Indistinguishability Obfuscation from Functional Encryption

NIR BITANSKY, Tel Aviv University
VINOD VAIKUNTANATHAN, MIT

Indistinguishability obfuscation (IO) is a tremendous notion, powerful enough to give rise to almost any known cryptographic object. Prior candidate IO constructions were based on specific assumptions on algebraic objects called multi-linear graded encodings.

We present a generic construction of indistinguishability obfuscation from public-key functional encryption with succinct encryption circuits and subexponential security. This shows the equivalence of indistinguishability obfuscation and public-key functional encryption, a primitive that has previously seemed to be much weaker, lacking the power and the staggering range of applications of indistinguishability obfuscation.

Our main construction can be based on functional encryption schemes that support a *single functional key*, and where the encryption circuit grows sub-linearly in the circuit-size of the function. We further show that sublinear succinctness in circuit-size for single-key schemes can be traded with sublinear succinctness in the number of keys (also known as the *collusion-size*) for multi-key schemes. We also show that, under the Learning with Errors assumption, our techniques imply that any indistinguishability obfuscator can be converted into one where the size of obfuscated circuits is twice that of the original circuit plus an additive overhead that is polynomial in its depth, input length, and the security parameter.

CCS Concepts: • **Security and privacy** → **Mathematical foundations of cryptography**;

Additional Key Words and Phrases: Obfuscation, functional encryption

ACM Reference format:

Nir Bitansky and Vinod Vaikuntanathan. 2018. Indistinguishability Obfuscation from Functional Encryption. *J. ACM* 65, 6, Article 39 (November 2018), 37 pages.
<https://doi.org/10.1145/3234511>

1 INTRODUCTION

Program obfuscation, aiming to turn programs into “unintelligible” ones while preserving functionality, has been a holy grail in cryptography for over a decade. While heuristic methods of obfuscation are widely used in practice, our theoretical understanding of obfuscation is still in

An extended abstract of this article appears in the proceedings of FOCS 2015.

Part of this research was done while at MIT and supported by NSF Grants No. CNS-1350619 and No. CNS-1414119 and by the NEC Corporation. Vinod Vaikuntanathan was supported in part by NSF Grants No. CNS-1350619 and No. CNS-1414119 Alfred P. Sloan Research Fellowship Microsoft Faculty Fellowship by the NEC Corporation and a Steven and Renee Finn Career Development Chair from MIT. This work was also sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and the U.S. Army Research Office under contracts W911NF-15-C-0226 and FA8750-11-2-0225.

Authors’ addresses: N. Bitansky, Tel Aviv University, P.O. Box 39040, Tel Aviv 6997801, Israel; email: nirbitan@tau.ac.il; V. Vaikuntanathan, MIT, 32 Vassar Street, G-696, Cambridge MA 02139, USA; email: vinodv@csail.mit.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

0004-5411/2018/11-ART39 \$15.00

<https://doi.org/10.1145/3234511>

its early stages. Rather unfortunately, the most natural and intuitively appealing notion of obfuscation, namely, *virtual-black-box* (VBB) obfuscation [13, 51], was shown to have strong limitations [13, 14, 45, 51]. Furthermore, except for very restricted function classes (see, for example, References [27, 31, 53, 54, 69, 76]), no candidate construction with any form of meaningful security was known for a long time.

This changed dramatically when Garg, Gentry, Halevi, Raykova, Sahai, and Waters [36] demonstrated a candidate obfuscation algorithm for all circuits and conjectured that it satisfies an apparently weak notion of *indistinguishability obfuscation* (IO) [13, 47], requiring only that the obfuscations of any two circuits of the same size and the same functionality (namely, the same truth table) are computationally indistinguishable. Since then, a sequence of works, pioneered by Sahai and Waters [74], have demonstrated that IO is not such a weak notion after all, leading to a plethora of applications and even resolving long-standing open problems. The number of cryptographic primitives that we do not know how to construct from IO is small and dwindling fast.¹

The tremendous power of IO also begets its reliance on strong and untested computational assumptions. Despite significant progress [41, 71], known IO constructions prior to this work [8, 12, 25, 36, 41, 71, 78] were based on the hardness of little-studied problems on multi-linear maps [35]. Thus, an outstanding foundational question in cryptography is

Can we base indistinguishability obfuscation on solid cryptographic foundations?

1.1 This Work

In this work, aiming to make progress in the above direction, we show how to construct indistinguishability obfuscation from an apparently weaker primitive: *public-key functional encryption*. In a functional encryption scheme [18, 20, 70, 73], it is possible to produce functional keys FSK_f for functions f (represented as circuits throughout this article). Given an encryption of an input x , computed using a public key PK and the functional key FSK_f , anyone can compute $f(x)$, but nothing more about x itself.

In the past few years, functional encryption (FE) schemes with different efficiency and security features were constructed from various computational assumptions. A central measure of interest (in general, and in the specific context of this work) is the size of ciphertexts, or more generally, the encryption time. Here, the ideal requirement is that the time to encrypt depends only on the underlying plaintext x , but this requirement may be relaxed in several meaningful ways, such as allowing dependence on the size of the circuits computing the corresponding functions, just the size of their output, or the number of generated functional keys.

Functional encryption, on the face of it, seems much less powerful than IO, and sure enough, it has not had nearly as many applications. Seemingly, IO derives its power from the fact that it allows *anyone* to compute meaningfully with a hidden object (say, a circuit) with no additional help. In contrast, FE does allow us to encrypt circuits,² but to evaluate the circuit on an input, one needs a functional key associated with the input! Not surprisingly, the power of FE seems to be limited to achieving a notion of “obfuscation on a leash” or “token-based obfuscation,” as defined by Goldwasser, Kalai, Popa, Vaikuntanathan, and Zeldovich [46].

¹Strictly speaking, we need the assumption that IO exists, plus a very mild (and minimal) complexity-theoretic assumption that $\text{NP} \not\subseteq \text{ioBPP}$ [59].

²Given FE for a sufficiently expressive class, we can switch the roles of circuits and inputs, going through a universal circuit.

Perhaps surprisingly, we show:

THEOREM 1.1 (INFORMAL). *Assuming the existence of a sub-exponentially-secure public-key functional encryption scheme for all circuits, where encryption time is polynomial in the input-size and sub-linear in the circuit-size, there exists indistinguishability obfuscation for all circuits.*

Furthermore, in the above theorem, it suffices to start from a scheme that supports only a single functional key and satisfies a mild selective-security indistinguishability-based guarantee. We can further relax the above to allow the encryption to also depend polynomially on circuit-depth (or even exponentially, assuming pseudo-random functions in NC^1).

We also show that the requirement for sub-linear dependence on circuit size can be traded, when moving to *multi-key* functional encryption schemes with a sub-linear dependence on the number of derived keys. Informally, in such multi-key schemes, it is guaranteed that an adversary in the possession of a set of functional keys, learns nothing on the encrypted message beyond the combination of function outputs corresponding to the set of keys.

In fact, we show a generic bootstrapping theorem that captures both the transformation from the multi-key setting to the single-key setting and the removal of depth dependence in single-key schemes.

THEOREM 1.2 (INFORMAL). *Assuming the existence of multi-key functional encryption schemes for all circuits, where encryption time is polynomial in the input-size and circuit-size, but sub-linear in the number of released keys, or single-key functional encryption schemes with sub-linear dependence on circuit-size and polynomial dependence on circuit-depth, there exist single-key functional encryption schemes with sub-linear dependence on circuit-size (and no further dependence on circuit-depth).*

This transformation, in particular, allows obtaining new IO candidates from existing multi-key functional encryption schemes such as the one by Garg et al. [37] (in its subexponentially-hard version).

Another corollary that follows from our techniques and previous results on FE with succinct keys [19] is that obfuscation size can always be reduced to linear in the function's circuit-size plus some overhead in circuit-depth.

COROLLARY 1.3 (INFORMAL). *Assuming sub-exponential hardness of the Learning with Errors problem and IO, there exists IO such that an obfuscation of any circuit C of depth d and input length n is of size $2|C| + \text{poly}(n, d, \lambda)$.*

Interpretation. Functional encryption schemes satisfying the succinctness properties required in Theorem 1.2 were previously known based on indistinguishability obfuscation [36] or the stronger notion of differing-inputs obfuscation [22]. Thus, our result establishes the equivalence of functional encryption and IO, up to some sub-exponential security loss. The question of basing IO on more standard assumptions still stands but is now reduced to improving the state of the art in functional encryption.

It is rather tempting to be pessimistic and to interpret our result as a lower-bound, showing that improving functional encryption based on standard assumptions may be very hard or perhaps straight out impossible. Our take on the result is quite optimistic. We hope that the construction would eventually lead to IO from more standard assumptions, or improved assumptions on multilinear graded encodings.³ Indeed, in the past few years, we have seen a remarkable progress in constructions of functional encryption based on standard assumptions [48, 49, 72]. The state-of-the-art scheme based on a standard assumption is that of Goldwasser, Kalai, Popa, Vaikuntanathan,

³Below, we mention subsequent work that has already partially fulfilled this hope.

and Zeldovich [46], relying on the sub-exponential learning with errors assumption. The construction achieves ciphertext size that only grows polynomially with the circuit output size and depth; thus, for circuits with, say, a single output bit, ciphertexts may indeed be sub-linear in circuit size, but this will not be the case for circuits with long outputs. Interestingly, the latter construction achieves a strong simulation-based security guarantee, under which sub-linear growth in the output size (let alone circuit-size) is actually impossible [2, 11, 46]. Reducing the dependence on the output (under an indistinguishability-based notion) has been a tantalizing problem. Now, this question becomes of central importance in the quest to achieve indistinguishability obfuscation.

Gorbunov, Vaikuntanathan, and Wee [50] showed how to construct predicate encryption schemes for all circuits (with *a priori* bounded depth) from sub-exponential hardness of the Learning with Errors problem (LWE). In their scheme, the ciphertext size is polynomial in the input length and the depth of the circuit and otherwise independent of the circuit size and output size. A predicate encryption scheme can be interpreted as a functional encryption scheme with a “weak attribute-hiding” property (see References [1, 50, 56] for more details). Strengthening this to “full attribute hiding” will give us a functional encryption scheme that satisfies the requirements of Theorem 1.2 and is yet another frontier in achieving indistinguishability obfuscation from LWE.

We next explain the main ideas standing behind our construction.

1.2 Main Ideas

Our starting point is a natural *input extension* approach: given an obfuscator $\mathcal{O}_{n-1}$ for circuits with input length $n - 1$, design an obfuscator $\mathcal{O}_n$ for circuits with input length n . Intuitively, this way we can get obfuscation for circuits with arbitrary polynomial input length—recursively apply the input extension step polynomially many times. The base case is trivial—for circuits C with a single input bit, simply define the obfuscation $\mathcal{O}_1(C)$ to be the corresponding truth table $(C(0), C(1))$.

A crucial feature of any such input extension procedure is the blowup it incurs in complexity. Indeed, a trivial input extension procedure,

$$\mathcal{O}_n(C(x_1, \dots, x_n)) := \mathcal{O}_{n-1}(C(x_1, \dots, x_{n-1}, 0) \circ C(x_1, \dots, x_{n-1}, 1)),$$

that first creates an $(n - 1)$ -bit input circuit with a doubly-long output (by taking two copies of the original circuit and fixing the last bit x_n to either 0 or 1) blows up the obfuscation size, at each step, by at least a multiplicative factor of two. Accordingly, such a procedure can only be applied logarithmically many times (indeed, it is equivalent to simply writing the truth table of the circuit). To avoid such blowup, we must ensure that the total size of circuits obfuscated in each recursive step does not outgrow the size of previous circuits (except perhaps by an additive amount).

At high-level, this work is dedicated to developing such an input-extension procedure, based on functional encryption.

From Token-based Obfuscation to Efficient Input-extension. The basic idea behind our input-extension procedure is founded on the concept of *token-based obfuscation* and its connection to function-hiding functional encryption. A token-based obfuscation algorithm consists of an obfuscation algorithm $\text{Tok.Obf}(C)$ that given a circuit C produces an obfuscation $\tilde{C}$ and a secret key SK . Unlike the standard notion of obfuscation, which would allow evaluating $\tilde{C}$ on any input x to learn $C(x)$, here evaluation requires a token $\tilde{x}$ corresponding to x . The token $\tilde{x}$ can be generated by an encoding algorithm $\text{Tok.Enc}(\text{SK}, x)$ using the secret key SK (for simplicity of exposition, we shall assume that Tok.Enc is deterministic). Security is guaranteed against any adversary that does not possess the secret key and only gets the obfuscation $\tilde{C}$, as well as a polynomial number of encoded inputs $\tilde{x}$ of its choice. For the notion to be non-trivial, the complexity of Tok.Enc is required to only depend on the input x , and not on the circuit C .

Intuitively (and for now thinking about an obfuscation as an opaque black-box), token-based obfuscation suggests a simple input-extension procedure:

$$\begin{aligned} \mathcal{O}_n(C(x_1, \dots, x_n)) &:= \text{Tok.Obf}(C(x_1, \dots, x_n)), \\ \mathcal{O}_{n-1}(\text{Tok.Enc}(\text{SK}, x_1, \dots, x_{n-1}, 0) \circ \text{Tok.Enc}(\text{SK}, x_1, \dots, x_{n-1}, 1)); \end{aligned}$$

namely, to obfuscate a circuit C with n -bit inputs $(x_1, \dots, x_n)$, obfuscate C using the token based obfuscation, and then use the obfuscator $\mathcal{O}_{n-1}$, to obfuscate a bit-fixing variant of the token generator,

$$\text{Tok.Enc}(\text{SK}, x_1, \dots, x_{n-1}, 0) \circ \text{Tok.Enc}(\text{SK}, x_1, \dots, x_{n-1}, 1),$$

that given $(x_1, \dots, x_{n-1})$ generates two encodings corresponding to fixing x_n to either 0 or 1.

Crucially, since the complexity of $\text{Tok.Enc}(\text{SK}, x)$ only grows with the encoded input $x \in \{0, 1\}^n$, the circuit recursively obfuscated by $\mathcal{O}_{n-1}$ (and then $\mathcal{O}_{n-2}$ and so on) is now bounded, through all steps, by a fixed polynomial $\text{poly}(n)$ in the input length n . Unwinding the recursion, the complexity of $\mathcal{O}_n$ will now be bounded by $\text{poly}(|C|) + n \cdot \text{poly}(n)$.

From Functional Encryption to Token-based Obfuscation. As observed in References [26, 46], token-based obfuscation can be constructed from any *symmetric-key* functional encryption scheme that has a succinct encryption circuit. Concretely, they show that it is possible to harness the existing message-hiding of functional encryption to also guarantee function hiding. Here a functional key FSK_C is guaranteed to hide the circuit C and can be viewed as a token-based obfuscation of C . The encryption algorithm $\text{Enc}(\text{SK}, \cdot)$, with the corresponding private encryption key SK , is then viewed as the token generator.

Combined with the token-based input extension procedure, this suggests a strategy for constructing obfuscation based on (symmetric-key) functional encryption with a succinct encryption circuit. Materializing this high-level strategy requires of course a more careful examination of the security guaranteed at each and every step. Assuming all the involved primitives satisfy an *ideal* (simulation-based) security guarantee would indeed allow implementing this strategy and eventually lead to an ideal obfuscation guarantee (known as virtual black-box security). However, ideal security is known to be impossible for either (succinct) functional encryption or obfuscation [2, 13].

The hope is that starting with a weaker indistinguishability-based guarantee for functional encryption would still allow to carry through the above strategy, leading to indistinguishability obfuscation. This turns out to encounter certain difficulties, which eventually lead to our requirement of public-key functional encryption (rather than symmetric-key), as well as our sub-exponential security requirement. We next overview these challenges and the way they are dealt with.

Under the Hood. A natural first attempt to achieve our goal is to mimic the ideal solution. Namely, starting from a (symmetric-key) function-hiding functional encryption scheme, to obfuscate any circuit C with input $\mathbf{x} = x_1 \dots x_n$, generate the functional key FSK_C , and add an obfuscation,

$$i\mathcal{O}(\text{Enc}(\text{SK}, x_1, \dots, x_{n-1}, 0) \circ \text{Enc}(\text{SK}, x_1, \dots, x_{n-1}, 1)),$$

of the corresponding (bit fixing) encryption circuit. This clearly satisfies the required functionality, but it is not clear how to prove security based on IO. While the function-hiding guarantee of symmetric-key FE holds in the presence of an encryption *oracle*, it may not hold when the adversary is presented with an actual circuit implementing this oracle. In particular, an indistinguishability obfuscation of the encryption circuit, which includes the secret encryption key, may potentially leak the secret key. (As a matter of fact, it may even be that the encryption function is *unobfuscatable* [13] in the sense that any circuit implementation thereof would leak the secret key.)

To overcome the above difficulty, we would like to obtain FE with a more robust function hiding guarantee that holds even in the presence of an indistinguishability obfuscation of the encryption

algorithm. Perhaps the first thought that comes to mind is to use *public-key* function-hiding FE, in which case the encryption circuit can be published without compromising security. However, it is easy to see that public-key function-hiding FE is too strong in the sense that it directly implies IO, which is the very thing we are trying to construct.⁴

Instead, we define a new notion of *puncturable functional encryption* that lies somewhere between the notions of symmetric-key and public-key for function-hiding FE. In particular, this notion may not provide full-fledged function hiding in the presence of the encryption key, but does satisfy certain weaker properties. Once we define this notion, we show how to obtain it from plain (non-function-hiding) public-key FE and show that the guaranteed security properties are already enough to prove the security of the IO input-extension procedure described above.

Puncturable FE. Roughly speaking, the notion of puncturable FE (PFE) satisfies two properties: *key indistinguishability* and *puncturing*. The first property asserts that for any two functions f_0 and f_1 , of the same size and functionality, it is possible to generate a fake functional key $\text{FSK}^* = \text{FSK}_{f_0, f_1}^*$ together with two fake encryption keys $\text{EK}_0^*, \text{EK}_1^*$ so that a real encryption key EK and a functional key FSK_{f_b} corresponding to f_b are indistinguishable from the fake functional key FSK^* together with the corresponding fake encryption key EK_b^* :

$$\text{EK}, \text{FSK}_{f_0} \approx_c \text{EK}_0^*, \text{FSK}^*, \quad \text{EK}_1^*, \text{FSK}^* \approx_c \text{EK}, \text{FSK}_{f_1}.$$

The gap between such an FE and a function hiding FE is that given the fake functional key FSK^* it may still be possible to distinguish EK_0^* from EK_1^* (indeed this is the case in our scheme).

