stf}, (\text{crs}, G), S, C, K]$, where K is a freshly generated key for the PRF F . Hence, by the perfect completeness of NIZK, on input of $((y_1, \text{aux}_1), \dots, (y_\ell, \text{aux}_\ell))$, A evaluates the circuit C on input of (y_1, \dots, y_ℓ) with random coins $F(K, (y_1, \dots, y_\ell))$. Therefore, the view of \mathcal{A} in the games \mathbf{G}_1 and \mathbf{G}_2 only differs in the fact that \mathbf{G}_1 supplies C with true random coins whereas \mathbf{G}_2 supplies C with $F(K, (y_1, \dots, y_\ell))$ as randomness. As F is a perfect PRF, the distribution $\{F(K, (y_1, \dots, y_\ell)) \mid K \stackrel{\$}{\leftarrow} F.\text{KeyGen}(1^\lambda)\}$ is identical to the uniform distribution over the image of F . Therefore, the view of \mathcal{A} in \mathbf{G}_1 and \mathbf{G}_2 is distributed identically.

By construction, all $S' \leftarrow \text{Encode}(\text{pp}, S)$ respect the support of S . Furthermore, by construction, perfect completeness of NIZK and perfect correctness of iO, for all $C' \leftarrow \text{Obfuscate}(\text{pp}, S, C)$ and all $(y_1, \text{aux}_1, \dots, y_\ell, \text{aux}_\ell)$ produced as in $\text{Exp}^{0\text{-rcorr}}$, $C'(y_1, \text{aux}_1, \dots, y_\ell, \text{aux}_\ell) \in \text{Supp}(C(y_1, \dots, y_\ell))$.

Security. Let $D \in \mathbf{S}^{\text{d-Ind}}$ be an arbitrary dynamic-input indistinguishable circuit sampler over \mathcal{C} . To prove that dpiO satisfies indistinguishability (Definition 12), we proceed over a series of hybrids. Toward contradiction, assume that there is a PPT adversary \mathcal{A} distinguishing $\text{Exp}_{\mathcal{A}}^{0\text{-ind}}(1^\lambda)$ from $\text{Exp}_{\mathcal{A}}^{1\text{-ind}}(1^\lambda)$ with non-negligible advantage ε over the random guess after making a polynomial number Q of queries to the oracle.

Game \mathbf{G}_0 . In this game, the challenger samples $b \stackrel{\$}{\leftarrow} \{0, 1\}$, and sets up the experiment $\text{Exp}_{\mathcal{A}}^{b\text{-ind}}(1^\lambda)$. More precisely, \mathcal{A} has access to the public parameters pp and an oracle $\mathcal{O}_b^{\text{ind}}[\text{pp}, D_\lambda]$, that on input of a sampler with input S , draws a sample (C_0, C_1, z) from D and outputs (C_0, C_1, z) together with an obfuscation $\text{Obfuscate}(\text{pp}, S, C_b)$. \mathcal{A} outputs a guess b' and the challenger returns 1 if $b' = b$. By assumption, $\Pr[\text{out}_0 = 1] = \varepsilon$.

Game \mathbf{G}_1 . In this game, the challenger samples G as $G \xleftarrow{s} \text{ELF.Gen}(M, t)$, where t is a polynomial such that any PPT algorithm of circuit size s has advantage at most $\varepsilon/2$ in distinguishing $\text{ELF.Gen}(M, M)$ from $\text{ELF.Gen}(M, t)$. The advantage of \mathcal{A} in this game is therefore lower bounded by $\varepsilon/2$: $\Pr[\text{out}_1 = 1] \geq \varepsilon/2$.

Game \mathbf{G}'_1 . This game proceeds exactly as \mathbf{G}_1 , except that after sampling $b \xleftarrow{s} \{0, 1\}$, the challenger always sets up the experiment $\text{Exp}_{\mathcal{A}}^{1\text{-ind}}(1^\lambda)$. The challenger still returns 1 iff $b' = b$.

By using a standard hybrid argument over the oracle queries, we prove that $|\Pr[\text{out}_1 = 1] - \Pr[\text{out}'_1 = 1]| \leq Q \cdot \text{negl}(\lambda)$, where Q is a polynomial in λ .

