

Indistinguishability Obfuscation from Trilinear Maps and Block-Wise Local PRGs

Huijia Lin*

Stefano Tessaro †

University of California, Santa Barbara
{rachel.lin, tessaro}@cs.ucsb.edu

Abstract

We consider the question of finding the lowest degree L for which L -linear maps suffice to obtain IO. The current state of the art (Lin, EUROCRYPT'16, CRYPTO '17; Lin and Vaikunthanathan, FOCS'16; Ananth and Sahai, EUROCRYPT '17) is that L -linear maps (under suitable security assumptions) suffice for IO, assuming the existence of pseudo-random generators (PRGs) with output locality L . However, these works cannot answer the question of whether $L < 5$ suffices, as no polynomial-stretch PRG with locality lower than 5 exists.

In this work, we present a new approach that relies on the existence of PRGs with *block-wise locality* L , i.e., every output bit depends on at most L (disjoint) *input blocks*, each consisting of up to $\log \lambda$ input bits. We show that the existence of PRGs with block-wise locality is plausible for any $L \geq 3$, and also provide:

- A construction of a general-purpose indistinguishability obfuscator from L -linear maps and a subexponentially-secure PRG with block-wise locality L and polynomial stretch.
- A construction of general-purpose functional encryption from L -linear maps and any slightly super-polynomially secure PRG with block-wise locality L and polynomial stretch.

All our constructions are based on the SXDH assumption on L -linear maps and subexponential Learning With Errors (LWE) assumption, and follow by instantiating our new generic bootstrapping theorems with Lin's recently proposed FE scheme (CRYPTO '17). Inherited from Lin's work, our security proof requires algebraic multilinear maps (Boneh and Silverberg, Contemporary Mathematics), whereas security when using noisy multilinear maps is based on a family of more complex assumptions that hold in the generic model.

Our candidate PRGs with block-wise locality are based on Goldreich's local functions, and we show that the security of instantiations with block-wise locality $L \geq 3$ is backed by similar validation as constructions with (conventional) locality 5. We further complement this with hardness amplification techniques that further weaken the pseudorandomness requirements.

*Huijia Lin was supported in part by NSF grants CNS-1528178, CNS-1514526, CNS-1652849 (CAREER).

†Stefano Tessaro was supported in part by NSF grants CNS-1423566, CNS-1528178, CNS-1553758 (CAREER), and IIS-152804.

Contents

1	Introduction	1
1.1	Block-Wise Locality	3
1.2	From Block-Wise Locality to IO and FE	4
1.3	Subsequent works	6
2	Preliminaries	7
2.1	μ -Hardness and μ -Indistinguishability	7
2.2	Indistinguishability Obfuscation	8
2.2.1	$i\ell$ -bit-Input IO	8
2.3	Puncturable Pseudo-Random Functions	9
2.4	Randomized Encodings	9
2.5	Functional Encryption	10
2.5.1	Public-Key Functional Encryption	10
2.5.2	FE for P/poly, NC ¹ and Compactness	11
2.6	Zero-Testing FE for Arithmetic Functions	12
2.7	Degree- D Asymmetric Multilinear Maps with SXDH Assumption	13
3	Block-Wise Local PRGs	13
3.1	Pseudorandom Generators, Locality, and Block-Wise Locality	13
3.2	Graph-Based Block-Wise local Functions	14
3.3	Pseudorandom and Unpredictability Generators	16
3.4	Block-Wise local Small-Bias Generators	18
3.5	Hardness Amplification via the XOR Construction	19
3.6	The Extraction Construction	21
4	IO from Block-Wise Locality-$(L, \log \lambda)$ PRG and L-Linear Maps	23
4.1	Step 1: Constructing Weakly-Compact FE	24
4.2	Step 2: Tight Construction of IO from Weakly-Compact FE	29
4.2.1	IO in the CRS model	31
4.2.2	IO in the CRS Model from Weakly-Compact FE	32
5	FE from $\omega(\log \lambda)$-Bit-Input IO for P/poly	38
5.1	From 1-key to Collusion-Resistant FE, Generically	39

1 Introduction

Indistinguishability obfuscation (IO), first defined in the seminal work of Barak *et al.* [BGI⁺01a], aims to obfuscate functionally equivalent programs into indistinguishable ones while preserving functionality. IO is an extraordinarily powerful object that has been shown to enable a large set of new cryptographic applications. All existing IO constructions [GGH⁺13b, BR14, BGK⁺14, PST14, AGIS14, GLSW15, Zim15, AB15, GMM⁺16, Lin16, LV16, Lin17, AS16, DGG⁺16] rely on *multilinear maps* or *graded encodings*. In particular, the power of an L -linear map – first made explicit by Boneh and Silverberg [BS02] – stems from the fact that it essentially allows to evaluate degree- L polynomials on secret encoded values, and to test whether the output of such polynomials is zero or not.

The case $L = 2$ corresponds to bilinear maps, which can be efficiently instantiated from elliptic curves. In contrast, the instantiation of L -linear maps with $L \geq 3$ has turned to be a far more challenging problem. Garg, Gentry, and Halevi [GGH13a] proposed in particular *noisy* (i.e., *approximate*) versions of L -linear maps for $L \geq 3$, and gave the first candidate construction. Unfortunately, vulnerabilities [CHL⁺15, CGH⁺15, MSZ16, CGH17, ADGM17] were later demonstrated against this and subsequent candidates [CLT13, LSS14, GGH15, CLT15]. Of course, this does not mean that the resulting constructions are insecure. In fact, this has motivated the search for IO constructions which withstand all existing attacks [GMM⁺16].

IO from low-degree multilinear maps. This paper addresses the problem of finding the smallest L such that degree- L multilinear maps are sufficient for constructing IO. This fits within the more general goal of ultimately assessing whether bilinear maps are sufficient. While first-generation IO constructions all required polynomial-degree multilinear maps, a series of recent works [Lin16, LV16, Lin17, AS16] reduced the required degree to $L = 5$, assuming the existence of PRGs with output locality 5 and subexponential LWE, and under suitable assumption on the 5-linear maps. However, these works left open the question of whether multilinear maps with degree $L < 5$ are sufficient.

Further reducing the degree is important. On the one hand, *if* IO can be achieved from bilinear maps, this is going to take us one step closer. On the other hand, even if bilinear maps would not suffice, it is potentially easier to find secure algebraic instantiations for low degree multilinear maps. Moreover, we want to understand the precise power these maps would enable.

Our contributions, in a nutshell This paper presents a new paradigm for IO constructions which admits instantiations with L -linear maps for $L \geq 3$, provided the SXDH assumption holds for the L -linear map. While this falls short of achieving IO from bilinear maps, our result shifts the focus on the fact that the gap between two- and three-linear maps is a seemingly fundamental barrier to be overcome. In particular, under the assumptions needed for our construction to be secure, this shows that building three-linear maps is as difficult as getting full-blown IO.

We fundamentally rely on the recent line of works on building IO from constant-degree multilinear maps [Lin16, LV16, Lin17, AS16], which all rely on so-called *local* pseudo-random generators (PRGs) – a PRG with *locality* L has every output bit depend on L input bits. It is known that if PRGs with locality L and polynomial stretch exist, then IO can be constructed from L -linear maps [Lin17, AS16]. Unfortunately, we do not even have locality-4 (polynomial stretch) PRGs [CM01, MST03], and candidate PRGs only exist starting from locality 5 [Gol01, MST03, OW14]. To circumvent the lower bound on PRG locality, we propose a new, relaxed, notion of locality, called *block-wise locality*. We build upon Lin’s [Lin17] recent IO construction, but show that in order to obtain IO from

L -linear maps, it suffices to use PRGs with block-wise locality L . As we will discuss below, such PRGs can exist for L as low as three.

Block-wise locality and IO We say that a PRG mapping $n \times \ell$ input bits to m output bits has *block-wise locality* L and *block-size* ℓ , if when viewing its input (i.e., the seed) as a matrix of $n \times \ell$ bits, every output bit depends on at most L columns in the matrix (as opposed to L input bits), as depicted in Figure 1. Observe that that the actual locality of such PRGs can go up to $L \times \ell$, yet, it has the special structure that all these input bits come from merely L input columns. This special structure is the key feature that allows for replacing local PRGs with block-wise-local PRGs, in the following applications.

- *Application I:* If there exists a *subexponentially-secure* PRG with block-wise-locality L , and any block-size $\ell = O(\log \lambda)$, then we can construct general-purpose IO from L -linear maps.
- *Application II:* If the block-wise local PRG is only *slightly superpolynomially secure*, we can still build special-purpose IO for circuits with super-logarithmic length inputs, which implies full-fledged Functional Encryption (FE), from L -linear maps.

All our constructions come with security reductions to (1) the security of block-wise-local PRGs, (2) the SXDH assumption on L -linear maps, and (3) the subexponential Learning With Errors (LWE), where 2) and 3) have the same level of hardness as that of the PRG.

Concurrently, we investigate the existence of block-wise local PRGs. In particular, we propose candidates following the common paradigm for candidate local PRGs [CM01, MST03, App12, OW14, AL16], which are variants of Goldreich’s functions [Gol00]. We simply replace every PRG input bit with a column of ℓ input bits. Such a block-wise local PRG is parameterized by a bipartite expander graph and a predicate (or potentially a set of predicates) over $L \times \ell$ input bits. We discuss the security of these candidates, against known attacks, in relation to the choice of graph and predicate. Furthermore, aiming at weakening the assumption on our candidates, we present two hardness amplification techniques that amplify respectively the weaker *next-bit-unpredictability* property and *pseudo-min-entropy generation* property to different levels of pseudo-randomness guarantees.

Instantiating the underlying multilinear maps We note that the results of this paper, per se, are merely new bootstrapping theorems, which do not rely, by themselves, on multilinear maps. More specifically, we show how to bootstrap a FE scheme for computing degree- L polynomials to an IO scheme, using a PRG with block-wise-locality L , and then rely on Lin’s [Lin17] FE construction. Some remarks on instantiations of the underlying multilinear maps are in order.

Concretely, the FE scheme from [Lin17] relies on algebraic L -linear maps, for which to date no candidate for $L \geq 3$ is known to exist. The alternative approach would be to instantiate them with existing noisy multilinear-map candidates. As discussed in [Lin17, Section 2.6], the existing proof would however fail in this case, in addition to the SXDH assumption itself being false on exiting noisy multilinear-map candidates. Still, a proof for ideal multilinear maps would be valid, but it is not known whether (1) existing cryptanalytic attacks can be adapted to break a construction, or (2) whether a proof in a weak ideal model as in [GMM⁺16] is possible.

Background on previous versions of this work In a previous version of this paper, we incorrectly claimed that our approach can be extended to bilinear maps. Two subsequent works, one by Barak *et al.* [BBKK17], the other by Lombardi and Vaikuntanathan [LV17], have presented attacks




Figure 1: Left: PRG with locality $L = 3$. Right: PRG with block-wise locality $L = 3$ and block size ℓ .

against PRGs with block-wise locality two. Strictly speaking, these results leave a narrow window of expansion factors open where block-wise PRGs could exist, but we are not aware whether our approach could be modified to use such low-stretch PRGs, or whether the attacks can be extended. We discuss these results more in detail further below in Section 1.3.

In contrast, attacks for $L \geq 3$ appear out of reach, as our assumption is implied by that made by recent works in the area of local PRGs and PRFs, c.f. e.g. the pseudorandomness assumptions from the recent work by Applebaum and Raykov [AR16] — and in fact, our amplification results show that even less needs to be achieved by the local function.

1.1 Block-Wise Locality

A $(n \times \ell, m)$ -PRG maps $n \times \ell$ input bits to m output bits. As introduced above, a PRG has block-wise locality L and block-size ℓ , if when viewing the input as a $n \times \ell$ matrix, every output bits depend on input bits in at most L columns. Such a function is fully specified by the input-output dependency graph G describing which input columns each output bit depends on, and the set of predicates $\{P_j\}_{j \in [m]}$ that each output bit is evaluated through.

In all our applications, we consider block-wise local PRGs with sufficiently large polynomial input- and output-lengths, n and m (in the security parameter λ) and logarithmic block-size $\ell = O(\log(\lambda))$. In this setting, a PRG has polynomial-stretch if $m = n^{1+\alpha}$ for some positive constant $\alpha > 0$. For convenience, below we assume such parameters are fixed in our discussion.

When compared with traditional local PRGs (which can be thought as the special case with block size $\ell = 1$), the advantage of block-wise local PRGs is that while they will still permit instantiations with L -linear maps in our applications, their output bits depend on $L \times \ell$ input bits, and hence we can use more complex, say logarithmic-degree, predicates. For this reason, known lower bounds on the locality of PRGs do not apply to block-wise locality, even when $L < 5$, when the block size satisfies $\ell = \Omega(\log(\lambda))$. Effectively, such PRGs can be seen as operating on input symbols with polynomial alphabet size. Moreover, the lower bounds in [CM01, MST03] show that for conventional locality, PRGs with polynomial stretch require $L \geq 5$, but they crucially rely on the fact that any locality-4 predicate is correlated with two of its input bits to rule out the existence of locality-4 PRGs. In contrast, a PRG with block-wise locality L can use predicates that depend on $L \log \lambda$ input bits; setting the predicate to be uncorrelated with any subset of $\log \lambda$ input bits circumvents the lower bound argument in [CM01, MST03].

Block-wise local PRGs via local PRGs Every function with block-wise locality L and block size ℓ is a function with locality $L\ell$. Therefore, the rich literature on the security of Goldreich’s local functions (see Applebaum’s survey [App15]) provides guidelines on how to choose candidate block-wise local PRGs, more specifically, the dependency graph G and predicates $\{P_j\}$. In par-

ticular, the graph G should be (k, c) -expanding, i.e., every subset of $k' \leq k$ output bits depends on at least $c \times k'$ input columns, for appropriately large k and c . We show that for $L \geq 3$, a large $1 - o(1)$ fraction of graphs G is $(n^{1-\eta}, (1-\eta)L)$ -expanding. This in turn means that we can think of this as an instance of Goldreich's function with locality $L\ell$ built from a graph which is $(n^{1-\eta}, (1-\eta)L\ell)$ -expanding, thus taking us back to the classical setting studied in the literature.

Using this analogy, we can show for example that for block-wise locality 3 and block size 2, for most graphs G , the resulting function withstands all linear attacks with sub-exponential bias ϵ when using the predicate outputting $x_1^0 \oplus x_2^0 \oplus x_3^0 \oplus (x_1^1 \wedge x_2^1)$ on input three columns $(x_1^0, x_1^1), (x_2^0, x_2^1), (x_3^0, x_3^1)$. This is a criterion that has been adopted so far to validate PRG security of local functions.

Moving even one step further, Applebaum and Raykov [AR16] recently postulated the following (even stronger) pseudorandomness assumption on functions with logarithmic locality:

Assumption 1 (Informal). *For locality $D = O(\log \lambda)$, and arbitrarily polynomial output length $m = n^{1+\alpha}$, there exist a suitable predicate, P' , such that, for any dependency graph G' that is $(n^{1-\eta}, (1-\eta)D)$ -expanding for some $0 < \eta < 1/2$, the locality- D function specified by P' and G' is $2^{-n^{1-\eta}}$ -pseudorandom against $2^{n^{1-\eta}}$ -time distinguishers.*

In our setting, for block-wise locality $L \geq 3$ and block-size $\log \lambda$, we show that when choosing the dependency graph G at random, the obtained block-wise local function can be thought as a function with locality $D = L \log \lambda$ satisfying the properties specified by the Applebaum-Raykov assumption, with $1 - o(1)$ probability. In particular, such functions withstand myopic inversion attacks (cf. e.g. [CEMT09]). In fact, our applications only need pseudorandomness to hold for output length $m = n^{1+\alpha}$ for some arbitrarily small constant $\alpha > 0$, and against polynomial time attackers, thus a much weaker requirement than what is guaranteed by the Applebaum-Raykov assumption.

For the case $L = 2$, the assumption that a block-wise local PRG exists is not backed by any of the past results, and indeed, recent works (following up on an earlier version of this paper) show that blockwise-local PRGs with sufficient stretch do not exist. We discuss this further below in Section 1.3.

Amplification In order to validate our assumptions even further, we present two transformations meant to enhance security of functions with block-wise locality. We consider two different techniques:

- *Amplification Technique I* produces a PRG construction with *quasi-polynomial* indistinguishability-gap (to polynomial-time distinguishers), from any *unpredictable generator* satisfying just *polynomial next-bit unpredictability* (i.e., the probability of predicting any output bit given previous output bits is at most $\frac{1}{2} + \frac{1}{\text{poly}(\lambda)}$, albeit for predictors in quasi-polynomial time). Though such PRGs are not strong enough for constructing IO, it suffices for constructing FE from L -linear maps; see the next section.
- *Amplification Technique II* produces a PRG construction with *sub-exponential* indistinguishability-gap, from certain special *pseudo-min-entropy-generator* whose output has sufficiently-high pseudo-min-entropy.

1.2 From Block-Wise Locality to IO and FE

We now move to an overview of our constructions from block-wise local PRGs.

IO from subexponentially secure block-wise-local PRGs Recent IO constructions from low-degree multilinear maps [LV16, Lin17, AS16] follow a common two-step approach: They first implement appropriate FE schemes, and then transform them into an IO scheme; we refer to the second step as the (FE-to-IO) *bootstrapping step*. In more detail, they use locality- L PRGs in the bootstrapping step in order to start with FE schemes that support only computation of *degree- L polynomials*; they then show that such FE schemes can be constructed from L -linear maps. In this work, following the blueprint and technique in [Lin17], we show how to replace the use of local PRGs with block-wise local PRGs within the bootstrapping step.

Theorem 1 (Bootstrapping using block-wise local PRGs). *Let L be any positive integer. There is a construction of IO for P/poly from the following primitives:*

- *Public-key fully-selectively-secure (collusion-resistant) FE for degree- L polynomials whose encryption time is linear in the input length (i.e., $\text{poly}(\lambda)N$);*
or with a secret-key FE scheme with the same properties, assuming additionally the subexponential hardness of LWE with subexponential modulus-to-noise ratio.
- *a PRG with block-wise locality L , block-size $\log \lambda$, and $n^{1+\alpha}$ -stretch for some positive constant α .*

where both FE and PRG need to have subexponential security.

The type of secret-key FE schemes for degree- L polynomials needed above was constructed by Lin [Lin17] assuming the SXDH assumption on L -linear maps.

Theorem 2 ([Lin17]). *Let L be any positive integer. Assuming the SXDH assumption on asymmetric L -linear maps, there is a construction of secret-key fully-selectively-secure (collusion-resistant) FE schemes for degree- L polynomials whose encryption time is linear in the input length (i.e., $\text{poly}(\lambda)N$). Moreover, the security reduction has a polynomial security loss.*

Therefore, combining our new bootstrapping theorem with Lin’s FE construction, we obtain IO from the subexponential SXDH assumption on L -linear maps, subexponentially-secure PRG with block-wise locality L , and subexponential LWE.

The power of super-polynomially secure block-wise local PRGs While constructing full-fledged IO for all polynomial-sized programs requires block-wise local PRGs with *subexponentially-security*, we ask what can be built from PRGs with weaker (slightly) *superpolynomial-security*. In particular, such PRGs can be obtained using the aforementioned *amplification technique I*, from unpredictable generator satisfying just polynomial next-bit unpredictability. To this end, we first give a parameterized version of Theorem 1 showing that if the PRG and L -linear maps are $(2^{-i\ell} \text{negl})$ -secure, then we can build IO schemes for circuits with $i\ell$ -bit inputs.

Theorem 3 (Parameterized version of Theorem 1). *Let L be any positive integer. Then, there is a construction of IO for the class of polynomial-sized circuits with $i\ell$ -bit inputs from the same primitives as in Theorem 1, and if FE and PRG are $(2^{-(i\ell+\kappa)} \text{negl})$ -secure, the resulting IO scheme is $(2^{-\kappa} \text{negl})$ -secure.*

Therefore, as discussed above, from slightly superpolynomially secure L -linear maps, a PRG with block-wise locality L , and subexponential LWE, we obtain IO for circuits with super-logarithmic, $\omega(\log \lambda)$, length inputs, and if the primitives are quasi-polynomially secure, we obtain IO for circuits with poly-logarithmic $\log^{1+\varepsilon}(\lambda)$ length inputs. Such IO schemes are already sufficient for two types of natural applications of IO:

- *Type 1*: Applications where IO is used to obfuscate a circuit with short inputs. For instance, for building FHE without relying on circular security [CLTV15], and constructing succinct randomized encoding for bounded space Turing machines [BGL⁺15]. In these applications, IO is used to obfuscate a circuit that receive as input an *index* from an arbitrary polynomial range.
- *Type 2*: Applications where the input length of the obfuscated circuit is determined by the security parameter of some other primitive. Then, by assuming exponential security of the other primitive, the input length can be made poly-logarithmic. For instance, as observed in [BNPW16, KS17], in the construction of public key encryption from one-way functions via IO, if assuming exponentially secure one-way functions, then IO for circuits with $\omega(\log \lambda)$ bit inputs suffices for the application.

We further show that IO for circuits with super-logarithmic length inputs implies full-fledged functional encryption.

Theorem 4 (Functional Encryption from $\omega(\log \lambda)$ -Input IO). *Let $i\ell$ be any super-logarithmic polynomial, that is, $i\ell = \omega(\log \lambda)$. Assume IO for the class of polynomial-sized circuits with $i\ell$ -bit inputs and public key encryption, both with $(2^{-i\ell} \text{negl})$ -security. Then, there exist collusion resistant (compact) public-key functional encryption for P/poly, satisfying adaptive-security.*

Combining the above two theorems, we immediately have that the existence of a PRG with block-wise locality L and L -linear maps, both with slightly super-polynomial security (and assuming subexponential LWE), implies the existence of full-fledged functional encryption, and all its applications, including, for instance, non-interactive key exchange (NIKE) for unbounded users [GPSZ16], trapdoor permutations [GPSZ16], PPAD hardness [BPR15, GPS16], publicly-verifiable delegation schemes in the CRS model [PRV12], and secure traitor tracing scheme [GGH⁺13b, BSW06, CFN94], which further implies hardness results in differential privacy [DNR⁺09, Ull13].

1.3 Subsequent works

Two recent works by Lombardi and Vaikuntanathan (LV) [LV17], and Barak, Brakerski, Komargodski, and Kothari (BBKK) [BBKK17] essentially rule out the existence of PRGs with block-wise locality $L = 2$, except for a very narrow window of expansion, as we explain next.

The LV attack The LV attack considers generators whose output bits are evaluated using the same predicate P , and whose dependency graph G is chosen at random. LV show that for any predicate P and a $1 - o(1)$ fraction of the graphs, the output can be efficiently distinguished from random, if its length reaches $\tilde{\Omega}(n2^\ell)$, where recall that $\ell = O(\log \lambda)$ is the block size. Their attack relies on two important ingredient. The first ingredient consists of techniques for refuting random L -CSPs over large q -ary alphabets, which corresponds to PRGs with block-wise locality- L and block-size $\ell = \log q$. Allen, O’Donnell, and Witmer [AOW15] presented an efficient algorithm for this, which succeeds when the number of constraints is roughly $\tilde{\Omega}(n^{L/2} \text{poly}(q)/\varepsilon^2)$, where ε controls the “quality” of refutation. The second ingredient is a novel structural lemma showing that any locality-2 balanced predicate P over alphabet \mathbb{Z}_q must be $(1/2 + O(1)/\sqrt{q})$ -correlated with a locality-2 predicate Q over the constant-sized alphabet \mathbb{Z}_{16} . Roughly speaking, to distinguish the output of a PRG with predicate P , they apply the refutation technique on CSPs w.r.t. the predicate Q correlated with P . This allows them to rule out PRGs with output length as short as $\tilde{\Omega}(n2^\ell)$.

The BBKK attack BBKK considered the more general case where the generators use an arbitrary set of predicates $\{P_j\}$ and arbitrary dependency graph G . They show that PRGs with block-wise locality 2 and output length $\tilde{\Omega}(n2^{2\ell})$ do not exist. The bound on the output length can be improved to $\tilde{\Omega}(n2^\ell)$ for the case where G is randomly chosen, and so is the predicate (in particular, the predicate is the same for all output bits). In fact, they proved a more general lower bound: There is no PRG whose outputs are evaluated using polynomials of degree at most d involving at most s monomials, and of output length $\tilde{O}(sn^{\lceil d/2 \rceil})$. Note that every block-wise locality L PRG can be written as such a generator, with $n2^\ell$ input bits, and using polynomials of degree L and at most $2^{L\ell}$ monomials. Their result is based on semidefinite programming and in particular the sum of squares (SOS) hierarchy.

BV and BBKK essentially rule out the existence of PRGs with block-wise locality 2, except for the corner case where the generator can use a set of different predicates $\{P_j\}$, a specific or random graph, and the output length is $\tilde{O}(n2^{(1+\varepsilon)\ell})$, for some $0 < \varepsilon < 1$. However, it is unclear to us whether PRGs with such small expansion is sufficient for constructing IO, or whether the attacks can be extended to cover this case.

Outline of this Paper

We review standard security notions (including security definitions for functional encryption) in Section 2. Section 3 discusses candidate constructions of block-wise local PRGs. Section 4 discusses our bootstrapping method using block-wise local PRGs. Finally, in Section 5, we discuss constructions of functional-encryption schemes in Section 5.