The second property, called *puncturing*, is meant to bridge this gap. Our notion of puncturing is inspired by that of *puncturable pseudorandom functions* [21, 23, 57] and roughly says the following. First, for any message x , encryptions under the two fake keys are indistinguishable,

$$\text{Enc}(\text{EK}_0^*, x), \text{FSK}^* \approx_c \text{Enc}(\text{EK}_1^*, x), \text{FSK}^*.$$

Moreover, there is an efficient way to generate a so called punctured version $\text{EK}_0^*\{x\}, \text{EK}_1^*\{x\}$ of the keys $\text{EK}_0^*, \text{EK}_1^*$, so that the above indistinguishability holds even in the presence of *both* punctured keys. In terms of functionality, the punctured keys still allow to compute encryptions for all messages $x' \neq x$.

Why is Puncturable FE Enough? While the above puncturing requirement still does not say that the fake keys EK_0^* and EK_1^* are indistinguishable, it is sufficient for showing that indistinguishability obfuscations $iO(\text{Enc}(\text{EK}_0^*, \cdot))$ and $iO(\text{Enc}(\text{EK}_1^*, \cdot))$ of the corresponding encryption algorithms are in fact indistinguishable (even in the presence of the fake functional key FSK^*), provided that the underlying puncturable FE and IO are subexponentially secure. This is shown using a hybrid technique commonly used in the context of IO [15, 29, 30, 60].

Roughly speaking, we can move from an obfuscation of one circuit to an obfuscation of the other through a sequence of hybrids $\{\mathcal{H}_x\}_{x \in \{0,1\}^n}$ ranging over all messages x of length n . In $\mathcal{H}_x$, the obfuscated circuit uses the key EK_1^* to encrypt all messages x' that are lexicographically smaller than x , and uses the key EK_0^* to encrypt all other messages. Going from one hybrid to the next, we consider a circuit where both keys are punctured at the point x , and the encryption of x is hardwired. The new circuit computes the exact same function as the original one, and thus this change cannot be detected when the circuit is (indistinguishability) obfuscated. Now, we can switch an encryption under one key to an encryption under the other relying on the security of the puncturable PFE.

⁴Roughly speaking, an obfuscation of a circuit C consists of a functional key for C and the public-encryption algorithm. Now, anyone could encrypt any input x , and obtain $C(x)$ by invoking functional decryption.

The distinguishing gap between the hybrids additively grows as we move between one hybrid to the next, and overall degrades by a factor of 2^n . This is indeed the reason we need to start from subexponentially secure IO and puncturable FE (where the security parameter will be an appropriate polynomial in n).

The Final Construction. Overall, our input-extension step has the following form. To generate an obfuscation $iO_n(C)$ of a circuit C with n -bit input $\mathbf{x} = x_1 \dots x_n$, generate a (real) functional key FSK_C together with a (real) encryption key EK and add an obfuscation

$$iO_{n-1}(\text{Enc}(\text{EK}, x_1, \dots, x_{n-1}, 0) \circ \text{Enc}(\text{EK}, x_1, \dots, x_{n-1}, 1))$$

of the corresponding (bit fixing) encryption circuit, defined on $n - 1$ bits. To evaluate the circuit on input $\mathbf{x}$, the evaluator runs the obfuscation on the $(n - 1)$ -bit prefix of $\mathbf{x}$ and chooses the encryption $\text{Enc}(\text{EK}, \mathbf{x})$ according to the last bit of $\mathbf{x}$. Then it simply performs functional decryption.

Indistinguishability between obfuscations of two functionally equivalent circuits C_0 and C_1 is shown by first using the key-indistinguishability property to switch real keys $\text{EK}, \text{FSK}_{C_0}$, corresponding to an obfuscation of C_0 , into fake keys $\text{EK}_0, \text{FSK}^*$. Then, we switch to the fake keys $\text{EK}_1, \text{FSK}^*$ using the puncturing property following the hybrid argument outlined above. Finally, we can switch back to real keys $\text{EK}, \text{FSK}_{C_1}$ corresponding to an encryption of C_1 , using key indistinguishability again.

In terms of security, each invocation of the input extension step incurs a multiplicative loss of 2^n in the distinguishing gap, and overall when extending a trivial one-bit-input IO to n -bit-input IO, the overall degradation is roughly $2^{\sum_{i \in [n]} i} \approx 2^{n^2}$.

Constructing Puncturable FE and the Need for Public-key Encryption. Finally, we describe how to convert any public-key FE scheme FE into a new scheme PFE that is puncturable. The construction relies on similar techniques to those used for function hiding in the symmetric-key setting [26].

Concretely, to generate a key for a function f , we use the underlying scheme FE to generate a functional key FSK_g , for an augmented function g . The circuit $g = g[\text{CT}_0, \text{CT}_1]$ has two (plain) symmetric-key encryptions CT_0, CT_1 , under two independently chosen secret keys SK_0, SK_1 , hardwired into its code. In the actual scheme, both ciphertexts encrypt the circuit f . The function g expects as input, not only an input x for f but also a key SK_b . Given those, it decrypts the corresponding ciphertext CT_b , and applies the decrypted circuit to the input x .

The (real) encryption key EK then consists of the encryption key EK_{FE} of the underlying FE along with a secret key SK_b for a randomly chosen b , as well as a key K for a pseudorandom function. Encryption is done using the encryption of the underlying FE, encrypting not only the message x but also the secret key SK_b ; any randomness r required for the encryption algorithm is derived by applying the pseudorandom function to x , namely $r = \text{PRF}_K(x)$.

A fake encryption key EK_0^* is distributed identically to a real key EK , consisting of $\text{EK}_{\text{FE}}, \text{SK}_{b_0}, K$, for a random bit b_0 . The second key EK_1^* only differs in the symmetric key, which is chosen to be the second key SK_{b_1} , where $b_1 = 1 - b_0$. As expected, the fake functional key FSK^* corresponds now to $g[\text{CT}_0, \text{CT}_1]$ where CT_{b_0} encrypts f_0 and CT_{b_1} encrypts f_1 (rather than both encrypting the same function).

Proving key indistinguishability and puncturing naturally extends the ideas used to prove function-hiding in Reference [26]. (There, the adversary, never actually sees the encryption key, and is only given an oracle that computes encryptions.) Key indistinguishability, follows directly from the security of the symmetric encryption scheme; indeed, the only difference between the real keys $\text{EK}, \text{FSK}_{f_0}$ and fake keys $\text{EK}_0^*, \text{FSK}^*$ is that in the latter the ciphertext CT_{b_1} embedded in FSK^* encrypts f_1 instead of f_0 . Since EK_0^* only includes the secret key SK_{b_0} and is independent of

SK_{b_1} , semantic security applies. A symmetric argument holds for the real keys EK, FSK_{f_1} and fake keys EK_1^*, FSK^* .

We then need to prove that encryptions of any message x under the fake key EK_0^* are indistinguishable from ones under EK_1^* , even in the presence of a punctured version of the keys. Let us first understand why such indistinguishability holds when the adversary does not obtain any encryption keys. In this case, which is analogous to Reference [26], we can rely on the security of the underlying FE. Specifically, recall that the only difference between encryptions under the two keys is that when encrypting with EK_0^* , we use the underlying encryption to encrypt (x, SK_{b_0}) , whereas when encrypting under EK_0^* , we encrypt (x, SK_{b_1}) . Note, however, that since f_0 and f_1 compute the same function, the function $g[CT_0, CT_1]$ does distinguish between two such inputs. Indeed, in the first case it outputs $f_0(x)$ and in the second $f_1(x)$. This, in particular, means that two such encryptions are indistinguishable.

It is left to show that we can produce a punctured version of the encryption keys $EK_0^*\{x\}, EK_1^*\{x\}$ such that the above would hold even in their presence. This basically implies that we have to: (1) generate a punctured version of the encryption key EK_{FE} in the underlying (non-function-hiding) functional encryption scheme, and (2) generate a punctured version of the pseudorandom functional key K . The second requirement can be dealt with easily relying on puncturable pseudorandom functions [21, 23, 57]—these are exactly pseudorandom functions where it is possible to puncture the secret key at any point x , so that given the punctured key $K\{x\}$ the value $r = PRF_K(x)$ is pseudorandom.

The question is how to satisfy the first requirement, namely, ensure puncturing for the functional encryption key EK_{FE} . This is exactly where we rely on public-key functional encryption. Indeed, with public-key functional encryption the encryption key $EK_{FE} = PK$ is trivially puncturable. That is, given the key PK itself (unchanged) indistinguishability is still guaranteed. This is in contrast to symmetric-key FE where in the presence of the encryption key, indistinguishability may no longer hold.⁵

1.3 Concurrent Work

We mention several concurrent and independent works:

- Ananth and Jain [4] also show how to construct indistinguishability obfuscation from sub-exponentially secure public-key functional encryption, under a similar assumption on the running-time of the encryption algorithm (which they term *compactness*). The two works take a somewhat different perspective to the problem. At high-level, Ananth and Jain show that functional encryption schemes as above can be converted into a multi-input functional encryption, a notion defined by Goldwasser et al. [44] that is known to imply indistinguishability obfuscation. The core step of their construction is a transformation from n -input FE to $(n + 1)$ -input FE, which is analogous to our recursive step of basing $(n + 1)$ -bit-input IO on n -bit-input IO. Our proof of security is perhaps more simple and concise, which we attribute to the fact that in each recursive step we fully exploit the expressive power of the IO guarantee, compared to the less expressive (multi-input) FE guarantee. In particular, we are able directly invoke previous proof techniques developed for IO.
- Brakerski, Komargodski, and Segev [24] show how to convert any (single-input) symmetric-key functional encryption scheme into an $O(1)$ -input symmetric-key scheme (or doubly

⁵The subsequent work [58] shows how to construct a puncturable symmetric-key FE directly from plain symmetric-key FE, thereby allowing to base the entire transformation on symmetric-key FE.

logarithmic-input assuming sub-exponential security), which is not known to be sufficient to go all the way to IO polynomially large inputs.

- Ananth, Jain, and Sahai [6] show how IO can be bootstrapped to always have linear-size overhead. By developing new techniques, they improve on the above Corollary 1.3 in two aspects. First, they avoid the LWE assumption. Second, they avoid the polynomial dependence of the obfuscated circuit-size on the depth of the original circuit.
- Ananth, Jain, and Sahai [5] also show how to transform any collusion-resistant FE into a single-key FE scheme with succinct encryption circuits.

1.4 Followup Work

We mention several subsequent works that have relied on our result, or have extended it:

- Lin, Pass, Seth, and Telang [65] show a different transformation from (public-key) functional encryption to IO. While their transformation shares much of the structure of our transformation, it has different features such as better (but still sub-exponential) security loss, and admits a very elegant description in the language of succinct randomized encodings [15, 29, 60]. Their description of the transformation from public-key FE to IO shares much of the same high-level structure as the classical Goldreich-Goldwasser-Micali transformation from a pseudorandom generator to a pseudorandom function.
- Lin, Pass, Seth, and Telang [64], in another work, introduce a relaxation of IO called Exponential Indistinguishability Obfuscation (XIO) that only requires that the size of an obfuscated circuit is sub-linear in the size of its truth table. Based on our result and the learning with errors (LWE) assumption, they show that this relaxation suffices for obtaining full-fledged IO.
- A progression of works [7, 62, 63, 66, 67] has shown how to reduce the degree of multilinear maps required for constructing IO (assuming also the existence of appropriate local pseudorandom generators). The core of these works is a construction of FE, which is then bootstrapped to IO using our transformation.
- Bitansky, Nishimaki, Passalegue, and Wichs [16] show how to obtain IO from sub-exponentially secure *symmetric-key* functional encryption and plain public-key encryption. In fact, they show how these primitives together imply a public-key functional encryption scheme and then invoke our transformation. Subsequently, Kitagawa, Nishimaki, and Tanaka [58] removed the assumption of public-key encryption, showing that subexponentially secure symmetric-key functional encryption implies public-key functional encryption.
- Garg, Pandey, Srinivasan, and Zhandry [38, 39] show how many of the applications of IO (such as the hardness of PPAD and multiparty key exchange) can be based instead directly on polynomially secure functional encryption. Their reduction invokes a variant of our input-extension technique, but avoids the sub-exponential security loss. A general framework that captures these works was subsequently introduced by Liu and Zhandry [68].
- Li and Micciancio [61] and Garg and Srinivasan [40] independently show a construction of a multi-key (collusion-resistant) functional encryption starting from any polynomially secure single-key functional encryption scheme with succinct encryption circuits. Such a transformation follows from our results as IO implies a collusion-resistant functional encryption scheme by the results of Reference [36], except that we lose a sub-exponential security factor that comes from invoking our transformation. The results of References [40, 61] avoid this loss. This, together with our result that transforms a polynomially secure

collusion-resistant FE into a single-key FE scheme with succinct encryption circuits, shows that both variants of FE are in fact equivalent.

2 DEFINITIONS

We review basic concepts and present the basic definitions used throughout the article.

2.1 Standard Computational Concepts

We rely on the standard notions of Turing machines and Boolean circuits.

- We say that a Turing machine is PPT if it is probabilistic and runs in polynomial time.
- For (a deterministic and time-bounded) algorithm A , we denote by $A(\cdot)$ the corresponding boolean circuit (computing the same function). We denote by $A(x, \cdot)$ the circuit where a prefix of the input is fixed to x .
- A polynomial-size circuit family C is a sequence of circuits $C = \{C_\lambda\}_{\lambda \in \mathbb{N}}$, such that each circuit C_λ is of polynomial size $\lambda^{O(1)}$ and has $\lambda^{O(1)}$ input and output bits.
- We follow the standard habit of modeling any efficient adversary as a family of polynomial-size circuits. For an adversary $\mathcal{A}$ corresponding to a family of polynomial-size circuits $\{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, we often omit the subscript λ , when it is clear from the context. When we write $\mathcal{A}(x)$ for a circuit $\mathcal{A}$, without specifying the length of x , we assume that $\mathcal{A}$ only reads an ℓ -bit prefix of x , where ℓ is the input-size of the circuit $\mathcal{A}$.
- A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is negligible if $f(\lambda) = \lambda^{-\omega(1)}$.
- For random variables X and Y , distinguisher $\mathcal{D}$, and $0 < \mu < 1$, we write $X \approx_{\mathcal{D}, \mu} Y$ if

$$|\Pr[\mathcal{D}(X) = 1] - \Pr[\mathcal{D}(Y) = 1]| \leq \mu.$$

- Two ensembles of random variables $\mathcal{X} = \{X_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are said to be computationally indistinguishable, denoted by $\mathcal{X} \approx_c \mathcal{Y}$, if for all polynomial-size distinguishers $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function μ such that for all λ ,

$$X_\lambda \approx_{\mathcal{D}, \mu(\lambda)} Y_\lambda.$$

For a concrete function δ , we denote by $\mathcal{X} \approx_\delta \mathcal{Y}$ the case that the above distinguishing gap is bounded by $\delta^{\Omega(1)}$. That is, for every polynomial-size distinguisher $\mathcal{D}$ there exists a constant $\alpha \leq 1$ such that $X_\lambda \approx_{\mathcal{D}, \delta(\lambda)^\alpha} Y_\lambda$.

- Our constructions rely on primitives that are subexponentially secure, or more generally δ -secure for some negligible δ . Typically, in the literature, this addresses adversaries of subexponential size that break the underlying system with subexponentially small probability. Like many other works in the context of indistinguishability obfuscation, for our purpose, it will be sufficient to require that polynomial-size (rather than subexponential-size) adversaries break the system with subexponentially small probability.

2.2 Single-Key FE with Succinct Encryption

In this work, we consider a restricted notion of single-key functional encryption schemes where the function is known in setup time. This is in contrast to the typical stronger definition in the literature where the function is not known in setup time and the encryption key and a master decryption key are generated independently of the function. (Note that since we aim to construct IO from functional encryption, considering such a weaker notion only strengthens the result.) Furthermore, we will require certain succinctness of the encryption circuit, which will play an essential role in our constructions.

Such a scheme FE, for a function class $\mathcal{F}$ (represented by boolean circuits) and message space $\{0, 1\}^*$, consists of three PPT algorithms (FE.Setup, FE.Enc, FE.Dec) with the following syntax:

- $\text{FE.Setup}(1^\lambda, f)$: takes as input a security parameter λ in unary and function $f \in \mathcal{F}$ and outputs a public key PK and a functional key FSK_f .
- $\text{FE.Enc}(\text{PK}, m)$: takes as input a public key PK and a message $m \in \{0, 1\}^*$ and outputs an encryption of m . We shall sometimes address the randomness r used in encryption explicitly, which we denote by $\text{FE.Enc}(\text{PK}, m; r)$.
- $\text{FE.Dec}(\text{FSK}_f, \text{CT})$: takes as input a functional key FSK_f and a ciphertext CT and outputs $\hat{m}$.

We next define the required correctness, security, and efficiency properties.

Definition 2.1 (Single-key, Selectively Secure, Public-Key FE with Succinct Encryption). A tuple of PPT algorithms $\text{FE} = (\text{FE.Setup}, \text{FE.Enc}, \text{FE.Dec})$ is a single-key, selectively secure, public-key functional encryption scheme with succinct encryption, for function class $\mathcal{F}$, and message space $\{0, 1\}^*$, if it satisfies the following:

- (1) **Correctness**: for every $\lambda, n \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and function $f \in \mathcal{F}$, with domain $\{0, 1\}^n$,

$$\Pr \left[f(m) \leftarrow \text{FE.Dec}(\text{FSK}_f, \text{CT}) \mid \begin{array}{l} (\text{PK}, \text{FSK}_f) \leftarrow \text{FE.Setup}(1^\lambda, f) \\ \text{CT} \leftarrow \text{FE.Enc}(\text{PK}, m) \end{array} \right] = 1.$$

- (2) **Selective security**: for any polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\lambda)$ such that for any $\lambda, n \in \mathbb{N}$, any $m_0, m_1 \in \{0, 1\}^n$, and function $f \in \mathcal{F}$ such that $f(m_0) = f(m_1)$,

$$\text{PK}, \text{FSK}_f, \text{FE.Enc}(\text{PK}, m_0) \approx_{\mathcal{A}, \mu} \text{PK}, \text{FSK}_f, \text{FE.Enc}(\text{PK}, m_1),$$

where $(\text{PK}, \text{FSK}_f) \leftarrow \text{FE.Setup}(1^\lambda, f)$.

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

- (3) **Succinct encryption circuit**: there exists a polynomial Φ and a constant $0 < \varepsilon \leq 1$, such that for any input-size n , circuit-size s , and s -size function $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$,

$$|\text{FE.Enc}(\text{PK}, \cdot)| \leq s^{1-\varepsilon} \cdot \Phi(n, \lambda),$$

where $(\text{PK}, \text{FSK}_f) \leftarrow \text{FE.Setup}(1^\lambda, f)$.