Game $\mathbf{G}_{1,q}$. This game is identical to \mathbf{G}_1 except for the fact that the first q oracle queries are answered using an obfuscation Λ_q of C_1 instead of C_b . Hence, $\Pr[\text{out}_{1,0} = 1] = \Pr[\text{out}_1 = 1]$ and $\Pr[\text{out}_{1,Q} = 1] = \Pr[\text{out}'_1 = 1]$, where Q is the number of adversarial oracle queries.

As $|\Pr[\text{out}_1 = 1] - \Pr[\text{out}'_1 = 1]| \leq \sum_{q=1}^Q |\Pr[\text{out}_{1,q} = 1] - \Pr[\text{out}_{1,q+1} = 1]|$, it suffices to upper bound the distinguishing gap between $\mathbf{G}_{1,q}$ and $\mathbf{G}_{1,q+1}$.

We observe that due to the (almost) perfect soundness of NIZK, the obfuscated circuit in the q -th oracle answer simulates the randomized computation of the circuit $C_{q,0}$ only on well-formed inputs, i.e. on outputs of S_q using random coins from the range of G . As ELF is in extremely lossy mode, this set of well-formed inputs is *extremely sparsified*. Therefore, by the strong regularity of ELF, we can enumerate over all possible outputs at all input positions $j \in [\ell]$. Let $B_{q,j}$ be the set of all well-formed inputs for input position j :

$$B_{q,j} := \{S_q(\text{stf}(y_1, \dots, y_{j-1}), x; G(r)) \mid x \in \mathcal{I}_\lambda, r \in \{0, 1\}^M, y_k \in B_k \text{ for } k \in [j-1]\}.$$

The set $B_{q,j}$ contains at most $|\mathcal{I}| \cdot t^{j-1}$ elements. Further, let $\gamma_{q,1} < \dots < \gamma_{q,\bar{t}}$ be the ordered enumeration of all ℓ -tuples in $B_q := \prod_{j=1}^{\ell} B_{q,j}$.¹² Hence, the total number of well-formed inputs $\bar{t} = \prod_{j=1}^{\ell} |B_{q,j}| \leq (|\mathcal{I}| \cdot t^{\ell-1})^\ell \leq |\mathcal{I}|^\ell \cdot t^{\ell^2}$ is polynomial in λ (given that ℓ is a constant, and $|\mathcal{I}|$ and t are polynomial).

Towards proving indistinguishability between $\mathbf{G}_{1,q}$ and $\mathbf{G}_{1,q+1}$, we conduct a hybrid argument over all well-formed inputs for the obfuscation Λ_q and gradually replace the evaluation of circuit $C_{q,b}$ with $C_{q,1}$. From here on, our proof strategy is similar to the one employed in [25]. However, we only need to consider polynomially many hybrids (as we assume $|\mathcal{I}|$ to be polynomial), hence we only lose a polynomial factor to the underlying assumptions.

Game $\mathbf{G}_{1,q,i}$. In game $\mathbf{G}_{1,q,i}$ the oracle answers the q -th query using an obfuscation of the circuit

$$\bar{C}'[\text{stf}, (\text{crs}, G), S_q, C_{q,b}, C_{q,1}, K_q, \gamma_{q,i}]$$

that is defined in Figure 7 using iO .

¹² We remark that the values of each set B_j can be computed efficiently by evaluating S on all possible inputs from $\mathcal{I} \times (\prod_{k=1}^{j-1} B_k)$ and all possible images in the range of G . Furthermore, it is possible to enumerate the image of G in polynomial time because G is strongly regular.

<pre> Circuit $\bar{C}'[\text{stf}, (\text{crs}, G), S, C_0, C_1, K, \gamma_i](x)$ ----- parse $x = ((y_1, \text{aux}_1), \dots, (y_\ell, \text{aux}_\ell))$ $\text{state}_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ if $\neg(\forall j \in [\ell]: \text{NIZK.Verify}(\text{crs}, (G, S, \text{state}_j, y_j), \text{aux}_j) = 1)$ then return \perp $\gamma := (y_1, \dots, y_\ell)$ if $\gamma < \gamma_i$ then $r := F(K, \gamma)$; return $C_1(\gamma; r)$ if $\gamma = \gamma_i$ then $r := F(K, \gamma)$; return $C_b(\gamma; r)$ if $\gamma > \gamma_i$ then $r := F(K, \gamma)$; return $C_b(\gamma; r)$ </pre>
--

Fig. 7. Definition of the circuit \bar{C}' .

