

Indistinguishability Obfuscation: from Approximate to Exact

Nir Bitansky*

Vinod Vaikuntanathan[†]

Abstract

We show general transformations from subexponentially-secure approximate indistinguishability obfuscation (IO) where the obfuscated circuit agrees with the original circuit on a $1/2 + \epsilon$ fraction of inputs, into exact indistinguishability obfuscation where the obfuscated circuit and the original circuit agree on all inputs (except for a negligible probability over the coin tosses of the obfuscator). As a step towards our results, which is of independent interest, we also obtain an approximate-to-exact transformation for functional encryption. At the core of our techniques is a method for “fooling” the obfuscator into giving us the correct answer, while preserving the indistinguishability-based security. This is achieved based on various types of secure computation protocols that can be obtained from different standard assumptions.

Put together with the recent results of Canetti, Kalai and Paneth (TCC 2015), Pass and Shelat (Eprint 2015), and Mahmoody, Mohammed and Nemathaji (Eprint 2015), we show how to convert indistinguishability obfuscation schemes in various ideal models into exact obfuscation schemes in the plain model.

Keywords: Functional Encryption, Program Obfuscation, Secure Function Evaluation.

*MIT. E-mail: nirbitan@csail.mit.edu. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

[†]MIT. E-mail: vinodv@csail.mit.edu. Research supported in part by NSF Grants CNS-1350619 and CNS-1414119, Alfred P. Sloan Research Fellowship, Microsoft Faculty Fellowship, the NEC Corporation, and a Steven and Renee Finn Career Development Chair from MIT. This work was done in part while the author was visiting the Simons Institute for the Theory of Computing, supported by the Simons Foundation and by the DIMACS/Simons Collaboration in Cryptography through NSF grant CNS-1523467.

1 Introduction

Program obfuscation, the science of making programs “unintelligible” while preserving functionality, has been a holy grail in cryptography for over a decade. While the most natural and intuitively appealing notion of obfuscation, namely *virtual-black-box* (VBB) obfuscation [BGI⁺12], was shown to have strong limitations [BGI⁺12, GK05, BCC⁺14], the recent work of Garg, Gentry, Halevi, Raykova, Sahai and Waters [GGH⁺13b, SW14] opened new doors by demonstrating that the weaker notion of *indistinguishability obfuscation* (IO) is both very useful and potentially achievable. Since then, a veritable flood of applications has made indistinguishability obfuscation virtually “crypto-complete”.

On the flip side, the tremendous power of IO also begets its reliance on strong and untested computational assumptions. Indeed, it has been a major cryptographic quest to come up with a construction of IO with a security proof based on well-studied computational assumptions. Garg et al. [GGH⁺13b] gave us the first *candidate* construction of IO, however the construction came as-is, without a security proof.

We have recently seen several beautiful works [PST14, GLSW14, AJ15, BV15] that shed light on how a security proof for IO will look like. Pass, Seth and Telang show security of an IO construction based on a “semantic security” assumption on multi-linear maps [GGH12]; Gentry, Lewko, Sahai and Waters [GLSW14] (following [GLW14]) show security based on the “multilinear subgroup elimination assumption” on multi-linear maps; Ananth and Jain [AJ15] and Bitansky and Vaikuntanathan [BV15] show how to construct IO from any functional encryption scheme. Unfortunately, the first two of these works are based on the mathematical abstraction of multi-linear maps which has had a troubled history so far (with several constructions [GGH13a, CLT13, BWZ14, GGHZ14, GGH15, CLT15] and matching attacks [GGH13a, LS14, CHL⁺15, CGH⁺15, HJ15]), and the last two rely on functional encryption for which the only known constructions, yet again, use multi-linear maps.

Yet another line of work on proving security of obfuscation works in so-called *idealized models*. In a typical idealized model, both the construction and the adversary have access to an oracle that implements a certain functionality; in the random oracle model [BR93], this is a random function; in the generic group model [Sho97], this is the functionality of a group; and the most recent entrant to this club, namely the ideal multilinear map model, is an abstraction of the functionality of multilinear maps. Several works [BR14b, BR14a, BGK⁺14, AB15, Zim15] along this route prove security of (different) constructions of obfuscation (even in the sense of virtual black-box security) in various ideal multi-linear map models.

An even more recent line of work, initiated by Canetti, Kalai, and Paneth [CKP15], investigates how to transform constructions of obfuscation in idealized models into ones in the plain model, where there are no oracles. Indeed, this may lead to an aesthetically appealing avenue to constructing obfuscation schemes:

1. *Construct an obfuscation scheme in an appropriate idealized model; and*
2. *“De-idealize” it: translate the ideal model obfuscation scheme into a scheme in the real world.*

Even if eventual constructions of obfuscation schemes do not initially proceed along these lines, we believe that this two-step process is a conceptually appealing route to eventual, mature, constructions of obfuscation schemes. Indeed, constructions in ideal models, while not immediately deployable, typically give us an abstract, high level, understanding.

Concretely, the work of [CKP15] essentially shows that any obfuscator in the random oracle model can be converted to an obfuscator in the plain model with the same security properties. Pass and Shelat [PS15] and subsequently, Mahmoody, Mohammed and Nematihaji [MMN15] extend this to the generic group and ring models respectively, as well as ideal multilinear maps model with *bounded multi-linearity*.

However, the resulting obfuscators suffer from a major drawback: *they only have approximate correctness*. That is, the plain model obfuscator may error on a polynomially large fraction of inputs (or more generally with some polynomial probability when inputs are taken from a given samplable distribution). Roughly speaking, these results proceed by isolating a list of “heavy oracle queries”, that is, queries that arise in the evaluation of the obfuscated circuit on a large fraction of inputs. Once the (polynomially large set of) heavy queries are identified, the result of the oracle queries on this set is published as part of the obfuscated circuit. This approach will inherently miss the queries made by a rare set of inputs, resulting in an incorrect evaluation.

While these transformations already have interesting negative consequences (regarding the impossibility of VBB in these idealised models), the lack of correctness presents a serious obstacle towards fulfilling the above two-step plan towards constructing IO in the plain model. Indeed, it is far from clear that applications of IO will work when we only have *approximate IO* at our disposal. Certainly, one could go through the applications of IO one-by-one, and attempt to re-derive them from approximate IO, but in the absence of automated theorem provers¹, this seems neither particularly efficient nor aesthetically pleasing.

This motivates us to ask:

Is there an approximate-to-exact transformation for indistinguishability obfuscation?

In other words, we are asking for “one transformation to rule them all”, a generic way to compile any approximate obfuscation scheme into a perfectly correct obfuscation scheme, automatically enabling to recover all the applications of IO even given only approximately correct obfuscation.

In this work, we provide exactly such a transformation, under minimal additional assumptions. Let us now describe our results in detail.

1.1 Our Results

We say that an obfuscator \widetilde{IO} is (\mathcal{X}, α) -correct for a given input sampler \mathcal{X} and $\alpha \in [0, 1]$ (which may depend on the security parameter), if it is correct with probability at least α over inputs sampled by \mathcal{X} . Security is defined as in the standard setting of (exact) indistinguishability obfuscation. We shall refer to such an obfuscator as an *approximate indistinguishability obfuscator*.

Our main result is that approximate IO with subexponential-security for a certain class of samplers can be converted under standard assumptions into *(almost) exact IO* where for any circuit, with overwhelming probability over the coins of the obfuscator algorithm the resulting obfuscation is correct on *all* inputs. We present two routes towards this results based on different assumptions and with different parameters.

Theorem 1.1 (informal). *Assuming DDH, there exists an input sampler \mathcal{X}_1 and a transformation that for any $\alpha \geq \frac{1}{2} + \lambda^{-O(1)}$, converts any (\mathcal{X}_1, α) -correct sub-exponentially secure IO scheme for P/poly into an (almost) exact IO scheme for P/poly .*

¹Graduate students do not count.

Theorem 1.2 (informal). *Assuming weak PRFs in NC^1 , there exists an input sampler \mathcal{X}_2 , polynomial poly_2 , and a transformation that for any $\alpha \geq 1 - \frac{1}{\text{poly}_2(\lambda)}$, converts any (\mathcal{X}_2, α) -correct sub-exponentially secure IO scheme for P/poly into an (almost) exact IO scheme for P/poly .*

Since the works of [CKP15, PS15, MMN15] apply to any efficient sampler \mathcal{X} and any α that is polynomially bounded away from 1, we obtain the following main corollary

Corollary (Main Theorems + [CKP15, PS15, MMN15]). *Assume that there is an indistinguishability obfuscator in either the random oracle model, the ideal generic group/ring model, or ideal multilinear maps model with bounded multi-linearity. Then, there is an (almost) exact obfuscator in the plain model.*

We also show how to transform approximate functional encryption into exact functional encryption, where approximate FE is defined analogously to approximate IO with respect to a distribution on the message space and decryption errors. Beside being of independent interest, this transformation will also be used as a building block to obtain the second theorem above.

Theorem 1.3 (Informal). *Assuming weak PRFs in NC^1 , there exists a message sampler \mathcal{X} , constant η , and a transformation that for any $\alpha \geq 1 - \eta$, converts any (\mathcal{X}, α) -correct FE scheme for P/poly into an (almost) exact scheme FE scheme for P/poly .*

We note that our theorems result in IO (or FE) that may still output an erroneous obfuscator, but only with some negligible probability over the coins of the obfuscator (FE setup algorithm). This is analogous to the setting of correcting decryption errors in plain public key encryption [DNR04], and as far as we know is sufficient in all applications.

We now proceed to provide an overview of our techniques.

1.2 Overview of Our Techniques

The starting point of our constructions comes from the notion of random self-reducibility [AFK89]. That is, imagine that you have an error-prone algorithm A that computes a (Boolean) function F correctly on a $1/2 + \varepsilon$ fraction of inputs, and suppose that there is an efficient randomizer that maps an input x into a uniformly random input $r = r(x)$ such that given $F(r)$, one can efficiently recover $F(x)$. To turn this into an error-less algorithm that is correct on all inputs, one proceeds as follows:

- On input x , apply the randomizer $N \gg \varepsilon^{-2}$ times to sample uniformly random and independent inputs r_1, r_2, \dots, r_N ;
- Run the error-prone algorithm on r_1, r_2, \dots, r_N to obtain answers $A(r_1), A(r_2), \dots, A(r_N)$;
- Output the majority of the answers.

Since the r_i are uniformly random inputs to A , we know that with high probability, (roughly) $1/2 + \Omega(\varepsilon)$ of the answers $A(r_i)$ are correct, and thus enable us to recover the correct value of $F(x)$. A majority vote then seals the deal.

In our setting, F is an arbitrary function, which is likely *not* random self-reducible. Nevertheless, we show how to make the *essence* of this idea work, using various notions of (two-party and multi-party) non-interactive secure function evaluation (SFE) [Yao86, BGW88, Gen09]. Indeed, some

form of non-interactive SFE (or homomorphic encryption) has been used in several instances in the literature to obtain a (sometimes computational) random self-reducibility [AIK06, CKV10, BP12, BGJ⁺15]. The rough idea is that if we can get the obfuscator to homomorphically evaluate a given function on encryptions for some fixed distribution on inputs, then it must also behave correctly with roughly the same probability on encryptions of any arbitrary input. This, however, should be done with care to ensure that homomorphic evaluation does not harm the security of the obfuscator. We next go into more details on how we carry out this agenda.

