

Constraint-Hiding Constrained PRFs for NC¹ from LWE

Ran Canetti^{1,2} and Yilei Chen¹

¹ Boston University
{canetti, chenyl}@bu.edu
² Tel Aviv University

Abstract. Constraint-hiding constrained PRFs (CHCPRFs), initially studied by Boneh, Lewi and Wu [PKC 2017], are constrained PRFs where the constrained key hides the description of the constraint. Envisioned with powerful applications such as searchable encryption, private-detectable watermarking and symmetric deniable encryption, the only known candidates of CHCPRFs are based on indistinguishability obfuscation or multilinear maps with strong security properties.

In this paper we construct CHCPRFs for all NC¹ circuits from the Learning with Errors assumption. The construction draws heavily from the graph-induced multilinear maps by Gentry, Gorbunov and Halevi [TCC 2015], as well as the existing lattice-based PRFs. In fact, our construction can be viewed as an instance of the GGH15 approach where security can be reduced to LWE.

We also show how to build from CHCPRFs reusable garbled circuits (RGC), or equivalently private-key function-hiding functional encryptions with 1-key security. This provides a different approach of constructing RGC from that of Goldwasser et al. [STOC 2013].

1 Introduction

Constrained PRFs [15,39,16] are pseudorandom functions with a special mode that outputs a constrained key defined by a predicate C . The constrained key CK_C preserves the functionality over the inputs x s.t. $C(x) = 1$, while leaving the function values on inputs x s.t. $C(x) = 0$ pseudorandom. In the standard formulation of constrained PRFs, the constrained key is not required to hide the predicate C . In fact, many constructions of constrained PRFs do reveal the constraint. A quintessential example is GGM’s puncturable PRF [33] where CK explicitly reveals the punctured points.

The notion of *constraint-hiding constrained PRF* (CHCPRF), proposed by Boneh, Lewi and Wu [12], makes the additional guarantee that the constraining predicate C remains hidden, even given the constrained key. Such an additional property allows the primitive to provide fairly natural constructions of searchable encryption, watermarking, deniable encryption, and others. However, they only propose candidates of CHCPRFs based on strong assumptions, like indistinguishability obfuscation (iO) or heuristic assumptions on candidate multilinear maps (multilinear-DDH or subgroup elimination).

This work. We further investigate the notion of CHCPRF, propose constructions based on standard cryptographic assumptions, and demonstrate more applications.

We first propose an alternative, simulation-based definition for CHCPRF. While for the cases addressed in our constructions the new style is (almost) equivalent to the indistinguishability-based one from [12], the new formulation provides a different viewpoint on the primitive.

Our main result is a construction of CHCPRF for all NC^1 circuit constraints based on the Learning with Errors (LWE) assumption [48]:

Theorem 1. *Assuming the intractability of LWE, there are CHCPRFs with 1-key simulation-based security, for all constraints recognizable by NC^1 circuits.*

The construction combines the graph-induced multilinear maps by Gentry, Gorbunov and Halevi [31], their candidate obfuscator, and the lattice-based PRFs of [6,11,5,19]. At the heart of our technical contribution is identifying a restricted (yet still powerful) variant of the GH15 maps, whose security can be reduced to LWE. This involves formulating new “LWE-hard” secret distributions that handle the permutation matrices underlying Barrington’s construction.

In addition, we construct function-hiding private-key functional encryptions (equivalently, reusable garbled circuits [34]) from CHCPRFs. This gives a construction of reusable garbled circuits from LWE that is very different from that of [34]:

Theorem 2. *For a circuit class \mathcal{C} , assuming 1-key simulation-based CHCPRFs for constraints in \mathcal{C} , and CPA secure private-key encryption whose decryption circuit is in \mathcal{C} , there exist 1-key secure reusable garbled circuits for \mathcal{C} .*

1.1 CHCPRFs, functional encryption and obfuscation

We propose a simulation-based definitional approach for CHCPRF, and compare this approach to the indistinguishability-based approach of Boneh et al [12].

Defining CHCPRFs. A constrained PRF consists of three algorithms: Master secret key generation, constrained key generation, and function evaluation. We first note that in order to have hope to hide the constraint, the function evaluation algorithm should return a random-looking value v even if evaluated on a constrained input x , as opposed to returning \perp as in the standard formulation. Furthermore, we require that the value of the original function on x remains pseudorandom even given the constrained key and the value v .

The definition of CHCPRF is aimed at capturing three requirements: (1) the constrained keys preserve functionality on inputs that do not match the constraint; (2) the function values at constrained points remain pseudorandom given the constrained key; (3) the constrained key does not reveal any information on the constraining function.

Boneh et al [12] give a number of indistinguishability-based definitions that vary in strength, depending on the level of adaptivity of the adversary in choosing

the constraints and evaluation points, as well as on the number of constrained keys that the adversary is allowed to see. We take an alternative approach and give a simulation-based definition. We also compare the definitions, and show equivalence and derivations in a number of cases.

Here is a sketch of the non-adaptive single-key variant of our simulation-based definition. The definition captures all three requirements via a single interaction: We require that, for any polytime adversary, there exists a polytime simulator such that the adversary can distinguish between the outcomes of the following two experiments only with negligible probability:

- In the real experiment, the system first generates a master secret key K . The adversary can then query a constraint circuit C and many inputs $x^{(1)}, \dots, x^{(t)}$. In return, it obtains $\text{CK}_C, x^{(1)}, \dots, x^{(t)}, y^{(1)}, \dots, y^{(t)}$, where CK_C is a key constrained by C , and $y^{(i)}$ is the result of evaluating the original, unconstrained function with key K at point $x^{(i)}$. (This is so regardless of whether $x^{(i)}$ meets the constraint or not.)
- In the ideal experiment, the simulator samples a master secret key K^S . Once received a constraint query, the simulator obtains only the description length of C and creates a simulated constrained key CK^S . Once received input queries $x^{(1)}, \dots, x^{(t)}$, the simulator also t indicator bits $d^{(1)}, \dots, d^{(t)}$ where the $d^{(i)}$ denotes whether $x^{(i)}$ is in the constraint, and generates simulated values $y^{(1)S}, \dots, y^{(t)S}$. If $d^{(i)} = 0$, then the simulated $y^{(i)S}$ is uniformly random. The output of the experiment is $\text{CK}^S, x^{(1)}, \dots, x^{(t)}, y^{(1)S}, \dots, y^{(t)S}$.

Secret-key functional encryption from simulation-based CHCPRFs. We sketch our construction of functional encryption from CHCPRFs. Functional encryption [13] allows the evaluator, given a functional decryption key, to learn the value of the function applied to encrypted data without learning anything else. With CHCPRFs in hand, it is rather simple to construct a private-key functional encryption scheme that is both function-private and input-private. Our functional encryption scheme proceeds as follows:

- Key generation: The master key for the scheme is a key K for a CHCPRF, and a key SK for a CPA-secure symmetric encryption scheme (Enc, Dec).
- Encrypt a message m : $\text{CT} = (c, t)$, where $c = \text{Enc}_{\text{SK}}(m)$, and $t = \text{CHCPRF}_K(c)$.
- Functional decryption key: The functional decryption key for a binary function f is a constrained-key $\text{CK}_{\hat{f}}$ for the function $\hat{f}(c) = f(\text{Dec}_{\text{SK}}(c))$. That is, \hat{f} has SK hardwired; it decrypts its input c and applies f to the plaintext.
- Functional decryption: Given ciphertext $\text{CT} = (c, t)$ and the constrained decryption key $\text{CK}_{\hat{f}}$, output 1 if $t = \text{CHCPRF}_{\text{CK}_{\hat{f}}}(c)$, 0 otherwise.

Correctness of decryption follows from the correctness and constrainability of the CHCPRF, and secrecy follows from the constraint-hiding property.

This construction is conceptually different from the previous construction [34]. In particular, the size of ciphertext (for 1-bit output) is the size of a symmetric encryption ciphertext plus the security parameter, independent of the depth of the circuit.

Two-key CHCPRFs imply obfuscation. It is natural to consider an extension of the CHCPRF definition to the case where the adversary may obtain multiple constrained keys derived from the same master key. Indeed in [12] some applications of this extended notion are presented.

We observe that this extended notion in fact implies full fledged program obfuscation: To obfuscate a circuit C , choose a key K for a CHCPRF, and output two constrained keys: The constrained key $\text{CK}[C]$, and the constrained key $\text{CK}[I]$, where I is the circuit that always outputs 1. To evaluate $C(x)$ check whether $\text{CHCPRF}_{\text{CK}[C]}(x) = \text{CHCPRF}_{\text{CK}[I]}(x)$.

Again, correctness of evaluation follows from the correctness and constrainability of the CHCPRF. The level of security for the obfuscation depends on the definition of CHCPRF in use. Specifically, the natural extension of the above simulation-based definition to the two-key setting implies that the above simple obfuscation method is VBB (which in turn means that the known impossibility results for VBB obfuscation carry over to this variant of two-key CHCPRF). The indistinguishability-based definition of [12] implies that the above obfuscation method is IO.