The circuits $\bar{C}[\text{stf}, (\text{crs}, G), S_q, C_{q,b}, K_q]$ and $\bar{C}'[\text{stf}, (\text{crs}, G), S_q, C_{q,0}, C_{q,1}, K_q, \gamma_{q,1}]$ are functionally equivalent (on input $x = ((y_1, \text{aux}_1), \dots, (y_\ell, \text{aux}_\ell))$), both return $C_{q,b}(y_1, \dots, y_\ell)$ with randomness $F(K_q, (y_1, \dots, y_\ell))$. Hence, this game hop is justified by the indistinguishability property of iO, more formally there exists a PPT adversary \mathcal{B} such that $|\Pr[\text{out}_{1,q} = 1] - \Pr[\text{out}_{1,q,1} = 1]| \leq \text{Adv}_{\text{iO}}(\mathcal{B})$.

We aim to reduce the game hop from $\mathbf{G}_{1,q,i}^b$ to $\mathbf{G}_{1,q,i+1}^b$ to the dynamic-input indistinguishability of the circuit sampler D_λ . For this purpose, we first need to supply $C_{q,b}$ with true randomness. Hence, we define an other series of hybrids between $\mathbf{G}_{1,q,i}$ and $\mathbf{G}_{1,q,i+1}$.

Game $\mathbf{G}_{1,q,i,1}$. This game is identical to $\mathbf{G}_{1,q,i}$ except for the fact that we use a punctured PRF key $K_q\{\gamma_{q,i}\} \stackrel{\$}{\leftarrow} \text{F.Punct}(K_q, \gamma_{q,i})$ and obfuscate the circuit

$$\bar{C}''[\text{stf}, (\text{crs}, G), C_{q,0}, C_{q,1}, K_q\{\gamma_{q,i}\}, Y := C_{q,b}(\gamma_{q,i}; F(K_q, \gamma_{q,i})), \gamma_{q,i}]$$

defined in Figure 8 using iO.

As F preserves the functionality under punctured keys, the circuits $\bar{C}'[\text{stf}, (\text{crs}, G), S_q, C_{q,0}, C_{q,1}, K_q, \gamma_{q,i}]$ and $\bar{C}''[\text{stf}, (\text{crs}, G), S_q, C_{q,0}, C_{q,1}, K_q\{\gamma_{q,i}\}, Y := C_{q,b}(\gamma_{q,i}; F(K_q, \gamma_{q,i})), \gamma_{q,i}]$ are functionally equivalent. Hence, there exists a PPT adversary \mathcal{B} such that $|\Pr[\text{out}_{1,q,i} = 1] - \Pr[\text{out}_{1,q,i,1} = 1]| \leq \text{Adv}_{\text{iO}}(\mathcal{B})$.

We note that the view of \mathcal{A} in game $\mathbf{G}_{1,q,i,1}$ does not depend on the PRF key K . This enables to exploit the selective security of F.

<pre> Circuit $\bar{C}''[\text{stf}, (\text{crs}, G), S, C_0, C_1, K\{\gamma_i\}, Y, \gamma_i](x)$ ----- parse $x = ((y_1, \text{aux}_1), \dots, (y_\ell, \text{aux}_\ell))$ $\text{state}_j := \text{stf}(y_1, \dots, y_{j-1}, \perp, \dots, \perp)$ if $\neg(\forall j \in [\ell]: \text{NIZK.Verify}(\text{crs}, (G, S, \text{state}_j, y_j), \text{aux}_j) = 1)$ then return \perp $\gamma := (y_1, \dots, y_\ell)$ if $\gamma < \gamma_i$ then $r := F(K\{\gamma_i\}, \gamma)$; return $C_1(\gamma; r)$ if $\gamma = \gamma_i$ then return Y if $\gamma > \gamma_i$ then $r := F(K\{\gamma_i\}, \gamma)$; return $C_b(\gamma; r)$ </pre>

Fig. 8. Definition of the circuit \bar{C}'' .

Game $\mathbf{G}_{1,q,i,2}$. In this game we replace the randomness $F(K_q, (\gamma_{q,i}))$ by true randomness, i.e. we produce Y as follows: $Y := C_{q,b}(\gamma_{q,i}; R)$. This game hop is justified by the selective PRF property, more formally $|\Pr[\text{out}_{1,q,i,1} = 1] - \Pr[\text{out}_{1,q,i,2} = 1]| \leq \text{Adv}_{\text{s-PRF}}(\mathcal{B})$ for some PPT adversary \mathcal{B} .