- Encryption is **fully succinct** if $\varepsilon = 1$.
- Encryption is **weakly depth-succinct** if

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq s^{1-\varepsilon} \cdot \Phi(n, d, \lambda).$$

- Encryption is **very weakly depth-succinct** if

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq s^{1-\varepsilon} \cdot \Phi(n, 2^d, \lambda).$$

In Section 4, we show that any weakly depth-succinct scheme implies a succinct scheme. Furthermore, assuming the existence of pseudorandom functions in NC^1 this extends to very-weakly depth-succinct schemes. Accordingly, throughout most of the article, we restrict attention to succinct schemes (and do not address depth).

2.3 Indistinguishability Obfuscation

We define indistinguishability obfuscation (IO) with respect to a given class of circuits. The definition is formulated as in Reference [13].

Definition 2.2 (Indistinguishability Obfuscation). A PPT algorithm iO is said to be an *indistinguishability obfuscator* for a class of circuits $\mathcal{C}$ if it satisfies the following:

- (1) **Functionality:** for any $C \in \mathcal{C}$ and security parameter λ ,

$$\Pr_{iO} [\forall x : iO(C, 1^\lambda)(x) = C(x)] = 1.$$

- (2) **Indistinguishability:** for any polynomial-size distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ there exists a negligible function $\mu(\lambda)$, such that for any two circuits $C_0, C_1 \in \mathcal{C}$ that compute the same function and are of the same size,

$$iO(C_0, 1^\lambda) \approx_{\mathcal{D}, \mu} iO(C_1, 1^\lambda),$$

where the probability is over the coins of iO .

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

2.4 Puncturable Pseudorandom Functions

We consider a simple case of puncturable pseudorandom functions (PRFs) where any PRF may be punctured at a single point. The definition is formulated as in Reference [74] and is satisfied by the GGM [43] pseudorandom function [21, 23, 57].

Definition 2.3 (Puncturable PRFs). Let $k(\lambda)$ be a polynomially bounded length function. An efficiently computable family of functions,

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

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

- (1) **Punctured-key correctness:** 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 polynomial-size distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ there exists a negligible function $\mu(\lambda)$, such that for all $\lambda \in \mathbb{N}$, and any $x^* \in \{0, 1\}^*$,

$$K\{x^*\}, \text{PRF}_K(x^*) \approx_{\mathcal{D}, \mu} K\{x^*\}, r,$$

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

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

Remark 2.4. The definition requires that punctured keys are of the same size as plain keys and that the same evaluation algorithm PRF works for both. We note that this requirement is w.l.o.g. by appropriate padding.

Remark 2.5. The original [43] construction of pseudorandom functions is described for a fixed input length. Here, we define and use pseudorandom functions over arbitrary strings. A slight variant of the [43] construction is known to achieve this [42] and admits exactly the same puncturing properties (and proof of security).

2.5 One-Time Symmetric Encryption with Local Decryption

A one-time symmetric encryption scheme Sym with local decryption consists of a tuple of two PPT algorithms $(\text{Sym.Enc}, \text{Sym.Dec})$ with the following syntax:

- $\text{Sym.Enc}(\text{SK}, m)$ takes as input a symmetric key $\text{SK} \in \{0, 1\}^\lambda$, where λ is the security parameter, and a message $m \in \{0, 1\}^n$, and outputs a size- n ciphertext CT .
- $\text{Sym.Dec}(\text{SK}, \text{CT}_i, i)$ takes as input the key SK a ciphertext bit CT_i and the index $i \in [n]$, and outputs the decrypted message bit m_i .

Throughout, we interpret $\text{Sym.Dec}(\text{SK}, \text{CT})$, where no set of bits is specified as decrypting the entire ciphertext $\text{Sym.Dec}(\text{SK}, \text{CT}) = \text{Sym.Dec}(\text{SK}, \text{CT}_1, 1), \dots, \text{Sym.Dec}(\text{SK}, \text{CT}_n, n)$.

Definition 2.6 (One-Time Symmetric Encryption with Local Decryption). A pair of deterministic polynomial-time algorithms $(\text{Sym.Enc}, \text{Sym.Dec})$ is a one-time symmetric encryption scheme for message space $\{0, 1\}^*$ if it satisfies the following:

- (1) **Correctness:** for every security parameter $\lambda \in \mathbb{N}$, message length $n \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and index $i \in [n]$,

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

We say that Sym is **shallow** if it has a decryption circuit of depth $\log \lambda$.

- (2) **One-time indistinguishability:** for any polynomial-size distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ there exists a negligible function $\mu(\lambda)$, such that for all $\lambda \in \mathbb{N}$, and any equal size messages m_0, m_1 ,

$$\text{Sym.Enc}(\text{SK}, m_0) \approx_{\mathcal{D}, \mu} \text{Sym.Enc}(\text{SK}, m_1),$$

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

We say that Sym is δ -secure, for some concrete negligible function $\delta(\lambda)$, if for all polynomial-size adversaries the above distinguishing gap $\mu(\lambda)$ is smaller than $\delta(\lambda)^{\Omega(1)}$.

Such encryption schemes follow from any pseudorandom function (or local pseudorandom generator). Given pseudorandom functions in NC^1 (which are known under various standard assumptions [17]), there exist such schemes that are also shallow.

FACT 2.7 ([43, 52]). *Assuming (δ -secure) one-way functions, there exists a (δ -secure) one-time symmetric encryption with local decryption. Assuming pseudorandom functions in NC^1 , the scheme is shallow.*

2.6 Decomposable Randomized Encodings

We rely on the notion of decomposable randomized encodings (REs) from References [9, 10, 55, 77]. There are different variants of REs in the literature differing in their security and decomposition (or locality) properties. Here, we define the properties that will be used in our constructions, which are satisfied by Yao's garbled circuit [77] in conjunction with appropriate pseudorandom functions.

Such a scheme RE consists of two polynomial-time algorithms $(\text{RE.Enc}, \text{RE.Dec})$ with the following syntax:

- $\text{RE.Enc}(f, 1^\lambda)$: takes as input a circuit $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$ and a security parameter 1^λ and outputs a new encoder circuit $\hat{f} : \{0, 1\}^n \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^*$.
- $\text{RE.Dec}(\hat{y})$ takes as input an encoding $\hat{y}$ and outputs y .

We next define the required correctness, security, and decomposability properties.

Definition 2.8 (Decomposable Randomized Encoding). A pair of polynomial-time algorithms $\text{RE} = (\text{RE.Enc}, \text{RE.Dec})$ is a decomposable randomized encoding if it satisfies the following:

- (1) **Correctness:** for every $\lambda, n \in \mathbb{N}$, circuit $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$, and input $x \in \{0, 1\}^n$,

$$\Pr \left[\text{RE.Dec}(\hat{f}(x; r)) = f(x) \mid \begin{array}{l} \hat{f} \leftarrow \text{RE.Enc}(f, 1^\lambda) \\ r \leftarrow \{0, 1\}^\lambda \end{array} \right] = 1.$$

- (2) **Privacy:** there exists a PPT simulator RE.Sim such that for any polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\lambda)$ such that for any $\lambda, n \in \mathbb{N}$, any $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and $x \in \{0, 1\}^n$,

$$\hat{f}(x; r) \approx_{\mathcal{A}, \mu} \text{RE.Sim}(f, f(x)),$$

where $\hat{f} \leftarrow \text{RE.Enc}(f, 1^\lambda)$ and $r \leftarrow \{0, 1\}^\lambda$.

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

- (3) **Decomposability:** there exists a polynomial Φ such that for any input size $n \in \mathbb{N}$ and size- s circuit $f : \{0, 1\}^n \rightarrow \{0, 1\}^*$, the encoder circuit $\hat{f} \leftarrow \text{RE.Enc}(f, 1^\lambda)$ can be decomposed into $\ell = s \cdot \Phi(n, \lambda)$ circuits,

$$\hat{f}(x; r) = (\hat{f}_1(x; r), \dots, \hat{f}_\ell(x; r)),$$

each of size $\Phi(n, \lambda)$ and with single output bit.

We say that RE is **shallow** if the depth of each $\hat{f}_i$ is $\log \lambda$.

Such randomized encodings can be constructed from one-time symmetric encryption (Definition 2.6) and pseudorandom functions (which can in turn be constructed from one-way functions) based on Yao's garbled circuit. Assuming pseudorandom functions in NC^1 , the randomized encodings are shallow.

FACT 2.9 ([10, 77]). *Assuming (δ -secure) one-way functions there exists a (δ -secure) decomposable randomized encoding. Assuming also pseudorandom functions in NC^1 , it is shallow.*

3 THE TRANSFORMATION

In this section, we describe the transformation from FE to IO and analyze it. The transformation consists of two steps. First, in Section 3.1, we define and construct a notion of puncturable functional encryption with a succinct encryption circuit that satisfies a weak form of function hiding. Then, in Section 3.2, we construct IO from such FE schemes.

3.1 Puncturable Functional Encryption

We now define the abstraction of Puncturable Functional Encryption (PFE) that will be used in our transformation.

Such a scheme PFE, for a function class $\mathcal{F}$ (represented by boolean circuits) and message space $\{0, 1\}^*$, consists of three polynomial-time algorithms (PFE.Setup, PFE.Enc, PFE.Dec), the first probabilistic and the other two deterministic, with the following syntax:

- $\text{PFE.Setup}(1^\lambda, f)$: takes as input a security parameter λ in unary and function $f \in \mathcal{F}$ with domain $\{0, 1\}^n$, and outputs an encryption key EK and a functional key FSK_f .
- $\text{PFE.Enc}(\text{EK}, m)$: takes as input an encryption key EK and a message $m \in \{0, 1\}^n$ and outputs a ciphertext CT encrypting m .

- $\text{PFE.Dec}(\text{FSK}_f, \text{CT})$: takes as input a functional key FSK_f and a ciphertext CT and outputs $\widehat{m}$.

We next define the required correctness, security, and efficiency properties. To define the security properties, we require the existence of two additional algorithms:

- $\text{PFE.Setup}^*(1^\lambda, f_0, f_1)$: a PPT algorithm that takes as input the security parameter λ , and two functions $f_0, f_1 \in \mathcal{F}$ with domain $\{0, 1\}^n$ and outputs a pair of keys $(\text{EK}_0^*, \text{EK}_1^*)$ and a fake functional key FSK^* .
- $\text{PFE.Punc}(\text{EK}_0^*, \text{EK}_1^*, m)$: a deterministic polynomial-time algorithm that takes as input keys $(\text{EK}_0^*, \text{EK}_1^*)$ and a message $m \in \{0, 1\}^n$ and outputs a punctured keys $(\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\})$.

Definition 3.1 (Puncturable Functional Encryption). $\text{PFE} = (\text{PFE.Setup}, \text{PFE.Enc}, \text{PFE.Dec})$ is a puncturable functional encryption scheme with succinct encryption, for function class $\mathcal{F}$ and message space $\{0, 1\}^*$, if it satisfies the following:

- (1) **Correctness**: for every $\lambda, n \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and function $f \in \mathcal{F}$ with domain $\{0, 1\}^n$,

$$\Pr \left[f(m) = \text{PFE.Dec}(\text{FSK}_f, \text{CT}) \mid \begin{array}{l} (\text{EK}, \text{FSK}_f) \leftarrow \text{PFE.Setup}(1^\lambda, f) \\ \text{CT} = \text{PFE.Enc}(\text{EK}, m) \end{array} \right] = 1.$$

- (2) **Punctured-key Correctness**: for every $\lambda, n \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and functions $f_0, f_1 \in \mathcal{F}$, and $b \in \{0, 1\}$,

$$\Pr \left[\forall m' \in \{0, 1\}^n \setminus \{m\} : \text{PFE.Enc}(\text{EK}_b^*\{m\}, m') = \text{PFE.Enc}(\text{EK}_b^*, m') \right] = 1,$$

where $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$ and $(\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}) = \text{PFE.Punc}(\text{EK}_0^*, \text{EK}_1^*, m)$.

- (3) **Key indistinguishability**: for any polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\lambda)$, such that for any $\lambda \in \mathbb{N}$, equal-size functions $f_0, f_1 \in \mathcal{F}$, and $b \in \{0, 1\}$:

$$\text{EK}_b, \text{FSK}_{f_b} \approx_{\mathcal{A}, \mu} \text{EK}_b^*, \text{FSK}^*,$$

where $(\text{EK}_b, \text{FSK}_{f_b}) \leftarrow \text{PFE.Setup}(1^\lambda, f_b)$ and $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$.

- (4) **Indistinguishability at punctured points**: for any polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\lambda)$, such that for any $\lambda, n \in \mathbb{N}$, $m \in \{0, 1\}^n$, and equal-size functions $f_0, f_1 \in \mathcal{F}$ such that $f_0(m) = f_1(m)$,

$$\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}, \text{FSK}^*, \text{PFE.Enc}(\text{EK}_0^*\{m\}, m) \approx_{\mathcal{A}, \mu}$$

$$\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}, \text{FSK}^*, \text{PFE.Enc}(\text{EK}_1^*\{m\}, m),$$

where $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$ and $(\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}) = \text{PFE.Punc}(\text{EK}_0^*, \text{EK}_1^*, m)$.

We further say that PFE is δ -secure, for some concrete negligible function $\delta(\lambda)$, if for all polynomial-size adversaries the above distinguishing gaps $\mu(\lambda)$ are smaller than $\delta(\lambda)^{\Omega(1)}$.

- (5) **Succinct encryption circuit**: there exists a polynomial Φ and a constant $0 < \varepsilon \leq 1$, such that for any input-size n , circuit-size s , s -size functions $f, f' : \{0, 1\}^n \rightarrow \{0, 1\}^*$, $m \in \{0, 1\}^n$, and $b \in \{0, 1\}$,

$$|\text{PFE.Enc}(\text{EK}_b^*, \cdot)| = |\text{PFE.Enc}(\text{EK}_b^*\{m\}, \cdot)| = |\text{PFE.Enc}(\text{EK}, \cdot)| \leq s^{1-\varepsilon} \cdot \Phi(n, \lambda),$$

where $(\text{EK}, \text{FSK}) \leftarrow \text{PFE.Setup}(1^\lambda, f)$, $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f, f')$,

Fig. 1. A puncturable functional encryption.

$$(EK_0^*\{m\}, EK_1^*\{m\}) \leftarrow \text{PFE.Punc}(EK_0^*, EK_1^*, m).$$

Encryption is **fully succinct** if $\varepsilon = 1$.

Puncturable FE from Single-key FE. We now show that any single-key, selectively secure public-key FE scheme (Definition 2.1) implies a puncturable FE scheme that has essentially the same succinctness properties as the original scheme.

Ingredients. We rely on the following primitives:

- A single-key, selectively secure, public-key functional encryption scheme FE for all circuits.
- A one-time symmetric encryption scheme Sym.
- A puncturable pseudo-random function family $\mathcal{PRF}$.

The constructed scheme PFE consists of the algorithms

$$(\text{PFE.Setup}, \text{PFE.Enc}, \text{PFE.Dec}, \text{PFE.Setup}^*, \text{PFE.Punc}),$$

described in Figure 1.

THEOREM 3.2. *PFE is a puncturable functional encryption for all circuits. If FE is succinct, then so is PFE. If FE, Sym, $\mathcal{PRF}$ are all δ -secure, then so is PFE.*

PROOF. We prove that the constructed scheme satisfies the properties of a PFE (Definition 3.1).

Correctness: Fix a security parameter $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and $f \in \mathcal{F}$ with domain $\{0, 1\}^n$. By the functionality of FE and correctness of Sym,

$$\begin{aligned} \text{PFE.Dec}(\text{FSK}_f, \text{FCT}) &= \text{FE.Dec}(\text{FSK}_F, \text{FCT}) \\ &= F(m, \text{SK}_\beta, \beta) \\ &= U(\text{Sym.Dec}(\text{SK}_\beta, \text{CT}_\beta), m) \\ &= U(f, m) = f(m), \end{aligned}$$

where $(\text{EK}, \text{FSK}_f) \leftarrow \text{PFE.Setup}(1^\lambda, f)$, $\text{FCT} = \text{PFE.Enc}(\text{EK}, m)$, $\text{EK} = (\text{PK}, \text{K}, \text{SK}_\beta, \beta)$, F is the function underlying FSK_f , and CT_0, CT_1 are the ciphertexts corresponding to F .

Punctured-key correctness: Fix a security parameter $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and $f_0, f_1 \in \mathcal{F}$ with domain $\{0, 1\}^n$. By the punctured-key correctness of $\mathcal{PRF}$, for any $m' \in \{0, 1\}^n \setminus \{m\}$ and $b \in \{0, 1\}$,

$$\begin{aligned} \text{PFE.Enc}(\text{EK}_b^*, m') &= \text{FE.Enc}(\text{PK}, (m', \text{SK}_{\beta_b}, \beta_b); \text{PRF}_{\text{K}\{m\}}(m')) \\ &= \text{FE.Enc}(\text{PK}, (m', \text{SK}_{\beta_b}, \beta_b); \text{PRF}_{\text{K}}(m')) \\ &= \text{PFE.Enc}(\text{EK}_b^*, m'), \end{aligned}$$

where $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$, $\text{EK}_b^* = (\text{PK}, \text{K}, \text{SK}_{\beta_b}, \beta_b)$, $(\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}) = \text{PFE.Punc}(\text{EK}_0^*, \text{EK}_1^*, m)$, and $\text{K}\{m\} = \text{Punc}(\text{K}, m)$.

Key indistinguishability: Fix any polynomial-size $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$. By the indistinguishability of Sym, there exists a negligible $\mu_{\text{Sym}}(\lambda)$, such that for any $\lambda, n \in \mathbb{N}$, $m \in \{0, 1\}^n$, equal-size $f_0, f_1 \in \mathcal{F}$ with domain $\{0, 1\}^n$, and $b \in \{0, 1\}$,

$$\text{EK}, \text{FSK}_{f_b} \approx_{\mathcal{A}, \mu_{\text{Sym}}} \text{EK}_b^*, \text{FSK}^*,$$

where $(\text{EK}, \text{FSK}_{f_b}) \leftarrow \text{PFE.Setup}(1^\lambda, f_b)$ and $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$.

Indeed, EK and EK_b are identically distributed, whereas the only difference between FSK_{f_b} and FSK^* is in how the underlying function is generated. In the first, we generate F such that both ciphertexts $\{\text{CT}_\alpha = \text{Sym.Enc}(\text{SK}_\alpha, f_b)\}_{\alpha \in \{0, 1\}}$ encrypt the same function f_b . In the second, $\text{CT}_{1-\beta_b} = \text{Sym.Enc}(\text{SK}_{1-\beta_b}, f_{1-b})$ encrypts f_{1-b} . The key EK , respectively, EK_b^* , only include the key SK_{β_b} and not the key $\text{SK}_{1-\beta_b}$. Thus, the indistinguishability of Sym applies.