1.2 Overview of our construction

Our construction of CHCPRFs draws heavily from the multilinear maps by Gentry, Gorbunov and Halevi [31], and the lattice-based PRFs of Banerjee, Peikert and Rosen and others [6,11,5,19]. We thus start with a brief review of the relevant parts of these works.

Recap GGH15. The GGH15 multilinear encoding is depicted by a DAG that defines the rule of homomorphic operations and zero-testing. For our purpose it is sufficient to consider the following special functionality (which corresponds to a graph of ℓ nodes and two parallel edges from node i to node $i+1$, see Figure 1.1a): We would like to encode $2\ell + 1$ secrets $s_1^0, s_1^1, \dots, s_\ell^0, s_\ell^1, s_T$ over some finite group G , in such a way that an evaluator who receives the encodings can test, for any given $x \in \{0, 1\}^\ell$, whether $s_T = \prod_{i=1}^\ell s_i^{x_i}$, and at the same time the encodings hide “everything else” about the secrets. (Indeed, “everything else” might have different meanings in different contexts.)

To do that, GGH15 take the group G to actually be a ring R_q , where R denotes the base ring (typical choices include $R = \mathbb{Z}^{n \times n}$ or $R = \mathbb{Z}[x]/(\Phi_n(x))$, where n is a parameter related to the lattice dimension, and Φ_n is the n^{th} cyclotomic polynomial), and q is the modulus. The encoder then samples $\ell + 1$ hard Ajtai-type matrices $\{\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_\ell, \mathbf{A}_{\ell+1} \leftarrow R_q^{1 \times m}\}$ with trapdoors [2,3,44], and associates each matrix with the corresponding node of the graph. These matrices and their trapdoors are treated as (universal) public and secret parameters, respectively. We refer to the indices $1 \dots \ell + 1$ as *levels*.

The 2ℓ secrets are associated with the 2ℓ edges of the graph in the natural way. Encoding a secret s_i^b is done in two steps: First create an LWE sample for the secret s_i^b under the matrix \mathbf{A}_{i+1} , namely $\mathbf{Y}_i^b = s_i^b \mathbf{A}_{i+1} + \mathbf{E}_i^b$. Next, sample a preimage \mathbf{D}_i^b of \mathbf{Y}_i^b under the matrix \mathbf{A}_i , using the trapdoor of \mathbf{A}_i .

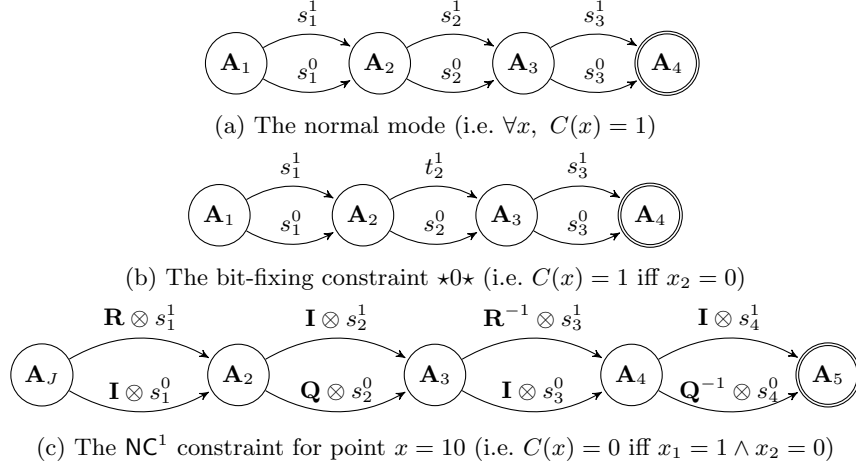


Fig. 1.1: Examples of the GGH15-based PRFs

That is, $\mathbf{A}_i \mathbf{D}_i^b = \mathbf{Y}_i^b$ and \mathbf{D}_i^b is sampled from discrete Gaussian distribution of small width. The encoder then lets \mathbf{D}_i^b be the encoding of s_i^b . The encoding T of s_T , where $s_T = \prod_{i=1}^{\ell} s_i^{x_i}$ for some $x \in \{0, 1\}^{\ell}$, is defined as $T = F(x)$, where $F(x) = \mathbf{A}_1 \prod_{i=1}^{\ell} \mathbf{D}_i^{x_i}$. Finally, the values $\mathbf{A}_1, \mathbf{D}_1^0, \mathbf{D}_1^1, \dots, \mathbf{D}_{\ell}^0, \mathbf{D}_{\ell}^1, T$ are given to the evaluator. To test a given $x' \in \{0, 1\}^{\ell}$, the evaluator computes $F(x')$ and checks whether $F(x') - T$ is a matrix with small entries.

To see why this works out functionality-wise consider the following equation:

$$F(x) = \mathbf{A}_1 \prod_{i=1}^{\ell} \mathbf{D}_i^{x_i} = \prod_{i=1}^{\ell} s_i^{x_i} \mathbf{A}_{\ell+1} + \underbrace{\sum_{i=1}^{\ell} \left(\prod_{j=1}^{i-1} s_j^{x_j} \cdot \mathbf{E}_i^{x_i} \cdot \prod_{k=i+1}^{\ell} \mathbf{D}_k^{x_k} \right)}_{\mathbf{E}_x} \pmod{q}. \quad (1)$$

Indeed, if the secrets s_i^b are set with small norm, then the entire \mathbf{E}_x term can be viewed as a small error term, so the dominant factor, $\prod_{i=1}^{\ell} s_i^{x_i} \mathbf{A}_{\ell+1}$, will be purely determined by the multiplicative relationship of the secrets. As for security, observe that the encoding \mathbf{D}_i^b of each secret s_i^b amounts to an LWE encoding of s_i^b , and furthermore the encoding of $s_x = \prod_{i=1}^{\ell} s_i^{x_i}$ is also in the form of an LWE instance $\mathbf{A}_{\ell+1}, \prod_{i=1}^{\ell} s_i^{x_i} \mathbf{A}_{\ell+1} + \mathbf{E}_x \pmod{q}$. Of course, being in the form of LWE does not amount to a clear security property that is based on LWE. We discuss this point further below.

The power and danger in the GGH15 approach. The GGH15 encoding embeds the plaintext s into the *secret* term of the LWE instance, unlike in other LWE-based systems (e.g. Regev [48] or dual-Regev [32]) where the plaintext is associated with the error term or the \mathbf{A} matrix. While the graph structure and trapdoor sampling mechanism enables homomorphic evaluations on the LWE

secrets, analyzing the security becomes tricky. Unlike the traditional case where the LWE secrets s are independent and random, here the LWE secrets, representing plaintexts, are taken from distributions that are potentially structured or correlated with each other.

Such dependencies make it hard to prove security of the trapdoor sampling: Recall that the encoding \mathbf{D}_i of some secret \hat{s}_i (possibly obtained from an evaluation over correlated secrets) is the preimage of $\mathbf{Y}_i := \hat{s}_i \mathbf{A}_{i+1} + \mathbf{E}$ sampled by the trapdoor of \mathbf{A}_i . For instance, in the extreme case where $\hat{s}_i = 0$, then the public encoding \mathbf{D}_i becomes a “weak trapdoor” of \mathbf{A}_i , which endangers the secrets encoded on the edges heading to \mathbf{A}_i [31].

Consequently, to safely use the GGH15 encoding, one has to consider the joint distribution of all the LWE secrets s_i^b , and demonstrate that the trapdoor sampling algorithm remains secure even with respect to these secrets. We demonstrate how to do that in a specific setting, by showing that there exists a “simulated” way to sample the encodings without knowing the secrets or trapdoors, and the resulting sample is indistinguishable from the real one.

LWE-based PRFs. The example of the “subset product” type encoding may remind the readers of the lattices-based pseudorandom functions [6,11,5,19]. Indeed, recall the basic construction of Banerjee et al [6, Section 5.1]. For modulus $2 \leq p < q$ chosen such that q/p is exponential in the input length ℓ . The secret keys of the PRF are exactly 2ℓ LWE secrets $s_1^0, s_1^1, \dots, s_\ell^0, s_\ell^1$ and a uniform matrix \mathbf{A} over R_q . To evaluate, compute $F(x) = \lfloor \prod_{i=1}^{\ell} s_i^{x_i} \mathbf{A} \rfloor_p$ where $\lfloor v \rfloor_p$ means multiplying v by p/q and rounding to the nearest integer. Rounding plays a crucial role in the security proof, since it allows to add fresh small noise terms without changing the functionality whp, hence one can inductively obtain fresh LWE instances on any level.

Our construction for bit-fixing constraints. A bit-fixing constraint is specified by a string $\mathbf{c} \in \{0, 1, \star\}^\ell$, where 0 and 1 are the matching bits and \star denotes the wildcards. The constrain predicate C outputs 1 if the input matches \mathbf{c} .

