Universal Computational Extractors and Multi-Bit AIPO from Lattice Assumptions

Yilei Chen * Xinyu Mao ⁺

February 26, 2025

Abstract

We put forth a new primitive called *obliviously programmable function* (*OPF*) to construct two random-oracle-like primitives:

- Universal computational extractors (UCEs), introduced by Bellare, Hoang, and Keelveedhi [BHK13], can securely replace random oracles in various applications, including KDMsecure encryption, deterministic encryption, RSA-OAEP, universal hardcore bits, etc.
- Multi-bit point obfuscation with auxiliary input (MB-AIPO). It enables upgrading CPAsecure public-key encryption (PKE) into a CCA-secure one [MH14] and serves as a tool to instantiate the random oracles used in the Fujisaki-Okamoto transform for lossy PKEs [MOZ23].

Despite their usefulness, constructing UCEs and MB-AIPO in the standard model is challenging. The existing constructions of both primitives [BM14a, BM14b] use indistinguishability obfuscation (iO) plus point functions with auxiliary input (AIPO).

OPF can replace the use iO in the constructions of UCE and MB-AIPO. We use OPF plus AIPO to construct

- UCE with one query for strongly unpredictable sources,
- MB-AIPO for strongly unpredictable distributions and
- PKE scheme that is IND-CPA secure in the presence of computationally uninvertible leakage on the secret key.

We then construct OPF for NC¹ circuits from lattice assumptions based on the GGH15 encodings [GGH15], without using iO. In sum, we give new constructions of the above three primitives under the following assumptions: (1) LWE with subexponential hardness; (2) private-coin evasive LWE assumption for specific samplers; (3) the existence of AIPO in NC¹. As a byproduct, we construct an 'NC¹-universal AIPO' under the assumptions (1) and (2).

1 Introduction

The Random Oracle Methodology refers to the popular paradigm of designing cryptographic schemes that consists of two steps: One first designs a scheme whose security can be proven in

^{*}Tsinghua University and Shanghai Qi Zhi Institute. Supported by Tsinghua University start-up funding and Shanghai Qi Zhi Institute Innovation Program SQZ202405. Email: chenyilei@mail.tsinghua.edu.cn

[†]University of Southern California. Email: xinyumao@usc.edu

the random oracle (RO) model; then, the random oracle is instantiated by a good 'cryptographic hash function' (e.g., SHA-3), hoping the resulting scheme is still safe. Well-known applications of the RO methodology include the Fiat-Shamir transform [FS87] and the Fujisaki-Oakamoto transform [FO99]. However, the RO methodology is only a rule of thumb and has been proven to be theoretically unsound: In a seminal work, Canetti et al. [CGH04] designed a scheme that is secure in the random oracle model but insecure when the random oracle is replaced by *any* function.

Even with these negative results, the random oracle methodology remains popular as people deem the known counterexamples as artificially contrived. The hope is that in natural and practical cases, the random oracle can be safely instantiated. A natural remedy is to identify 'RO-like' properties that are sufficient for important applications and then construct hash functions with such properties under well-formed assumptions. Along this line, a number of security notions have been proposed in existing literature, such as point obfuscation [Can97], correlation intractability [CGH04], correlated-input security [GOR11], and universal computational extractors (UCEs) [BHK13]. In this paper, we focus on the construction of point obfuscation and UCEs.

Point obfuscation. A point function $\mathbf{1}_x$ maps the string x to 1 and all other strings to 0. Obfuscating point functions, or point obfuscation, can be easily achieved in the RO model by outputting the value of the RO at the target point. Bitansky and Paneth [BP12] considered the presence of leakage about the point x, introducing the notion of *point obfuscation with auxiliary input* (AIPO). Matsuda and Hanaoka [MH14] generalized the notion to *multi-bit* point obfuscation with auxiliary input (MB-AIPO) where the function maps the string x to some value m and all other strings to \bot , and showed how to use it to upgrade a CPA-secure public-key encryption into a CCA-secure one.

UCE security. *Universal computational extractors*, introduced by Bellare, Hoang, and Keelveedhi (BHK, [BHK13]), enable instantiations of random oracles in various applications, e.g., universal hardcore function, deterministic encryption, RSA-OAEP, etc. For a keyed function family H, UCE security is defined by a two-stage game as follows. An adversary in this game is a pair of algorithms (S, D), where S is called the *source* and D the *distinguisher*. S has access to an oracle HASH that is either H.Eval(hk, \cdot) or a random oracle; D receives a message L from S and is given hk; D has to guess to which oracle S has access to. The adversary wins if D guesses correctly.

Without any restrictions, we can easily design a winning adversary: S queries HASH on a random point x and send L = (x, y) to D, where y is the oracle answer; then D simply checks whether H.Eval(hk, x) = y. Therefore, we restrict the source to be in some set S, and allow D to be any PPT algorithm. H is said to be S-UCE-secure if for every $S \in S$ and PPT D, the winning probability of (S, D) in the UCE game is negligible.

BHK first considered the set of all computationally unpredictable sources, denoted by S^{cup} . Roughly speaking, $S \in S^{cup}$ if no PPT predictor, given *L*, can predict the queries made by *S* to HASH. However, UCE security for computationally unpredictable sources S^{cup} cannot be achieved in the standard model, assuming indistinguishability obfuscation (iO) exists [BFM14]. Then, Brzuska and Mittelbach (BM, [BM14b]) strengthened the restriction to be *strongly* computationally unpredictable, meaning that the oracle answers to the source are also given to the predictor. Assuming the existence of iO and AIPO, they constructed a UCE-secure function for sources that are (1) strongly computationally unpredictable and (2) only make one query to HASH. We denote the set of such sources as S_1^{scup} . The status of MB-AIPO is similar to that of UCE. BM [BM14a] proved that MB-AIPO for all unpredictable distributions does not exist, assuming iO; they also consider the relaxation to strongly unpredictable distributions (strongly unpredictability roughly says given auxiliary information *and the value* m, it is computationally hard to guess the point x), and give a construction under the same assumptions (iO plus AIPO).

As iO appears to be a much stronger primitive than UCE and MB-AIPO, a natural question arises:

Can we construct UCE-secure functions and MB-AIPO without using iO (in the standard model)?

1.1 Our Results

In this paper, we give a simple construction of S_1^{scup} -UCE-secure function and MB-AIPO (for strongly computationally unpredictable distributions) from lattice assumptions based on the GGH15 encoding [GGH15], yielding other interesting primitives under the same assumptions as well. Our main result is

Theorem 1.1. Under the subexponential LWE assumption and the private-coin evasive LWE assumption for specific samplers, if there exists an AIPO in NC¹, then there exist

- an S_1^{scup} -UCE-secure function with input length ℓ_{in} ,
- an MB-AIPO for strongly unpredictable distributions, and
- a PKE scheme that is IND-CPA secure in the presence of computationally uninvertible leakage on the secret key.

Here, we need the existence of AIPO where the obfuscated program can be computed by a circuit family in NC¹, i.e., there exists $k \in \mathbb{N}$ such that on inputs of length n, the AIPO always outputs a boolean circuit of depth at most $k \log n$. Such AIPO exists assuming some non-standard yet plausible variants of LWE or DDH, which we will explain later.

According to [BHK13], an S_1^{scup} -UCE-secure function family is a universal hardcore function family; that is, it is a hardcore for any one-way function.

Corollary 1.2. Under the same assumption as theorem 1.1, there exists a universal hardcore function family that outputs a polynomial number of bits.

Among previous constructions of universal hardcore function, only the constructions due to Zhandry avoid relying explicitly on iO: One [Zha16] uses extractable witness PRF, and the other [Zha19] uses extremely lossy functions. Extractable witness PRF contains strong knowledge assumptions and the only instantiation of extremely lossy functions is based on the exponential hardness of DDH. Our construction is the first lattice-based one without using iO.

Another interesting byproduct is a 'universal AIPO' construction. This is an 'iO-like' property. Let P_x be a simple circuit computing the point function $\mathbf{1}_x$. Consider padding P_x to size s(|x|) and then obfuscating the padded circuit using iO. This construction is an AIPO assuming the *existence* of an AIPO that always outputs a circuit of size at most s(n) on the input of size n. Our result achieves such an universal AIPO without using iO. **Theorem 1.3** (Universal AIPO for NC¹). Under the subexponential LWE assumption and the private-coin evasive LWE assumption for specific samplers, there exists an explicit family of algorithms $\{AIPO_{k,s}\}_{k\in\mathbb{N},s\in\text{poly}}$ with the following property. (1) If there exists an AIPO such that on inputs of length n, it always outputs a boolean circuit of depth at most $k \log n$ and size at most s(n), then $AIPO_{k,s}$ is an AIPO. (2) $AIPO_{k,s}$ runs in time $poly(n^k, s(n))$ on inputs of length n.

Since subexponential LWE is rather a standard assumption, we discuss the other two assumptions in what follows.

On the evasive LWE assumption. Evasive LWE is a new assumption proposed by Wee [Wee22]. Suppose that we have some joint distributions over matrices **P**, **S** and auxiliary information aux sampled by some PPT sampler Samp. The evasive LWE assumption postulates that, for a uniformly random (and secret) matrix **B**,

if
$$(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathsf{aux}) \approx_c (\mathbf{U}, \mathbf{U}', \mathsf{aux})$$

then $(\mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux}) \approx_c (\mathbf{U}, \mathbf{B}^{-1}(\mathbf{P}), \mathsf{aux})$ (1)

where **U**, **U**' are uniformly random matrices, and **E**, **E**' are chosen from the LWE error distribution with appropriate parameters. Essentially the above says that given SB + E, getting the additional component $B^{-1}(P)$ is no more useful than just getting the product $(SB + E) \cdot B^{-1}(P) \approx SP + E'$. Evasive LWE has proven to be useful in constructing advanced primitives such as witness encryption [Tsa22, VWW22], attribute-based encryption [HLL23], and succinct non-interactive arguments of knowledge [MPV24]. There are significant differences in the formulation of the evasive LWE assumption in different works. Indeed, Brzuska, Ünal, and Woo gave counterexamples for some formulations of evasive LWE [BÜW24], including the one in eq. (1), which they call *private-coin evasive LWE*; 'private-coin' means that the randomness used by the sampler Samp is not included in the distributions in the assumption. In light of these counterexamples, we view evasive LWE as a heuristic and refrain from making general evasive LWE assumptions. Instead, we specify the class of samplers relevant to our security proof and postulate that private-coin evasive LWE holds *for this specific class of samplers*. We note that private-coin evasive LWE for the same class of samplers is used in [VWW22] to construct witness encryption and null-iO. Formal definitions can be found in section 5.2.

On the AIPO assumption. We consider a variant of point obfuscators. Loosely speaking, we require the obfuscation of any point function to be indistinguishable from the obfuscation of the all-zero function. Construction of AIPO in NC¹ with statistically unpredictable auxiliary inputs is known from standard LWE [GKPV10]. We conjecture that the same construction (with minor changes in the parameters) is an AIPO for computationally unpredictable auxiliary inputs.

It is worth noting that the AIPO is only used in security proof; the constructions only involve lattice computations. That is, we only need the existence of an AIPO in NC¹. Therefore, regardless of the status of the AIPO assumption, our construction is still a good candidate for UCE and MB-AIPO in light of its simplicity.

1.2 Technical Overview

We use the construction of UCE as an example to demonstrate our main ideas. Our starting point is the construction of UCE based on iO and AIPO due to BM [BM14b]. Their construction of UCE uses

$$\begin{array}{ccc} y^* = \mathsf{F}\left(msk, x^*\right) & & & \\ \mathsf{hk} = \mathsf{iO}(\mathsf{F}(msk, \cdot)) & & & \mathsf{hk} = \mathsf{iO}(P_{x^*, y^*, k_{x^*}}) & & \\ \mathsf{hk} = \mathsf{iO}(\mathsf{F}(msk, \cdot)) & & & \mathsf{hk} = \mathsf{iO}(\mathsf{F}(msk, \cdot)) & \\ \end{array} \\ \end{array}$$

Figure 1: Three major steps of the proof in [BM14b]. We only focus on the value y^* (given to the source) and hk (given to the distinguisher) in the UCE game. The goal is to change y^* into uniformly random while keeping hk unchanged.

the iO of a puncturable pseudorandom function (PRF); that is, the hash key $hk \leftarrow iO(F(msk, \cdot))$ is an obfuscated program, where F is a puncturable PRF and msk is the master key of F. The security proof essentially uses iO to privately switch the real PRF key to a punctured key that involves the query information from the source. This hints that privately constrained PRFs (PCPRFs) [BLW17] may be good candidates for UCEs, where the hash key hk is simply the constrained key for an empty circuit. Given that there exist a few constructions of PCPRFs from lattice assumptions without using iO (see e.g. [BKM17, CC17, BTVW17, PS18]), we might be able to construct a UCE without iO using ideas from those PCPRFs.

We start with a recap explaining how BM proved that the iO of puncturable PRF is a UCE. We assume readers have some familiarity of the iO and puncturable PRF methodology (for readers who are not familiar with the iO and puncturable PRF methodology, [SW14, BM14b] are good references). Suppose that the source *S* makes only one query, denoted by x^* , and receives the oracle answer $y^* = F(msk, x^*)$. The proof comprises three major steps as shown in fig. 1.

- 1. First, switch y^* to a uniformly random value, and meanwhile, hk is changed to $iO(P_{x^*,y^*,k_{x^*}})$ where k_{x^*} is the punctured key at point x^* and $P_{x^*,y^*,k_{x^*}}(\cdot)$ is the following program:
 - On input *x*, if $x = x^*$, output y^* ; otherwise, output $F(k_{x^*}, x)$.

This step uses the pseudorandomness of $F(x^*)$ given k_{x^*} .

- 2. Next, hk is changed while keeping y* uniformly random. AIPO is introduced in this step: Instead of hardcoding x* into P_{x*,y*,kx*}, it suffices to hardcode a point obfuscation of x* denoted by p_{x*} ← AIPO(x*), because we only need to check whether x = x* or not. Moreover, we can use msk instead of k_{x*}, since F(k_{x*}, x) = F(msk, x) for all x ≠ x*. More precisely, we replace P_{x*,y*,kx*} by
 - $P_{p_{x^*},y^*}$: On input x, if $p_{x^*}(x) = 1$, output y^* ; otherwise, output F(msk, x).
- 3. Finally, hk is switched back to $iO(F(msk, \cdot))$. By the strong unpredictability of *S* and the security of AIPO, it is hard to find x^* given p_{x^*}, y^* , and the message *L* sent by source. Note that x^* is the only input that $F(msk, \cdot)$ and $P_{p_{x^*}, y^*}(\cdot)$ differs. Then we can finish the proof by a result from [BCP14] stating that iO for all circuits in P/poly is also a differing-inputs obfuscator for circuits that differ on at most polynomially many inputs.

Nevertheless, if we wish to use a PCPRF G instead, we will get stuck even in the first step: In a wishful thinking, we want to first switch the key to the punctured key k_{x^*} while using $G(msk, x^*)$ as the oracle answer in the first stage of UCE game; then, we can replace the oracle answer $G(msk, x^*)$ by a random value due to pseudorandomness. However, to switch from msk (i.e., constrained key for an empty circuit) to k_{x^*} relying on the constraint-hiding property, the adversary (for G) has no

access to $G(msk, x^*)$, and thus it cannot properly reply $G(msk, x^*)$ to the query by the source. To fix this issue, we conceive the following strategy:

• Imagine that the punctured key is generated by first sampling y^* (uniformly at random) and then producing a key k_{x^*} such that $G(k_{x^*}, x^*) = y^*$. This way, in the reduction from the constraint-hiding property, the reduction can compute $y^* = G(k_{x^*}, x^*)$ and reply y^* to the source.

That is, we want to program the value at x^* to be y^* , a random value sampled before the generation of the puncture key k_{x^*} .

Even if we manage to do such programming, the AIPO also causes a problem: In the second step, we want to use $p_{x^*} \leftarrow \text{AIPO}(x^*)$ instead of x^* to generate the key. Hence, we have to program the value at x^* to be y^* without explicitly knowing x^* — we are only given p_{x^*} , a circuit that computes the point function $\mathbf{1}_{x^*}$. Therefore, we roughly require the following functionality:

• Given a circuit *C* that computes the point function $\mathbf{1}_{x^*}$ and a value y^* , one can generate a 'programmed key' k_C such that $G(k_C, x^*) = y^*$.

This functionality is supported by the PCPRF constructed from iO [BLW17], but not by the existing lattice-based PCPRFs [BKM17, CC17, BTVW17, PS18]. This motivates our definition of a new primitive called *obliviously programmable function*. By oblivious we mean that the programmed key is generated only given a circuit *C* that computes $\mathbf{1}_{x^*}$, without explicitly knowing x^* .

Obliviously programmable function (OPF). Let C be a circuit class. An OPF $OPF(\cdot, \cdot)$ for C, like PRFs, is a keyed function where the first input is viewed as the key. It provides an algorithm OPF.Program(C, y) that takes as input a circuit $C \in C$ and a value y, and outputs a programmed key k_C . The following properties are required.

- <u>Correctness.</u> If *C* computes the point function $\mathbf{1}_x$, then $OPF(k_C, x) = y$.
- Privacy. k_C computationally hides C if (1) C computes a point function or an all-zero function;
 (2) the programmed value y is uniformly random.
- Value-Hiding for the all-zero function. If C computes the all zero function, k_C computationally hides the value y.

The value-hiding property intuitively says that if *C* computes the all-zero function, the value *y* has no effect. In section 4, we shall prove that $H(hk, x) \stackrel{\text{def}}{=} \mathsf{OPF}(hk, x)$ is a UCE, with the following key generation algorithm:

The key generation algorithm H.Gen chooses a value *y* uniformly at random, and outputs hk ← OPF.Program(C_∅, *y*), where C_∅ ∈ C is any (fixed) circuit that computes the all-zero function.¹

Theorem 1.4. (Theorem 4.3, informally) Assume that there exists an AIPO that always outputs a circuit in the circuit class C (e.g., NC¹). If OPF is an OPF for C, then H defined above is S_1^{scup} -UCE-secure.

¹Actually, by the value-hiding property, we can set y to be any fixed value.