Indistinguishability at punctured points: Fix any polynomial-size $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$. By pseudorandomness at punctured points of $\mathcal{PRF}$ and selective security of FE, there exist a negligible $\mu_{\text{PRF}}(\lambda), \mu_{\text{FE}}(\lambda)$, such that for any $\lambda, n \in \mathbb{N}$, $m \in \{0, 1\}^n$, $f_0, f_1 \in \mathcal{F}$ such that $f_0(m) = f_1(m)$, and $b \in \{0, 1\}$,

$$\begin{aligned} &\{\text{EK}_\alpha^*\{m\}\}_{\alpha \in \{0, 1\}}, \text{FSK}^*, \text{PFE.Enc}(\text{EK}_b^*, m) \\ &= \{\text{PK}, \text{K}\{m\}, \text{SK}_{\beta_\alpha}, \beta_\alpha\}_{\alpha \in \{0, 1\}}, \text{FSK}^*, \text{FE.Enc}(\text{PK}, (m, \text{SK}_{\beta_b}, \beta_b); \text{PRF}_{\text{K}}(m)) \approx_{\mathcal{A}, \mu_{\text{PRF}}} \\ &\quad \{\text{PK}, \text{K}\{m\}, \text{SK}_{\beta_\alpha}, \beta_\alpha\}_{\alpha \in \{0, 1\}}, \text{FSK}^*, \text{FE.Enc}(\text{PK}, (m, \text{SK}_{\beta_b}, \beta_b); \text{r}) \approx_{\mathcal{A}, \mu_{\text{FE}}} \\ &\quad \{\text{PK}, \text{K}\{m\}, \text{SK}_{\beta_\alpha}, \beta_\alpha\}_{\alpha \in \{0, 1\}}, \text{FSK}^*, \text{FE.Enc}(\text{PK}, (m, \text{SK}_0, 0); \text{r}), \end{aligned}$$

where $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE.Setup}^*(1^\lambda, f_0, f_1)$, $(\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}) = \text{PFE.Punc}(\text{EK}_0^*, \text{EK}_1^*, m)$, $\text{K}\{m\} = \text{Punc}(\text{K}, m)$, and $\text{r} \leftarrow \{0, 1\}^\lambda$.

The first indistinguishability follows from pseudorandomness at punctured points. The second indistinguishability follows from the selective security of FE and the fact that the function F^*

underlying FSK^* satisfies

$$F^*(m, \text{SK}_{\beta_b}, \beta_b) = U(\text{Sym.Dec}(\text{SK}_{\beta_b}, \text{CT}_{\beta_b}), m) = U(f_b, m) = f_b(m) = f_0(m) = F^*(m, \text{SK}_0, 0).$$

Overall, the distribution $\text{EK}_0^*\{m\}, \text{EK}_1^*\{m\}, \text{FSK}^*, \text{PFE.Enc}(\text{EK}_b^*, m))$ is indistinguishable from a distribution that is independent of b , and thus the two distributions corresponding to $b \in \{0, 1\}$ are indistinguishable.

For the last two properties, it follows readily that if $\text{Sym}, \mathcal{PRF}, \text{FE}$ are δ -secure for some concrete negligible δ then PFE is δ -secure.

Succinct encryption: Assume that FE is succinct. That is, there exist a polynomial Φ_{FE} and constant $0 < \varepsilon \leq 1$ such that the size of its encryption circuit is bounded by

$$s^{1-\varepsilon} \cdot \Phi_{\text{FE}}(n, \lambda),$$

where n, s are the input-size and circuit-size of the function chosen during the setup phase.

We show that PFE is also succinct with parameters $(\varepsilon, \Phi_{\text{PFE}})$ for a fixed polynomial Φ_{PFE} . First, observe that fake keys EK^* and fake punctured keys $\text{EK}^*\{m\}$, for any message $m \in \{0, 1\}^n$, are of the same size as real keys EK^* , and the corresponding encryption circuit is of the same size:

$$|\text{PFE.Enc}(\text{EK}, \cdot)| = |\text{PFE.Enc}(\text{EK}^*, \cdot)| = |\text{PFE.Enc}(\text{EK}^*\{m\}, \cdot)|.$$

We now bound the size of $\text{PFE.Enc}(\text{EK}, \cdot)$, which is of the form $\text{FE.Enc}(\text{PK}, (\cdot, \text{SK}, \beta); \text{PRF}_K(\cdot))$.

Observe that

- $\text{PRF}_K(\cdot)$, for messages of size n , can be computed by a circuit of size $\Phi_{\text{PRF}}(n, \lambda)$ for a fixed polynomial Φ_{PRF} .
- Given $r = \text{PRF}_K(m)$ as input, $\text{FE.Enc}(\text{PK}, (m, \text{SK}, \beta); r)$ can be computed by a circuit of size $s'^{1-\varepsilon} \Phi_{\text{FE}}(n', \lambda)$, where in PFE , messages m of size n translate to messages (m, SK, β) of size n' and circuits $f(m)$ of size s translate to circuits $F(m, \text{SK}, \beta)$ of size s' in FE .

Observe that $n' = n + \lambda + 1$. It is left to bound s' . F chooses CT_β according to β and applies $U(\text{Sym.Dec}(\text{SK}_\beta, \text{CT}_\beta), m)$. By Fact 2.7, each CT_β encrypting f is of size s and applying Sym.Dec requires size $s \cdot \Phi_{\text{Sym}}(\lambda)$ for a fixed polynomial Φ_{Sym} . Applying a universal circuit for an input circuit of size s requires size $s \cdot \Phi_U(\log s) \leq s \cdot \Phi_U(\lambda)$ for a fixed polynomial Φ_U [75]. Thus, overall,

$$s' \leq s \cdot \Phi_{\text{Sym}, U}(\lambda),$$

for a fixed polynomial $\Phi_{\text{Sym}, U}(\lambda)$ that aggregates the above.

In conclusion, we can bound the size of the encryption circuit $\text{PFE.Enc}(\text{EK}, \cdot)$ by

$$\Phi_{\text{PRF}}(n, \lambda) + s'^{1-\varepsilon} \Phi_{\text{FE}}(n', \lambda) \leq s^{1-\varepsilon} (\Phi_{\text{PRF}}(n, \lambda) + \Phi_{\text{Sym}, U}(\lambda) \cdot \Phi_{\text{FE}}(n + \lambda + 1, \lambda)),$$

establishing our requirement with $\Phi_{\text{PFE}}(n, \lambda) := \Phi_{\text{PRF}}(n, \lambda) + \Phi_{\text{Sym}, U}(\lambda) \cdot \Phi_{\text{FE}}(n + \lambda + 1, \lambda)$. $\square$

3.2 From Puncturable Functional Encryption to Indistinguishability Obfuscation

In this section, we show how to transform any puncturable functional encryption with succinct encryption, such as the one constructed in the previous section, into an indistinguishability obfuscator.

Ingredient. We rely on a puncturable functional encryption scheme,

$$\text{PFE} = (\text{PFE.Setup}, \text{PFE.Enc}, \text{PFE.Dec}, \text{PFE.Setup}^*, \text{PFE.Punc}).$$

Notation: Throughout, bit-strings will be boldfaced, whereas bits will not. In particular, for a string $\mathbf{x} \in \{0, 1\}^n$, we denote by $\mathbf{x}_j$ its j -long prefix and by x_j the j th bit in the string.

Fig. 2. The obfuscator.

Fig. 3. The recursive obfuscation procedure.

Fig. 4. The recursive evaluation procedure.

The Obfuscator. The obfuscator, formally given in Figure 2, is parameterized by a function $\tilde{\lambda}(n, \lambda) \geq \lambda$ (of the circuit input-length and the security parameter), which will define the security parameter with which PFE will be invoked. The security and the complexity of the obfuscator will depend on the choice of the function.

A High-level Description. Given a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and security parameter λ , the obfuscator $iO_{\tilde{\lambda}}(C, 1^{\lambda})$, first computes a new security parameter $\tilde{\lambda} = \tilde{\lambda}(n, \lambda)$, and invokes a recursive obfuscation procedure $rO.Obf(n, C, 1^{\tilde{\lambda}})$, formally described in Figure 3. A corresponding recursive evaluation procedure $rO.Eval$ is described right after in Figure 4.

The recursive obfuscation procedure $rO.Obf(i, C_i, 1^{\tilde{\lambda}})$ extends obfuscation for circuits with $i-1$ input-bits to obfuscation for circuits with i input-bits. To this end, it generates an obfuscation of an encryption circuit $\mathcal{E}_{i-1}$ and a corresponding functional key FSK_i , under the puncturable functional encryption scheme PFE. The encryption circuit $\mathcal{E}_{i-1}$ takes a prefix $\mathbf{x}_{i-1} \in \{0, 1\}^{i-1}$ and generates two encryptions—one for each possible continuation $\mathbf{x}_{i-1}0$ and $\mathbf{x}_{i-1}1$. The corresponding functional key FSK_i allows to evaluate the circuit C_i on the encrypted input.

Unrolling this recursive process, we obtain “a tree of encryptions,” where for each node corresponding to some input-prefix $\mathbf{x}_i$, there is a corresponding encryption $CT_{\mathbf{x}_i}$ of that prefix, and the node’s children can be computed using the corresponding functional key FSK_i . The leaves, correspond to encryptions of the full input $\mathbf{x} = \mathbf{x}_n$, and the last functional key FSK_n corresponds

to the original circuit C . Overall, to evaluate the obfuscation at an input $\mathbf{x}$, one expands the tree along the path corresponding to $\mathbf{x}$, until obtaining $C(\mathbf{x})$ at the final level.

THEOREM 3.3. *Let λ be a security parameter, let n denote circuit-input size, and let $\tilde{\lambda}(n, \lambda)$ be a function. Then if PFE is $\delta(\tilde{\lambda})$ -secure, for security parameter $\tilde{\lambda}$, $iO_{\tilde{\lambda}}$, using security parameter λ , is a $2^{n^2} \delta(\tilde{\lambda})$ -secure indistinguishability obfuscator for all circuits.*

The theorem immediately implies that subexponentially secure puncturable functional encryption (which in turn, can be constructed from subexponentially secure public-key functional encryption) implies an indistinguishability obfuscator.

COROLLARY 3.4. *If PFE is $2^{-\tilde{\lambda}^\alpha}$ -secure, for security parameter $\tilde{\lambda}$, then $iO_{\tilde{\lambda}}$ is $\lambda^{-\omega(1)}$ -secure, for security parameter λ , and any function $\tilde{\lambda}(n, \lambda) = \omega((n^2 + \log \lambda)^{\alpha-1})$.*

Remark 3.5 (Technical Remark on the Proof). One may think of our obfuscator as iteratively applying an input-extension procedure that takes an obfuscator for $i - 1$ bits and produces an obfuscator for i bits. Each such step incurs a certain security loss in the distinguishing gap, which may of course depend on the size of the specific analyzed adversary. We formally analyze the obfuscator as a whole, rather than analyzing a single input-extension step. This allows us to capture the overall loss of an n -step extension process for any specific adversary (or rather any specific adversary size).

PROOF OF THEOREM 3.3. We prove that the constructed scheme is an indistinguishability obfuscator.

Functionality. Correctness follows the intuition outlined above—an evaluation corresponds to gradually constructing an encryption of the input $\mathbf{x} = \mathbf{x}_n$, by using the corresponding function at every level of the tree, and eventually applying the last functional key corresponding to the obfuscated circuit C .

We formally prove correctness by induction on the input length $1 \leq i \leq n$. Specifically, we prove that for any circuit $C_i : \{0, 1\}^i \rightarrow \{0, 1\}^*$, letting $\tilde{\mathcal{E}}_i = rO.\text{Obf}(i, C_i, 1^{\tilde{\lambda}})$, it holds that for any input $\mathbf{x}_i \in \{0, 1\}^i$, $rO.\text{Eval}(i, \tilde{\mathcal{E}}_i, \mathbf{x}_i) = C_i(\mathbf{x}_i)$. Correctness then follows by considering the full input $\mathbf{x} = \mathbf{x}_n$ and the obfuscated circuit $C_n = C$.

For the base case $i = 1$, it holds by definition that $\tilde{\mathcal{E}}_1 = C_1(0), C_1(1)$ and for any $\mathbf{x}_1 \in \{0, 1\}$, $rO.\text{Eval}(1, \tilde{\mathcal{E}}_1, \mathbf{x}_1) = C_1(\mathbf{x}_1)$. We now assume that the required correctness holds for $i - 1$ and show that it holds for i .

By definition,

$$\tilde{\mathcal{E}}_i = rO.\text{Obf}(i, C_i, 1^{\tilde{\lambda}}) = (\tilde{\mathcal{E}}_{i-1}, \text{FSK}_i),$$

where $\tilde{\mathcal{E}}_{i-1} = rO.\text{Obf}(i - 1, \mathcal{E}_{i-1}^\ell, 1^{\tilde{\lambda}})$, $(\text{EK}_i, \text{FSK}_i) \leftarrow \text{PFE.Setup}(1^{\tilde{\lambda}}, C_i)$, and $\mathcal{E}_{i-1}^\ell(\mathbf{x}_{i-1})$ is a circuit that computes two encryptions $\{\text{CT}_{\mathbf{x}_{i-1}b} = \text{PFE.Enc}(\text{EK}_i, \mathbf{x}_{i-1}b)\}_{b \in \{0, 1\}}$ (padded to size $\ell(|\mathcal{E}_{i-1}|)$).

By the induction hypothesis, we know that $rO.\text{Eval}(i - 1, \tilde{\mathcal{E}}_{i-1}, \mathbf{x}_{i-1}) = (\text{CT}_{\mathbf{x}_{i-1}0}, \text{CT}_{\mathbf{x}_{i-1}1})$; namely, the obfuscation $\tilde{\mathcal{E}}_{i-1}$ preserves the functionality of the obfuscated circuit $\tilde{\mathcal{E}}_{i-1}$. It then, follows by definition that

$$rO.\text{Eval}(i, \tilde{\mathcal{E}}_i, \mathbf{x}_i) = \text{PFE.Dec}(\text{FSK}_i, \text{CT}_{\mathbf{x}_i}),$$

which by the functionality of PFE is equal to $C_i(\mathbf{x}_i)$ as required.

Security. Let $s(\lambda), n(\lambda)$ be any two polynomially bounded functions and let $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$ be any polynomial-size distinguisher that works on obfuscations $iO(C, 1^\lambda)$ for circuits C of size $s(\lambda)$, defined on $\{0, 1\}^{n(\lambda)}$.

We aim to bound

$$\begin{aligned}\delta_{iO}(\lambda) &:= \max_{C_0, C_1} \left| \Pr \left[\mathcal{D}(iO_{\tilde{\lambda}}(C_0, 1^\lambda)) = 1 \right] - \Pr \left[\mathcal{D}(iO_{\tilde{\lambda}}(C_1, 1^\lambda)) = 1 \right] \right| \\ &= \max_{C_0, C_1} \left| \Pr \left[\mathcal{D}(rO.\text{Obf}(n, C_0, 1^{\tilde{\lambda}})) = 1 \right] - \Pr \left[\mathcal{D}(rO.\text{Obf}(n, C_1, 1^{\tilde{\lambda}})) = 1 \right] \right|,\end{aligned}$$

where C_0, C_1 are any two circuits defined on $\{0, 1\}^{n(\lambda)}$ of the same functionality and size $s(\lambda)$.

For every $\lambda \in \mathbb{N}$, $1 \leq i \leq n(\lambda)$, and size- s circuit C , we may view

$$rO.\text{Obf}(n, C, 1^{\tilde{\lambda}}) = (rO.\text{Obf}(i, C_i, 1^{\tilde{\lambda}}), z)$$

as consisting of an obfuscation $rO.\text{Obf}(i, C_i, 1^{\tilde{\lambda}})$ for inputs of size i along with additional information z (capturing the functional keys $\text{FSK}_{i+1}, \dots, \text{FSK}_n$). We shall denote by $\zeta_i = \zeta_i(\tilde{\lambda})$ the length of such a z .

We now define $\delta_n(\lambda) := \delta_{iO}(\lambda)$ and for $i < n(\lambda)$,

$$\delta_i(\lambda) := \max_{C_0, C_1, z} \left| \Pr \left[\mathcal{D}(rO.\text{Obf}(i, C_0, 1^{\tilde{\lambda}}), z) = 1 \right] - \Pr \left[\mathcal{D}(rO.\text{Obf}(i, C_1, 1^{\tilde{\lambda}}), z) = 1 \right] \right|,$$

where C_0 and C_1 are any two circuits defined on $\{0, 1\}^i$ of the same size and functionality and $z \in \{0, 1\}^{\zeta_i}$. $\square$

PROPOSITION 3.6. *If PFE is γ -secure, then there exists a function $\delta_{\text{PFE}}(\tilde{\lambda}) = \gamma(\tilde{\lambda})^{\Omega(1)}$ such that for any $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $\tilde{\lambda} = \tilde{\lambda}(n, \lambda)$, $i \in \{2, \dots, n\}$:*

- $\delta_i(\lambda) \leq 2^{i+1} \cdot (\delta_{i-1}(\lambda) + \delta_{\text{PFE}}(\tilde{\lambda}))$.
- $\delta_1(\lambda) = 0$.

Before proving the proposition, we show that it concludes the security analysis:

Claim 3.7. $\delta_n \leq O(2^{n^2} \cdot \delta_{\text{PFE}})$.

PROOF. To prove the claim, we show by induction on i that

$$\delta_i \leq \delta_{\text{PFE}} \sum_{j=1}^{i-1} \prod_{k=0}^{j-1} 2^{i+1-k}.$$

By Proposition 3.6, $\delta_1 = 0$ and thus satisfies the above. Assuming the above holds for $i-1 \geq 1$, and using the proposition again:

$$\begin{aligned}\delta_i &\leq 2^{i+1}(\delta_{i-1} + \delta_{\text{PFE}}) \\ &\leq 2^{i+1}\delta_{\text{PFE}} + 2^{i+1}\delta_{\text{PFE}} \sum_{j=1}^{i-2} \prod_{k=0}^{j-1} 2^{i-k} \\ &= 2^{i+1}\delta_{\text{PFE}} + \delta_{\text{PFE}} \sum_{j=1}^{i-2} \prod_{k=0}^j 2^{i+1-k} \\ &= 2^{i+1}\delta_{\text{PFE}} + \delta_{\text{PFE}} \sum_{j=2}^{i-1} \prod_{k=0}^{j-1} 2^{i+1-k} \\ &= \delta_{\text{PFE}} \sum_{j=1}^{i-1} \prod_{k=0}^{j-1} 2^{i+1-k}.\end{aligned}$$

We deduce accordingly,

$$\delta_n \leq \delta_{\text{PFE}} \sum_{j=1}^{n-1} \prod_{k=0}^{j-1} 2^{n+1-k} \leq \delta_{\text{PFE}} (n-1) 2^{\sum_{j=3}^{n+1} j} \leq O(\delta_{\text{PFE}} 2^{n^2}). \quad \square$$

We now turn to prove the proposition.