Game $\mathbf{G}_{1,q,i,3}$. Game $\mathbf{G}_{1,q,i,3}$ is the same as $\mathbf{G}_{1,q,i,2}$ except for the fact that Y is produced using the circuit $C_{q,1}$, i.e. $Y := C_{q,1}(\gamma_{q,i}; R)$. This game hop is justified by the fact that the circuit sampler D_λ is a dynamic-input indistinguishable sampler.

Game $\mathbf{G}_{1,q,i,4}$. This game is the same as $\mathbf{G}_{1,q,i,3}$ with the difference that we again use pseudorandom coins to compute Y , i.e. $Y := C_{q,1}(\gamma_{q,i}; F(K_q, \gamma_{q,i}))$. For every PPT adversary \mathcal{A} there exists a PPT adversary \mathcal{B} such that $|\Pr[\text{out}_{1,q,i,3} = 1] - \Pr[\text{out}_{1,q,i,4} = 1]| \leq \text{Adv}_{s\text{-cPRF}}(\mathcal{B})$.

As the pPRF F preserves functionality under punctured keys, the two circuits $\bar{C}''[\text{stf}, (\text{crs}, G), S_q, C_{q,0}, C_{q,1}, K_q\{\gamma_{q,i}\}, Y := C_{q,1}(\gamma_{q,i}; F(K_q, \gamma_{q,i})), \gamma_{q,i}]$ and $\bar{C}'[\text{stf}, (\text{crs}, G), S_q, C_{q,0}, C_{q,1}, K_q, \gamma_{q,i+1}]$ are functionally equivalent. Therefore, we have that $|\Pr[\text{out}_{1,q,i,4} = 1] - \Pr[\text{out}_{1,q,i+1} = 1]| \leq \text{Adv}_{i\text{O}}(\mathcal{B})$.

Summing up, the advantage to distinguish \mathbf{G}_1 and $\mathbf{G}_{1,Q}$ is bounded by $|\mathcal{I}|^\ell \cdot t^{\ell^2} \cdot \text{negl}(\lambda)$. As ℓ is constant and $|\mathcal{I}|, t$ are polynomial, this quantity is negligible. As the circuit obfuscated in $\mathbf{G}_{1,Q}$ is now functionally equivalent to the circuit obfuscated in \mathbf{G}_1^1 , the game hop to \mathbf{G}'_1 is justified by the indistinguishability property of $i\text{O}$. More formally there exists a PPT adversary \mathcal{B} such that $|\Pr[\text{out}_{1,Q} = 1] - \Pr[\text{out}'_1 = 1]| \leq \text{Adv}_{\mathcal{B}}^{\text{iO}}(\lambda)$. This implies that the advantage of \mathcal{A} in game \mathbf{G}'_1 is lower bounded by $\varepsilon/2 - \text{negl}(\lambda)$, which is non-negligible. However, the view of \mathcal{A} in \mathbf{G}'_1 is perfectly independent of b , hence its advantage in this game cannot be non-zero; therefore, we reach a contradiction, which concludes the proof. \square

4.3 Extension

We sketch a straightforward extension of our above construction. It follows easily by inspection that the same proof strategy would work even if the ℓ sources, which sample inputs accorded to an encoding of a sampler S with respect to public parameters pp , are not required anymore to use the *same* public parameters. The ℓ sources could even each use different public parameters $(\text{pp}_1, \dots, \text{pp}_\ell)$. The modified proof for this scenario would proceed by first switching the ELFs in $(\text{pp}_1, \dots, \text{pp}_\ell)$ to an extremely-lossy mode, through a sequence of ℓ hybrids. Each extremely-lossy mode is chosen so that \mathcal{A} has at most $\varepsilon/2\ell$ advantage in distinguishing it from the injective mode. By a union bound, \mathcal{A} has therefore advantage at most $\varepsilon/2$ in distinguishing the all-injective modes from the all-lossy modes. Then, enumerating over all possible valid inputs to an obfuscated circuit takes polynomial time as before, as each input of a source comes from a set of polynomial size. Therefore, the exact same sequence of hybrids proves security, with a polynomial loss in the underlying primitives. To adapt the security properties of our definition of dpiO to this multi-parameter setting, it suffices to let all experiments initially sample and send to the adversary ℓ public parameters $(\text{pp}_1, \dots, \text{pp}_\ell)$ instead of one. In the simulatability of encodings definition (resp. in the indistinguishability definition), the adversary is allowed to specify under which public parameters it wants to receive a (real or simulated) sample

$(y, \text{aux}; y')$ (resp. under which public parameters it wants C_b to be obfuscated in the indistinguishability experiment).