Realizing OPF via GGH15 encodings. We start with a brief introduction of GGH15 encodings. In this framework, circuits are encoded as matrix branching programs (MBPs). A read-once MBP Γ is specified by $\mathbf{v} \in \{0, 1\}^w$ and $\{\mathbf{M}_{i,b} \in \{0, 1\}^{w \times w}\}_{i \in [h], b \in \{0, 1\}}$. On inputs $\mathbf{x} \in \{0, 1\}^h$, the output of the MBP is 1 if $\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} = \mathbf{0}$ and otherwise 0, where $\mathbf{M}_{\mathbf{x}} \stackrel{\text{def}}{=} \prod_{i \in [h]} \mathbf{M}_{i,x_i}$. To encode such a MBP, we first construct $\{\widehat{\mathbf{S}}_{i,b}\}_{i \in [h], b \in \{0, 1\}}$:

$$\widehat{\mathbf{S}}_{1,b} = \left(\mathbf{I}_n \mid \mathbf{v}^{\top} \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b}\right), \widehat{\mathbf{S}}_{i,b} = \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix} \text{ for } i = 2, \dots, h,$$

where $\mathbf{S}_{j,b} \leftarrow \mathcal{D}_{\sigma}^{n \times n}$. Then GGH15 encodings of such an MBP is given by

$$\mathsf{GGH}.\mathsf{Encode}(\{\widehat{\mathbf{S}}_{i,b}\}) = \left\{ \underbrace{\widehat{\mathbf{S}}_{1,b}\mathbf{A}_1}_{\sim \sim \sim \sim \sim}, \mathbf{A}_1^{-1}(\underbrace{\widehat{\mathbf{S}}_{2,b}\mathbf{A}_2}_{\sim \sim \sim \sim \sim}), \ldots, \mathbf{A}_{h-1}^{-1}(\underbrace{\widehat{\mathbf{S}}_{h,b}\mathbf{A}_h}_{\sim \sim \sim \sim \sim}) \right\}_{b \in \{0,1\}}$$

where $\mathbf{A}_i \leftarrow \mathbb{Z}_q^{(n+nw) \times m}$ for some $m = \Theta(nw \log q)$ and wavy underline is put in place of noise terms. Given the encoding and $\mathbf{x} \in \{0, 1\}^h$, we can approximate $\widehat{\mathbf{S}}_{\mathbf{x}} \mathbf{A}_h$, where $\widehat{\mathbf{S}}_{\mathbf{x}} \stackrel{\text{def}}{=} \prod_{i \in [h]} \widehat{\mathbf{S}}_{i,x_i}$. Intuitively, to generate a programmed key for Γ , we let the encoding be the programmed key k_{Γ}

Intuitively, to generate a programmed key for Γ , we let the encoding be the programmed key k_{Γ} and let OPF.Eval (k_{Γ}, \mathbf{x}) be the approximation of $\widehat{\mathbf{S}}_{\mathbf{x}}\mathbf{A}_h$ given by the encoding. But how to program the value at \mathbf{x}^* (assuming Γ computes the point function $\mathbf{1}_{\mathbf{x}^*}$)?

Note that we add a seemingly useless I_n -track in the \hat{S} -matrices. Remarkably, it is this modification that allows us to program on the target point. To see this, observe that

$$\begin{aligned} \mathsf{OPF}.\mathsf{Eval}(k_{\Gamma},\mathbf{x}) &\approx \widehat{\mathbf{S}}_{\mathbf{x}} \mathbf{A}_{h} = \left(\mathbf{I}_{n} \mid \mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} \otimes \mathbf{S}_{\mathbf{x}}\right) \cdot \mathbf{A}_{h} \\ &= \overline{\mathbf{A}_{h}} + \left(\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} \otimes \mathbf{S}_{\mathbf{x}}\right) \cdot \mathbf{A}_{h}, \end{aligned} \tag{2}$$

where $\overline{\mathbf{A}_h}$ denotes the top *n* rows of \mathbf{A}_h and $\underline{\mathbf{A}_h}$ is the rest part. If Γ computes the point function $\mathbf{1}_{\mathbf{x}^*}$, i.e., $\mathbf{v}^\top \mathbf{M}_{\mathbf{x}^*} = \mathbf{0}$, then OPF.Eval $(k_{\Gamma}, \mathbf{x}) \approx \widehat{\mathbf{S}}_{\mathbf{x}^*} \mathbf{A}_h = \overline{\mathbf{A}_h}$. Hence, we can program the value at \mathbf{x}^* by controlling $\overline{\mathbf{A}_h}$.

For privacy, we need to show the encoding computationally hides Γ *provided that* Γ *computes a point function or the all-zero function*. This is reminiscent of witness encryption or null-iO, where security is required only when the underlying circuit computes the all-zero function. In our case, we need to show security also for circuits that compute point functions. To this end, we observe that though [VWW22] aims at constructing witness encryption and null-iO, they in fact give a general reduction assuming evasive LWE: In order to prove the encodings are pseudorandom (thus hiding Γ), it suffices to show that the evaluated products $\{\widehat{\mathbf{S}}_{\mathbf{x}}\mathbf{A}_h + \mathbf{E}_{\mathbf{x}}\}_{\mathbf{x}\in\{0,1\}^h}$ are pseudorandom, where $\mathbf{E}_{\mathbf{x}}$ are independent noises. Continuing eq. (2), we have

$$\begin{split} \widehat{\mathbf{S}}_{\mathbf{x}} \mathbf{A}_h + \mathbf{E}_{\mathbf{x}} &= \overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}) \cdot (\mathbf{I} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \underline{\mathbf{A}_h} + \mathbf{E}_{\mathbf{x}} \\ &\approx \overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}) \cdot \underbrace{\overbrace{(\mathbf{I} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \mathbf{A}_h}^{\text{pseudorandom}}}_{\mathbf{X}}. \end{split}$$

Since there is at most one point \mathbf{x}^* with $\mathbf{v}^\top \mathbf{M}_{\mathbf{x}^*} = \mathbf{0}$, we have that

• the value at \mathbf{x}^* approximately equals to $\overline{\mathbf{A}_h}$ so it is uniformly random as long as $\overline{\mathbf{A}_h}$ is;

• the value at $\mathbf{x} \neq \mathbf{x}^*$ is pseudorandom since $(\mathbf{I} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \underline{\mathbf{A}_h}$ is pseudorandom and $(\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}) \neq \mathbf{0}$.

The value-hiding property follows from a similar argument. The difference is that we want to show the programmed value $\overline{\mathbf{A}_h}$ is computationally hidden, provided that Γ computes the all-zero function. This is true because if Γ computes the all-zero function, then for all $\mathbf{x} \in \{0,1\}^h$, $\widehat{\mathbf{S}}_{\mathbf{x}}\mathbf{A}_h + \mathbf{E}_{\mathbf{x}}$ is randomized by $(\mathbf{I} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \mathbf{A}_h$, hence the evaluated products are still pseudorandom

even if $\overline{\mathbf{A}_h}$ is chosen and fixed by the distinguisher. Consequently, the encoding is pseudorandom by the aforementioned reduction.

We can generalize the above proof strategy to work with a larger class called read-c MBPs for any polynomial c. Any NC¹ circuit can be translated into such MBPs; therefore, we get an OPF for NC¹:

Theorem 1.5. Assuming subexponential LWE and private-coin evasive LWE for specific samplers, there exists an OPF for NC^1 .

Now the first item of theorem 1.1 follows from theorem 1.5 and theorem 1.4.

1.3 Discussion

OPF v.s. PCPRF. OPF is different from PCPRF in both syntax and security.

- Syntax. OPF has no master key, unlike PCPRF. The programmed circuit is chosen before key generation.
- Security. The definition of OPF does not explicitly require an OPF to be a PRF when given black-box access; it is more about programmability. It allows us to program an input-value pair (x^*, y^*) into the function where x^* is not explicitly known only a circuit computing the point function $\mathbf{1}_{x^*}$ is given. Interestingly, though in our construction, the OPF is indeed a PRF when programmed on the all-zero function, but it seems difficult to prove 'every OPF programmed on the all-zero function is a PRF' from the definitions.

Our OPF construction is somewhat similar to the construction of PCPRFs by Chen, Vaikuntanathan, and Wee (CVW, [CVW18]). The CVW construction also uses a two-track structure in the \hat{S} -matrices. While we put all identity matrices in the top track to support programming, the CVW construction puts $S_{i,b}$ in the top track of $\hat{S}_{i,b}$ so that the value computed by the top track (i.e., $S_x \overline{A_h}$) is the evaluation of the constrained PRF with the master key.

We would like to emphasize that the OPF we construct from lattice assumptions is not a constrained PRF, in the sense that we do not support a general constraining algorithm. Constructing an obliviously programmable constrained PRF without using iO remains an open problem.

Programming on more points? Our result suggests that for a keyed function H, if we can program one point in H obliviously, then assuming AIPO, we show that H behaves like an RO with one query. It is conceivable that if we manage to program *q* points obliviously, then assuming obfuscation for functions that take the value 1 on at most *q* values, we can prove H behaves like an RO with *q*-queries, e.g., proving H is a UCE that supports *q*-queries. Indeed, this is realized in [BM14b] using iO plus *composable virtual grey box point obfuscators*.

Another direction is to construct OPF for larger circuit classes, e.g., P/poly, without using iO. If we can do so then there is no need to restrict the AIPO to be in NC¹. One potential approach is to start from the PCPRFs in [BTVW17, PS18], which support P/poly constraints.

Remove the evasive LWE assumption? Another interesting open problem is proving our construction of UCE is secure without using the evasive LWE assumption. Evasive LWE is a plausible but yet strong and unfalsifiable assumption. Let us briefly explain why we currently need to assume evasive LWE by looking back to the previous applications of the GGH15 encoding. For the constructions of PCPRFs and lockable obfuscations [CC17, GKW17, WZ17, CVW18] from the GGH15 encoding, their security properties can be converted *locally* into each level of the GGH15 encoding, so they can be proven from standard LWE. For the constructions of witness encryption [Tsa22, VWW22] from GGH15, the security property (i.e., there is no witness) is *global*, and it is not clear how to convert it to local properties in GGH15, so they use evasive LWE instead. In our proof, 'the AIPO has only one input evaluated to 1' is a global property, and we don't know how to convert it to each level of GGH15 since we don't know which input evaluates to 1 in the AIPO, so we can only argue security using evasive LWE. However, there might be a chance of finding a smart proof technique that allows us to base UCE (and witness encrytion) from GGH15 encoding on standard LWE, and we leave it as an open problem.

2 Preliminaries

Notation. We use lowercase bold symbols for vectors (e.g., **v**) and uppercase for matrices (e.g., **A**). **I**_n denotes the identity matrix of dimension *n*. For a set *S*, we use $\mathcal{U}(S)$ to denote the uniform distribution over *S*. We use \leftarrow to denote sampling from a distribution or choosing an element from a set uniformly at random. For two distributions χ_1 and χ_2 , we write $\chi_1 \approx_s \chi_2$ if χ_1 and χ_2 are statistically close; and $\chi_1 \approx_c \chi_2$ means that they are computationally indistinguishable. In experiments or games, $[\![E]\!]$ equals 1 if the event *E* happens and otherwise 0; Expt \Rightarrow 1 means the experiment Expt outputs 1. For $x \in \mathbb{Z}_q$, let $\lfloor x \rceil_p^{\text{def}} \lfloor x \cdot \frac{p}{q} \rfloor$ denote the rounding of *x* to \mathbb{Z}_p ; and $\lfloor A \rceil_p$ is the matrix resulting from rounding every entry of **A** to \mathbb{Z}_p . For a function $\nu : \mathbb{N} \to [0, 1]$, we write $\nu = \operatorname{negl}(\lambda)$ if for every $c \in \mathbb{N}$, $\nu(\lambda) \leq 1/(c\lambda^c)$ for sufficiently large λ .

2.1 Lattice Background

Gaussian. For $\sigma > 0$, we use \mathcal{D}_{σ} to denote the Guassian over \mathbb{Z} with standard deviation σ , namely, $\forall x \in \mathbb{Z}, \mathcal{D}_{\sigma}(x) \propto e^{-\pi x/\sigma^2}$.

Learning with Error. We recall the learning with errors problem.

Definition 2.1 (Decisional learning with errors (LWE) [Reg09]). For $n, m \in \mathbb{N}$ and modulus $q \ge 2$, distributions θ, π, χ over \mathbb{Z}_q for secret vectors, public matrices, and error vectors respectively. The LWE_{*n*,*m*,*q*, θ,π,χ assumption states that}

$$(\mathbf{s}^{\top}\mathbf{A} + \mathbf{e}^{\top} \mod q, \mathbf{A}) \approx_{c} (\mathbf{u}^{\top}, \mathbf{A})_{c}$$

where $\mathbf{s} \leftarrow \theta^n, \mathbf{A} \leftarrow \pi^{n \times m}, \mathbf{e} \leftarrow \chi^m, \mathbf{u} \leftarrow \mathcal{U}(\mathbb{Z}_q^m)$. The *subexponential LWE* assumption states that for some $\delta > 0$, the above indistinguishability holds against adversaries running in $2^{n^{\delta}}$ with advantage at most $2^{-n^{\delta}}$ with the following parameters:

$$m = \operatorname{poly}(n), \theta = \pi = \mathcal{U}(\mathbb{Z}_q), \chi = \mathcal{D}_{\sigma}, q \leq 2^{n^o} \cdot \sigma$$

Lemma 2.2 ([BLMR13]). Subexponential LWE assumption with $\pi = D_{2\sqrt{n}}$ (and other parameters the same) are implied by the subexponential LWE assumption (with $\pi = U(\mathbb{Z}_q)$).

Trapdoor and preimage sampling. We recall the background of lattice trapdoor and the capability of using the trapdoor to sample a short preimage of the Ajtai function.

Lemma 2.3 ([Ajt99, AP09, MP12]). There is a PPT algorithm TrapSam $(1^n, 1^m, q)$ that, given the modulus $q \ge 2$, dimensions n, m such that $m \ge 2n \log q$, outputs $\mathbf{A} \approx_s \mathcal{U}(\mathbb{Z}_q^{n \times m})$ with a trapdoor τ .

Given any $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} \in \mathbb{Z}_q^n$, $\sigma > 0$, we use $\mathbf{A}^{-1}(\mathbf{y}, \sigma)$ to denote the distribution of a vector \mathbf{d} sampled from \mathcal{D}_{σ}^m conditioned on $\mathbf{A}\mathbf{d} = \mathbf{y} \pmod{q}$. We sometimes suppress σ when the context is clear.

Lemma 2.4 ([GPV08]). There is a PPT algorithm that for $\sigma \ge 2\sqrt{n \log q}$, given $(\mathbf{A}, \tau) \leftarrow \mathsf{TrapSam}(1^n, 1^m, q)$, $\mathbf{y} \in \mathbb{Z}_a^n$, outputs a sample from $\mathbf{A}^{-1}(\mathbf{y}, \sigma)$.

2.2 Matrix Branching Programs

Below we introduce the terminologies for matrix branching programs.

Definition 2.5 (Matrix branching program, MBP). A matrix branching program Γ with width w, length h and input length ℓ consists of the following data:

$$\Gamma = \left(\iota : [h] \to [\ell], \mathbf{v} \in \{0, 1\}^w, \left\{\mathbf{M}_{i, b} \in \{0, 1\}^{w \times w}\right\}_{i \in [h], b \in \{0, 1\}}\right)$$

Here, ι is called the index-to-input map, and it naturally defines an input-to-index map $\varpi : \{0,1\}^{\ell} \to \{0,1\}^{h}$ by letting $\varpi(\mathbf{x})_{i} = x_{\iota(i)}, \forall i \in [h], \mathbf{x} \in \{0,1\}^{\ell}$.

This branching program is computing the function $f_{\Gamma}: \{0,1\}^{\ell} \to \{0,1\}$, defined as

$$f_{\Gamma}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{v}^{\top} \mathbf{M}_{\varpi(\mathbf{x})} = \mathbf{v}^{\top} \prod_{i \in [h]} \mathbf{M}_{i, x_{\iota(i)}} \neq \mathbf{0}; \\ 1 & \text{if } \mathbf{v}^{\top} \mathbf{M}_{\varpi(\mathbf{x})} = \mathbf{v}^{\top} \prod_{i \in [h]} \mathbf{M}_{i, x_{\iota(i)}} = \mathbf{0}. \end{cases}$$

For simplicity, we write $\Gamma(\mathbf{x})$ instead of $f_{\Gamma}(\mathbf{x})$ henceforth.

Write $c \stackrel{\text{def}}{=} h/\ell$. Γ is called a read-c matrix branching program if $\varpi : \{0,1\}^{\ell} \to \{0,1\}^{h}$ outputs c copies of x, namely, $\varpi(\mathbf{x}) = \underbrace{\mathbf{x} | \mathbf{x} | \cdots | \mathbf{x}}_{c \text{ times}}$. In this paper, we shall focus on read-once branching pro-

grams, i.e., c = 1 with ι, ϖ being the identity function, where we simply write $\mathbf{M}_{\mathbf{x}} \stackrel{\text{def}}{=} \prod_{i \in [h]} \mathbf{M}_{i,x_i}$.

Definition 2.6. For $\ell, w, c \in \mathbb{N}$, let $\mathsf{MBP}_{\ell,w}^c$ denote the set of read-*c* matrix branching programs with input length ℓ and width *w*.

Note that read-*c* MBPs are closely related to read-once MBPs: Given a read-*c* MBP Γ with input length ℓ , we can repeat the input *c* times then feed it into a read-once MBP Γ' , where Γ' is the same as Γ except that its index-to-input map is the identity function (so its input length is *h*). However, the correctness is only about Γ' on valid inputs, namely, inputs that are *c*-copies of an ℓ -bit string. To generalize our result to read-*c*-branching programs, we need to ensure that $\Gamma'(\mathbf{x}') = 0$ (i.e., $\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} \neq 0$) for all invalid input \mathbf{x}' . This is done by the following lemma, proven in appendix A.

Lemma 2.7. Let Γ be a read-c MBP with width w, length h, and input length $\ell = h/c$. Let repeat $(\mathbf{x}) = \mathbf{x}|\mathbf{x}| \cdots |\mathbf{x}$. Then there exists a read-once MBP Γ' with the following properties.

c times

- 1. Γ' has width h + w and length h.
- 2. For all $\mathbf{x} \in \{0,1\}^{\ell}$, $\Gamma(\mathbf{x}) = \Gamma'(\operatorname{repeat}(\mathbf{x}))$.
- 3. For all invalid $\mathbf{x}' \in \{0,1\}^h$, i.e., $\mathbf{x}' \neq \text{repeat}(\mathbf{x})$ for any $\mathbf{x} \in \{0,1\}^\ell$, it holds that $\Gamma'(\mathbf{x}') = 0$.

In particular, if Γ computes a point function or all-zero function, then so is Γ' .

2.3 Point Obfuscation

For a point $x \in \{0,1\}^*$, define the point function $\mathbf{1}_x : \{0,1\}^{|x|} \to \{0,1\}$ as $\mathbf{1}_x(u) \stackrel{\mathsf{def}}{=} \begin{cases} 1 & \text{if } u = x \\ 0 & \text{otherwise} \end{cases}$.

We consider a variant of point function obfuscators. Loosely speaking, we require the obfuscation of any point function to be indistinguishable from the obfuscation of the all-zero function. Note that for many distributional VBB obfuscators for point functions or even evasive functions (e.g. [WZ17, GKW17]) appeared in the literature, their simulators are essentially simulating an obfuscated null circuit that is indistinguishable from the real obfuscated circuit, so those constructions immediately satisfy our definition.

Formally, let us first define unpredictable distributions which are used in the definition of obfuscators for point functions.

Definition 2.8 (Unpredictable distribution). A distribution ensemble on $D = \{D_{\lambda} = (Z_{\lambda}, X_{\lambda})\}_{\lambda \in \mathbb{N}}$ is computationally unpredictable if for every poly-size circuit family $\{C_{\lambda}\}_{\lambda \in \mathbb{N}}$ and for all sufficiently large λ , $\mathbf{Pr}_{(z,x)\leftarrow D_{\lambda}}[C_{\lambda}(z) = x] = \operatorname{negl}(\lambda)$.

Definition 2.9 (Auxiliary input point obfuscation for unpredictable distributions (AIPO)). A PPT algorithm AIPO is a *point obfuscator for computationally unpredictable distributions* if it satisfies the following properties.