The combination of GGH15 and lattice-based PRFs inspires us to construct CHCPRFs for bit-fixing constraints. In fact, after rounding $F(x)$ in Equation (1), the functionality of $\lfloor F(x) \rfloor$ is equivalent to (up to the rounding error) both the BPR PRF [6, Section 5.1] and a variant of the PRF in [11, Section 5.1]. If we take the 2ℓ LWE secrets $s_1^0, s_1^1, \dots, s_\ell^0, s_\ell^1$ as master secret key, the encodings $\mathbf{A}_1, \mathbf{D}_1^0, \mathbf{D}_1^1, \dots, \mathbf{D}_\ell^0, \mathbf{D}_\ell^1$ as the evaluation key in the normal mode. An intuitive constraining algorithm is simply replacing the LWE secret of the constrained bit with an independent random element t , and reproduce its encoding \mathbf{D}_t . As an example, Figure 1.1a and Figure 1.1b illustrate the normal mode and constrained mode of a bit-fixing PRF.

We show that the key and the outputs from both the normal mode and the constrained mode (both modes use trapdoor sampling) are indistinguishable from an oblivious sampling procedure without using the trapdoors. The proof proceeds level-by-level (from level ℓ to level 1). Within each level i , there are

two steps. The first step uses the computational hardness of LWE: observe that the LWE samples associated on \mathbf{A}_{i+1} are with independent secrets, and \mathbf{A}_{i+1} is trapdoor-free in that hybrid distribution by induction, so the LWE samples are indistinguishable from uniformly random. The second step uses a statistical sampling lemma by Gentry, Peikert and Vaikuntanathan [32], which says the preimage of uniform outputs can be sampled without using the trapdoor of \mathbf{A}_i . The proof strategy is first illustrated by Brakerski et al. where they construct an evasive conjunction obfuscator from GGH15 [20].

We note that this construction and analysis imply that a variant of the PRF from [11] also satisfies 1-key bit-fixing constraint hiding. Although the PRF from [11] does not involve the trapdoor sampling procedure and is much simpler as a bit-fixing CHCPRF, understanding the GGH15-based version is beneficial for understanding the CHCPRF for NC^1 coming next.

Embedding a general constraint in the PRF keys. We move on towards embedding a general constraint in the key. Consider in particular the task of puncturing the key at a single point without revealing the point, which is essential to the applications like watermarking and deniable encryption mentioned in [12]. Indeed, even that simple function seems to require some new idea.

To preserve the graph structure while handling general constraints, Barrington’s Theorem [8] comes into the picture. Recall that Barrington’s Theorem converts any depth- d Boolean circuits into an oblivious branching program of length $z \leq 4^d$ composed of permutation matrices $\{\mathbf{B}_i^b\}_{b \in \{0,1\}, i \in [z]}$ of dimension w (by default $w = 5$). Evaluation is done via multiplying the matrices selected by input bits, with the final output $\mathbf{I}^{w \times w}$ or a w -cycle \mathbf{P} recognizing 1 or 0 respectively.

To embed permutation matrices in the construction, we set the secret term

for the normal mode as $\mathbf{S}_i^b = \mathbf{I}^{w \times w} \otimes s_i^b = \begin{bmatrix} s_i^b & 0 \\ & \ddots \\ 0 & s_i^b \end{bmatrix}$ (where \otimes is the tensor

product operator); in the constrained mode as $\mathbf{S}_i^b = \mathbf{B}_i^b \otimes s_i^b$. This provides the functionality of constraining all NC^1 circuits. See Figure 1.1c for an example of 2-bit point constraint $x_1 x_2 \in \{0, 1\}^2$, where x_1 controls the 1st and 3rd branches, x_2 controls the 2nd and 4th branches, \mathbf{Q} and \mathbf{R} represent different w -cycles.

We then analyze whether the permutation matrix structures are hidden in the constrained key, and whether the constrained outputs are pseudorandom. The first observation is that the tensor product of a permutation matrix \mathbf{B} and any hard LWE secret distribution s forms a hard LWE distribution, i.e. \mathbf{A} , $(\mathbf{B} \otimes s) \cdot \mathbf{A} + \mathbf{E}$ is indistinguishable from uniformly random. This means both the secret and the permutation matrices are hidden in the constrained key.

Still, the rounded constrained output $\left[(\mathbf{P} \otimes \prod_{i=1}^{\ell} s_i^{x_i}) \cdot \mathbf{A}_{z+1} \right]$ is a fixed permutation of the original value. so the adversary can left-multiply \mathbf{P}^{-1} to obtain the original output. To randomize the constrained outputs, we adapt the “book-end” idea from the GGH15 candidate obfuscator. That is, we multiply the output

on the left by a small random vector $\mathbf{J} \in R^{1 \times w}$. By a careful reduction to standard LWE, one can show that $\mathbf{A}, \mathbf{J}\mathbf{A} + \mathbf{E}, \mathbf{J}(\mathbf{P} \otimes 1_R)\mathbf{A} + \mathbf{E}'$ is indistinguishable from uniformly random.

With these two additional hard LWE distributions in the toolbox, we can base NC^1 CHCPRF on LWE via the same two-step proof strategy (i.e. LWE+GPV in each level) used in the bit-fixing construction.

1.3 More on related work

More background on multilinear maps and the implication of this work. The notion of cryptographic multilinear maps was introduced by Boneh and Silverberg [14]. Currently there are three main candidates [30,26,31], with a number of variants. However, what security properties hold for the candidates remains unclear. In particular, none of the candidates is known to satisfy the multilinear DDH or subgroup elimination assumptions that are sufficient for the CHCPRFs by Boneh et al [12] (see [30,24,37,25] for the attacks on these assumptions).

Note that even our result does not imply that GGH15 satisfies the traditional assumptions like multilinear DDH, but at least it demonstrates a safe setting. To what extent can the safe setting be generalized remains an open problem. Indeed, a central task in the study of the existing candidate multilinear maps is to identify settings where they can be used based on standard cryptographic assumptions [36].

Relations to the GGH15 candidate program obfuscator. Our construction for NC^1 constraints is strongly reminiscent of the candidate obfuscator from GGH15 [31, Section 5.2]. In particular, the “secrets” in the CHCPRF corresponds to the “multiplicative bundling scalars” from the GGH15 obfuscator. Under the restriction of releasing only 1 branch (either the functional branch or the dummy branch), our result implies that the “scalars” and permutation matrices can be hidden (without using additional safeguards such as the Kilian-type randomization and padded randomness on the diagonal).

In contrast, the recent cryptanalysis of the GGH15 obfuscator [23] shows that when releasing both the functional key and the dummy key, one can extract the bundling scalars even if the obfuscator is equipped with all the safeguards.

It might be instructive to see where our reduction to LWE fail if one attempts to apply our proof technique to the two-key setting. The point is that in this case, the adversary obtains LWE samples \mathbf{Y}, \mathbf{Y}' with correlated secrets; Therefore it is not clear how to simulate the Gaussian samples of \mathbf{D} conditioned on $\mathbf{A}\mathbf{D} = \mathbf{Y}$ or of \mathbf{D}' conditioned on $\mathbf{A}'\mathbf{D}' = \mathbf{Y}'$, without knowing the trapdoors of \mathbf{A} and \mathbf{A}' .

1.4 Concurrent work

In an independent work, Boneh, Kim and Montgomery [10] build CHCPRF from LWE, for the special case of input puncturing constraints. Their construction is

very different from ours. In particular, their starting point is the (non-hiding) constrained PRF by Brakerski and Vaikuntanathan [19].

While they analyze their construction with respect to the indistinguishability-based definition, they also consider a simulation-based definition that is significantly stronger than the one here. They show that it is impossible to realize that definition for general functions. To do that, they use the same construction of functional encryption from CHCPRFs as the one presented here.

2 Preliminaries

Notations and terminology. Let $\mathbb{R}, \mathbb{Z}, \mathbb{N}$ be the set of real numbers, integers and positive integers. The notation R is often used to denote some base ring. The concrete choices of R are $\mathbb{Z}^{n \times n}$ (the integer matrices) and $\mathbb{Z}[x]/(x^n + 1)$ (where n is a power of 2). We denote $R/(qR)$ by R_q . The rounding operation $[a]_p : \mathbb{Z}_q \rightarrow \mathbb{Z}_p$ is defined as multiplying a by p/q and rounding the result to the nearest integer.

For $n \in \mathbb{N}$, $[n] := \{1, \dots, n\}$. A vector in \mathbb{R}^n is represented in column form, and written as a bold lower-case letter, e.g. \mathbf{v} . For a vector \mathbf{v} , the i^{th} component of \mathbf{v} will be denoted by v_i . A matrix is written as a bold capital letter, e.g. \mathbf{A} . The i^{th} column vector of \mathbf{A} is denoted \mathbf{a}_i .

The length of a vector is the ℓ_p -norm $\|\mathbf{v}\|_p = (\sum v_i^p)^{1/p}$. The length of a matrix is the norm of its longest column: $\|\mathbf{A}\|_p = \max_i \|\mathbf{a}_i\|_p$. By default we use ℓ_2 -norm unless explicitly mentioned. When a vector or matrix is called “small” (or “short”), we refer to its norm (resp. length). The thresholds of “small” will be precisely parameterized in the article and are not necessary negligible functions.