It can prove convenient to simplify the construction in some applications to allow different sources to use different public parameters. Let us illustrate the syntax we adopt on an example: if $(\text{Setup}, \text{Encode}, \text{Obfuscate})$ is a 5-source dpiO scheme, we denote by $\text{Obfuscate}(\text{pp}_1[1-3], \text{pp}_2[4,5], S, C)$ an obfuscation of a circuit C , whose first three inputs should be sampled with respect to pp_1 , and whose last two inputs should be sampled with respect to pp_2 . We will also sometimes slightly abuse our notation, noting that an ℓ -source dpiO scheme directly implies an i -source dpiO scheme for $i \leq \ell$, and allow an ℓ -source scheme to obfuscate a circuit C that takes $i < \ell$ inputs.

5 Leveled Homomorphic Encryption

In this section we show that our notion of dpiO from Section 3 can be applied to construct leveled homomorphic encryption in a similar way as in [25]. This construction leads to a transformation which operates on an encryption scheme E , satisfying IND-CPA security (and possibly other security properties, e.g., KDM security), and produces a leveled homomorphic encryption scheme that retains the security properties of E . We recall the definition of IND-CPA secure encryption schemes in the full version [1].

Let stf_λ be the trivial state function, i.e. $\text{stf}: (y_1, y_2) \mapsto \perp$ for each $(y_1, y_2) \in (\{0, 1\}^\lambda \cup \{\perp\})^2$. Let $E = (E.\text{KeyGen}, E.\text{Enc}, E.\text{Dec})$ be an IND-CPA-secure public-key encryption scheme. Let the class \mathcal{SI} contain all samplers S^{pk} that on input of a state state and an input $x \in \mathcal{I} := \{0, 1\}$, produce an encryption $y := E.\text{Enc}(\text{pk}, x)$ and $y' := \perp$ ignoring state , where pk is a public key in the range of $E.\text{KeyGen}(1^\lambda)$. Let \mathcal{C} be the class of polynomially sized randomized circuits and let $\mathbf{S}^{\text{d-lnd}}$ be the class of dynamic-input indistinguishable samplers over \mathcal{C} .

Theorem 15. *Let $(\text{Setup}, \text{Encode}, \text{Obfuscate})$ be a 2-source dpiO scheme for $(\text{stf}, \mathcal{SI}, \mathcal{C}, \mathbf{S}^{\text{d-lnd}})$ and let E be an IND-CPA secure public-key encryption scheme. Then, LHE as defined in Figure 9 is an IND-CPA secure LHE scheme.*

The proof strategy is similar as in [25]. Here we provide an informal sketch of the proof and refer the reader to the full version [1] for the full proof. On a high level, we want to reduce the security of LHE to the security of the underlying encryption scheme E . However, the evaluation key ek contains information (even though obfuscated) on the secret keys of each level. For the purpose of invoking the security of E on the challenge ciphertext, we need to remove this dependency on sk_0 . Therefore, we gradually (starting from level L) replace the obfuscations of the circuits C with an obfuscation of *trapdoor* circuits tC that simply output samples produced by the encoded sampler S' on input of 0 (hence, not needing any information on decryption keys). These two circuits only differ in the fact that they sample from the same encoded sampler S' using (possibly) different inputs. Due to the simulatability of encodings and the IND-CPA security of E , the two circuits are dynamic-input indistinguishable. Hence, by the indistinguishability

```

LHE.KeyGen( $1^\lambda, 1^L$ )
-----
for  $i \in \{0, \dots, L\}$  do
   $(pk_i, sk_i) \xleftarrow{\$} E.KeyGen(1^\lambda)$ 
   $pp_i \xleftarrow{\$} Setup(1^\lambda)$ 
   $S'^{pk_i} \leftarrow Encode(pp_i, S^{pk_i})$ 
for  $i \in \{1, \dots, L\}$  do
   $A_i \xleftarrow{\$} Obfuscate(pp_{i-1}, S^{pk_{i-1}}, C[S'^{pk_i}, sk_{i-1}])$ 
 $pk := S'^{pk_0}, sk := sk_L, ek := (A_1, \dots, A_L)$ 
return  $(pk, ek, sk)$ 

LHE.Enc( $pk, m \in \{0, 1\}$ )    LHE.Dec( $sk, c$ )    LHE.Eval( $ek, C, (c_1, \dots, c_l)$ )
-----
parse  $pk =: S'^{pk_0}$         parse  $sk =: sk_L$     for  $i \in \{1, \dots, L\}$  do
 $(y, aux, y') \xleftarrow{\$} S'^{pk_0}(\perp, m)$   parse  $c =: (y, aux)$     foreach gate  $g$  on level  $i$  do
return  $c \leftarrow (y, aux)$     return  $E.Dec(sk_L, y)$     // let  $\alpha_g, \beta_g$  denote the respective inputs
                                      $\gamma_g := A_i(\alpha_g, \beta_g)$ 