2 Preliminaries

Let \mathbb{Z} and \mathbb{N} denote the set of integers, and positive integers, respectively. Let $[n]$ denote the set $\{1, 2, \dots, n\}$. We use \mathcal{R} to denote either a ring, or an ensemble of rings $\mathcal{R} = \{\mathcal{R}_\lambda\}$, which will be clear in the context.

We denote by PPT probabilistic polynomial time Turing machines. The term *negligible* is used for denoting functions that are (asymptotically) smaller than any inverse polynomial. More precisely, a function $\nu(\star)$ from non-negative integers to reals is called *negligible* if for every constant $c > 0$ and all sufficiently large n , it holds that $\nu(n) < n^{-c}$.

We use boldface to denote vectors, for example, $\mathbf{u}, \mathbf{v}, \mathbf{c}$ etc., and use u_i, v_i, c_i to denote the i^{th} elements in the vectors.

2.1 μ -Hardness and μ -Indistinguishability

Definition 1 (μ -Hard One-Way Functions). *Let $\mu : \mathbb{N} \rightarrow [0, 1]$ be a function. A one-way function f is μ -hard if for every family of polynomial-sized adversaries $\{A_\lambda\}_{\lambda \in \mathbb{N}}$, and every sufficiently large security parameter $\lambda \in \mathbb{N}$, it holds that*

$$\Pr[x \xleftarrow{\$} \{0, 1\}^n; y = f(x) : f(A_\lambda(y)) = y] \leq \mu(\lambda)$$

Definition 2 (μ -indistinguishability). *Let $\mu : \mathbb{N} \rightarrow [0, 1]$ be a function. A pair of distribution ensembles $\{X_\lambda\}_{\lambda \in \mathbb{N}}, \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are μ -indistinguishable if for every family of polynomial-sized distinguishers $\{D_\lambda\}_{\lambda \in \mathbb{N}}$, and every sufficiently large security parameter $\lambda \in \mathbb{N}$, it holds that*

$$|\Pr[x \xleftarrow{\$} X_\lambda : D(1^\lambda, x, z) = 1] - \Pr[y \xleftarrow{\$} Y_\lambda : D(1^\lambda, y, z) = 1]| \leq \mu(\lambda)$$

Definition 3 (Computational and Sub-exponential Indistinguishability). *A pair of distribution ensembles $\{X_\lambda\}_{\lambda \in \mathbb{N}}$, $\{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable if they are $1/p$ -indistinguishable for every polynomial p , and are sub-exponentially indistinguishable if they are μ -indistinguishable for some sub-exponentially small $\mu(\lambda) = 2^{-\lambda^\varepsilon}$ with a constant $\varepsilon > 0$.*

Note that the above definition of sub-exponential indistinguishability is weaker than standard sub-exponential hardness assumptions that consider distinguishers running in sub-exponential time.

Below, we provide definitions of standard cryptographic primitives using the terminology of μ -indistinguishability, which implicitly defines variants with polynomial or sub-exponential security. As a matter of convention, we will drop μ when μ is a negligible function, and say sub-exponential security when μ is a sub-exponentially small function.

2.2 Indistinguishability Obfuscation

We recall the notion of indistinguishability obfuscation for a class of circuit defined by [BGI⁺01b].

Definition 4 (Indistinguishability Obfuscator ($i\mathcal{O}$) for a circuit class). *A uniform PPT machine $i\mathcal{O}$ is an indistinguishability obfuscator for a class of circuits $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$, if the following conditions are satisfied:*

Correctness: *For all security parameters $\lambda \in \mathbb{N}$, for every $C \in \mathcal{C}_\lambda$, and every input x , we have that*

$$\Pr[C' \leftarrow i\mathcal{O}(1^\lambda, C) : C'(x) = C(x)] = 1$$

where the probability is taken over the coin-tosses of the obfuscator $i\mathcal{O}$.

μ -Indistinguishability: *For every ensemble of pairs of circuits $\{C_{0,\lambda}, C_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ satisfying that $C_{b,\lambda} \in \mathcal{C}_\lambda$, $|C_{0,\lambda}| = |C_{1,\lambda}|$, and $C_{0,\lambda}(x) = C_{1,\lambda}(x)$ for every x , the following ensembles of distributions are μ -indistinguishable:*

$$\begin{aligned} & \left\{ C_{1,\lambda}, C_{2,\lambda}, i\mathcal{O}(1^\lambda, C_{1,\lambda}) \right\}_{\lambda \in \mathbb{N}} \\ & \left\{ C_{1,\lambda}, C_{2,\lambda}, i\mathcal{O}(1^\lambda, C_{2,\lambda}) \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

Definition 5 (IO for P/poly). *A uniform PPT machine $i\mathcal{O}_{\text{P/poly}}(\star, \star)$ is an indistinguishability obfuscator for P/poly if it is an indistinguishability obfuscator for the class $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ of circuits of size at most λ .*

2.2.1 $i\ell$ -bit-Input IO

In this work, we consider IO for polynomial-sized circuits with bounded input-length. Formally, we define IO for $i\ell$ -bit-input circuits, or referred to as $i\ell$ -bit-input IO, as follows.

Definition 6 ($i\ell$ -bit-input IO for P/poly). *A uniform PPT machine $i\mathcal{O}(\star, \star)$ is an $i\ell$ -bit-input indistinguishability obfuscator for P/poly if it is an indistinguishability obfuscator for the class $\{\mathcal{C}_\lambda^{i\ell}\}_{\lambda \in \mathbb{N}}$ of circuits with size at most λ and input-length at most $i\ell(\lambda) \leq \lambda$.*

2.3 Puncturable Pseudo-Random Functions

We recall the definition of puncturable pseudo-random functions (PPRF) from [SW14]. Since in this work, we only use puncturing at one point, the definition below is restricted to puncturing only at one point instead of at a polynomially many points.

Definition 7 (Puncturable PRFs). *Let n be a computable polynomial. A puncturable family of PRFs with domains $\{0, 1\}^{n(\lambda)}$ is given by a triple of uniform PPT machines $\text{PPRF} = (\text{PRF.Gen}, \text{PRF.Punc}, F)$ satisfying the following conditions:*

Correctness: *For every $\lambda \in \mathbb{N}$, and every output K of $\text{PRF.Gen}(1^\lambda)$, every input $i \in \{0, 1\}^{n(\lambda)}$, and $K\{i\} = \text{PRF.Punc}(K, i)$, we have that $F(K\{i\}, x) = F(K, x)$ for all $x \neq i$.*

μ -pseudorandomness at punctured point: *For every ensemble $\{i_\lambda \in \{0, 1\}^{n(\lambda)}\}$, the following ensembles (where $i = i_\lambda$) are μ -indistinguishable.*

$$\left\{ K \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda), K\{i\} = \text{PRF.Punc}(K, i) : K\{i\}, i, F(K, i) \right\}$$

$$\left\{ K \stackrel{\$}{\leftarrow} \text{PRF.Gen}(1^\lambda), K\{i\} = \text{PRF.Punc}(K, i) : K\{i\}, i, U_\lambda \right\}$$

As observed by [BW13, BGI14, KPTZ13], the GGM tree-based construction of PRFs [GGM86] from one-way functions yields PPRFs. Furthermore, their construction incurs only a polynomial security loss, and hence, if the underlying one-way functions are μ -hard, then the resulting PPRF is μ -pseudorandom.

2.4 Randomized Encodings

In this section, we recall the traditional definition of randomized encodings with simulation security [IK02, AIK06].

Definition 8 (Randomized encoding scheme for circuits). *A randomized encoding scheme \mathbf{RE} consists of two PPT algorithms,*

- $\hat{C}_x \stackrel{\$}{\leftarrow} \text{REnc}(1^\lambda, C, x)$: *On input a security parameter 1^λ , circuit C , and input x , REnc generates an encoding \hat{C}_x .*
- $y = \text{REval}(\hat{C}_x)$: *On input \hat{C}_x produced by REnc , REval outputs y .*

Correctness: *The two algorithms REnc and REval satisfy the following correctness condition: For all security parameters $\lambda \in \mathbb{N}$, circuit C , input x , it holds that,*

$$\Pr[\hat{C}_x \stackrel{\$}{\leftarrow} \text{REnc}(1^\lambda, C, x) : \text{Eval}(\hat{C}_x) = C(x)] = 1$$

μ -Simulation Security: *There exists a PPT algorithm RSim , such that, for every ensemble $\{C_\lambda, x_\lambda\}_\lambda$ where $|C_\lambda|, |x_\lambda| \leq \text{poly}(\lambda)$, the following ensembles are μ -indistinguishable for all $\lambda \in \mathbb{N}$.*

$$\left\{ \hat{C}_x \stackrel{\$}{\leftarrow} \text{REnc}(1^\lambda, C, x) : \hat{C}_x \right\}_{\lambda \in \mathbb{N}}$$

$$\left\{ \hat{C}_x \stackrel{\$}{\leftarrow} \text{RSim}(1^\lambda, C(x), 1^{|C|}, 1^{|x|}) : \hat{C}_x \right\}_{\lambda \in \mathbb{N}}$$

where $C = C_\lambda$ and $x = x_\lambda$.

Furthermore, let \mathcal{C} be a complexity class, we say that randomized encoding scheme \mathbf{RE} is in \mathcal{C} , if the encoding algorithm REnc can be implemented in that complexity class.

2.5 Functional Encryption

We provide the definition of a public-key functional encryption (FE) scheme with indistinguishability-based security which originally appeared in [BSW12, O'N10]. Below we define public key FE first, and then note the difference with secret key FE.

2.5.1 Public-Key Functional Encryption

Syntax Let $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$ be ensembles of sets. Let $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, where every function in the set \mathcal{F}_λ maps inputs in \mathcal{X}_λ to outputs in \mathcal{Y}_λ .

A public-key functional encryption scheme **FE** for $\{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ consists of four PPT algorithms (FE.Setup, FE.KeyGen, FE.Enc, FE.Dec).

- *Setup*: FE.Setup($1^\lambda, \text{pp}$) is an algorithm that on input a security parameter and some public parameter (e.g., description of bilinear pairing groups) outputs a master public key and a master secret key (MPK, MSK).
- *Key Generation*: FE.KeyGen(MSK, f) on input the master secret key MSK and the description of a function $f \in \mathcal{F}_\lambda$, outputs a secret key SK_f .
- *Encryption*: FE.Enc(MPK, x) on input the master public key MPK and a message $x \in \mathcal{X}_\lambda$, outputs an encryption CT of x .
- *Decryption*: FE.Dec(SK, CT) on input the secret key associated with f and an encryption of x , outputs $y \in \mathcal{Y}_\lambda$.

Correctness: We define perfect correctness here. For every $\lambda, f \in \mathcal{F}_\lambda, x \in \mathcal{X}_\lambda$, it holds that,

$$\Pr \left[\begin{array}{l} (\text{MPK}, \text{MSK}) \stackrel{\$}{\leftarrow} \text{FE.Setup}(1^\lambda, \text{pp}) \\ \text{CT} \stackrel{\$}{\leftarrow} \text{FE.Enc}(\text{MPK}, x) \\ \text{SK} \stackrel{\$}{\leftarrow} \text{FE.KeyGen}(\text{MSK}, f) \end{array} : f(x) = \text{FE.Dec}(\text{SK}, \text{CT}) \right] = 1$$

Indistinguishability Security. Indistinguishability security of a functional encryption requires that no adversary can distinguish the FE encryption of one input x_0 from that of another x_1 , if the adversary only obtains secret keys for functions that yield the same outputs on x_0 and x_1 , that is, for every secret key SK_f , it holds that $f(x_0) = f(x_1)$. In the adaptive setting, the two challenge inputs (x_0, x_1) and all functions f are chosen adaptively by the adversary. In the weaker fully-selective setting, the adversary is restricted to choose (x_0, x_1) and all functions f statically.

Definition 9 (Adap-security). A public-key FE scheme $\mathbf{FE} = (\text{FE.Setup}, \text{FE.KeyGen}, \text{FE.Enc}, \text{FE.Dec})$ for $\{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is μ -Adap-secure, if for every PPT adversary A , and every sufficiently large security parameter $\lambda \in \mathbb{N}$, the adversary's advantage in the following games is bounded by $\mu(\lambda)$

$$\text{Adv}_A^{\mathbf{FE}} = \left| \Pr[\text{Adap}_A^{\mathbf{FE}}(1^\lambda, 0) = 1] - \Pr[\text{Adap}_A^{\mathbf{FE}}(1^\lambda, 1) = 1] \right| \leq \mu(\lambda)$$

$\text{Adap}_A^{\mathbf{FE}}(1^\lambda, b)$ proceeds as follows:

1. **Key Generation.** The challenger CH samples $(\text{MPK}, \text{MSK}) \stackrel{\$}{\leftarrow} \text{FE.Setup}(1^\lambda, \text{pp})$ and sends MPK to the adversary.

2. **Function Queries I.** Repeat the following for an arbitrary number of times determined by A : Upon A choosing a function query $f \in \mathcal{F}_\lambda$, CH sends A a function key $SK_f \xleftarrow{\$} \text{FE.KeyGen}(\text{MSK}, f)$.
3. **Message Queries.** Upon A choosing a pair of messages (x_0, x_1) , CH sends A a ciphertext $\text{CT} \xleftarrow{\$} \text{FE.Enc}(\text{MPK}, x_b)$.
4. **Function Queries II.** Repeat the second step, for an arbitrary number of times determined by A .
5. Finally A outputs a bit b' which is also the output of the experiment.

Restriction: Every function query f must satisfy that $f(x_0) = f(x_1)$.

Definition 10 (Full-Sel-security). We say that **FE** is μ -Full-Sel-secure if the condition in Definition 9 holds for modified experiments $\text{Full-Sel}_A^{\text{FE}}(1^\lambda, b)$ that proceeds identically to $\text{Adap}_A^{\text{FE}}(1^\lambda, b)$ except that the adversaries choose challenge messages (x_0, x_1) and all function queries $\{f\}$ at the beginning of the experiment.

Note that our notion of fully-selective security is weaker than the notion of selective security in some papers in the literature (e.g., [GKP⁺13, ABSV15]), which only requires the adversaries to choose challenge inputs x_0, x_1 statically, but allows the adversaries to choose challenge function inputs adaptively. Intuitively, the notion of fully-selective security is sufficient for applications that are non-interactive, for instance, building IO from FE as in [BV15, AJ15].

Definition 11 (1-key FE). We say that **FE** is a μ -Adap-secure (or μ -Full-Sel-secure) 1-key FE scheme if it satisfies the security requirements in Definition 9 (or, respectively, Definition 10) against adversaries that ask for at most one function key query.

2.5.2 FE for P/poly, NC^1 and Compactness

Definition 12 (FE schemes for families of function classes). Let $\mathbb{F} = \{\mathcal{F}^I\}_{I \in \mathcal{I}}$ be a family of function classes. We say that $\mathcal{FE} = \{\mathbf{FE}^I\}_{I \in \mathcal{I}}$ is a family of (1-key) FE schemes for \mathbb{F} with μ -Adap-security or μ -Full-Sel-security if for every function class $\mathcal{F}^I = \{\mathcal{F}_\lambda^I\}_{\lambda \in \mathbb{N}}$, \mathbf{FE}^I is a (1-key) FE scheme for \mathcal{F}^I with μ -Adap-security or μ -Full-Sel-security.

Moreover, define the following special cases:

- **FE for P/poly** is a family of FE schemes for $\mathbb{F} = \{\mathcal{F}^{N,D,S}\}_{N \in \mathcal{N}, D \in \mathcal{D}, S \in \mathcal{S}}$, where $\mathcal{N}, \mathcal{D}, \mathcal{S}$ are the sets of all polynomials and $\mathcal{F}^{N,D,S}$ is the class of binary functions that can be computed by circuits with $N(\lambda)$ -bit inputs, $S(\lambda)$ size, and $D(\lambda)$ depth.
- **FE for NC^1** is a family of FE schemes for $\mathbb{F} = \{\mathcal{F}^{N,D,S}\}_{N \in \mathcal{N}, D \in \mathcal{D}, S \in \mathcal{S}}$ as defined above but with \mathcal{D} the set of all logarithmic functions.

Compactness In the above definition of families of FE schemes, algorithms in scheme $\mathbf{FE}^{N,D,S}$ could run in polynomial time depending on polynomials N, D, S . In the literature, stronger efficiency requirements have been considered. In particular, the works of [AJ15, BV15] defined compact FE schemes for NC^1 , which requires the encryption time to be independent of the circuit size S of the functions.

Definition 13 (Compactness of FE schemes for NC^1). Let $\mathcal{FE} = \{\mathbf{FE}^{N,D,S}\}$ be a family of FE schemes for NC^1 .

Compactness: We say that the functional encryption scheme \mathcal{FE} is compact if for every logarithmic function D , there is a polynomial p , such that, for every polynomials N, S , the encryption algorithm of $\mathbf{FE}^{N,D,S}$ runs in time $p(\lambda, N(\lambda), \log S(\lambda))$.

$(1 - \varepsilon)$ -Sublinear Compactness (a.k.a. $(1 - \varepsilon)$ -Weakly Compactness): We say that \mathcal{FE} is $(1 - \varepsilon)$ -sublinearly compact, if for every logarithmic function D , there is a polynomial p , such that, for every polynomials N, S , the encryption algorithm of $\mathbf{FE}^{N,D,S}$ runs in time $p(\lambda, N(\lambda)) \cdot S(\lambda)^{1-\varepsilon}$.

2.6 Zero-Testing FE for Arithmetic Functions

For any ring \mathcal{R} , we refer to functions mapping from \mathcal{R}^* to \mathcal{R}^* as arithmetic functions in \mathcal{R} . Many previous works (e.g. [ABCP15, BJK15]) constructed FE schemes for classes of arithmetic functions in \mathcal{R} with a relaxed correctness guarantee, namely, decryption does not reveal the output (in \mathcal{R}) entirely, but only reveals whether the output is zero or not. We refer to this relaxed correctness guarantee as *zero-testing correctness*, and FE schemes with such relaxed correctness as *zero-testing FE*. We stress that though the correctness requirement is relaxed, the security requirements, namely IND-security and function hiding, remain the same. Therefore, zero-testing FE is strictly weaker than standard FE.

Definition 14 (Zero-testing FE). Let $\mathcal{R} = \{\mathcal{R}_\lambda\}$ be an ensemble of rings, and $\{\mathcal{F}_\lambda\}$ a class of functions where \mathcal{F}_λ maps from $\mathcal{X}_\lambda \subseteq \mathcal{R}_\lambda^*$ to $\mathcal{Y}_\lambda \subseteq \mathcal{R}_\lambda^*$. We say that \mathbf{FE} is a (1-key) zero-testing FE scheme for $\{\mathcal{F}_\lambda\}$ with μ -Adap-security or μ -Full-Sel-security, if it is a FE scheme for $\{\mathcal{F}_\lambda\}$ with the same security guarantee as in Definition 9 or 10 respectively, and the following relaxed correctness guarantee.

- **Zero-Testing Correctness:** For every $\lambda, f \in \mathcal{F}_\lambda, x \in \mathcal{X}_\lambda$, it holds that,

$$\Pr \left[\begin{array}{l} (\text{MPK}, \text{MSK}) \xleftarrow{\$} \text{FE.Setup}(1^\lambda, \text{pp}) \\ \text{CT} \xleftarrow{\$} \text{FE.Enc}(\text{MPK}, x) \\ \text{SK} \xleftarrow{\$} \text{FE.KeyGen}(\text{MSK}, f) \end{array} : \text{ZT}(f(x)) = \text{FE.Dec}(\text{SK}, \text{CT}) \right] = 1$$

where ZT is a predicate that outputs 1 iff its input is the zero element in \mathcal{R}_λ , and in the case of secret key FE, $\text{MPK} = \text{MSK}$.

Zero-Testing FE for Degree- d Polynomials

Definition 15 (Zero-testing FE schemes for families of arithmetic function classes). Let $\mathbb{F} = \{\mathcal{F}^I\}_{I \in \mathcal{I}}$ be a family of arithmetic function classes. A family $\mathcal{FE} = \{\mathbf{FE}^I\}_{I \in \mathcal{I}}$ of (1-key) zero-testing FE schemes for \mathbb{F} is defined identically as in Definition 12 except that every scheme \mathbf{FE}^I has zero-testing correctness.

Moreover, define the following special cases:

- **Zero-testing FE for degree- d polynomials in \mathcal{R}** is a family of zero-testing FE schemes for $\mathbb{F} = \{\mathcal{F}^N\}$ where \mathcal{F}^N is the set of degree- d polynomials mapping from $\mathcal{R}_\lambda^{N(\lambda)}$ to \mathcal{R}_λ .

Definition 16 (Linear efficiency). Let $\mathcal{FE} = \{\mathbf{FE}^N\}$ be a family of FE schemes for degree- d polynomials or inner products in \mathcal{R} . We say that \mathcal{FE} has **linear efficiency** if there exists a polynomial function p , such that, for every polynomial N , the encryption algorithm of \mathbf{FE}^N runs in time $N(\lambda) \text{poly}(\lambda)$.

In the rest of the paper, whenever we talk about FE for arithmetic functions, in particular, FEs for degree- d polynomials, over a family of non-binary ring \mathcal{R} , we mean by default a zero-testing FE.

2.7 Degree- D Asymmetric Multilinear Maps with SXDH Assumption

Introduced by Boneh and Silverberg [BS02], asymmetric Multilinear Maps (MMaps) naturally generalize asymmetric bilinear maps to higher degree. Let \mathcal{G} denote a group generator that on input 1^λ outputs $(p, G_1, \dots, G_D, G_{D+1}, \mathbf{pair})$, where G_1, \dots, G_D, G_{D+1} are cyclic groups with order p (prime or composite). G_1 to G_D are referred to as the source groups and G_{D+1} the target group. Assume without loss of generality that the description of the source groups contain generators g_1, \dots, g_D of G_1, \dots, G_D . In addition, the following properties hold.

- *Admissible:* $\mathbf{pair} : G_1 \times \dots \times G_D \rightarrow G_{D+1}$ is efficiently computable and $g_{D+1} = \mathbf{pair}(g_1 \dots, g_D)$ generates G_{D+1} .
- *Multilinear:* For any $a_1, \dots, a_D \in \mathbb{Z}_p$, $\mathbf{pair}(g_1^{a_1}, \dots, g_D^{a_D}) = \mathbf{pair}(g_1, \dots, g_D)^{a_1 a_2 \dots a_D} = g_{D+1}^{a_1 a_2 \dots a_D}$.

We denote by $\mathcal{R}_\lambda = (\mathbb{Z}_p, +, \times)$ the ring corresponding to the exponent space of these multilinear pairing groups.

The Bracket Notation For clarity of notions, we use the following bracket notations to denote group elements.

$$\forall l \in [D + 1], \quad [a]_l = g_l^a$$

We refer to $[a]_l$ as an encoding of a in group G_l , or with label l . Under this notation, the generator in group $l \in [D + 1]$ is represented as $[1]_l = g_l$. We also use the following vector notation to represent vectors of group elements succinctly: For any $\mathbf{v} = (v_1, \dots, v_m) \in \mathbb{Z}_p^m$, and $l \in \{0, 1, T\}$:

$$[\mathbf{v}]_l = [v_1]_l \dots [v_m]_l$$

The SXDH Assumption The SXDH assumption states that the standard DDH assumption holds in each of the source groups. Formally, for every source group G_l for $l \in [D]$, the following two ensembles are μ -indistinguishable.

$$\left\{ \text{pp} = (p, G_1, \dots, G_D, G_{D+1}, \mathbf{pair}) \stackrel{\$}{\leftarrow} \mathcal{G}(1^\lambda), a, b \stackrel{\$}{\leftarrow} \mathbb{Z}_p : (\text{pp}, [a]_l, [b]_l, [ab]_l) \right\}_\lambda$$

$$\left\{ \text{pp} = (p, G_1, \dots, G_D, G_{D+1}, \mathbf{pair}) \stackrel{\$}{\leftarrow} \mathcal{G}(1^\lambda), a, b, r \stackrel{\$}{\leftarrow} \mathbb{Z}_p : (\text{pp}, [a]_l, [b]_l, [r]_l) \right\}_\lambda$$

3 Block-Wise Local PRGs

In this section, we introduce the notion of a block-wise local PRG. We start with formal definitions, in Section 3.1, which we refer to throughout the rest of the paper. Then, the remaining sub-sections will discuss a graph-based framework for block-wise local functions, and discuss candidates.

3.1 Pseudorandom Generators, Locality, and Block-Wise Locality

We review the notion of a PRG family, and its locality.

Definition 17 (Family of Pseudo-Random Generators (PRGs)). *Let n and m be polynomials. A family of $(n(\lambda), m(\lambda))$ -PRG is an ensemble of distributions $\mathbf{PRG} = \{\mathbf{PRG}_\lambda\}$ satisfying the following properties:*

Syntax: For every $\lambda \in \mathbb{N}$, every PRG in the support of \mathbf{PRG}_λ defines a function mapping $n(\lambda)$ bits to $m(\lambda)$ bits.

Efficiency: There is a uniform Turing machine M satisfying that for every $\lambda \in \mathbb{N}$, every PRG in the support of \mathbf{PRG}_λ , and every $x \in \{0, 1\}^{n(\lambda)}$, $M(\text{PRG}, x) = \text{PRG}(x)$.