- <u>Correctness</u>. On input x, it outputs a polynomial-size circuit that computes the point function $\mathbf{1}_x$; on input $(1^{\lambda}, \text{null})$, it outputs a polynomial-size circuit that computes the all-zero function of input length λ , where null is a special symbol.
- Indistinguishability from a null program. For every (efficiently samplable) unpredictable distribution \mathcal{B}_1 over $\{0,1\}^{poly(\lambda)} \times \{0,1\}^{\lambda}$ and every PPT algorithm \mathcal{B}_2 , it holds that

$$\left| \Pr\left[\text{AIPO}_{\mathsf{AIPO},(\mathcal{B}_1,\mathcal{B}_2)}(1^{\lambda}) \Rightarrow 1 \right] - \frac{1}{2} \right| = \texttt{negl}(\lambda)$$

where the experiment AIPO_{AIPO} ($\mathcal{B}_1, \mathcal{B}_2$) (1^{λ}) proceeds as follows.

(z, x) ← B₁(1^λ);
 b ← {0,1};
 C₀ ← AIPO(x), C₁ ← AIPO(1^λ, null);
 b' ← B₂(z, C_b); the experiment outputs 1 iff b = b'.

The circuit complexity of AIPO. Let $C = \{C_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a circuit class, where each $C \in C_{\lambda}$ has input length $\ell_{in}(\lambda)$. We say AIPO is in C if for all $\lambda \in \mathbb{N}, x \in \{0,1\}^{\ell_{in}(\lambda)}$, AIPO(x) always output a circuit in C_{λ} .

Candidate AIPO in NC^1 . Although we only need the *existence* of AIPO in NC^1 , let us mention that *concrete* candidate constructions of AIPO in NC^1 are known directly from the AI-DHI assumption [Can97], or with reductions from other assumptions such as deterministic public-key encryption [BS16], some of which are computable in NC^1 (e.g. [BS14, XXZ12]])

Let us also mention another candidate AIPO in NC¹ from a variant of the LWE assumption. The construction is very simple: To obfuscate a point $\mathbf{s} \in \{0,1\}^n$, we treat \mathbf{s} as an LWE secret, and outputs $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{y} = \mathbf{A}^\top \mathbf{s} + \mathbf{e} \mod q$ as the point obfuscation for \mathbf{s} , where q is a prime in poly(n), $m \ge 20n \log q$, $\mathbf{e} \in \mathbb{Z}^m$ satisfies $\|\mathbf{e}\|_{\infty} < B$ for some bound B such that 1/10 < B/q < 1/5. The evaluation algorithm takes a point $\mathbf{x} \in \{0,1\}^n$, outputs 1 iff $\|\mathbf{A}^\top \mathbf{x} - \mathbf{y} \mod q\|_{\infty} < B$. For certain statistical unpredictable auxiliary inputs, finding the secret is hard based on standard LWE [BD20]. Furthermore, we conjecture the security holds when (\mathbf{s}, z) , where z is the auxiliary input, is sampled from a computationally unpredictable distribution. We don't know how to prove (via a reduction from the standard LWE) the hardness of LWE with a computationally unpredictable source, but we are not aware of any counterexamples either. Note that we choose $B/q \in O(1)$ since a smaller ratio would lead to an attack that succeeds with super-polynomial advantage slightly better than guessing; more details could be found in appendix B.

Let us also remark that even if the underlying AIPO is only secure against a certain subclass of computational unpredictable source C (instead of all computational unpredictable sources), our result will imply UCE and MB-AIPO for corresponding strongly computational unpredictable sources.

2.4 Multi-Bit Point Obfuscation

For a point $x \in \{0,1\}^*$ and value $m \in \{0,1\}^*$, the point function with multi-bit output (MBPF) $p_{x,m}$ is defined as $p_{x,m}(u) \stackrel{\text{def}}{=} \begin{cases} m & \text{if } u = x \\ \bot & \text{otherwise} \end{cases}$. Obfuscating such functions is a natural extension of point obfuscation. Here, we also consider the presence of auxiliary input.

Definition 2.10 (Strongly unpredictable distribution). A distribution ensemble on $D = \{D_{\lambda} = (Z_{\lambda}, X_{\lambda}, M_{\lambda})\}_{\lambda \in \mathbb{N}}$ is *strongly computationally unpredictable* if for every poly-size circuit family $\{C_{\lambda}\}_{\lambda \in \mathbb{N}}$ and for all sufficiently large λ , it holds that $\Pr_{(z,x,m) \leftarrow D_{\lambda}}[C_{\lambda}(z,m) = x] = \operatorname{negl}(\lambda)$.

Definition 2.11 (Multi-Bit AIPO (MB-AIPO)). Let ℓ be a length function. A PPT algorithm AIPO is an ℓ -bit point obfuscator for (strongly) unpredictable distributions if it satisfies the following properties.

- <u>Correctness</u>. On input $(x, m) \in \{0, 1\}^{\lambda} \times \{0, 1\}^{\ell(\lambda)}$, it outputs a polynomial-size circuit that computes the function $p_{x,m}$.
- Value-hiding. For every efficiently samplable (strongly) unpredictable distribution \mathcal{B}_1 over $\overline{\{0,1\}^{\mathsf{poly}(\lambda)} \times \{0,1\}^{\lambda} \times \{0,1\}^{\ell(\lambda)}}$ and every PPT algorithm \mathcal{B}_2 ,

$$\left| \Pr\left[\text{MBAIPO}_{\mathsf{MBAIPO},(\mathcal{B}_1,\mathcal{B}_2)}(1^{\lambda}) \Rightarrow 1 \right] - \frac{1}{2} \right| = \texttt{negl}(\lambda),$$

where the experiment MBAIPO_{AIPO,($\mathcal{B}_1, \mathcal{B}_2$)}(1^{λ}) proceeds as follows.

- 1. $(z, x, m_0) \leftarrow \mathcal{B}_1(1^{\lambda});$
- 2. $b \leftarrow \{0,1\}, m_1 \leftarrow \{0,1\}^{\ell(\lambda)};$
- 3. $C \leftarrow \mathsf{MBAIPO}(x, m_b);$
- 4. $b' \leftarrow \mathcal{B}_2(z, C)$; the experiment outputs 1 iff b = b'.

3 Obliviously Programmable Functions (OPF)

Syntax. Let $\ell_{in} = \ell_{in}(\lambda)$ be a length function. Let $C = \{C_{\lambda}\}_{\lambda \in \mathbb{N}}$ be a circuit class such that every circuit $C \in C_{\lambda}$ has input length $\ell_{in}(\lambda)$, namely, $C : \{0, 1\}^{\ell_{in}(\lambda)} \to \{0, 1\}$. *C* is called a *point circuit* if $|C^{-1}(1)| \stackrel{\text{def}}{=} |\{x \in \{0, 1\}^{\ell_{in}(\lambda)} : C(x) = 1\}| \leq 1$. We assume that there is a simple representation of point functions in C_{λ} , i.e., for every $x \in \{0, 1\}^{\ell_{in}(\lambda)}$, there exists a simple and explicit $P_x \in C_{\lambda}$ that computes $\mathbf{1}_x$.

An *obliviously programmable function* (*OPF*) for C consists of two algorithms $\Pi = (\text{Program}, \text{Eval})$ and specifies a codomain $\mathcal{R} = {\mathcal{R}_{\lambda}}_{\lambda \in \mathbb{N}}$.

- 1. Program $(1^{\lambda}, C, y) \mapsto k_C$. The programming algorithm Program takes as input a point circuit $C \in C_{\lambda}$ and $y \in \mathcal{R}_{\lambda}$, outputs a key k_C .
- 2. $\operatorname{Eval}(k, x) \mapsto y$. The evaluation algorithm Eval, on input a key k and a point $x \in \{0, 1\}^{\ell_{\operatorname{in}}(\lambda)}$, outputs $y = \operatorname{Eval}(k, x) \in \mathcal{R}_{\lambda}$.

Correctness. There exists a negligible function ν such that for all $\lambda \in \mathbb{N}$, $C \in C_{\lambda}$ and $y \in \mathcal{R}_{\lambda}$, if *C* computes the point function $\mathbf{1}_{x^*}$, then

$$\Pr_{k_C \leftarrow \Pi.\mathsf{Program}(1^\lambda,C,y)} \left[\mathsf{Eval}(k_C,x^*) = y\right] \geq 1 - \nu(\lambda)$$

The security of OPF consists of the following two features.

Privacy. We require $k_C \leftarrow \Pi$.Program(C, y) computationally hides C if (1) C is a point circuit and (2) y is chosen uniformly at random. Concretely, for every PPT adversary A,

$$\left| \Pr\left[\operatorname{Priv}_{\Pi, \mathcal{A}}(1^{\lambda}) \Rightarrow 1 \right] - 1/2 \right| = \operatorname{negl}(\lambda),$$

where the experiment $\operatorname{Priv}_{\Pi,\mathcal{A}}(1^{\lambda})$ is defined as follows:

- 1. On input 1^{λ} , \mathcal{A} submits two point circuits $C_0, C_1 \in \mathcal{C}_{\lambda}$ to the challenger.
- 2. The challenger samples $b \leftarrow \{0,1\}, y \leftarrow \mathcal{R}_{\lambda}$, and send $k^* := \Pi.\mathsf{Program}(1^{\lambda}, C_b, y)$ back to \mathcal{A} .
- 3. On receiving k^* , \mathcal{A} outputs b'; the experiment output 1 iff b = b'.

Value-Hiding (when programming on the all-zero function). We require that $k \leftarrow \Pi$. Program(C, y) computationally hides the value y whenever C computes the all-zero function. Formally, for every PPT adversary A,

$$\left| \Pr\left[\operatorname{VH}_{\Pi,\mathcal{A}}(1^{\lambda}) \Rightarrow 1 \right] - 1/2 \right| = \operatorname{negl}(\lambda),$$

where the experiment $VH_{\Pi,\mathcal{A}}(1^{\lambda})$ is defined as follows:

- 1. On input 1^{λ} , \mathcal{A} submits (C, y_0, y_1) to the challenger, where $y_0, y_1 \in \mathcal{R}_{\lambda}$ and C computes the all-zero function.
- 2. The challenger samples $b \leftarrow \{0, 1\}$, and sends $k^* := \Pi$. Program $(1^{\lambda}, C, y_b)$ back to \mathcal{A} .
- 3. On receiving k^* , \mathcal{A} outputs b'; the experiment outputs 1 iff b = b'.

4 Constructions from OPF and AIPO

This section presents several constructions from OPF plus AIPO.

4.1 UCE

UCE is a security notion defined for a keyed function family. We first recall its syntax. A keyed function family $H = \{H_{\lambda} : \mathcal{K}_{\lambda} \times \{0,1\}^{H.\ell_{in}(\lambda)} \to \mathcal{R}_{\lambda}\}_{\lambda \in \mathbb{N}}$ consists of a pair of algorithms (H.Gen, H.Eval) with the following syntax.

- Gen(1^λ) → hk ∈ K_λ. The key generation algorithm Gen outputs a key hk on input security parameter 1^λ.
- Eval(hk, x) $\mapsto y \in \mathcal{R}_{\lambda}$. When hk $\leftarrow \text{Gen}(1^{\lambda}), x \in \{0, 1\}^{\mathsf{H}.\ell_{\mathsf{in}}(\lambda)}$, the evaluation algorithm outputs a value $y \in \mathcal{R}_{\lambda}$.

The UCE security is defined by a two-stage game. The first player has access to an oracle HASH that is either $H.Eval(hk, \cdot)$ or a random oracle. The second player receives a message from the first player and is given hk; it has to guess to which oracle the first player has access.

Definition 4.1. Let (S, D) be an adversary, where *S* is called the source, and *D* is called the distinguisher. We associate them with the game in fig. 2. Define the advantage as $\operatorname{Adv}_{H,(S,D)}^{\operatorname{uce}}(\lambda) \stackrel{\text{def}}{=} \left| \operatorname{Pr} \left[\operatorname{UCE}_{S,D}^{\mathsf{H}}(\lambda) \Rightarrow 1 \right] - \frac{1}{2} \right|$. We say H is *S*-UCE-secure, denoted by $\mathsf{H} \in \operatorname{UCE}[S]$, if for every source $S \in S$ and PPT *D*, $\operatorname{Adv}_{\mathsf{H},(S,D)}^{\operatorname{uce}}(\lambda) = \operatorname{negl}(\lambda)$.

$UCE_{S,D}^{H}(\lambda)$		Hash(x)		
1:	$\beta \leftarrow \{0,1\}, hk \leftarrow H.Gen(1^\lambda)$	1:	if $x \notin Q$ then	
2 :	$L \gets S^{Hash}$	2:	$Q \leftarrow Q \cup \{x\}$	
3:	$\beta' \gets D(L,hk)$	3:	if $\beta = 0$ then $T[x] \leftarrow \mathcal{R}_{\lambda}$	
4:	return $\llbracket \beta = \beta' \rrbracket$	4:	$\mathbf{else} \; T[x] \gets H(hk, x) \mathbf{endif}$	
		5:	endif	
		6:	return $T[x]$	

Figure 2: Games for defining UCE security.

Pre	$d_P^S(\lambda)$	SPre	$\mathrm{ed}_P^S(\lambda)$	Has	$\operatorname{SH}(x)$
1:	$hk \gets H.Gen(1^\lambda)$	1:	$hk \gets H.Gen(1^\lambda)$	1:	if $x \notin Q$ then
2 :	$L \gets S^{Hash}$	2:	$L \gets S^{Hash}$	2:	$Q \leftarrow Q \cup \{x\}$
3 :	$x \leftarrow P(L)$	3:	$x \gets P(L, T. values)$	3:	$T[x] \leftarrow \mathcal{R}_{\lambda}$
4:	return $[\![x \in Q]\!]$	4:	return $[\![x \in Q]\!]$	4:	endif
				5:	return $T[x]$

Figure 3: Games for defining unpredictable sources. Here, *T*.values denote the set of all oracle answers (not query-answer pairs).

Definition 4.2. Consider the games in fig. 3. We say *S* is *unpredictable* if for every PPT predictor P, $\left| \Pr\left[\operatorname{Pred}_{P}^{S}(\lambda) \Rightarrow 1 \right] - \frac{1}{2} \right| = \operatorname{negl}(\lambda)$. Moreover, *S* is called *strongly unpredictable* if for every PPT predictor *P*,

$$\left| \Pr\left[\operatorname{SPred}_{P}^{S}(\lambda) \Rightarrow 1 \right] - \frac{1}{2} \right| = \operatorname{negl}(\lambda).$$

Define S^{cup} and S^{scup} as the sets of unpredictable and strongly unpredictable sources, respectively. Furthermore, given a number $q \in \mathbb{N}$ and a class of sources S, let

 $\mathcal{S}_q \stackrel{\text{def}}{=} \{ S \in \mathcal{S} : S \text{ makes at most } q \text{ queries to HASH} \}.$

Construction 1. Let $\Pi = (\Pi$.Program, Π .Eval) be an OPF for circuit class C with input length $\ell_{in} = \ell_{in}(\lambda)$ and codomain $\{\mathcal{R}_{\lambda}\}_{\lambda \in \mathbb{N}}$. Consider the following construction of a keyed function family H = (H.Gen, H.Eval):

- H.Gen $(1^{\lambda}) \mapsto hk: y \leftarrow \mathcal{R}_{\lambda}, k_{\emptyset} \leftarrow \Pi.Program(C_{\emptyset}, y)$, where C_{\emptyset} denotes the all-zero function; output hk := k_{\emptyset} .
- H.Eval(hk, x) \mapsto y: output $y := \Pi$.Eval(hk, x).

Theorem 4.3. Suppose there exists an OPF for some circuit class C where there exists an AIPO computable in C; let Π be such an OPF and H be the construction in construction **1**. Then $H \in UCE[S_1^{scup}]$.

Proof. Let $S \in S_1^{scup}$ be a source and D be a PPT distinguisher. We start with $Game_0$, as shown in fig. 4, which is exactly the UCE game $UCE_{S,D}^{H}$ when the hidden bit β is fixed to 1 (i.e., HASH returns

the real hash value). Since *S* only queries HASH once, we use a variable x^* to record this query. Our goal is to replace $y^* := H.Eval(hk, x)$ by $y^* \leftarrow \mathcal{R}_{\lambda}$ (as in the last game Game₅) via a sequence of undetectable change.

$Game_0,Game_5$	Hash(x)
1: $x^* := \bot$ 2: $hk \leftarrow \boxed{ \begin{array}{c} H.Gen(1^{\lambda}) \\ \hline 1: y \leftarrow \mathcal{R}_{\lambda} \\ 2: k_{\emptyset} \leftarrow \Pi.Program(1^{\lambda}, C_{\emptyset}, y) \\ \hline 3: hk := k_{\emptyset} \end{array} }$	1: $x^* := x$ 2: $y^* := H.Eval(hk, x^*)$ // Game 0 3: $y^* \leftarrow \mathcal{R}_{\lambda}$ // Game 5 4: return y^*
3: $L \leftarrow S^{\text{Hash}}$ 4: $\beta' \leftarrow D(\text{hk}, L)$ 5: return $\llbracket 1 = \beta' \rrbracket$	

Figure 4: Game₀ and Game₅.

	$Game_0$	$Game_1$	Game ₂	
	1: $x^* := \bot$	$1: x^* := \bot$	1: $x^* := \bot$	
	$2: L \leftarrow S^{\mathrm{Hash}}$	$2: L \gets S^{Hash'}$	$2: L \leftarrow S^{\text{RO}}$	
	3: hk := k_{\emptyset}	3: $hk := k_{x^*}$	3: $k_{x^*} \leftarrow \Pi.Program(1^\lambda,AIPO(x^*),y^*))$	
	$4: \beta' \leftarrow D(hk,L)$	$4: \beta' \leftarrow D(hk,L)$	4: hk := k_{x^*}	
	5: return $\llbracket 1 = \beta' \rrbracket$	5: return $\llbracket 1 = \beta' \rrbracket$	5: $\beta' \leftarrow D(hk, L)$	
			6: return $\llbracket 1 = \beta' \rrbracket$	
Gam	ie ₃	Game ₄	Game ₅	
1:	$x^* := \bot$	$1: x^*:=\bot$	1: $x^* := \bot$	
2:	$msk \leftarrow \Pi.Gen(1^{\lambda})$	2: $msk \leftarrow \Pi.0$	$Gen(1^{\lambda}) \qquad \qquad 2: msk \leftarrow \Pi.Gen(1^{\lambda})$	
3:	$L \leftarrow S^{\mathrm{RO}}$	$3: L \leftarrow S^{\mathrm{RO}}$	$3: L \leftarrow S^{ ext{RO}}$	
4:	$k_{\emptyset} \leftarrow \Pi.Program(msk,AIPO(null),$	y^*) 4: $y \leftarrow \mathcal{R}_{\lambda}$	$4: y \leftarrow \mathcal{R}_{\lambda}$	
5:	$hk:=k_{\emptyset}$	$5: k_{\emptyset} \leftarrow \Pi.Prop$	$ogram(msk,AIPO(null),y) 5: k_\emptyset \leftarrow \Pi.Program(msk,C_\emptyset,g)$	y)
6:	$\beta' \gets D(hk,L)$	6: hk := k_{\emptyset}	$6: hk:=k_{\emptyset}$	
7:	return $\llbracket 1 = \beta' \rrbracket$	7: $\beta' \leftarrow D(hk)$	$\mathbf{x}, L) \qquad \qquad 7: \beta' \leftarrow D(hk, L)$	
		8 : return [[1 =	$= \beta'] $ 8: return $\llbracket 1 = \beta' \rrbracket$	
	Hash(x)	Hash'(x)	$\operatorname{RO}(x)$	
	$1: x^* := x$	$1: x^* :=$	x 1: $x^* := x$	
	2: $y \leftarrow \mathcal{R}_{\lambda}$	2: $y \leftarrow \mathcal{R}$	\mathcal{R}_{λ} 2: $y^* \leftarrow \mathcal{R}_{\lambda}$	
	3: $k_{\emptyset} \leftarrow \Pi.Program(1)$	$^{\lambda}, C_{\emptyset}, y)$ 3: $k_{x^*} \leftarrow$	- $\Pi.Program(1^\lambda,AIPO(x^*),y))$ 3 : return y^*	
	$4: y^*:=\Pi.Eval(k_u,x)$	$^{*}) \qquad \qquad 4: y^{*}:= 1$	$\Pi.Eval(k_{x^*},x^*)$	
	5: return y^*	5: return	n y^*	

Figure 5: Games in the proof of theorem 4.3.