2.1 Matrix branching programs

Definition 1 (Matrix branching programs). A width- w , length- z matrix branching program over ℓ -bit inputs consists of an index-to-input map, a sequence of pairs of matrices \mathbf{B}_i^b , and a non-identity matrix \mathbf{P} representing 0: $\text{BP} = \{\iota : [z] \rightarrow [\ell], \{\mathbf{B}_i^b \in \{0, 1\}^{w \times w}\}_{i \in [z], b \in \{0, 1\}}, \mathbf{P} \in \{0, 1\}^{w \times w} \setminus \{\mathbf{I}\}\}$. The program computes the function $f_{\text{BP}} : \{0, 1\}^\ell \rightarrow \{0, 1\}$, defined as

$$f_{\text{BP}}(x) = \begin{cases} 1 & \text{if } \prod_{i \in [z]} \mathbf{B}_i^{x_{\iota(i)}} = \mathbf{I} \\ 0 & \text{if } \prod_{i \in [z]} \mathbf{B}_i^{x_{\iota(i)}} = \mathbf{P} \\ \perp & \text{elsewhere} \end{cases}$$

A set of branching programs $\{\text{BP}\}$ is called **oblivious** if all the programs in the set have the same index-to-input map ι .

Theorem 3 (Barrington’s theorem [8]). For $d \in \mathbb{N}$, and for any set of depth- d fan-in-2 Boolean circuits $\{C\}$, there is an oblivious set of width-5 length- 4^d branching programs $\{\text{BP}\}$ with a index-to-input map ι , where each BP is composed of permutation matrices $\{\mathbf{B}_i^b \in \{0, 1\}^{5 \times 5}\}_{i \in [z], b \in \{0, 1\}}$, a 5-cycle \mathbf{P} , and ι .

2.2 Lattices

An n -dimensional lattice Λ is a discrete additive subgroup of \mathbb{R}^n . Given n linearly independent basis vectors $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^n\}$, the lattice generated by \mathbf{B} is $\Lambda(\mathbf{B}) = \Lambda(\mathbf{b}_1, \dots, \mathbf{b}_n) = \left\{ \sum_{i=1}^n x_i \cdot \mathbf{b}_i, x_i \in \mathbb{Z} \right\}$. We have the quotient group \mathbb{R}^n / Λ of cosets $\mathbf{c} + \Lambda = \{\mathbf{c} + \mathbf{v}, \mathbf{v} \in \Lambda\}$, $\mathbf{c} \in \mathbb{R}^n$. Let $\tilde{\mathbf{B}}$ denote the Gram-Schmidt orthogonalization of \mathbf{B} .

Gaussian on lattices. For any $\sigma > 0$, define the Gaussian function on \mathbb{R}^n centered at \mathbf{c} with parameter σ :

$$\forall \mathbf{x} \in \mathbb{R}^n, \rho_{\sigma, \mathbf{c}}(\mathbf{x}) = e^{-\pi \|\mathbf{x} - \mathbf{c}\|^2 / \sigma^2}$$

For any $\mathbf{c} \in \mathbb{R}^n$, $\sigma > 0$, and n -dimensional lattice Λ , define the discrete Gaussian distribution over Λ as:

$$\forall \mathbf{x} \in \Lambda, D_{\Lambda + \mathbf{c}, \sigma}(\mathbf{x}) = \frac{\rho_{\sigma, \mathbf{c}}(\mathbf{x})}{\rho_{\sigma, \mathbf{c}}(\Lambda)}$$

Lemma 1 ([47,45]). *Let \mathbf{B} be a basis of an m -dimensional lattice Λ , and let $\sigma \geq \|\tilde{\mathbf{B}}\| \cdot \omega(\log n)$, then $\Pr_{\mathbf{x} \leftarrow D_{\Lambda, \sigma}}[\|\mathbf{x}\| \geq \sigma \cdot \sqrt{m} \vee \mathbf{x} = \mathbf{0}] \leq \text{negl}(n)$.*

Gentry, Peikert and Vaikuntanathan [32] show how to sample statistically close to discrete Gaussian distribution in polynomial time for sufficiently large σ (the algorithm is first proposed by Klein [40]). The sampler is upgraded in [18] so that the output is distributed exactly as a discrete Gaussian.

Lemma 2 ([32,18]). *There is a p.p.t. algorithm that, given a basis \mathbf{B} of an n -dimensional lattice $\Lambda(\mathbf{B})$, $\mathbf{c} \in \mathbb{R}^n$, $\sigma \geq \|\tilde{\mathbf{B}}\| \cdot \sqrt{\ln(2n+4)}/\pi$, outputs a sample from $D_{\Lambda + \mathbf{c}, \sigma}$.*

We then present the trapdoor sampling algorithm and the corollary of GPV lemma in the general ring R .

Lemma 3 ([2,3,44]). *There is a p.p.t. algorithm $\text{TrapSam}(R, 1^n, 1^m, q)$ that, given the base ring R , modulus $q \geq 2$, lattice dimension n , and width parameter m (under the condition that $m = \Omega(\log q)$ if $R = \mathbb{Z}^{n \times n}$, $m = \Omega(n \log q)$ if $R = \mathbb{Z}[x]/(x^n + 1)$), outputs $\mathbf{A} \leftarrow U(R_q^{1 \times m})$ with a trapdoor τ .*

Lemma 4 ([32]). *There is a p.p.t. algorithm $\text{PreimgSam}(\mathbf{A}, \tau, \mathbf{y}, \sigma)$ that with all but negligible probability over $(\mathbf{A}, \tau) \leftarrow \text{TrapSam}(R, 1^n, 1^m, q)$, for sufficiently large $\sigma = \Omega(\sqrt{n \log q})$, the following distributions are statistically close:*

$$\{\mathbf{A}, \mathbf{x}, \mathbf{y} : \mathbf{y} \leftarrow U(R_q), \mathbf{x} \leftarrow \text{PreimgSam}(\mathbf{A}, \tau, \mathbf{y}, \sigma)\} \approx_s \{\mathbf{A}, \mathbf{x}, \mathbf{y} : \mathbf{x} \leftarrow \gamma_\sigma, \mathbf{y} = \mathbf{A}\mathbf{x}\}$$

where γ_σ represents $D_{\mathbb{Z}_{nm}^{1 \times n}, \sigma}$ if $R = \mathbb{Z}^{n \times n}$; represents $D_{R^m, \sigma}$ if $R = \mathbb{Z}[x]/(x^n + 1)$.

When the image is a matrix $\mathbf{Y} = [\mathbf{y}_1 || \dots || \mathbf{y}_\ell]$, we abuse the notation for the preimage sampling algorithm, use $\mathbf{D} \leftarrow \text{PreimgSam}(\mathbf{A}, \tau, \mathbf{Y}, \sigma)$ to represent the concatenation of ℓ samples from $\mathbf{d}_i \leftarrow \text{PreimgSam}(\mathbf{A}, \tau, \mathbf{y}_i, \sigma)_{i \in [\ell]}$.

2.3 General learning with errors problems

The learning with errors (LWE) problem, formalized by Regev [48], states that solving noisy linear equations, in certain rings and for certain error distributions, is as hard as solving some worst-case lattice problems. The two typical forms used in cryptographic applications are (standard) LWE and RingLWE. The latter is introduced by Lyubashevsky, Peikert and Regev [42].

We formulate them as the General learning with errors problems similar to those of [17], with more flexibility in the secret distribution and the base ring.

Definition 2 (General learning with errors problem). *The (decisional) general learning with errors problem (GLWE) is parameterized by the base ring R , dimension parameters k, ℓ, m for samples, dimension parameter n for lattices, modulus q , the secret distribution η over $R^{k \times \ell}$, and the error distribution χ over $R^{\ell \times m}$. The $\text{GLWE}_{R,k,\ell,m,n,q,\eta,\chi}$ problem is to distinguish the following two distributions: (1) LWE samples $s \leftarrow \eta$, $\mathbf{A} \leftarrow U(R_q^{\ell \times m})$, $\mathbf{E} \leftarrow \chi^{k \times m}$, output $(\mathbf{A}, s\mathbf{A} + \mathbf{E}) \in (R_q^{\ell \times m} \times R_q^{k \times m})$; (2) uniform distributions $U(R_q^{\ell \times m} \times R_q^{k \times m})$.*

We define $\text{GLWE}_{R,k,\ell,m,n,q,\eta,\chi}$ -hardness for secret distributions. The subscripts are dropped if they are clear from the context.

Definition 3. *A secret distribution η is called $\text{GLWE}_{R,k,\ell,m,n,q,\eta,\chi}$ -hard if no p.p.t. adversary distinguishes the two distributions in the $\text{GLWE}_{R,k,\ell,m,n,q,\eta,\chi}$ problem with $1/2$ plus non-negligible probability.*

Here are the connections of decisional LWE/RingLWE to the worst-case lattice problems, in the language of GLWE-hardness. For the LWE problem we present the version where the secret is a square matrix.