μ -Indistinguishability: The following ensembles are μ -indistinguishable

$$\left\{ \text{PRG} \stackrel{\$}{\leftarrow} \mathbf{PRG}_\lambda; s \stackrel{\$}{\leftarrow} \{0, 1\}^{n(\lambda)} : (\text{PRG}, \text{PRG}(s)) \right\}_{\lambda \in \mathbb{N}} \\ \approx_\mu \left\{ \text{PRG} \stackrel{\$}{\leftarrow} \mathbf{PRG}_\lambda; r \stackrel{\$}{\leftarrow} \{0, 1\}^{m(\lambda)} : (\text{PRG}, r) \right\}_{\lambda \in \mathbb{N}}$$

Definition 18 (Block-Wise Locality of PRGs). Let n, m, L , and ℓ be polynomials. We say that a family of $(n(\lambda)\ell(\lambda), m(\lambda))$ -PRGs has block-wise locality- $(L(\lambda), \ell(\lambda))$ if for every λ and every PRG in the support of \mathbf{PRG}_λ , inputs of PRG are viewed as $n(\lambda) \times \ell(\lambda)$ matrices of bits, and every output bit of PRG depends on input bits contained in at most $L(\lambda)$ columns.

3.2 Graph-Based Block-Wise local Functions

In this section, we discuss candidate PRGs with block-wise locality d , where d can be as small as two. Here, we start with the notational framework and then move on to discussing concrete assumptions on them in Section 3.3.

Goldreich’s function We will consider local functions based on Goldreich’s construction [Gol00], which have been the subject for extensive study (cf. e.g. Applebaum’s survey [App15]).

Recall first that an $[n, m, d]$ -hypergraph is a collection $G = (S_1, \dots, S_m)$ where the *hyperedges* S_i are elements of $[n]^d$, i.e., $S_i = (i_1, \dots, i_d)$, where $i_j \in [n]$ (note that we allow for potential repetitions, merely for notational convenience). We use hypergraphs to build functions as follows.

Definition 19 (Goldreich’s function). Let $\mathbf{G} = \{\mathbf{G}_\lambda\}_{\lambda \in \mathbb{N}}$ be an ensemble such that \mathbf{G}_λ is a distribution on $[n(\lambda), m(\lambda), d(\lambda)]$ -hypergraphs, for polynomial functions m, n, d . Also let $P = \{P_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of predicates, where P_λ operates on $d(\lambda)$ -bit strings. Then, define the function ensemble $\mathbf{GF}^{\mathbf{G}, P} = \{\mathbf{GF}_\lambda^{\mathbf{G}, P}\}_{\lambda \in \mathbb{N}}$, where $\mathbf{GF}_\lambda^{\mathbf{G}, P}$ samples first a graph $G = (S_1, \dots, S_m) \stackrel{\$}{\leftarrow} \mathbf{G}_\lambda$, and then outputs the function $\text{GF}_{G, P} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ such that for all n -bit x ,

$$\text{GF}_{G, P}(x) = (y_1, \dots, y_m), \quad y_i = P(x[S_i]),$$

where $x[S]$ denotes the d -bit sub-string obtained by concatenating the bits at positions indexed by S .¹

Functions with block-wise locality We want to extend the notation used above to consider the case where an edge of G does not solely give a pointer to individual bits to be injected in the computation, but rather, to “chunks” consisting of ℓ -bit strings, and the predicate is applied to the concatenation of these bits. The resulting function clearly then satisfies block-wise locality d with block size ℓ .

¹The notion could be block-wise to the cases where predicates are drawn by a distribution, and possibly differ from each output bit. We are going to dispense with such extensions, which are straightforward but easily lead to notational overhead.

Definition 20 (Block-wise local graph-based function). Let $\mathbf{G} = \{\mathbf{G}_\lambda\}_{\lambda \in \mathbb{N}}$ be such that \mathbf{G}_λ is a distribution on $[n(\lambda), m(\lambda), d(\lambda)]$ -hypergraphs, for polynomial functions m, n, d . Also let $\ell(\lambda)$ be a polynomial function, and $P = \{P_\lambda\}_{\lambda \in \mathbb{N}}$ a family of predicates, where P_λ operates on $(d(\lambda) \times \ell(\lambda))$ -bit strings. Then, define the function ensemble $\mathbf{GF}^{\mathcal{G}, P, \ell} = \{\mathbf{GF}_\lambda^{\mathcal{G}, P, \ell}\}_{\lambda \in \mathbb{N}}$, where $\mathbf{GF}_\lambda^{\mathcal{G}, P, \ell}$ samples first a graph $G = (S_1, \dots, S_m) \stackrel{\$}{\leftarrow} \mathbf{G}_\lambda$, and then outputs the function $\text{GF}_{G, P, \ell} : \{0, 1\}^{n \cdot \ell} \rightarrow \{0, 1\}^m$ such that for all $(n \times \ell)$ -bit inputs $\mathbf{x} = (\mathbf{x}[1], \dots, \mathbf{x}[n])$, where $\mathbf{x}[1], \dots, \mathbf{x}[n] \in \{0, 1\}^\ell$,

$$\text{GF}_{G, P, \ell}(x) = (y_1, \dots, y_m), \quad y_i = P(\mathbf{x}[S_i]),$$

where $\mathbf{x}[S]$ denotes the $d \cdot \ell$ -bit sub-string obtained by concatenating ℓ -bit input chunks indexed by S .

We typically refer to the graph G describing $\text{GF}_{G, P, \ell}$ as the *base graph*. This is because $\text{GF}_{G, P, \ell}$ can be seen as a special case of Goldreich's function defined above, for a suitable graph. Namely, the base graph G can be extended to an $[n \cdot \ell, m, d\ell]$ -hypergraph \bar{G} naturally, where each edge $S_i = (i_1, \dots, i_d)$ from G is mapped into a new hyper-edge \bar{S}_i with $d \cdot \ell$ elements such that

$$\bar{S}_i = ((i_1 - 1) \cdot \ell + 1, \dots, i_1 \cdot \ell, \dots, (i_d - 1) \cdot \ell + 1, \dots, i_d \cdot \ell),$$

then clearly $\text{GF}_{G, P, \ell} = \text{GF}_{\bar{G}, P, 1} = \text{GF}_{\bar{G}, P}$. This view will be convenient to connect back to the body of work on studying the security of Goldreich's function on suitable graphs, for which our block-wise local designs serve as a special case.

Expansion properties In general, we will want to instantiate our framework with functions where the base graph G is a good expander graph. Recall the following.

Definition 21. $G = (S_1, \dots, S_m)$ is a (k, c) -expander (or, equivalently, is (k, c) -expanding) if for all sets $J \subseteq [m]$ with $|J| \leq k$, we have $|\bigcup_{j \in J} S_j| \geq c \cdot |J|$.

Ideally, we will want in fact \bar{G} to be a good expander (in order to resort to large body of analyses for such functions). This will follow by making the base graph a good expander. In particular, the following simple fact stems from the observation that when going from G to \bar{G} , we have $|\bar{S}_j| = \ell |S_j|$, and hence the (relative) expansion factors of G and \bar{G} are identical.

Lemma 1. Let G be an $[n, m, d]$ -hypergraph which is $(k, (1 - \gamma)d)$ -expanding. Then, for any block-size ℓ , the resulting $[n \cdot \ell, m, d\ell]$ -hypergraph \bar{G} is $(k, (1 - \gamma)d\ell)$ -expanding.

In general, if we have high degree (say $O(\log \lambda)$), we can prove the existence (at least probabilistically) of very good expanders with expansion rate very close to the degree. Unfortunately, our construction of \bar{G} imposes some structure, and the actual expansion factor is dictated by the graph G with much lower degree d . The following lemma establishes the existence of good expander graphs, which we summarize below in a corollary with more useful parameters. While the proof of the lemma is folklore (we take notational inspiration from the one in [ABR16]), we give a more careful analysis tailored at a tight characterization for low degrees.

Lemma 2 (Strong expansion lemma). Let $d \geq 2$, and let $\gamma \in (0, 1)$ and $\beta \in (0, 1/2)$ be such that $d\gamma = 1 + \beta$. Further, let $1 \leq \Delta \leq n^\beta / \log(n)$. Then, there exists a constant $\alpha > 0$ such that a random $[n, m = \Delta n, d]$ -hypergraph G is a $(k = \alpha n / \Delta^{1/\beta}, d(1 - \gamma))$ -expander with probability $1 - o(1)$.

Proof. We pick a random $[n, \Delta n, d]$ -hypergraph $G = (S_1, \dots, S_m)$. Then, by a standard argument, with $c = (1 - \gamma)d$, the probability that G is not (k, c) -expanding is upper bounded by

$$\begin{aligned} \sum_{r=1}^k \binom{\Delta n}{r} \binom{n}{c \cdot r} \left(\frac{cr}{n}\right)^{d \cdot r} &\leq \sum_{r=1}^k \left(\frac{e\Delta n}{r}\right)^r \left(\frac{en}{c \cdot r}\right)^{cr} \left(\frac{cr}{n}\right)^{d \cdot r} \\ &= \sum_{r=1}^k \left(\frac{e^{c+1}\Delta n}{r}\right)^r \left(\frac{cr}{n}\right)^{\gamma dr} \\ &= \sum_{r=1}^k \left(\frac{e^{c+1}\Delta c^{\gamma d}}{(n/r)^{\gamma d - 1}}\right)^r = \sum_{r=1}^k \left(\frac{C \cdot \Delta}{(n/r)^{\gamma d - 1}}\right)^r, \end{aligned}$$

where $C = C_{\gamma, d} = e^{c+1}c^{\gamma d}$ is a constant which only depends on γ and d . Now, with $p_r = \left(\frac{C \cdot \Delta}{(n/r)^{\gamma d - 1}}\right)^r$, note that because $1 \leq \Delta \leq n^\beta / \log(n)$ and $\gamma d - 1 = \beta$,

$$p_r \leq \left(\frac{C \Delta r^\beta}{n^\beta}\right)^r \leq \left(\frac{C r^\beta}{\log(n)}\right)^r.$$

To compute $\sum_{r=1}^k p_r$, we partition its summands into three different sets:

- For $r = 1, \dots, \lfloor 1/\beta \rfloor$, we have that the sum of the p_r 's in this range is $o(1)$, because $\lfloor 1/\beta \rfloor \geq 1$ is constant.
- For $\lfloor 1/\beta \rfloor \leq r \leq 10 \log(n)$, then we have

$$p_r \leq \left(\frac{C 10^\beta \log(n)^\beta}{\log(n)}\right)^{1/\beta} \leq O\left(\frac{1}{\log^{\frac{1}{\beta}-1}(n)}\right),$$

and because $1/\beta - 1 > 1$, the sum of the p_r 's in this range is also $o(1)$.

- For $10 \log(n) \leq r \leq n/(2C \cdot \Delta)^{1/\beta}$, we have $p_r \leq (\frac{1}{2})^{10 \log(n)} \leq O(n^{-10})$. However, note that the number of r 's in this range is at most $O(n)$, and thus the sum of the p_r 's in this range is also $o(1)$.

This concludes the proof. □

Corollary 1. *For every γ and d such that $1 < \gamma d < 1.5$, and every $\eta \in (0, 1)$, there exists a $[n, n^{1+\zeta}, d]$ -hypergraph (for some $\zeta > 0$) which is a $(n^{1-\eta}, (1 - \gamma)d)$ -expander.*

Proof. With $\beta = \gamma d - 1$, set $\Delta = n^{\eta\beta} / \log(n)$. Then, $\Delta^{1/\beta} < n^\eta$ for large enough n , and thus $k > n^{1-\eta}$. □

3.3 Pseudorandom and Unpredictability Generators

We are interested in the question of finding $[n, m, d]$ -hypergraphs for $m = n^{1+\alpha}$ and a constant $d \geq 2$ such that $\text{GF}_{G, P, \ell}$ is a good PRG, for $\ell = O(\log \lambda)$. We consider a parameterized assumption on such functions (in terms of unpredictability), and discuss it briefly. Below, in Sections 3.5 and 3.6, we are then going to show how strong indistinguishability follows from (potentially) weaker versions of this assumption.

Unpredictability generator and assumptions Let $\mathbf{UG} = \{\mathbf{UG}_\lambda\}_{\lambda \in \mathbb{N}}$ be a function ensemble, where \mathbf{UG}_λ is a distribution on functions from $n(\lambda)$ to $m(\lambda)$ bits, for some polynomial functions m and n .

Definition 22 (Unpredictability generator). *We say that \mathbf{UG} is an (s, δ) -unpredictability generator (or (s, δ) -UG, for short) if for all (non-uniform) adversaries $A = \{A_\lambda\}_{\lambda \in \mathbb{N}}$ with size at most $s(\lambda)$ and all sequences of indices $i(\lambda) \in \{0, \dots, i(\lambda) - 1\}$, we have*

$$\Pr \left[\begin{array}{l} x \stackrel{\$}{\leftarrow} \{0, 1\}^{n(\lambda)} \\ \mathbf{UG} \stackrel{\$}{\leftarrow} \mathbf{UG}_\lambda \end{array} : A_\lambda(\mathbf{UG}, \mathbf{UG}_{\leq i(\lambda)}(x)) = \mathbf{UG}_{i(\lambda)+1}(x) \right] \leq \frac{1}{2} + \delta(\lambda),$$

where $\mathbf{UG}_{\leq j}(x)$ and $\mathbf{UG}_j(x)$ denote the first j bits and the j -th bit of $\mathbf{UG}(x)$, respectively.

Note that by a standard argument, being a (s, δ) -UG implies being a (family of) $(s, O(m \cdot \delta))$ -PRGs. We now consider the following assumption, which parametrizes the fact that $\mathbf{GF}_{G,P,\ell}$ is a good PRG.

Definition 23 (BLUG-assumption). *Let $n, \ell, s : \mathbb{N} \rightarrow \mathbb{N}$, and let $d \geq 2$ and $\alpha > 0$ be constants. Also, let $\delta : \mathbb{N} \rightarrow [0, 1]$. Then, the (d, ℓ) -BLUG(n, α, s, δ) assumption is the assumption that there exists a family $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$ of $[n(\lambda), n(\lambda)^{1+\alpha}, d]$ hypergraphs, and a family $P = \{P_\lambda\}_{\lambda \in \mathbb{N}}$ of predicates on $(d(\lambda) \times \ell(\lambda))$ -bit strings such that $\mathbf{GF}^{G,P,\ell}$ is an (s, δ) -UG.*

We are being a bit informal here, in the sense that obviously we would like $\mathbf{GF}^{G,P,\ell}$ to additionally be efficiently computable in a uniform sense. Our candidates will not have this property, as we are only able to infer the existence of suitable G 's probabilistically. There are two ways of thinking about the resulting ensemble: Either non-uniformly – the graph G_λ is given as advice for security parameter λ – but usually we actually show that a $1 - o(1)$ fraction of the $[n, n^{1+\alpha}, d]$ -hypergraphs are good choices. In that case, we replace G with \mathbf{G} where \mathbf{G}_λ chooses a random $[n(\lambda), n(\lambda)^{1+\alpha}, d(\lambda)]$ -hypergraph G , which is *bad* with vanishing probability $o(1)$. This is of course not good enough, yet the problem can often be by-passed in an application-dependent way, by considering the fact that the *end scheme* using $\mathbf{GF}^{G,P,\ell}$ will also be insecure with probability $o(1)$. One can then consider $\omega(1)$ -instances of this scheme, each using an independent instance from $\mathbf{GF}^{G,P,\ell}$, and then combine them with a combiner, if it exists.

Our constructions below require $(d, O(\log(\lambda)))$ -BLUG($n, \alpha, \text{poly}(\lambda), 2^{-\omega(\log \lambda)}$) to be true for some $n(\lambda) = \text{poly}(\lambda)$ and $\alpha > 0$. For stronger results, we are going to replace $2^{-\omega(\log \lambda)}$ with $2^{-\lambda^\epsilon}$ for some $\epsilon > 0$. Below, in Sections 3.5 and 3.6, we will discuss whether this assumption can be implied by (qualitatively) weaker properties. We will show in particular that $(d, O(\log^{1-\epsilon}(\lambda)))$ -BLUG($n, \alpha, 2^{\omega(\log \lambda)}, 1/\lambda^{\Omega(1)}$) implies $(d, O(\log(\lambda)))$ -BLUG($n, \alpha, \text{poly}(\lambda), 2^{-\omega(\log \lambda)}$).

Here, we briefly discuss what can be expected to start with.

The case $d \geq 3$. For the case $d \geq 3$, a good candidate to study is the case where $\ell = O(\log(\lambda))$ and $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$ is such that G_λ is an $[n(\lambda), n(\lambda)^{1+\alpha}, d]$ -hypergraph which is a good $(n^{1-\gamma}, (1-\gamma)d)$ -expander where $\gamma < \frac{1}{2}$, which exists (for some suitable $\alpha > 0$) by Corollary 1. The corresponding \overline{G}_λ are then in turn also $(n^{1-\gamma}, (1-\gamma)d\ell)$ -expanders by Lemma 1.

Applebaum and Raykov [AR16] recently justify the assumption that for suitable predicates, P , the function family $\mathbf{GF}^{\overline{G},P}$ is one way and a PRG against adversary running in time $2^{n^{1-\gamma}}$, which cannot succeed with probability larger than $2^{-n^{1-\gamma}}$. In the same paper, they also give a decision-to-search reduction for such functions, which however applies only for degrees where we can

accommodate some γ with $3\gamma < 1$. In particular, such functions withstand existing attacks, such as myopic inversion attacks [CEMT09]. Also, the degree of P can be high, e.g., $O(\log(\lambda))$, and this prevents a number of attacks exploiting weakness of the predicate [CM01, BQ12].

Also, as we show in the next section, it is possible to adopt the techniques from [ABR16] to show that we can get good ϵ -biased generators (for a sub-exponential ϵ) with block-wise locality $(3, 2)$. This has been the main technique in validating PRG assumptions on graph-based local functions [MST03, ABR16, OW14].

The special case $d = 2$. The case $d = 2$ is particularly important, as it does allow instantiations from bilinear maps in our applications. Note that algebraic attacks are mitigated here – in contrast to the case of plain locality, i.e., $\ell = 1$, we can set $\ell = O(\log \lambda)$ and achieve sufficiently high algebraic degree of the predicate P . Unfortunately, this is not sufficient to prove pseudorandomness, as shown by recent attacks [BBKK17, LV17], which we have discussed above in Section 1.3.

3.4 Block-Wise local Small-Bias Generators

Several works [CM01, MST03, AL16, ABR16] have focused on studying weaker properties achieved by local generators. In particular, a standard statement towards validating their security is that of showing that they meet the definition of being a *small-bias generator*.

Definition 24. We say $SB : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is an ϵ -small biased generator if $\max_{J \subseteq [n], J \neq \emptyset} |\Pr[x \stackrel{\$}{\leftarrow} \{0, 1\}^n : \bigoplus_{j \in J} SB_j(x) = 1] - \frac{1}{2}| \leq \epsilon$, where $SB_j(x)$ denotes the j -th bit of $SB(x)$.

We show that $\text{GF}_{G,Q,2}$ is a good small-biased generator for a sub-exponential ϵ , where G is an $[n, m, 3]$ -hypergraph, and Q is the predicate which given three 2-bit blocks $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ where $\mathbf{x}_i = (x_i^l, x_i^h)$, outputs

$$Q(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) = x_1^l \oplus x_2^l \oplus x_3^l \oplus (x_1^h \wedge x_2^h) .$$

Another convenient way to think about $\text{GF}_{G,Q,2}$ is as

$$\text{GF}_{G,Q,2}((x_1^l, x_1^h), \dots, (x_n^l, x_n^h)) = \text{GF}_{G,Q^l}(x_1^l, \dots, x_n^l) \oplus \text{GF}_{G,Q^h}(x_1^h, \dots, x_n^h) ,$$

where $Q^l(x_1, x_2, x_3) = x_1 \oplus x_2 \oplus x_3$ and $Q^h(x_1, x_2, x_3) = x_1 \wedge x_2$. To show that $\text{GF}_{G,Q,2}$ has small bias, the main idea is fairly straightforward. Indeed, current analyses of local small-biased generators give two separate analyses for so called “light tests” and “heavy tests”, where the “weight” of a test amounts to the cardinality of $|J|$. For standard locality, withstanding both at the same time forces the graph degree to be at least five, since the predicate needs to be “non-degenerate” for the construction to withstand tests (and the theorem of [ABR16] to apply), and all predicates up to $d = 4$ are degenerate (cf. e.g. [CM01]). This will not be a problem here, as we only target block-wise locality, and thus effectively the predicate can be non-degenerate. The proof will in fact show that for most graphs G , GF_{G,Q^l} resists light tests, whereas GF_{G,Q^h} resists heavy tests, and thus their xor resists all tests for most graphs. The proof easily extends to any Q^l and Q^h which resist light and heavy tests, respectively.

Lemma 3. For all $\delta > 0$ and $\alpha < \frac{1-\delta}{4}$, for a fraction of $1 - o(1)$ of all $[n, n^{1+\alpha}, 3]$ -hypergraphs G , and Q as defined above, $\text{GF}_{G,Q,2}$ is an $\left(e^{-\frac{n^\delta}{4}}\right)$ -biased generator.

Proof. Our proof relies on the machinery introduced by [ABR16]. In particular, we will distinguish between light and heavy tests, depending on whether the set $|J|$ indexing the bits to be xored is large or not. Then, for a random x , denote $y = \text{GF}_{G,Q,2}(x) = y^h \oplus y^l$, where $y^h = \text{GF}_{G,Q^h}(x)$ and $y^l = \text{GF}_{G,Q^l}(x)$. We have

$$\bigoplus_{j \in J} y_j = \left(\bigoplus_{j \in J} y_j^h \right) \oplus \left(\bigoplus_{j \in J} y_j^l \right),$$

and since $\bigoplus_{j \in J} y_j^h$ is independent from $\bigoplus_{j \in J} y_j^l$, with $b^v(J) = \left| \Pr \left[x \stackrel{\$}{\leftarrow} \{0,1\}^n : \bigoplus_{j \in J} y_j^v = 1 \right] - \frac{1}{2} \right|$ for $v \in \{l, h\}$, we have

$$b(J) := \left| \Pr \left[x \stackrel{\$}{\leftarrow} \{0,1\}^n : \bigoplus_{j \in J} y_j = 1 \right] - \frac{1}{2} \right| \leq 2b^l(J) \cdot b^h(J) \leq 2 \min\{b^l(J), b^h(J)\}. \quad (1)$$

We are going to show that for a suitable G , $b^l(J)$ is small for light tests, i.e., small J , whereas $b^h(J)$ will be shown to be small for heavy tests. Let us start discussing the former case.

Let $\Delta = n^\alpha \leq \sqrt{n}/\log(n)$. First, assume we pick $G = (S_1, \dots, S_m)$ as a random $[n, \Delta n, 3]$ hypergraph. We say that G is k -linear if for all sets $J \subseteq [m]$ with $|J| \leq k$, the incidence vectors $\{v_j\}_{j \in J}$ of the sets $\{S_j\}_{j \in J}$ (which are in particular vectors in \mathbb{F}_2^n) are linearly independent. Then, [ABR16] show that G is k -linear with probability $1 - o(1)$, for $k = k(n, \Delta) = \Omega(n/\Delta^2)$. Then, [ABR16] also show in particular that in this case, because the predicate $Q^l(x_1, x_2, x_3) = x_1 \oplus x_2 \oplus x_3$ is 2-resilient (i.e., its Fourier coefficients with weight ≤ 2 are all 0) the bits y_1^l, \dots, y_m^l are k -wise independent. In particular, this means that for every $J \subseteq [m]$ with $|J| \leq k$, we have $b^l(J) = 0$, and therefore $b(J) = 0$ by (1).

Now, we bound $b^h(J)$ for a sufficiently large J . In particular, again assume that G is a random $[n, n\Delta, 3]$ -graph. Then, because Q^h has algebraic degree 2, for every $\delta > 0$, [ABR16] show that with probability $1 - o(1)$ over the choice of G , for every J such that $|J| \geq \Delta^2 n^\delta$, we have

$$b^h(J) \leq \frac{1}{2} e^{-n^{\delta/2}/4},$$

and thus also $b(J) \leq e^{-n^{\delta/2}/4}$ for such J 's.

Now, note that with probability $1 - o(1)$, we also must have a graph G for which both light and heavy tests have small biases. To conclude, we only need to verify that all cardinalities of J are covered, which is certainly true if $\Delta^2 n^\delta \leq k = \Omega(n/\Delta^2)$. This equivalently means that $n^{4\alpha} \leq \Omega(n^{1-\delta})$, which holds whenever $\alpha < \frac{1-\delta}{4}$, as assumed in the theorem statement. \square

3.5 Hardness Amplification via the XOR Construction

In this paper, we rely on the assumption that $\text{GF}^{G,P,\ell}$ is a good PRG for an appropriate family G of expanders. However, we want to add additional justification to our assumptions. Here, in particular, we discuss how weak unpredictability for graph-based block-wise local functions can be amplified to super-polynomially small unpredictability generically. This means in particular that block-wise local PRGs have strong self-amplifying properties, and that for any G and P , in order to invalidate our assumption, we need to find an attack which succeeds in predicting the next bit with large (i.e., polynomial) advantage over $\frac{1}{2}$. For otherwise, the lack of such an attack would imply that for the same G and (a related) P' and ℓ' , $\text{GF}^{G,P',\ell'}$ is a strong PRG.