1. Game₀ in fig. 5 is the same as fig. 4, but we move the generation of k_{\emptyset} into HASH, which is just



Figure 6: Outline of the proof.

a conceptual change.

- 2. Game₁. Game₁ is identical to Game₀ except that
 - hk := k_{x^*} where $k_{x^*} \leftarrow \Pi$.Program $(1^{\lambda}, \mathsf{AIPO}(x^*), y)$.

 $Game_0 \approx_c Game_1$ readily follows from the privacy of Π .

- 3. Game₂. Game₂ is identical to Game₁ except that HASH directly returns the programmed value instead of computing Π .Eval (k_{x^*}, x^*) . Also, we move the generation of k_{x^*} out of the oracle, so that the oracle behaves exactly as a random oracle. By the correctness of Π , Game₁ \approx_c Game₂.
- 4. Game₃. Game₃ is identical to Game₂ except that hk is replaced by $k_{\emptyset} \leftarrow \Pi$.Program $(msk, AIPO(null), y^*)$. We shall prove Game₂ \approx_c Game₃ via the security of AIPO and the assumption that S is strongly unpredictable.
- 5. Game₄. Game₄ is identical to Game₃ except that $k_{\emptyset} \leftarrow \Pi$.Program $(1^{\lambda}, \text{AIPO(null}), y)$, where $y \leftarrow \mathcal{R}_{\lambda}$ is a fresh random value (like in Game₀). Note that when programming on the allzero function, the value y is hidden by the programmed key. Therefore, Game₃ \approx_c Game₄ readily follows from the value-hiding property of Π .
- 6. Game₅. Finally, Game₅ is identical to Game₄ except that hk = k_∅ is generated in the original way (k_∅ ← Π.Program(msk, C_∅, y)), removing AIPO. The two programmed keys are indistinguishable according to the privacy of Π, which is similar to Game₁ ≈_c Game₀. Hence, Game₄ ≈_c Game₅

Note that Game₅ is exactly the UCE game with $\beta = 0$; hence we conclude that $\mathbf{Adv}_{\mathsf{H},(S,D)}^{\mathsf{uce}}(\lambda)$ is negligible.

The full proof of the game-hoppings can be found in appendix C due to limited space. See fig. 6 for an outline. \Box

4.2 MB-AIPO and Universal AIPO

We construct an MB-AIPO for *strongly* unpredictable distribution assuming OPF and the existence of AIPO. The same notion of MB-AIPO was constructed from iO and the existence of AIPO in [BM14a].

Construction 2. Let ℓ be a length function. Let $\Pi = (\Pi.\text{Program}, \Pi.\text{Eval})$ be an OPF for circuit class C with input length λ and codomain $\{0,1\}^{\ell(\lambda)}$, and let AIPO be an AIPO in C. Define an algorithm MBAIPO as follows: On input $(x, m) \in \{0,1\}^{\lambda} \times \{0,1\}^{\ell}$,

- 1. sample $r \leftarrow \{0,1\}^{\ell}$ and generate $C \leftarrow \mathsf{AIPO}(x)$;
- 2. generate $k \leftarrow \Pi$.Program $(1^{\lambda}, C, r)$;
- 3. output the program $P[k, r \oplus m, C]$ described below.

Program P[k, w, C]

- Hardwired: A key $k, w \in \{0, 1\}^{\ell}$, and a circuit *C*.
- Input: $x \in \{0, 1\}^{\lambda}$.
- **Operations:** If C(x) = 1, output Π . Eval $(k, x) \oplus w$; otherwise, output \bot .

Theorem 4.4 (MB-AIPO for strongly unpredictable distribution). Suppose there exists an OPF for some circuit class C such that there exists an AIPO computable in C, then Construction 2 is an ℓ -bit MB-AIPO for strongly unpredictable distributions.

The proof of theorem 4.4 can be found in appendix C.

Application: Leakage Resilient Public-key Encryption. Observe that in construction 2, MBAIPO can generate the obfuscated program P[k, w, C] using $C \leftarrow AIPO(x)$ alone, without knowing x explicitly. Therefore, by an argument similar to [BM14a], we have the following PKE construction. The formal security definition and proof are in appendix C.

Construction 3. Let ℓ , Π , AIPO be the same as in construction 2. Define PKE = (PKE.Gen, PKE.Enc, PKE.Dec) as follows.

- PKE.Gen (1^{λ}) : Choose a secret key $x \leftarrow \{0,1\}^{\lambda}$ uniformly at random, and output $C \leftarrow AIPO(x)$ as public key.
- PKE.Enc(C, m): On input public key C and message $m \in \{0, 1\}^{\ell}$, sample $r \leftarrow \{0, 1\}^{\ell}$ and generate $k \leftarrow \Pi$.Program $(1^{\lambda}, C, r)$; output $(k, r \oplus m)$ as ciphertext.
- PKE.Dec(x, (k, z)): On input a secret key *s* and a ciphertext (k, z), output Π .Eval $(k, x) \oplus z$.

Theorem 4.5. Suppose there exists an OPF for some circuit class *C* such that there exists an AIPO computable in *C*, then Construction **3** is IND-CPA secure in the presence of computationally uninvertible leakage on the secret key.

Universal AIPO. If the OPF ensures the evaluation is equal to the programmed value y^* only at the programmed point x^* , we can check whether $x = x^*$ by checking whether the value at x is equal to y^* . This gives a construction of AIPO, assuming the existence of AIPO. The construction itself does not use AIPO; only the security proof relies on the existence of AIPO. This makes the construction, in some sense, universal.

Definition 4.6. Let Π be an OPF for circuit class C with codomain \mathcal{R} . We say Π satisfies *value uniqueness* if there exists a negligible function ν such that for all $\lambda \in \mathbb{N}, C \in C_{\lambda}, y \in \mathcal{R}_{\lambda}$, if C computes the point functions $\mathbf{1}_{x^*}$, then

 $\Pr_{k_C \leftarrow \Pi.\mathsf{Program}(1^\lambda,C,y)} \left[\forall x \neq x^* \ \Pi.\mathsf{Eval}(x) \neq y \right] \geq 1 - \nu(\lambda).$

Construction 4. Let $\Pi = (\Pi. \text{Program}, \Pi. \text{Eval})$ be an OPF for circuit class C with input length λ and codomain $\{0, 1\}^{\ell}$. Suppose that Π satisfies value uniqueness. Define an algorithm AIPO as follows: On input $x \in \{0, 1\}^{\lambda}$,

- 1. Sample $r \leftarrow \{0,1\}^{2\lambda}$ and generate $k \leftarrow \Pi$. Program $(1^{\lambda}, P_x, r)$, where $P_x \in \mathcal{C}$ is a simple circuit computing $\mathbf{1}_x$.
- 2. Output the program P[k, r] described below.

Program P[k, r]

- Hardwired: A key k and $r \in \{0, 1\}^{\ell}$.
- Input: $x \in \{0, 1\}^{\lambda}$.
- **Operations:** If Π . Eval $(k, x) \neq r$, output \bot ; otherwise, output 1.

Theorem 4.7 (Universal AIPO from OPF). Let *C* be the circuit class in construction 4. If there exists an AIPO in *C*, then construction 4 is an AIPO.

Proof. The correctness follows from the correctness and value uniqueness of Π . The security proof is similar to that of theorem 4.4.

5 Obliviously Programmable Function from Lattice Assumptions

In this section, we present a construction of OPF based on GGH15 encodings [GGH15]. We draw on LWE with subexponential hardness and evasive LWE assumption to prove the security of our construction.

Theorem 5.1. Let λ be the security parameter and $\ell_{in}(\lambda), w(\lambda) = \lambda^{O(1)}$. Assume subexponential LWE and private-coin evasive LWE with respect to specific samplers (defined in section 5.2). Then there exists an OPF with input length ℓ_{in} for $MBP^1_{\ell_{in},w}$, where $MBP^1_{\ell_{in},w} = \left\{ MBP^1_{\ell_{in}(\lambda),w(\lambda)} \right\}_{\lambda \in \mathbb{N}}$. (Recall that $MBP^1_{\ell,w}$ denotes the set of read-once matrix branching programs with input length ℓ and width w.) Moreover, the construction satisfies value uniqueness (definition 4.6).

We formally state the evasive LWE assumption in section 5.2 and use it to prove the security property of GGH 15 encodings we need.

Our construction can be generalized to any read-*c* (matrix) branching programs (see appendix A for more detail).

Theorem 5.2. Let λ be the security parameter and $\ell_{in}(\lambda), c(\lambda), w(\lambda) = \lambda^{O(1)}$. Under the same assumption as in theorem 5.1, there exists an OPF with input length ℓ_{in} for MBP^c_{ℓ_{in}, w} (satisfying value uniqueness).

Such branching programs can represent the NC¹ circuit class. Together with theorem 4.3, theorem 4.4, and theorem 4.5, we have our main theorem (theorem 1.1). Theorem 1.3 follows from theorem 5.2 and theorem 4.7.

5.1 OPF Construction Based on GGH15 Encodings

GGH15 encodings. We first describe generalized GGH15 encodings, following [GGH15, CVW18, VWW22].

Construction 5 (GGH15 encodings). The randomized algorithm GGH.Encode takes the following inputs

- parameters 1^{λ} , $h, m, q, \hat{n}_0, \hat{n} \in \mathbb{N}$ and Gaussian parameters $\sigma_1, \sigma_2, \sigma_3, \sigma_4$,
- matrices $\left\{\widehat{\mathbf{S}}_{1,b} \in \mathbb{Z}_q^{\widehat{n}_0 \times \widehat{n}}, \widehat{\mathbf{S}}_{2,b}, \dots, \widehat{\mathbf{S}}_{h,b} \in \mathbb{Z}_q^{\widehat{n} \times \widehat{n}}\right\}_{b \in \{0,1\}}, \mathbf{A}_h \in \mathbb{Z}_q^{\widehat{n} \times m}$,

and proceeds as follows.

- 1. Sample $(\mathbf{A}_i, \tau_i) \leftarrow \mathsf{TrapGen}(1^\lambda, q)$ for $i \in [h-1]$.
- 2. Sample $\mathbf{E}_{1,b} \leftarrow \mathcal{D}_{\sigma_1}^{\widehat{n}_0 \times \widehat{n}}$ and $\mathbf{E}_{2,b}, \dots, \mathbf{E}_{h,b} \leftarrow \mathcal{D}_{\sigma_4}^{\widehat{n} \times \widehat{n}}$ for $i = 2, \dots, b, b \in \{0, 1\}$.
- 3. Compute

$$\mathbf{C}_b := \widehat{\mathbf{S}}_{1,b} \mathbf{A}_1 + \mathbf{E}_{1,b} \text{ and } \mathbf{D}_{i,b} := \mathbf{A}_{i-1}^{-1} (\widehat{\mathbf{S}}_{i,b} \mathbf{A}_i + \mathbf{E}_{i,b}, \sigma_3)$$

for $i = 2, ..., h, b \in \{0, 1\}$, where $\mathbf{A}_{i-1}^{-1}(\cdot, \sigma_3)$ is computed using trapdoor τ_{i-1} .

4. Outputs $C_0, C_1, \{D_{i,b}\}_{i=2,...,h,b\in\{0,1\}}$.

The functionality of GGH15 encodings allows us to approximate $\widehat{\mathbf{S}}_{\mathbf{x}} \mathbf{A}_h$ by $\mathbf{C}_{x_1} \cdot \prod_{i=2}^h \mathbf{D}_{1,x_i}$, which is stated formally in the next lemma.

Lemma 5.3 (Correctness of GGH15 encodings, see, e.g., [CVW18], Lemma 5.3). Let $\mathbf{A}_h \in \mathbb{Z}_q^{\widehat{n} \times m}$. For all $\mathbf{x} \in \{0, 1\}^h$, with high probability over

$$\mathbf{C}_{0}, \mathbf{C}_{1}, \left\{\mathbf{D}_{i,b}\right\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \mathsf{GGH}.\mathsf{Encode}(\left\{\widehat{\mathbf{S}}_{i,b}\right\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_{h}),$$

it holds that

$$\left\| \mathbf{C}_{x_1} \cdot \prod_{i=2}^{h} \mathbf{D}_{1,x_i} - \prod_{i=1}^{h} \widehat{\mathbf{S}}_{i,x_i} \cdot \mathbf{A}_h \right\|_{\infty} \le h \cdot \sigma_1 \cdot \left((\sigma_3 + \sigma_4) m \cdot \max_{i \in [h], b \in \{0,1\}} \left\| \widehat{\mathbf{S}}_{i,b} \right\|_{\infty} \right)^h,$$

where $\left\|\widehat{\mathbf{S}}_{i,b}\right\|_{\infty}$ is the largest absolute value among all entries of $\widehat{\mathbf{S}}_{i,b}$.

The OPF construction. For simplicity, we present the construction for read-once MBPs. The construction can be generalized to read-*c* MBPs, which is presented in appendix **A**.

Construction 6. $\Pi = (\Pi. \text{Program}, \Pi. \text{Eval})$. Input length $\ell_{\text{in}} = h$ and codomain $\mathcal{R}_{\lambda} = \mathbb{Z}_{2}^{n \times m}$. Let $\widehat{n}_{0} \stackrel{\text{def}}{=} n, \widehat{n} \stackrel{\text{def}}{=} n + nw$. For a matrix $\mathbf{U} \in \mathbb{Z}_{q}^{\widehat{n} \times t}$, we use $\overline{\mathbf{U}} \in \mathbb{Z}_{q}^{n \times t}$ to denote the top n rows of \mathbf{U} and $\underline{\mathbf{U}} \in \mathbb{Z}_{q}^{nw \times t}$ the bottom nw rows.

- Program $(1^{\lambda}, C, \mathbf{Y} \in \mathbb{Z}_2^{n \times m}) \mapsto k_C.$
 - 1. Parse *C* as a read-once MBP $\Gamma = (\mathbf{v}, {\mathbf{M}_{i,b}}_{i \in [h], b \in \{0,1\}}).$
 - 2. Sample $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0,1\}$ and set

$$\widehat{\mathbf{S}}_{1,b} = \begin{pmatrix} \mathbf{I}_n \mid \mathbf{v}^\top \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b} \end{pmatrix}, \widehat{\mathbf{S}}_{i,b} = \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix},$$

for $i = 2, ..., h, b \in \{0, 1\}$.

3. Sample $\mathbf{A}_h \leftarrow \mathbb{Z}_q^{\widehat{n} \times m}$ conditioned on $\left\lfloor \overline{\mathbf{A}_h} \right\rfloor_2 = \mathbf{Y}$ and output

$$k_C := \mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \mathsf{GGH}.\mathsf{Encode}(\left\{\widehat{\mathbf{S}}_{i,b}\right\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_h)$$

• $\operatorname{Eval}(k_C, \mathbf{x} \in \{0, 1\}^h) \mapsto \mathbf{Y} \in \mathbb{Z}_2^{n \times m}$. Parse $k_C = \mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b \in \{0,1\}}$ and output $\mathbf{Y} := \left\lfloor \mathbf{C}_{x_1} \prod_{i=2}^h \mathbf{D}_{i,x_i} \right\rfloor_2$.

Parameter setting. Our parameter setting is similar to that of [VWW22]. The adversary runs in time $poly(\lambda)$. We rely on $2^{n^{\delta}}$ -hardness for LWE. We have the following requirements for our parameters:

$$2^{n^{\delta}} > \max \left\{ 2^{h\lambda^{2}}, q/\sigma_{4} \right\}$$
LWE hardness

$$\sigma_{2} = \lambda^{h} \cdot \sigma_{4} \cdot \lambda^{\omega(1)}, \ \sigma_{1} = \sigma_{2} \cdot \lambda^{\omega(1)}$$
noise flooding

$$q \ge \lambda^{\omega(1)} \cdot B \text{ where } B = h \cdot \sigma_{1} \cdot ((\sigma_{3} + \sigma_{4})m \cdot \lambda\sqrt{n})^{h}$$
correctness

$$\sigma_{3} = 2\sqrt{n(w+1)\log q}, \ m = 2n(w+1)\log q$$
trapdoor sampling

A possible setting satisfying the constraints above is

$$n = (h^2 \lambda)^{1/\delta}, \quad q = 2^{n^\delta} = 2^{h^2 \lambda}, \quad \sigma_4 = \Theta(n), \quad m = 2n(w+1)\log q$$

and consequently $B = (poly(\lambda, h))^h$.

Correctness. Correctness of construction 6 readily follows from the correctness of GGH15 encoding and the parameter setting above. Suppose that

$$k_C = \mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \Pi.\mathsf{Program}(1^\lambda, C, \mathbf{Y}).$$

With overwhelming probability, we have $\max_{i \in [h], b \in \{0,1\}} \left\| \widehat{\mathbf{S}}_{i,b} \right\|_{\infty} \le \lambda \sqrt{n}$, Hence, by lemma 5.3, for all $\mathbf{x} \in \{0,1\}^{h}$, it holds (with high probability) that

$$\mathbf{C}_{1,x_1}\prod_{i=2}^{h}\mathbf{D}_{i,x_i}\approx \widehat{\mathbf{S}}_{\mathbf{x}}\mathbf{A}_h = \left(\mathbf{I}_n \mid \mathbf{v}^{\top}\mathbf{M}_{\mathbf{x}}\otimes \mathbf{S}_{\mathbf{x}}\right) \cdot \mathbf{A}_h = \overline{\mathbf{A}_h} + \left(\mathbf{v}^{\top}\mathbf{M}_{\mathbf{x}}\otimes \mathbf{S}_{\mathbf{x}}\right) \cdot \underline{\mathbf{A}}_h,$$

where the \approx is up to an additive factor of $B = h \cdot \sigma_1 \cdot ((\sigma_3 + \sigma_4)m \cdot \lambda \sqrt{n})^h$.