Lemma 5 (LWE [48,46,18]). *Let n be an integer, $R = \mathbb{Z}^{n \times n}$. q be an integer modulus, $0 < \sigma < q$ such that $\sigma > 2\sqrt{n}$. If there exists an efficient (possibly quantum) algorithm that breaks $\text{GLWE}_{R,1,1,m,n,q,U(R_q),D_{\mathbb{Z},\sigma}^{n \times n}}$, then there exists an efficient (possibly quantum) algorithm for approximating SIVP and GapSVP in the ℓ_2 norm, in the worst case, to within $\tilde{O}(nq/\sigma)$ factors.*

Lemma 6 (RingLWE [42,29,41]). *Let n be a power of 2, $R = \mathbb{Z}[x]/(x^n + 1)$. Let q be a prime integer s.t. $q \equiv 1 \pmod{n}$. $0 < \sigma < q$, $\sigma > \omega(\sqrt{\log(n)})$, $\sigma' > n^{3/4}m^{1/4}\sigma$. If there exists an efficient (possibly quantum) algorithm that breaks $\text{GLWE}_{R,1,1,m,n,q,U(R_q),D_{R,\sigma'}}$, then there exists a polynomial time quantum algorithm for solving SVP for ideal-lattices over R , in the worst case, to within $\tilde{O}(\sqrt{n}q/\sigma)$ factors.*

For proper choices of parameters, error distributions of small norm can be used as hard secret distribution (usually called Hermit-normal-form LWE).

Lemma 7 (HNF-LWE [4,18]). *For R, m, n, q, σ chosen as was in Lemma 5, $\text{GLWE}_{R,1,1,m',n,q,D_{\mathbb{Z},\sigma}^{n \times n},D_{\mathbb{Z},\sigma}^{n \times n}}$ is as hard as $\text{GLWE}_{R,1,1,m,n,q,U(R_q),D_{\mathbb{Z},\sigma}^{n \times n}}$ for $m' \leq m - (16n + 4 \log \log q)$.*

Lemma 8 (HNF-RingLWE [43]). *For $R, m, n, q, \sigma, \sigma'$ chosen as in Lemma 6, $\text{GLWE}_{R,1,1,m-1,n,q,D_{R,\sigma'},D_{R,\sigma'}}$ is as hard as $\text{GLWE}_{R,1,1,m,n,q,U(R_q),D_{R,\sigma'}}$.*

Pseudorandom functions based on GLWE. We adapt theorems from the PRF construction of Boneh, Lewi, Montgomery, and Raghunathan [11, Theorems 4.3, 5.1]. The result was originally stated for LWE. We observe that it holds for general rings under proper choices of parameter. A proof sketch is described in [21].

Lemma 9 (Adapted from [11]). *Let $\ell \in \mathbb{N}$ be the bit-length of the input. $m, n, q, p \in \mathbb{N}$, $\sigma, B \in \mathbb{R}$ s.t. $0 < \sigma < q$, $B \geq \sigma\sqrt{m}$, $q/p > B^\ell$. $\eta = U(R_q)$, γ_σ is a distribution over $R^{m \times m}$ parameterized by σ , χ_σ is a distribution over $R^{1 \times m}$ parameterized by σ . $\|\gamma_\sigma\|, \|\chi_\sigma\| \leq \sigma\sqrt{m}$.*

Consider the function $f : \{0, 1\}^\ell \rightarrow R_p^{1 \times m}$, $f_{\mathbf{U}}(x) = \left[\mathbf{U} \prod_{i=1}^\ell \mathbf{D}_i^{x_i} \right]_p$, where $\mathbf{U} \leftarrow U(R_q^{1 \times m})$ is the private parameter, $\{\mathbf{D}_i^b \leftarrow \gamma_\sigma\}_{b \in \{0,1\}, i \in [\ell]}$ is the public parameter.

If there is an efficient algorithm that given input $\mathbf{A} \leftarrow U(R_q^{1 \times m})$, outputs $\mathbf{U} \in R_q^{1 \times m}, \mathbf{D} \in R^{m \times m}$ that are statistically close to $U(R_q^{1 \times m}) \times \gamma_\sigma$ and $\mathbf{U}\mathbf{D} = \mathbf{A}$; then f is a PRF assuming the hardness of $\text{GLWE}_{R,1,1,m,n,q,\eta,\chi_\sigma}$.

3 GLWE-hard distributions: extension package

We prove GLWE-hardness for the following “structural” secret distributions. They are used in the analysis of Construction 11.

Lemma 10. *Fix a permutation matrix $\mathbf{B} \in \{0, 1\}^{w \times w}$. If a secret distribution η over R is $\text{GLWE}_{R,1,1,w^2m,n,q,\eta,\chi}$ -hard, then the secret distribution $\mathbf{B} \otimes \eta$ is $\text{GLWE}_{R^{w \times w},1,1,m,n,q,\mathbf{B} \otimes \eta,\chi^{w \times w}}$ -hard.*

Proof. For a permutation matrix $\mathbf{B} \in \{0, 1\}^{w \times w}$, suppose there is a p.p.t. distinguisher between samples from

$$(\mathbf{B}, \mathbf{A}, (\mathbf{B} \otimes s)\mathbf{A} + \mathbf{E}), \text{ where } \mathbf{A} \leftarrow U(R_q^{w \times wm}), s \leftarrow \eta, \mathbf{E} \leftarrow \chi^{w \times wm}$$

and samples from the uniform distribution $(\mathbf{B}, U(R_q^{w \times wm}), U(R_q^{w \times wm}))$, then we build an attacker for $\text{GLWE}_{R,1,1,w^2m,n,q,\eta,\chi}$.

The attacker is given an $\text{GLWE}_{R,1,1,w^2m,n,q,\eta,\chi}$ instance

$$(\mathbf{A}', \mathbf{Y}') = (\mathbf{A}_1 || \dots || \mathbf{A}_w, \mathbf{Y}_1 || \dots || \mathbf{Y}_w), \text{ where } \mathbf{A}_i, \mathbf{Y}_i \in R^{1 \times wm}, i \in [w].$$

It then rearranges the blocks as $(\mathbf{U}, \mathbf{V}) \in R^{w \times wm} \times R^{w \times wm}$, where the i^{th} (blocked) row of \mathbf{U} is \mathbf{A}_i , the i^{th} (blocked) row of \mathbf{V} is \mathbf{Y}_i . The attacker then sends $(\mathbf{B}, \mathbf{U}, (\mathbf{B} \otimes 1_R)\mathbf{V})$ to the distinguisher. Observe that $(\mathbf{B}, \mathbf{U}, (\mathbf{B} \otimes 1_R)\mathbf{V})$ is from the $\mathbf{B} \otimes \eta$ secret distribution if $(\mathbf{A}', \mathbf{Y}')$ is from the η secret distribution, or from the uniform distribution if $(\mathbf{A}', \mathbf{Y}')$ is from the uniform distribution. Hence the attacker wins with the same probability as the distinguisher. \square

Lemma 11. *Let $w \in [2, \infty) \cap \mathbb{Z}$. Fix a permutation matrix $\mathbf{C} \in \{0, 1\}^{w \times w}$ that represents a w -cycle. If a secret distribution η over R is $\text{GLWE}_{R,1,1,wm,n,q,\eta,\chi}$ -hard, then $(\eta^{1 \times w}, \eta^{1 \times w} \times (\mathbf{C} \otimes 1_R))$ is $\text{GLWE}_{R,2,w,m,n,q,(\eta^{1 \times w}, \eta^{1 \times w} \times (\mathbf{C} \otimes 1_R)),\chi}$ -hard.*

Proof. Let $\mathbf{H} = [h_1, h_2, \dots, h_w]$ where $\{h_i \leftarrow \eta\}_{i \in [w]}$. Let

$$\mathcal{H} := \{(\mathbf{A}_j, \mathbf{Y}_{i,j} = h_i \mathbf{A}_j + \mathbf{E}_{i,j}) \mid \mathbf{A}_j \leftarrow U(R_q^{1 \times m}), h_i \leftarrow \eta, \mathbf{E}_{i,j} \leftarrow \chi, i, j \in [w]\}$$

be the rearranging of w independent GLWE samples from $\text{GLWE}_{R,1,1,wm,n,q,\eta,\chi}$. \mathcal{H} is indistinguishable from the uniform distribution $\mathcal{U} := \{(\mathbf{A}_j, \mathbf{Y}_{i,j}) \mid \mathbf{A}_j \leftarrow U(R_q^{1 \times m}), \mathbf{Y}_{i,j} \leftarrow U(R_q^{1 \times m}), i, j \in [w]\}$ due to standard GLWE.

We show that if there is an attacker D' that distinguishes

$$(\mathbf{A}, \mathbf{H}\mathbf{A} + \mathbf{E}, \mathbf{H}(\mathbf{C} \otimes \mathbf{1}_R)\mathbf{A} + \mathbf{E}'),$$

where $\mathbf{E}, \mathbf{E}' \leftarrow \chi^{1 \times m}$ from

$$U(R_q^{w \times m} \times R_q^{1 \times m} \times R_q^{1 \times m}),$$

then there is a distinguisher D for (a subset of) \mathcal{H} and \mathcal{U} .