To this end, we use a simple construction xoring the outputs of generators, which has already been studied to amplify PRG security [DIJK09, MT10]. Our analysis resembles the one from [DIJK09], but is given for completeness. Also, a more general construction, with xoring replaced by a general extractor, was considered by Applebaum [App12]. The use of xor, however, is instrumental to preserve block-wise locality. The main drawback of this construction is that it can *at best* ensure $2^{-\Omega(\log^{1+\theta} \lambda)}$ distinguishing gap for some $\theta \in (0, 1]$ while retaining block size $\ell = O(\log \lambda)$. In Section 3.6, we explain a different approach which relies on a different assumption. and potentially guarantees $2^{-\lambda^{\Omega(1)}}$ distinguishing gap.

The XOR construction Let $\mathbf{UG} = \{\mathbf{UG}_\lambda\}_{\lambda \in \mathbb{N}}$ be an (s, δ) -UG, where \mathbf{UG}_λ is a distribution on functions $\{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}$. For an additional parameter $k = k(\lambda) \geq 1$, we define the ensemble $\mathbf{UG}^k = \{\mathbf{UG}_\lambda^k\}_{\lambda \in \mathbb{N}}$, where \mathbf{UG}_λ^k samples functions $\mathbf{UG}_1, \dots, \mathbf{UG}_k \stackrel{\$}{\leftarrow} \mathbf{UG}_\lambda$ and output the description of a function $\mathbf{UG}^k : \{0, 1\}^{n \times k} \rightarrow \{0, 1\}^m$ which, on input $x = x^1 \parallel \dots \parallel x^k$, where $x^i \in \{0, 1\}^{n(\lambda)}$, outputs

$$\mathbf{UG}^k(x) = \mathbf{UG}_1(x^1) \oplus \dots \oplus \mathbf{UG}_k(x^k).$$

We show the following theorem, whose proof relies on Yao's XOR Lemma [Yao82, GNW11].

Theorem 5 (Security of the XOR Construction). *If \mathbf{UG} is a (s, δ) -UG and $k = k(\lambda)$ is polynomial in λ , then \mathbf{UG}^k is a (s', ϵ) -PRG, where*

$$\epsilon(\lambda) \leq (2\delta(\lambda))^{k(\lambda)}, \quad s'(\lambda) = \Theta\left(\frac{\delta(\lambda)^{2k} \cdot s(r)}{k \log(k/\delta(\lambda))}\right).$$

Proof. Let $A = \{A_\lambda\}_\lambda$ be a predictor family size s' such that for some $i = i(\lambda)$ guesses with probability

$$\pi_{A,i}(\lambda) = \Pr \left[\begin{array}{l} x \stackrel{\$}{\leftarrow} \{0, 1\}^{n \times k} \\ \mathbf{UG}^k \stackrel{\$}{\leftarrow} \mathbf{UG}_\lambda^k \end{array} : A_\lambda(\mathbf{UG}^k, \mathbf{UG}_{\leq i}^k(x)) = \mathbf{UG}_i^k(x) \right].$$

Now, we can build another adversary family $B = \{B_\lambda\}_{\lambda \in \mathbb{N}}$ which takes as input $\mathbf{UG}_1, \dots, \mathbf{UG}_k$ from the range of \mathbf{UG}_λ , as well as $y_1 = \mathbf{UG}_{1, \leq i}(x^1), \dots, y_k = \mathbf{UG}_{k, \leq i}(x^k)$ for $x_1, \dots, x_k \stackrel{\$}{\leftarrow} \{0, 1\}^n$, and outputs a bit

$$B_\lambda(\mathbf{UG}_1, \dots, \mathbf{UG}_k, y_1, \dots, y_k) = A_\lambda(\mathbf{UG}^k = (\mathbf{UG}_1, \dots, \mathbf{UG}_k), y_1 \oplus \dots \oplus y_k),$$

and thus, by construction,

$$\pi_{A,i}(\lambda) = \Pr \left[\begin{array}{l} \mathbf{UG}_1, \dots, \mathbf{UG}_k \stackrel{\$}{\leftarrow} \mathbf{UG}_\lambda \\ x^1, \dots, x^r \stackrel{\$}{\leftarrow} \{0, 1\}^n \\ y_j = \mathbf{UG}_{j, \leq i}(x^j) \text{ for } j = 1, \dots, k \\ b = \mathbf{UG}_{1, i+1}(x^1) \oplus \dots \oplus \mathbf{UG}_{k, i+1}(x^k) \end{array} : B_\lambda(\mathbf{UG}_1, \dots, \mathbf{UG}_k, y_1, \dots, y_k) = b \right].$$

To continue, we rely on a (concrete) version of Yao's XOR Lemma, which we state here, with parameters obtained from Levin's proof. See [GNW11] for further details. (Note that g is not sampled from a distribution in the following statement, but this can be simulated by making g part of the randomness x . Also, for the uniform setting, there are no efficiency requirements on g and P .)

Theorem 6 (XOR Lemma). Let $g : \{0, 1\}^r \rightarrow \{0, 1\}^*$ and $P : \{0, 1\}^r \rightarrow \{0, 1\}$ such that for all adversaries A with size $s(r)$ we have

$$\Pr \left[x \stackrel{\$}{\leftarrow} \{0, 1\}^r : A(g(x)) = P(x) \right] \leq \frac{1}{2} + \delta(r) .$$

Then, for all $k = k(r)$, all $\gamma = \gamma(r)$ and all adversaries B with size $s'(r)$,

$$\Pr \left[x_1, \dots, x_r \stackrel{\$}{\leftarrow} \{0, 1\}^r : B(g(x_1), \dots, g(x_r)) = P(x_1) \oplus \dots \oplus P(x_r) \right] \leq \frac{1}{2} + 2^{k-1} \delta^k + \gamma(r) , \quad (2)$$

where $s'(r) = \Theta \left(\frac{\gamma^2 \cdot s(r)}{k \cdot \log(k/\gamma)} \right)$.

Therefore, by the assumption in the theorem statement that UG is (s, δ) -UG, the statement follows by setting $\gamma = \delta^k$ and noting that $2^{k-1} + 1 \leq 2^k$ for all $k \geq 1$. \square

Block-wise local instantiation We instantiate the construction with parameter k when $\text{UG} = \mathbf{GF}^{G, P, \ell}$ for a family of $[n, m, d]$ -hypergraphs $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$, some $\ell = \ell(\lambda)$, and a family P of $(d \times \ell)$ -bit predicates. Since the resulting function UG_λ^k uses k instances of the *same* function $\mathbf{GF}_{G_\lambda, P_\lambda, \ell}$, it can equivalently be thought as having the form (up to re-arranging the order of the input bits) $\mathbf{GF}_{G_\lambda, P_\lambda^k, \ell(\lambda) \cdot k(\lambda)}$, where the predicate P^k on input $d(k \cdot \ell)$ -bit blocks $\mathbf{x}_1, \dots, \mathbf{x}_d$, it interprets each of them as k ℓ -bit blocks $\mathbf{x}_i = \mathbf{x}_{i,1} \parallel \dots \parallel \mathbf{x}_{i,k}$ and outputs

$$P^k(\mathbf{x}_1, \dots, \mathbf{x}_d) = P(\mathbf{x}_{1,1}, \dots, \mathbf{x}_{d,1}) \oplus \dots \oplus P(\mathbf{x}_{k,1}, \dots, \mathbf{x}_{k,d}) .$$

To instantiate our transformation, we assume that for some $\ell(\lambda) = \Omega(\log^{1-\theta}(\lambda))$ and a family of $[n(\lambda), m(\lambda), d]$ -hypergraphs $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$, the function family $\text{UG} = \mathbf{GF}^{G, P, \ell}$ is a $(s(\lambda) = 2^{\log^3(\lambda)}, \delta(\lambda) = \lambda^{-\Omega(1)})$ -UG. Now, set $k(\lambda) = \log^\theta(\lambda)$. Then, UG^k is by the above $(d, O(\log(\lambda)))$ -block-wise local, and it is also (s', ϵ) -UG for $s'(\lambda) = \text{poly}(\lambda)$, and

$$\epsilon(\lambda) = (2\delta(\lambda))^{k(\lambda)} = 2^{-\Omega(\log^{1+\theta}(\lambda))} .$$

In other words, we have just established the following corollary.

Corollary 2. For any $\beta > 0$, $d \geq 2$, and $\theta \in (0, 1]$, if the $(d, O(\log^{1-\theta}(\lambda)))$ -BLUG $(n, \beta, 2^{\log^3(\lambda)}, 1/\lambda^{\Omega(1)})$ assumption holds, then the assumption $(d, O(\log(\lambda)))$ -BLUG $(n, \beta, \text{poly}(\lambda), 2^{-\Omega(\log^{1+\theta}(\lambda))})$ also holds true.

3.6 The Extraction Construction

The XOR construction guarantees that finding a graph and predicate for which $\mathbf{GF}^{G, P, \ell}$ is even only mildly unpredictable (for slightly super-polynomial predictors) implies already a block-wise local with block size $O(\log(\lambda))$ and inverse super-polynomial distinguishing gap. However, note that sub-exponential distinguishing advantage while being $(d, O(\log(\lambda)))$ -block-wise local is out of reach, as this require $k = O(\log(\lambda))$, which in turns can only gives us (even assuming an “ideal” version of the XOR Lemma) distinguishing advantage $\delta^{O(\log(\lambda))} = 2^{-O(\log^2 \lambda)}$, since $\delta = 1/\lambda^{O(1)}$.

Here, we give a second construction of a block-wise local PRG with polynomial stretch $m = n^{1+\alpha}$ that uses an instantiation of $\mathbf{GF}^{G, P, \ell}$ which merely ensures its output has a (sufficient) amount of pseudo-min-entropy. We stress that we have no reason to focus on such an assumption other than the fact that this may appear easier to reach for a given graph and predicate.

The extraction construction The following assumption is a weakening of the notion of a PRG, to a pseudo-min-entropy generator (PMEG). Recall that a random variable X has *min-entropy* k if every value is taken with probability at most 2^{-k} .

Definition 25 (Family of Pseudo-Min-Entropy Generators (PMEGs)). *Let n and $\Delta \leq m$ be polynomials, and $\mu : \mathbb{N} \rightarrow [0, 1]$ be a function. A family of $(n(\lambda), m(\lambda), \mu(\lambda), \Delta(\lambda))$ -PMEGs is an ensemble of distributions $\mathbf{PMEG} = \{\mathbf{PMEG}_\lambda\}$ satisfying the same syntax and efficiency requirements as a PRG, but moreover:*

μ -Indistinguishability: *There exists a family of distributions $\{X_\lambda\}_{\lambda \in \mathbb{N}}$ such that X_λ is over the $m(\lambda)$ -bit strings,*

$$\begin{aligned} & \left\{ \mathbf{PMEG} \stackrel{\$}{\leftarrow} \mathbf{PMEG}_\lambda; s \stackrel{\$}{\leftarrow} \{0, 1\}^{n(\lambda)} : (\mathbf{PMEG}, \mathbf{PMEG}(s)) \right\}_{\lambda \in \mathbb{N}} \\ & \approx_\mu \left\{ \mathbf{PMEG} \stackrel{\$}{\leftarrow} \mathbf{PMEG}_\lambda; r \stackrel{\$}{\leftarrow} X_\lambda : (\mathbf{PRG}, r) \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

and moreover, X_λ has min-entropy $m(\lambda) - \Delta(\lambda)$.

We use the following lemma by Dodis and Smith [DS05], which shows that xoring a sufficiently high-min-entropy source with the output of a small-bias generator yields uniform randomness.

Lemma 4. *Let $\text{SB} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be an ϵ -biased generator and let X be a random variable taking values in $\{0, 1\}^m$ with min-entropy at least $m - \Delta$ for some $\Delta \geq 0$. Then,*

$$\text{SD}((\text{SB}(U_n) \oplus X), U_m) \leq \frac{\epsilon}{\sqrt{2}} \cdot 2^{\Delta/2},$$

where U_k denotes the uniform distribution on k -bit strings, and SD denotes statistical distance.

If we use the instantiation of an ϵ -biased generator guaranteed to exist by Lemma 3 as $\text{GF}_{G, Q, 2}$, we see that for $m = n^{1+\alpha}$, with a small loss allowing notational simplifications, we can have $\delta = 2^{-n^{1-5\alpha}/4}$. In particular, the statistical distance in Lemma 4 remains sub-exponential as long as the entropy loss is $h = o(n^{1-5\alpha})$. In the following, for n and α , we denote by $\mathcal{G}_{n, \alpha, d}$ the set of $[n, n^{1+\alpha}, d]$ -hypergraphs G for which $\text{GF}_{G, Q, 2}$ is a $2^{-n^{1-5\alpha}/4}$ -biased generator. (Lemma 3 needs $d = 3$, but we can trivially extend the result to $d \geq 4$ by ignoring any extra input.) Lemma 3 implies that $\mathcal{G}_{n, \alpha, d}$ includes a $1 - o(1)$ fraction of $[n, n^{1+\alpha}, d]$ -hypergraphs.

Now, let $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of $[n(\lambda), n(\lambda)^{1+\alpha}, d]$ -hypergraphs, with $d \geq 3$, and let $\ell(\lambda) = O(\log(\lambda))$. Moreover, let $P = \{P_\lambda\}_{\lambda \in \mathbb{N}}$ be a family of predicates on $(d \times \ell)$ -bit strings. Also, for the predicate Q defines as above, we consider the family of predicates $P \oplus Q = \{R_\lambda\}_{\lambda \in \mathbb{N}}$, where the corresponding predicate R_λ for security parameter $\lambda \in \mathbb{N}$ outputs, on input $d(\ell(\lambda) + 2)$ -bit strings $\mathbf{x}_1, \dots, \mathbf{x}_d$,

$$P_\lambda(\mathbf{x}_1[1 \dots \ell], \dots, \mathbf{x}_d[1 \dots \ell]) \oplus Q(\mathbf{x}_1[\ell + 1, \ell + 2], \mathbf{x}_2[\ell + 1, \ell + 2], \mathbf{x}_3[\ell + 1, \ell + 2]).$$

Then, note that evaluating $\text{GF}_{G_\lambda, R_\lambda, \ell(\lambda)+2}$ on a random input, is the same as evaluating $\text{GF}_{G_\lambda, P_\lambda, \ell(\lambda)}$ and $\text{GF}_{G_\lambda, Q, 2}$ on independent random inputs, and xoring the results. This means in particular that by Lemma 4, if the graph family G yields a small-bias generator, and $\mathbf{GF}^{G, P, \ell}$'s output has enough computational min-entropy, then we obtain a good PRG. This is summarized by the following lemma.

Lemma 5. Let $d \geq 3$, let $n(\lambda)$ be a polynomial, $\ell(\lambda) = O(\log(\lambda))$, and let $\alpha > 0$. If there exists a family $G = \{G_\lambda\}_{\lambda \in \mathbb{N}}$ of $[n(\lambda), n(\lambda)^{1+\alpha}, d]$ -hypergraphs, and a family $P = \{P_\lambda\}_{\lambda \in \mathbb{N}}$ of predicates on $(d \cdot \ell)$ -bit strings such that:

1. $G_\lambda \in \mathcal{G}_{n(\lambda), \alpha, d}$, and
2. $\mathbf{GF}^{G, P, \ell}$ is a $(n(\lambda), n(\lambda)^{1+\alpha}, \nu(\lambda), \Delta(\lambda))$ -PMEG for $\Delta(\lambda) = o(n(\lambda)^{1-5\alpha})$.

Then, $\mathbf{GF}^{G, P \oplus Q, \ell+2}$ is a PRG with sub-exponential distinguishing gap.

4 IO from Block-Wise Locality- $(L, \log \lambda)$ PRG and L -Linear Maps

In this section, we prove the following bootstrapping theorem.

Theorem 7 (Bootstrapping via block-wise local PRGs). Let $\mathcal{R} = \{\mathcal{R}_\lambda\}$ be any family of rings, ε be any positive constant, L any positive integer, n any sufficiently large polynomial, and $i\ell$ and κ any polynomials. There is a construction of $i\ell(\lambda)$ -bit-input IO for P/poly, from the following primitives:

- A family of $(n(\lambda) \times \log \lambda, n(\lambda)^{1+\varepsilon})$ -PRGs with block-wise locality $(L, \log \lambda)$.
- A public-key FE for degree- L polynomials in \mathcal{R} , with linear efficiency and Full-Sel-security; or with a secret-key FE with the same properties, assuming additionally LWE with subexponential modulo-to-noise ratio.

The IO scheme is $(2^{-\kappa(\lambda)} \text{negl}(\lambda))$ -secure, if the PRG and FE schemes are $(2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda))$ -secure, and LWE is $(2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda))$ -hard.

Theorem 7 follows the same approach as Lin’s recent bootstrapping theorem [Lin17], but modifies it in two ways. First, it uses block-wise local PRGs to replace local PRGs. Second, it makes explicit the relation between the security level (more precisely, the maximal distinguishing gap) of the underlying PRG and FE, and the input-length and security level of the resulting IO — if the underlying primitives are $2^{-i\ell+\kappa}$ negl-secure, then the resulting IO scheme is for $i\ell$ -bit-input circuits and 2^κ negl-security. Such relations are implicit in previous works, and not as tight as shown here.

Overview of Proof of Theorem 7 To show the theorem, similar to previous works [LV16, Lin17], we take two steps:

Step 1 Construct a single-key public-key (or secret-key) FE schemes $\mathbf{CFE} = \{\mathbf{CFE}^{N, D, S}\}$ for P/poly, with $(1 - \varepsilon)$ -sublinear compactness and $2^{-i\ell+\kappa}$ negl-Full-Sel-security, starting from a public-key (or secret-key) FE for degree- L polynomials in \mathcal{R} , with linear efficiency and Full-Sel-security.

Previously, the work of [LV16] showed how to achieve this transformation from a locality- L PRGs and FE for computing degree $3L + 2$ polynomials. Following that, the two recent works of [Lin17, AS16] used a *pre-processing technique* to relax the requirement on the underlying FE to supporting only degree- L polynomials. In this work, we extend their *pre-processing technique* even further, in order to relax the requirement on the underlying PRGs from having locality L to having block-wise locality $(L, \log \lambda)$.

In particular, our approach is a fairly straightforward extension of [Lin17], albeit notationally tedious, once we make the following observation. For any block-wise locality L function PRG(\mathbf{x}) with input \mathbf{x} of dimension $\ell \times n$, one can compute the PRG as a degree L function in an arbitrary ring, provided all the monomials defined on the variables of each ℓ -bit column are pre-computed.

More concretely, let PRG_i denote the function that computes the i^{th} output bit of PRG. We know that PRG_i depends on L input columns, $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_L}$, each of size ℓ bits. We can arithmetize PRG_i and write it as a sum of multilinear monomials over the bits in these input columns (recall that all computations are taking place in a corresponding *ring*, so all expressions are within this ring) – in particular, with \mathcal{M}_i being the set of monomials appearing in the computation of PRG_i ,

$$\text{PRG}_i(\mathbf{x}) = \sum_{M \in \mathcal{M}_i} M(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_L}).$$

Furthermore, every monomial $M(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_L})$ can be written as the product of L monomials M_1, \dots, M_L , defined over each of the columns,

$$\text{PRG}_i(\mathbf{x}) = \sum_{M \in \mathcal{M}_i} M_1(\mathbf{x}_{i_1}) \times \dots \times M_L(\mathbf{x}_{i_L})$$

Now suppose that we have pre-computed *all possible* monomials over every column. In particular, let $\text{Mnml}(\mathbf{x}_i)$ denote the set of all 2^ℓ multilinear monomials over bits in \mathbf{x}_i . Then $\text{PRG}_i(\mathbf{x})$ can be computed in degree L , because there exists a degree- L function PRG'_i such that

$$\text{PRG}'_i(\text{Mnml}(\mathbf{x}_1), \dots, \text{Mnml}(\mathbf{x}_n)) = \text{PRG}_i(\mathbf{x}).$$

Given this, we can have a way of computing PRG'_i using the underlying FE for degree L polynomials. Also, because we have $\ell(\lambda) = O(\log \lambda)$, the domain size of PRG has now increased only by a $\text{poly}(\lambda)$ multiplicative factor when transforming it into PRG' . While this is the main idea, some care must be taken to ensure that this trick fits together with the rest of the preprocessing in [Lin17]. We describe this step in full detail in Section 4.1.

In the case that the obtained FE scheme **CFE** is a secret-key one, we invoke the result of [BNPW16] to transform it into a public key FE scheme with the same properties, assuming LWE with subexponential modulus-to-noise ratio.

Since our transformation from FE for low-degree computations to weakly-compact FE for P/poly in Section 4.1 incurs only a polynomial security loss, and so does the transformation of [BNPW16], the resulting weakly-compact FE has essentially the same level of security as that of underlying primitives.

Step 2. Apply an FE-to-IO transformation to obtain $i\ell$ -bit-input IO for P/poly, with $2^{-\kappa}$ negl-security.

The literature already offers three FE-to-IO transformations [BV15, AJ15, LPST16] that start from a public key FE scheme $\text{CFE} = \{\text{CFE}^{N,D,S}\}$ as described above w.r.t. any positive constant ε . In this work, we reduce the security loss incurred in the transformation so as to start with $2^{-i\ell+\kappa}$ negl-secure FE (as opposed to $2^{-O(i\ell^2)+\kappa}$ negl-secure or $2^{-O(\log \lambda)i\ell+\kappa}$ negl-secure FE as in previous works). To do so, we present a new FE-to-IO transformation inspired by that of [LPST16] and present a tight analysis. We describe this step below in Section 4.2.

4.1 Step 1: Constructing Weakly-Compact FE

Proposition 1. *Let \mathcal{R} , ε , L , and n be defined as in Theorem 7, and $\bar{\kappa}$ be any polynomial. There is a construction of 1-key weakly-compact public-key FE for P/poly from the following primitives:*

- A family of $(n(\lambda) \times \log \lambda, n(\lambda)^{1+\varepsilon})$ -PRGs with block-wise locality $(L, \log \lambda)$.
- Public-key FE for degree- L polynomials in \mathcal{R} , with linear efficiency and Full-Sel-security; or secret-key FE with the same properties, assuming additionally LWE with subexponential modulus-to-noise ratio.

The weakly-compact FE is $(2^{-\bar{\kappa}(\lambda)} \text{negl}(\lambda))$ -Full-Sel-secure, if the underlying PRG and FE are $(2^{-\bar{\kappa}(\lambda)} \text{negl}(\lambda))$ -secure and LWE is $(2^{-\bar{\kappa}(\lambda)} \text{negl}(\lambda))$ -hard.

It was shown in [Lin17] that 1-key weakly-compact FE for P/poly can be constructed from locality- L PRG and (unbounded collusion) FE for degree- L polynomials. Their construction of weakly-compact FE follows from the blue-print of previous works [Lin16, LV16], which uses FE for low degree polynomials to compute a randomized encoding of a computation in P/poly, with pseudo-randomness generated through a local PRG. The locality of RE and PRG ensures that their composition can be computed in low degree. However, the straightforward composition of RE and PRG leads to a computation with degree $3L + 2$. The key idea in [Lin17] and the concurrent work of [AS16] is that part of the RE computation can already be done at encryption time, that is, by asking the encryptor to pre-process the inputs (of the computation in P/poly) and seeds of PRG, and encrypt the pre-processed values, the composition of RE and PRG can be computed in just degree L from the pre-processed values, at decryption time — This is called the *preprocessing technique*. We take this technique one step further: By also performing part of the PRG computation at encryption time, we can replace local PRG with block-wise local PRG (with appropriate parameters) at “no cost”.

Below, we first briefly review the blueprint of [LV16], then describe the pre-processing idea of [Lin17] and how to use it to accommodate PRG with block-wise locality.

The General Blueprint of [LV16] To construct 1-key weakly-compact FE for P/poly, Lin and Vaikuntanathan [LV16] (LV) first observed that, using the Trojan Method [CIJ⁺13], it suffices to construct 1-key weakly-compact FE for NC^1 functions with some fixed depth $D(\lambda) = O(\log \lambda)$; denote this class of functions as NC_D^1 .

Next, to bootstrap a low-degree FE scheme to FE for NC_D^1 , the idea is using randomized encoding to “compress” any function $h(\mathbf{x}) \in \text{NC}_D^1$ into a function $g(\mathbf{x}, \mathbf{s}) = \text{REnc}(f, \mathbf{x}; \text{PRG}(\mathbf{s}))$ with small degree in \mathcal{R} . The reason that local PRG is used is that the locality of a Boolean function bounds the degree of computing this function in any ring. Then, plugging-in randomized encodings with small locality like that of [AIK04] the overall degree of g is small. For the security proof to work out, the actual functions used in the LV construction are more complicated and has form

$$g(\mathbf{x}, \mathbf{s}, \mathbf{s}', b) = (1 - b)(\text{REnc}(f, \mathbf{x}; \text{PRG}(\mathbf{s}))) + b(\mathbf{CT} \oplus \text{PRG}(\mathbf{s}')),$$

where \mathbf{CT} is a ciphertext hardwired in the secret key, and serves as “space” to hide values in the secret key in the security proof.