If C(x) = 1, meaning that $\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} = \mathbf{0}$, we get

$$\mathsf{Eval}(k_C, \mathbf{x}) = \left[\mathbf{C}_{x_1} \prod_{i=2}^{h} \mathbf{D}_{i, x_i} \right]_2 \stackrel{(*)}{=} \left[\overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \underline{\mathbf{A}_h} \right]_2 = \left[\overline{\mathbf{A}_h} \right]_2, \tag{3}$$

where (*) holds with probability $1 - \operatorname{negl}(\lambda)$ since we set $q \ge \lambda^{\omega(1)}B$. In the running of Program, it is guaranteed that $\lfloor \overline{\mathbf{A}}_h \rfloor_2 = \mathbf{Y}$, and hence we have $\operatorname{Eval}(k_C, \mathbf{x}) = \mathbf{Y}$ as required.

5.2 Security of GGH15 Encodings from Evasive LWE

We now prove the pseudorandomness of GGH15 encoding with respect to the distributions of $\{\widehat{\mathbf{S}}_{i,b}\}_{i\in[h],b\ in\{0,1\}}$, \mathbf{A}_h involved in our construction: lemma 5.4 and lemma 5.5.

We start by formally stating the private-coin evasive LWE assumption.

Private-coin evasive LWE. Let the parameters

$$param = (q, n, m, m_P, t, \sigma_B, \sigma_P, \sigma)$$

be paramatrized by λ . Let Samp be a PPT algorithm that on input 1^{λ} , outputs

$$\mathbf{S} \in \mathbb{Z}_{a}^{t imes n}, \mathbf{P} \in \mathbb{Z}_{a}^{n imes m_{P}},$$
aux $\in \{0, 1\}^{*}$.

Define the following advantage functions (for adversaries A_0, A_1):

$$\begin{aligned} \mathbf{Adv}_{\mathcal{A}_{0}}^{\mathsf{pre}}(\lambda) &\stackrel{\mathsf{def}}{=} \left| \mathbf{Pr} \left[\mathcal{A}_{0}([\mathbf{SB} + \mathbf{E}], [\mathbf{SP} + \mathbf{E'}], \mathsf{aux}) = 1 \right] - \mathbf{Pr} \left[\mathcal{A}_{0}(\mathbf{C}, \mathbf{C'}, \mathsf{aux}) = 1 \right] \right|, \\ \mathbf{Adv}_{\mathcal{A}_{1}}^{\mathsf{post}}(\lambda) \stackrel{\mathsf{def}}{=} \left| \mathbf{Pr} \left[\mathcal{A}_{1}([\mathbf{SB} + \mathbf{E}], \mathbf{D}, \mathsf{aux}) = 1 \right] - \mathbf{Pr} \left[\mathcal{A}_{1}(\mathbf{C}, \mathbf{D}, \mathsf{aux}) = 1 \right] \right|, \end{aligned}$$

where

$$\begin{split} & (\mathbf{S}, \mathbf{P}, \mathsf{aux}) \leftarrow \mathsf{Samp}\left(1^{\lambda}\right), \\ & \mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_B}^{t \times m}, \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_P}^{t \times m_P}, \\ & \mathbf{C} \leftarrow \mathbb{Z}_q^{t \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{t \times m_P}, \qquad \mathbf{D} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \sigma). \end{split}$$

The EvLWE^{priv}_{param,Samp} assumption states that there exists some polynomial $p(\cdot)$ such that for every PPT A_1 , there exists another PPT A_0 satisfying

$$\mathbf{Adv}_{\mathcal{A}_0}^{\mathsf{pre}}(\lambda) \geq \mathbf{Adv}_{\mathcal{A}_1}^{\mathsf{post}}(\lambda) / p(\lambda) + \mathsf{negl}(\lambda) \text{ and } \mathsf{time}(\mathcal{A}_0) \leq \mathsf{time}(\mathcal{A}_1) \cdot p(\lambda).$$

We consider the same class of samplers used in [VWW22]. Fix parameters $q, n, m, h, \hat{n}, \sigma_1, \sigma_2, \sigma_3, \sigma_4$ in GGH15 encodings. For j = 2, 3, ..., h, let

$$\mathsf{param}_j \stackrel{\mathsf{def}}{=} (q, n, m, 2m, 2^{j-1} \widehat{n}, \sigma_1, \sigma_2, \sigma_3),$$

and the sampler Samp_{VWW,j} outputs the following:

$$\begin{split} \mathbf{S} &:= \left\{ \widehat{\mathbf{S}}_{\mathbf{x}} \right\}_{\mathbf{x} \in \{0,1\}^{j-1}} \in \mathbb{Z}_q^{2^{j-1}\widehat{n} \times m}, \\ \mathbf{P} &:= \left(\widehat{\mathbf{S}}_{j,0} \mathbf{A}_j + \mathbf{E}_{j,0} \| \widehat{\mathbf{S}}_{j,1} \mathbf{A}_j + \mathbf{E}_{j,1} \right) \in \mathbb{Z}_q^{\widehat{n} \times 2m} \\ \mathsf{aux} &:= \left\{ \mathbf{A}_{i-1}^{-1} \left(\widehat{\mathbf{S}}_{i,b} \mathbf{A}_j + \mathbf{E}_{i,b}, \sigma_3 \right) \right\}_{i \ge j+1, b \in \{0,1\}}, \left\{ \widehat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}. \end{split}$$

where

- \$\hfrac{\mathbf{S}}{\mathbf{s}_{i,b}\$ for \$i \in [h]\$, \$b \in {0,1}\$ are arbitrary matrices (which in our context represents a branching program), and \$\begin{bmatrix} \mathbf{S}_{\mathbf{x}} \Begin{bmatrix} \$_{x \in {\{0,1\}}^{j-1}\$}\$ denotes stacking the \$2^{j-1}\$ matrices vertically;
- $\mathbf{E}_{i,b} \leftarrow \mathcal{D}_{\sigma_4}^{\widehat{n} \times m}$ for all $i \ge j$ and $b \in \{0, 1\}$;
- \mathbf{A}_i for $i \ge j$ are uniformly random matrices.

Let $\mathsf{EvLWE}_{\mathsf{VWW}}^{\mathsf{priv}}$ denote the assumption that $\mathsf{'EvLWE}_{\mathsf{param}_j,\mathsf{Samp}_{\mathsf{VWW},j}}^{\mathsf{priv}}$ holds for all $j \in \{2, 3, \ldots, h\}$ and all choices of $\{\widehat{\mathbf{S}}_{i,b}\}_{i \in [h], b \in \{0,1\}}$ '—this is the precise assumption we use.

The pseudorandomness of GGH15 encodings in our constructions is captured by the following two lemmas.

Lemma 5.4. Let $\Gamma = \left(\mathbf{v} \in \{0,1\}^w, \left\{ \mathbf{M}_{i,b} \in \{0,1\}^{w \times w} \right\}_{i \in [h], b \in \{0,1\}} \right)$ be a read-once MBP such that $|\Gamma^{-1}(1)| \leq 1$. Let $\mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,...,h,b \in \{0,1\}}$ be generated as follows.

1. Sample $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0,1\}$, $\mathbf{A}_h \leftarrow \mathbb{Z}_q^{\widehat{m} \times m}$ and set

$$\widehat{\mathbf{S}}_{1,b} := \begin{pmatrix} \mathbf{I}_n \mid \mathbf{v}^{\top} \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b} \end{pmatrix}, \widehat{\mathbf{S}}_{i,b} := \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix},$$

2. *Output* $\mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \mathsf{GGH}.\mathsf{Encode}(\{\widehat{\mathbf{S}}_{i,b}\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_h).$

Then, assuming subexponential LWE and EvLWE^{priv}_{VWW}, we have

$$\{\mathbf{C}_b\}_{b\in\{0,1\}}, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}} \approx_c \{\mathcal{U}(\mathbb{Z}_q^{n\times m})\}_{b\in\{0,1\}}, \{\mathcal{D}_{\sigma_3}^{m\times m}\}_{i=2,\dots,h,b\in\{0,1\}}.$$

Lemma 5.5. Fix $\overline{\mathbf{A}}_h \in \mathbb{Z}_q^{n \times m}$. Let $\Gamma = \left(\mathbf{v} \in \{0,1\}^w, \{ \mathbf{M}_{i,b} \in \{0,1\}^{w \times w} \}_{i \in [h], b \in \{0,1\}} \right)$ be a read-once MBP that computes the all-zero function. Let $\mathbf{C}_0, \mathbf{C}_1, \{ \mathbf{D}_{i,b} \}_{i=2,\dots,h,b \in \{0,1\}}$ be generated as follows.

1. Sample $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0, 1\}$, $\underline{\mathbf{A}_h} \leftarrow \mathbb{Z}^{nw \times m}$, and set

$$\widehat{\mathbf{S}}_{1,b} := ig(\mathbf{I}_n \mid \mathbf{v}^{ op} \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b}ig), \widehat{\mathbf{S}}_{i,b} := ig(egin{matrix} \mathbf{I}_n & & \ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix}, \mathbf{A}_h := ig(egin{matrix} \mathbf{A}_h \ & \mathbf{A}_h \end{pmatrix}$$

2. *Output* $\mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow (\{\widehat{\mathbf{S}}_{i,b}\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_h).$

Then, assuming subexponential LWE and EvLWE^{priv}_{VWW}, we have

$$\{ \mathbf{C}_b \}_{b \in \{0,1\}}, \{ \mathbf{D}_{i,b} \}_{i=2,\dots,h,b \in \{0,1\}}, \mathbf{A}_h \\ \approx_c \left\{ \mathcal{U}(\mathbb{Z}_q^{n \times m}) \right\}_{b \in \{0,1\}}, \left\{ \mathcal{D}_{\sigma_3}^{m \times m} \right\}_{i=2,\dots,h,b \in \{0,1\}}, \overline{\mathbf{A}_h}.$$

The proofs of the two lemmas above follow the proof structure of [VWW22].

1. First, relying on evasive LWE, proving the pseudorandomness of GGH15 encodings

$$\mathbf{C}_{0}, \mathbf{C}_{1}, \left\{\mathbf{D}_{i,b}\right\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \mathsf{GGH}.\mathsf{Encode}(\left\{\widehat{\mathbf{S}}_{i,b}\right\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_{h})$$

is reduced to proving that all the evaluation products $\{\widehat{\mathbf{S}}_{\mathbf{x}'}\mathbf{A}_j + \mathbf{E}_{\mathbf{x}'}\}_{j\in[h],\mathbf{x}'\in\{0,1\}^j}$ are pseudorandom. Here, $\mathbf{E}_{\mathbf{x}'}$ are independent errors and $\mathbf{A}_j \leftarrow \mathbb{Z}_q^{n\times m}$ for $j \in [h-1]$, but the distribution of \mathbf{A}_h could be tailored. This is done by lemma 5.6, which is directly from [VWW22].

2. It remains to show that all the evaluation products are indeed pseudorandom; we call this a 'precondition'. The only difference between lemma 5.4 and lemma 5.5 is that we have different distributions for A_h . Hence, we verify the preconditions respectively in the two claims at the end of this subsection.

The reduction step is by the following lemma.

Lemma 5.6 (Lemma 5.1 in [VWW22]). Fix some distributions for $\{\widehat{\mathbf{S}}_{i,b}\}_{i\in[h],b\in\{0,1\}}$ and let \mathcal{E} be an efficiently samplable (and publicly known) distribution over $\mathbb{Z}_q^{\widehat{n}\times m}$. Suppose that for all $j \in [h]$, we have

$$\left\{ \begin{bmatrix} \widehat{\mathbf{S}}_{\mathbf{x}'} \mathbf{A}_j + \mathbf{E}_{\mathbf{x}'} \end{bmatrix} \right\}_{\mathbf{x}' \in \{0,1\}^j}, \left\{ \widehat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}$$

$$\approx_c \left\{ \mathcal{U}(\mathbb{Z}_q^{\widehat{n}_0 \times m}) \right\}_{\mathbf{x}' \in \{0,1\}^j}, \left\{ \widehat{\mathbf{S}}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}$$

$$(4)$$

where $\mathbf{E}_{\mathbf{x}'} \leftarrow \mathcal{D}_{\sigma_1}^{\hat{n}_0 \times m}$, $\mathbf{A}_j \leftarrow \mathbb{Z}_q^{\hat{n} \times m}$ for $j \in [h-1]$, and $\mathbf{A}_h \leftarrow \mathcal{E}$. Then, assuming subexponential LWE and $\mathsf{EvLWE}_{\mathsf{VWW}}^{\mathsf{priv}}$, we have

$$\{\mathbf{C}_b\}_{b\in\{0,1\}}, \{\mathbf{D}_{i,b}\}_{i=2,\dots,h,b\in\{0,1\}}, \approx_c \{\mathcal{U}(\mathbb{Z}_q^{n\times m})\}_{b\in\{0,1\}}, \{\mathcal{D}_{\sigma_3}^{m\times m}\}_{i=2,\dots,h,b\in\{0,1\}},$$

where

$$\left\{\mathbf{C}_{b}\right\}_{b\in\left\{0,1\right\}},\left\{\mathbf{D}_{i,b}\right\}_{i=2,\dots,h,b\in\left\{0,1\right\}}\leftarrow\mathsf{GGH}.\mathsf{Encode}\left(\left\{\widehat{\mathbf{S}}_{i,b}\right\}_{i\in\left[h\right],b\in\left\{0,1\right\}},\mathbf{A}_{h}\right),\mathbf{A}_{h}\leftarrow\mathcal{E}.$$

Remark 5.7. A few comments are in order regarding lemma 5.6.

- 1. It is required that the precondition eq. (4) holds with hardness $2^{h^2\lambda}$, namely, any adversary running within $2^{h^2\lambda}$ times has advantage at most $2^{-h^2\lambda}$.
- 2. For parameters in the proof of this lemma, it is also required that $\sigma_1 = \sigma_2 \cdot \lambda^{\omega(1)}$ and $\sigma_2 = \lambda^h \cdot \sigma_4 \cdot \lambda^{\omega(1)}$ for noise flooding.

The following two claims verify the preconditions for two different distributions of A_h .

Claim 5.8 (Precondition 1). Fix a MBP

$$\Gamma = \left(\mathbf{v} \in \{0, 1\}^{w}, \left\{ \mathbf{M}_{i, b} \in \{0, 1\}^{w \times w} \right\}_{i \in [h], b \in \{0, 1\}} \right)$$

such that $|\Gamma^{-1}(1)| \leq 1$. Let

$$\widehat{\mathbf{S}}_{1,b} := \begin{pmatrix} \mathbf{I}_n \mid \mathbf{v}^{ op} \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b} \end{pmatrix}, \widehat{\mathbf{S}}_{i,b} := \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix}$$

where $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0, 1\}$. Then, by the LWE assumption, for every $j \in [h]$ we have

$$\left\{ \begin{bmatrix} \widehat{\mathbf{S}}_{\mathbf{x}'} \mathbf{A}_j + \mathbf{E}_{\mathbf{x}'} \end{bmatrix} \right\}_{\mathbf{x}' \in \{0,1\}^j}, \{\mathbf{S}_{i,b}\}_{i \in [h], b \in \{0,1\}}$$

$$\approx_c \left\{ \mathcal{U}(\mathbb{Z}_q^{n \times m}) \right\}_{\mathbf{x}' \in \{0,1\}^j}, \{\mathbf{S}_{i,b}\}_{i \in [h], b \in \{0,1\}},$$

$$(5)$$

where $\mathbf{E}_{\mathbf{x}'} \leftarrow \mathcal{D}_{\sigma_1}^{n \times m}$, $\mathbf{A}_j \leftarrow \mathbb{Z}_q^{\widehat{n} \times m}$ for $j \in [h]$.

Proof. Fix an arbitrary $j \in [h]$. For all $\mathbf{x}' \in \{0, 1\}^j$, we have

$$F(\mathbf{x}') \stackrel{\text{def}}{=} \widehat{\mathbf{S}}_{\mathbf{x}'} \mathbf{A}_j + \mathbf{E}_{\mathbf{x}'} = \left(\mathbf{I}_n \mid \mathbf{v}^\top \mathbf{M}_{\mathbf{x}'} \otimes \mathbf{S}_{\mathbf{x}'}\right) \cdot \mathbf{A}_h + \mathbf{E}_{\mathbf{x}'} \\ = \overline{\mathbf{A}_h} + \left(\mathbf{v}^\top \mathbf{M}_{\mathbf{x}'} \otimes \mathbf{S}_{\mathbf{x}}\right) \cdot \underline{\mathbf{A}_h} + \mathbf{E}_{\mathbf{x}'} \\ = \overline{\mathbf{A}_h} + \left(\mathbf{v}^\top \mathbf{M}_{\mathbf{x}'} \otimes \mathbf{I}_n\right) \cdot \left(\mathbf{I}_w \otimes \mathbf{S}_{\mathbf{x}'}\right) \cdot \underline{\mathbf{A}_h} + \mathbf{E}_{\mathbf{x}'} \\ \approx_s \overline{\mathbf{A}_h} + \left(\mathbf{v}^\top \mathbf{M}_{\mathbf{x}'} \otimes \mathbf{I}_n\right) \cdot \underbrace{\left(\left(\mathbf{I}_w \otimes \mathbf{S}_{\mathbf{x}'}\right) \cdot \underline{\mathbf{A}_h} + \mathcal{D}_{\sigma_2}^{nw \times m}\right)}_{\stackrel{\text{def}}{=} G(\mathbf{x}')} + \mathbf{E}_{\mathbf{x}'},$$

where the last step is by noise flooding ($\sigma_1 = \sigma_2 \cdot \lambda^{\omega(1)}$). Next, by the security of the BLMR PRF [BLMR13], $G(\mathbf{x}')$ is pseudorandom, i.e.,

$$\left\{ \begin{bmatrix} (\mathbf{I}_w \otimes \mathbf{S}_{\mathbf{x}'}) \cdot \underline{\mathbf{A}}_h + \mathcal{D}_{\sigma_2}^{nw \times m} \end{bmatrix} \right\}_{\mathbf{x}' \in \{0,1\}^j}, \{\mathbf{S}_{i,b}\}_{i \in [h], b \in \{0,1\}}$$

 $\approx_c \left\{ \mathcal{U}(\mathbb{Z}_q^{nw \times m}) \right\}_{\mathbf{x}' \in \{0,1\}^j}, \{\mathbf{S}_{i,b}\}_{i \in [h], b \in \{0,1\}}.$

This relies on $((\mathbf{I}_w \otimes \mathbf{S})\mathbf{A} + \mathbf{E}, \mathbf{S}) \approx_c (\mathcal{U}(\mathbb{Z}_q^{nw \times m}), \mathbf{S})$, which follows from LWE via a simple reduction to the case of lemma 2.2.

• $\underline{j \in [h-1]}$. Since $|\Gamma^{-1}(1)| \leq 1$, we have $\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}'} \neq \mathbf{0}$ for all $\mathbf{x}' \in \{0,1\}^j$. Therefore, $(\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}_n) \cdot G(\mathbf{x}') \approx_c \mathcal{U}(\mathbb{Z}_q^{nw \times m})$ and hence $\{F(\mathbf{x}')\}_{\mathbf{x}' \in \{0,1\}^j} \approx_c \{\mathcal{U}(\mathbb{Z}_q^{nw \times m})\}_{\mathbf{x}' \in \{0,1\}^j}$.