Our First Construction. Our first construction uses a two-party non-interactive secure function evaluation protocol with security against malicious senders. For simplicity, let us describe this approach in the language of fully homomorphic encryption (FHE). Let $(\text{Enc}, \text{Dec}, \text{Eval})$ be a (secret-key) FHE scheme (not necessarily compact). (We will assume that the randomness of the key generation algorithm acts as the secret key, and avoid explicitly dealing with the key generation algorithm.)

To exactly obfuscate a circuit C , we use the approximate obfuscator $\widetilde{i\mathcal{O}}$ to obfuscate the circuit Eval_C which, on input an encryption of some x , homomorphically computes an encryption of $C(x)$. Assume that $\widetilde{i\mathcal{O}}(\text{Eval}_C)$ is correct on a $1/2 + \varepsilon$ fraction of encryptions of 0^n . The key observation is that semantic security of the encryption scheme means that $\widetilde{i\mathcal{O}}(\text{Eval}_C)$ is also correct on a $1/2 + \varepsilon - \lambda^{-\omega(1)}$ fraction of encryptions of any x , that is, it outputs $\text{Eval}_C(\text{Enc}(x)) = \text{Enc}(C(x))$. Taking a majority of many invocations then gives us correctness *for every input* x .

The problem with this idea is the security of the final obfuscator. Indeed, $\text{Eval}_C(\text{Enc}(x))$ may reveal information about the circuit C beyond the output $C(x)$. The problem goes even further: since the evaluator in this setting is untrusted, she can try to run the obfuscated circuits with malformed encryptions, potentially making the problem much worse. The solution is to rely on a *maliciously function-hiding* homomorphic encryption scheme. Such an object can be constructed using Yao’s garbled circuits combined with an oblivious transfer (OT) protocol secure against malicious receivers (such as the Naor-Pinkas protocol based on the DDH assumption [NP01]). The evaluation procedure, however, is randomized, which we can derandomize relying on a pseudo-random function.

While the above works perfectly assuming ideal VBB obfuscation, this is not necessarily the case for IO. Nevertheless, we observe that we can use $\widetilde{i\mathcal{O}}$ to obfuscate this randomized circuit using the machinery of probabilistic IO [CLTV15]. This allows us to show that indistinguishability obfuscation is maintained, but requires going through an exponential number of hybrids requiring sub-exponential security from $\widetilde{i\mathcal{O}}$ (and the some of other involved primitives).

Our Second Construction. Our second construction goes through the notion of functional encryption (FE). In a (public-key, selectively indistinguishability-based) FE scheme, the owner of a functional secret key FSK_F can “decrypt” a ciphertext $\text{FE.Enc}(\text{MPK}, m)$ to learn $F(m)$ (but should learn nothing else about m). In an approximately correct FE scheme, the decryption algorithm could err on encryptions of certain messages m , but should be correct with probability $1 - \varepsilon$ on inputs x drawn from a (sampleable) distribution \mathcal{X} .

We show how to transform any approximately correct FE scheme into an exact FE scheme. Here the main advantage over the setting of approximate IO is that we are only concerned with honestly generated encrypted messages and are not concerned with function hiding. In particular, we can relax the assumptions required for the SFE and rely on (a non-interactive) information-theoretic version of the Ben-Or-Goldwasser-Wigderson [BGW88] multi-party computation protocol for NC^1 .

This construction also provides an alternative route for the IO transformation. Concretely, we show that starting from approximate IO, we can first apply the transformation of Garg et al. [GGH⁺13b] to obtain approximate FE. For this to work, we need show how to obtain (exact) NIZKs and public-key encryption directly from approximate IO, which are necessary transformation. Then in the second step, we apply we apply our exact-to-approximate transformation for FE, and finally invoke a transformation from (exact) FE to IO [AJ15, BV15] (for which we need to rely on certain compactness properties of our transformation).

Overall, this leads to a construction of (exact) IO from subexponentially-secure approximate IO and one-way functions² for NC¹, and can be bootstrapped to all circuits assuming in addition weak PRFs in NC¹ [ABSV14].

Organization. In Section 2, we define the required tools for our transformations, including the forms of SFE that we rely on. In Section 3, we describe our first basic transformation from approximate to exact IO. In Section 4, we describe our transformation from approximate to exact FE. In Section 5, we describe our second transformation for IO, going through our transformation for FE.

2 Preliminaries

The cryptographic definitions in the paper follow the convention of modeling security against non-uniform adversaries. An efficient adversary \mathcal{A} is modeled as a sequence of circuits $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, such that each circuit \mathcal{A}_λ is of polynomial size $\lambda^{O(1)}$ with $\lambda^{O(1)}$ input and output bits. We often omit the subscript λ when it is clear from the context.

When we refer to a randomized algorithm \mathcal{A} , we typically do not explicitly denote its random coins, and use the notation $s \leftarrow \mathcal{A}$ or $s \leftarrow \mathcal{A}(x)$ if \mathcal{A} has an extra input x . When we want to be explicit regarding the coins, we shall denote $s \leftarrow \mathcal{A}(r)$, or $s \leftarrow \mathcal{A}(x; r)$, respectively.

Whenever we refer to a circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}$, we mean that each set \mathcal{C}_λ consists of Boolean circuits of size at most $\text{poly}(\lambda)$ for some polynomial $\text{poly}(\cdot)$, defined on the domain $\{0, 1\}^{n(\lambda)}$. When referring to inputs $x \in \{0, 1\}^{n(\lambda)}$, we often omit λ from the notation.

2.1 Non-Interactive Secure Function Evaluation

We consider two-message secure function evaluation (SFE) protocols. Typically, such a protocol consists of two parties (A, B) and has the following syntax. Party A is given input x , encrypts x and sends the encrypted input to B . B given as additional input a function f , homomorphically evaluates f on the encrypted x , and returns the result to A , who can then decrypt the result $f(x)$. The protocol is required to ensure input-privacy for A and function privacy for B (on top of correctness).

Definition 2.1 (Secure Function Evaluation). *A scheme $\text{SFE} = (\text{Enc}, \text{Eval}, \text{Dec})$, where Enc, Eval are probabilistic and Dec is deterministic, is a two-message secure function evaluation protocol for circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}$, where \mathcal{C}_λ is defined over $\{0, 1\}^{n(\lambda)}$, if the following requirements hold:*

²Which can be replaced with a mild complexity assumption [KMN⁺14].

- **Correctness:** for any $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$ and input $x \in \{0, 1\}^n$ in the domain of C it holds that:

$$\Pr \left[\text{Dec}(\mathbf{R}, \widehat{\text{CT}}) = C(x) \mid \begin{array}{l} (\text{CT}, \mathbf{R}) \leftarrow \text{Enc}(x) \\ \widehat{\text{CT}} \leftarrow \text{Eval}(\text{CT}, C) \end{array} \right] \geq 1 - \mu(\lambda) ,$$

for some negligible $\mu(\cdot)$, where the probability is over the coin tosses of Enc and Eval .

- **Input Hiding:** for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, and equal size inputs $x_0, x_1 \in \{0, 1\}^n$:

$$|\Pr[\mathcal{D}(\text{CT}_0) = 1] - \Pr[\mathcal{D}(\text{CT}_1) = 1]| \leq \mu(\lambda) ,$$

where $\text{CT}_b \leftarrow \text{Enc}(x_b)$.

- **Malicious Function Hiding:** there exists a (possibly inefficient) function Ext , such that for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, maliciously chosen CT^* , and equal size circuits $C_0, C_1 \in \mathcal{C}_\lambda$ that agree on $x = \text{Ext}(\text{CT})$:

$$|\Pr[\mathcal{D}(\widehat{\text{CT}}_0) = 1] - \Pr[\mathcal{D}(\widehat{\text{CT}}_1) = 1]| \leq \mu(\lambda) ,$$

where $\widehat{\text{CT}}_b \leftarrow \text{Eval}(\text{CT}^*, C_b)$.

We say that the scheme is δ -function-hiding, for some concrete negligible function $\delta(\cdot)$, if for all poly-size distinguishers, the above indistinguishability gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

Remark 2.2 (strong function privacy). For our most basic transformation from approximate IO to exact IO, we will require $2^{-\sigma(\lambda)} \cdot \lambda^{-\omega(1)}$ -function-hiding, where $\sigma(\lambda)$ is the size of encryptions in the scheme. Below, we discuss an instantiation, based on the DDH assumption, that has perfect function-hiding, and thus satisfies this requirement.

Distributed secure function evaluation. We will also consider a notion of two-message distributed function evaluation (DSFE). Such a protocol consists of $k + 2$ parties (A, B_1, \dots, B_k, C) and has the following syntax. Party A , given input x , shares x into k shares and sends the shares to B_1, \dots, B_k . The parties B_1, \dots, B_k given as additional input a function f , homomorphically and non-interactively evaluate f on each share, and send the evaluated shares to C , who can then decrypt and obtain the result $f(x)$.

The protocol is required to ensure that each individual share sent by A in the second message hides all information regarding the input x . We also require that C gains no information on the input, except for the output of the function (formally, we will require an indistinguishability-based guarantee analogous to that of functional encryption.) Furthermore, we will require that correctness holds even if some τ fraction of the parties B_1, \dots, B_k are faulty.

Definition 2.3 (Distributed Secure Function Evaluation). *A scheme DSFE = (Enc, Eval, Dec), where Enc is probabilistic and Eval, Dec are deterministic, is a (k, τ) -secure distributed function evaluation protocol for circuit class $\mathcal{C} = \{\mathcal{C}_\lambda\}$, where \mathcal{C}_λ is defined over $\{0, 1\}^n$ for $n = n(\lambda)$, $k = k(\lambda)$, and $\tau = \tau(\lambda)$, if the following requirements hold:*

- **Correctness in the presence of faults:** for any $\lambda \in \mathbb{N}$, $C \in \mathcal{C}_\lambda$ and input $x \in \{0, 1\}^n$ in the domain of C and any set $S \in [k]$ of size smaller than τn , and functions $\{\text{Err}_i : i \in S\}$ it holds that:

$$\Pr \left[\text{Dec}(R, \widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k) = C(x) \mid \begin{array}{l} (\text{CT}_1, \dots, \text{CT}_k, R) \leftarrow \text{Enc}(x) \\ \forall i \in [k] \setminus S : \widehat{\text{CT}}_i = \text{Eval}(\text{CT}_i, C) \\ \forall i \in S : \widehat{\text{CT}}_i \leftarrow \text{Err}_i(\text{CT}_i) \end{array} \right] \geq 1 - \mu(\lambda) ,$$

for some negligible $\mu(\cdot)$, where the probability is over the coin-tosses of Enc .

- **Input Hiding:** for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, and equal size inputs $x_0, x_1 \in \{0, 1\}^n$ and any $i \in [k]$:

$$|\Pr[\mathcal{D}(\text{CT}_{0,i}) = 1] - \Pr[\mathcal{D}(\text{CT}_{1,i}) = 1]| \leq \mu(\lambda) ,$$

where $\text{CT}_{b,i}$ denotes the i -th ciphertext output by $\text{Enc}(x_b)$.