PROOF OF PROPOSITION 3.6. First, to see that $\delta_1 = 0$, note that for any C defined on $\{0, 1\}$,

$$r\mathcal{O}.\text{Obf}(1, C, 1^{\tilde{\lambda}}) = (C(0), C(1)),$$

by definition, and thus for any two C_0, C_1 with the same functionality,

$$r\mathcal{O}.\text{Obf}(1, C_0, 1^{\tilde{\lambda}}) \equiv r\mathcal{O}.\text{Obf}(1, C_1, 1^{\tilde{\lambda}}).$$

We now prove the main part of the proposition. Fix $i \in \{2, \dots, n(\lambda)\}$, and let C_0, C_1 be any two circuits defined on $\{0, 1\}^i$ of equal size and fix any auxiliary input $z \in \{0, 1\}^{\ell_i}$. We will bound

$$\left| \Pr \left[\mathcal{D}(r\mathcal{O}.\text{Obf}(i, C_0, 1^{\tilde{\lambda}}), z) = 1 \right] - \Pr \left[\mathcal{D}(r\mathcal{O}.\text{Obf}(i, C_1, 1^{\tilde{\lambda}}), z) = 1 \right] \right|.$$

For this purpose, we consider the following sequence hybrid experiments.

$\mathcal{H}_b$ for $b \in \{0, 1\}$: This hybrid corresponds to the (real) distribution, where

$$r\mathcal{O}.\text{Obf}(i, C_b, 1^{\tilde{\lambda}}) = (\tilde{\mathcal{E}}, \text{FSK}),$$

where $\tilde{\mathcal{E}} = r\mathcal{O}.\text{Obf}(i-1, \mathcal{E}^\ell[\text{EK}], 1^{\tilde{\lambda}})$, $(\text{EK}, \text{FSK}) \leftarrow \text{PFE}.\text{Setup}(1^{\tilde{\lambda}}, C_b)$, and $\mathcal{E}^\ell[\text{EK}](\mathbf{x})$ is a circuit that for any $\mathbf{x} \in \{0, 1\}^{i-1}$ computes two encryptions,

$$\mathcal{E}[\text{EK}](\mathbf{x}) = \{\text{PFE}.\text{Enc}(\text{EK}, \mathbf{x}\beta)\}_{\beta \in \{0, 1\}},$$

padded to size $\ell(|\mathcal{E}[\text{EK}]|)$.

$\mathcal{H}_b^*$ for $b \in \{0, 1\}$: This hybrid is identical to $\mathcal{H}_b$, except that we sample fake keys $(\text{EK}_0^*, \text{EK}_1^*, \text{FSK}^*) \leftarrow \text{PFE}.\text{Setup}^*(1^{\tilde{\lambda}}, C_0, C_1)$ and use $\text{EK}_b^*, \text{FSK}^*$ instead of EK and FSK.

$\mathcal{G}_y^*$ for $y \in [2^i + 1]$: This hybrid is identical to $\mathcal{H}_0^*$, except that $\tilde{\mathcal{E}}$ is an obfuscation of a hybrid circuit $\mathcal{E}_y[\text{EK}_0^*, \text{EK}_1^*]$ instead of the circuit $\mathcal{E}[\text{EK}_0^*]$. The hybrid circuit computes two encryptions, but uses EK_0^* for all plaintexts $\mathbf{x}\beta \geq y$ and EK_1^* for all plaintexts $\mathbf{x}\beta < y$, where we naturally interpret $\mathbf{x}\beta$ as an integer in $[2^i]$ according to lexicographic order.

That is, for any $\mathbf{x} \in \{0, 1\}^{i-1}$, the circuit $\mathcal{E}_y[\text{EK}_0^*, \text{EK}_1^*]$ is defined by

$$\mathcal{E}_y[\text{EK}_0^*, \text{EK}_1^*](\mathbf{x}) = \left\{ \begin{array}{ll} \text{PFE}.\text{Enc}(\text{EK}_0^*, \mathbf{x}\beta), & \text{if } \mathbf{x}\beta \geq y; \\ \text{PFE}.\text{Enc}(\text{EK}_1^*, \mathbf{x}\beta), & \text{if } \mathbf{x}\beta < y. \end{array} \right\}_{\beta \in \{0, 1\}}$$

and is padded to size $\ell = \ell(|\mathcal{E}[\text{EK}]|)$.

$\mathcal{G}_{y,b}^*$ for $y \in [2^i + 1], b \in \{0, 1\}$: This hybrid is identical to $\mathcal{G}_y^*$, except that $\tilde{\mathcal{E}}$ is an obfuscation of a hybrid circuit $\mathcal{E}_{y,b}^\ell[\text{EK}_0^*\{y\}, \text{EK}_1^*\{y\}, \text{CT}_{y,b}]$ instead of the circuit $\mathcal{E}_y[\text{EK}_0^*, \text{EK}_1^*]$. Here the keys $(\text{EK}_0^*\{y\}, \text{EK}_1^*\{y\}) = \text{PFE}.\text{Punc}(\text{EK}_0^*, \text{EK}_1^*, y)$ are punctured at y , and the ciphertext $\text{CT}_{y,b} = \text{PFE}.\text{Enc}(\text{EK}_b^*, y)$ is hardwired and consists of an encryption of y under EK_b^* . The circuit behaves similarly to $\mathcal{E}_y[\text{EK}_0^*, \text{EK}_1^*]$, only that for all $y' \neq y$ it uses the punctured keys, whereas for y it outputs the hardwired ciphertext $\text{CT}_{y,b}$.

That is, for any $\mathbf{x} \in \{0, 1\}^{i-1}$, the circuit $\mathcal{E}_{y,b}^\ell[\text{EK}_0^\star\{\mathbf{y}\}, \text{EK}_1^\star\{\mathbf{y}\}, \text{CT}_{y,b}]$ is defined by

$$\mathcal{E}_{y,b}[\text{EK}_0^\star\{\mathbf{y}\}, \text{EK}_1^\star\{\mathbf{y}\}, \text{CT}_{y,b}](\mathbf{x}) = \begin{cases} \text{PFE.Enc}(\text{EK}_0^\star\{\mathbf{y}\}, \mathbf{x}\beta), & \text{if } \mathbf{x}\beta > \mathbf{y}; \\ \text{CT}_{y,b}, & \text{if } \mathbf{x}\beta = \mathbf{y}; \\ \text{PFE.Enc}(\text{EK}_1^\star\{\mathbf{y}\}, \mathbf{x}\beta), & \text{if } \mathbf{x}\beta < \mathbf{y}. \end{cases}_{\beta \in \{0,1\}}$$

and is padded to size $\ell = \ell(|\mathcal{E}[\text{EK}]|)$. $\square$

CLAIM 3.8. Assuming that PFE is γ -secure, there exists function $\delta_{\text{PFE}}(\tilde{\lambda}) = \gamma(\tilde{\lambda})^{\Omega(1)}$ such that for any $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $\tilde{\lambda} = \tilde{\lambda}(\lambda, n)$, any $i \in [n]$, and any two equal-size circuits $C_0, C_1 : \{0, 1\}^i \rightarrow \{0, 1\}^*$:

- (1) $\mathcal{H}_0 \approx_{\mathcal{D}, \delta_{\text{PFE}}} \mathcal{H}_0^\star$.
- (2) $\mathcal{H}_0^\star = \mathcal{G}_1^\star$.
- (3) $\mathcal{G}_y^\star \approx_{\mathcal{D}, \delta_{i-1}} \mathcal{G}_{y,0}^\star$, for all $\mathbf{y} \in [2^i + 1]$.
- (4) $\mathcal{G}_{y,0}^\star \approx_{\mathcal{D}, \delta_{\text{PFE}}} \mathcal{G}_{y,1}^\star$, for all $\mathbf{y} \in [2^i + 1]$.
- (5) $\mathcal{G}_{y,1}^\star \approx_{\mathcal{D}, \delta_{i-1}} \mathcal{G}_{y+1}^\star$, for all $\mathbf{y} \in [2^i]$.
- (6) $\mathcal{G}_{2^i+1}^\star = \mathcal{H}_1^\star$.
- (7) $\mathcal{H}_1^\star \approx_{\mathcal{D}, \delta_{\text{PFE}}} \mathcal{H}_1$.

Before proving the claim, we show that it concludes the proof of Proposition 3.6:

$$\begin{aligned} \delta_i &= \left| \Pr[\mathcal{D}(\mathcal{H}_0) = 1] - \Pr[\mathcal{D}(\mathcal{H}_1) = 1] \right| \\ &= \left| \Pr[\mathcal{D}(\mathcal{H}_0^\star) = 1] - \Pr[\mathcal{D}(\mathcal{H}_1^\star) = 1] \right| + \sum_{b \in \{0,1\}} \left| \Pr[\mathcal{D}(\mathcal{H}_b) = 1] - \Pr[\mathcal{D}(\mathcal{H}_b^\star) = 1] \right| \\ &= \left| \Pr[\mathcal{D}(\mathcal{G}_1^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{2^i}^\star) = 1] \right| + \sum_{b \in \{0,1\}} \left| \Pr[\mathcal{D}(\mathcal{H}_b) = 1] - \Pr[\mathcal{D}(\mathcal{H}_b^\star) = 1] \right| \\ &\leq \sum_{y \in [2^i]} \left| \Pr[\mathcal{D}(\mathcal{G}_y^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{y+1}^\star) = 1] \right| + \sum_{b \in \{0,1\}} \left| \Pr[\mathcal{D}(\mathcal{H}_b) = 1] - \Pr[\mathcal{D}(\mathcal{H}_b^\star) = 1] \right| \\ &\leq \sum_{y \in [2^i]} \left(\left| \Pr[\mathcal{D}(\mathcal{G}_y^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{y,0}^\star) = 1] \right| \right. \\ &\quad \left. + \left| \Pr[\mathcal{D}(\mathcal{G}_{y,0}^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{y,1}^\star) = 1] \right| \right. \\ &\quad \left. + \left| \Pr[\mathcal{D}(\mathcal{G}_{y,1}^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{y+1}^\star) = 1] \right| \right) + \sum_{b \in \{0,1\}} \left| \Pr[\mathcal{D}(\mathcal{H}_b) = 1] - \Pr[\mathcal{D}(\mathcal{H}_b^\star) = 1] \right| \\ &\leq 2^i(2\delta_{i-1} + \delta_{\text{PFE}}) + 2\delta_{\text{PFE}} = 2^{i+1}\delta_{i-1} + (2^i + 2)\delta_{\text{PFE}} \leq 2^{i+1}(\delta_{i-1} + \delta_{\text{PFE}}). \end{aligned}$$

PROOF OF CLAIM 3.8. The indistinguishability between the hybrids is established by applying the γ -security of the underlying puncturable functional encryption PFE and the bound δ_{i-1} on the distinguishing gap for circuits on $i-1$ input bits.

$\mathcal{H}_b \approx \mathcal{H}_b^\star$ for $b \in \{0, 1\}$ (items 1, 7): Here, the difference between the hybrids is that in $\mathcal{H}_b$ we sample real keys $(\text{EK}, \text{FSK}) \leftarrow \text{PFE.Setup}(1^{\tilde{\lambda}}, C_b)$, whereas in $\mathcal{H}_b^\star$ we sample fake keys $(\text{EK}_0^\star, \text{EK}_1^\star, \text{FSK}^\star) \leftarrow \text{PFE.Setup}^\star(1^{\tilde{\lambda}}, C_0, C_1)$ and use $\text{EK}_b^\star$. By the key indistinguishability of PFE,

$$\text{EK}, \text{FSK} \approx_\gamma \text{EK}_b^\star, \text{FSK}^\star.$$

Thus,

$$\mathcal{H}_b \approx_{\mathcal{D}, \gamma^\alpha} \mathcal{H}_b^\star,$$

for some constant $\alpha \leq 1$ that depends only on $\mathcal{D}$.

$\mathcal{H}_0^\star = \mathcal{G}_1^\star, \mathcal{G}_{2^{i+1}}^\star = \mathcal{H}_1^\star$ (items 2, 6): The equality between the hybrids follows by their definition. Indeed, in $\mathcal{H}_0^\star$ we always use $\text{EK}_0^\star$ as in $\mathcal{G}_1^\star$, and in $\mathcal{H}_1^\star$ we always use $\text{EK}_1^\star$ as in $\mathcal{G}_{2^{i+1}}^\star$.

$\mathcal{G}_{y,b}^\star \approx \mathcal{G}_{y+b}^\star$ for $b \in \{0, 1\}$ (items 3, 5): Here the difference between $\mathcal{G}_{y+b}^\star$ and $\mathcal{G}_{y,b}^\star$ is in the obfuscated circuit $\tilde{\mathcal{E}}$. In the first, $\tilde{\mathcal{E}} = r\mathcal{O}.\text{Obf}(i-1, \mathcal{E}_{y+b}^\ell[\text{EK}_0^\star, \text{EK}_1^\star], 1^{\tilde{\lambda}})$, whereas in the second, $\tilde{\mathcal{E}} = r\mathcal{O}.\text{Obf}(i-1, \mathcal{E}_{y+b,b}^\ell[\text{EK}_0^\star, \text{EK}_1^\star, \text{CT}_{y+b,b}], 1^{\tilde{\lambda}})$.

The two circuits compute the exact same function, but in two different ways:

- The first circuit uses the keys $(\text{EK}_0^\star, \text{EK}_1^\star)$ to encrypt all plaintexts $\mathbf{x} \neq \mathbf{y} + b$, whereas the second circuit uses the punctured keys $(\text{EK}_0^\star\{\mathbf{y} + b\}, \text{EK}_1^\star\{\mathbf{y} + b\})$. By punctured-key correctness of PFE, the two result in the same ciphertexts.
- For the plaintext $\mathbf{y} + b$, the first circuit computes $\text{CT}_{y+b,b} = \text{PFE}.\text{Enc}(\text{EK}_b, \mathbf{y} + b)$ on its own using the key $\text{EK}_b^\star$, whereas the second circuit uses a hardwired $\text{CT}_{y+b,b}$.

The two circuits are padded to the same size $\ell = \ell(|\mathcal{E}[\text{EK}]|)$. We can thus bound

$$\begin{aligned} & \left| \Pr[\mathcal{D}(\mathcal{G}_{y+b}^\star) = 1] - \Pr[\mathcal{D}(\mathcal{G}_{y,b}^\star) = 1] \right| \\ &= \left| \Pr[\mathcal{D}(r\mathcal{O}.\text{Obf}(i-1, \mathcal{E}_{y+b}^\ell, \text{FSK}^\star, z) = 1] - \Pr[\mathcal{D}(r\mathcal{O}.\text{Obf}(i-1, \mathcal{E}_{y,b}^\ell, \text{FSK}^\star, z) = 1] \right| \\ &\leq \max_{C'_0, C'_1, z'} \left| \Pr[\mathcal{D}(r\mathcal{O}.\text{Obf}(i-1, C'_0, 1^{\tilde{\lambda}}), z') = 1] - \Pr[\mathcal{D}(r\mathcal{O}.\text{Obf}(i-1, C'_1, 1^{\tilde{\lambda}}), z') = 1] \right| = \delta_{i-1}, \end{aligned}$$

where above C_0 and C_1 are any two circuits defined on $\{0, 1\}^{i-1}$ of the same functionality and size ℓ , and we view $(z, \text{FSK}^\star)$ as $z' \in \{0, 1\}^{\zeta_{i-1}}$.

It follows that

$$\mathcal{G}_{y+b}^\star \approx_{\mathcal{D}, \delta_{i-1}} \mathcal{G}_{y,b}^\star.$$

$\mathcal{G}_{y,0}^\star \approx \mathcal{G}_{y,1}^\star$ (item 4): Here the difference between the hybrids is in the hardwired ciphertext $\text{CT}_{y,b} = \text{PFE}.\text{Enc}(\text{EK}_b, \mathbf{y})$, where $b = 0$ in the first hybrid and $b = 1$ in the second. By indistinguishability at punctured points of PFE:

$$\text{EK}_0^\star\{\mathbf{y}\}, \text{EK}_1^\star\{\mathbf{y}\}, \text{FSK}^\star, \text{CT}_{y,0} \approx_\delta \text{EK}_0^\star\{\mathbf{y}\}, \text{EK}_1^\star\{\mathbf{y}\}, \text{FSK}^\star, \text{CT}_{y,1}.$$

Thus,

$$\mathcal{G}_{y,0}^\star \approx_{\gamma^{\alpha'}} \mathcal{G}_{y,1}^\star.$$

for some constant $\alpha' \leq 1$ that depends only on $\mathcal{D}$.

The Function δ_{PFE} . The function $\delta_{\text{PFE}}(\tilde{\lambda}) = \gamma(\tilde{\lambda})^{\min\{\alpha, \alpha'\}}$ thus satisfies the requirement of the Claim 3.8. This concludes the proof of the claim. $\square$

The Padding Parameter. The padding parameter $\ell(|\mathcal{E}|)$ is chosen to account for the maximal-size circuit among the circuits $\mathcal{E}, \mathcal{E}_y, \mathcal{E}_{y,b}$ considered in the above hybrids. Observe that the size of the circuits $\mathcal{E}_y, \mathcal{E}_{y,b}$ is indeed linear in the size of the (real or fake) encryption circuits

$$|\text{PFE}.\text{Enc}(\text{EK}_b^\star, \cdot)| = |\text{PFE}.\text{Enc}(\text{EK}_b^\star\{\mathbf{y}\}, \cdot)| = |\text{PFE}.\text{Enc}(\text{EK}, \cdot)| = |\mathcal{E}|.$$

This concludes the proof of Proposition 3.6 and the proof of security. $\square$

Efficiency. We now analyze the efficiency of the obfuscator relying on succinct encryption of the underlying PFE. Assume that PFE is succinct. That is, there exists a polynomial Φ_{PFE} and a constant $0 < \varepsilon \leq 1$ such that the size of the corresponding encryption circuit is bounded by

$$s^{1-\varepsilon} \cdot \Phi_{\text{PFE}}(n, \lambda),$$

where n, s are the input-size and circuit-size of the functions chosen during the (possibly fake) setup phase.