• $\underline{j = h}$. If Γ computes the all-zero function, the argument in the first item still goes for j = h. Now assume that $\Gamma^{-1}(1) = \{\mathbf{x}^*\}$. Then

$$F(\mathbf{x}) = \begin{cases} \overline{\mathbf{A}_h} + \mathbf{E}_{\mathbf{x}} & \text{if } \mathbf{x} = \mathbf{x}^* \\ \overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}_n) \cdot G(\mathbf{x}) + \mathbf{E}_{\mathbf{x}} & \text{if } \mathbf{x} \neq \mathbf{x}^* \end{cases}$$

Since $\overline{\mathbf{A}_h}$ is uniformly distributed and $(\mathbf{v}^{\top}\mathbf{M}_{\mathbf{x}} \otimes \mathbf{I}_n) \cdot G(\mathbf{x}) \approx_c \mathcal{U}(\mathbb{Z}_q^{nw \times m})$ for all $\mathbf{x} \neq \mathbf{x}^*$, we have $\{F(\mathbf{x})\}_{\mathbf{x} \in \{0,1\}^h} \approx_c \{\mathcal{U}(\mathbb{Z}_q^{nw \times m})\}_{\mathbf{x} \in \{0,1\}^h}$, proving eq. (5) for j = h.

Claim 5.9 (Precondition 2). *Fix* $\overline{\mathbf{A}}_h \in \mathbb{Z}_q^{n \times m}$. *Let*

$$\Gamma = \left(\mathbf{v} \in \{0, 1\}^{w}, \left\{ \mathbf{M}_{i, b} \in \{0, 1\}^{w \times w} \right\}_{i \in [h], b \in \{0, 1\}} \right)$$

be an MBP that computes the all-zero function and let

$$\widehat{\mathbf{S}}_{1,b} := \begin{pmatrix} \mathbf{I}_n \mid \mathbf{v}^{ op} \mathbf{M}_{1,b} \otimes \mathbf{S}_{1,b} \end{pmatrix}, \widehat{\mathbf{S}}_{i,b} := \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix}, \mathbf{A}_h := \begin{pmatrix} \overline{\mathbf{A}_h} \\ \underline{\mathbf{A}_h} \end{pmatrix},$$

where $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0,1\}$, $\underline{\mathbf{A}_h} \leftarrow \mathbb{Z}_q^{\widehat{n} \times m}$. Then, by the LWE assumption, for every $j \in [h]$ we have

$$\left\{ \left| \widehat{\mathbf{S}}_{\mathbf{x}'} \mathbf{A}_{j} + \mathbf{E}_{\mathbf{x}'} \right| \right\}_{\mathbf{x}' \in \{0,1\}^{j}}, \left\{ \mathbf{S}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}, \overline{\mathbf{A}}_{h}$$

$$\approx_{c} \left\{ \overline{\mathcal{U}}(\mathbb{Z}_{q}^{\widehat{n}_{0} \times m}) \right\}_{\mathbf{x}' \in \{0,1\}^{j}}, \left\{ \mathbf{S}_{i,b} \right\}_{i \in [h], b \in \{0,1\}}, \overline{\mathbf{A}}_{h},$$

$$= \widehat{\mathbf{A}}_{i} = \mathbf{A}_{i}$$
(6)

where $\mathbf{E}_{\mathbf{x}'} \leftarrow \mathcal{D}_{\sigma_1}^{n \times m}$, $\mathbf{A}_j \leftarrow \mathbb{Z}_q^{\widehat{n} \times m}$ for $j \in [h-1]$.

Proof. For $j \le h - 1$, the proof is the same as that of the previous claim (Precondition 1). For j = h, note that when $\mathbf{v}^{\top} \mathbf{M}_{\mathbf{x}} \ne \mathbf{0}$ for all $\mathbf{x} \in \{0, 1\}^h$, so $F(\mathbf{x})$ is completely randomized by $G(\mathbf{x})$, and hence knowing $\overline{\mathbf{A}_h}$ is not helpful.

Proving the pseudorandomness of GGH15 encodings: lemma 5.4 and lemma 5.5.

Proof of lemma 5.4. Since $|\Gamma^{-1}(1)| \le 1$, by the first claim (Precondition 1), the precondition of lemma 5.6 holds with \mathcal{E} being the uniform distribution over $\mathbb{Z}_q^{\widehat{n} \times m}$. Then the lemma follows from lemma 5.6.

Proof of lemma 5.5. The second claim (Precondition 2) shows that the precondition of lemma 5.6 holds for the assumed distribution of $\{\widehat{\mathbf{S}}_{i,b}\}_{i \in [h], b \in \{0,1\}}$ and \mathbf{A}_h . Note that \mathcal{E} is publicly known and hence $\overline{\mathbf{A}_h}$ can be given to the distinguisher. Then the lemma follows from lemma 5.6.

5.3 Security Proof of Our OPF Construction

Finally, we show that construction $\frac{6}{6}$ is indeed an OPF, proving theorem 5.1.

Theorem 5.10 (Security of construction 6). Under subexponential LWE assumption and evasive LWE assumption, construction 6 is an obliviously programmable function for $MBP^1_{\ell_{in},w}$.

Proof. Correctness is proven in section 5.1. The proof of value uniqueness is omitted and can be found in appendix C. Here we prove that the two security properties of OPF are satisfied.

Privacy. Note that when $\mathbf{Y} \leftarrow \mathbb{Z}_2^{n \times m}$, \mathbf{A}_h sampled in Π .Program $(1^{\lambda}, \Gamma, \mathbf{Y})$ is also uniformly distributed, and thus satisfies the condition of lemma 5.4. By lemma 5.4, for any $\Gamma \in \mathsf{MBP}^1_{\ell_{in},w}$, if $|\Gamma^{-1}(1)| \leq 1$, then $k_{\Gamma} \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma, \mathbf{Y})$ is pseudorandom, where $\mathbf{Y} \leftarrow \mathbb{Z}_2^{n \times m}$. Therefore, for arbitrary (Γ_0, Γ_1) with $|\Gamma_0^{-1}(1)| \leq 1$, $|\Gamma_1^{-1}(1)| \leq 1$, no PPT adversary can distinguish $k_{\Gamma_0} \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma_0, \mathbf{Y})$ from $k_{\Gamma_1} \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma_1, \mathbf{Y})$ for $\mathbf{Y} \leftarrow \mathbb{Z}_2^{n \times m}$.

Value-Hiding when programming on the all-zero function. Let Γ_{\emptyset} be a MBP that computes the all-zero function. Note that $\mathbf{Y} = [\overline{\mathbf{A}_h}]_2$ is a deterministic function of $\overline{\mathbf{A}_h}$. By lemma 5.5, for any fixed $\mathbf{Y}, k_{\Gamma} \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma, \mathbf{Y})$ is pseudorandom even if the adversary knows \mathbf{Y} . Therefore, for arbitrary $(\mathbf{Y}_0, \mathbf{Y}_1)$, no PPT adversary can distinguish $k_0 \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma_{\emptyset}, \mathbf{Y}_0)$ from $k_1 \leftarrow \Pi$.Program $(1^{\lambda}, \Gamma_{\emptyset}, \mathbf{Y}_1)$.

Acknowledgments

We thank the authors of [BÜW24] for sharing an early version of [BÜW24] with us. We also thank Minki Hhan for explaining the dual attack for solving LWE to us. We also thank anonymous reviewers for their helpful comments. Y.C. is supported by Tsinghua University start-up funding and Shanghai Qi Zhi Institute Innovation Program SQZ202405. X. M. is supported by NSF CAREER award 2141536.

References

- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *ICALP*, volume 1644 of Lecture Notes in Computer Science, pages 1–9. Springer, 1999. 10 [AP09] Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. In STACS, volume 3 of LIPIcs, pages 75–86. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2009. 10 [BCP14] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In *Theory* of cryptography conference, pages 52–73. Springer, 2014. 5 [BD20] Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In EUROCRYPT (2), volume 12106 of Lecture Notes in Computer Science, pages 551–575. Springer, 2020. 12 [BFM14] Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Indistinguishability obfuscation and uces: The case of computationally unpredictable sources. In Advances in Cryptology–CRYPTO 2014: 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I 34, pages 188–205. Springer, 2014. 2 [BG14] Shi Bai and Steven D. Galbraith. Lattice decoding attacks on binary LWE. In ACISP, volume 8544 of Lecture Notes in Computer Science, pages 322–337. Springer, 2014. 32
- [BHK13] Mihir Bellare, Viet Tung Hoang, and Sriram Keelveedhi. Instantiating random oracles via uces. In Advances in Cryptology–CRYPTO 2013: 33rd Annual Cryptology Conference,

Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part II, pages 398–415. Springer, 2013. 1, 2, 3

- [BKM17] Dan Boneh, Sam Kim, and Hart William Montgomery. Private puncturable prfs from standard lattice assumptions. In EUROCRYPT (1), volume 10210 of Lecture Notes in Computer Science, pages 415–445, 2017. 5, 6
- [BLMR13] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *Annual Cryptology Conference*, pages 410–428. Springer, 2013. 10, 25
- [BLW17] Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. In *Public Key Cryptography* (2), volume 10175 of *Lecture Notes in Computer Science*, pages 494–524. Springer, 2017. 5, 6
- [BM14a] Christina Brzuska and Arno Mittelbach. Indistinguishability obfuscation versus multibit point obfuscation with auxiliary input. In Advances in Cryptology–ASIACRYPT 2014: 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, ROC, December 7-11, 2014, Proceedings, Part II 20, pages 142– 161. Springer, 2014. 1, 3, 17, 18
- [BM14b] Christina Brzuska and Arno Mittelbach. Using indistinguishability obfuscation via uces. In ASIACRYPT (2), volume 8874 of Lecture Notes in Computer Science, pages 122– 141. Springer, 2014. 1, 2, 4, 5, 8
- [BP12] Nir Bitansky and Omer Paneth. Point obfuscation and 3-round zero-knowledge. In *Theory of Cryptography Conference*, pages 190–208. Springer, 2012. 2
- [BS14] Zvika Brakerski and Gil Segev. Better security for deterministic public-key encryption: The auxiliary-input setting. *J. Cryptol.*, 27(2):210–247, 2014. 12
- [BS16] Mihir Bellare and Igors Stepanovs. Point-function obfuscation: A framework and generic constructions. In *TCC (A2)*, volume 9563 of *Lecture Notes in Computer Science*, pages 565–594. Springer, 2016. 12
- [BTVW17] Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained prfs (and more) from LWE. In *Theory of Cryptography - 15th International Conference, TCC 2017, Baltimore, MD, USA, November 12-15, 2017, Proceedings, Part I,* pages 264–302, 2017. 5, 6, 9
- [BÜW24] Chris Brzuska, Akin Ünal, and Ivy K. Y. Woo. Evasive lwe assumptions: Definitions, classes, and counterexamples. In *ASIACRYPT* 2024. Springer-Verlag, 2024. 4, 27
- [Can97] Ran Canetti. Towards realizing random oracles: Hash functions that hide all partial information. In Advances in Cryptology—CRYPTO'97: 17th Annual International Cryptology Conference Santa Barbara, California, USA August 17–21, 1997 Proceedings 17, pages 455–469. Springer, 1997. 2, 12
- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained PRFs for NC¹ from LWE. In *EUROCRYPT 2017, Part I*, pages 446–476, 2017. **5**, **6**, **9**

- [CGH04] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited. J. ACM, 51(4):557–594, 2004. 2
- [CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In CRYPTO (2), volume 10992 of Lecture Notes in Computer Science, pages 577–607. Springer, 2018. 8, 9, 20
- [FO99] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In Advances in Cryptology—CRYPTO'99: 19th Annual International Cryptology Conference Santa Barbara, California, USA, August 15–19, 1999 Proceedings, pages 537–554. Springer, 1999. 2
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Advances in cryptology*—*CRYPTO '86 (Santa Barbara, Calif., 1986)*, volume 263 of *Lecture Notes in Comput. Sci.*, pages 186–194. Springer, Berlin, 1987.
 2
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *TCC 2015, Part II*, pages 498–527, 2015. 1, 3, 19, 20
- [GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In *ICS*, pages 230–240. Tsinghua University Press, 2010. 4
- [GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In FOCS, pages 612–621, 2017. 9, 11
- [GOR11] Vipul Goyal, Adam O'Neill, and Vanishree Rao. Correlated-input secure hash functions. In *Theory of cryptography*, volume 6597 of *Lecture Notes in Comput. Sci.*, pages 182–200. Springer, Heidelberg, 2011. 2
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008. 10
- [HLL23] Yao-Ching Hsieh, Huijia Lin, and Ji Luo. Attribute-based encryption for circuits of unbounded depth from lattices. In 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS), pages 415–434. IEEE, 2023. 4
- [MH14] Takahiro Matsuda and Goichiro Hanaoka. Chosen ciphertext security via point obfuscation. In *Theory of Cryptography: 11th Theory of Cryptography Conference, TCC 2014, San Diego, CA, USA, February 24-26, 2014. Proceedings 11*, pages 95–120. Springer, 2014. 1, 2
- [MOZ23] Alice Murphy, Adam O'Neill, and Mohammad Zaheri. Instantiability of classical random-oracle-model encryption transforms. In Advances in Cryptology–ASIACRYPT 2022: 28th International Conference on the Theory and Application of Cryptology and Information Security, Taipei, Taiwan, December 5–9, 2022, Proceedings, Part IV, pages 323–352. Springer, 2023. 1
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In EUROCRYPT, volume 7237 of Lecture Notes in Computer Science, pages 700– 718. Springer, 2012. 10

- [MPV24] Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge snarks for up. In Annual International Cryptology Conference, pages 38– 71. Springer, 2024. 4
- [PS18] Chris Peikert and Sina Shiehian. Privately constraining and programming PRFs, the LWE way. In *PKC*, 2018. 5, 6, 9
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009. 9
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *STOC*, pages 475–484. ACM, 2014. 5
- [Tsa22] Rotem Tsabary. Candidate witness encryption from lattice techniques. In *CRYPTO* (1), volume 13507 of *Lecture Notes in Computer Science*, pages 535–559. Springer, 2022. **4**, 9
- [VWW22] Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Witness encryption and nullio from evasive lwe. In *International Conference on the Theory and Application of Cryptology* and Information Security, pages 195–221. Springer, 2022. 4, 7, 9, 20, 21, 23, 24
- [Wee22] Hoeteck Wee. Optimal broadcast encryption and cp-abe from evasive lattice assumptions. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 217–241. Springer, 2022. 4
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In FOCS, pages 600–611, 2017. 9, 11
- [XXZ12] Xiang Xie, Rui Xue, and Rui Zhang. Deterministic public key encryption and identitybased encryption from lattices in the auxiliary-input setting. In *SCN*, volume 7485 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2012. 12
- [Zha16] Mark Zhandry. How to avoid obfuscation using witness PRFs. In *Theory of cryptography. Part II*, volume 9563 of *Lecture Notes in Comput. Sci.*, pages 421–448. Springer, Berlin, 2016. 3
- [Zha19] Mark Zhandry. The magic of elfs. Journal of Cryptology, 32:825–866, 2019. 3

A Generalization to Read-c MBPs

We first prove the lemma that allows us to represent a read-c MBP by a read-once MBP and ensures all invalid inputs are evaluated to zero.

Lemma A.1 (lemma 2.7 restated). Let Γ be a read-c MBP with width w, length h, and input length $\ell = h/c$. Let repeat(\mathbf{x}) = $\underbrace{\mathbf{x} | \mathbf{x} | \cdots | \mathbf{x}}_{c \text{ times}}$. Then there exists a read-once MBP Γ' with the following properties.

- 1. Γ' has width h + w and length h.
- 2. For all $\mathbf{x} \in \{0,1\}^{\ell}$, $\Gamma(\mathbf{x}) = \Gamma'(\operatorname{repeat}(\mathbf{x}))$.
- 3. For all invalid $\mathbf{x}' \in \{0,1\}^h$, i.e., $\mathbf{x}' \neq \text{repeat}(\mathbf{x})$ for any $\mathbf{x} \in \{0,1\}^\ell$, it holds that $\Gamma'(\mathbf{x}') = 0$.

In particular, if Γ computes a point function or all-zero function, then so is Γ' .

Proof. Let $\Gamma = \left\{ \mathbf{v} \in \{0,1\}^w, \left\{ \mathbf{M}_{i,b} \in \{0,1\}^{w \times w} \right\}_{i \in [h], b \in \{0,1\}} \right\}$ be a read-*c* MBP (and thus we omit ι). Let

$$\mathbf{V}^{(0)} \stackrel{\text{def}}{=} \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 1 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{pmatrix}, \mathbf{V}^{(1)} \stackrel{\text{def}}{=} \begin{pmatrix} 0 & 0 & 0 & \cdots & 0 & 0 \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \in \{0, 1\}^{c \times c}.$$

Claim A.2. For all $z \in \{0,1\}^c$, $\prod_{i \in [c]} \mathbf{V}^{(z_i)} = \mathbf{0}$ if and only if z is all-zero or all-one.

For $i \in [h], j \in [\ell], b \in \{0, 1\}$, define $U_{i,b,j} \in \{0, 1\}^{c \times c}$ as follows.

$$\mathbf{U}_{i,b,j} = \begin{cases} \mathbf{V}^b, & \text{if } i \equiv j \pmod{\ell}; \\ \mathbf{I}_c, & \text{otherwise.} \end{cases}$$

For $i \in [h], b \in \{0, 1\}$, set

$$\mathbf{M}'_{i,b} = \text{diag}(\mathbf{U}_{i,b,1}, \mathbf{U}_{i,b,2}, \dots, \mathbf{U}_{i,b,\ell}, \mathbf{M}_{i,b}) \in \{0,1\}^{(h+w) \times (h+w)}$$

Let $\mathbf{v}' = (1 \dots 1 | \mathbf{v}) \in \{0, 1\}^{h+w}$. We claim that the read-once MBP

$$\Gamma' = \left\{ \mathbf{v}', \left\{ \mathbf{M}'_{i,b} \in \{0,1\}^{(h+w) \times (h+w)} \right\}_{i \in [h], b \in \{0,1\}} \right\}$$

satisfies the said properties.

Note that $(\mathbf{v}')^{\top}\mathbf{M}'_{\mathbf{x}'} = (\mathbf{1}^{\top}\mathbf{U}^{(1)}|\cdots|\mathbf{1}^{\top}\mathbf{U}^{(\ell)}|\mathbf{v}^{\top}\mathbf{M}_{\mathbf{x}'})$, where $\mathbf{U}^{(j)} \stackrel{\text{def}}{=} \prod_{i \in [h]} \mathbf{U}_{i,x'_i,j}$. And for all $j \in [\ell]$

$$\mathbf{U}^{(j)} = \prod_{i \in [h]} \mathbf{U}_{i,x'_i,j} = \prod_{i \in [h]: i \equiv j \pmod{\ell}} \mathbf{U}_{i,x'_i,j} = \prod_{k=1}^c \mathbf{V}^{(x'_{(k-1)c+j})}.$$

Therefore, by the claim above, $\mathbf{1}^{\top}\mathbf{U}^{(j)} = \mathbf{0}$ if and only if

$$x'_{j} = x'_{j+\ell} = x'_{j+2\ell} = \dots = x'_{j+(c-1)\ell}$$

Since this holds for all $j \in [\ell]$, we have the desired properties.