- **Residual Input Hiding:** for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, inputs $x_0, x_1 \in \{0, 1\}^n$, and circuit $C \in \mathcal{C}_\lambda$ such that $C(x_0) = C(x_1)$:

$$\left| \Pr[\mathcal{D}(R_0, \widehat{\text{CT}}_{0,1}, \dots, \widehat{\text{CT}}_{0,k}) = 1] - \Pr[\mathcal{D}(R_1, \widehat{\text{CT}}_{1,1}, \dots, \widehat{\text{CT}}_{1,k}) = 1] \right| \leq \mu(\lambda) ,$$

where for $(b, i) \in \{0, 1\} \times [k]$, $\widehat{\text{CT}}_{b,i} = \text{Eval}(\text{CT}_{b,i}, C)$, and $(\text{CT}_{b,1}, \dots, \text{CT}_{b,k}, R_b) \leftarrow \text{Enc}(x_b)$.

Remark 2.4 (difference from SFE). There are two main differences from SFE. The first is in security, in the above we do not require any type of function-hiding, but require residual input-hiding. The second is the functionality: we allow distributed evaluation (with some resilience to faults). The second difference is not essential, and is considered in order to reduce the underlying computational assumptions. In particular, a (non-distributed) SFE with residual input-hiding implies DSFE with $k = 1, \tau = 0$.

Remark 2.5 (deterministic Eval). Jumping ahead, we remark that we will use distributed SFE in a setting where the encryptor is always honest. Since we are not requiring any privacy against the encryptor, we may assume w.l.o.g that Eval is deterministic. Indeed, we can always sample any required randomness as part of the encryption process and embed it in the shares $\text{CT}_1, \dots, \text{CT}_k$.

2.1.1 Instantiations

We now mention known instantiations of SFE and DSFE schemes, which we can rely on.

SFE. As mentioned above, for our application, we will require rather strong function-hiding. To instantiate the scheme we may rely on the SFE protocol obtained by using the oblivious transfer protocol of Naor and Pinkas [NP01] that is based on DDH and is secure against unbounded receivers in conjunction with an information-theoretic variant of Yao's garbled circuit [Yao86] for NC^1 [IK02]. The resulting SFE scheme is for classes of circuits in NC^1 , which will suffice for our purposes. Alternatively, we can use a strong enough computational variant of Yao based on sub-exponential one-way functions, resulting in a construction for all polynomial-size circuits.

More generally, the Naor-Pinkas OT can be replaced with any OT that had statistical function-hiding. In the CRS model, such two-message protocols exist from other standard assumptions as well [PVW08]. While our main transformation is described using SFE in the plain model, it can be naturally extended to the CRS setting Remark 3.6.

DSFE. An information-theoretically secure DSFE scheme for circuit classes in NC^1 can be obtained based on a non-interactive variant of the BGW protocol [BGW88] similar to that used in [GVW12]. For the sake of completeness, we now outline this variant.

Given a class of circuits $\mathcal{C} = \{\mathcal{C}_\lambda\}$ in NC^1 defined on inputs in $\{0, 1\}^{n(\lambda)}$, we can interpret it as a class of arithmetic circuits where any circuit C is defined over inputs $(x_1, \dots, x_n) \in \mathbb{F}^n$, and computes a polynomial of total degree at most $D = 2^d$, where $d = d(\lambda)$ is the maximal depth of any circuit in \mathcal{C}_λ .

At a high-level, sharing the inputs in the scheme corresponds to encoding them using the Shamir secret-sharing scheme (that is, as random Reed-Solomon polynomials), evaluation corresponds to homomorphic evaluation over the polynomials, and residual input hiding is guaranteed by adding a random zero polynomial to the evaluated shares.

Concretely, the scheme is defined as follows. Fix a field \mathbb{F} , such that $|\mathbb{F}| \geq 3D + 1$, and let $k = 3D + 1$. Let $\alpha_1, \dots, \alpha_k$ be k distinct elements in the field.

- $\text{Enc}(x_1, \dots, x_n)$:
 1. sample n random degree-one polynomials $p_1(\cdot), \dots, p_n(\cdot)$, where $p_i(0) = x_i$,
 2. sample a random degree D polynomial $z(\cdot)$ such that $z(0) = 0$.
 3. set $\text{CT}_j = p_1(\alpha_j), \dots, p_n(\alpha_j), z(\alpha_j)$.
 4. output $\text{CT}_1, \dots, \text{CT}_k$.
- $\text{Eval}(\text{CT}_j, C)$:
 1. parse $\text{CT}_j = \pi_1, \dots, \pi_n, \gamma$,
 2. consider the univariate polynomial $E(\cdot) = C(p_1(\cdot), \dots, p_n(\cdot))$ (that has degree at most D), and compute homomorphically $\eta = C(\pi_1, \dots, \pi_n)$. (The result is meant to be $E(\alpha_j)$.)
 3. output $\widehat{\text{CT}}_j = \eta + \gamma$.
- $\text{Dec}(\widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k)$:
 1. parse $(\widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k)$ as a Reed-Solomon codeword in \mathbb{F}^k , and decode a polynomial $\tilde{E}(\cdot)$,
 2. output $\tilde{E}(0)$.

We claim that the above scheme is $(k, \frac{1}{3})$ -secure. The analysis follows the standard BGW analysis (detailed in [AL11]). Very roughly, to show correctness, note that by the homomorphic properties of the Reed-Solomon code the correct polynomial E is such that $E(0) = C(x_1, \dots, x_n)$, and this also holds for $E(\cdot) + z(\cdot)$. Furthermore, decoding such that $\tilde{E} = E + z$ is guaranteed as long as there are at most D faults. Input-hiding follows from the fact that each individual CT_j is distributed uniformly at random on \mathbb{F} . Residual input-hiding follows by the fact that after adding z , the new polynomial $E + z$ is a random polynomial with free coefficient $C(x_1, \dots, x_n)$, and thus can be completely simulated from this value. For more details, see [BGW88, AL11].

Remark 2.6 (complexity of encryption). One measure of interest, when considering our application to correcting errors in functional encryption, will be the complexity of the encryption procedure. Here we note that this size is $O(nk \log D) = O(nD \log D) = n \cdot 2^{O(d)}$; namely, it does not grow with the size of the circuits, but does grow exponentially with the maximal depth d of the circuits. (As will be discussed later on, this will still be good enough in our context, to bootstrap functional encryption to indistinguishability obfuscation, as shown in [AJ15, BV15].)

One point to notice is that the above is not entirely accurate if the output of the circuit C is a large $\ell = \ell(\lambda)$. Indeed, naïvely to guarantee residual input-privacy, we will need to generate ℓ separate polynomials z_1, \dots, z_ℓ , meaning the encryption size will grow linearly with ℓ . This can be avoided by shifting the randomness to the evaluation procedure (which will slightly complicate our transformation). Alternatively, this can be avoided assuming the existence of a pseudo-random generator, by adding to the ciphertexts a seed, and having Eval use it to generate the multiple polynomials z_1, \dots, z_ℓ .

2.2 Symmetric Encryption

A symmetric encryption scheme Sym consists of a tuple of two PPT algorithms ($\text{Sym.Enc}, \text{Sym.Dec}$). The encryption algorithm takes as input a symmetric key $\text{SK} \in \{0, 1\}^\lambda$, where λ is the security parameter, and a message $m \in \{0, 1\}^*$ of polynomial size in the security parameter, and outputs a ciphertext SCT . The decryption algorithm takes as input (SK, SCT) , and outputs the decrypted message m . For this work, we only require one-time security.

Definition 2.7 (One-Time Symmetric Encryption). *A pair of PPT algorithms ($\text{Sym.Enc}, \text{Sym.Dec}$) is a one-time symmetric encryption scheme for message space $\{0, 1\}^*$ if it satisfies:*

1. **Correctness:** For every security parameter λ and message $m \in \{0, 1\}^*$,

$$\Pr \left[\text{Sym.Dec}(\text{SK}, \text{SCT}) = m \mid \begin{array}{l} \text{SK} \leftarrow \{0, 1\}^\lambda \\ \text{SCT} \leftarrow \text{Sym.Enc}(\text{SK}, m) \end{array} \right] = 1 .$$

2. **Indistinguishability:** for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, and any equal size messages m_0, m_1 ,

$$|\Pr[\mathcal{D}(\text{Sym.Enc}(\text{SK}, m_0)) = 1] - \Pr[\mathcal{D}(\text{Sym.Enc}(\text{SK}, m_1)) = 1]| \leq \mu(\lambda) ,$$

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

We further say that Sym is δ -secure, for some concrete negligible function $\delta(\cdot)$, if for all polysize distinguishers the above indistinguishability gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

A symmetric encryption scheme meeting this definition can be constructed from any pseudo-random generator, and thus any one-way function [HILL99].

2.3 Puncturable Pseudorandom Functions

We consider a simple case of puncturable pseudo-random functions (PRFs) where any PRF may be punctured at a single point. The definition is formulated as in [SW14], and is satisfied by the Goldreich-Goldwasser-Micali PRF construction [GGM86, BW13, KPTZ13, BGI14].

Definition 2.8 (Puncturable PRFs). *Let n, k be polynomially bounded length functions. An efficiently computable family of functions*

$$\mathcal{PRF} = \left\{ \text{PRF}_K : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda \mid K \in \{0, 1\}^{k(\lambda)}, \lambda \in \mathbb{N} \right\} ,$$

associated with an efficient (probabilistic) key sampler $\text{Gen}_{\mathcal{PRF}}$, is a puncturable PRF if there exists a poly-time puncturing algorithm Punc that takes as input a key K , and a point x^ , and outputs a punctured key $K\{x^*\}$, so that the following conditions are satisfied:*

1. **Functionality is preserved under puncturing:** *For every $x^* \in \{0, 1\}^*$,*

$$\Pr_{K \leftarrow \text{Gen}_{\mathcal{PRF}}(1^\lambda)} [\forall x \neq x^* : \text{PRF}_K(x) = \text{PRF}_{K\{x^*\}}(x) \mid K\{x^*\} = \text{Punc}(K, x^*)] = 1 .$$

2. **Indistinguishability at punctured points:** *for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for all $\lambda \in \mathbb{N}$, and any $x^* \in \{0, 1\}^*$,*

$$|\Pr[\mathcal{D}(x^*, K\{x^*\}, \text{PRF}_K(x^*)) = 1] - \Pr[\mathcal{D}(x^*, K\{x^*\}, u) = 1]| \leq \mu(\lambda) ,$$

where $K \leftarrow \text{Gen}_{\mathcal{PRF}}(1^\lambda)$, $K\{x^\} = \text{Punc}(K, x^*)$, and $u \leftarrow \{0, 1\}^\lambda$.*

We further say that \mathcal{PRF} is δ -secure, for some concrete negligible function $\delta(\cdot)$, if for all polysize distinguishers the above indistinguishability gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

Remark 2.9 (uniform output). For some of our constructions, it will be convenient to assume that for any fixed x , $\text{PRF}_K(x)$ is distributed uniformly at random (when K is sampled at random). It is not hard to see that such a puncturable PRF can be easily obtained from any puncturable PRF by adding a random string U to the key and XORing U to every output.

3 Correcting Errors in Indistinguishability Obfuscation

In this section, we define approximate IO and show how to transform any approximate IO to (almost) perfectly correct IO, based on SFE.

3.1 Approximate and Exact IO

We start by defining indistinguishability obfuscation (IO) with (almost) perfect correctness. The definition is formulated as in [BGI⁺12].