```

Fig. 9. Description of the LHE scheme LHE. The circuit C is defined in Figure 10.

$C[S'^{pk}, sk'](x_\alpha, x_\beta)$ <hr/> $\alpha \leftarrow E.Dec(sk', x_\alpha)$ $\beta \leftarrow E.Dec(sk', x_\beta)$ $(y, aux, y') \xleftarrow{\$} S'^{pk}(\perp, \alpha \bar{\wedge} \beta)$ return (y, aux)	$tC[S'^{pk}](x_\alpha, x_\beta)$ <hr/> $(y, aux, y') \xleftarrow{\$} S'^{pk}(\perp, 0)$ return (y, aux)
--	--

Fig. 10. Definition of the circuits C and tC .

property of dpiO for $\mathbf{S}^{\text{d-Ind}}$, an honest evaluation key and an evaluation key consisting only of trapdoor circuits are indistinguishable.

Given these modifications, the challenge ciphertext c^* consists of an encryption of a bit b under pk_0 accompanied by some auxiliary information produced by the corresponding encoded sampler. This auxiliary information might leak information on the bit b and thereby prevents to directly employ the IND-CPA security of E . However, as dpiO satisfies simulatability of encodings, this auxiliary information can be simulated without knowledge of b and, hence, contains no information about b . Therefore, by the IND-CPA security of E , LHE is IND-CPA secure. Given our construction of dpiO from Section 4, we obtain the following corollary:

Corollary 16. *Assuming polynomially secure indistinguishability obfuscation and extremely lossy functions, there exists a leveled homomorphic encryption scheme.*

Note that IND-CPA secure cryptosystems, as required in our construction, can be constructed from (polynomially secure) IO and one-way function (the latter being implied by ELFs). Previously, constructions of LHE were only known from the learning with error assumption, or from *subexponentially secure* indistinguishability obfuscation (together with lossy encryption, which can be based e.g. on DDH). Using the generic transformation from leveled homomorphic encryption to fully homomorphic encryption from [25], we also get:

Corollary 17. *Assuming slightly-superpolynomially secure indistinguishability obfuscation and extremely lossy functions, there exists a fully homomorphic encryption scheme.*

Due to space limitations we state here two corollaries concerning FHE and KDM security and refer the reader to the full version [1] for a detailed discussion.

Corollary 18. *Assuming polynomially-secure indistinguishability obfuscation and extremely lossy functions, there exists a fully homomorphic encryption scheme.*

Corollary 19. *Assuming polynomially-secure indistinguishability obfuscation and eDDH, there exists a fully KDM-secure encryption scheme.*

References

1. Agrikola, T., Couteau, G., Hofheinz, D.: The usefulness of sparsifiable inputs: How to avoid subexponential iO. Cryptology ePrint Archive, Report 2018/470 (2018), <https://eprint.iacr.org/2018/470>
2. Ananth, P., Boneh, D., Garg, S., Sahai, A., Zhandry, M.: Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689 (2013), <http://eprint.iacr.org/2013/689>
3. Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 308–326. Springer, Heidelberg (Aug 2015)
4. Ananth, P., Sahai, A.: Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In: Coron, J., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part I. LNCS, vol. 10210, pp. 152–181. Springer, Heidelberg (Apr / May 2017)
5. Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer, Heidelberg (Aug 2001)
6. Barak, B., Haitner, I., Hofheinz, D., Ishai, Y.: Bounded key-dependent message security. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 423–444. Springer, Heidelberg (May / Jun 2010)
7. Bellare, M., Stepanovs, I., Waters, B.: New negative results on differing-inputs obfuscation. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 792–821. Springer, Heidelberg (May 2016)
8. Benhamouda, F., Lin, H.: k-round multiparty computation from k-round oblivious transfer via garbled interactive circuits. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 500–532. Springer, Heidelberg (Apr / May 2018)
9. Bitansky, N.: Verifiable random functions from non-interactive witness-indistinguishable proofs. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 567–594. Springer, Heidelberg (Nov 2017)
10. Bitansky, N., Canetti, R., Kalai, Y.T., Paneth, O.: On virtual grey box obfuscation for general circuits. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part II. LNCS, vol. 8617, pp. 108–125. Springer, Heidelberg (Aug 2014)
11. Bitansky, N., Paneth, O.: On the impossibility of approximate obfuscation and applications to resettable cryptography. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 241–250. ACM Press (Jun 2013)