We show that for any polynomial $\tilde{\lambda}(n, \lambda) = \text{poly}(n, \lambda)$, there exists a fixed polynomial Φ_{iO} such that $iO_{\tilde{\lambda}}(C, 1^\lambda)$ runs in time $\Phi_{iO}(|C|, \lambda)$. First, observe that the recursive $rO.\text{Obf}(i, C_i, 1^\lambda)$ only invokes the efficient PFE algorithms, and thus there exists a polynomial Φ_r such that for all $\lambda, n \in \mathbb{N}, i \in [n]$, the running time of $rO.\text{Obf}(i, C_i, 1^\lambda)$ is bounded by $\Phi_r(|C_i|, \tilde{\lambda}|i|)$.

We prove the following:

CLAIM 3.9. *Let $\ell(s) = c \cdot s$ be the padding function and let $\Phi_{\text{PFE}} = \Phi_{\text{PFE}}(n, \tilde{\lambda})$. Then, for all $i \in [n]$,*

$$|C_i| \leq |C| \cdot (c\Phi_{\text{PFE}})^{1/\varepsilon}.$$

Before proving the claim, we show that it concludes the efficiency analysis. Indeed, it implies that the total running time of the obfuscator is bounded by a fixed polynomial,

$$\Phi_{iO}(|C|, \lambda) \leq n \cdot \max_{i \in [n]} \Phi_r(|C_i|, \tilde{\lambda}(n, \lambda)) \leq |C| \cdot \Phi_r(|C| \cdot (c\Phi_{\text{PFE}})^{1/\varepsilon}, \tilde{\lambda}(|C|, \lambda)).$$

PROOF OF CLAIM 3.9. We prove by induction on i that

$$|C_{n-i}| \leq |C| \cdot (c\Phi_{\text{PFE}})^{\sum_{j=0}^{i-1} (1-\varepsilon)^j}.$$

This concludes the proof, since $\sum_{j=0}^{\infty} (1-\varepsilon)^j = \varepsilon^{-1}$.

For $i = 0$, $C_n = C$ and the bound holds.

Assume the bounds hold for $i > 0$. By the induction hypothesis, the succinctness of PFE, and the definition of $rO.\text{Obf}$, it holds that for any $i > 0$,

$$\begin{aligned} |C_{n-i}| &= |\mathcal{E}_i^\ell| = \ell(|\mathcal{E}_i|) \leq c \cdot |C_{n-i+1}|^{1-\varepsilon} \Phi_{\text{PFE}} \\ &\leq |C|^{1-\varepsilon} (c\Phi_{\text{PFE}})^{(1-\varepsilon) \sum_{j=0}^{i-2} (1-\varepsilon)^j} c\Phi_{\text{PFE}} \\ &\leq |C| \cdot (c\Phi_{\text{PFE}})^{\sum_{j=0}^{i-1} (1-\varepsilon)^j}. \end{aligned} \quad \square$$

This concludes the proof of Theorem 3.3.

3.3 IO with Linear Overhead

In this section, we observe that a variant of our construction, combined with known results from the literature, implies that any IO scheme can be turned into an IO scheme where the obfuscation only grows linearly in the size of the circuit. Specifically, concretely, under the Learning with Errors (LWE) Assumption, we show that any IO can be transformed into one where an obfuscation of a size- s , depth- d circuit, with input length n , is of size $2s + \text{poly}(n, d, \lambda)$.

This is achieved by relying on a single application of the recursive obfuscator (instead of n applications) and instantiating our puncturable functional encryption with a public-key functional encryption scheme with succinct keys as the one constructed by Boneh et al. [19]:

PROPOSITION 3.10 (FE WITH SUCINCT KEYS [19]). *Assuming subexponential LWE, there exists a subexponentially secure, single-key, public-key, functional encryption scheme, where there exists a polynomial Φ_{FE} such that for any depth- d circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, letting $(\text{PK}, \text{FSK}_C) \leftarrow \text{FE.Setup}(1^\lambda, C)$, it holds that:*

- $\text{FSK}_C = (C, \text{fsk})$.
- The key fsk and the encryption circuit $\text{FE.Enc}(\text{PK}, \cdot)$ are of size at most $\Phi_{\text{FE}}(n, m, d, \lambda)$.

We prove:

COROLLARY 3.11. *Assuming subexponential LWE and subexponentially secure IO, there exists IO where there exists a polynomial Φ , such that given any size- s , depth- d circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, a corresponding obfuscation is of size $2s + \Phi(n, m, d, \lambda)$.*

PROOF. First, we observe that by plugging-in the FE scheme given by Proposition 3.10 into the construction from Section 3.1, we obtain a puncturable FE that inherits the succinctness of the underlying FE scheme up to a mild loss. $\square$

CLAIM 3.12. *Assuming subexponential LWE, there exists a subexponentially secure puncturable functional encryption scheme, where there exists a polynomial Φ_{PFE} such that for any depth- d circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, letting $(\text{EK}, \text{FSK}_C) \leftarrow \text{PFE.Setup}(1^\lambda, C)$, it holds that:*

- FSK_C is of size at most $2|C| + \Phi_{\text{PFE}}(n, m, d, \lambda)$.
- The encryption circuit $\text{PFE.Enc}(\text{EK}, \cdot)$ is of size at most $\Phi_{\text{PFE}}(n, m, d, \lambda)$.

PROOF. Let FE be the scheme given by Proposition 3.10 and let PFE^{FE} be the puncturable functional encryption scheme from Section 3.1 when instantiated with FE. In PFE^{FE} , a functional key FSK for a circuit C consists of a functional key FSK_F generated under FE for a circuit $F = F[\text{CT}_0, \text{CT}_1]$ that is parameterized by two symmetric-key encryptions CT_0, CT_1 of C . The circuit F can be efficiently constructed given only CT_0, CT_1 . Also, recall that in the given FE scheme, FSK has the form (F, fsk) where fsk is of size $\Phi(n, m, d, \lambda)$.

We now consider a slight variant of PFE^{FE} , which we shall denote by

$$\widehat{\text{PFE}}^{\text{FE}} = (\text{PFE.Setup}, \text{PFE.Enc}, \text{PFE.Dec}, \text{PFE.Setup}^*, \text{PFE.Punc}),$$

where the functional key $\widehat{\text{FSK}}$ for C consists of $(\text{CT}_0, \text{CT}_1, \text{fsk})$ and the functional decryption process $\text{PFE.Dec}(\widehat{\text{FSK}}, \text{CT})$ first parses $\widehat{\text{FSK}} = (\text{CT}_0, \text{CT}_1, \text{fsk})$, then computes the secret key $\text{FSK} = (F[\text{CT}_0, \text{CT}_1], \text{fsk})$, and finally applies $\text{FE.Dec}(\text{FSK}, \text{CT})$.

Correctness and Security. First, observe that the scheme $\widehat{\text{PFE}}^{\text{FE}}$ satisfies the same functionality and security properties of PFE^{FE} (proven in Section 3.1). Indeed, the two schemes only defer in the format of functional keys and a functional key FSK under PFE^{FE} can be efficiently simulated from a functional key $\widehat{\text{FSK}}$ under $\widehat{\text{PFE}}^{\text{FE}}$.

Function-key Succinctness. We now analyze the succinctness of $\widehat{\text{PFE}}^{\text{FE}}$. By the above construction, the size of a functional key for a circuit C is bounded by

$$|\widehat{\text{FSK}}| = 2|\text{CT}_0| + |\text{fsk}| = 2|C| + \Phi_{\text{FE}}(n, m, d, \lambda),$$

where we use the fact that in (one-time) symmetric-key encryption schemes, the ciphertext-size equals the message-size (Fact 2.7).

Encryption Succinctness. The difference from the succinctness analysis in Section 3.1 is that there the size of the encryption circuit of FE grows (sublinearly) with circuit-size, whereas here it does not grow with the circuit size, but does grow (polynomially) with the output length and depth.

The encryption circuit in $\widehat{\text{PFE}}^{\text{FE}}$ has the form

$$\text{PFE.Enc}(\text{EK}, \cdot) = \text{FE.Enc}(\text{PK}, (\cdot, \text{SK}, \beta); \text{PRF}_K(\cdot)).$$

Observe that

- $\text{PRF}_K(\cdot)$, for messages of size n , can be computed by a circuit of size $\Phi_{\text{PRF}}(n, \lambda)$ for a fixed polynomial Φ_{PRF} .

- Given $r = \text{PRF}_K(m)$ as input, $\text{FE.Enc}(\text{PK}, (m, \text{SK}, \beta); r)$ can be computed by a circuit of size $\Phi_{\text{FE}}(n', m', d', \lambda)$, where in PFE, messages m of size n translate to messages (m, SK, β) of size n' and circuits $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ of size s and depth d translate to circuits $F : \{0, 1\}^{n'} \rightarrow \{0, 1\}^{m'}$ of size s' and depth d' in FE.

Observe that $n' = n + \lambda + 1$ and $m' = m$. It is left to bound s' (here, it suffices that $s' = s^{O(1)}$) and d' . F chooses CT_β according to β and applies $U(\text{Sym.Dec}(\text{SK}_\beta, \text{CT}_\beta), m)$. By Fact 2.7, each CT_β encrypting f is of size s and applying Sym.Dec requires size $s \cdot \Phi_{\text{Sym}}(\lambda)$ and depth $\Phi_{\text{Sym}}(\lambda)$ for a fixed polynomial Φ_{Sym} . Applying a depth-universal circuit for an input circuit of size s and depth d requires size $s^3 \cdot \Phi_U(\log s) \leq s^3 \cdot \Phi_U(\lambda)$ and depth $c_U \cdot d$ for a fixed polynomial Φ_U and constant c_U [33]. Thus, overall, there exists a polynomial Φ such that

$$s' \leq \Phi(s, \lambda), \quad d' \leq \Phi(d, \lambda),$$

for a fixed polynomial Φ that aggregates the above.

In conclusion, we can bound the size of the encryption circuit $\text{PFE.Enc}(\text{EK}, \cdot)$ by

$$\Phi_{\text{PRF}}(n, \lambda) + \Phi_{\text{FE}}(n', m', d', \lambda) \leq \Phi_{\text{PRF}}(n, \lambda) + \Phi_{\text{FE}}(n + \lambda + 1, m, \Phi(d, \lambda), \lambda),$$

establishing our requirement with $\Phi_{\text{PFE}}(n, m, d, \lambda) := \Phi_{\text{PRF}}(n, \lambda) + \Phi_{\text{FE}}(n + \lambda + 1, m, \Phi(d, \lambda), \lambda)$. $\square$

We next show that any $2^{\tilde{\lambda}^\alpha}$ -secure indistinguishability obfuscator iO can be combined with a $2^{\tilde{\lambda}^\alpha}$ -secure puncturable functional encryption $\widehat{\text{PFE}}$ as the one given by Claim 3.12 (where in both $\tilde{\lambda}$ is the security parameter) to obtain a new indistinguishability obfuscator iO' where the size of an obfuscation of a size s depth d circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is of size $2s + \Phi(n, m, d, \lambda)$ for a fixed polynomial Φ .

The Construction. We consider a single iteration of the recursive transformation from Section 3.2. That is, we define

$$iO'(C, 1^\lambda) = rO.\text{Obf}(n, C, 1^{\tilde{\lambda}}),$$

where $rO.\text{Obf}(n, \cdot, 1^{\tilde{\lambda}})$ is defined as in Section 3.1 and is instantiated with $\widehat{\text{PFE}}$ (from Claim 3.12), and redefine

$$rO.\text{Obf}(n - 1, \cdot, 1^{\tilde{\lambda}}) = iO(\cdot, 1^{\tilde{\lambda}}).$$

In addition, we define $\tilde{\lambda}(\lambda, n) = \omega(n + \log \lambda)^{\alpha-1}$.

The evaluation procedure is changed accordingly. That is, given an obfuscation $\tilde{C}$, we define

$$iO'.\text{Eval}(\tilde{C}, \mathbf{x}) = rO.\text{Eval}(n, \tilde{C}, \mathbf{x}),$$

where $rO.\text{Eval}(n, \tilde{C}, \mathbf{x})$ is defined as in Section 3.1 and is instantiated with $\widehat{\text{PFE}}$ (from Claim 3.12), and redefine

$$rO.\text{Eval}(n - 1, \tilde{E}, \mathbf{x}_{n-1}) = iO.\text{Eval}(\tilde{E}, \mathbf{x}_{n-1}).$$

The correctness of the scheme follows readily, we focus on succinctness and security.

Efficiency and Succinctness. In the new scheme, an obfuscation

$$iO'(C, 1^\lambda) = (\tilde{\mathcal{E}}, \widehat{\text{FSK}})$$

consists of a functional key $\widehat{\text{FSK}}$ and an obfuscation $\tilde{\mathcal{E}} \leftarrow iO(\mathcal{E}^\ell, 1^{\tilde{\lambda}})$ of the circuit $\mathcal{E}(\cdot) = \text{PFE.Enc}(\text{EK}, \cdot)$ padded to size $\ell(|\mathcal{E}|)$, where $(\text{EK}, \widehat{\text{FSK}}) \leftarrow \text{PFE.Setup}^*(1^\lambda, C)$ and $\ell(s) = c \cdot s$, for an absolute constant c .

By the succinctness of PFE and the efficiency of iO :

$$\begin{aligned} |iO'(C, 1^\lambda)| &= |\widehat{\text{FSK}}| + |iO(\mathcal{E}^\ell, 1^{\tilde{\lambda}})| \\ &\leq 2|C| + \Phi_{\text{PFE}}(n, m, d, \tilde{\lambda}) + \Phi_{iO}(|\mathcal{E}^\ell|) \\ &\leq 2|C| + \Phi_{\text{PFE}}(n, m, d, \tilde{\lambda}) + \Phi_{iO}(c \cdot \Phi_{\text{PFE}}(n, m, d, \tilde{\lambda})), \end{aligned}$$

where Φ_{iO} is a fixed polynomial that bounds the running time of iO .

This establishes the succinctness requirement given by Corollary 3.4, with

$$\Phi(n, m, d, \lambda) := \Phi_{\text{PFE}}(n, m, d, \tilde{\lambda}(n, \lambda)) + \Phi_{iO}(c \cdot \Phi_{\text{PFE}}(n, m, d, \tilde{\lambda}(n, \lambda))).$$

The fact that the obfuscator runs in polynomial time follows readily from the fact that all underlying algorithms run in polynomial time and are applied at most once.

Security. The security of the scheme follows from the security proof of the recursive scheme in Section 3.2. Specifically, for any polynomial-size distinguisher $\mathcal{D} = \{\mathcal{D}_\lambda\}_{\lambda \in \mathbb{N}}$, we define $\delta_n(\lambda)$ and $\delta_{n-1}(\lambda)$ as in Section 3.2. (The first is the maximal distinguishing advantage of $\mathcal{D}$ against the constructed obfuscator iO' and the second is against $rO.\text{Obf}(n-1, \cdot, 1^{\tilde{\lambda}})$, which in our case is the underlying obfuscator $iO(\cdot, 1^{\tilde{\lambda}})$).

By Proposition 3.6, the subexponential-security of PFE and iO , and our choice of $\tilde{\lambda}(n, \lambda)$,

$$\begin{aligned} \delta_n(\lambda) &\leq 2^{n+1} \cdot (\delta_{n-1}(\lambda) + \delta_{\text{PFE}}(\tilde{\lambda})) \\ &\leq 2^{n+1} \cdot (2^{\Omega(-\tilde{\lambda}^\alpha)} + 2^{\Omega(-\tilde{\lambda}^\alpha)}) \leq 2^{-\omega(\log \lambda)}. \end{aligned}$$

□

4 A BOOTSTRAPPING THEOREM

In this section, we show how to transform any *multi-key* functional encryption scheme with certain weak succinctness into a succinct single-key functional encryption scheme (which, in particular, is suitable for our FE to IO transformation from Section 3).

The notion of multi-key functional encryption schemes that we consider is a natural generalization of the single-key notion (Section 2.2), where security is guaranteed even in the presence of multiple functional keys. In the literature this notion is commonly referred to as *collusion resistance*. We in fact consider a definition that is somewhat weaker than the typical definition in the literature where all functions are known at setup time. (Note that since we aim to construct succinct single-key functional encryption from multi-key functional encryption, considering such a weaker notion only strengthens the result.)

In terms of succinctness, the common requirement for such schemes is that the complexity of encryption may grow with the circuit-size of functions, but not with the overall number of keys generated. Here, we will even allow sublinear dependence on the number of functions (we call this requirement *weak size-succinctness*). We show that any such scheme can be transformed to a succinct single-key scheme.

The transformation can be applied to functional encryption schemes from the literature, such as the one by Garg, Gentry, Halevi, and Zhandry [37]. This gives rise to an IO construction based on a subexponential variant of the assumptions in Reference [37] on multi-linear graded encodings.

We further show that the same transformation can be applied to transform any succinct single/multi-key scheme where encryption complexity scales (polynomially/exponentially) with *circuit depth* into a succinct scheme with no dependence on the depth (beyond the sublinear dependence on circuit size).

We start with the relevant definition and then proceed to the transformation.

4.1 Multi-Key FE with Succinct Encryption

A multi-key FE scheme MFE, for a function class $\mathcal{F}$ (represented by boolean circuits) and message space $\{0, 1\}^*$, consists of three PPT algorithms (MFE.Setup, MFE.Enc, MFE.Dec) with the following syntax:

- MFE.Setup($1^\lambda, f_1, \dots, f_\ell$): takes as input a security parameter λ in unary and ℓ functions $f_1, \dots, f_\ell \in \mathcal{F}$ and outputs a public key PK and functional keys $\text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}$.
- MFE.Enc(PK, m): takes as input a public key PK and a message $m \in \{0, 1\}^*$ and outputs an encryption of m . We shall sometimes address the randomness r used in encryption explicitly, which we denote by MFE.Enc(PK, $m; r$).
- MFE.Dec(FSK_f, CT): takes as input a functional key FSK_f , a ciphertext CT, and outputs $\hat{m}$.

We next define the required correctness, security, and efficiency properties.

Definition 4.1 (Multi-key, Selectively Secure, Public-Key FE with Succinct Encryption). A tuple of PPT algorithms $\text{MFE} = (\text{MFE.Setup}, \text{MFE.Enc}, \text{MFE.Dec})$ is a multi-key, selectively secure, public-key functional encryption scheme with succinct encryption, for function class $\mathcal{F}$, and message space $\{0, 1\}^*$, if it satisfies the following:

- (1) **Correctness:** for every $\lambda, n, \ell \in \mathbb{N}$, message $m \in \{0, 1\}^n$, functions $f = (f_1, \dots, f_\ell) \in \mathcal{F}^\ell$, with domain $\{0, 1\}^n$, and every $i \in [\ell]$,

$$\Pr \left[f_i(m) \leftarrow \text{MFE.Dec}(\text{FSK}_{f_i}, \text{CT}) \mid \begin{array}{l} (\text{PK}, \text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}) \leftarrow \text{MFE.Setup}(1^\lambda, f) \\ \text{CT} \leftarrow \text{MFE.Enc}(\text{PK}, m) \end{array} \right] = 1.$$

- (2) **Selective security:** for any polynomial-size adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\lambda)$ such that for any $\lambda, n, \ell \in \mathbb{N}$, any $m_0, m_1 \in \{0, 1\}^n$, and functions $f = (f_1, \dots, f_\ell) \in \mathcal{F}^\ell$ such that $f(m_0) = f(m_1)$,

$$\text{PK}, \text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}, \text{MFE.Enc}(\text{PK}, m_0) \approx_{\mathcal{A}, \mu} \text{PK}, \text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}, \text{MFE.Enc}(\text{PK}, m_1),$$

where $(\text{PK}, \text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}) \leftarrow \text{MFE.Setup}(1^\lambda, f)$.