Definition 3.1 (Indistinguishability obfuscation). *A PPT algorithm $i\mathcal{O}$ is said to be an indistinguishability obfuscator for a class of circuits $\mathcal{C} = \{C_\lambda\}$, if it satisfies:*

1. **(Almost) Perfect Correctness:** *for any security parameter λ and $C \in \mathcal{C}_\lambda$,*

$$\Pr_{i\mathcal{O}} [\forall x : i\mathcal{O}(C, 1^\lambda)(x) = C(x)] \geq 1 - \mu(\lambda) ,$$

for some negligible $\mu(\cdot)$.

2. **Indistinguishability:** for any polysize distinguisher \mathcal{D} there exists a negligible function $\mu(\cdot)$, such that for any two circuits $C_0, C_1 \in \mathcal{C}$ that compute the same function and are of the same size:

$$\left| \Pr[\mathcal{D}(i\mathcal{O}(C_0, 1^\lambda)) = 1] - \Pr[\mathcal{D}(i\mathcal{O}(C_1, 1^\lambda)) = 1] \right| \leq \mu(\lambda) ,$$

where the probability is over the coins of \mathcal{D} and $i\mathcal{O}$.

We further say that $i\mathcal{O}$ is δ -secure, for some concrete negligible function $\delta(\cdot)$, if for all polysize distinguishers the above indistinguishability gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

We now define an approximate notion of correctness that allows the obfuscated circuit to err with some probability over inputs taken from some samplable distribution.

Definition 3.2 ((α, \mathcal{X}) -correct IO). For $\alpha(\lambda) \in [0, 1]$ and an ensemble of input samplers $\mathcal{X} = \{\mathcal{X}_\lambda\}$, we say that $i\mathcal{O}$ is (α, \mathcal{X}) -correct if instead of (almost) perfect correctness, it satisfies the following relaxed requirement:

1. **Approximate Correctness:** for any security parameter λ , $C \in \mathcal{C}_\lambda$,

$$\Pr \left[i\mathcal{O}(C, 1^\lambda)(x) = C(x) \mid x \leftarrow \mathcal{X}_\lambda \right] \geq \alpha(\lambda) ,$$

where the probability is also over the coins of $i\mathcal{O}$.

3.2 The Transformation

We now describe a transformation from approximately correct IO to (almost) perfectly correct IO and analyze it. The transformation is based on SFE satisfying a strong function-hiding guarantee. We discuss an instantiation based on standard computational assumptions in Section 3.3.

In Section 5, we discuss an alternative transformation through functional encryption based on weaker computational assumptions.

Ingredients. In the following, let λ denote a security parameter, let $\varepsilon < 1$ be some constant, $\eta(\lambda) = \lambda^{-\Omega(1)}$ and let $\mathcal{C} = \{\mathcal{C}_\lambda\}$ denote a circuit class. We rely on the following primitives:

- A secure function evaluation scheme SFE for \mathcal{C} that is $2^{-\omega(\sigma(\lambda) + \log \lambda)}$ -function-hiding, where $\sigma(\lambda)$ is the length of fresh ciphertexts generated by the encryption algorithm Enc for security parameter λ (and inputs of size $n = n(\lambda)$ in the domain of \mathcal{C}_λ).
- A $2^{-\tilde{\lambda}^\varepsilon}$ -secure puncturable pseudo-random function family \mathcal{PRF} , where the security parameter is $\tilde{\lambda} = \omega(\sigma(\lambda) + \log \lambda)^{1/\varepsilon}$.
- A $(\frac{1}{2} + \eta(\lambda), \mathcal{X})$ -correct, $2^{-\tilde{\lambda}^\varepsilon}$ -secure indistinguishability obfuscator $\widetilde{i\mathcal{O}}$ for $\overline{\mathcal{C}}$, where the security parameter is $\tilde{\lambda} = \omega(\sigma(\lambda) + \log \lambda)^{1/\varepsilon}$. The sampler class \mathcal{X} depends on SFE and the class $\overline{\mathcal{C}}$ depends on SFE, \mathcal{PRF} , and \mathcal{C} . Both \mathcal{X} and $\overline{\mathcal{C}}$ are specified below as part of the description of the constructed (exact) obfuscator $i\mathcal{O}$.

The exact obfuscator $i\mathcal{O}$:

Given a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}$ and security parameter λ , the obfuscator $i\mathcal{O}(C, 1^\lambda)$

1. computes a new security parameter $\tilde{\lambda} = \omega(\sigma(\lambda) + \log \lambda)^{1/\varepsilon}$, where $\sigma(\lambda)$ is the length of ciphertexts as defined above, as well as a parameter $N = \frac{\omega(n + \log \lambda)}{\eta^2(\lambda)}$,
2. samples puncturable PRF seeds K_1, \dots, K_N , where $K \leftarrow \text{Gen}_{\mathcal{PRF}}(1^{\tilde{\lambda}})$,
3. computes the augmented C -evaluation circuits C_{K_1}, \dots, C_{K_N} defined in Figure 1,
4. outputs N approximate obfuscations $\tilde{C}_1, \dots, \tilde{C}_N$, where $\tilde{C}_i \leftarrow \tilde{\mathcal{O}}(C_{K_i}, 1^{\tilde{\lambda}})$, and N random strings r_1, \dots, r_N , where $r_i \leftarrow \{0, 1\}^\lambda$.

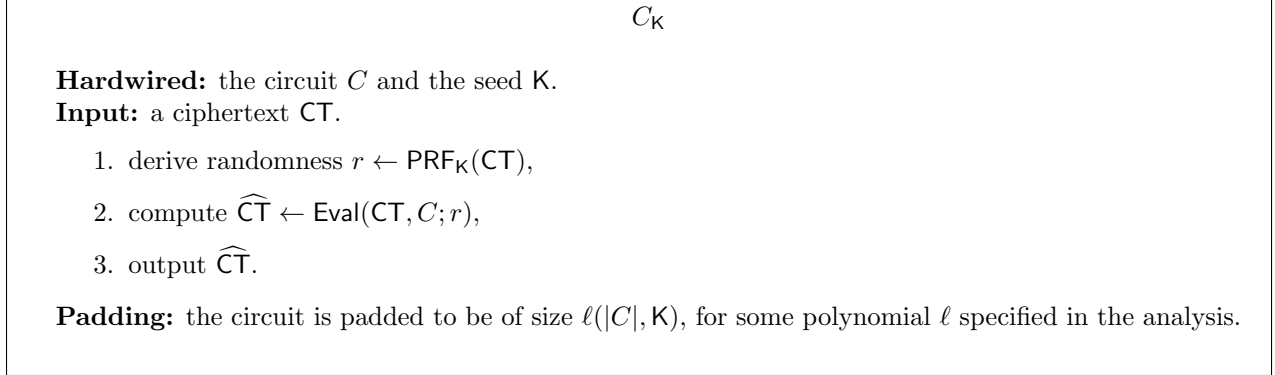


Figure 1: The augmented C -evaluation circuit.

Remark 3.3 (deterministic evaluator). Publishing the random strings r_i is done to match the usual obfuscation syntax where the evaluation is deterministic. We may also let the evaluator sample this randomness.

Remark 3.4 (reusing K). In the above construction we can reuse K in all N approximate obfuscations. Sampling independent seeds simplifies the analysis.

We next describe the how the obfuscation $(\tilde{C}_1, r_1), \dots, (\tilde{C}_N, r_N)$ is evaluated.

Evaluation:

Given an obfuscation $\{\tilde{C}_i, r_i \mid i \in [N]\}$, an input $x \in \{0, 1\}^n$, and security parameter λ :

1. For $i \in [N]$,
 - compute $(\text{CT}_i, R_i) \leftarrow \text{Enc}(x; r_i)$,
 - compute $\widehat{\text{CT}}_i = \tilde{C}_i(\text{CT}_i)$,
 - $y_i = \text{Dec}(R_i, \widehat{\text{CT}}_i)$.
2. Output $y = \text{majority}(y_1, \dots, y_N)$.

The ensemble of samplers \mathcal{X} consists of samplers \mathcal{X}^0 that sample encryptions from $\text{Enc}(0^n)$ whereas the class $\bar{\mathcal{C}}$ consists of circuits C_K as defined in Figure 1.

Theorem 3.5. *$i\mathcal{O}$ is an (almost) perfectly correct indistinguishability obfuscator.*

The intuition behind the proof is outlined in the introduction. We now turn to the actual proof.

Proof. We first prove that the new obfuscator is (almost) exactly correct, and then prove that it is secure.

Almost exact correctness. For any $\lambda, n = n(\lambda)$, input $x \in \{0, 1\}^n$, let us denote $\mathcal{X}^x(r) := \text{Enc}(x; r)$ a sampler for encryptions of x . Then, by the input-hiding guarantee of SFE, and the approximate correctness of $i\tilde{\mathcal{O}}$, we claim that the approximate obfuscation is correct on encryptions of an arbitrary $x \in \{0, 1\}^n$ as on encryptions of 0^n . That is, there exists a negligible $\mu(\lambda)$ such that

$$\begin{aligned} & \Pr \left[\tilde{C}(\text{CT}) = C_{\mathbf{K}}(\text{CT}) \mid \text{CT} \leftarrow \mathcal{X}^x \right] \geq \\ & \Pr \left[\tilde{C}(\text{CT}) = C_{\mathbf{K}}(\text{CT}) \mid \text{CT} \leftarrow \mathcal{X}^0 \right] - \mu(\lambda) \geq \\ & \frac{1}{2} + \eta(\lambda) - \mu(\lambda) , \end{aligned}$$

where in both of the above $\mathbf{K} \leftarrow \text{Gen}_{\mathcal{P}\mathcal{R}\mathcal{F}}(1^{\tilde{\lambda}})$, $\tilde{C} \leftarrow i\tilde{\mathcal{O}}(C_{\mathbf{K}}, 1^{\tilde{\lambda}})$.

It now follows that any one of the N decryption attempts is correct with probability noticeably larger than half. Concretely,

$$\begin{aligned} & \Pr \left[\text{Dec}(\mathbf{R}, \widehat{\text{CT}}) = C(x) \mid \begin{array}{l} \text{CT}, \mathbf{R} \leftarrow \text{Enc}(x) \\ \widehat{\text{CT}} = \tilde{C}(\text{CT}) \end{array} \right] \geq \\ & \Pr \left[\text{Dec}(\mathbf{R}, \widehat{\text{CT}}) = C(x) \mid \begin{array}{l} \text{CT}, \mathbf{R} \leftarrow \text{Enc}(x) \\ \widehat{\text{CT}} = C_{\mathbf{K}}(\text{CT}) \end{array} \right] \cdot \Pr \left[\tilde{C}(\text{CT}) = C_{\mathbf{K}}(\text{CT}) \mid \text{CT} \leftarrow \mathcal{X}^x \right] = \\ & \Pr \left[\text{Dec}(\mathbf{R}, \widehat{\text{CT}}) = C(x) \mid \begin{array}{l} \text{CT}, \mathbf{R} \leftarrow \text{Enc}(x) \\ r = \text{PRF}_{\mathbf{K}}(\text{CT}) \\ \widehat{\text{CT}} = \text{Eval}(\text{CT}, C; r) \end{array} \right] \cdot \Pr \left[\tilde{C}(\text{CT}) = C_{\mathbf{K}}(\text{CT}) \mid \text{CT} \leftarrow \mathcal{X}^x \right] \geq \\ & (1 - \mu'(\lambda)) \cdot \left(\frac{1}{2} + \eta(\lambda) - \mu(\lambda) \right) , \end{aligned}$$

where in all of the above $\mathbf{K} \leftarrow \text{Gen}_{\mathcal{P}\mathcal{R}\mathcal{F}}(1^{\tilde{\lambda}})$, $\tilde{C} \leftarrow i\tilde{\mathcal{O}}(C_{\mathbf{K}}, 1^{\tilde{\lambda}})$, and $\mu'(\cdot)$ is some negligible function (corresponding to the negligible decryption error of SFE). In the last step, we relied on the fact that for any fixed CT, $\text{PRF}_{\mathbf{K}}(\text{CT})$ is distributed uniformly at random (Remark 2.9), and the (almost) perfect correctness of SFE.