To do so, we simulate the $(\eta^{1 \times w}, \eta^{1 \times w} \times (\mathbf{C} \otimes \mathbf{1}_R))$ samples from \mathcal{H} or \mathcal{U} by setting $\mathbf{A} \in R^{w \times m}$ where the j^{th} row of \mathbf{A} is \mathbf{A}_j , $\mathbf{Y} := \sum_{j \in [w]} \mathbf{Y}_{j,j}$, and $\mathbf{Z} := \sum_{j \in [w]} \mathbf{Y}_{\zeta(j),j}$, where $\zeta(j) : [w] \rightarrow [w]$ outputs the row number of the 1-entry in the j^{th} column of \mathbf{C} . Note that being a w -cycle indicates that the 1-entries in \mathbf{C} disjoint with the 1-entries in $\mathbf{I}^{w \times w}$. Observe that the sample $(\mathbf{A}, \mathbf{Y}, \mathbf{Z})$ is from the secret distribution $(\eta^{1 \times w}, \eta^{1 \times w} \times (\mathbf{C} \otimes \mathbf{1}_R))$ if transformed from \mathcal{H} , or from the uniform distribution if transformed from \mathcal{U} . Hence the distinguisher D' wins with the same probability as the attacker D . \square

4 Constraint-hiding constrained PRFs

This section provides the definitions of constraint-hiding constrained PRFs. We first recall the indistinguishability-based definition from [12], then give our simulation-based definition, and discuss the relations among these two definitions and program obfuscation.

4.1 The indistinguishability-based definition

We first recall the indistinguishability-based definition for CHCPRF from [12].

Definition 4 (Indistinguishability-based CHCPRF [12]).

Consider a family of functions $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ where $\mathcal{F}_\lambda = \{F_k : D_\lambda \rightarrow R_\lambda\}_{k \in \mathbb{N}}$, along with a triple of efficient functions $(\text{Gen}, \text{Constrain}, \text{Eval})$. For a constraint family $\mathcal{C} = \{C_\lambda : D_\lambda \rightarrow \{0,1\}\}_{\lambda \in \mathbb{N}}$; the key generation algorithm $\text{Gen}(1^\lambda)$ generates the master secret key MSK , the constraining algorithm $\text{Constrain}(1^\lambda, \text{MSK}, C)$ takes the master secret key MSK , a constraint C , outputs the constrained key CK ; the evaluation algorithm $\text{Eval}(k, x)$ takes a key k , an input x , outputs $F_k(x)$.

We say that \mathcal{F} is an **indistinguishability-based CHCPRF** for \mathcal{C} if it satisfies the following properties:

Functionality preservation over unconstrained inputs. For input $x \in D_\lambda$ s.t. $C(x) = 1$, $\Pr[\text{Eval}(\text{MSK}, x) = \text{Eval}(\text{CK}, x)] \geq 1 - \text{negl}(\lambda)$, where the probability is taken over the randomness in algorithms **Gen** and **Constrain**.

Pseudorandomness for constrained inputs. Consider the following experiment between a challenger and an adversary. The adversary can ask 3 types of oracle queries: constrained key oracle, evaluation oracle, and challenge oracle. For $b \in \{0, 1\}$, the challenger responds to each oracle query in the following manner:

- *Constrained key oracle.* Given a circuit $C \in \mathcal{C}$, the challenger outputs a constrained key $\text{CK} \leftarrow \text{Constrain}(1^\lambda, \text{MSK}, C)$.
- *Evaluation oracle.* Given an input $x \in D_\lambda$, the challenger outputs $y \leftarrow \text{Eval}(\text{MSK}, x)$.
- *Challenge oracle.* Given an input $x_c \in D_\lambda$, the challenger outputs $y \leftarrow \text{Eval}(\text{MSK}, x_c)$ if $b = 1$; outputs $y \leftarrow U(R_\lambda)$ if $b = 0$.

The queries from the adversary satisfy the conditions that $C(x_c) = 0$, and x_c is not sent among evaluation queries. At the end of the experiment, the adversary chooses b' and wins if $b' = b$. The scheme satisfies the pseudorandomness property if the winning probability of any p.p.t. adversary is bounded by $1/2 + \text{negl}(\lambda)$.

Indistinguishability-based constraint-hiding. Consider the following experiment between a challenger and an adversary. The adversary can ask 2 types of oracle queries: constrained key oracle or evaluation oracle. For $b \in \{0, 1\}$, the challenger responds to each oracle query in the following manner:

- *Constrained key oracle.* Given a pair of circuits $C_0, C_1 \in \mathcal{C}$, the challenger outputs a constrained key for C_b : $\text{CK} \leftarrow \text{Constrain}(1^\lambda, \text{MSK}, C_b)$.
- *Evaluation oracle.* Given an input $x \in D_\lambda$, the challenger outputs $y \leftarrow \text{Eval}(\text{MSK}, x)$.

For a circuit $C \in \mathcal{C}$, denote $S(C) := \{x \in D_\lambda : C(x) = 1\}$. Suppose the adversary asks h pairs of circuit constraints $\{C_0^{(g)}, C_1^{(g)}\}_{g \in [h]}$, the queries are **admissible** if (1) $\forall i \neq j \in [h]$, $S(C_0^{(i)}) \cap S(C_0^{(j)}) = S(C_1^{(i)}) \cap S(C_1^{(j)})$; (2) for all input evaluation queries x , for all $g \in [h]$, $C_0^{(g)}(x) = C_1^{(g)}(x)$.

At the end of the experiment, the adversary chooses b' and wins if $b' = b$. The scheme satisfies the constraint-hiding property if the winning probability of any p.p.t. adversary is bounded by $1/2 + \text{negl}(\lambda)$.

4.2 The simulation-based definition

Next we give the simulation-based definition. We first present a definition that is central to the discussions and constructions in the paper, then mention its variants.

Definition 5 (Simulation-based CHCPRF). Consider a family of functions $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ with the same syntax as in Definition 4. We say that \mathcal{F} is

simulation-based CHCPRF for family \mathcal{C} of circuits if for any polytime stateful algorithm Adv , there is a polytime stateful algorithm Sim such that:

$$\{\text{Experiment } \text{REAL}_{\text{Adv}}(1^\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Experiment } \text{IDEAL}_{\text{Adv}, \text{Sim}}(1^\lambda)\}_{\lambda \in \mathbb{N}}$$

The ideal and real experiments are defined as follows for adversaries Adv and Sim . Both algorithms are stateful.

<p><i>Experiment</i> $\text{REAL}_{\text{Adv}}(1^\lambda)$</p> <p>$\text{MSK} \leftarrow \text{Gen}(1^\lambda)$,</p> <p><i>Repeat</i> :</p> <p style="padding-left: 2em;">$\text{Adv} \rightarrow (x, d_x); y = \text{Eval}(\text{MSK}, x)$</p> <p style="padding-left: 2em;">$\text{Adv} \leftarrow y$</p> <p>$\text{Adv} \rightarrow C$;</p> <p style="padding-left: 2em;">if $d_x \neq C(x)$ for some x then <i>Output</i> \perp</p> <p style="padding-left: 2em;">else $\text{Adv} \leftarrow \text{Constrain}(\text{MSK}, C)$</p> <p><i>Repeat</i> :</p> <p style="padding-left: 2em;">$\text{Adv} \rightarrow x; y = \text{Eval}(\text{MSK}, x)$</p> <p style="padding-left: 2em;">$\text{Adv} \leftarrow y$</p> <p>$\text{Adv} \rightarrow b$; <i>Output</i> b</p>	<p><i>Experiment</i> $\text{IDEAL}_{\text{Adv}, \text{Sim}}(1^\lambda)$</p> <p>$\text{Sim} \leftarrow 1^\lambda$</p> <p><i>Repeat</i> :</p> <p style="padding-left: 2em;">$\text{Adv} \rightarrow (x, d_x); y = \text{Sim}(x, d_x)$</p> <p style="padding-left: 2em;">if $d_x = 0$ then $y = U(R)$; $\text{Adv} \leftarrow y$</p> <p>$\text{Adv} \rightarrow C$;</p> <p style="padding-left: 2em;">if $d_x \neq C(x)$ for some x then <i>Output</i> \perp</p> <p style="padding-left: 2em;">else $\text{Adv} \leftarrow \text{Sim}(1^{ C })$</p> <p><i>Repeat</i> :</p> <p style="padding-left: 2em;">$\text{Adv} \rightarrow x; y = \text{Sim}(x, C(x))$</p> <p style="padding-left: 2em;">if $C(x) = 0$ then $y = U(R)$; $\text{Adv} \leftarrow y$</p> <p>$\text{Adv} \rightarrow b$; <i>Output</i> b</p>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

That is, in the experiments the adversary can ask a single constraint query and polynomially many input queries, in any order. For input queries x made before the circuit query, Adv is expected to provide a bit b_x indicating whether $C(x) = 1$. In the real experiment Adv obtains the unconstrained function value at x . In the ideal experiment Sim learns the indicator bit d_x ; if $d_x = 1$ then Adv gets a value generated by Sim , and if $d_x = 0$ then Adv obtains a random value from the range R of the function. Once Adv makes the constraint query $C \in \mathcal{C}_\lambda$, both experiments verify the consistency of the indicator bits d_x for all the inputs x queried by Adv so far. If any inconsistency is found then the experiment halts. Next, in the real experiment Adv obtains the constrained key generated by the constraining algorithm; in the ideal experiment Adv obtains a key generated by Sim , whereas Sim is given only the size of C . The handling of input queries made by Adv after the circuit query is similar to the ones before, with the exception that the indicator bit d_x is no longer needed and Sim obtains the value of $C(x)$ instead. The output of the experiment is the final output bit of Adv .