A formal description of the LV public key FE scheme $\mathbf{CFE}^{N,D,S}$ for NC^1 circuits with input-length $N = N(\lambda)$, depth $D = D(\lambda) = O(\log \lambda)$, and size $S = S(\lambda)$ is in Figure 2. (The secret-key case has almost identical construction.) The scheme uses the following tools:

- Full-Sel-secure (collusion resistant) FE schemes for degree- $(3L + 2)$ polynomials in some \mathcal{R} , $\{\mathbf{FE}^{N'} = (\text{FE.Setup}, \text{FE.KeyGen}, \text{FE.Enc}, \text{FE.Dec})\}$, with linear efficiency.
- A $(n, n^{1+\alpha})$ -pseudorandom generator PRG with locality L , for a sufficiently large polynomial input length $n = n(\lambda)$ and any positive constant α .
- The AIK randomized encoding scheme in NC^0 [AIK04]; denote the encoding algorithm as $\text{AIK}(f, \mathbf{x}; \mathbf{r})$.

We refer the reader to [LV16] for the correctness and security of the scheme. The compactness of the scheme \mathbf{CFE} follows from the following two facts:

Single-key Compact FE Scheme CFE by [LV16]

SETUP: $\text{CFE.Setup}(1^\lambda)$ samples $(\text{MPK}, \text{MSK}) \xleftarrow{\$} \text{FE.Setup}(1^\lambda)$.

ENCRYPTION: $\text{CFE.Enc}(\text{MPK}, \mathbf{x})$ samples $\mathbf{s}, \mathbf{s}' \xleftarrow{\$} \{0, 1\}^\Gamma$ for $\Gamma = S^{1/1+\alpha} \text{poly}(\lambda)$, and generates

$$\text{CT} \xleftarrow{\$} \text{FE.Enc}(\text{MPK}, (\mathbf{x}, \mathbf{s}, \mathbf{s}', 0))$$

KEY GENERATION: $\text{CFE.KeyGen}(\text{MSK}, h)$ does the following:

- Sample $\text{CT} \xleftarrow{\$} \{0, 1\}^\ell$, where ℓ is set below.
- Define function g as follows: On input \mathbf{x} of length N , two PRG seeds \mathbf{s}, \mathbf{s}' each of length Γ , and a bit b ,

$g(\mathbf{x}, \mathbf{s}, \mathbf{s}', b)$ does the following:

- For every $i \in [S]$, let $h_i(\mathbf{x})$ denote the function that computes the i^{th} output bit of $h(\mathbf{x})$. Since $h \in \text{NC}_D^1$, h_i has depth $D(\lambda) = O(\log \lambda)$ and size $2^{D(\lambda)} = \text{poly}(\lambda)$.
- If $b = 0$, compute $\mathbf{r} = \text{PRG}(\mathbf{s})$, whose output has length $\Gamma^{1+\alpha} = S \text{poly}(\lambda)$; divide the output into S equally long portions and denote by $\mathbf{r}[i]$ the i^{th} portion. For every $i \in [S]$, compute the AIK encoding $\Pi[i]$ of computation (h_i, \mathbf{x}) as follows:

$$\forall i \in [S], \quad \Pi[i] = \text{AIK}(h_i, \mathbf{x}; \mathbf{r}[i]).$$

Output $\Pi = \{\Pi[i]\}_i$; set $\ell = |\Pi|$.

- If $b = 1$, output $\Pi = \text{CT} \oplus \text{PRG}(\mathbf{s}')$.
- For every $l \in [\ell]$, generate a secret key $\text{SK}_l \xleftarrow{\$} \text{FE.KeyGen}(\text{MSK}, g_l)$ for g_l that computes the l^{th} output bit of g .

Output $\text{SK} = \{\text{SK}_l\}_{l \in [\ell]}$.

DECRYPTION: $\text{CFE.Dec}(\text{SK}, \text{CT})$ computes $\Pi = \{\text{FE.Dec}(\text{SK}_l, \text{CT})\}_{l \in [\ell]}$, parses $\Pi = \{\Pi[i]\}$, and decodes every $\Pi[i]$ using the AIK decoding algorithm to obtain the output $h(\mathbf{x})$.

Figure 2: Single-key Compact FE CFE by [LV16]

1. The length of the input $(\mathbf{x}, \mathbf{s}, \mathbf{s}', 0)$ encrypted using **FE** is $N+2\Gamma+1 = N+S(\lambda)^{1/(1+\alpha)} \text{poly}(\lambda)$.
2. **FE** has linear efficiency.

Putting them together, we have,

$$\begin{aligned} \text{Time}_{\text{CFE.Enc}}(\text{MPK}, \mathbf{x}) &= \text{Time}_{\text{FE.Enc}}(\text{MPK}, (\mathbf{x}, \mathbf{s}, \mathbf{s}', 0)) \\ &= \text{poly}(\lambda)|(\mathbf{x}, \mathbf{s}, \mathbf{s}', 0)| = S(\lambda)^{1/(1+\alpha)} \text{poly}(\lambda, N) \end{aligned}$$

which is sublinear in the function size as desired. Furthermore, to see why degree- $(3L+2)$ FE suffices for the construction, note that the construction uses the underlying FE to generate keys computing the function g in Figure 2, and hence it suffices to argue that g can be computed in degree $3L+2$. By definition of g , when $b=1$, the output can be computed in degree L as the PRG can be computed in degree L in \mathcal{R} (XOR with **CT** does not incur additional degree as **CT** are constants hardwired in the function g); when $b=0$, the output can be computed in degree $3L+1$, since the AIK randomized encoding has degree 3 in the random bits (*i.e.* PRG output) and 1 in the input \mathbf{x} . Therefore, g has exactly degree $3L+2$, as selection by b can be done with one multiplication.

The Idea of Preprocessing in [Lin17] Towards reducing the degree of the underlying FE and accommodating PRGs with block-wise locality- $(L, \log \lambda)$, the idea is letting the encryptor preprocess the input $(\mathbf{x}, \mathbf{s}, \mathbf{s}', b)$ to produce certain intermediate values, from which the output of function g can be computed in exactly degree L . To see this, the output of g is viewed as corresponding to S AIK randomized encodings for functions $\{h_i\}_{i \in [S]}$. If the l^{th} output bit belongs to the i^{th} randomized encoding for h_i with random tape $\mathbf{r}[i]$, the function g_l computing it can be written as a sum of monomials as follows:

$$\begin{aligned} g_l(\mathbf{x}, \mathbf{s}, \mathbf{s}', b) &= (1-b)g_{l0}(\mathbf{x}, \mathbf{s}) + bg_{l1}(\mathbf{s}') \\ &= (1-b) \sum_{i_0, i_1, i_2, i_3} c_{i_0, i_1, i_2, i_3} \mathbf{x}_{i_0} \mathbf{r}[i_1] \mathbf{r}[i_2] \mathbf{r}[i_3] + b \sum_j c_j \mathbf{r}'_j \end{aligned} \quad (3)$$

where $\mathbf{r}[i]$ is the i^{th} portion in $\mathbf{r} = \text{PRG}(\mathbf{s})$, and $\mathbf{r}' = \text{PRG}(\mathbf{s}')$. This is because in the case of $b=0$, the output is a bit in the AIK encoding of h_i and hence has degree 1 in the input \mathbf{x} and degree 3 in $\mathbf{r}[i]$, while in the case of $b=1$, the output has degree 1 in \mathbf{r}' .

When PRG has locality L , the straightforward way of computing a degree-3 monomial $\mathbf{r}[i_1] \mathbf{r}[i_2] \mathbf{r}[i_3]$ from the seed \mathbf{s} requires degree $3L$. The works of [Lin17, AS16] showed how to reduce the degree to just L . First, they use a different way to compute each $\mathbf{r}[i]$. View the seed \mathbf{s} as a $Q \times \Gamma'$ matrix with $Q = Q(\lambda) = \text{poly}(\lambda)$ rows and $\Gamma' = S^{1/1+\alpha}$ columns; apply PRG on each row of \mathbf{s} to expand the seed matrix into a $Q \times S$ matrix \mathbf{r} of pseudo-random bits. That is, denote the q^{th} row of \mathbf{s} and \mathbf{r} as \mathbf{s}_q and \mathbf{r}_q ; $\mathbf{r}_q = \text{PRG}(\mathbf{s}_q)$. Finally, set the random tape for computing the i^{th} AIK encoding to be the i^{th} column $\mathbf{r}[i]$ of \mathbf{r} .

In [Lin17], they used PRGs with locality L . Let $\text{PRG}[i]$ denote the function computing the i^{th} output bit of PRG, and let $\text{Nbr}(i) = \{\gamma_1, \dots, \gamma_L\}$ be the indexes of the L seed bits that the i^{th} output bit depends on. Therefore,

$$\begin{aligned} \mathbf{r}[i_1] \mathbf{r}[i_2] \mathbf{r}[i_3] &= \text{PRG}[i_1](\mathbf{s}_{i_1}) \text{PRG}[i_2](\mathbf{s}_{i_2}) \text{PRG}[i_3](\mathbf{s}_{i_3}) \\ &= \sum_{\substack{\text{Monomials} \\ X, Y, Z \text{ in PRG}[i]}} \begin{pmatrix} X(s_{i_1, \gamma_1}, \dots, s_{i_1, \gamma_L}) \\ \times Y(s_{i_2, \gamma_1}, \dots, s_{i_2, \gamma_L}) \\ \times Z(s_{i_3, \gamma_1}, \dots, s_{i_3, \gamma_L}) \end{pmatrix}. \end{aligned} \quad (4)$$

Suppose that one has pre-computed all degree ≤ 3 monomials over bits in each column $s[\gamma]$ of s .

$$\text{Define } \text{Mnml}^{\leq 3}(A) := \{a_i a_j a_k \mid a_i, a_j, a_k \in A \cup \{1\}\}$$

Given $\text{Mnml}^{\leq 3}(s[\gamma])$ for every $\gamma \in \text{Nbr}(i)$, one can compute $\mathbf{r}[i]_{i_1} \mathbf{r}[i]_{i_2} \mathbf{r}[i]_{i_3}$ in Equation (4) using just degree L . Similarly, given $\text{Mnml}^{\leq 3}(s[\gamma])$ for all $\gamma \in [\Gamma']$, one can compute *any* degree 3 monomials over bits in $\mathbf{r}[i]$ for *any* i , sufficient for the computation of g .

Furthermore, the size of each set $\text{Mnml}^{\leq 3}(s[\gamma])$ is bounded by $(Q+1)^3 = \text{poly}(\lambda)$, and thus the size of their union for all γ is bounded by $\Gamma' \text{poly}(\lambda) = S^{1/1+\alpha} \text{poly}(\lambda)$ — only a polynomial factor (in λ) larger than the original seed s itself. Therefore the encryptor can afford to precompute all these monomials and encrypt them, without compromising the weak-compactness of the resulting FE for NC_D^1 scheme.

This Work: Handling Block-Wise Local PRG. Our new observation is that the above technique naturally extends to accommodate block-wise local PRGs. Consider a family of $(n(\lambda) \times \log \lambda, n(\lambda)^{1+\alpha})$ -PRGs with block-wise locality- $(L, \log \lambda)$. As before, we think of the seed of such PRGs as a vector \mathbf{t} of length n , where every element t_i is a block of $\log \lambda$ bits, and each output bit $\text{PRG}[i](\mathbf{t})$ depends on at most L blocks.

Correspondingly, think of the seed matrix s described above as consisting of $Q \times \Gamma'$ blocks of $\log \lambda$ bits. When $\mathbf{r}[i]$ is computed using block-wise local PRGs, the degree-3 monomial $\mathbf{r}[i]_{i_1} \mathbf{r}[i]_{i_2} \mathbf{r}[i]_{i_3}$ in Equation (4) now depends on a set of blocks $\{s_{i_t, \gamma_s}\}_{t \in [3], s \in [L]}$. Though the actual locality of the PRG is $L \log \lambda$, due to its special structure, we can still pre-process the seed s to enable computing any degree-3 monomial over $\mathbf{r}[i]$ for any i using degree L , in the following two steps.

1. Precompute all *multilinear* monomials over bits in each block $s_{q, \gamma}$ in s .

$$\text{Define } \text{Mnml}(A) := \{a_{i_1} a_{i_2} \cdots a_{i_q} \mid q \leq |A| \text{ and } \forall j, k \ a_{i_j} \neq a_{i_k} \in A\} .$$

More precisely, precompute $\text{Mnml}(s_{q, \gamma})$ for all $q \in [Q]$ and $\gamma \in [\Gamma']$. Note that each set $\text{Mnml}(s_{q, \gamma})$ has exactly size λ .

2. For every column $\gamma \in [\Gamma']$, take the union of monomials over blocks in column γ , that is, $\cup_q \text{Mnml}(s_{q, \gamma})$. Then, precompute all degree- ≤ 3 monomials over this union, that is, $\text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q, \gamma}))$, for each γ . Observe that from $\{\text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q, \gamma}))\}_{\gamma \in [\Gamma']}$, one can again compute *any* degree-3 monomial in $\mathbf{r}[i]$ for *any* i in just degree L .

Furthermore, since $|\text{Mnml}(s_{q, \gamma})| = \lambda$ for any q, γ , the number of monomials in $\text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q, \gamma}))$ is bounded by $(Q\lambda + 1)^3 = \text{poly}(\lambda)$. Therefore, the total size of pre-computed monomials is

$$\left| \{\text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q, \gamma}))\}_{\gamma \in [\Gamma']} \right| \leq \Gamma' \text{poly}(\lambda) = S^{1/1+\alpha} \text{poly}(\lambda) , \quad (5)$$

which is still sublinear in the circuit size S and does not compromise the weak-compactness of the resulting FE for NC_D^1 scheme.

Putting Things Together So far, we showed how to “compress” the computation of degree 3 monomials over $\mathbf{r}[i]$, for any i , into a degree- L computation. To compute function g in Equation (3) in degree L , we need to additionally pre-compute multiplications with \mathbf{x} and b . As described in [Lin17], this can be done easily by pre-computing the following:

$$\mathbf{V}_1 = \{\text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q, \gamma}))\}_{\gamma \in [\Gamma']} \otimes (\mathbf{x} \| b \| 1)$$

(where the sets of monomials are first interpreted as a vector before taking tensor product.) Given the tensor product, one can compute *any* monomial with degree ≤ 3 in $r[i]$ for any i , degree ≤ 1 in \mathbf{x} , and degree ≤ 1 in b , in just degree L , which is sufficient for computing the first additive term in g_l in Equation (3). Similarly, to compute the second additive term in g_l , it suffices to precompute all multilinear monomials over every block in s' (of length Γ), and compute their tensor product with $b||1$, that is,

$$\mathbf{V}_2 = \{\text{Mnml}(s'_\gamma)\}_{\gamma \in [\Gamma]} \otimes (b||1)$$

In summary, for every $l \in [\ell]$, there exists a degree- L polynomial P_l that on input $(\mathbf{V}_1, \mathbf{V}_2)$ outputs $g_l(\mathbf{x}, s, s', b)$.

$$\text{Define } P_l := \text{the degree-}L \text{ polynomial s.t. } P_l(\mathbf{V}_1, \mathbf{V}_2) = g_l(\mathbf{x}, s, s', b) \quad (6)$$

Moreover, we show that both \mathbf{V}_1 and \mathbf{V}_2 have length sublinear in the circuit size. First, combining Equation (5) with the fact that $|\mathbf{x}||b||1| = N + 2$, we have that

$$|\mathbf{V}_1| \leq S^{1/1+\alpha} \text{poly}(\lambda) \times (N + 2) = S^{1/1+\alpha} \text{poly}(\lambda, N). \quad (7)$$

The size of \mathbf{V}_2 is

$$|\mathbf{V}_2| = \lambda \times \Gamma \times 2 \leq S^{1/1+\alpha} \text{poly}(\lambda). \quad (8)$$

Finally, to construct a 1-key weakly-compact FE scheme for NC_D^1 from FE for just degree L polynomials. We modify the LV construction as follows: 1) Instead of encrypting (\mathbf{x}, s, s', b) , the encryptor pre-computes and encrypts $\mathbf{V}_1||\mathbf{V}_2$ as described above, and 2) instead of generating secret keys for functions $\{g_l\}_{l \in [\ell]}$ which have degree $3L + 2$, generate secret keys for $\{P_l\}_{l \in [\ell]}$ which have only degree L . This way, at decryption time, the decryptor computes the correct output $\{P_l(\mathbf{V}_1||\mathbf{V}_2) = g_l(\mathbf{x}, s, s', b)\}$. The resulting new compact FE scheme **CFE** is described in Figure 3 (with key difference from the LV scheme highlighted). The compactness of the new scheme follows directly from the fact that the encrypted input $\mathbf{V}_1, \mathbf{V}_2$ have length sublinear in $S(\lambda)$, and that the degree- L FE scheme has linear efficiency. Moreover, its correctness and security follows from the same proof as that in [LV16]; since their security proof incur only a polynomial security loss, we conclude Proposition 1.

4.2 Step 2: Tight Construction of IO from Weakly-Compact FE

Proposition 2. *Let $i\ell$ and κ be defined as in Theorem 7. Assume the existence of 1-key weakly-compact public-key FE for P/poly , with $2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda)$ -security. There exists $i\ell(\lambda)$ -bit-input IO for P/poly , with $2^{-\kappa(\lambda)} \text{negl}(\lambda)$ -security.*

Using weakly-compact public key FE for P/poly , we construct IO for P/poly using induction over the input-length of the circuits that can be obfuscated. That is, in the $i\ell^{\text{th}}$ induction step, we define the behavior of the scheme when obfuscating circuits with input-length $i\ell$, by recursively invoking the scheme itself for obfuscating circuits with input length $i\ell - 1$ and utilizing the weakly compact FE scheme. In fact, for the induction to work out, we need to use a different “interface” for the induction hypothesis, namely, instead of inductively defining an indistinguishability obfuscator in the plain model, we inductively define an indistinguishability obfuscator in the CRS model. Below, we start with defining this notion.

Our Single-key Compact FE Scheme CFE

SETUP: $\text{CFE.Setup}(1^\lambda)$ samples $(\text{MPK}, \text{MSK}) \xleftarrow{\$} \text{FE.Setup}(1^\lambda)$, and $\text{PRG} \xleftarrow{\$} \text{PRG}_\lambda$.

ENCRYPTION: $\text{CFE.Enc}(\text{MPK}, \mathbf{x})$ samples

- a PRG seed \mathbf{s} viewed as a $Q \times \Gamma'$ matrix for $Q = \text{poly}(\lambda)$ and $\Gamma' = S^{1/1+\alpha}$, where each element $s_{q,\gamma}$ in \mathbf{s} is a block of $\log \lambda$ bits, and
- another PRG seed \mathbf{s}' viewed as a vector of length $\Gamma = S^{1/1+\alpha} \text{poly}(\lambda)$, where again each element s'_γ in \mathbf{s}' is a block of $\log \lambda$ bits.

Pre-Compute the following for $b = 0$:

$$\mathbf{V}_1 = \left\{ \text{Mnml}^{\leq 3}(\cup_q \text{Mnml}(s_{q,\gamma})) \right\}_{\gamma \in [\Gamma']} \otimes (\mathbf{x} \| b \| 1) \quad (9)$$

$$\mathbf{V}_2 = \left\{ \text{Mnml}(s'_\gamma) \right\}_{\gamma \in [\Gamma]} \otimes (b \| 1) \quad (10)$$

Finally generate:

$$\text{CT} \xleftarrow{\$} \text{FE.Enc}(\text{MPK}, (\mathbf{V}_1, \mathbf{V}_2))$$

KEY GENERATION: $\text{CFE.KeyGen}(\text{MSK}, h)$ does the following:

- Sample $\text{CT} \xleftarrow{\$} \{0, 1\}^\ell$, where ℓ is set below.
- Define function g as follows: On input \mathbf{x} of length N , PRG seeds \mathbf{s} and \mathbf{s}' of dimensions described above, and a bit b .

$g(\mathbf{x}, \mathbf{s}, \mathbf{s}', b)$ does the following:

- For every $i \in [S]$, let $h_i(\mathbf{x})$ denote the function that computes the i^{th} output bit of $h(\mathbf{x})$. Since $h \in \text{NC}_D^1$, h_i has depth $D(\lambda) = O(\log \lambda)$ and size $2^{D(\lambda)} = \text{poly}(\lambda)$.

- If $b = 0$, do:

Expand each row of \mathbf{s} using PRG to obtain a $Q \times S$ matrix \mathbf{r} of pseudo-random bits. That is, let s_i denote the i^{th} row of \mathbf{s} ; the i^{th} row \mathbf{r}_i of \mathbf{r} is $\text{PRG}(s_i)$. Denote by $\mathbf{r}[i]$ the i^{th} column of matrix \mathbf{r} , which has length $Q = \text{poly}(\lambda)$.

For every $i \in [S]$, compute the AIK encoding $\Pi[i]$ of computation (h_i, \mathbf{x}) as follows:

$$\forall i \in [S], \quad \Pi[i] = \text{AIK}(h_i, \mathbf{x}; \mathbf{r}[i]).$$

Output $\Pi = \{\Pi[i]\}_i$; set $\ell = |\Pi|$.

- If $b = 1$, output $\Pi = \text{CT} \oplus \text{PRG}(\mathbf{s}')$.

- For every $l \in [\ell]$, let P_l be the degree- L polynomial that on input $(\mathbf{V}_1, \mathbf{V}_2)$ in Equations (9) and (10) computes the l^{th} output bit of $g(\mathbf{x}, \mathbf{s}, \mathbf{s}', b)$.

For every l , generate a secret key $\text{SK}_l \xleftarrow{\$} \text{FE.KeyGen}(\text{MSK}, P_l)$ for P_l .

Output $\text{SK} = \{\text{SK}_l\}_{l \in [\ell]}$.

DECRYPTION: $\text{CFE.Dec}(\text{SK}, \text{CT})$ computes $\Pi = \{\text{FE.Dec}(\text{SK}_l, \text{CT})\}_{l \in [\ell]}$, parses $\Pi = \{\Pi[i]\}$, and decodes every $\Pi[i]$ using the AIK decoding algorithm to obtain the output $h(\mathbf{x})$.

Figure 3: Single-key Compact FE CFE from block-wise locality- L PRG and degree- L FE

4.2.1 IO in the CRS model

Definition 26. An indistinguishability obfuscator in the CRS model for a class of circuits $\{\mathcal{C}_\lambda\}_{\lambda \in \mathbb{N}}$ is a tuple of uniform PPT machines $(\text{SetupIO}, i\mathcal{O}, \text{EvalIO})$ satisfying the following conditions.

Syntax: The algorithms runs as follows:

- Setup: SetupIO on input a security parameter 1^λ , outputs a reference string σ , and a public parameter pp .
- Obfuscation with public parameter: $i\mathcal{O}$ on input a public parameter pp and a circuit $C \in \mathcal{C}_\lambda$, outputs an obfuscated circuit \hat{C} .
- Evaluation with the reference string: EvalIO on input an obfuscated circuit \hat{C} , an input x , and the reference string σ , outputs a string y .

Correctness: For all security parameters $\lambda \in \mathbb{N}$, for every $C \in \mathcal{C}_\lambda$, and every input x , we have that

$$\Pr[(\sigma, \text{pp}) \xleftarrow{\$} \text{SetupIO}(1^\lambda), \hat{C} \leftarrow i\mathcal{O}(\text{pp}, C) : \text{EvalIO}(\sigma, \hat{C}, x) = C(x)] = 1$$

μ -Indistinguishability: For every ensemble of pairs of circuits $\{C_{0,\lambda}, C_{1,\lambda}\}_{\lambda \in \mathbb{N}}$ satisfying that $C_{b,\lambda} \in \mathcal{C}_\lambda$, $C_{0,\lambda}$ and $C_{1,\lambda}$ have the same size, input-length, and truth table, the following ensembles of distributions are μ -indistinguishable:

$$\left\{ (\sigma, \text{pp}) \xleftarrow{\$} \text{SetupIO}(1^\lambda) : \sigma, \text{pp}, C_{0,\lambda}, C_{1,\lambda}, i\mathcal{O}(\text{pp}, C_{0,\lambda}) \right\}_{\lambda \in \mathbb{N}}$$

$$\left\{ (\sigma, \text{pp}) \xleftarrow{\$} \text{SetupIO}(1^\lambda) : \sigma, \text{pp}, C_{0,\lambda}, C_{1,\lambda}, i\mathcal{O}(\text{pp}, C_{1,\lambda}) \right\}_{\lambda \in \mathbb{N}}$$

Toward proving Proposition 2 for constructing il -bit-input IO for P/poly in the plain model, we first construct il -bit-input IO for P/poly in the CRS model as defined below. We say that a polynomial p is upper bounded by another q , if it holds that for all $\lambda \in \mathbb{N}$, $p(\lambda) \leq q(\lambda)$.

Definition 27. Let il be any polynomial upper bounded by λ . An il -bit-input indistinguishability obfuscator in the CRS model for P/poly is a tuple of uniform PPT machines $(\text{SetupIO}, i\mathcal{O}, \text{EvalIO})$ satisfying the following condition: SetupIO takes two inputs $(1^\lambda, 1^l)$, and for every polynomial l upper bounded by il , $(\text{SetupIO}^l(1^\lambda) = \text{SetupIO}(1^\lambda, 1^{l(\lambda)}), i\mathcal{O}, \text{EvalIO})$ is an IO scheme in the CRS model for the class of circuits $\mathcal{C}^l = \{\mathcal{C}_\lambda^l\}_{\lambda \in \mathbb{N}}$ containing all circuits with size at most λ and input-length at most $l(\lambda)$.

When $il(\lambda) = \lambda$, we say that the scheme is a bounded-input indistinguishability obfuscator for P/poly in the CRS model.