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

- (3) **Succinct encryption circuit:** there exists a polynomial Φ and a constant $0 < \varepsilon \leq 1$, such that for any input-size n , circuit-size s , depth d and s -size, d -depth functions $f = (f_1, \dots, f_\ell) : \{0, 1\}^n \rightarrow \{0, 1\}^*$,

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq (\ell \cdot s)^{1-\varepsilon} \cdot \Phi(n, \lambda),$$

where $(\text{PK}, \text{FSK}_{f_1}, \dots, \text{FSK}_{f_\ell}) \leftarrow \text{MFE.Setup}(1^\lambda, f)$.

- Encryption is **fully succinct** if $\varepsilon = 1$.
- Encryption is **weakly depth-succinct** if

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq (\ell \cdot s)^{1-\varepsilon} \cdot \Phi(n, d, \lambda).$$

- Encryption is **very weakly depth-succinct** if

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq (\ell \cdot s)^{1-\varepsilon} \cdot \Phi(n, 2^d, \lambda).$$

- Encryption is **weakly size-succinct** if

$$|\text{MFE.Enc}(\text{PK}, \cdot)| \leq \ell^{1-\varepsilon} \cdot \Phi(s, \lambda).$$

We first observe that for all notions of succinctness, but weak size-succinctness, single-key, and multi-key FE are equivalent (essentially tautological).

CLAIM 4.2. *The notions of (δ -secure) single-key FE (Definition 2.1) and multi-key FE (Definition 2.1) are equivalent with respect to succinctness/weak-depth-succinctness/very-weak-depth-succinctness.*

PROOF. First, note that any multi-key FE is a single-key FE where $\ell = 1$, with respect to the above succinctness notions. To see that any single-key FE implies a multi-key FE MFE, consider the MFE scheme where any tuple of size- s , depth- d functions $f_1, \dots, f_\ell : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is parsed as a size- $s \cdot \ell$, depth- d single function $f : \{0, 1\}^n \rightarrow \{0, 1\}^{m \times \ell}$.

We then define the following:

- $\text{MFE.Setup}(1^\lambda, f_1, \dots, f_\ell) = \text{PK}, (\text{FSK}_f, 1), \dots, (\text{FSK}_f, \ell)$, where $\text{PK}, \text{FSK}_f \leftarrow \text{FE.Setup}(1^\lambda, f)$.
- $\text{MFE.Enc}(\text{PK}, m) = \text{FE.Enc}(\text{PK}, m)$.
- $\text{MFE.Dec}((\text{FSK}, i), \text{CT}) = m_i$, where $m_1 \dots m_\ell = \text{FE.Dec}(\text{FSK}, \text{CT})$.

The functionality, security, and succinctness/weak-depth-succinctness/very-weak-depth-succinctness of MFE follow readily from those of FE. $\square$

In contrast to the above, *weak size-succinctness* is only meaningful for $\ell \gg 1$. For $\ell = 1$, it corresponds to a single-key scheme with no succinctness at all. Indeed, the notion of full collusion-resistance common in the literature (for instance, in Reference [37]) only implies weak size-succinctness.

From Weak Succinctness to Succinctness. We now prove a bootstrapping theorem that shows that any multi-key FE scheme with (one of several forms of) weak succinctness implies succinct single-key FE (and also succinct multi-key FE by Claim 4.2).

The transformation is similar in spirit to previous randomized-encoding-based bootstrapping schemes from the literature [3, 48]. Roughly speaking, we rely on the fact that any function $f(x)$ represented by an arbitrary circuit has a randomized encoding $\hat{f}(x; r)$ that can be *decomposed* into multiple functions $\{\hat{f}_i(x; r)\}$ of lower complexity, both in terms of circuit-size and circuit-depth. Such an encoding can be efficiently decoded to $f(x)$ and reveals nothing on x but the result $f(x)$. This intuitively means that to generate a functional key for any function f it is sufficient to generate functional keys for the encoding functions $\{\hat{f}_i\}$, while including the randomness required for encoding as part of the encrypted message. In particular, if the complexity of encryption depends on the complexity of functions, we have now managed to reduce it.

The actual construction is somewhat more complicated to facilitate the security proof. Here, we also rely on a common proof technique based on *the trapdoor paradigm* [3, 32, 34] (sometimes termed *the Trojan method*).

4.2 The Transformation

We turn to describe the actual transformation.

Ingredients. We rely on the following primitives:

- A multi-key, selectively secure, public-key functional encryption scheme MFE for all circuits.
- A one-time symmetric encryption scheme Sym.
- A decomposable randomized encoding scheme RE.

The constructed scheme FE consists of algorithms

$$(\text{FE.Setup}, \text{FE.Enc}, \text{FE.Dec})$$

described in Figure 5.

Fig. 5. A single-key functional encryption with succinct encryption.

THEOREM 4.3. *FE is a single-key functional encryption for all circuits. If MFE is weakly (depth/size) succinct, or MFE is very weakly depth-succinct and both RE and Sym are shallow, then FE is succinct. If MFE, RE, and Sym are all δ -secure, then so is FE.*

The above theorem, in particular, implies that any single key FE that is weakly depth-succinct implies a succinct FE, and assuming also pseudorandom functions in NC^1 , so does any very weakly depth-succinct FE.

COROLLARY 4.4 (OF CLAIM 4.2 AND FACTS 2.7,2.9). *Any (δ -secure) single-key functional encryption that is weakly depth-succinct can be transformed into a (δ -secure) succinct scheme. Assuming pseudorandom functions in NC^1 , very weak depth-succinctness is sufficient.*

PROOF. We prove that the constructed scheme satisfies the properties of single-key FE with succinct encryption.

Correctness: Fix a security parameter $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}^n$, and $f \in \mathcal{F}$ with domain $\{0, 1\}^n$. By the functionality of MFE and correctness of RE,

$$\begin{aligned} \text{FE.Dec}(\text{FSK}_f, \text{FCT}) &= \text{RE.Dec}\left(\left\{\text{MFE.Dec}(\text{FSK}_{\hat{F}_i}, \text{FCT})\right\}_{i \in [\ell]}\right) \\ &= \text{RE.Dec}\left(\left\{\hat{F}_i(m, r, \perp, 0)\right\}_{i \in [\ell]}\right) \\ &= \text{RE.Dec}\left(\left\{\hat{f}_i(m; r)\right\}_{i \in [\ell]}\right) = f(m), \end{aligned}$$

where $(\text{PK}, \text{FSK}_f) \leftarrow \text{FE.Setup}(1^\lambda, f)$, $\text{FSK}_f = \{\text{FSK}_{\hat{F}_i}\}_{i \in [\ell]}$, $\hat{F}_i$ is the function underlying $\text{FSK}_{\hat{F}_i}$, $\text{FCT} = \text{FE.Enc}(\text{PK}, m)$, and $\{\hat{f}_i\} \leftarrow \text{RE.Enc}(f)$.

Selective security: Fix any polynomial-size $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$, we prove that there exists negligible $\mu_{\text{FE}}(\lambda)$ such that for any security parameter λ , function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and two messages $m_0, m_1 \in \{0, 1\}^n$ such that $f(m_0) = f(m_1)$,

$$\text{FSK}_f, \text{FE.Enc}(\text{PK}, m_0) \approx_{\mathcal{A}, \mu} \text{FSK}_f, \text{FE.Enc}(\text{PK}, m_1),$$

where $(\text{PK}, \text{FSK}_f) \leftarrow \text{FE.Setup}(1^\lambda, f)$.

For this purpose, we consider the following sequence hybrid experiments.

$\mathcal{H}_b^0$ for $b \in \{0, 1\}$: This hybrid corresponds to the (real) distribution:

$$\begin{aligned} (\text{PK}, \text{FSK}_f) &= (\text{PK}, \text{FSK}_{\hat{F}_1}, \dots, \text{FSK}_{\hat{F}_\ell}) \leftarrow \text{MFE.Setup}(1^\lambda, \hat{F}_1, \dots, \hat{F}_\ell), \\ \text{FCT} &\leftarrow \text{MFE.Enc}(\text{PK}, (m_b, r, \perp, 0)), \end{aligned}$$

where $\hat{F}_i = \hat{F}_i[\text{CT}_i]$ is a circuit that according to a choice input bit β , either computes the i th encoding $\hat{f}_i(m, r)$ or decrypts CT_i . In this hybrid, $\beta = 0$ and the circuit performs the first operation. The ciphertext $\text{CT} = \text{CT}_1 \dots, \text{CT}_\ell \leftarrow \text{Sym.Enc}(\text{SK}, 0^\ell)$ encrypts zeros.

$\mathcal{H}_b^1$ for $b \in \{0, 1\}$: This hybrid is identical to $\mathcal{H}_b^0$ except that $\text{CT} \leftarrow \text{Sym.Enc}(\text{SK}, \hat{f}_1(m_b; r), \dots, \hat{f}_\ell(m_b; r))$ is an encryption of the encoding $\hat{f}(m; r) = \hat{f}_1(m; r), \dots, f_\ell(m; r)$ instead of 0^ℓ .

$\mathcal{H}_b^2$ for $b \in \{0, 1\}$: This hybrid is identical to $\mathcal{H}_b^1$ except that $\text{FCT} \leftarrow \text{MFE.Enc}(\text{PK}, (\perp, \perp, \text{SK}, 1))$ encrypts $(\perp, \perp, \text{SK}, 1)$ instead of $(m_b, r, \perp, 0)$.

$\mathcal{H}^3$: This hybrid is identical to $\mathcal{H}_b^2$ except that $\text{CT} \leftarrow \text{Sym.Enc}(\text{SK}, \hat{y}_1, \dots, \hat{y}_\ell)$ is an encryption of a simulated encoding $y_1, \dots, \hat{y}_\ell \leftarrow \text{RE.Sim}(f, f(m_0))$ instead of $\hat{f}(m_b; r)$.

CLAIM 4.5. *There exist negligible functions $\mu_{\text{Sym}}(\lambda), \mu_{\text{MFE}}(\lambda), \mu_{\text{RE}}(\lambda)$, such that*

$$\mathcal{H}_b^0 \approx_{\mathcal{A}, \mu_{\text{Sym}}} \mathcal{H}_b^1 \approx_{\mathcal{A}, \mu_{\text{MFE}}} \mathcal{H}_b^2 \approx_{\mathcal{A}, \mu_{\text{RE}}} \mathcal{H}^3.$$

Furthermore, if Sym, MFE, RE are δ -secure, then the functions μ above can be replaced with $\delta^{\Omega(1)}$.

Before proving the claim, note that it concludes the security proof. Indeed, setting $\mu_{\text{FE}} = 2(\mu_{\text{Sym}} + \mu_{\text{MFE}} + \mu_{\text{RE}})$, it implies indistinguishability of the two real experiments: $\mathcal{H}_0^0 \approx_{\mathcal{A}, \mu_{\text{FE}}} \mathcal{H}_0^1$.

PROOF OF CLAIM 4.5. We prove indistinguishability between each two subsequent hybrids.

$\mathcal{H}_b^0 \approx \mathcal{H}_b^1$: Here the difference is in the symmetric encryption $\text{CT} = \text{CT}_1, \dots, \text{CT}_\ell$ underlying the functions $\hat{F}_1[\text{CT}_1], \dots, \hat{F}_\ell[\text{CT}_\ell]$. In $\mathcal{H}_b^0$, we encrypt 0^ℓ and in $\mathcal{H}_b^1$, we encrypt $\hat{f}(m_b, r)$.

Thus, by the indistinguishability of Sym, there exists μ_{Sym} , such that

$$\mathcal{H}_b^0 \approx_{\mathcal{A}, \mu_{\text{Sym}}} \mathcal{H}_b^1.$$

$\mathcal{H}_b^1 \approx \mathcal{H}_b^2$: Here the difference is in the functional ciphertext FCT. In $\mathcal{H}_b^1$, we encrypt $(m_b, r, \perp, 0)$ and in $\mathcal{H}_b^2$, we encrypt $(\perp, \perp, \text{SK}, r)$. For every function $\hat{F}_i$,

$$\hat{F}_i(m_b, r, \perp, 0) = \hat{F}_i(m_b; r) = \text{Sym.Dec}(\text{SK}, \text{CT}_i, i) = \hat{F}_i(\perp, \perp, \text{SK}, 1).$$

Thus, by the selective security of MFE, there exists μ_{MFE} , such that

$$\mathcal{H}_b^1 \approx_{\mathcal{A}, \mu_{\text{MFE}}} \mathcal{H}_b^2.$$

$\mathcal{H}_b^2 \approx \mathcal{H}^3$: Here the difference is again in the symmetric encryption $\text{CT} = \text{CT}_1, \dots, \text{CT}_\ell$ underlying the functions $\hat{F}_1[\text{CT}_1], \dots, \hat{F}_\ell[\text{CT}_\ell]$. In $\mathcal{H}_b^2$, we encrypt the encoding $\hat{f}(m_b, r)$, and in $\mathcal{H}^3$ we encrypt the simulated encoding $\hat{y} \leftarrow \text{RE.Sim}(f, y)$, where $y = f(m_0) = f(m_b)$.

Thus, by the privacy of RE, there exists μ_{RE} , such that

$$\mathcal{H}_b^2 \approx_{\mathcal{A}, \mu_{\text{RE}}} \mathcal{H}^3.$$

In all of the above, it follows readily that if Sym, RE, MFE are δ -secure for some concrete negligible δ , then FE is δ -secure.

This concludes the proof of the claim and of the security of the constructed FE. $\square$

Succinct encryption: We first show that if MFE is weakly size-succinct, then FE is succinct. This also implies that FE is succinct if MFE is weakly depth-succinct, as weak depth-succinctness implies weak size-succinctness. We then show that FE is succinct also if MFE is very weakly depth-succinct and RE and Sym are shallow.

Assume that MFE is weakly size-succinct. That is, there exist a polynomial Φ_{MFE} and a constant $0 < \varepsilon \leq 1$ such that the size of the encryption circuit corresponding to ℓ functions is bounded by

$$\ell^{1-\varepsilon} \cdot \Phi_{\text{MFE}}(\hat{s}, \lambda),$$

where $\hat{s}$ is the circuit-size of each of the ℓ functions chosen during the setup phase. We show that FE is also succinct with parameters $(\varepsilon, \Phi_{\text{FE}})$ for a fixed polynomial Φ_{FE} .

Fix a size- s function f . We analyze the size and depth of the corresponding functions $\hat{F}_1, \dots, \hat{F}_\ell$. Recall that each function $\hat{F}_i = \hat{F}_i[\text{CT}_i]$ either decrypts a bit, or computes one encoding bit. Specifically, it has the form

$$\hat{F}_i(m, r, \text{SK}, b) = b \otimes \text{Sym.Dec}(\text{SK}, \text{CT}_i, i) \oplus (1 - b) \otimes \hat{f}_i(m; r),$$

where $\otimes, \oplus$ denote multiplication and addition modulu 2.

Thus,

$$\hat{s} := |\hat{F}_i| \leq O(|\text{Sym.Dec}(\cdot, \text{CT}_i, i)| + |\hat{f}_i(\cdot; \cdot)|) \leq \Phi_{\text{Sym, RE}}(n, \lambda),$$

where $\Phi_{\text{Sym, RE}}$ is a fixed polynomial that depends only on the schemes Sym and RE. Furthermore, by the decomposability of RE, it is the case that

$$\ell \leq s \cdot \Phi_{\text{RE}}(n, \lambda),$$

where Φ_{RE} is a fixed polynomial that depends only on RE.

Applying the weak size-succinctness of MFE, it follows that in FE, the size of an encryption circuit corresponding to a size- s function f is bounded by

$$(s \cdot \Phi_{\text{RE}}(n, \lambda))^{1-\varepsilon} \Phi_{\text{MFE}}(\Phi_{\text{Sym, RE}}(n, \lambda), \lambda).$$

This implies succinctness of FE with parameters (ε, Φ) when setting

$$\Phi(n, \lambda) = \Phi_{\text{RE}}(n, \lambda) \Phi_{\text{MFE}}(\Phi_{\text{Sym, RE}}(n, \lambda), \lambda).$$

Assume now that MFE is very weakly depth-succinct. That is, there exists a polynomial Φ_{MFE} and constant $0 < \varepsilon \leq 1$ such that the size of the encryption circuit corresponding to ℓ functions is bounded by

$$(\ell \hat{s})^{1-\varepsilon} \cdot \Phi_{\text{MFE}}(n, 2^{\hat{d}}, \lambda),$$

where $\hat{s}$ is the circuit-size and $\hat{d}$ is the circuit-depth of each of the ℓ functions chosen during the setup phase. Assume also that both Sym and RE are shallow.

Recalling again the definition of the functions $\hat{F}_1, \dots, \hat{F}_\ell$ note that the depth of each $\hat{F}_i$ is at most

$$c \cdot \max \{d_{\text{Sym}}, d_{\text{RE}}\} \leq 2c \cdot \log \lambda,$$

where d_{Sym} and d_{RE} are the depth of the decryption circuit in Sym and of each $\hat{f}_i$ in RE, c is an absolute constant, and we rely on the fact that both schemes are shallow.

Applying the very weak depth-succinctness of MFE, it follows that in FE, the size of an encryption circuit corresponding to a size- s function f is bounded by

$$(s \cdot \Phi_{\text{RE}}(n, \lambda) \cdot \Phi_{\text{Sym, RE}}(n, \lambda))^{1-\varepsilon} \Phi_{\text{MFE}}(n, 2^{2c \log \lambda}, \lambda).$$

This implies succinctness of FE with parameters (ε, Φ) when setting

$$\Phi(n, \lambda) = \Phi_{\text{RE}} \cdot \Phi_{\text{Sym, RE}}(n, \lambda) \cdot \Phi_{\text{MFE}}(n, \lambda^{2^c}, \lambda).$$

□

ACKNOWLEDGMENTS

We thank the reviewers of JACM for their insightful comments and valuable suggestions. We thank Daniel Wichs for pointing out an error in a previous version of the article. Nir Bitansky is a member of the Check Point Institute of Information Security and is supported by the Alon Young Faculty Fellowship and by Len Blavatnik and the Blavatnik Family Foundation. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA and ARO.