By a Chernoff bound, for large enough λ , and any $x \in \{0, 1\}^n$, the probability that the majority value y among all decrypted y_1, \dots, y_N is incorrect is bounded by

$$\Pr [y \neq C(x)] \leq 2^{-\Omega(N \cdot \eta^2(\lambda))} \leq 2^{-n + \omega(\log \lambda)} .$$

The required correctness for all input then follows by a union bound over all inputs in $\{0, 1\}^n$.

Security analysis. Consistently with the notation above, for $\mathbf{K} \leftarrow \text{Gen}_{\mathcal{P}\mathcal{R}\mathcal{F}}(1^{\tilde{\lambda}})$, and a circuit $C \in \mathcal{C}_{\lambda}$, we denote by $\tilde{C} \leftarrow i\tilde{\mathcal{O}}(C_{\mathbf{K}}, 1^{\tilde{\lambda}})$ the corresponding approximate obfuscation of the (derandomized) evaluation circuit. Recalling that each obfuscation is made of polynomially-many independent copies of (\tilde{C}, r) , and that r is a random string, independent of \tilde{C} , it suffices (by a standard hybrid

argument) to show that for any polysize distinguisher there exists a negligible $\mu(\cdot)$, such that for any $C_0, C_1 \in \mathcal{C}_\lambda$ that compute the same function it holds that

$$\left| \Pr[\mathcal{D}(\tilde{C}_0) = 1] - \Pr[\mathcal{D}(\tilde{C}_1) = 1] \right| \leq \mu(\lambda) .$$

The above follows from the fact that the output of the two underlying obfuscated circuits on any point $\text{CT} \in \{0, 1\}^{\sigma(\lambda)}$ is indistinguishable even given C_0, C_1 . Indeed, because the circuits C_0, C_1 compute the same function, by the function-hiding of SFE, for any ciphertext $\text{CT} \in \{0, 1\}^{\sigma(\lambda)}$, the evaluated ciphers $\text{Eval}(\text{CT}, C_0)$ and $\text{Eval}(\text{CT}, C_1)$ are indistinguishable. Canetti, Lin, Tessaro, and Vaikuntanathan [CLTV15] show that (sub-exponential) IO in conjunction with (sub-exponential) puncturable PRFs are sufficient in this setting, which they formalize by *probabilistic IO* notion. For the sake of completeness, we next sketch the argument.

We consider a sequence of $2^\sigma + 1$ hybrids $\{\mathcal{H}_{\text{CT}}\}_{\text{CT} \in \{0, \dots, 2^\sigma\}}$, where we naturally identify integers in $[2^\sigma]$ with strings in $\{0, 1\}^\sigma$. In \mathcal{H}_{CT} , we obfuscate a circuit $\mathbb{C}_{\text{CT}}(\text{CT}')$ that computes $C_{0,\kappa}$ for all $\text{CT}' > \text{CT}$ and $C_{1,\kappa}$ for all $\text{CT}' \leq \text{CT}$; the circuit is padded to size ℓ (as in Figure 1).

We first note that \mathbb{C}_0 computes the same function as $C_{0,\kappa}$ and that \mathbb{C}_{2^σ} computes the same function as $C_{1,\kappa}$, and thus by the IO security,

$$\begin{aligned} \left| \Pr[\mathcal{D}(\tilde{C}_0) = 1] - \Pr[\mathcal{D}(\mathcal{H}_0) = 1] \right| &\leq 2^{-\tilde{\lambda}^\epsilon} , \\ \left| \Pr[\mathcal{D}(\mathcal{H}_{2^\sigma}) = 1] - \Pr[\mathcal{D}(\tilde{C}_1) = 1] \right| &\leq 2^{-\tilde{\lambda}^\epsilon} . \end{aligned}$$

We show that for any $\text{CT} \in [2^\sigma]$,

$$\left| \Pr[\mathcal{D}(\mathcal{H}_{\text{CT}-1}) = 1] - \Pr[\mathcal{D}(\mathcal{H}_{\text{CT}}) = 1] \right| \leq O(2^{-\tilde{\lambda}^\epsilon}) .$$

This follows a standard puncturing argument with respect to the point CT , consisting of:

- puncturing PRF_κ at CT , and hardwiring $C_{0,\kappa}(\text{CT}) = \text{Eval}(\text{CT}, C_0; \text{PRF}_\kappa(\text{CT}))$, which relies on IO security,
- replacing $\text{PRF}_\kappa(\text{CT})$ with true randomness, which relies on pseudorandomness at punctured points,
- replacing $\text{Eval}(\text{CT}, C_0)$ with $\text{Eval}(\text{CT}, C_1)$, which relies on function hiding.
- reversing the above steps.

Each of the steps induces a loss of $2^{-\tilde{\lambda}^\epsilon} = 2^{-\omega(\sigma(\lambda) + \log \lambda)}$ in the indistinguishability gap.

This completes the security analysis and the proof that the constructed $i\mathcal{O}$ is a secure indistinguishability obfuscator. □

Remark 3.6 (SFE in the CRS model). The above construction can be naturally extended to rely also on non-interactive SFE schemes in the CRS model (rather than the plain model). Indeed, the CRS can be generated by the (honest) obfuscator.

3.3 Instantiating the Scheme

As discussed in Section 2.1.1, we can instantiate the SFE based on the DDH assumption and an information-theoretic variant of Yao. This will result in IO for NC^1 , and can then be bootstrapped assuming LWE [GGH⁺13b]. Alternatively, we can use a strong computational variant of Yao based on sub-exponential one-way functions, resulting in IO for all poly-size circuits. We note that the work of Komargodsky et al. [KMN⁺14], shows that (approximate) sub-exponential IO (which we assume anyhow) in conjunction with the assumption that $\text{Ntime}(2^{n^\epsilon}) \neq \text{coRTIME}(2^{n^\epsilon})$, implies subexponential one-way functions.

We can thus state the following theorem

Theorem 3.7. *Assuming sub-exponentially secure approximate IO for P/poly , DDH, and LWE (or $\text{Ntime}(2^{n^\epsilon}) \neq \text{coRTIME}(2^{n^\epsilon})$), there exists (almost) perfectly correct IO for P/poly .*

Alternative instantiations of the above under more computational assumptions [PVW08] can be obtained when extending the scheme to SFE in the CRS model.

4 Correcting Errors in Functional Encryption

In this section, we define approximate FE and show how to transform any approximate FE to (almost) perfectly correct FE, based on DSFE. For the sake of concreteness, we focus on the public-key setting. We also focus on selective-security, which can be generically boosted to adaptive security [ABSV14].

4.1 Approximate and Exact FE

We recall the definition of public-key functional encryption (FE) with selective indistinguishability-based security [BSW12, O’N10], and extend the definition to the case of approximate correctness.

A public-key functional encryption (FE) scheme FE , for a function class $\mathcal{F} = \{\mathcal{F}_\lambda\}$ (represented by boolean circuits) and message space $\mathcal{M} = \{\{0, 1\}^{n(\lambda)} : \lambda \in \mathbb{N}\}$, consists of four PPT algorithms (FE.Setup, FE.Gen, FE.Enc, FE.Dec) with the following syntax:

- $\text{FE.Setup}(1^\lambda)$: Takes as input a security parameter λ in unary and outputs a (master) public key and a secret key (MPK, MSK).
- $\text{FE.Gen}(\text{MSK}, f)$: Takes as input a secret key MSK, a function $f \in \mathcal{F}_\lambda$ and outputs a functional key FSK_f .
- $\text{FE.Enc}(\text{MPK}, m)$: Takes as input a public key MPK, a message $m \in \{0, 1\}^{n(\lambda)}$ and outputs an encryption of m .
- $\text{FE.Dec}(\text{FSK}_f, \text{FCT})$: Takes as input a functional key FSK_f , a ciphertext FCT and outputs \hat{m} .

We next recall the required security properties as well the common (almost) perfect correctness requirement.

Definition 4.1 (Selectively-secure public-key FE). *A tuple of PPT algorithms $\text{FE} = (\text{FE.Setup}, \text{FE.Gen}, \text{FE.Enc}, \text{FE.Dec})$ is a selectively-secure public-key functional encryption scheme, for function class $\mathcal{F} = \{\mathcal{F}_\lambda\}$, and message space $\mathcal{M} = \{\{0, 1\}^{n(\lambda)} : \lambda \in \mathbb{N}\}$, if it satisfies:*

1. **(Almost) Perfect Correctness:** for every $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}^{n(\lambda)}$, and function $f \in \mathcal{F}_\lambda$,

$$\Pr \left[f(m) \leftarrow \text{FE.Dec}(\text{FSK}_f, \text{FCT}) \mid \begin{array}{l} (\text{MPK}, \text{MSK}) \leftarrow \text{FE.Setup}(1^\lambda) \\ \text{FSK}_f \leftarrow \text{FE.Gen}(\text{MSK}, f) \\ \text{FCT} \leftarrow \text{FE.Enc}(\text{MPK}, m) \end{array} \right] \geq 1 - \mu(\lambda),$$

for some negligible $\mu(\cdot)$.

2. **Selective-security:** for any polysize adversary \mathcal{A} , there exists a negligible function $\mu(\lambda)$ such that for any $\lambda \in \mathbb{N}$, it holds that

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

where for each $b \in \{0, 1\}$ and $\lambda \in \mathbb{N}$ the experiment $\text{Expt}_{\mathcal{A}}^{\text{FE}}(1^\lambda, b)$, modeled as a game between the challenger and the adversary \mathcal{A} , is defined as follows:

- (a) The adversary submits the challenge message-pair $m_0, m_1 \in \{0, 1\}^{n(\lambda)}$ to the challenger.
- (b) The challenger executes $\text{FE.Setup}(1^\lambda)$ to obtain (MPK, MSK) . It then executes $\text{FE.Enc}(\text{MPK}, m_b)$ to obtain FCT . The challenger sends (MPK, FCT) to the adversary.
- (c) The adversary submits function queries to the challenger. For any submitted function query $f \in \mathcal{F}_\lambda$, if $f(m_0) = f(m_1)$, the challenger generates and sends $\text{FSK}_f \leftarrow \text{FE.Gen}(\text{MSK}, f)$. In any other case, the challenger aborts.
- (d) The output of the experiment is the output of \mathcal{A} .

We further say that FE is δ -secure, for some concrete negligible function $\delta(\cdot)$, if for all polysize adversaries the above indistinguishability gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

We now define an approximate notion of correctness that allows decryption to error with some probability over encryption of messages taken from some given distribution.