Remark 1. One may also consider a stronger definition than Definition 5 where the adversary is not required to provide the indicator bits d_x in the queries prior to providing the constraint. However we note that this stronger definition is unachievable if the number of input queries before the constraint query is unbounded, due to an ‘‘incompressibility’’ argument similar to the one from [1].

Remark 2. The simulation-based definition can also be generalized to the setting where the adversary queries multiple constrained keys. That is, once received

each constrained key query, the simulator has to simulate a constrained key, given only the size of the constraining circuit. We further discuss this strong variant shortly.

4.3 Relations among the definitions

We discuss the relation among the definitions of CHCPRF and program obfuscation.

Multiple-key CHCPRFs implies obfuscation. We show that the simulation-based CHCPRF for 2 keys implies virtual black-box obfuscation (VBB), which is impossible to obtain for general functionalities [35,7]. For the indistinguishability-based definition proposed in [12], achieving 2-key security implies indistinguishability obfuscation [7].

Recall the definitions for VBB obfuscation (we present the strongest variant in [7]) and indistinguishability obfuscation.

Definition 6 (Obfuscation [35,7]). *A probabilistic algorithm O is an obfuscator for a class of circuit \mathcal{C} if the following conditions hold:*

- (Preservation of the function) For all inputs x , $\Pr[C(x) = O(C(x))] > 1 - \text{negl}(\lambda)$.
- (Polynomially slowdown) There is a polynomial p s.t. $|O(C)| < p(|C|)$.
- (Strong virtual black-box obfuscation) For any p.p.t. adversary Adv , there is a p.p.t. simulator Sim s.t. for all C , $\{\text{Adv}(1^\lambda, O(C))\} \approx_c \{\text{Sim}^C(1^\lambda, |C|)\}$.
- (Indistinguishability obfuscation) For functionally equivalent circuits C_0, C_1 , $O(C_0) \approx_c O(C_1)$.

Construction 4 (Obfuscator from 2-key CHCPRFs) *Given a CHCPRF, we construct an obfuscator for C by create a constrained key $\text{CK}[C]$, and a constrained key $\text{CK}[I]$ where I is the circuit that always outputs 1. To evaluate $C(x)$, output 1 if $\text{CHCPRF}_{\text{CK}[C]}(x) = \text{CHCPRF}_{\text{CK}[I]}(x)$, 0 otherwise.*

Theorem 5. *If 2-key simulation-secure CHCPRF exists for circuit class \mathcal{C} , then strong VBB obfuscation exists for circuit class \mathcal{C} .*

Proof. The simulator for the VBB obfuscator (does not have to make oracle queries to C) runs the simulator for CHCPRF, produce simulated constraint keys for $\text{CK}^S[C]$, $\text{CK}^S[I]$, which are indistinguishable from the real constrained keys $\text{CK}[C]$, $\text{CK}[I]$ that are used to construct the obfuscator. \square

Corollary 1 ([35,7]). *There are circuit classes for which 2-key simulation-secure CHCPRF does not exist.*

Theorem 6. *If 2-key indistinguishability-based CHCPRF exists for circuit class \mathcal{C} , then indistinguishability obfuscation exists for circuit class \mathcal{C} .*

Proof. For a circuit C , the obfuscator outputs $\text{CK}[C]$, $\text{CK}[I]$. For functionally equivalent circuits C_0 and C_1 , $S(C_0) \cap S(I) = S(C_1) \cap S(I)$. By indistinguishability constraint-hiding, $(\text{CK}[C_0], \text{CK}[I]) \approx_c (\text{CK}[C_1], \text{CK}[I])$. \square

Simulation and indistinguishability-based definitions for CHCPRF. Next we discuss the relation of the simulation and indistinguishability-based definitions for CHCPRF, under 1-key security. The two definitions are equivalent in the 1-key setting, for the corresponding order of queries and adaptivity. Below we state the theorems for the non-adaptive version of the definitions, then discuss their generalizations to the adaptive setting.

We first show that the simulation based definition implies the indistinguishability based definition.

Theorem 7. *If a CHCPRF satisfies the non-adaptive simulation-based definition, then it satisfies the non-adaptive indistinguishability-based definition.*

The proof of this theorem is via a standard hybrid argument and we describe the proof in [21]. The implication holds for the adaptive setting. In particular, the standard simulation definition from Definition 5 implies Definition 4 where the predicates on the input queries are committed; for the stronger simulation definition discussed in Remark 1, it implies the fully adaptive variant of Definition 4.

In the 1-key setting, the indistinguishability definition implies the simulation based definition.

Theorem 8. *If a CHCPRF satisfies 1-key non-adaptive indistinguishability-based definition, it satisfies the 1-key non-adaptive simulation-based definition.*

Proof. For a CHCPRF \mathcal{F} that satisfies Definition 4 for one constrained key query, we construct a simulator as per Definition 5. The simulator picks an all-1 circuit $C^S = I$ such that $I(x) = 1, \forall x \in D_\lambda$, and use the indistinguishability-secure constraining algorithm to derive a constrained key CK^S for C^S . Once the simulator obtains the inputs and the indicators $\{x^{(k)}, d^{(k)}\}_{k \in [t]}$, if $d^{(k)} = 1$, outputs $\text{Eval}(\text{CK}^S, x^{(k)})$; if $d^{(k)} = 0$, outputs $y \leftarrow U(R_\lambda)$.

We first prove constraint-hiding. Suppose there is an adversary A' that distinguishes the simulated distribution from the real distribution, we build an adversary A that breaks the indistinguishability definition for \mathcal{F} . A sends constrained circuit queries $C_0 = C$ and $C_1 = I$, obtains $\text{CK}[C_b]$. Then A sends input queries. For $x^{(k)}$ s.t. $C(x^{(k)}) = I(x^{(k)}) = 1$, the output is $\text{Eval}(\text{CK}[C_b], x^{(k)})$; for $x^{(k)}$ s.t. $C(x^{(k)}) \neq I(x^{(k)})$, it is an inadmissible query so A samples a uniform random output on its own. Then A forwards $\text{CK}[C_b]$, inputs and outputs to A' . The choice of A' for the real or the simulated distribution corresponds to $b = 0$ or 1, hence the advantage of A is equivalent to A' .

The proof for pseudorandomness of constrained outputs is analogous. \square

The theorem extends to the setting where the input queries can be made after the constraint query.

5 The constructions

In Sections 5.1 and 5.2 we present the bit-fixing and NC^1 CHCPRFs.

5.1 Bit-fixing CHCPRFs

Definition 7 (Bit-fixing constraint [15]). A bit-fixing constraint is specified by a string $\mathbf{c} \in \{0, 1, \star\}^\ell$, where 0 and 1 are the fixing bits and \star denotes the wildcards. $C(x) = 1$ if the input matches \mathbf{c} , namely $((x_1 = c_1) \vee (c_1 = \star)) \wedge \dots \wedge ((x_\ell = c_\ell) \vee (c_\ell = \star))$.

We start with a brief overview of the construction and then give the details. For a PRF with ℓ -bit input, the key-generation algorithm samples 2ℓ secrets from GLWE-hard distributions with small Euclidean norm $\{s_i^b \leftarrow \eta\}_{b \in \{0,1\}, i \in [\ell]}$, places them in a chain of length ℓ and width 2, and uses the GGH15 methodology to encode the chain. The evaluation key consists of the resulting \mathbf{A}_1 matrix and the \mathbf{D} matrices $\{\mathbf{D}_i^b\}_{b \in \{0,1\}, i \in [\ell]}$.

The evaluation algorithm selects the path according to the input, computes the product of \mathbf{D} matrices along the path $\prod_{i=1}^\ell \mathbf{D}_i^{x_i}$, then multiplies \mathbf{A}_1 on the left. The unrounded version of the output $\mathbf{A}_1 \prod_{i=1}^\ell \mathbf{D}_i^{x_i}$ is close to $\prod_{i=1}^\ell s_i^{x_i} \mathbf{A}_{\ell+1}$, where “close” hides the cumulated error terms. Finally, the resulting subset product is rounded by p where $2 \leq p < q$, $q/p > B$ with B being the maximum error bound. Rounding is required for correctness and security.