It follows from a very simple argument that il -bit-input IO for P/poly in the CRS model actually implies that in the plain model.

Claim 1. Let il be any polynomial upper bounded by λ . An il -bit-input indistinguishability obfuscator for P/poly in the CRS model $(\text{SetupIO}, i\mathcal{O}, \text{EvalIO})$ with μ -security, implies an il -bit-input indistinguishability obfuscator for P/poly in the plain model $i\mathcal{O}'$ with μ -security.

Proof. The obfuscator $i\mathcal{O}'$ on input 1^λ and a circuit C with size $|C| < \lambda$ and input-length l works as follows: It samples a fresh reference string and a public parameter $(\sigma, \text{pp}) \xleftarrow{\$} \text{SetupIO}(1^\lambda, 1^l)$ using the input-length of C as the input-length bound, obfuscate C to obtain $\hat{C} \xleftarrow{\$} i\mathcal{O}(\text{pp}, C)$, and outputs $\tilde{C} = (\sigma, \text{pp}, \hat{C})$ as the new obfuscated circuit. The correctness and security of the new obfuscator follows directly from that of the obfuscator in the CRS model. \square

4.2.2 IO in the CRS Model from Weakly-Compact FE

We now construct an $i\ell$ -bit-input indistinguishability obfuscator for P/poly in the CRS model (SetupIO, $i\mathcal{O}$, EvalIO), from a public-key weakly-compact FE scheme. To do so, we define the behavior of SetupIO, $i\mathcal{O}$, EvalIO and analyze their correctness, efficiency and security, by induction over the input-length $i\ell$ the algorithms are executed with. More precisely, the construction makes use of the following building blocks:

- Public key FE schemes $\mathbf{CFE} = \{\mathbf{CFE}^{N,D,S}\}$ for P/poly with $(1 - \varepsilon)$ -sublinear compact for arbitrary constant $\varepsilon > 0$, where the scheme $\mathbf{CFE}^{N,D,S} = (\mathbf{CFE}.\text{Setup}, \mathbf{CFE}.\text{KeyGen}, \mathbf{CFE}.\text{Enc}, \mathbf{CFE}.\text{Dec})$ is a scheme that handles circuits with input-length $N = N(\lambda)$, depth $D = D(\lambda)$, and size $S = S(\lambda)$. Such FE schemes for any polynomials N, D, S are constructed in Section 4.1.
- Two PRGs $\text{PRG}_1, \text{PRG}_2$ mapping λ -bit inputs to sufficiently long outputs, of $m_1(\lambda)$ -bit and $m_2(\lambda)$ -bit respectively. The mapping takes time $\text{poly}(\lambda)m_1(\lambda)$ and $\text{poly}(\lambda)m_2(\lambda)$ respectively. (Such PRGs can be implemented for example using a PRF.)

For any $\lambda \in \mathbb{N}$, we define the behavior of SetupIO, $i\mathcal{O}$, EvalIO when executed with security parameter 1^λ , by induction on the input-length $i\ell$ they are executed with (more precisely, when SetupIO is invoked with input $(1^\lambda, 1^{i\ell})$, and $i\mathcal{O}$ and EvalIO are executed with the produced public parameter and reference string). The scheme is described in Figure 4, which recursively calls itself. Observe that scheme is well-defined, since when it is executed with an input length $i\ell + 1$, it recursively invokes itself with input length $i\ell$; in the base case when $i\ell = 1$, the IO scheme is trivially defined, by letting the obfuscator outputs the only two possible outputs $C(0), C(1)$.

Next, we analyze the correctness, efficiency, and security of the IO scheme in the CRS model defined in Figure 4 by induction, in Lemma 6, 7, and 8.

Lemma 6 (Induction on Correctness). *For any $\lambda \in \mathbb{N}$ and any $i\ell \in \mathbb{N}$, if the IO scheme (SetupIO, $i\mathcal{O}$, EvalIO) in Figure 4 satisfy correctness when executed with security parameter 1^λ and input-length $1^{i\ell}$, for any circuit C with size at most λ and input-length $i\ell$, then the scheme also satisfies correctness when executed with 1^λ and $1^{i\ell+1}$, for any circuit C with size at most λ and input-length $i\ell + 1$.*

Proof. When executed with $1^\lambda, 1^{i\ell+1}$, the setup algorithm SetupIO recursively calls itself with input $(1^\lambda, 1^{i\ell})$ to obtain $(\sigma^{i\ell}, \text{pp}^{i\ell})$. Moreover, it also calls the setup algorithm of the FE scheme to obtain (MPK, MSK), and generates a secret key SK for a function f , which on input $(C, s_0, s_1, 0^\lambda, 0)$ obfuscates the circuit C with the first input-bit fixed to 0 and 1, namely $C^{(0)} = C(0, \star)$ and $C^{(1)} = C(1, \star)$, using $\text{pp}^{i\ell}$ (and pseudorandomness generated by a PRG). Finally, the reference string and public parameter are set to $\sigma^{i\ell+1} = (\sigma^{i\ell}, \text{SK})$ and $\text{pp}^{i\ell+1} = (\text{pp}^{i\ell}, \text{MPK})$.

The obfuscation of a circuit C with size at most λ and input-length $i\ell + 1$ generated with public parameter $\text{pp}^{i\ell+1}$ is simply a FE ciphertext CT of $(C, s_0, s_1, 0^\lambda, 0)$ with fresh random PRG seeds.

Evaluation decrypts the ciphertext CT using SK in $\sigma^{i\ell+1}$. By the correctness of the FE scheme, decryption outputs $f(C, s_0, s_1, 0^\lambda, 0) = (\hat{C}^{(0)}, \hat{C}^{(1)})$, where $\hat{C}^{(d)}$ is an obfuscation of $C^{(d)}$ generated using $\text{pp}^{i\ell}$. Thus, by the correctness of the IO scheme when executed with input length $i\ell$, for any $x \in \{0, 1\}^{i\ell+1}$, evaluation outputs the right output $C_{x_1}(x_{\geq 2}) = C(x)$, which concludes the proof of the lemma. \square

Lemma 7 (Induction on Efficiency). *There exists sufficiently large universal polynomials P, Q , and T , such that, the following holds for the IO scheme (SetupIO, $i\mathcal{O}$, EvalIO) in Figure 4: For every λ and input-length $i\ell$,*

An IO scheme (SetupIO, $i\mathcal{O}$, EvalIO) in the CRS model

Base Case — Input-length $i\ell = 1$

- SetupIO($1^\lambda, 1$) outputs $(\sigma^1, \text{pp}^1) = (\text{base}, \text{base})$.
- $i\mathcal{O}(\text{pp}^1 = \text{base}, C)$ outputs $(C(0), C(1))$.
- $i\mathcal{O}(\sigma^1 = \text{base}, (y_0, y_1), x)$ outputs \perp if x is not a single bit, and otherwise outputs y_x .

Induction — From Input-length $i\ell$ to $i\ell + 1$

- SetupIO($1^\lambda, 1^{i\ell+1}$) does:
 - Recursively setup the scheme for input-length $i\ell$ to obtain $(\sigma^{i\ell}, \text{pp}^{i\ell}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell})$.
 - Sample a random string $\text{CT} \stackrel{\$}{\leftarrow} \{0, 1\}^\ell$ of length ℓ set below.
 - Sample master keys of the weakly-compact FE scheme $(\text{MPK}, \text{MSK}) \stackrel{\$}{\leftarrow} \text{CFE.Setup}(1^\lambda)$.
 - Generate a FE secret key $\text{SK}_f \stackrel{\$}{\leftarrow} \text{CFE.KeyGen}(\text{MSK}, f)$ for function f defined as follows:

$f(C, s_0, s_1, s', b)$ does the following, on input a circuit C of size at most S and input length at most $l + 1$, PRG seeds s_0, s_1, s' of length λ , and a bit b .

 - * If $b = 0$, let $C^{(d)} = C(d, \star)$ denote the circuit C with the first input bit fixed to d . Recursively call the IO scheme to obfuscate $C^{(0)}$ and $C^{(1)}$, which have ℓ -bit input, using pseudo-randomness generated through a PRG,

$$\left\{ \hat{C}^{(d)} = i\mathcal{O}(\text{pp}^{i\ell}, C^{(d)}; \text{PRG}_1(s_d)) \right\}_{d \in \{0,1\}} .$$

Output $z = (\hat{C}^{(0)}, \hat{C}^{(1)})$; set $\ell = |z|$.

 - * If $b = 1$, output $z = \text{CT} \oplus \text{PRG}_2(s')$.
- Output $\sigma^{i\ell+1} = (\sigma^{i\ell}, \text{SK}_f)$ and $\text{pp}^{i\ell+1} = (\text{pp}^{i\ell}, \text{MPK})$.

- $i\mathcal{O}(\text{pp}^{i\ell+1} = (\text{pp}^{i\ell}, \text{MPK}), C)$ does:

Sample two PRG seeds $s_0, s_1 \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$, and encrypt $(C, s_0, s_1, 0^\lambda, 0)$ using FE,

$$\text{CT} \stackrel{\$}{\leftarrow} \text{CFE.Enc}(\text{MPK}, (C, s_0, s_1, 0^\lambda, 0))$$

Output the ciphertext $\hat{C} = \text{CT}$ as the obfuscated circuit.

- EvalIO($\sigma^{i\ell+1} = (\sigma^{i\ell}, \text{SK}), \hat{C} = \text{CT}, x$) does:

- Decrypt the FE ciphertext and secret key $z = \text{CFE.Dec}(\text{SK}, \text{CT})$; parse z as $(\hat{C}^{(0)}, \hat{C}^{(1)})$.
- Recursively call the IO scheme to evaluate $y = \text{EvalIO}(\sigma^{i\ell}, \hat{C}_{x_1}, x_{\geq 2})$, where $x_{\geq 2}$ are the last l bits of x .

Figure 4: An IO scheme (SetupIO, $i\mathcal{O}$, EvalIO) in the CRS model from weakly-compact public-key FE.

- Suppose that the scheme when executed with $1^\lambda, 1^{i\ell}$ for arbitrary circuit C with size at most λ and input-length $i\ell$, and input $x \in \{0, 1\}^{i\ell}$, satisfy the following efficiency

$$\begin{aligned} (\sigma^{i\ell}, \text{pp}^{i\ell}) &\stackrel{\S}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell}) & \text{Time}_{\text{SetupIO}}(1^\lambda, 1^{i\ell}) &\leq i\ell \times Q(\lambda), \\ \hat{C} &\stackrel{\S}{\leftarrow} i\mathcal{O}(\text{pp}^{i\ell}, C) & \text{Time}_{i\mathcal{O}}(\text{pp}^{i\ell}, C) &\leq P(\lambda), \\ y = \text{EvalIO}(\sigma^{i\ell}, \hat{C}, x) & & \text{Time}_{\text{EvalIO}}(\sigma^{i\ell}, \hat{C}, x) &\leq i\ell \times T(\lambda). \end{aligned}$$

- Then, they satisfy the following efficiency when executed with $1^\lambda, 1^{i\ell+1}$ for arbitrary circuit C with size at most λ and input-length $i\ell + 1$, and input $x \in \{0, 1\}^{i\ell+1}$,

$$\begin{aligned} (\sigma^{i\ell+1}, \text{pp}^{i\ell+1}) &\stackrel{\S}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell+1}) & \text{Time}_{\text{SetupIO}}(1^\lambda, 1^{i\ell+1}) &\leq (i\ell + 1) \times Q(\lambda), \\ \hat{C} &\stackrel{\S}{\leftarrow} i\mathcal{O}(\text{pp}^{i\ell+1}, C) & \text{Time}_{i\mathcal{O}}(\text{pp}^{i\ell+1}, C) &\leq P(\lambda), \\ y = \text{EvalIO}(\sigma^{i\ell+1}, \hat{C}, x) & & \text{Time}_{\text{EvalIO}}(\sigma^{i\ell+1}, \hat{C}, x) &\leq (i\ell + 1) \times T(\lambda). \end{aligned}$$

Proof. To analyze the efficiency of the IO scheme in an execution with $1^\lambda, 1^{i\ell+1}$, circuit C , and input x as specified in the lemma statement, we start with analyzing the size S and input-length N of the circuit that computes the function f for which the FE secret key is generated. The function f takes input of form $(\bar{C}, s_0, s_1, s', b)$ for any circuit with size $|\bar{C}| \leq \lambda$ and input-length $i\ell + 1$, and PRG seeds of length λ . Thus, its input-length is

$$N = N(\lambda) = 4\lambda + 1.$$

If the input bit $b = 0$, f obfuscates $\bar{C}^{(0)}$ and $\bar{C}^{(1)}$ using the IO scheme with a public parameter $\text{pp}^{i\ell}$ generated for input-length $i\ell$ and pseudorandomness produced by $\text{PRG}_1(s_0), \text{PRG}_1(s_1)$. By our hypothesis on the efficiency of the IO scheme for input-length $i\ell$, and the efficiency of PRG_1 , this step takes at most time $\text{poly}(\lambda) \times P(\lambda)$. If $b = 1$, f “decrypts” a hardwired one-time-pad ciphertext CT using pseudorandom pad generated by $\text{PRG}_2(s')$. The length of the ciphertext CT is the same as the length of the obfuscation of $\bar{C}^{(0)}$ and $\bar{C}^{(1)}$, and hence is bounded by $2P(\lambda)$. Thus, by the efficiency of PRG_2 , this step takes at most time $\text{poly}(\lambda) \times P(\lambda)$. In summary, f can be computed by a circuit with size

$$S \leq \text{poly}(\lambda) \times P(\lambda).$$

We now analyze the runtime of SetupIO , $i\mathcal{O}$, and EvalIO , when executed with $1^\lambda, 1^{i\ell+1}$, C with input-length $i\ell + 1$, and $x \in \{0, 1\}^{i\ell+1}$.

$$\begin{aligned} \text{Time}_{\text{SetupIO}}(1^\lambda, 1^{i\ell+1}) &= \text{Time}_{\text{SetupIO}}(1^\lambda, 1^{i\ell}) + \text{Time}_{\text{CFE.Setup}}(1^\lambda) + \text{Time}_{\text{CFE.KeyGen}}(\text{MSK}, f) + \text{poly}(\lambda) \\ &= i\ell \times Q(\lambda) + \text{poly}(\lambda) + \text{poly}(\lambda, S) + \text{poly}(\lambda) \\ &\leq i\ell \times Q(\lambda) + \text{poly}(\lambda, \text{poly}(\lambda)P(\lambda)) + \text{poly}(\lambda) \\ &\leq (i\ell + 1) \times Q(\lambda) \end{aligned}$$

where the second equality follows from the hypothesis on the runtime of SetupIO with input-length $i\ell$ and the efficiency of the FE scheme, the first inequality follows from plugging in the bound on S , and the last inequality follows if $Q(\lambda)$ is sufficiently large (comparing to P).

$$\begin{aligned} \text{Time}_{i\mathcal{O}}(\text{pp}^{i\ell+1}, C) &= \text{Time}_{\text{CFE.Enc}}(\text{MPK}, (C, s_0, s_1, 0^\lambda, 0)) + \text{poly}(\lambda) \\ &= \text{poly}(\lambda, N)S^{1-\varepsilon} + \text{poly}(\lambda) \\ &\leq \text{poly}(\lambda)(\text{poly}(\lambda)P(\lambda))^{1-\varepsilon} + \text{poly}(\lambda) \\ &\leq P(\lambda) \end{aligned}$$

where the second equality follows from the $(1 - \varepsilon)$ -sublinear compactness of the FE scheme, the first inequality follows from plugging the bound on S , and the last inequality holds when P is sufficiently large.

$$\begin{aligned}
\text{Time}_{\text{EvalIO}}(\sigma^{i\ell+1}, \hat{C}, x) &= \text{Time}_{\text{CFE.Dec}}(\text{SK}, \text{CT}) + \text{Time}_{\text{EvalIO}}(\sigma^{i\ell}, \hat{C}_{x_1}, x_{\geq 2}) + \text{poly}(\lambda) \\
&= \text{poly}(\lambda, S) + i\ell \times T(\lambda) + \text{poly}(\lambda) \\
&\leq \text{poly}(\lambda, \text{poly}(\lambda)P(\lambda)) + i\ell \times T(\lambda) + \text{poly}(\lambda) \\
&\leq (i\ell + 1) \times T(\lambda)
\end{aligned}$$

where the second equality follows from the efficiency of FE decryption, and the hypothesis on the runtime of EvalIO for evaluating obfuscated circuits with $i\ell$ -bit inputs, the first inequality follows from plugging in the bound on S , and last inequality holds when T is sufficiently large (comparing to P).

We conclude the proof. \square

Lemma 8 (Induction on Security). *Let $\mu, \mu_{\text{FE}}, \mu_{\text{PRG}}$ be any functions from \mathbb{N} to $[0, 1]$. Suppose that the following holds:*

- The FE scheme **CFE** is μ_{FE} -Full-Sel-secure.
- The PRGs PRG_1 and PRG_2 are both μ_{PRG} -indistinguishable.
- For every efficient distinguisher D , every $\lambda \in \mathbb{N}$, $i\ell < \lambda \in \mathbb{N}$, C_0, C_1 with size $|C_0| = |C_1| \leq \lambda$, input-length $i\ell$, and identical truth table, the following holds.

$$\begin{aligned}
&\left| \Pr[(\sigma^{i\ell}, \text{pp}^{i\ell}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell}) : D(\sigma^{i\ell}, \text{pp}^{i\ell}, C_0, C_1, i\mathcal{O}(\text{pp}^{i\ell}, C_0)) = 1] \right. \\
&= \left. \Pr[(\sigma^{i\ell}, \text{pp}^{i\ell}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell}) : D(\sigma^{i\ell}, \text{pp}^{i\ell}, C_0, C_1, i\mathcal{O}(\text{pp}^{i\ell}, C_1)) = 1] \right| \leq \mu(\lambda)
\end{aligned}$$

Then, for every efficient distinguisher D , $\lambda \in \mathbb{N}$, $i\ell + 1 < \lambda \in \mathbb{N}$, C_0, C_1 with size $|C_0| = |C_1| \leq \lambda$, input-length $i\ell + 1$, and identical truth table, the following holds.

$$\begin{aligned}
&\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell+1}) : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, i\mathcal{O}(\text{pp}^{i\ell+1}, C_0)) = 1] \right. \\
&- \left. \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell+1}) : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, i\mathcal{O}(\text{pp}^{i\ell+1}, C_1)) = 1] \right| \\
&\leq 2\mu(\lambda) + 2\mu_{\text{FE}}(\lambda) + 6\mu_{\text{PRG}}(\lambda)
\end{aligned}$$

Proof. Fix any efficient distinguisher D , $\lambda, i\ell + 1 < \lambda, C_0, C_1$ with size $|C_0| = |C_1| \leq \lambda$, input-length $i\ell + 1$, and identical truth table, we prove the lemma through a sequence of hybrids.

Hybrid Distribution \mathcal{D}_0^b : Sample $(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C})$ using the IO scheme honestly, obfuscating the circuit C_b , that is,

$$\mathcal{D}_0^b : (\sigma^{i\ell+1}, \text{pp}^{i\ell+1}) \stackrel{\$}{\leftarrow} \text{SetupIO}(1^\lambda, 1^{i\ell+1}); \hat{C} \stackrel{\$}{\leftarrow} i\mathcal{O}(\text{pp}^{i\ell+1}, C_b).$$

Parse $\sigma^{i\ell+1} = (\sigma^{i\ell}, \text{SK})$, $\text{pp}^{i\ell+1} = (\text{pp}^{i\ell}, \text{MPK})$, and $\hat{C} = \text{CT}$.

Hybrid Distribution \mathcal{D}_1^b : Sample $(\sigma^{i\ell+1} \text{pp}^{i\ell+1}, \hat{C})$ identically to \mathcal{D}_0^b , except that the secret key SK in $\sigma^{i\ell+1}$ is generated differently.

In \mathcal{D}_0^b , the SK is associated with a function f that is hardwired with a random string **CT**, and the obfuscated circuit \hat{C} is a ciphertext CT encrypting input $(C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0)$. During evaluation, CT is decrypted by SK producing $f(C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0)$, which by definition of f is

$$f(C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0) = \left(\hat{C}^{(d)} = i\mathcal{O}(\text{pp}^{i\ell}, C_b^{(d)}; \text{PRG}_1(\mathbf{s}_d)) \right)_{d \in \{0,1\}}. \quad (11)$$

In \mathcal{D}_1^b , **CT** is no longer a random string, and instead is set to a one-time-pad ciphertext encrypting $(\hat{C}^{(0)}, \hat{C}^{(1)})$ as described above using pad $\text{PRG}_2(\mathbf{s}')$, that is,

$$\mathbf{CT} = (\hat{C}^{(0)}, \hat{C}^{(1)}) \oplus \text{PRG}_2(\mathbf{s}').$$

The only difference between \mathcal{D}_0^b and \mathcal{D}_1^b is whether **CT** is a random string or a one-time-pad ciphertext with pad $\text{PRG}_2(\mathbf{s}')$. Since the seed \mathbf{s}' is not used anywhere else in \mathcal{D}_0^b and \mathcal{D}_1^b , by the μ_{PRG} -indistinguishability of PRG_2 , we have that the distinguisher D distinguishes these two distributions with advantage at most $\mu_{\text{PRG}}(\lambda)$, that is,

$$\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \xleftarrow{\$} \mathcal{D}_0^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] - \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \xleftarrow{\$} \mathcal{D}_1^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right| \leq \mu_{\text{PRG}}(\lambda).$$

Hybrid \mathcal{D}_2^b : Sample $(\sigma^{i\ell+1} \text{pp}^{i\ell+1}, \hat{C})$ identically to \mathcal{D}_1^b , except that the obfuscated circuit \hat{C} is generated differently.

In \mathcal{D}_1^b , \hat{C} is a FE ciphertext CT encrypting $(C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0)$,

$$\hat{C} = \text{CT} \xleftarrow{\$} \text{CFE.Enc}(\text{MPK}, (C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0)).$$

In \mathcal{D}_2^b , \hat{C} becomes a ciphertext CT of $(C_b, 0^\lambda, 0^\lambda, \mathbf{s}', 1)$,

$$\hat{C} = \text{CT} \xleftarrow{\$} \text{CFE.Enc}(\text{MPK}, (C_b, 0^\lambda, 0^\lambda, \mathbf{s}', 1)).$$

Note that in both \mathcal{D}_1^b and \mathcal{D}_2^b , **CT** is set to $(\hat{C}^{(0)}, \hat{C}^{(1)}) \oplus \text{PRG}_2(\mathbf{s}')$. With such **CT**, by definition of f ,

$$f(C_b, \mathbf{s}_0, \mathbf{s}_1, 0^\lambda, 0) = f(C_b, 0^\lambda, 0^\lambda, \mathbf{s}', 1).$$

Therefore, by the μ_{FE} -Full-Sel-security of **CFE**, we have that the distinguisher D distinguishes distributions \mathcal{D}_1^b and \mathcal{D}_2^b with advantage at most $\mu_{\text{FE}}(\lambda)$, that is,

$$\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \xleftarrow{\$} \mathcal{D}_1^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] - \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \xleftarrow{\$} \mathcal{D}_2^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right| \leq \mu_{\text{FE}}(\lambda).$$

Hybrid \mathcal{D}_3^b : Sample $(\sigma^{i\ell+1}\text{pp}^{i\ell+1}, \hat{C})$ identically to \mathcal{D}_2^b , except that again the string **CT** hardwired in the function f associated with the FE secret key SK in $\sigma^{i\ell+1}$ is generated differently.

In \mathcal{D}_2^b , **CT** is the one-time-pad encryption of the obfuscation $(\hat{C}^{(0)}, \hat{C}^{(1)})$ of $(C_b^{(0)}, C_b^{(1)})$ using pseudorandomness generated by $\text{PRG}_1(\mathbf{s}_0)$ and $\text{PRG}_1(\mathbf{s}_1)$, as described in Equation (11).

In \mathcal{D}_3^b , **CT** is the one-time-pad encryption of still obfuscation $(\tilde{C}^{(0)}, \tilde{C}^{(1)})$ of $(C_b^{(0)}, C_b^{(1)})$, but generated with truly random coins $\mathbf{r}_0, \mathbf{r}_1$,

$$\left(\tilde{C}^{(d)} = i\mathcal{O}(\text{pp}^{i\ell}, C_b^{(d)}; \mathbf{r}_d) \right)_{d \in \{0,1\}}, \quad \mathbf{CT} = (\tilde{C}^{(0)}, \tilde{C}^{(1)}) \oplus \text{PRG}_2(\mathbf{s}').$$

The only difference between \mathcal{D}_2^b and \mathcal{D}_3^b is whether the encrypted obfuscation are generated using pseudorandomness $\text{PRG}_1(\mathbf{s}_0), \text{PRG}_1(\mathbf{s}_1)$ or using true randomness $\mathbf{r}_0, \mathbf{r}_1$. Since in both distributions $\mathbf{s}_0, \mathbf{s}_1$ are not used anywhere else, it follows from the μ_{PRG} -indistinguishability of the PRGs that

$$\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_2^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right. \\ \left. - \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_3^b : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right| \leq 2\mu_{\text{PRG}}(\lambda).$$

Next, observe that the only difference between \mathcal{D}_3^0 and \mathcal{D}_3^1 lies in that in the former, $(\tilde{C}^{(0)}, \tilde{C}^{(1)})$ obfuscate $(C_0^{(0)}, C_0^{(1)})$, whereas in the latter they obfuscate $(C_1^{(0)}, C_1^{(1)})$. By the fact that C_0 and C_1 are equivalent (in the sense of having the same size, input-length, and truth table), so are $C_0^{(d)}$ and $C_1^{(d)}$ for any $d \in \{0, 1\}$. Therefore, for any d , by our hypothesis on the security of the IO scheme for circuits with $i\ell$ -bit inputs, the obfuscation of $C_0^{(d)}$ and $C_1^{(d)}$ with input-length $i\ell$ is $\mu(\lambda)$ indistinguishable. Hence,

$$\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_3^0 : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right. \\ \left. - \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_3^1 : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right| \leq 2\mu(\lambda)$$

Finally, it follows from a hybrid argument that

$$\left| \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_0^0 : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right. \\ \left. - \Pr[(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, \hat{C}) \stackrel{\$}{\leftarrow} \mathcal{D}_0^1 : D(\sigma^{i\ell+1}, \text{pp}^{i\ell+1}, C_0, C_1, \hat{C}) = 1] \right| \leq 2\mu(\lambda) + 2\mu_{\text{FE}}(\lambda) + 6\mu_{\text{PRG}}(\lambda)$$

This concludes the proof of the lemma. \square

With Lemma 6, 7, and 8, we now prove Proposition 2.