Construction 7. $\Pi = (\Pi.\operatorname{Program}, \Pi.\operatorname{Eval})$. Input length ℓ_{in} and codomain $\mathcal{R}_{\lambda} = \mathbb{Z}_2^{n \times m}$. Let $h \stackrel{\text{def}}{=} c \cdot \ell_{\text{in}}, w' \stackrel{\text{def}}{=} w + h, \widehat{n}_0 \stackrel{\text{def}}{=} n, \widehat{n} \stackrel{\text{def}}{=} nw' + n$.

- $\mathsf{Program}(1^{\lambda}, C, \mathbf{Y} \in \mathbb{Z}_p^{n \times m}) \mapsto k_C$
 - 1. Parse *C* as a read-*c* MBP Γ . Let $\Gamma' = (\mathbf{v}', \{\mathbf{M}'_{i,b}\}_{i \in [h], b \in \{0,1\}})$ be the read-once MBP representation of Γ given by lemma 2.7.
 - 2. Sample $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n_1}}^{n \times n}$ for $i \in [h], b \in \{0, 1\}$ and set

$$\widehat{\mathbf{S}}_{1,b} = \begin{pmatrix} \mathbf{I}_n \mid (\mathbf{v}')^{\top} \mathbf{M}'_{1,b} \otimes \mathbf{S}_{1,b} \end{pmatrix}, \widehat{\mathbf{S}}_{i,b} = \begin{pmatrix} \mathbf{I}_n & \\ & \mathbf{M}'_{i,b} \otimes \mathbf{S}_{i,b} \end{pmatrix},$$

for $i = 2, \ldots, h, b \in \{0, 1\}$.

3. Sample $\mathbf{A}_h \leftarrow \mathbb{Z}_q^{\widehat{n} \times m}$ conditioned on $\left\lfloor \overline{\mathbf{A}_h} \right\rfloor_2 = \mathbf{Y}$ and output

$$k_c := \mathbf{C}_0, \mathbf{C}_1, \left\{\mathbf{D}_{i,b}\right\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \mathsf{GGH}.\mathsf{Encode}(\left\{\widehat{\mathbf{S}}_{i,b}\right\}_{i\in[h],b\in\{0,1\}}, \mathbf{A}_h).$$

• $\text{Eval}(k_C, \mathbf{x} \in \{0, 1\}^h) \mapsto \mathbf{Y} \in \mathbb{Z}_2^{n \times m}$. Parse $k_C = \mathbf{C}_0, \mathbf{C}_1, \{\mathbf{D}_{i,b}\}_{i=2,...,h,b \in \{0,1\}}$ and let $\mathbf{x}' = \text{repeat}(\mathbf{x}) \in \{0,1\}^h$. Output

$$\mathbf{Y} := \left\lfloor \mathbf{C}_{x_1'} \prod_{i=2}^h \mathbf{D}_{i,x_i'} \right\rceil_2$$

Analysis.

• <u>Correctness</u>. Since for all $\mathbf{x} \in \{0, 1\}^{\ell_{in}}$ and $\mathbf{x}' = \text{repeat}(\mathbf{x}) \in \{0, 1\}^{h}$, we have

$$\Gamma(\mathbf{x}) = 1 \iff \Gamma'(\mathbf{x}') = 1 \iff (\mathbf{v}')^\top \mathbf{M}'_{\mathbf{x}'} = \mathbf{0}.$$

Hence, correctness follows from the same calculation as in section 5.1.

- Security. Note that $\Gamma'^{-1}(1) = \Gamma^{-1}(1)$, and thus lemma 5.4, lemma 5.5 works perfectly for Γ' . Therefore, the argument in the proof of theorem 5.10 still goes, with Γ being replaced by Γ' .
- Efficiency. The only efficiency loss is that we need to augment the width from w to $w' = \overline{w + c \cdot \ell_{in}}$, which is still polynomial in λ . In the setting of parameters, we shall use w' in place of w.

B The Dual Attack for LWE

In this section, we recall the dual attack [BG14] for solving LWE with binary secret for certain parameters with slightly super-polynomial advantage than guessing. More concretely, suppose the modulus q is polynomial in n, $\mathbf{A} \in \mathbb{Z}_q^{r \times m}$, the secret is sampled from $\{0,1\}^r$, and the error bound is q^c for some c < 1, say c = 1/2. If we use BKZ with block-size $\log(n)$ to find a short kernel of \mathbf{A} , then we can compute a $r = (\log n)^2 / \log \log n$ dimensional secret in poly(n) time. This gives a poly(n) time attack for LWE where the secret space is slightly super-polynomial in n.

For BKZ to work, we choose a subset of m' LWE samples where

- 1. δ is the approximate root Hermite factor of BKZ with blocksize $\beta = \log n$. It satisfies $\log \delta = \Theta(\log \beta/\beta) = \Theta(\log \log n/\log n)$.
- 2. $m' = \sqrt{r \log q / \log \delta}$.
- 3. $2^{2\sqrt{r \log q \log \delta}} \ll q^{(1-c)}$, so that the *m*'-dimensional vector obtained by BKZ, denoted by **v**, satisfies $|\langle \mathbf{v}, \mathbf{e} \rangle| < q/4$, where **e** represents the LWE error term.
- 4. $\delta^r \ll q^{(1-c)}$ so that *r*-dimension LWE (with the same secret/error bound) can be solved by BKZ.

All conditions are satisfied with the choice $r = k \log q / \log \delta$ for some small constant k. This gives $r = k \log q / \log \delta \approx (\log n)^2 / \log \log n$.

One way of avoiding the attack is to set the error bound to be very large, say $|e| \in O(q)$ for each entry. This is the bound we choose for our candidate AIPO.

C Missing Proofs

C.1 UCE

Here we present the complete proof of theorem 4.3.

Theorem C.1 (Theorem 4.3 restated). Let Π , H be as in construction 1. Assume that there exists an AIPO in C, then $H \in UCE[S_1^{scup}]$.

Let $S \in S_1^{scup}$ be a source and D be a PPT distinguisher. Consider the following sequence of games with full description in fig. 5.

- 1. Game₀. Game₀ is the UCE game UCE^H_{S,D} when the hidden bit β is fixed to 1. We have moved the generation of k_{\emptyset} into HASH, which is just a conceptual change.
- 2. Game₁. Game₁ is identical to Game₀ except that
 - hk := k_{x^*} where $k_{x^*} \leftarrow \Pi$.Program $(1^{\lambda}, \mathsf{AIPO}(x^*), y)$.

 $Game_0 \approx_c Game_1$ readily follows from the privacy of Π (lemma C.2).

3. Game₂. Game₂ is identical to Game₁ except that HASH directly returns the programmed value instead of computing Π .Eval (k_{x^*}, x^*) . Also, we move the generation of k_{x^*} out of the oracle, so that the oracle behaves exactly as a random oracle. By the correctness of Π ,

$$|\mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_1 \Rightarrow 1]| = \mathtt{negl}(\lambda).$$

4. Game₃. Game₃ is identical to Game₂ except that hk is replaced by

 $k_{\emptyset} \leftarrow \Pi.\mathsf{Program}(msk,\mathsf{AIPO}(\mathsf{null}),y^*).$

We shall prove $Game_2 \approx_c Game_3$ via the security of AIPO (lemma C.3).

5. $Game_4$. $Game_4$ is identical to $Game_3$ except that

$$k_{\emptyset} \leftarrow \Pi.\mathsf{Program}(1^{\lambda},\mathsf{AIPO}(\mathsf{null}), y),$$

where $y \leftarrow \mathcal{R}_{\lambda}$ is a fresh random value (like in Game₀). Note that when programming on the all-zero function, the value y is hidden by the programmed key. Therefore, Game₄ \approx_c Game₃ readily follows from the value-hiding property of Π (lemma C.4).

6. Game₅. Finally, Game₅ is identical to Game₄ except that $hk = k_{\emptyset}$ is generated in the original way, removing AIPO:

$$k_{\emptyset} \leftarrow \Pi$$
.Program (msk, C_{\emptyset}, y) .

The two programmed keys are indistinguishable according to the privacy of Π , which is similar to is similar to Game₁ \approx_c Game₀. Hence,

$$|\mathbf{Pr}[\mathsf{Game}_5 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_4 \Rightarrow 1]| = \mathtt{negl}(\lambda).$$

Note that Game₅ is exactly the UCE game with $\beta = 0$. It remains to prove the following lemmas on game-hopping.

Lemma C.2. By the privacy of Π , $|\mathbf{Pr}[\mathsf{Game}_1 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_0 \Rightarrow 1]| = \mathsf{negl}(\lambda)$.

Proof. Consider the following adversary A that aims to break the privacy of Π .

- 1. Simulate $S^{(\cdot)}$ until *S* issues the oracle query x^* .
- 2. Submit the challenge $(C_0 = C_{\emptyset}, C_1 = \mathsf{AIPO}(x^*))$ and receive k^* from the challenger where k^* is generated in the following way: $b \leftarrow \{0, 1\}, y \leftarrow \mathcal{R}_{\lambda}, k^* \leftarrow \Pi.\mathsf{Program}(1^{\lambda}, C_b, y)$.
- 3. Forward $y := \Pi$.Eval (k^*, x^*) to S^{HASH} and get L.
- 4. Simulate $D(hk = k^*, L)$, and let β' denote the output of D.
- 5. Output $b' := \beta'$.

When b = 0, A simulates Game₀; when b = 1, A simulates Game₁. Therefore,

$$\Pr\left[\mathsf{Game}_1 \Rightarrow 1\right] - \Pr\left[\mathsf{Game}_0 \Rightarrow 1\right] = 2 \left| \operatorname{Priv}_{\Pi, \mathcal{A}}(1^{\lambda}) - 1/2 \right| = \operatorname{\mathsf{negl}}(\lambda).$$

Lemma C.3. By strong unpredicatbility of S and the security of AIPO,

 $|\mathbf{Pr}[\mathsf{Game}_3 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1]| = \mathtt{negl}(\lambda).$

Proof. Consider the adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ that aims to break the security of AIPO, as shown in fig. 7. The output of \mathcal{B}_1 is an unpredictable ensemble since *S* is an unpredictable source. Recall that the AIPO experiment AIPO_{AIPO,B}(1^{λ}) proceeds as follows:

(x, z) ← B₁(1^λ);
 b ← {0, 1};

- 3. $C_0 \leftarrow \mathsf{AIPO}(x), C_1 \leftarrow \mathsf{AIPO}(\mathsf{null})$
- 4. $b' \leftarrow \mathcal{B}_2(z, C_b)$; the experiment outputs 1 iff b = b'.

By the assumption that *S* is strongly unpredictable, B_1 is computationally unpredictable, and hence the security of the AIPO guarantees that

$$\left| \Pr\left[\operatorname{AIPO}_{\mathsf{AIPO},\mathcal{B}}(1^{\lambda}) \Rightarrow 1 \right] - 1/2 \right| = \operatorname{negl}(\lambda).$$

Note that conditioned on b = 0, the AIPO game is exactly Game₂, and conditioned on b = 1, it is exactly Game₃. Therefore,

$$\begin{aligned} |\mathbf{Pr}[\mathsf{Game}_3 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1]| &\leq 2 \left| \mathbf{Pr}\left[\mathrm{AIPO}_{\mathsf{AIPO},\mathcal{B}}(1^{\lambda}) \Rightarrow 1 \right] - 1/2 \right| \\ &= \mathtt{negl}(\lambda). \end{aligned}$$

${\cal B}_1(1^\lambda)$					
1:	$y^* \leftarrow \mathcal{R}_\lambda$				
2:	Simulate S^{Hash} with y^* as oracle answer				
3:	let x^\ast and L be the query and output of S respectively.				
4:	return $(x^*, z = (L, y^*))$				
$\mathcal{B}_2(z)$	(z, C)				
1:	Parse $z = (L, y^*)$				
2:	$k_C \leftarrow \Pi.Program(1^{\lambda}, C, y^*)$				
3:	return $D(hk = k_C, L)$				

Figure 7: Adversary \mathcal{B} for AIPO.

Lemma C.4. *By the value-hiding property of* Π *,*

 $|\mathbf{Pr}[\mathsf{Game}_4 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_3 \Rightarrow 1]| = \mathtt{negl}(\lambda).$

Proof. Consider the following adversary A that aims to break the value-hiding property of Π .

- 1. Simulate $S^{(\cdot)}$ until *S* issue the oracle query x^* and reply with $y_0 \leftarrow \mathcal{R}_{\lambda}$. Let *L* be the output of *S*.
- 2. Samples $y_1 \leftarrow \mathcal{R}_{\lambda}$ and $C \leftarrow \mathsf{AIPO}(\mathsf{null})$. Submit the challenge (C, y_0, y_1) and receive k^* from the challenger where k^* is generated as follows: $b \leftarrow \{0, 1\}, k^* \leftarrow \Pi.\mathsf{Program}(1^{\lambda}, C, y_b)$.
- 3. Simulate $D(hk = k^*, L)$, and let β' denote the output of D.
- 4. Output $b' := \beta'$.

When b = 0, A simulates Game₃; when b = 1, A simulate Game₄. Therefore,

$$\Pr\left[\mathsf{Game}_3 \Rightarrow 1\right] - \Pr\left[\mathsf{Game}_4 \Rightarrow 1\right]| = 2 \left| \mathsf{VH}_{\Pi,\mathcal{A}}(1^{\lambda}) - 1/2 \right| = \texttt{negl}(\lambda).$$

Т

C.2 MB-AIPO

Theorem C.5 (Theorem 4.4 restated). Construction 2 is an ℓ -bit MB-AIPO for strongly unpredictable distributions.

Proof. We start with proving correctness. Let (k, w, C) be the output of MBAIPO(x, m). By the correctness of AIPO, for $u \neq x$, then C(u) = 0 and hence $P[k, w, C](u) = \bot$. Meanwhile, $P[k, w, C](x) = \Pi$. Eval $(k, x) \oplus w = m$ by the correctness of Π and the relation $w = \Pi$. Eval $(k, x) \oplus m$.

Next, we prove MB-AIPO is value-hiding. Let $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ be an admissible adversary for MB-AIPO, i.e., \mathcal{B}_1 is a sampler for some strongly unpredictable distribution. We start with Game_0 in fig. 9, which is the experiment MBAIPO_{MBAIPO,B} when the hidden bit *b* is fixed to 0. Game_6 is the experiment MBAIPO_{MBAIPO,B} when the hidden bit *b* is fixed to 1. We shall prove $\mathsf{Game}_0 \approx_c \mathsf{Game}_6$ via a sequence of game-hoppings.

- 1. Game₁. Game₁ is identical to Game₀ except that we generate k by programming on C' where C' is output by a fresh execution of AIPO(x). We have Game₀ \approx_c Game₁ by the privacy of Π . Formally, Consider the following adversary A attacking the privacy of Π .
 - \mathcal{A} runs $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda}), C, C' \leftarrow \mathsf{AIPO}(x)$ and submits (C, C') to the challenger.
 - \mathcal{A} either receives k_{β} , where $k_0 \leftarrow \Pi$.Program $(1^{\lambda}, C, r)$, $k_1 \leftarrow \Pi$.Program $(1^{\lambda}, C', r)$ for some $r \leftarrow \{0, 1\}^{\ell}$, and $\beta \leftarrow \{0, 1\}$ is the challenge bit.
 - Then, it can retrieve $r' = \Pi$. Eval(k, x) and forwards $\beta' \leftarrow \mathcal{B}_2(z, P[k, r' \oplus m, C])$ as its guess of β .

By the correctness of Π , r = r' with overwhelming probability. Conditioned on r = r', if $\beta = 0$, \mathcal{A} perfectly simulate Game₀; if $\beta = 1$, \mathcal{A} perfectly simulate Game₁. Hence,

$$|\mathbf{Pr}[\mathsf{Game}_1 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_0 \Rightarrow 1]| = 2 \left| \operatorname{Priv}_{\Pi,\mathcal{A}}(1^{\lambda}) - \frac{1}{2} \right| = \operatorname{negl}(\lambda).$$

- 2. Game₂. Game₂ is identical to Game₁ except that C' is generate by $C' \leftarrow AIPO(1^{\lambda}, null)$. We prove Game₁ \approx_c Game₂ using the security of AIPO in lemma C.6.
- 3. Game₃. Game₂ is identical to Game₁ except that the programmed value is switched to a fresh random value m'. We have Game₂ \approx_c Game₃ by the value-hiding property of Π . Specifically, Consider the following adversary \mathcal{A}' attacking the value-hiding property of Π .
 - \mathcal{A}' generates $C' \leftarrow \mathsf{AIPO}(1^{\lambda}, \mathsf{null})$, samples $r, r' \leftarrow \{0, 1\}^{\ell}$, and submits (C', r, r') to the challenger.
 - \mathcal{A}' receives k_{β} , where $k_0 \leftarrow \Pi$.Program $(1^{\lambda}, C, r)$, $k_1 \leftarrow \Pi$.Program $(1^{\lambda}, C', r)$ for some $r \leftarrow \{0, 1\}^{\ell}$, and $\beta \leftarrow \{0, 1\}$ is the challenge bit.
 - \mathcal{A}' forwards $\beta' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C])$ as its guess of β .
 - If $\beta = 0$, \mathcal{A}' perfectly simulate Game₂; if $\beta = 1$, \mathcal{A}' perfectly simulate Game₃. Hence,

$$|\mathbf{Pr}[\mathsf{Game}_3 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1]| = 2 \left| \mathsf{VH}_{\Pi,\mathcal{A}'}(1^{\lambda}) - \frac{1}{2} \right| = \mathtt{negl}(\lambda).$$

- 4. Game₄. Game₄ is identical to Game₃ except that C' is generate by $C' \leftarrow \mathsf{AIPO}(x)$. Analogous to Game₁ \approx_c Game₂, we have Game₃ \approx_c Game₄.
- 5. Game₅. Game₅ is identical to Game₄ except that we generate k by programming on C. We have Game₄ \approx_c Game₅ by the privacy of Π , similar to Game₀ \approx_c Game₁.
- 6. Game₆. Game₆ is identical to Game₅ up to renaming variables: To get Game₆ from Game₅, one first replace *r* by *m*' and rename *r*' to be *r*.

It remains to prove the following lemma on game-hopping.

Lemma C.6. By the security of AIPO, we have

$$|\mathbf{Pr}[\mathsf{Game}_1 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1]| = \mathtt{negl}(\lambda).$$

Proof. Consider the adversary $\mathcal{D} = (\mathcal{D}_1, \mathcal{D}_2)$ that aims to break the security of AIPO as shown in fig. 8. we first show \mathcal{D} is an admissible AIPO adversary, i.e., \mathcal{D}_1 is an unpredictable distribution. Since $\mathcal{B}_1 = (Z_\lambda, X_\lambda, M_\lambda)$ is strongly unpredictable, by definition, $\mathcal{E} \stackrel{\text{def}}{=} (Z'_\lambda = (Z_\lambda, M_\lambda), X_\lambda)$ is unpredictable. Note that \mathcal{D}_1 is obtained by adding AIPO (X_λ) to the auxiliary input of \mathcal{E} . By the following claim, we assert that \mathcal{D}_1 is unpredictable.