REFERENCES

- [1] Shweta Agrawal, David Mandell Freeman, and Vinod Vaikuntanathan. 2011. Functional encryption for inner product predicates from learning with errors. In *Proceedings of the 17th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'11) (Lecture Notes in Computer Science)*, Dong Hoon Lee and Xiaoyun Wang (Eds.), Vol. 7073. Springer, 21–40. DOI : http://dx.doi.org/10.1007/978-3-642-25385-0_2
- [2] Shweta Agrawal, Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2013. Functional encryption: New perspectives and lower bounds, see Reference Canetti and Garay [28], 500–518. DOI : http://dx.doi.org/10.1007/978-3-642-40084-1_28
- [3] Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. 2015. The Trojan method in functional encryption: From selective to adaptive security, generically. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'15)*.
- [4] Prabhanjan Ananth and Abhishek Jain. 2015. Indistinguishability obfuscation from compact functional encryption. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'15)*.
- [5] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. 2015. Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. *IACR Cryptol. ePrint Arch.* 2015 (2015), 730. <http://eprint.iacr.org/2015/730>.
- [6] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. 2017. Indistinguishability obfuscation for turing machines: Constant overhead and amortization. In *Proceedings of the 37th Annual International Cryptology Conference (CRYPTO'17)*. 252–279. DOI : http://dx.doi.org/10.1007/978-3-319-63715-0_9
- [7] Prabhanjan Ananth and Amit Sahai. 2017. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In *Proceedings of the 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'17)*. 152–181. DOI : http://dx.doi.org/10.1007/978-3-319-56620-7_6
- [8] Benny Applebaum and Zvika Brakerski. 2015. Obfuscating circuits via composite-order graded encoding. In *Proceedings of the Theory of Cryptography Conference (TCC'15)*.
- [9] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2004. Cryptography in NC^0 . In *Proceedings of the 45th Symposium on Foundations of Computer Science (FOCS'04)*. 166–175. DOI : <http://dx.doi.org/10.1109/FOCS.2004.20>
- [10] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. 2006. Computationally private randomizing polynomials and their applications. *Comput. Complex.* 15, 2 (2006), 115–162.
- [11] Benny Applebaum, Yuval Ishai, Eyal Kushilevitz, and Brent Waters. 2013. Encoding functions with constant online rate or how to compress garbled circuits keys, see Reference Canetti and Garay [28], 166–184. DOI : http://dx.doi.org/10.1007/978-3-642-40084-1_10
- [12] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. 2014. Protecting obfuscation against algebraic attacks. In *Proceedings of the 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'14)*. 221–238. DOI : http://dx.doi.org/10.1007/978-3-642-55220-5_13
- [13] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. 2012. On the (im)possibility of obfuscating programs. *J. ACM* 59, 2 (2012), 6.
- [14] Nir Bitansky, Ran Canetti, Henry Cohn, Shafi Goldwasser, Yael Tauman Kalai, Omer Paneth, and Alon Rosen. 2014. The impossibility of obfuscation with auxiliary input or a universal simulator. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'14)*. 71–89.
- [15] Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. 2015. Succinct randomized encodings and their applications. In *Proceedings of the Symposium on Theory of Computing (STOC'15)*.
- [16] Nir Bitansky, Ryo Nishimaki, Alain Passelègue, and Daniel Wichs. 2016. From cryptomania to obfustopia through secret-key functional encryption. *IACR Cryptol. ePrint Arch.* 2016 (2016), 558. Retrieved from <http://eprint.iacr.org/2016/558>.
- [17] Andrej Bogdanov and Alon Rosen. 2017. Pseudorandom functions: Three decades later. In *Tutor. Found. Cryptogr.* 79–158. DOI : http://dx.doi.org/10.1007/978-3-319-57048-8_3

- [18] Dan Boneh, Giovanni Di Crescenzo, Rafail Ostrovsky, and Giuseppe Persiano. 2004. Public key encryption with keyword search. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'04)*. 506–522.
- [19] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. 2014. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Proceedings of the 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'14)*. 533–556. DOI: http://dx.doi.org/10.1007/978-3-642-55220-5_30
- [20] Dan Boneh, Amit Sahai, and Brent Waters. 2012. Functional encryption: A new vision for public-key cryptography. *Commun. ACM* 55, 11 (2012), 56–64. DOI: <http://dx.doi.org/10.1145/2366316.2366333>
- [21] Dan Boneh and Brent Waters. 2013. Constrained pseudorandom functions and their applications. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'13) (Lecture Notes in Computer Science)*, Kazue Sako and Palash Sarkar (Eds.), Vol. 8270. Springer, 280–300.
- [22] Elette Boyle, Kai-Min Chung, and Rafael Pass. 2014. On extractability obfuscation. In *Proceedings of the Theory of Cryptography Conference (TCC'14)*. 52–73.
- [23] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. 2014. Functional signatures and pseudorandom functions. In *Proceedings of the International Conference on Practice and Theory of Public Key Cryptography (PKC'14) (Lecture Notes in Computer Science)*, Hugo Krawczyk (Ed.), Vol. 8383. Springer, 501–519. DOI: <http://dx.doi.org/10.1007/978-3-642-54631-0>
- [24] Zvika Brakerski, Ilan Komargodski, and Gil Segev. 2016. Multi-input functional encryption in the private-key setting: Stronger security from weaker assumptions. In *Proceedings of the 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'16)*. 852–880. DOI: http://dx.doi.org/10.1007/978-3-662-49896-5_30
- [25] Zvika Brakerski and Guy N. Rothblum. 2014. Virtual black-box obfuscation for all circuits via generic graded encoding. In *Proceedings of the 11th Theory of Cryptography Conference (TCC'14) (Lecture Notes in Computer Science)*, Yehuda Lindell (Ed.), Vol. 8349. Springer, 1–25. DOI: http://dx.doi.org/10.1007/978-3-642-54242-8_1
- [26] Zvika Brakerski and Gil Segev. 2015. Function-private functional encryption in the private-key setting. In *Proceedings of the Theory of Cryptography Conference (TCC'15)*.
- [27] Ran Canetti. 1997. Towards realizing random oracles: Hash functions that hide all partial information. In *Proceedings of the 17th Annual International Cryptology Conference (CRYPTO'97) (Lecture Notes in Computer Science)*, Burton S. Kaliski Jr. (Ed.), Vol. 1294. Springer, 455–469. DOI: <http://dx.doi.org/10.1007/BFb0052255>
- [28] Ran Canetti and Juan A. Garay (Eds.). 2013. *Proceedings of the 33rd Annual Cryptology Conference (CRYPTO'13)*. Lecture Notes in Computer Science, Vol. 8043. Springer. DOI: <http://dx.doi.org/10.1007/978-3-642-40084-1>
- [29] Ran Canetti, Justin Holmgren, Abhishek Jain, and Vinod Vaikuntanathan. 2015. Succinct garbling and indistinguishability obfuscation for RAM programs. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC'15)*. 429–437. DOI: <http://dx.doi.org/10.1145/2746539.2746621>
- [30] Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. 2015. Obfuscation of probabilistic circuits and applications. In *Proceedings of the Theory of Cryptography Conference (TCC'15)*.
- [31] Ran Canetti, Guy N. Rothblum, and Mayank Varia. 2010. Obfuscation of hyperplane membership. In *Proceedings of the 7th Theory of Cryptography Conference (TCC'10) (Lecture Notes in Computer Science)*, Daniele Micciancio (Ed.), Vol. 5978. Springer, 72–89. DOI: http://dx.doi.org/10.1007/978-3-642-11799-2_5
- [32] Angelo De Caro, Vincenzo Iovino, Abhishek Jain, Adam O'Neill, Omer Paneth, and Giuseppe Persiano. 2013. On the achievability of simulation-based security for functional encryption, see Reference Canetti and Garay [28], 519–535. DOI: http://dx.doi.org/10.1007/978-3-642-40084-1_29
- [33] Stephen A. Cook and H. James Hoover. 1985. A depth-universal circuit. *SIAM J. Comput.* 14, 4 (1985), 833–839. DOI: <http://dx.doi.org/10.1137/0214058>
- [34] Uriel Feige and Adi Shamir. 1989. Zero knowledge proofs of knowledge in two rounds. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'89)*. 526–544.
- [35] Sanjam Garg, Craig Gentry, and Shai Halevi. 2013. Candidate multilinear maps from ideal lattices. In *Proceedings of the 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'13)*. 1–17. DOI: http://dx.doi.org/10.1007/978-3-642-38348-9_1
- [36] Sanjam Garg, Craig Gentry, Shai Halevi, Amit Sahai, Mariana Raikova, and Brent Waters. 2013. Candidate Indistinguishability Obfuscation and Functional Encryption for all circuits. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS'13)*.
- [37] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. 2016. Functional encryption without obfuscation. In *Proceedings of the 13th International Theory of Cryptography Conference (TCC'16)*. 480–511. DOI: http://dx.doi.org/10.1007/978-3-662-49099-0_18
- [38] Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. 2016. Revisiting the cryptographic hardness of finding a Nash equilibrium. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'16)*.

- [39] Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. 2017. Breaking the sub-exponential barrier in obfuscation. In *Proceedings of the 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'17)*. 156–181. DOI : http://dx.doi.org/10.1007/978-3-319-56617-7_6
- [40] Sanjam Garg and Akshayaram Srinivasan. 2016. Single-key to multi-key functional encryption with polynomial loss. In *Proceedings of the 14th International Theory of Cryptography Conference (TCC'16)*. 419–442. DOI : http://dx.doi.org/10.1007/978-3-662-53644-5_16
- [41] Craig Gentry, Allison B. Lewko, Amit Sahai, and Brent Waters. 2015. Indistinguishability obfuscation from the multi-linear subgroup elimination assumption. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS'15)*.
- [42] Oded Goldreich. 2001. *The Foundations of Cryptography—Volume 1, Basic Techniques*. Cambridge University Press.
- [43] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1986. How to construct random functions. *J. ACM* 33, 4 (1986), 792–807. DOI : <http://dx.doi.org/10.1145/6490.6503>
- [44] Shafi Goldwasser, S. Dov Gordon, Vipul Goyal, Abhishek Jain, Jonathan Katz, Feng-Hao Liu, Amit Sahai, Elaine Shi, and Hong-Sheng Zhou. 2014. Multi-input functional encryption. In *Proceedings of the 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'14)*. 578–602. DOI : http://dx.doi.org/10.1007/978-3-642-55220-5_32
- [45] Shafi Goldwasser and Yael Tauman Kalai. 2005. On the impossibility of obfuscation with auxiliary input. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS'05)*. IEEE Computer Society, 553–562.
- [46] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. 2013. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the Symposium on Theory of Computing Conference (STOC'13)*. 555–564. DOI : <http://dx.doi.org/10.1145/2488608.2488678>
- [47] Shafi Goldwasser and Guy N. Rothblum. 2007. On best-possible obfuscation. In *Proceedings of the Theory of Cryptography Conference (TCC'07)*. 194–213.
- [48] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2012. Functional encryption with bounded collusions via multi-party computation. In *Proceedings of the Annual International Cryptology Conference (CRYPTO'12)*. 162–179.
- [49] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2013. Attribute-based encryption for circuits. In *Proceedings of the Symposium on Theory of Computing Conference (STOC'13)*. 545–554.
- [50] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. 2015. Predicate encryption for circuits from LWE. *IACR Cryptol. ePrint Arch.* 2015 (2015), 29. Retrieved from <http://eprint.iacr.org/2015/029>.
- [51] Satoshi Hada. 2000. Zero-knowledge and code obfuscation. In *Proceedings of the 6th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT'00)*. 443–457. DOI : http://dx.doi.org/10.1007/3-540-44448-3_34
- [52] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. 1999. A pseudorandom generator from any one-way function. *SIAM J. Comput.* 28, 4 (1999), 1364–1396. DOI : <http://dx.doi.org/10.1137/S0097539793244708>
- [53] Dennis Hofheinz, John Malone-Lee, and Martijn Stam. 2007. Obfuscation for cryptographic purposes. In *Proceedings of the 4th Theory of Cryptography Conference (TCC'07) (Lecture Notes in Computer Science)*, Salil P. Vadhan (Ed.), Vol. 4392. Springer, 214–232. DOI : http://dx.doi.org/10.1007/978-3-540-70936-7_12
- [54] Susan Hohenberger, Guy N. Rothblum, Abhi Shelat, and Vinod Vaikuntanathan. 2011. Securely obfuscating re-encryption. *J. Cryptol.* 24, 4 (2011), 694–719. DOI : <http://dx.doi.org/10.1007/s00145-010-9077-7>
- [55] Yuval Ishai and Eyal Kushilevitz. 2000. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS'00)*. IEEE Computer Society, 294–304.
- [56] Jonathan Katz, Amit Sahai, and Brent Waters. 2013. Predicate encryption supporting disjunctions, polynomial equations, and inner products. *J. Cryptol.* 26, 2 (2013), 191–224. DOI : <http://dx.doi.org/10.1007/s00145-012-9119-4>
- [57] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. 2013. Delegatable pseudorandom functions and applications. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'13)*, Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung (Eds.). ACM, 669–684. DOI : <http://dx.doi.org/10.1145/2508859.2516668>
- [58] Fuyuki Kitagawa, Ryo Nishimaki, and Keisuke Tanaka. 2017. Indistinguishability obfuscation for all circuits from secret-key functional encryption. *IACR Cryptol. ePrint Arch.* 2017 (2017), 361. Retrieved from <http://eprint.iacr.org/2017/361>.
- [59] Ilan Komargodski, Tal Moran, Moni Naor, Rafael Pass, Alon Rosen, and Eylon Yogev. 2014. One-way functions and (im)perfect obfuscation. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS'14)*. IEEE Computer Society, 374–383. DOI : <http://dx.doi.org/10.1109/FOCS.2014.47>
- [60] Venkata Koppula, Allison Bishop Lewko, and Brent Waters. 2015. Indistinguishability obfuscation for turing machines with unbounded memory. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC'15)*. 419–428. DOI : <http://dx.doi.org/10.1145/2746539.2746614>

- [61] Baiyu Li and Daniele Micciancio. 2016. Compactness vs. collusion resistance in functional encryption. In *Proceedings of the 14th International Theory of Cryptography Conference (TCC'16)*. 443–468. DOI : http://dx.doi.org/10.1007/978-3-662-53644-5_17
- [62] Huijia Lin. 2016. Indistinguishability obfuscation from constant-degree graded encoding schemes. In *Proceedings of the 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'16)*. 28–57. DOI : http://dx.doi.org/10.1007/978-3-662-49890-3_2
- [63] Huijia Lin. 2017. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In *Proceedings of the 37th Annual International Cryptology Conference (CRYPTO'17)*. 599–629. DOI : http://dx.doi.org/10.1007/978-3-319-63688-7_20
- [64] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. 2016. Indistinguishability obfuscation with non-trivial efficiency. In *Proceedings of the 19th IACR International Conference on Practice and Theory in Public-Key Cryptography (PKC'16)*. 447–462. DOI : http://dx.doi.org/10.1007/978-3-662-49387-8_17
- [65] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. 2016. Output-compressing randomized encodings and applications. In *Proceedings of the 13th International Theory of Cryptography Conference (TCC'16)*. 96–124. DOI : http://dx.doi.org/10.1007/978-3-662-49096-9_5
- [66] Huijia Lin and Stefano Tessaro. 2017. Indistinguishability obfuscation from trilinear maps and block-wise local PRGs. In *Proceedings of the 37th Annual International Cryptology Conference (CRYPTO'17)*. 630–660. DOI : http://dx.doi.org/10.1007/978-3-319-63688-7_21
- [67] Huijia Lin and Vinod Vaikuntanathan. 2016. Indistinguishability obfuscation from DDH-like assumptions on constant-degree graded encodings. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS'16)*. 11–20. DOI : <http://dx.doi.org/10.1109/FOCS.2016.11>
- [68] Qipeng Liu and Mark Zhandry. 2017. Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In *Proceedings of the 15th International Conference on Theory of Cryptography (TCC'17)*. 138–169. DOI : http://dx.doi.org/10.1007/978-3-319-70500-2_6
- [69] Ben Lynn, Manoj Prabhakaran, and Amit Sahai. 2004. Positive results and techniques for obfuscation. In *Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'04) (Lecture Notes in Computer Science)*, Christian Cachin and Jan Camenisch (Eds.), Vol. 3027. Springer, 20–39. DOI : http://dx.doi.org/10.1007/978-3-540-24676-3_2
- [70] Adam O'Neill. 2010. Definitional issues in functional encryption. *Cryptology ePrint Archive, Report 2010/556*. (2010).
- [71] Rafael Pass, Karn Seth, and Sidharth Telang. 2014. Indistinguishability obfuscation from semantically secure multilinear encodings. In *Proceedings of the 34th Annual Cryptology Conference (CRYPTO'14) (Lecture Notes in Computer Science)*, Juan A. Garay and Rosario Gennaro (Eds.), Vol. 8616. Springer, 500–517. DOI : http://dx.doi.org/10.1007/978-3-662-44371-2_28
- [72] Amit Sahai and Hakan Seyalioglu. 2010. Worry-free encryption: Functional encryption with public keys. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS'10)*. 463–472.
- [73] Amit Sahai and Brent Waters. 2005. Fuzzy identity-based encryption. In *Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'05) (Lecture Notes in Computer Science)*, Ronald Cramer (Ed.), Vol. 3494. Springer, 457–473. DOI : http://dx.doi.org/10.1007/11426639_27
- [74] Amit Sahai and Brent Waters. 2014. How to use indistinguishability obfuscation: Deniable encryption, and more. In *Proceedings of the Symposium on Theory of Computing (STOC'14)*, David B. Shmoys (Ed.). ACM, 475–484. DOI : <http://dx.doi.org/10.1145/2591796.2591825>
- [75] Leslie G. Valiant. 1976. Universal circuits (preliminary report). In *Proceedings of the 8th Annual ACM Symposium on Theory of Computing (STOC'76)*. 196–203.
- [76] Hoeteck Wee. 2005. On obfuscating point functions. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, Harold N. Gabow and Ronald Fagin (Eds.). ACM, 523–532. DOI : <http://dx.doi.org/10.1145/1060590.1060669>
- [77] Andrew Chi-Chih Yao. 1986. How to generate and exchange secrets (extended abstract). In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science (FOCS'86)*. IEEE Computer Society, 162–167. DOI : <http://dx.doi.org/10.1109/SFCS.1986.25>
- [78] Joe Zimmerman. 2015. How to obfuscate programs directly. In *Proceedings of the Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT'15)*.

Received October 2016; revised May 2018; accepted June 2018