Definition 4.2 ((α, \mathcal{X}) -correct FE). For $\alpha(\lambda) \in [0, 1]$ and an ensemble of samplers $\mathcal{X} = \{\mathcal{X}_\lambda\}$ with support $\mathcal{M} = \{\{0, 1\}^{n(\lambda)} : \lambda \in \mathbb{N}\}$, we say that FE is (α, \mathcal{X}) -correct if instead of (almost) perfect correctness, it satisfies the following relaxed requirement:

1. **Approximate Correctness:** for every $\lambda \in \mathbb{N}$, and function $f \in \mathcal{F}_\lambda$,

$$\Pr \left[f(m) \leftarrow \text{FE.Dec}(\text{FSK}_f, \text{FCT}) \mid \begin{array}{l} (\text{MPK}, \text{MSK}) \leftarrow \text{FE.Setup}(1^\lambda) \\ \text{FSK}_f \leftarrow \text{FE.Gen}(\text{MSK}, f) \\ m \leftarrow \mathcal{X}_\lambda \\ \text{FCT} \leftarrow \text{FE.Enc}(\text{MPK}, m) \end{array} \right] \geq \alpha(\lambda).$$

4.2 The Transformation

We now describe the transformation from approximately correct FE to (almost) perfectly correct FE and analyze it. The transformation is based on DSFE. We discuss instantiations in Section 4.3.

Ingredients. In the following, let λ denote a security parameter, and let $\mathcal{F} = \{\mathcal{F}_\lambda\}$ denote a function class. Consider functions $k(\lambda) \in \mathbb{N}$, and $\rho(\lambda), \eta(\lambda) \in [0, 1]$ such that $\eta = \frac{1}{2} - \sqrt{\rho} \in [\frac{1}{\lambda^{\mathcal{O}(1)}}, \frac{1}{2} - \frac{1}{\lambda^{\mathcal{O}(1)}}]$. We rely on the following primitives:

- A $(k, \sqrt{\rho})$ -secure distributed function evaluation scheme DSFE for \mathcal{C} . We shall further assume that when encrypting an input, the shares $\text{CT}_1, \dots, \text{CT}_k$ all have the same marginal distribution (i.e., $\text{CT}_i \equiv \text{CT}_j$).³
- A $(1-\rho, \mathcal{X})$ -correct (single-key, selectively-secure) functional encryption scheme $\widetilde{\text{FE}} = (\widetilde{\text{FE}}.\text{Setup}, \widetilde{\text{FE}}.\text{Gen}, \widetilde{\text{FE}}.\text{Enc}, \widetilde{\text{FE}}.\text{Dec})$ for $\overline{\mathcal{C}}$. The sampler class \mathcal{X} depends on DSFE and the class $\overline{\mathcal{F}}$ depends on DSFE, and \mathcal{F} . Both \mathcal{X} and $\overline{\mathcal{F}}$ are specified below as part of the description of the constructed (exact) scheme FE.

The exact scheme FE: The scheme FE, for function class $\mathcal{F} = \{\mathcal{F}_\lambda\}$ and message space $\mathcal{M} = \{\{0, 1\}^{n(\lambda)} : \lambda \in \mathbb{N}\}$, consists of the algorithms (FE.Setup, FE.Gen, FE.Enc, FE.Dec) defined as follows:

- FE.Setup(1^λ): let $N = \frac{\omega(n+\log \lambda)}{\eta^2}$. For $i \in [N]$, generate $(\widetilde{\text{MPK}}_i, \widetilde{\text{MSK}}_i) \leftarrow \widetilde{\text{FE}}.\text{Setup}(1^\lambda)$. The public key MPK and secret key MSK are accordingly set to be all of the sampled public keys $\{\widetilde{\text{MPK}}_i\}_{i \in [N]}$ and secret keys $\{\widetilde{\text{MSK}}_i\}_{i \in [N]}$.
- FE.Gen(MSK, f): For $i \in [N]$, sample $\text{SCT}_i \leftarrow \text{Sym}.\text{Enc}(\text{SK}, 0^{\ell \times k})$, where $\ell = \ell(\lambda)$ is a polynomial specified in the security analysis, and $\text{SK} \leftarrow \{0, 1\}^\lambda$. Consider the augmented f -evaluation function f_{SCT_i} as defined in Figure 2. Generate $\widetilde{\text{FSK}}_{\text{SCT}_i} \leftarrow \widetilde{\text{FE}}.\text{Gen}(\widetilde{\text{MSK}}_i, f_{\text{SCT}_i})$. The functional key FSK $_f$ will consists of the functional keys $\{\widetilde{\text{FSK}}_{\text{SCT}_i}\}_{i \in [N]}$.
- FE.Enc(MPK, m): For $i \in [N]$,
 1. Compute $(\text{CT}_{i,1}, \dots, \text{CT}_{i,k}, \text{R}_i) \leftarrow \text{DSFE}.\text{Enc}(m)$,
 2. For $j \in [k]$
 - let $\widetilde{m}_{i,j} = (\text{norm}, \text{CT}_{i,j}, \perp, \perp)$
 - generate $\widetilde{\text{FCT}}_{i,j} \leftarrow \widetilde{\text{FE}}.\text{Enc}(\widetilde{\text{MPK}}_i, \widetilde{m}_{i,j})$.

Output $\text{FCT} = \{\widetilde{\text{FCT}}_{i,1}, \dots, \widetilde{\text{FCT}}_{i,k}, \text{R}_i\}_{i \in [N]}$.

- FE.Dec(FSK $_f$, FCT):
 1. Parse $\text{FSK}_f = \{\widetilde{\text{FSK}}_{\text{SCT}_i}\}_{i \in [N]}$ and $\text{FCT} = \{\widetilde{\text{FCT}}_{i,1}, \dots, \widetilde{\text{FCT}}_{i,k}, \text{R}_i\}_{i \in [N]}$.
 2. For $i \in [N]$,
 - for $j \in [k]$, compute $\widehat{\text{CT}}_{i,j} \leftarrow \widetilde{\text{FE}}.\text{Dec}(\widetilde{\text{FSK}}_{\text{SCT}_i}, \widetilde{\text{FCT}}_{i,j})$.
 - compute $y_i = \text{DSFE}.\text{Dec}(\text{R}_i, \widehat{\text{CT}}_{i,1}, \dots, \widehat{\text{CT}}_{i,k})$.
 3. Output $y = \text{majority}(y_1, \dots, y_N)$.

The ensemble of samplers \mathcal{X} consists of samplers \mathcal{X}^0 that sample FE plaintexts of the form $\widetilde{m} = (\text{norm}, \text{CT}, \perp, \perp)$ where CT is the first of k ciphertext components sampled from $\text{DSFE}.\text{Enc}(0^n)$, i.e. it is a share of a zero-encryption in the underlying DSFE scheme. The class $\overline{\mathcal{F}}$ consists of circuits f_{SCT} as in Figure 2.

³This is just to simplify the construction and is satisfied the instantiation discussed in Section 2. In Remark 4.4, we explain how this assumption can be removed (at the cost of complicating the construction).

f_{SCT}

Hardwired: a symmetric key ciphertext SCT.

Input $\tilde{m} = (b, \text{CT}, \text{SK}, j)$:

- a flag bit b ,
 - a DSFE ciphertext CT_j ,
 - a symmetric encryption key SK.
 - index $j \in [k]$.
1. If $b = \text{norm}$ (normal mode of operation, ignoring inputs SK, j),
 - compute $\widehat{\text{CT}} = \text{Eval}(\text{CT}, f)$.
 2. If $b = \text{trap}$ (trapdoor mode of operation, ignoring inputs CT, K),
 - compute $(\widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k) = \text{Sym.Dec}(\text{SK}, \text{SCT})$,
 - let $\widehat{\text{CT}} := \widehat{\text{CT}}_j$.
 3. Output $\widehat{\text{CT}}$.

Figure 2: The augmented f -evaluation circuit.

Theorem 4.3. FE is an (almost) perfectly correct functional encryption scheme.

We now turn to the proof.

Proof. We first prove that the new obfuscator is (almost) exactly correct, and then prove that it is secure.

Almost exact correctness. For any $\lambda, n = n(\lambda)$, message $m \in \{0, 1\}^n$, let us denote \mathcal{X}^m a sampler for FE plaintexts of the form $\tilde{m} = (\text{norm}, \text{CT}, \perp, \perp)$ that is defined just like \mathcal{X}^0 except that CT is a share of an encryption of m sampled from $\text{DSFE.Enc}(m)$ in the underlying DSFE scheme, rather than a share of an encryption of 0^n .

Then, by the input-hiding guarantee of SFE, and the approximate correctness of $\widetilde{\text{FE}}$, we claim that, for any function $f \in \mathcal{F}$ and corresponding f_{SCT} , decryption in $\widetilde{\text{FE}}$ is correct on encryptions of an arbitrary $m \in \{0, 1\}^n$ as on encryptions of 0^n . That is, there exists a negligible $\mu(\lambda)$ such that

$$\begin{aligned} & \Pr \left[\widetilde{\text{FE}}.\text{Dec}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}) = f_{\text{SCT}}(\tilde{m}) \mid \tilde{m} \leftarrow \mathcal{X}^m \right] \geq \\ & \Pr \left[\widetilde{\text{FE}}.\text{Dec}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}) = f_{\text{SCT}}(\tilde{m}) \mid \tilde{m} \leftarrow \mathcal{X}^0 \right] - \mu \geq \\ & 1 - \rho - \mu, \end{aligned}$$

where in both $(\widetilde{\text{MPK}}, \widetilde{\text{MSK}}) \leftarrow \widetilde{\text{FE}}.\text{Setup}(1^\lambda), \widetilde{\text{FSK}}_{f_{\text{SCT}}} \leftarrow \widetilde{\text{FE}}.\text{Gen}(\widetilde{\text{MSK}}, f_{\text{SCT}}), \widetilde{\text{FCT}} \leftarrow \widetilde{\text{FE}}.\text{Enc}(\widetilde{\text{MPK}}, \tilde{m})$, all defined above in the construction of the exact scheme, and $\tilde{m} = (\text{norm}, \text{CT}, \perp, \perp)$.

We now consider alternative samplers \mathcal{X}_j^m that sample \tilde{m}_j just as in the canonical \mathcal{X}^m , except that CT is sampled as the the j th share of a DSFE encryption of m (rather than the first). Note that by our assumption that shares $\text{CT}_1, \dots, \text{CT}_k \leftarrow \text{DSFE.Enc}(m)$ have the same marginal distribution,

the samplers $\mathcal{X}^m, \mathcal{X}_1^m, \dots, \mathcal{X}_k^m$ all sample from the same distribution. In particular, they satisfy the above statement regarding the probability of correct decryption, satisfied by \mathcal{X}^m .

We shall denote by $\mathcal{X}_j^m | \text{CT}_j$ the corresponding sampler conditioned on CT being some CT_j . We now consider the joint sampler $(\tilde{m}_1, \dots, \tilde{m}_k) \leftarrow \mathcal{X}_{[k]}^m$ where first shares $(\text{CT}_1, \dots, \text{CT}_k)$ are sample from $\text{DSFE.Enc}(m)$, and then each \tilde{m}_j is sampled from $\mathcal{X}_j | \text{CT}_j$. Note that this sampler corresponds to the way that encryption is done in our actual scheme FE defined above (in any single instance out of N).