Proof of Proposition 2. Assume the existence of 1-key weakly-compact public-key FE for P/poly, with $2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda)$ -security. This implies the existence of PRGs with $2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda)$ -indistinguishability. That is,

$$\mu_{\text{FE}}(\lambda) = \mu_{\text{PRG}}(\lambda) = \delta = 2^{-i\ell(\lambda)+\kappa(\lambda)} \text{negl}(\lambda)$$

To show that there exists $i\ell(\lambda)$ -bit-input IO for P/poly, with $2^{-\kappa(\lambda)} \text{negl}(\lambda)$ -security, by Claim 1, it suffices to show that there exists $i\ell(\lambda)$ -bit-input IO for P/poly in the CRS model with $2^{-\kappa(\lambda)} \text{negl}(\lambda)$ -security. We argue that the construction in Figure 4 is such a scheme.

First observe that in the base case, when $(\text{SetupIO}, i\mathcal{O}, \text{EvalIO})$ are executed with $(1^\lambda, 1)$, the scheme is correct and efficient (for any circuit with size $\leq \lambda$ and single-bit inputs and any input bit x), and perfectly secure with distinguishing gap 0 (for any two single-bit-input circuits with the same size $\leq \lambda$ and truth tables). In other words, when $i\ell = 1$, the scheme satisfy the premises of Lemma 6, 7, and 8.

Therefore, applying these lemmas, we have that for every λ and every $l \leq i\ell(\lambda)$, when $(\text{SetupIO}, i\mathcal{O}, \text{EvalIO})$ are executed with $(1^\lambda, 1^l)$, the scheme is still correct, efficient, and has distinguishing gap μ_l bounded as below.

$$\begin{aligned} \mu_l &\leq 2\mu_{l-1} + 2\mu_{\text{FE}} + 6\mu_{\text{PRG}} = 2\mu_{l-1} + 8\delta \\ &\leq 2(2\mu_{l-2} + 8\delta) + 8\delta \leq \dots \leq 2^{l-1}\mu_1 + \left(\sum_{i=0}^{l-2} 2^i\right) 8\delta \\ &\leq 2^{l+2}\delta = 2^{l+2}2^{-i\ell+\kappa} \text{negl}(\lambda) = 2^{-(i\ell-l+\kappa)} \text{negl}(\lambda) \end{aligned}$$

The fifth inequality follows from the fact that $\mu_1 = 0$. Thus, $\mu_{i\ell} = 2^{-\kappa} \text{negl}(\lambda)$, which concludes the proposition. \square

5 FE from $\omega(\log \lambda)$ -Bit-Input IO for P/poly

In this section, we show Theorem 4, i.e., we prove via a new transformation that adaptively-secure collusion-resistant public-key functional encryption for P/poly is implied by IO for circuits with short, $\omega(\log \lambda)$ -bit, inputs and public key encryption, both with slightly super-polynomial security. Note that, in contrast, previous constructions of collusion-resistant FE for P/poly either rely on multilinear maps [GGHZ16], or require IO for all P/poly, including circuits with long (polynomial) inputs [GGH⁺13b].

Our proof generically transforms any 1-key (public key) FE scheme for any circuit class \mathcal{C} into a collusion-resistant (public key) FE scheme for the same circuit class, using IO for circuits with $\omega(\log \lambda)$ -bit inputs. The encryption time of the resulting FE schemes is polynomial in the encryption time of the original schemes, and hence if the original scheme is (non-)compact, so is the resulting FE scheme. The transformation also preserves the same type of security — namely Full-Sel- or Adap-security— and incurs a $2^{\omega(\log \lambda)}$ security loss.

More precisely, we prove the following below in Section 5.1.

Proposition 3. *Let \mathcal{C} be any circuit class, τ be any polynomial, and $i\ell$ be any polynomial such that $i\ell(\lambda) = \omega(\log \lambda) \leq \lambda$. Assume the existence of an $i\ell(\lambda)$ -bit-input indistinguishability obfuscator $i\mathcal{O}$ for P/poly. Then, any 1-key public-key FE schemes **OFE** for \mathcal{C} can be generically transformed into collusion-resistant FE schemes **CRFE** for \mathcal{C} , with the following properties:*

- The encryption time of **CRFE** is polynomial in the encryption time of **OFE**.
- If $i\mathcal{O}$ is $2^{-(i\ell(\lambda)+\tau(\lambda))} \text{negl}(\lambda)$ -secure and **OFE** is $2^{-(i\ell(\lambda)+\tau(\lambda))} \text{negl}(\lambda)$ -(Adap or Full-Sel)-secure, then **CRFE** is $2^{-\tau(\lambda)} \text{negl}(\lambda)$ -(Adap or Full-Sel)-secure.

It is known that adaptively-secure 1-key non-compact public-key FE for P/poly can be constructed from just pulic key encryption [GVW12].

Theorem 8 (1-Key Adap-Secure Public-Key FE for P/poly [GVW12]). *Let μ be any function from \mathbb{N} to $[0, 1]$. Assuming public key encryption with $\mu(\lambda) \text{negl}(\lambda)$ -security, there exist $\mu(\lambda) \text{negl}(\lambda)$ -Adap-secure 1-key non-compact public-key FE schemes for P/poly.*

Now, applying the transformation of Proposition 3 to the μ negl-Adap-secure 1-key FE schemes for P/poly with $\mu = 2^{-(i\ell+\tau)}$, yields $2^{-\tau}$ negl-Adap-secure *collusion-resistant* (non-compact public-key) FE for P/poly. Finally, note that it follows from [AJS15] that collusion-resistant non-compact FE schemes implies collusion-resistant compact FE schemes with the same level of security, which yields Theorem 4.

5.1 From 1-key to Collusion-Resistant FE, Generically

In this section, we prove Proposition 3. Let us fix any circuit class \mathcal{C} , and any $i\ell$ such that $i\ell(\lambda) = \omega(\log \lambda) \leq \lambda$. The resulting collusion-resistant FE scheme for \mathcal{C} , denoted **CRFE** = (CRFE.Setup, CRFE.KeyGen, CRFE.Enc, CRFE.Dec), relies on the following building blocks:

- An $i\ell$ -bit-input indistinguishability obfuscator $i\mathcal{O}$ for P/poly.
- A 1-key FE scheme **OFE** = (OFE.Setup, OFE.KeyGen, OFE.Enc, OFE.Dec) for \mathcal{C} .
- A puncturable PRF scheme PPRF = (PRF.Gen, PRF.Punc, F).

Given the above building blocks, to construct collusion resistant FE **CRFE** for \mathcal{C} , we start with the following intuition. *If efficiency were not a problem*, we could trivially construct a FE scheme that support releasing any polynomial number of secret keys, essentially by using a super-polynomial number of instances of **OFE**. Concretely, we would proceed as follows:

- Setup: Generate a super-polynomial number, $M = 2^{i\ell(\lambda)} = 2^{\omega(\lambda)}$, of **OFE** instances with master keys $\{(OMPK_i, OMSK_i) \stackrel{\$}{\leftarrow} \text{OFE.Setup}(1^\lambda)\}_{i \in [M]}$.
- Key Generation: To generate a key for a function f , sample an index at random $i_f \stackrel{\$}{\leftarrow} [M]$ and generate a secret key using the i_f^{th} master secret key $OSK_{i_f} \stackrel{\$}{\leftarrow} \text{OFE.KeyGen}(OMSK_{i_f}, f)$. Since there are at most a polynomial number of secret keys ever generated, the probability that every **OFE** instance is used to generate at most one secret key is overwhelming.
- Encryption: To encrypt any input x , simply encrypt the input x under all master public keys, $\{OCT_i \stackrel{\$}{\leftarrow} \text{OFE.Enc}(OMPK_i, x)\}_{i \in [M]}$. Given the set of ciphertexts, one can compute the output $f(x)$ of any function f for which a secret key OSK_{i_f} has been generated, by decrypting the appropriated ciphertext OCT_{i_f} using the secret key OSK_{i_f} .

Of course, the only problem with this FE scheme is that its setup and encryption algorithms run in super-polynomial time. To address this, we follow the previously adopted idea (e.g. [BGL⁺15, CLTV15]) of using IO to “compress” these super-polynomially many **OFE** instances into “polynomial size”. More precisely, instead of having the setup algorithm publish all M master public keys, let it generate an *obfuscated circuit* that on input $i \in [M]$ outputs the i^{th} master public key. Similarly, instead of having the encryption algorithm publish M ciphertexts, let it generate an obfuscated circuit that on input $i \in [M]$ outputs the i^{th} ciphertext under the i^{th} master public key. Since the inputs to the obfuscated circuits are indexes from the range $[M]$, which could be represented in $i\ell$ bits, it suffices to use $i\ell$ -bit-input IO. Furthermore, for “compression” to the possible, all M master public and secret keys, as well as all M ciphertexts, need to be sampled using pseudo-randomness generated by puncturable PRFs. The resulting obfuscated circuits have polynomial size, since generating individual master public keys and ciphertexts using pseudorandomness is efficient, and hence the new FE scheme becomes efficient. Finally, the security of the new FE scheme follows

Collusion Resistant FE Scheme CRFE for \mathcal{C}

SETUP: CRFE.Setup(1^λ) does:

- Sample a PPRF key $K^s \xleftarrow{\$} \text{PRF.Gen}(1^\lambda)$.
- Obfuscate the program $P_{\text{setup}}[0, K^s, \perp]$ described in Figure 6

$$\hat{P}_{\text{setup}} \xleftarrow{\$} \text{iO}(1^\kappa, P_{\text{setup}}[0, K^s, \perp, \perp]),$$

where the IO scheme is invoked with a security parameter $\kappa = \max(\lambda, |P_{\text{setup}}|)$. (As shown in Claim 3 below, $|P_{\text{setup}}| = \text{poly}(\lambda)$.)

- Output MPK = \hat{P}_{setup} and MSK = K^s .

ENCRYPTION: CRFE.Enc(MPK = \hat{P}_{setup} , x) does the following to encrypt an input $x \in \{0, 1\}^N$:

- Sample a PPRF key $K^e \xleftarrow{\$} \text{PRF.Gen}(1^\lambda)$.
- Obfuscate the program $P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e, x, \perp, \perp]$ described in Figure 7,

$$\text{CT} = \hat{P}_{\text{enc}} \xleftarrow{\$} \text{iO}(1^{\kappa'}, P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e, x, \perp, \perp]),$$

where the IO scheme is invoked with a security parameter $\kappa' = \max(\lambda, |P_{\text{enc}}|)$. (As shown in Claim 3 below, $|P_{\text{enc}}| = \text{poly}(\lambda, N)$.)

- Output the obfuscated circuit as the ciphertext CT = \hat{P}_{enc} .

KEY GENERATION: CRFE.KeyGen(MSK = K^s , f) a key for function $f \in \mathcal{C}$ as follows:

- Sample at random an index $i_f \xleftarrow{\$} [M]$.
- Generate a secret key of f under the i_f^{th} master secret key,

$$\begin{aligned} (\text{OMPK}_{i_f}, \text{OMSK}_{i_f}) &= \text{OFE.Setup}(1^\lambda; F(K^s, i_f)), \\ \text{OSK}_{i_f} &\xleftarrow{\$} \text{OFE.KeyGen}(\text{OMSK}_{i_f}, f). \end{aligned}$$

- Output SK = (i_f, OSK_{i_f}) .

DECRYPTION: CRFE.Dec(SK = (i_f, OSK_{i_f}) , CT = \hat{P}_{enc}) does:

- Compute the ciphertext of x under the i_f^{th} master public key,

$$\text{OCT}_{i_f} = \hat{P}_{\text{enc}}(i_f).$$

- Decrypt the obtained ciphertext using OSK_{i_f} ,

$$y = \text{OFE.Dec}(\text{OSK}_{i_f}, \text{OCT}_{i_f}).$$

- Output y .

Figure 5: Collusion Resistant FE Scheme CRFE for \mathcal{C} from $i\ell(\lambda) = \omega(\lambda)$ -bit-input IO

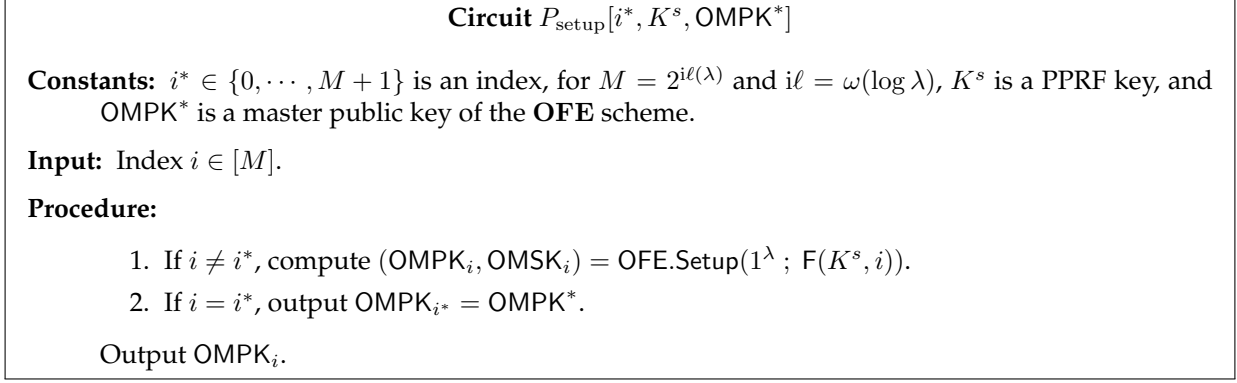


Figure 6: Circuit P_{setup} in the construction and analysis of **CRFE**

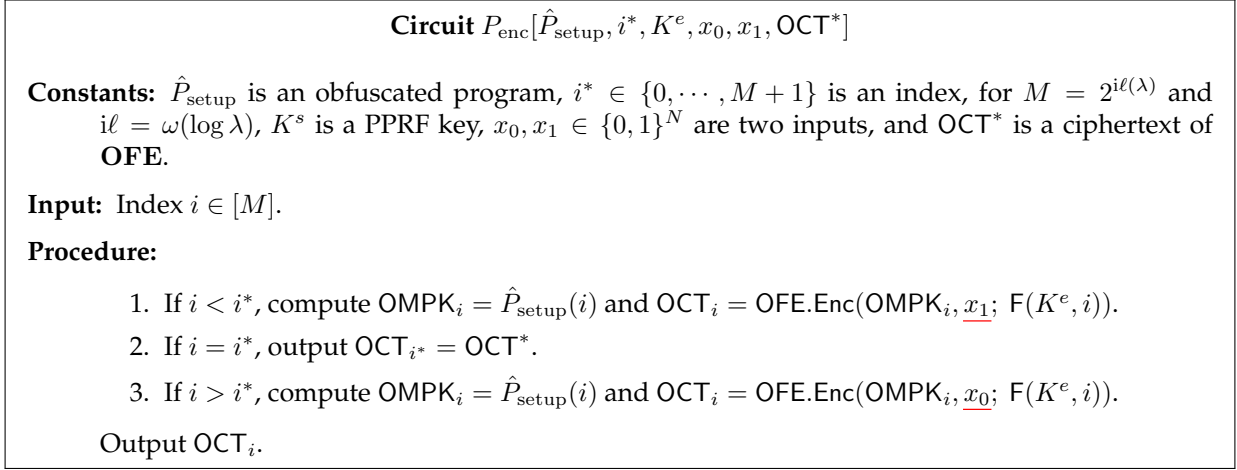


Figure 7: Circuit P_{enc} in the construction and analysis of **CRFE**

from the common “one-input-at-a-time” argument, which incurs a $2^{-|i|} = 2^{-i\ell}$ security loss. We formally describe the collusion-resistant FE scheme **CRFE** for \mathcal{C} in Figure 5.

Next, we proceed to analyzing the correctness, efficiency, and security of the **CRFE** schemes.

Claim 2. *The scheme **CRFE** in Figure 5 is correct.*

Proof. The correctness of **CRFE** follows from the correctness of the underlying FE scheme **OFE** and the IO scheme $i\mathcal{O}$. Fix any λ , any function $f \in \mathcal{C}$, and any input x of f . Consider executing the algorithms of **CRFE** with f and x .

- *Setup:* An honestly generated master public key of **CRFE** is an obfuscated circuit of the program P_{setup} . By construction in Figure 6, P_{setup} on input any $i \in [M]$ outputs an honestly generated master public key OMPK_i of the underlying **OFE** scheme. Let $\text{MPK} = \hat{P}_{\text{setup}}$. Then, by the correctness of the IO scheme, we have that $\hat{P}_{\text{setup}}(i)$ produces the master secret key OMPK_i .
- *Encryption:* A **CRFE** ciphertext CT of x is another obfuscated circuit of the program P_{enc} . By construction in Figure 7, P_{enc} on input any $i \in [M]$ outputs an honestly generated **OFE** ciphertext OCT_i of x under the key output by $\text{MPK}(i) = \hat{P}_{\text{setup}}(i)$, which is OMPK_i as argued above. Let $\text{CT} = \hat{P}_{\text{enc}}$. Then, by the correctness of the IO scheme, $\hat{P}_{\text{enc}}(i)$ produces the ciphertext OCT_i .

- *Key Generation:* A **CRFE** secret key SK of f contains a randomly chosen index $i_f \xleftarrow{\$} [M]$, and an **OFE** secret key OSK_{i_f} for f under the i_f^{th} master secret key MSK_{i_f} .
- *Decryption:* When decrypting CT using SK, the decryptor first evaluates $\hat{P}_{\text{enc}}(i_f)$ to obtain OCT_{i_f} as argued above. Next, the decryptor decrypts OCT_{i_f} using the secret key OSK_{i_f} contained in SK. By the correctness of the **OFE** scheme, the output is $f(x)$.

Therefore the scheme **CRFE** is correct. \square

Next, we show that algorithms of the new scheme **CRFE** is only polynomially slower than that of **OFE**.

Claim 3. *There exists a universal polynomial p (independent of \mathcal{C}), such that, algorithms of **CRFE** in Figure 5 run in time,*

$$\begin{aligned}
t_{\text{CRFE.Setup}}(\lambda, N, D, S) &\leq p(\lambda, t_{\text{OFE.Setup}}(\lambda, N, D, S)) \\
t_{\text{CRFE.KeyGen}}(\lambda, N, D, S) &\leq p(\lambda, t_{\text{OFE.KeyGen}}(\lambda, N, D, S)) \\
t_{\text{CRFE.Enc}}(\lambda, N, D, S) &\leq p(\lambda, t_{\text{OFE.Setup}}(\lambda, N, D, S), t_{\text{OFE.Enc}}(\lambda, N, D, S)) \\
t_{\text{CRFE.Dec}}(\lambda, N, D, S) &\leq p(\lambda, t_{\text{OFE.Setup}}(\lambda, N, D, S), t_{\text{OFE.Enc}}(\lambda, N, D, S), t_{\text{OFE.Dec}}(\lambda, N, D, S)),
\end{aligned}$$

where $t_{\star}(\lambda, N, S)$ denotes the runtime of algorithm \star , and $N = N(\lambda)$, $D = D(\lambda)$, $S = S(\lambda)$ are polynomial upper-bounds on the input-length, depth, and size of circuits that compute functions in \mathcal{C}_{λ} .

In particular, according the above claim, if the original FE scheme **OFE** is *compact*: There exists a polynomial q , such that, $t_{\text{OFE.Setup}}(\lambda, N, D, S) \leq q(\lambda)$, $t_{\text{OFE.Enc}}(\lambda, N, D, S) \leq q(\lambda, N)$, $t_{\text{OFE.KeyGen}}(\lambda, N, D, S) \leq q(\lambda, S)$, and $t_{\text{OFE.Dec}}(\lambda, N, D, S) \leq q(\lambda, S)$. Then the new FE scheme **CRFE** is also *compact* w.r.t. a different polynomial q' .

Proof. We first analyze the size of the programs P_{setup} and P_{enc} , which are obfuscated in **CRFE**. Observe that $P_{\text{setup}}(i)$ basically either invokes the setup algorithm of **OFE** with pseudo-randomness, or outputs a hardwired value. By the efficiency of the PPRF, we have that

$$|P_{\text{setup}}| = \text{poly}(t_{\text{OFE.Setup}}(\lambda, N, D, S)).$$

Furthermore, by the efficiency of the IO scheme, an obfuscation \hat{P}_{setup} of P_{setup} has size

$$\begin{aligned}
|\hat{P}_{\text{setup}}| &\leq \text{Time}_{i\mathcal{O}}(1^{\kappa}, P_{\text{setup}}) \\
&= \text{poly}(\kappa, |P_{\text{setup}}|) \\
&= \text{poly}(\max(\lambda, |P_{\text{setup}}|), |P_{\text{setup}}|) \\
&= \text{poly}(t_{\text{OFE.Setup}}(\lambda, N, D, S)).
\end{aligned}$$

Next, observe that $P_{\text{enc}}(i)$ in figure 6 basically either evaluates $\text{OMPK}_i = \hat{P}_{\text{setup}}(i)$ and generates a ciphertext of the input x under OMPK_i , or outputs a hardwired ciphertext. By the above analysis of the size of \hat{P}_{setup} , and the efficiency of PPRF, we have that

$$|P_{\text{enc}}| = \text{poly}(t_{\text{OFE.Setup}}(\lambda, N, D, S), t_{\text{OFE.Enc}}(\lambda, N, D, S)).$$

Next, we analyze the efficiency of the algorithms of **CRFE**.

- *Setup*: The runtime of the setup algorithm is dominated by the step of obfuscating the program P_{setup} , that is,

$$\begin{aligned}
t_{\text{CRFE.Setup}}(\lambda, N, D, S) &= \text{Time}_{i\mathcal{O}}(1^\kappa, P_{\text{setup}}) + \text{poly}(\lambda) \\
&= \text{poly}(\kappa, |P_{\text{setup}}|) + \text{poly}(\lambda) \\
&= \text{poly}(\max(\lambda, |P_{\text{setup}}|), |P_{\text{setup}}|) + \text{poly}(\lambda) \\
&\leq p(\lambda, t_{\text{OFE.Setup}}(\lambda, N, D, S)),
\end{aligned}$$

where the second equality follows from the efficiency of the IO scheme (and the last inequality holds when p is sufficiently large).

- *Encryption*: The runtime of the encryption algorithm is dominated by obfuscating the program P_{enc} , that is,

$$\begin{aligned}
t_{\text{CRFE.Enc}}(\text{MPK}, x) &= \text{Time}_{i\mathcal{O}}(1^{\kappa'}, P_{\text{enc}}) + \text{poly}(\lambda) \\
&= \text{poly}(\kappa', |P_{\text{enc}}|) + \text{poly}(\lambda) \\
&= \text{poly}(\max(\lambda, |P_{\text{enc}}|), |P_{\text{enc}}|) + \text{poly}(\lambda) \\
&\leq p(\lambda, \lambda, t_{\text{OFE.Setup}}(\lambda), t_{\text{OFE.Enc}}(\lambda, N, S)).
\end{aligned}$$

- *Key Generation*: The runtime of the key generation algorithm is dominated by generating a secret key of the **OFE** scheme. More precisely,

$$\begin{aligned}
t_{\text{CRFE.KeyGen}}(\text{MSK}, f) &= t_{\text{OFE.KeyGen}}(\lambda, N, S) + \text{poly}(\lambda) \\
&\leq p(\lambda, t_{\text{OFE.KeyGen}}(\lambda, N, S)).
\end{aligned}$$

- *Decryption*: The decryption algorithm involves evaluating the obfuscated circuit \hat{P}_{enc} contained in the ciphertext CT on input i_f , and decrypting the obtained **OFE** ciphertext OCT_{i_f} using the **OFE** secret key OSK_{i_f} contained in the secret key SK. Therefore,

$$\begin{aligned}
t_{\text{CRFE.Dec}}(\text{CT}, \text{SK}) &= |\hat{P}_{\text{enc}}| + t_{\text{OFE.Dec}}(\lambda, N, S) + \text{poly}(\lambda) \\
&\leq p(\lambda, t_{\text{OFE.Setup}}(\lambda), t_{\text{OFE.Enc}}(\lambda, N, S), t_{\text{OFE.Dec}}(\lambda, N, S))
\end{aligned}$$

This concludes the claim. □

Lemma 9. *If $i\mathcal{O}$ and PPRF are $2^{-(i\ell(\lambda)+\tau(\lambda))} \text{negl}(\lambda)$ -indistinguishable, and **OFE** is $2^{-(i\ell(\lambda)+\tau(\lambda))} \text{negl}(\lambda)$ -(Adap or Full-Sel)-secure, then, **CRFE** in Figure 5 is $2^{-\tau(\lambda)} \text{negl}(\lambda)$ -(Adap or Full-Sel)-secure.*

Proof. We prove the theorem for the case of Adap-security; the proof for the case of Full-Sel-security are syntactically identical.