12. Bitansky, N., Paneth, O.: ZAPs and non-interactive witness indistinguishability from indistinguishability obfuscation. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 401–427. Springer, Heidelberg (Mar 2015)
13. Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. In: Guruswami, V. (ed.) 56th FOCS. pp. 171–190. IEEE Computer Society Press (Oct 2015)
14. Black, J., Rogaway, P., Shrimpton, T.: Encryption-scheme security in the presence of key-dependent messages. In: Nyberg, K., Heys, H.M. (eds.) SAC 2002. LNCS, vol. 2595, pp. 62–75. Springer, Heidelberg (Aug 2003)
15. Blum, M., Feldman, P., Micali, S.: Non-interactive zero-knowledge and its applications (extended abstract). In: 20th ACM STOC. pp. 103–112. ACM Press (May 1988)
16. Boneh, D., Boyen, X.: Efficient selective-ID secure identity based encryption without random oracles. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 223–238. Springer, Heidelberg (May 2004)
17. Boneh, D., Boyen, X.: Secure identity based encryption without random oracles. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 443–459. Springer, Heidelberg (Aug 2004)
18. Boneh, D., Halevi, S., Hamburg, M., Ostrovsky, R.: Circular-secure encryption from decision Diffie-Hellman. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 108–125. Springer, Heidelberg (Aug 2008)
19. Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (Dec 2013)
20. Boyle, E., Chung, K.M., Pass, R.: On extractability obfuscation. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 52–73. Springer, Heidelberg (Feb 2014)
21. Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 337–367. Springer, Heidelberg (Apr 2015)
22. Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer, Heidelberg (Mar 2014)
23. Canetti, R.: Towards realizing random oracles: Hash functions that hide all partial information. In: Kaliski Jr., B.S. (ed.) CRYPTO’97. LNCS, vol. 1294, pp. 455–469. Springer, Heidelberg (Aug 1997)
24. Canetti, R., Chen, Y., Reyzin, L., Rothblum, R.D.: Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 91–122. Springer, Heidelberg (Apr / May 2018)
25. Canetti, R., Lin, H., Tessaro, S., Vaikuntanathan, V.: Obfuscation of probabilistic circuits and applications. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 468–497. Springer, Heidelberg (Mar 2015)
26. Canetti, R., Raghuraman, S., Richelson, S., Vaikuntanathan, V.: Chosen-ciphertext secure fully homomorphic encryption. In: Fehr, S. (ed.) PKC 2017, Part II. LNCS, vol. 10175, pp. 213–240. Springer, Heidelberg (Mar 2017)
27. Dodis, Y., Halevi, S., Rothblum, R.D., Wichs, D.: Spooky encryption and its applications. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part III. LNCS, vol. 9816, pp. 93–122. Springer, Heidelberg (Aug 2016)
28. Döttling, N., Nishimaki, R.: Universal proxy re-encryption. *Cryptology ePrint Archive, Report 2018/840* (2018), <https://eprint.iacr.org/2018/840>