Claim C.7. Let $\mathcal{E} = (Z'_{\lambda}, X'_{\lambda})$ be an unpredictable distribution. If we augment the auxiliary input in \mathcal{E} by adding $C \leftarrow \mathsf{AIPO}(X'_{\lambda})$ to Z'_{λ} , then the resulting distribution, denoted by $\widehat{\mathcal{E}}$, is still unpredictable.

Note that if the challenge bit *b* in the experiment AIPO_{AIPO,D} is 0, then *C*' is output by AIPO(*x*), and hence D perfectly simulate Game₁; otherwise, D perfectly simulate Game₂. Therefore,

$$\begin{aligned} |\mathbf{Pr}[\mathsf{Game}_1 \Rightarrow 1] - \mathbf{Pr}[\mathsf{Game}_2 \Rightarrow 1]| &\leq 2 \left| \mathbf{Pr} \left[\mathsf{AIPO}_{\mathsf{AIPO},\mathcal{D}}(1^{\lambda}) \Rightarrow 1 \right] - 1/2 \right| \\ &= \mathsf{negl}(\lambda). \end{aligned}$$

It remains to prove the claim above.

Proof of Claim. Let \mathcal{P} be a PPT adversary for breaking the unpredictability of $\widehat{\mathcal{E}}$. Let

$$\mathsf{adv}_{\mathcal{P}}(\lambda) \stackrel{\mathsf{def}}{=} \Pr_{(z',x) \leftarrow \widehat{\mathcal{E}}_{\lambda}} \left[\mathcal{P}(z') = x \right] = \Pr_{(z,x) \leftarrow \mathcal{E}_{\lambda}, C \leftarrow \mathsf{AIPO}(x)} \left[\mathcal{P}(z,C) = x \right].$$

Consider the adversary $\mathcal{T} = (\mathcal{T}_1 = \mathcal{E}, \mathcal{T}_2)$ for AIPO with \mathcal{T}_2 acts as follows: On input auxiliary information z and C, run $\mathcal{P}(z, C)$ and let x' be the output; if C(x') = 1, outputs 0; otherwise, output 1. Since $\mathcal{T}_1 = \mathcal{E}$ is unpredictable, \mathcal{T} is an admissible adversary. Note that

$$\Pr\left[\mathrm{AIPO}_{\mathsf{AIPO},\mathcal{T}}(\lambda) \Rightarrow 1\right] = \frac{1}{2} \cdot \mathsf{adv}_{\mathcal{P}}(\lambda) + \frac{1}{2}.$$

This is because if the challenge bit in the game AIPO_{AIPO, $\mathcal{T}(\lambda)$} is 0, \mathcal{T}_2 receives $C \leftarrow AIPO(x)$ from the challenger, and it guesses the challenge bit correctly (i.e., output 0) with probability $adv_{\mathcal{P}}(\lambda)$; in the other case, it receives $C \leftarrow AIPO(1^{\lambda}, null)$ outputs 1 with probability 1. Hence, by the security of AIPO, it must be that $adv_{\mathcal{P}}(\lambda) = negl(\lambda)$.

$\mathcal{D}_1(1^{\lambda})$		$\mathcal{D}_2(z',C')$		
1:	$(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$	1:	Parse $z' = (z, m, C)$	
2:	$C \leftarrow AIPO(x);$	2:	$r \leftarrow \{0,1\}^{\ell};$	
3:	$y^* \leftarrow \mathcal{R}_{\lambda};$	3:	$k \gets \Pi.Program(1^\lambda, C', r)$	
4:	z' := (z, m, C);	4:	return $\mathcal{B}_2(z, P[k, r \oplus m, C])$	
5:	return (z', x)			

Figure 8: Adversary \mathcal{D} for AIPO.

Game ₀ Game ₆	$Game_1$	$Game_2$
$1: (z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$	1: $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$	1: $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$
2: $r, m' \leftarrow \{0, 1\}^{\ell};$	2: $r \leftarrow \{0,1\}^{\ell};$	2: $r \leftarrow \{0,1\}^{\ell};$
3: $C \leftarrow AIPO(x);$	3: $C \leftarrow AIPO(x);$	3: $C \leftarrow AIPO(x);$
4: $k \leftarrow \Pi$.Program $(1^{\lambda}, C, r);$	$4: C' \leftarrow AIPO(x);$	$4: C' \leftarrow AIPO(1^{\lambda}, null);$
5: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$	5: $k \leftarrow \Pi.Program(1^{\lambda}, C', r);$	5: $k \leftarrow \Pi.Program(1^{\lambda}, C', r);$
6: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m', C])$:	$6: b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$	$6: b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$
7: return $[1 = b']$	7: return $\llbracket 1 = b' \rrbracket$	7: return $\llbracket 1 = b' \rrbracket$
$Game_3$	$Game_4$	$Game_5$
1: $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$	1: $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$	1: $(z, x, m) \leftarrow \mathcal{B}_1(1^{\lambda});$
2: $r, r' \leftarrow \{0, 1\}^{\ell};$	$2: r,r' \leftarrow \{0,1\}^{\lambda}$	2: $r, r' \leftarrow \{0, 1\}^{\lambda};$
3: $C \leftarrow AIPO(x);$	3: $C \leftarrow AIPO(x);$	3: $C \leftarrow AIPO(x);$
4: $C' \leftarrow AIPO(1^{\lambda}, null);$	4: $C' \leftarrow AIPO(x);$	$4: k \leftarrow \Pi.Program(1^{\lambda}, C, r');$
I = I = I		1/10 ($D[1 - 0]$)
5: $k \leftarrow \Pi.Program(1^{\wedge}, C', r')$	5: $k \leftarrow \Pi.\operatorname{Program}(1^{\wedge}, C', r');$	5: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$
5: $k \leftarrow \text{II.Program}(1^{\wedge}, C', r')$ 6: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$	5: $k \leftarrow \Pi.\operatorname{Program}(1^{\wedge}, C', r');$ 6: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$	5: $b' \leftarrow \mathcal{B}_2(z, P[k, r \oplus m, C]);$ 6: return $\llbracket 1 = b' \rrbracket$

Figure 9: Games in the proof of theorem 4.4.

C.3 Leakage-Resilient PKE

Definition C.8. Let $\mathsf{PKE} = (\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec})$ be a PKE scheme encrypting ℓ -bit message. For adversary $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$, consider the experiment $\mathsf{Expt}_{\mathsf{PKE}, \mathcal{A}}^{\mathsf{aux-ind-cpa}}$ defined in fig. 10. \mathcal{A} is admissible if for all PPT inverter \mathcal{I} ,

$$\Pr_{\substack{(\mathsf{pk},\mathsf{sk})\leftarrow\mathsf{Gen}(1^\lambda),z\leftarrow\mathcal{A}_0(\mathsf{sk})}}[\mathcal{I}(z)=\mathsf{sk}]=\mathtt{negl}(\lambda).$$

We say Π is *IND-CPA secure with hard-to-invert key-leakage* if for all admissible PPT adversary A, it holds that

$$\left| \Pr\left[\mathsf{Expt}_{\mathsf{PKE},\mathcal{A}}^{\mathsf{aux-ind-cpa}}(\lambda) \Rightarrow 1 \right] - \frac{1}{2} \right| = \mathtt{negl}(\lambda).$$

$$\begin{split} \hline & \mathsf{Expt}_{\mathsf{PKE},\mathcal{A}}^{\mathsf{aux-ind-cpa}}(\lambda) \\ \hline 1: \quad (\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{Gen}(1^{\lambda}) \\ 2: \quad z \leftarrow \mathcal{A}_0(\mathsf{sk}) \\ 3: \quad (m_0,\mathsf{st}) \leftarrow \mathcal{A}_1(\mathsf{pk},z) \\ 4: \quad m_1 \leftarrow \{0,1\}^{\ell(\lambda)} \\ 5: \quad b \leftarrow \{0,1\}, c \leftarrow \mathsf{Enc}(pk,m_b) \\ 6: \quad b' \leftarrow \mathcal{A}_2(pk,\mathsf{st},c) \\ 7: \quad \mathbf{return} \llbracket b = b' \rrbracket \end{split}$$

Figure 10: Experiment defining IND-CPA security under key leakage.

Theorem C.9 (Theorem 4.5 restated). Construction 3 is IND-CPA secure with hard-to-invert key-leakage.

Proof. We reduce the security to the security of MBAIPO in construction 2. Let $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2)$ be an admissible adversary. We construct an adversary $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ for MBAIPO as shown in fig. 11. First, Note that \mathcal{B}_1 is unpredictable since \mathcal{A} is admissible and the claim in the proof of lemma C.6;

$\mathcal{B}_1(1)$	$l^{\lambda})$	$\mathcal{B}_2(1)$	P[k, w, C], z' = (z, st))
1:	$x \leftarrow \{0,1\}^{\ell}$	1:	$b' \leftarrow \mathcal{A}_2(st, c = (k, w));$
2:	$z \leftarrow \mathcal{A}_0(x);$	2:	return b'
3:	$C' \gets AIPO(x)$		
4:	$(m, st) \leftarrow \mathcal{A}_1(C', z)$		
5:	return $(z' = (z, st), x, m)$		

Figure 11: Adversary \mathcal{D} for MB-AIPO.

hence, \mathcal{B} is an admissible adversary for MBAIPO. The game MBAIPO_{MBAIPO, $\mathcal{B}(\lambda)$} perfect simulates the experiment $\text{Expt}_{\mathsf{PKE},\mathcal{A}}^{\mathsf{aux-ind-cpa}}(\lambda)$, except that the $\mathsf{pk} = C'$ given to \mathcal{A}_1 and the C that is used for encryption are two independent outputs of $\mathsf{AIPO}(x)$. Here, we are in the same situation as in the proof of theorem 4.4 (i.e., from Game_0 to Game_1), and the proof is the same. Therefore, the advantage of \mathcal{A} translates to the advantage of \mathcal{B} against MBAIPO, and thus the advantage of \mathcal{A} must be negligible.

C.4 Value Uniqueness of Construction 6

We use A[i, j] and A[i :] to denote the (i, j) entry and the *i*-th row of **A** respectively.

Lemma C.10. Construction 6 satisfies value uniqueness provided that $m \ge 2n \log q$.

Proof. We will use the following lemma and prove it later.

Lemma C.11. Let $n \in \mathbb{N}$ and let q > n be a prime. For all $h \ge 1$, it holds that

$$\Pr_{\mathbf{S}_{i,b}\leftarrow\mathcal{D}_{2\sqrt{n}}^{n\times n}\text{ for }i\in[h],\,b\in\{0,1\}}\left[\forall \mathbf{x}\in\{0,1\}^{h} \ \mathbf{S}_{\mathbf{x}} \bmod q\neq\mathbf{0}\right]\geq 1-2^{h+1}\cdot n^{1-n/2}.$$

Suppose that

$$k_C = \mathbf{C}_0, \mathbf{C}_1, \left\{\mathbf{D}_{i,b}\right\}_{i=2,\dots,h,b\in\{0,1\}} \leftarrow \Pi.\mathsf{Program}(1^\lambda, C, \mathbf{Y}).$$

For all $\mathbf{x} \in \{0,1\}^h$, it holds (with high probability) that if C(x) = 0,

$$\mathsf{Eval}(k_C, \mathbf{x}) = \left[\mathbf{C}_{x_1} \prod_{i=2}^h \mathbf{D}_{i, x_i} \right]_2 = \left[\overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \underline{\mathbf{A}_h} \right]_2.$$

Recall that the programmed value is $|\overline{\mathbf{A}_h}|_{2'}$ and thus it suffices to show for any fixed $\overline{\mathbf{A}_h}$, with overwhelming probability over the choice of $\mathbf{S}_{i,b}$ and \mathbf{A}_h , it holds that

$$\forall \mathbf{x} \in \{0,1\}^h \ \mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \neq \mathbf{0} \bmod q \implies \lfloor \mathbf{W} \rceil_2 \neq \lfloor \overline{\mathbf{A}_h} \rceil_2, \tag{7}$$

where $\mathbf{W} \stackrel{\text{def}}{=} \overline{\mathbf{A}_h} + (\mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \otimes \mathbf{S}_{\mathbf{x}}) \cdot \underline{\mathbf{A}_h}$. Fix $\mathbf{x} \in \{0,1\}^h$ such that $\mathbf{z} \stackrel{\text{def}}{=} \mathbf{v}^\top \mathbf{M}_{\mathbf{x}} \neq 0$. WLOG, say $z_1 \neq 0 \mod q$. For $i \in [w]$, let $\mathbf{A}^{(i)} \in \mathbb{Z}_q^{n \times m}$ be the ((i-1)w+1)-th to the (iw)-th rows of $\underline{\mathbf{A}_h} \in \mathbb{Z}_q^{wn \times m}$. Then $\mathbf{W} = \overline{\mathbf{A}_h} + z_1 \mathbf{S}_{\mathbf{x}} \mathbf{A}^{(1)} + \cdots + z_w \mathbf{S}_{\mathbf{x}} \mathbf{A}^{(w)}$. By lemma C.11, $\mathbf{S}_{\mathbf{x}} \neq 0 \mod q$ with overwhelming probability. WLOG, say $\mathbf{S}_{\mathbf{x}}[1,1] \neq 0$. Then $\mathbf{W}[1,1] = \mathbf{S}_{\mathbf{x}}[1,1]\mathbf{A}^{(1)}[1,1] + v$ for some $v \in \mathbb{Z}$ that is independent of $\mathbf{A}^{(1)}[1,1]$. Thus

$$\Pr_{\mathbf{A}^{(1)}[1,1]\leftarrow\mathbb{Z}_q}\left[\left\lfloor\mathbf{W}[1,1]\right\rceil_2=\left\lfloor\overline{\mathbf{A}_h}[1,1]\right\rceil_2\right]\leq 2/3.$$

Note that $W[1,1], \ldots, W[1,m]$ are independent and the above argument applies to each W[1,j]. Hence,

$$\Pr_{\mathbf{A}^{(1)}}\left[\left\lfloor\mathbf{W}[1:]\right]_2 = \left\lfloor\overline{\mathbf{A}_h}[1:]\right]_2\right] \le \left(\frac{2}{3}\right)^m \le \left(\frac{3}{2}\right)^{2n\log q} \le q^{-1.16n}.$$

Finally, by a union bound over all **x**, eq. (7) holds for at most $2^h \cdot q^{-1.16n}$ fraction of $\underline{\mathbf{A}}_h$, which is a negligible probability in our parameter setting. This finishes the proof.

Proof of lemma C.11. Call a matrix *good* if every row is non-zero modulo q. It suffices to show that for all $h \ge 1$,

$$\Pr_{\left\{\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}\right\}_{i \in [h], b \in \{0,1\}}} \left[\forall \mathbf{x} \in \{0,1\}^h, \mathbf{S}_{\mathbf{x}} \text{ is good} \right] \ge 1 - 2^{h+1} \cdot n^{1-n/2}.$$
(8)

We shall prove eq. (8) by induction on *h*. The following claim is the crux of the proof.

Claim C.12. Let $\mathbf{T} \in \mathbb{Z}^{n \times n}$ be a good matrix. Then

$$\Pr_{\mathbf{S} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}} [\mathbf{TS} \text{ is not good}] \le n^{1-n/2}.$$

Proof of claim. Write $\mathbf{U} \stackrel{\text{def}}{=} \mathbf{TS}$. Since **T** is good, then there exists $\tau : [n] \to [n]$ such that $\mathbf{T}[i, \tau(i)] \mod q \neq 0$ for all $i \in [n]$.

Fix $i \in [n]$. For all $j \in [n]$,

$$\begin{split} & \Pr_{\mathbf{S} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}} \left[\mathbf{U}[i, j] \mod q = 0 \right] \\ &= \Pr_{\mathbf{S} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}} \left[\sum_{k \in [n]} \mathbf{T}[i, k] \mathbf{S}[k, j] \mod q = 0 \right] \\ &= \Pr_{\mathbf{S} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}} \left[\mathbf{T}[i, \tau(i)] \mathbf{S}[\tau(i), j] \equiv -\sum_{k \neq \tau(i)} \mathbf{T}[i, k] \mathbf{S}[k, j] \pmod{q} \right] \\ &= \sum_{s_1, \dots, s_{\tau(i)-1}, s_{\tau(i)+1}, \dots, s_n \leftarrow \mathcal{D}_{2\sqrt{n}}} \left[\mathcal{D}_{2\sqrt{n}} \left(\mathbf{T}[i, \tau(i)]^{-1} \cdot \left(-\sum_{k \neq \tau(i)} \mathbf{T}[i, k] s_k \right) + q \mathbb{Z} \right) \right] \\ &\leq \frac{1}{\sqrt{n}}. \end{split}$$

where $\mathbf{T}[i, \tau(i)]^{-1}$ is the multiplicative inverse in \mathbb{Z}_q . Here, the last inequality uses the fact that for all $a \in \mathbb{Z}$, $\mathcal{D}_{2\sqrt{n}}(a + q\mathbb{Z}) \leq \frac{1}{\sqrt{n}}$. Since $\mathbf{U}[i, 1], \ldots, \mathbf{U}[i, n]$ are independent, we have

$$\Pr_{\mathbf{S} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}} \left[\mathbf{U}[i:] \bmod q = \mathbf{0} \right] \le \frac{1}{\sqrt{n^n}}.$$

The claim follows from a union bound for all $i \in [n]$.

For h = 1, we apply the claim with $\mathbf{T} = \mathbf{I}_n$ to get

$$\Pr_{\mathbf{S}_{1,0},\mathbf{S}_{1,1}\leftarrow\mathcal{D}_{2\sqrt{n}}^{n\times n}}[\mathbf{S}_{1,0},\mathbf{S}_{1,1} \text{ is good}] \ge 1-2n^{1-n/2}.$$

Now assume $h \ge 2$. For any fixed $\{\mathbf{S}_{i,b}\}_{i\in[h-1],b\in\{0,1\}}$ such that $\mathbf{S}_{\mathbf{x}'}$ is good for all $\mathbf{x}' \in \{0,1\}^{h-1}$, by the claim and union bound over $\mathbf{x}' \in \{0,1\}^{h-1}$, we have

$$\Pr_{\mathbf{S}_{h,0},\mathbf{S}_{h,1}\leftarrow\mathcal{D}_{2\sqrt{n}}^{n\times n}}\left[\forall \mathbf{x}\in\{0,1\}^{h},\mathbf{S}_{\mathbf{x}}\text{ is good}\right]\geq 1-2^{h}\cdot n^{1-n/2}.$$
(9)

Let W_h denote the event '**S**_x is good for all $\mathbf{x} \in \{0, 1\}^{h}$ '. We have

$$\mathbf{Pr}[W_h] = \mathbf{Pr}[W_h \mid W_{h-1}] \, \mathbf{Pr}[W_{h-1}] \ge (1 - 2^h \cdot n^{1-n/2}) \cdot (1 - 2^h \cdot n^{1-n/2}) \\\ge 1 - 2^{h+1} \cdot n^{1-n/2},$$

where the probability is over $\mathbf{S}_{i,b} \leftarrow \mathcal{D}_{2\sqrt{n}}^{n \times n}$ for $i \in [h], b \in \{0, 1\}$. Here, the first inequality follows from eq. (9) and the induction hypothesis that $\Pr[W_{h-1}] \ge 1 - 2^h \cdot n^{1-n/2}$. This completes the proof.