Noting that the marginal distribution of each \tilde{m}_j sampled accordingly to $\mathcal{X}_{[k]}^m$ is the same as \mathcal{X}_j^m . It follows that the expected number of successful decryptions for a sample from $\mathcal{X}_{[k]}^m$ can be lower bounded

$$\begin{aligned} & \mathbb{E} \left[\left| \left\{ j \mid \widetilde{\text{FE.Dec}}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}_j) = f_{\text{SCT}}(\tilde{m}_j) \right\} \right| \mid (\tilde{m}_1, \dots, \tilde{m}_k) \leftarrow \mathcal{X}_{[k]}^m \right] = \\ & k \cdot \Pr \left[\widetilde{\text{FE.Dec}}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}_j) = f_{\text{SCT}}(\tilde{m}_j) \mid \tilde{m}_j \leftarrow \mathcal{X}^m \right] \geq \\ & k \cdot (1 - \rho - \mu) , \end{aligned}$$

where $(\widetilde{\text{MPK}}, \widetilde{\text{MSK}}) \leftarrow \widetilde{\text{FE.Setup}}(1^\lambda)$, $\widetilde{\text{FSK}}_{f_{\text{SCT}}} \leftarrow \widetilde{\text{FE.Gen}}(\widetilde{\text{MSK}}, f_{\text{SCT}})$, $\widetilde{\text{FCT}}_j \leftarrow \widetilde{\text{FE.Enc}}(\widetilde{\text{MPK}}, \tilde{m}_j)$.

It follows by averaging that with probability at least $1 - \sqrt{\rho} - \frac{2\mu}{\sqrt{\rho}}$ the number of successful decryptions as defined above is larger than $k \cdot (1 - \sqrt{\rho})$. In particular, (for large enough λ) the fraction of faults is below the threshold $\sqrt{\rho}$ allowing to reconstruct $f_{\text{SCT}}(\tilde{m})$.

Going to our actual encryption scheme FE, we now claim that each of the N independent decryption attempts is correct with probability noticeably larger than half.

Concretely,

$$\begin{aligned} & \Pr \left[\text{DSFE.Dec}(\text{R}, \widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k) = f(m) \mid \begin{array}{l} \text{CT}_1, \dots, \text{CT}_k, \text{R} \leftarrow \text{DSFE.Enc}(m) \\ \widehat{\text{CT}}_j = \widetilde{\text{FE.Dec}}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}_j) \end{array} \right] \geq \\ & \Pr \left[\text{DSFE.Dec}(\text{R}, \widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k) = f(m) \mid \begin{array}{l} \text{CT}_1, \dots, \text{CT}_k, \text{R} \leftarrow \text{DSFE.Enc}(m) \\ \left| \left\{ \widehat{\text{CT}}_j = f_{\text{SCT}}(\tilde{m}_j) \right\} \right| \geq \sqrt{\rho} \cdot k \end{array} \right] . \\ & \Pr \left[\left| \left\{ \widetilde{\text{FE.Dec}}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}_j) = f_{\text{SCT}}(\tilde{m}_j) \right\} \right| \geq \sqrt{\rho} \cdot k \mid \text{CT}_1, \dots, \text{CT}_k, \text{R} \leftarrow \text{DSFE.Enc}(m) \right] = \\ & \Pr \left[\text{DSFE.Dec}(\text{R}, \widehat{\text{CT}}_1, \dots, \widehat{\text{CT}}_k) = f(m) \mid \begin{array}{l} \text{CT}_1, \dots, \text{CT}_k, \text{R} \leftarrow \text{DSFE.Enc}(m) \\ \left| \left\{ \widehat{\text{CT}}_j = \text{DSFE.Eval}(\text{CT}_j, f) \right\} \right| \geq \sqrt{\rho} \cdot k \end{array} \right] . \\ & \Pr \left[\left| \left\{ \widetilde{\text{FE.Dec}}(\widetilde{\text{FSK}}_{f_{\text{SCT}}}, \widetilde{\text{FCT}}_j) = f_{\text{SCT}}(\tilde{m}_j) \right\} \right| \geq \sqrt{\rho} \cdot k \mid \text{CT}_1, \dots, \text{CT}_k \leftarrow \mathcal{X}_{[k]}^m \right] \geq \\ & (1 - \mu') \cdot \left(1 - \sqrt{\rho} - \frac{2\mu}{\sqrt{\rho}} \right) \geq \frac{1}{2} + \eta - \lambda^{-\omega(1)} , \end{aligned}$$

where in all of the above $(\widetilde{\text{MPK}}, \widetilde{\text{MSK}}) \leftarrow \widetilde{\text{FE.Setup}}(1^\lambda)$, $\widetilde{\text{FSK}}_{f_{\text{SCT}}} \leftarrow \widetilde{\text{FE.Gen}}(\widetilde{\text{MSK}}, f_{\text{SCT}})$, $\tilde{m}_j = (\text{norm}, \text{CT}_j, \perp, \perp)$, $\widetilde{\text{FCT}}_j \leftarrow \widetilde{\text{FE.Enc}}(\widetilde{\text{MPK}}, \tilde{m}_j)$, and $\mu'(\cdot)$ is some negligible function (corresponding to the negligible decryption error of DSFE).

By a Chernoff bound, for large enough λ , and any $m \in \{0, 1\}^n$, the probability that the majority value y among all decrypted y_1, \dots, y_N is incorrect is bounded by

$$\Pr [y \neq f(m)] \leq 2^{-\Omega(N \cdot \eta^2)} \leq 2^{-n + \omega(\log \lambda)} .$$

The required correctness for all input then follows by a union bound over all messages in $\{0, 1\}^n$.

Security analysis. Recall that the scheme FE constructed above is based on N independent copies of one basic scheme. It suffices (by a standard hybrid argument) to prove the security of a single instance. We prove this in a sequence of hybrids, showing that any adversary \mathcal{A} cannot tell the case that the challenge is an encryption of m_0 from the case that the challenge is an encryption of m_1 , for the corresponding (m_0, m_1) of his choice.

\mathcal{H}_1 : this corresponds to the usual security game where the challenge is an encryption of m_0 .

\mathcal{H}_2 : here, when generating a key FSK_f for a function f , and accordingly generating an (approximate) key $\widetilde{\text{FSK}}_{f_{\text{SCT}}}$ for the function SCT, the symmetric ciphertext SCT is not an encryption of $0^{\ell \times k}$ as in the previous hybrid, but rather an encryption of the DSFE evaluation corresponding to the challenge ciphertext. Concretely, consider the generation of the challenge ciphertext FCT^*

- $\text{FE.Enc}(\text{MPK}, m_0)$:
 1. Compute $(\text{CT}_1^*, \dots, \text{CT}_k^*, \text{R}^*) \leftarrow \text{DSFE.Enc}(m_0)$,
 2. For $j \in [k]$
 - let $\tilde{m}_j^* = (\text{norm}, \text{CT}_j^*, \perp, \perp)$
 - generate $\widetilde{\text{FCT}}_j^* \leftarrow \widetilde{\text{FE.Enc}}(\widetilde{\text{MPK}}, \tilde{m}_j^*)$.

Output $\text{FCT}^* = (\widetilde{\text{FCT}}_1^*, \dots, \widetilde{\text{FCT}}_k^*, \text{R}^*)$.

Then SCT encrypts $\widehat{\text{CT}}_{f,1}^*, \dots, \widehat{\text{CT}}_{f,k}^*$, where $\widehat{\text{CT}}_{f,j}^* = \text{DSFE.Eval}(\text{CT}_j^*, f)$.

Indistinguishability from the previous hybrid follows from the semantic-security of the symmetric encryption. (Note that at this point, a corresponding symmetric secret key SK is not present, in all encryptions the symmetric-key slot is set to \perp .)

\mathcal{H}_3 : here, we change the generation of the challenge ciphertext so to invoke the trapdoor mode rather than the normal mode. Concretely, for each $j \in [k]$, we generate $\tilde{m}_j^* = (\text{trap}, \perp, \perp, \text{SK}, j)$, where SK is the symmetric key corresponding SCT.

Indistinguishability from the previous hybrid follows from the security of the underlying scheme $\widetilde{\text{FE}}$. Indeed, at this point, for every function f for which a key $\widetilde{\text{FSK}}_{f_{\text{SCT}}}$ was generated,

$$f_{\text{SCT}}(\text{trap}, \perp, \perp, \text{SK}, j) = f_{\text{SCT}}(\text{norm}, \text{CT}_j^*, \perp, \perp) = \widehat{\text{CT}}_{f,j}^*$$

\mathcal{H}_4 : here, we change how the evaluations $\widehat{\text{CT}}_{f,j}^*$ are generated. Recall that in the previous hybrid $\widehat{\text{CT}}_{f,j}^* = \text{DSFE.Eval}(\text{CT}_j^*, f)$, where CT_j^* was generated as part of $(\text{CT}_1^*, \dots, \text{CT}_k^*, \text{R}^*) \leftarrow \text{DSFE.Enc}(m_0)$. Now, instead of encrypting m_0 in the latter we encrypt m_1 .

Indistinguishability now follows from the residual input privacy of the underlying DSFE, since $f(m_0) = f(m_1)$. (Recall, that this is guaranteed also in the presence of R^* , provided that $\text{CT}_1^*, \dots, \text{CT}_k^*$ are absent from the adversary's view, which is indeed the case in this hybrid.)

\mathcal{H}_5 - \mathcal{H}_8 : symmetrically follow the above hybrids in reverse order, until the usual security game where m_1 is encrypted in the challenge.

This completes the security analysis and proof that the constructed FE is secure. \square

Remark 4.4 (removing the assumption regarding equally-distributed shares). In the above construction we have assumed that DSFE shares CT_1, \dots, CT_k have the same marginal distribution. Roughly, to remove this assumption, we could have an instance of an approximate FE scheme \widetilde{FE}_i for each i with respect to the corresponding distribution on CT_i (whereas in the construction above we relied on one instance of an approximate FE defined with respect to the marginal distribution which was the same for all shares).

4.3 Instantiating the Scheme

As discussed in Section 2.1.1, we can instantiate the DSFE based an information-theoretic variant of BGW for NC^1 , resulting in an FE scheme for NC^1 . The scheme can then be generically bootstrapped to P/poly assuming weak PRFs in NC^1 [ABSV14].

We can thus state the following theorem

Theorem 4.5. *Assuming approximate FE for P/poly and weak PRFs in NC^1 , there exists (almost) perfectly correct FE.*

5 An Alternative Transformation for IO based on FE

Recall that the transformation from (subexponential) approximate IO to exact IO described in Section 3.2 required SFE with function hiding against malicious receivers, and was instantiated based on DDH and LWE (or $Ntime(2^{n^\epsilon}) \neq coRTIME(2^{n^\epsilon})$). In this section, we note that an alternative transformation based on weaker assumptions can be obtained when combining the above FE transformation with known results from the literature.

At high-level the idea consists of the following steps.

1. Start with a (subexponentially-secure) approximate IO and implement directly (subexponentially-secure) approximate FE with compact ciphertexts imitating the construction in the exact IO setting [GGH⁺13b].
2. Apply the above transformation from approximate FE to obtain exact FE with compact ciphertexts.
3. Apply a transformation from exact FE to (exact) IO [AJ15, BV15].