29. Farshim, P., Hesse, J., Hofheinz, D., Larraia, E.: Graded encoding schemes from obfuscation. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 371–400. Springer, Heidelberg (Mar 2018)
30. Garg, S., Gentry, C., Halevi, S., Raykova, M.: Two-round secure MPC from indistinguishability obfuscation. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 74–94. Springer, Heidelberg (Feb 2014)
31. Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS. pp. 40–49. IEEE Computer Society Press (Oct 2013)
32. Garg, S., Gentry, C., Halevi, S., Wichs, D.: On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 518–535. Springer, Heidelberg (Aug 2014)
33. Garg, S., Pandey, O., Srinivasan, A.: Revisiting the cryptographic hardness of finding a Nash equilibrium. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 579–604. Springer, Heidelberg (Aug 2016)
34. Garg, S., Pandey, O., Srinivasan, A., Zhandry, M.: Breaking the sub-exponential barrier in obfuscation. In: Coron, J., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part III. LNCS, vol. 10212, pp. 156–181. Springer, Heidelberg (Apr / May 2017)
35. Garg, S., Srinivasan, A.: Single-key to multi-key functional encryption with polynomial loss. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 419–442. Springer, Heidelberg (Oct / Nov 2016)
36. Garg, S., Srinivasan, A.: Two-round multiparty secure computation from minimal assumptions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 468–499. Springer, Heidelberg (Apr / May 2018)
37. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: 25th FOCS. pp. 464–479. IEEE Computer Society Press (Oct 1984)
38. Goldwasser, S., Kalai, Y.T.: On the impossibility of obfuscation with auxiliary input. In: 46th FOCS. pp. 553–562. IEEE Computer Society Press (Oct 2005)
39. Goldwasser, S., Rothblum, G.N.: On best-possible obfuscation. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 194–213. Springer, Heidelberg (Feb 2007)
40. Goyal, R., Hohenberger, S., Koppula, V., Waters, B.: A generic approach to constructing and proving verifiable random functions. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part II. LNCS, vol. 10678, pp. 537–566. Springer, Heidelberg (Nov 2017)
41. Groth, J., Ostrovsky, R., Sahai, A.: Perfect non-interactive zero knowledge for NP. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 339–358. Springer, Heidelberg (May / Jun 2006)
42. Hada, S.: Zero-knowledge and code obfuscation. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 443–457. Springer, Heidelberg (Dec 2000)
43. Hästad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM Journal on Computing* 28(4), 1364–1396 (1999)
44. Hofheinz, D., Kiltz, E.: Programmable hash functions and their applications. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 21–38. Springer, Heidelberg (Aug 2008)
45. Hofheinz, D., Malone-Lee, J., Stam, M.: Obfuscation for cryptographic purposes. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 214–232. Springer, Heidelberg (Feb 2007)
46. Hofheinz, D., Rao, V., Wichs, D.: Standard security does not imply indistinguishability under selective opening. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 121–145. Springer, Heidelberg (Oct / Nov 2016)

47. Hohenberger, S., Rothblum, G.N., shelat, a., Vaikuntanathan, V.: Securely obfuscating re-encryption. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 233–252. Springer, Heidelberg (Feb 2007)
48. Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 201–220. Springer, Heidelberg (May 2014)
49. Hohenberger, S., Waters, B.: Short and stateless signatures from the RSA assumption. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 654–670. Springer, Heidelberg (Aug 2009)
50. Ishai, Y., Pandey, O., Sahai, A.: Public-coin differing-inputs obfuscation and its applications. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015, Part II. LNCS, vol. 9015, pp. 668–697. Springer, Heidelberg (Mar 2015)
51. Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 669–684. ACM Press (Nov 2013)
52. Li, B., Micciancio, D.: Compactness vs collusion resistance in functional encryption. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 443–468. Springer, Heidelberg (Oct / Nov 2016)
53. Lin, H.: Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 599–629. Springer, Heidelberg (Aug 2017)
54. Lin, H., Tessaro, S.: Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 630–660. Springer, Heidelberg (Aug 2017)
55. Liu, Q., Zhandry, M.: Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 138–169. Springer, Heidelberg (Nov 2017)
56. Lynn, B., Prabhakaran, M., Sahai, A.: Positive results and techniques for obfuscation. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 20–39. Springer, Heidelberg (May 2004)
57. Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation from semantically-secure multilinear encodings. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 500–517. Springer, Heidelberg (Aug 2014)
58. Pass, R., shelat, a.: Impossibility of VBB obfuscation with ideal constant-degree graded encodings. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016-A, Part I. LNCS, vol. 9562, pp. 3–17. Springer, Heidelberg (Jan 2016)
59. Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) 46th ACM STOC. pp. 475–484. ACM Press (May / Jun 2014)
60. Waters, B.R.: Efficient identity-based encryption without random oracles. In: Cramer, R. (ed.) EUROCRYPT 2005. LNCS, vol. 3494, pp. 114–127. Springer, Heidelberg (May 2005)
61. Wee, H.: On obfuscating point functions. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 523–532. ACM Press (May 2005)
62. Zhandry, M.: The magic of ELFs. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 479–508. Springer, Heidelberg (Aug 2016)
63. Zimmerman, J.: How to obfuscate programs directly. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 439–467. Springer, Heidelberg (Apr 2015)