Fix any PPT attacker A , we need to show that the advantage of A in games $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$ and $\text{IND}_A^{\text{CRFE}}(1^\lambda, 1)$ is bounded by $2^{-\tau} \text{negl}$.

$$\text{Adv}_A^{\text{CRFE}} = \left| \Pr[\text{IND}_A^{\text{CRFE}}(1^\lambda, 0) = 1] - \Pr[\text{IND}_A^{\text{CRFE}}(1^\lambda, 1) = 1] \right| \leq 2^{-\tau(\lambda)} \text{negl}(\lambda)$$

Recall that the game $\text{IND}_A^{\text{CRFE}}(1^\lambda, b)$ proceeds in four stages: 1) The challenger samples a pair of master keys (MPK, MSK) $\xleftarrow{\$}$ $\text{CRFE.Setup}(1^\lambda)$ and sends MPK to A ; 2) A can obtain an arbitrary number of secret keys $\{\text{SK}_i\}$ for functions $\{f_i\}$ it chooses adaptively; 3) A chooses two challenge

messages (x_0, x_1) and receives the ciphertext CT of x_b ; 4) A again obtains secret keys of functions of its choice.

To bound the advantage of A , we define hybrids $\{H_{i^*}^0, \dots, H_{i^*}^3\}_{i^* \in [M+1]}$, and show that the advantage of A in distinguishing $H_{i^*}^j$ from $H_{i^*}^{j+1}$ for any $0 \leq j \leq 2$, as well as in distinguishing $H_{i^*}^3$ and $H_{i^*+1}^0$, is bounded by $\mu(\lambda) = 2^{-(i\ell(\lambda)+\tau(\lambda))} \text{negl}(\lambda)$. Furthermore, the advantage of A in distinguishing H_0^0 from $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$, and H_{M+1}^0 from $\text{IND}_A^{\text{CRFE}}(1^\lambda, 1)$ is also bounded by μ . Next, we formally describe the hybrids.

Hybrid $H_{i^*}^0$: This hybrid proceeds identically to $\text{IND}_A^{\text{CRFE}}(1^\lambda, b)$, except that, the master public key MPK and ciphertext CT are generated differently. Recall that by construction of **CRFE**, the master public key and ciphertext are obfuscated circuits of programs P_{setup} and P_{enc} respectively.

$$\begin{aligned} \text{MPK} &= \hat{P}_{\text{setup}} \stackrel{\$}{\leftarrow} i\mathcal{O}(1^\kappa, P_{\text{setup}}) \\ \text{CT} &= \hat{P}_{\text{enc}} \stackrel{\$}{\leftarrow} i\mathcal{O}(1^{\kappa'}, P_{\text{enc}}) \end{aligned}$$

In $\text{IND}_A^{\text{CRFE}}(1^\lambda, b)$, the two programs are $P_{\text{setup}} = P_{\text{setup}}[0, K^s, \perp]$ and $P_{\text{enc}} = P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e, x_b, \perp, \perp]$, where \hat{P}_{setup} is the obfuscation of P_{setup} .

In $H_{i^*}^0$, the master public key and ciphertext contain obfuscation of P_{setup} and P_{enc} hardwired with different constants as described below (recall that $K\{i\}$ denotes a PPRF key punctured at point i).

$$\begin{aligned} P_{\text{setup}} &= P_{\text{setup}}[i^*, K^s\{i^*\}, \text{OMPK}_{i^*}], \\ &\text{where } (\text{OMPK}_{i^*}, \text{OMSK}_{i^*}) \stackrel{\$}{\leftarrow} \text{OFE.Setup}(1^\lambda; F(K^s, i^*)), \text{ and} \end{aligned} \quad (12)$$

$$\begin{aligned} P_{\text{enc}} &= P_{\text{enc}}[\hat{P}_{\text{setup}}, i^*, K^e\{i^*\}, x_0, x_1, \text{OCT}_{i^*}], \\ &\text{where } \text{OCT}_{i^*} \stackrel{\$}{\leftarrow} \text{OFE.Enc}(\text{OMPK}_{i^*}, x_0; F(K^e, i^*)). \end{aligned} \quad (13)$$

The rest of the experiment proceeds identically to $\text{IND}_A^{\text{CRFE}}(1^\lambda, b)$. Finally, $H_{i^*}^0$ outputs the bit that A outputs.

Below, we show that

Claim 4. *If $i\mathcal{O}$ is μ -secure, then the outputs of $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$ and H_0^0 , and the outputs of $\text{IND}_A^{\text{CRFE}}(1^\lambda, 1)$ and H_{M+1}^0 , are μ -close.*

Proof. By definition of $P_{\text{setup}}[\star]$ and $P_{\text{enc}}[\star]$ in Figure 6 and 7 and the correctness of PPRF, we have that the programs obfuscated in $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$ are functionally equivalent to that obfuscated in H_0^0 , that is,

$$\begin{aligned} P_{\text{setup}}[0, K^s, \perp] &\equiv P_{\text{setup}}[0, K^s\{0\}, \text{OMPK}_0] \\ P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e, x_0, \perp, \perp] &\equiv P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e\{0\}, x_0, x_1, \text{OCT}_0] \end{aligned}$$

where OMSK_0 and OCT_0 are generated as in Equation (12) and (13). Therefore, if $i\mathcal{O}$ is μ -indistinguishable, the outputs of $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$ and H_0^0 are μ -close.

Similarly, the outputs of $\text{IND}_A^{\text{CRFE}}(1^\lambda, 1)$ and H_{M+1}^0 are also μ -close, following from the fact that the programs obfuscated in them are functionally equivalent.

$$\begin{aligned} P_{\text{setup}}[0, K^s, \perp] &\equiv P_{\text{setup}}[M+1, K^s\{M+1\}, \text{OMPK}_{M+1}] \\ P_{\text{enc}}[\hat{P}_{\text{setup}}, 0, K^e, x_1, \perp, \perp] &\equiv P_{\text{enc}}[\hat{P}_{\text{setup}}, M+1, K^e\{M+1\}, x_0, x_1, \text{OCT}_{M+1}] \end{aligned}$$

□

Hybrid $H_{i^*}^1$: This hybrid proceed identically to $H_{i^*}^0$ except that the master public key OMPK_{i^*} and ciphertext OCT_{i^*} hardwired in P_{setup} and P_{enc} are generated using true randomness, instead of pseudorandomness generated by PPRF.

$$\begin{aligned} (\text{OMPK}_{i^*}, \text{OMSK}_{i^*}) &\stackrel{\$}{\leftarrow} \text{OFE.Setup}(1^\lambda; \underline{U_{\text{poly}(\lambda)}}), \text{ and} \\ \text{OCT}_{i^*} &\stackrel{\$}{\leftarrow} \text{OFE.Enc}(\text{OMPK}_{i^*}, x_0; \underline{U_{\text{poly}(\lambda)}}). \end{aligned}$$

The rest of the experiment proceeds identically to $H_{i^*}^0$.

Since the only difference between $H_{i^*}^0$ and $H_{i^*}^1$ lies in whether OMPK_{i^*} and OCT_{i^*} are generated by pseudorandomness or true randomness, it follows from the μ -indistinguishability of PPRF that the outputs of $H_{i^*}^0$ and $H_{i^*}^1$ are μ -close.

Claim 5. *If PPRF is $\mu(\lambda)$ -secure, then the outputs of $H_{i^*}^0$ and $H_{i^*}^1$ are $\mu(\lambda)$ -close.*

Hybrid $H_{i^*}^2$: This hybrid proceeds identically to $H_{i^*}^1$ except that the ciphertext OCT_{i^*} hardwired P_{enc} encrypts x_1 , instead of x_0 .

$$\begin{aligned} (\text{OMPK}_{i^*}, \text{OMSK}_{i^*}) &\stackrel{\$}{\leftarrow} \text{OFE.Setup}(1^\lambda; U_{\text{poly}(\lambda)}), \text{ and} \\ \text{OCT}_{i^*} &\stackrel{\$}{\leftarrow} \text{OFE.Enc}(\text{OMPK}_{i^*}, \underline{x_1}; U_{\text{poly}(\lambda)}). \end{aligned}$$

Note that the only difference between $H_{i^*}^1$ and $H_{i^*}^2$ lies in whether OCT_{i^*} encrypts x_0 or x_1 . In both hybrids, the $(i^*)^{\text{th}}$ master keys $(\text{OMPK}_{i^*}, \text{OMSK}_{i^*})$ and OCT_{i^*} are generated honestly using true randomness. By construction of **CRFE**, the secret key SK of a function f contains an **OFE** secret key OSK_{i_f} of f under a randomly chosen master secret key OMPK_{i_f} from a super-polynomial number M of master secret keys $i_f \stackrel{\$}{\leftarrow} [M]$. Since the attacker A obtains at most a polynomial number of secret keys, the probability that OMSK_{i^*} is used to generate two secret key is negligible. Conditioned on this event happening, it follows from the μ -Adap-security of **OFE** that the ciphertext OCT_{i^*} in $H_{i^*}^1$ encrypting x_0 is indistinguishable to that in $H_{i^*}^2$ encrypting x_1 . Therefore,

Claim 6. *If **OFE** is $\mu(\lambda)$ -Adap-secure, then the outputs of $H_{i^*}^1$ and $H_{i^*}^2$ are $\mu(\lambda)$ -close.*

Hybrid $H_{i^*}^3$: This hybrid proceeds identically to $H_{i^*}^2$ except that the master public key OMPK_{i^*} and ciphertext OCT_{i^*} hardwired in P_{setup} and P_{enc} are generated using pseudorandomness generated by PPRF, instead of true randomness.

$$(\text{OMPK}_{i^*}, \text{OMSK}_{i^*}) \stackrel{\$}{\leftarrow} \text{OFE.Setup}(1^\lambda; \underline{\text{PPRF}(K^s, i^*)}), \text{ and} \quad (14)$$

$$\text{OCT}_{i^*} \stackrel{\$}{\leftarrow} \text{OFE.Enc}(\text{OMPK}_{i^*}, x_1; \underline{\text{PPRF}(K^e, i^*)}). \quad (15)$$

It again follows from the μ -indistinguishability of PPRF that

Claim 7. *If PPRF is $\mu(\lambda)$ -indistinguishable, then the outputs of $H_{i^*}^2$ and $H_{i^*}^3$ are $\mu(\lambda)$ -close.*

Furthermore, we note that for every $i^* \in [M]$, the programs P_{setup} and P_{enc} obfuscated in hybrids $H_{i^*}^3$ and $H_{i^*+1}^0$ are functionally equivalent.

$$\begin{aligned} P_{\text{setup}}[i^*, K^s\{i^*\}, \text{OMPK}_{i^*}] &\equiv P_{\text{setup}}[i^*+1, K^s\{i^*+1\}, \text{OMPK}_{i^*+1}] \\ P_{\text{enc}}[\hat{P}_{\text{setup}}, i^*, K^e\{i^*\}, x_0, x_1, \text{OCT}_{i^*}] &\equiv P_{\text{enc}}[\hat{P}_{\text{setup}}, i^*+1, K^e\{i^*+1\}, x_0, x_1, \text{OCT}_{i^*+1}] \end{aligned}$$

where on the left hand side of the equations, the hardwired master public key OMPK_{i^*} and ciphertext OCT_{i^*} are generated as in Equation (14) and (15), while on the right hand side, the mater public key OMPK_{i^*+1} and ciphertext OCT_{i^*+1} are generated as in Equation (12) and (13). By definition of $P_{\text{setup}}[\star]$, the P_{setup} programs with different hardwired constants described above are functionally equivalent because for any input $i \in [M]$, they both output OMPK_i generated honestly using OFE.Setup with pseudorandomness $F(K^s, i)$. By definition of $P_{\text{enc}}[\star]$, the P_{enc} programs above are also functionally equivalent because for any input $i \in [M]$ and $i \leq i^*$, they both produce OCT_i encrypting x_1 using pseudorandomness $F(K^e, i)$, and for any $i \in [M]$ and $i > i^*$, they both produce OCT_i encrypting x_0 using pseudorandomness $F(K^e, i)$. Therefore, it follows from the μ -indistinguishability of $i\mathcal{O}$ that the outputs of $H_{i^*}^3$ and $H_{i^*+1}^0$ are μ -close.

Claim 8. *If $i\mathcal{O}$ is μ -secure, then the outputs of $H_{i^*}^3$ and $H_{i^*+1}^0$ are μ -close.*

Using the above hybrids, we now conclude the lemma. Since there are in total $O(M)$ hybrids, it then follows from a hybrid argument that the advantage of A in distinguishing $\text{IND}_A^{\text{CRFE}}(1^\lambda, 0)$ and $\text{IND}_A^{\text{CRFE}}(1^\lambda, 1)$ is bounded by $O(M) \times \mu = 2^{-\tau(\lambda)} \text{negl}(\lambda)$. \square

Acknowledgements

The authors thank Benny Applebaum and Vinod Vaikuntanathan for many helpful discussions and insights.

Huijia Lin was supported in part by NSF grants CNS-1528178, CNS-1514526, and CNS-1652849 (CAREER). Stefano Tessaro was supported in part by NSF grants CNS-1423566, CNS-1528178, CNS-1553758 (CAREER), and IIS-152804.

References

- [AB15] Benny Applebaum and Zvika Brakerski. Obfuscating circuits via composite-order graded encoding. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 528–556, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany.
- [ABCP15] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. Simple functional encryption schemes for inner products. In Jonathan Katz, editor, *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, volume 9020 of *Lecture Notes in Computer Science*, pages 733–751. Springer, 2015.
- [ABR16] Benny Applebaum, Andrej Bogdanov, and Alon Rosen. A dichotomy for local small-bias generators. *Journal of Cryptology*, 29(3):577–596, July 2016.
- [ABSV15] Prabhanjan Ananth, Zvika Brakerski, Gil Segev, and Vinod Vaikuntanathan. From selective to adaptive security in functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 657–677, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.

- [ADGM17] Daniel Apon, Nico Döttling, Sanjam Garg, and Pratyay Mukherjee. Cryptanalysis of indistinguishability obfuscations of circuits over GGH13. In *ICALP 2017*, LNCS, 2017.
- [AGIS14] Prabhanjan Vijendra Ananth, Divya Gupta, Yuval Ishai, and Amit Sahai. Optimizing obfuscation: Avoiding Barrington’s theorem. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 14*, pages 646–658, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.
- [AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in nc^0 . In *FOCS*, pages 166–175, 2004.
- [AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Computationally private randomizing polynomials and their applications. *Computational Complexity*, 15(2):115–162, 2006.
- [AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of LNCS, pages 308–326, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [AJS15] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. *IACR Cryptology ePrint Archive*, 2015:730, 2015.
- [AL16] Benny Applebaum and Shachar Lovett. Algebraic attacks against random local functions and their countermeasures. In Daniel Wichs and Yishay Mansour, editors, *48th ACM STOC*, pages 1087–1100, Cambridge, MA, USA, June 18–21, 2016. ACM Press.
- [AOW15] Sarah R. Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In Venkatesan Guruswami, editor, *56th FOCS*, pages 689–708, Berkeley, CA, USA, October 17–20, 2015. IEEE Computer Society Press.
- [App12] Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 805–816, New York, NY, USA, May 19–22, 2012. ACM Press.
- [App15] Benny Applebaum. The cryptographic hardness of random local functions – survey. *Cryptology ePrint Archive*, Report 2015/165, 2015. <http://eprint.iacr.org/2015/165>.
- [AR16] Benny Applebaum and Pavel Raykov. Fast pseudorandom functions based on expander graphs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part I*, volume 9985 of LNCS, pages 27–56, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.
- [AS16] Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. *Cryptology ePrint Archive*, Report 2016/1097, 2016. <http://eprint.iacr.org/2016/1097>.

- [BBKK17] Boaz Barak, Zvika Brakerski, Ilan Komargodski, and Pravesh K. Kothari. Limits on low-degree pseudorandom generators (or: Sum-of-squares meets program obfuscation). Cryptology ePrint Archive, Report 2017/312, 2017. <http://eprint.iacr.org/2017/312>.
- [BGI⁺01a] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *Advances in Cryptology CRYPTO 2001*, pages 1–18. Springer, 2001.
- [BGI⁺01b] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer, Heidelberg, Germany.
- [BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In *PKC*, pages 501–519, 2014.
- [BGK⁺14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 221–238, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- [BGL⁺15] Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. Succinct randomized encodings and their applications. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 439–448. ACM, 2015.
- [BJK15] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. Function-hiding inner product encryption. In Tetsu Iwata and Jung Hee Cheon, editors, *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 470–491. Springer, 2015.
- [BNPW16] Nir Bitansky, Ryo Nishimaki, Alain Passelègue, and Daniel Wichs. From cryptomania to obfustopia through secret-key functional encryption. In Martin Hirt and Adam D. Smith, editors, *Theory of Cryptography - 14th International Conference, TCC 2016-B, Beijing, China, October 31 - November 3, 2016, Proceedings, Part II*, volume 9986 of *Lecture Notes in Computer Science*, pages 391–418, 2016.
- [BPR15] Nir Bitansky, Omer Paneth, and Alon Rosen. On the cryptographic hardness of finding a nash equilibrium. In Guruswami [Gur15], pages 1480–1498.
- [BQ12] Andrej Bogdanov and Youming Qiao. On the security of goldreich’s one-way function. *Computational Complexity*, 21(1):83–127, 2012.
- [BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 1–25, San Diego, CA, USA, February 24–26, 2014. Springer, Heidelberg, Germany.

- [BS02] Dan Boneh and Alice Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324:71–90, 2002.
- [BSW06] Dan Boneh, Amit Sahai, and Brent Waters. Fully collusion resistant traitor tracing with short ciphertexts and private keys. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 573–592, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.
- [BSW12] Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: a new vision for public-key cryptography. *Commun. ACM*, 55(11):56–64, 2012.
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 171–190, 2015.
- [BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *ASIACRYPT (2)*, pages 280–300, 2013.
- [CEMT09] James Cook, Omid Etesami, Rachel Miller, and Luca Trevisan. Goldreich’s one-way function candidate and myopic backtracking algorithms. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 521–538. Springer, Heidelberg, Germany, March 15–17, 2009.
- [CFN94] Benny Chor, Amos Fiat, and Moni Naor. Tracing traitors. In Yvo Desmedt, editor, *CRYPTO’94*, volume 839 of *LNCS*, pages 257–270, Santa Barbara, CA, USA, August 21–25, 1994. Springer, Heidelberg, Germany.
- [CGH⁺15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancrede Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 247–266, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [CGH17] Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part III*, volume 10212 of *Lecture Notes in Computer Science*, pages 278–307, 2017.
- [CHL⁺15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 3–12, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- [CIJ⁺13] Angelo De Caro, Vincenzo Iovino, Abhishek Jain, Adam O’Neill, Omer Paneth, and Giuseppe Persiano. On the achievability of simulation-based security for functional encryption. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II*, pages 519–535, 2013.

- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 476–493, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.
- [CLT15] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 267–286, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.
- [CLTV15] Ran Canetti, Huijia Lin, Stefano Tessaro, and Vinod Vaikuntanathan. Obfuscation of probabilistic circuits and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 468–497. Springer, 2015.
- [CM01] M. Cryan and P. B. Miltersen. On pseudorandom generators in nc_0 . In Proc. 26th MFCS, 2001.
- [DGG⁺16] Nico Döttling, Sanjam Garg, Divya Gupta, Peihan Miao, and Pratyay Mukherjee. Obfuscation from low noise multilinear maps. Cryptology ePrint Archive, Report 2016/599, 2016. <http://eprint.iacr.org/2016/599>.
- [DIJK09] Yevgeniy Dodis, Russell Impagliazzo, Ragesh Jaiswal, and Valentine Kabanets. Security amplification for interactive cryptographic primitives. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 128–145. Springer, Heidelberg, Germany, March 15–17, 2009.
- [DNR⁺09] Cynthia Dwork, Moni Naor, Omer Reingold, Guy N. Rothblum, and Salil P. Vadhan. On the complexity of differentially private data release: efficient algorithms and hardness results. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 381–390, Bethesda, MD, USA, May 31 – June 2, 2009. ACM Press.
- [DS05] Yevgeniy Dodis and Adam Smith. Correcting errors without leaking partial information. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 654–663, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013.
- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 40–49. IEEE Computer Society, 2013.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*,

- volume 9015 of *LNCS*, pages 498–527, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany.
- [GGHZ16] Sanjam Garg, Craig Gentry, Shai Halevi, and Mark Zhandry. Functional encryption without obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 480–511, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.
- [GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
- [GKP⁺13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 555–564, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.
- [GLSW15] Craig Gentry, Allison Bishop Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In Guruswami [Gur15], pages 151–170.
- [GMM⁺16] Sanjam Garg, Eric Miles, Pratyay Mukherjee, Amit Sahai, Akshayaram Srinivasan, and Mark Zhandry. Secure obfuscation in a weak multilinear map model. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 241–268, Beijing, China, October 31 – November 3, 2016. Springer, Heidelberg, Germany.
- [GNW11] Oded Goldreich, Noam Nisan, and Avi Wigderson. *On Yao’s XOR-Lemma*, pages 273–301. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
- [Gol00] Oded Goldreich. Candidate one-way functions based on expander graphs. *Electronic Colloquium on Computational Complexity (ECCC)*, 7(90), 2000.
- [Gol01] Oded Goldreich. *Foundations of Cryptography — Basic Tools*. Cambridge University Press, 2001.
- [GPS16] Sanjam Garg, Omkant Pandey, and Akshayaram Srinivasan. Revisiting the cryptographic hardness of finding a nash equilibrium. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 579–604, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.
- [GPSZ16] Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. Breaking the sub-exponential barrier in obfustopia. Cryptology ePrint Archive, Report 2016/102, 2016. <http://eprint.iacr.org/2016/102>.
- [Gur15] Venkatesan Guruswami, editor. *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*. IEEE Computer Society, 2015.
- [GVW12] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption with bounded collusions via multi-party computation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 162–179, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

- [IK02] Yuval Ishai and Eyal Kushilevitz. Perfect constant-round secure computation via perfect randomizing polynomials. In *ICALP*, pages 244–256, 2002.
- [KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *CCS*, pages 669–684, 2013.
- [KS17] Ilan Komargodski and Gil Segev. From minicrypt to obfustopia via private-key functional encryption. Cryptology ePrint Archive, Report 2017/080, 2017. <http://eprint.iacr.org/2017/080>.
- [Lin16] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 28–57, Vienna, Austria, May 8–12, 2016. Springer, Heidelberg, Germany.
- [Lin17] Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 PRGs. In *CRYPTO 2017*, *LNCS*, 2017.
- [LPST16] Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Output-compressing randomized encodings and applications. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 96–124, Tel Aviv, Israel, January 10–13, 2016. Springer, Heidelberg, Germany.
- [LSS14] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. Gghlite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256. Springer, 2014.
- [LV16] Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from ddh-like assumptions on constant-degree graded encodings. In *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, New Brunswick, NJ, USA, 9-11 October, 2016*, 2016.
- [LV17] Alex Lombardi and Vinod Vaikuntanathan. On the non-existence of blockwise 2-local prgs with applications to indistinguishability obfuscation. Cryptology ePrint Archive, Report 2017/301, 2017. <http://eprint.iacr.org/2017/301>.
- [MST03] Elchanan Mossel, Amir Shpilka, and Luca Trevisan. On e-biased generators in NC0. In *44th FOCS*, pages 136–145, Cambridge, MA, USA, October 11–14, 2003. IEEE Computer Society Press.
- [MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.

- [MT10] Ueli M. Maurer and Stefano Tessaro. A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak PRGs with optimal stretch. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 237–254, Zurich, Switzerland, February 9–11, 2010. Springer, Heidelberg, Germany.
- [O’N10] Adam O’Neill. Definitional issues in functional encryption. *Cryptology ePrint Archive*, Report 2010/556, 2010. <http://eprint.iacr.org/>.
- [OW14] Ryan O’Donnell and David Witmer. Goldreich’s PRG: evidence for near-optimal polynomial stretch. In *IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014*, pages 1–12, 2014.
- [PRV12] Bryan Parno, Mariana Raykova, and Vinod Vaikuntanathan. How to delegate and verify in public: Verifiable computation from attribute-based encryption. In Ronald Cramer, editor, *TCC 2012*, volume 7194 of *LNCS*, pages 422–439, Taormina, Sicily, Italy, March 19–21, 2012. Springer, Heidelberg, Germany.
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 500–517, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: Deniable encryption, and more. *Proc. of STOC 2014*, 2014.
- [Ull13] Jonathan Ullman. Answering $n_{2+o(1)}$ counting queries with differential privacy is hard. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 361–370, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.
- [Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press.
- [Zim15] Joe Zimmerman. How to obfuscate programs directly. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 439–467, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.