Fulfilling this high-level plan requires some care though. The transformation of Garg et al. [GGH⁺13b] from IO to FE naturally extends to the approximate setting, but relies on additional assumptions: (exact) public-key encryption and (exact, or rather complete) NIZKs. While these primitives are known based on exact IO [SW14, BP15], they do not work in the approximate setting. Nevertheless, we show how these constructions can be extended to imply the exact versions of the primitives (from approximate IO). A second issue that should be addressed is how the approximate FE to exact FE transformation affects the complexity of FE encryption. Indeed, the transformations of [AJ15, BV15] require certain compactness properties. We observe that the transformation, when instantiated with the BGW-based DSFE, satisfies the required compactness.

We now sketch the above steps.

Approximate FE from approximate IO. The starting point for this step is the Garg et al. [GGH⁺13b]. Roughly speaking to obtain FE from IO and PKE, and NIZKs, the transformation

works as follows. Encryptions consist of double encryptions under a plain (exact) public-key encryption scheme, together with an appropriate NIZK regarding their well-formedness. A function key for a function f , consists of an obfuscation that verifies the proof, decrypts (one of) the plaintext, applies f to the decrypted message, and outputs the result.

It follows rather directly that if we replace exact IO in this transformation with approximate IO (while still using exact PKE and NIZKs) the resulting scheme would be an approximate FE. Concretely to get approximate FE for a message sampler \mathcal{X} , we will start with approximate IO for an input sampler \mathcal{X}' that samples FE encryptions (with respect to the [GGH⁺13b] scheme) of random messages taken from \mathcal{X} . We next explain how to obtain exact PKE and NIZKs from approximate IO.

To obtain exact PKE, we can start with the PKE of Sahai and Waters [SW14] based on exact IO and one-way functions. Replacing exact IO with approximate IO in their transformation results in approximate PKE that is correct over random encryptions of random messages. This can be corrected using standard techniques [DNR04]. To obtain exact NIZKs, we examine the construction of Bitansky and Paneth [BP15] based on exact IO and one-way functions. We observe that in that construction replacing exact IO with approximate IO results in a NIZK that has negligible soundness and a bounded completeness error. The completeness error can then be arbitrarily decreased using parallel repetition. One caveat of the latter transformation is that it only works if the approximate IO errors on some small polynomial fraction of inputs, and not say a constant. We note though that in the de-idealized constructions of obfuscation [CKP15, PS15, MMN15] the error rate can be an arbitrary small polynomial.

Another thing to note is that the resulting constructions are not exact, but rather have an arbitrarily small error. Using standard techniques [Nao91, DNR04], we can make sure that the errors are shifted to the sampling of public parameters; that is with overwhelming probability over the choice of public parameters, all operations will be exactly correct. This suffices for our purposes. Finally, we note that (subexponential) one-way functions are known to follow from (subexponential) approximate IO and the assumption that $Ntime(2^{n^\epsilon}) \neq coRTIME(2^{n^\epsilon})$.

Compactness. The above results in an approximate FE scheme where the complexity of encryption is independent of the circuit and output size of the corresponding functions. To fulfill our approach we need to make sure that applying our transformation to exact FE still preserves certain compactness properties. Concretely, we note that our transformation inherits its compactness from the underlying DSFE scheme. As discussed in 4.3, using the BGW-based DSFE, incurs a $2^{O(d)}$ overhead in the complexity of encryption, where d is the maximal depth of any circuit in the class, but is otherwise as efficient. As shown in [BV15], this is still sufficient for constructing (exact) IO from (exact) FE.

We can thus state the following theorem

Theorem 5.1. *Assuming sub-exponentially-secure approximate IO for P/poly , weak PRFs in NC^1 , and $Ntime(2^{n^\epsilon}) \neq coRTIME(2^{n^\epsilon})$, there exists (almost) perfectly correct IO P/poly .*

References

- [AB15] Benny Applebaum and Zvika Brakerski. Obfuscating circuits via composite-order graded encoding. In Dodis and Nielsen [DN15], pages 528–556.

- [ABSV14] Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. The trojan method in functional encryption: From selective to adaptive security, generically. *IACR Cryptology ePrint Archive*, 2014:917, 2014.
- [AFK89] Martín Abadi, Joan Feigenbaum, and Joe Kilian. On hiding information from an oracle. *J. Comput. Syst. Sci.*, 39(1):21–50, 1989.
- [AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.
- [AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In *Crypto*, 2015.
- [AL11] Gilad Asharov and Yehuda Lindell. A full proof of the BGW protocol for perfectly-secure multiparty computation. *IACR Cryptology ePrint Archive*, 2011:136, 2011.
- [BCC⁺14] Nir Bitansky, Ran Canetti, Henry Cohn, Shafi Goldwasser, Yael Tauman Kalai, Omer Paneth, and Alon Rosen. The impossibility of obfuscation with auxiliary input or a universal simulator. In *CRYPTO*, pages 71–89, 2014.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6, 2012.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudo-random functions. In Hugo Krawczyk, editor, *PKC*, volume 8383 of *Lecture Notes in Computer Science*, pages 501–519. Springer, 2014.
- [BGJ⁺15] Nir Bitansky, Shafi Goldwasser, Abhishek Jain, Omer Paneth, Vinod Vaikuntanathan, and Brent Waters. Time-lock puzzles from randomized encodings. *IACR Cryptology ePrint Archive*, 2015:514, 2015.
- [BGK⁺14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT*, volume 8441 of *Lecture Notes in Computer Science*, pages 221–238. Springer, 2014.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*, pages 1–10, 1988.
- [BP12] Nir Bitansky and Omer Paneth. From the impossibility of obfuscation to a new non-black-box simulation technique. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 223–232, 2012.
- [BP15] Nir Bitansky and Omer Paneth. Zaps and non-interactive witness indistinguishability from indistinguishability obfuscation. In Dodis and Nielsen [DN15], pages 401–427.

- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *CCS '93, Proceedings of the 1st ACM Conference on Computer and Communications Security, Fairfax, Virginia, USA, November 3-5, 1993.*, pages 62–73. ACM, 1993.
- [BR14a] Zvika Brakerski and Guy N. Rothblum. Black-box obfuscation for d-cnfs. In Moni Naor, editor, *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*, pages 235–250. ACM, 2014.
- [BR14b] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *Theory of Cryptography - 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings*, volume 8349 of *Lecture Notes in Computer Science*, pages 1–25. Springer, 2014.
- [BSW12] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: a new vision for public-key cryptography. *Commun. ACM*, 55(11):56–64, 2012.
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In *FOCS*, 2015.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT (2)*, volume 8270 of *Lecture Notes in Computer Science*, pages 280–300. Springer, 2013.
- [BWZ14] Dan Boneh, David J. Wu, and Joe Zimmerman. Immunizing multilinear maps against zeroizing attacks. *IACR Cryptology ePrint Archive*, 2014:930, 2014.
- [CGH⁺15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrede Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO*, Lecture Notes in Computer Science. Springer, 2015. To appear.
- [CHL⁺15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT, Part I*, volume 9056 of *Lecture Notes in Computer Science*, pages 3–12. Springer, 2015.
- [CKP15] Ran Canetti, Yael Tauman Kalai, and Omer Paneth. On obfuscation with random oracles. In Dodis and Nielsen [DN15], pages 456–467.
- [CKV10] Kai-Min Chung, Yael Tauman Kalai, and Salil P. Vadhan. Improved delegation of computation using fully homomorphic encryption. In Tal Rabin, editor, *Advances in Cryptology - CRYPTO 2010, 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings*, volume 6223 of *Lecture Notes in Computer Science*, pages 483–501. Springer, 2010.

- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 476–493. Springer, 2013.
- [CLT15] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO*, *Lecture Notes in Computer Science*. Springer, 2015. To appear.
- [CLTV15] Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In *TCC*, 2015.
- [DN15] Yevgeniy Dodis and Jesper Buus Nielsen, editors. *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*. Springer, 2015.
- [DNR04] Cynthia Dwork, Moni Naor, and Omer Reingold. Immunizing encryption schemes from decryption errors. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, volume 3027 of *Lecture Notes in Computer Science*, pages 342–360. Springer, 2004.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [GGH12] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices and applications. *IACR Cryptology ePrint Archive*, 2012:610, 2012.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, Mariana Raikova, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 498–527. Springer, 2015.
- [GGHZ14] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Fully secure functional encryption without obfuscation. *IACR Cryptology ePrint Archive*, 2014:666, 2014.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GK05] Shafi Goldwasser and Yael Tauman Kalai. On the impossibility of obfuscation with auxiliary input. In *FOCS*, pages 553–562. IEEE Computer Society, 2005.

- [GLSW14] Craig Gentry, Allison B. Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. *IACR Cryptology ePrint Archive*, 2014:309, 2014.
- [GLW14] Craig Gentry, Allison B. Lewko, and Brent Waters. Witness encryption from instance independent assumptions. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 426–443. Springer, 2014.
- [GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In *CRYPTO*, pages 162–179, August 2012.
- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudo-random generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HJ15] Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. *IACR Cryptology ePrint Archive*, 2015:301, 2015.
- [IK02] Yuval Ishai and Eyal Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In *In Proc. 29th ICALP*, pages 244–256, 2002.
- [KMN⁺14] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. One-way functions and (im)perfect obfuscation. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 374–383. IEEE Computer Society, 2014.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *CCS*, pages 669–684. ACM, 2013.
- [LS14] Hyung Tae Lee and Jae Hong Seo. Security analysis of multilinear maps over the integers. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 224–240. Springer, 2014.
- [MMN15] Mohammad Mahmood, Ameer Mohammed, and Soheil Nematihaji. More on impossibility of virtual black-box obfuscation in idealized models. *Cryptology ePrint Archive*, Report 2015/632, 2015. <http://eprint.iacr.org/>.
- [Nao91] Moni Naor. Bit commitment using pseudorandomness. *J. Cryptology*, 4(2):151–158, 1991.
- [NP01] Moni Naor and Benny Pinkas. Efficient oblivious transfer protocols. In *SODA*, pages 448–457, 2001.

- [O’N10] Adam O’Neill. Definitional issues in functional encryption. Cryptology ePrint Archive, Report 2010/556, 2010.
- [PS15] Rafael Pass and Abhi Shelat. Impossibility of VBB obfuscation with ideal constant-degree graded encodings. *IACR Cryptology ePrint Archive*, 2015:383, 2015.
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 500–517. Springer, 2014.
- [PVW08] Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In *CRYPTO*, pages 554–571, 2008.
- [Sho97] Victor Shoup. Lower bounds for discrete logarithms and related problems. In Walter Fumy, editor, *Advances in Cryptology - EUROCRYPT ’97, International Conference on the Theory and Application of Cryptographic Techniques, Konstanz, Germany, May 11-15, 1997, Proceeding*, volume 1233 of *Lecture Notes in Computer Science*, pages 256–266. Springer, 1997.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *STOC*, pages 475–484. ACM, 2014.
- [Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th Annual Symposium on Foundations of Computer Science, Toronto, Canada, 27-29 October 1986*, pages 162–167. IEEE Computer Society, 1986.
- [Zim15] Joe Zimmerman. How to obfuscate programs directly. In *Eurocrypt*, 2015.