Construction 9 (Bit-fixing CHCPRFs) We construct a function family $\mathcal{F} = \{f : \{0, 1\}^\ell \rightarrow R_p^{1 \times m}\}$ equipped with algorithms (Gen, Constrain, Eval) and a set of vectors $\mathcal{C} = \{\mathbf{c} \in \{0, 1, \star\}^\ell\}$:

- **Gen**(1^λ) takes the security parameter λ , samples parameters $q, p, \sigma, m, \mathbf{A}_{\ell+1} \leftarrow U(R_q^m)$, $\{(\mathbf{A}_i, \tau_i) \leftarrow \text{TrapSam}(R, 1^n, 1^m, q)\}_{i \in [\ell]}$. Then, sample 2ℓ independent small secrets from GLWE-hard distributions $\{s_i^b \leftarrow \eta\}_{b \in \{0,1\}, i \in [\ell]}$. Next, encode the secrets as follows: first compute $\{\mathbf{Y}_i^b = s_i^b \mathbf{A}_{i+1} + \mathbf{E}_i^b, \mathbf{E}_i^b \leftarrow \chi^m\}_{i \in [\ell], b \in \{0,1\}}$, then sample $\{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{i \in [\ell], b \in \{0,1\}}$. Set $\text{MSK} := (\{\mathbf{A}_i\}_{i \in [1, \ell+1]}, \{\tau_i\}_{i \in [\ell]}, \{s_i^b, \mathbf{D}_i^b\}_{i \in [\ell], b \in \{0,1\}})$.
- **Constrain**(MSK, \mathbf{c}) takes MSK and the bit-matching vector \mathbf{c} , for $i \in [\ell]$, if $c_i \neq \star$ (i.e. specified as 0 or 1), replaces the original $s_i^{1-c_i}$ by a fresh $t_i^{1-c_i} \leftarrow \eta$, then updates the encodings on these secrets: $\mathbf{Y}_i^{1-c_i} = t_i^{1-c_i} \mathbf{A}_{i+1} + \mathbf{E}_i^{\prime 1-c_i}, \mathbf{E}_i^{\prime 1-c_i} \leftarrow \chi^m$, samples $\mathbf{D}_i^{1-c_i} \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^{1-c_i}, \sigma)$. Set $\text{CK} := (\mathbf{A}_1, \{\mathbf{D}_i^b\}_{i \in [\ell], b \in \{0,1\}})$.
- **Eval**(k, x) takes the key $k = (\mathbf{A}_1, \{\mathbf{D}_i^b\}_{i \in [\ell], b \in \{0,1\}})$ and the input x , outputs $\left\lfloor \mathbf{A}_1 \prod_{i=1}^\ell \mathbf{D}_i^{x_i} \right\rfloor_p$.

Remark 3. We occasionally call $i \in \{1, 2, \dots, \ell\}$ “levels”, from low to high.

Setting of parameters. Parameters shall be set to ensure both correctness (i.e. the preservation of functionality over unconstrained inputs) and security. Note that the approximation factors of the underlying worst-case (general or ideal) lattices problems are inherently exponential in ℓ .

Specifically, for $R = \mathbb{Z}^{n \times n}$, set $\eta = \chi = D_{\mathbb{Z}, \sigma}^{n \times n}$, $\gamma = D_{\mathbb{Z}^{nm}, \sigma}^{1 \times nm}$. The parameters are set to satisfy $m \geq 2 \log q$ due to Lemma 3; $q/p > (\sigma \cdot m)^\ell$ due to Lemma 1

for the correctness of rounding; $0 < \sigma < q$, $\sigma = 2\sqrt{n \log q}$, $nq/\sigma < 2^{\lambda^{1-\epsilon}}$ due to Lemmas 4, 5, 7, and 9. An example setting of parameters: $p = 2$, $\epsilon = 1/2$, $q = (32\ell n^2 \log n)^\ell$, $\lambda = n = (\log q)^2$.

For $R = \mathbb{Z}[x]/(x^n + 1)$, n being a power of 2, set $\eta = \chi = D_{R,\sigma}$, $\gamma = D_{R^m,\sigma}^{1 \times m}$. The parameters are set to satisfy $m \geq 2 \cdot n \log q$ due to Lemma 3; $q/p > (\sigma \cdot n^{3/4} m^{5/4})^\ell$ due to Lemma 1 for the correctness of rounding; $0 < \sigma < q$, $\sigma = 2\sqrt{n \log q}$, $nq/\sigma < 2^{\lambda^{1-\epsilon}}$ due to Lemmas 4, 6, 8, and 9. An example setting of parameters against the state-of-art ideal SVP algorithms [9,27,28]: $p = 2$, $\epsilon = 0.5001$, $q = (70\ell n^3 \log n)^\ell$, $\lambda = n = (\log q)^{2.1}$.

Theorem 10. *Assuming $\text{GLWE}_{R,1,1,m,n,q,\eta,\chi}$, Construction 9 is a simulation-secure bit-fixing CHCPRF.*

Functionality preservation on the unconstrained inputs. The constraining algorithm does not change any secrets on the unconstrained paths. So the functionality is perfectly preserved.

Security proof overview. The aim is to capture two properties: (1) pseudorandomness on the constrained inputs (2) the constrained key is indistinguishable from an obliviously sampled one.

We construct a simulator as follows: the simulator samples a key composed of \mathbf{A} matrices from uniform distribution and \mathbf{D} matrices from discrete-Gaussian distribution of small width. For the input-output pairs queried by the adversary, if the functionality is preserved on that point, then the simulator, knowing the input x , simply outputs the honest evaluation on the simulated key. If the input is constrained, it means at some level i , the secret $t_i^{x_i}$ in the constrained key is sampled independently from the original secret key $s_i^{x_i}$. Therefore the LWE instance $s_i^{x_i} \mathbf{A}_{i+1} + \mathbf{E}_i^{x_i}$, in the expression of the constrained output, provides a fresh random mask \mathbf{U} . The reduction moves from level $\ell + 1$ to level 1. At level 1, by the result of [11], the rounded output on x is pseudorandom if $C(x) = 0$.

Note that the evaluation algorithm only needs \mathbf{A}_1 but not the rest of the \mathbf{A} matrices. However, in the analysis we assume all the \mathbf{A} matrices are public.

Proof. The simulator samples all the $\{\mathbf{A}_j\}_{j \in [1, \ell+1]}$ matrices from random and $\{\mathbf{D}_i^b\}_{b \in \{0,1\}, i \in [\ell]}$ from γ , outputs the constrained key $(\mathbf{A}_1, \{\mathbf{D}_i^b\}_{i \in [\ell], b \in \{0,1\}})$. To respond the input queries, the simulator picks $\{y^{(k)}\}_{k \in [\ell]}$ according to $\{d^{(k)}\}_{k \in [\ell]}$: if $d^{(k)} = 1$ (i.e. the functionality is preserved on the constraint key at $x^{(k)}$), then outputs $y^{(k)} = \left[\mathbf{A}_1 \prod_{i=1}^{\ell} \mathbf{D}_i^{x_i^{(k)}} \right]_p$ (the honest evaluation on the simulated key); otherwise $y^{(k)} \leftarrow U(R_p^{1 \times m})$.

The proof consists of two parts. The first part (Lemma 12) shows that the real distribution is indistinguishable from a semi-simulated one, where all the \mathbf{D} matrices on the constrained key are sampled obliviously without knowing the constraint and the trapdoors of \mathbf{A} matrices, and all the outputs are derived from the simulated constrained key. The second part (Lemma 13) argues that the outputs are pseudorandom if they are in the constrained area.

In the first part, we define intermediate hybrid distributions $\{H_v\}_{v \in [0, \ell]}$. H_ℓ corresponds to the real constrained key and outputs, H_0 corresponds to the simulated constrained key and the semi-simulated outputs. The intermediate simulator in H_v knows the partial constraint from level 1 to v , and the level $w^{(k)} \in \{v, \dots, \ell\}$ where the input $x^{(k)}$ starts to deviate from the constraint vector.

Descriptions of H_v , $v \in [0, \ell]$: The simulator in H_v

1. Samples $\{(\mathbf{A}_j, \tau_j) \leftarrow \text{TrapSam}(R, 1^n, 1^m, q)\}_{j \in [v]}$ with trapdoors, $\{\mathbf{A}_{j'} \leftarrow U(R_q^m)\}_{j' \in [v+1, \ell+1]}$ from uniform;
2. Samples the GLWE secrets $\{s_i^b \leftarrow \eta\}_{b \in \{0,1\}, i \in [v]}$ below level v ; then, with part of the constraint vector $\mathbf{c}_{[v]}$ in hand, for $i \in [v]$, if $c_i \neq \star$, samples $t_i^{1-c_i} \leftarrow \eta$;
3. For $b \in \{0, 1\}$, $i \in [v]$, if t_i^b is sampled in the previous step, samples $\mathbf{Y}_i^b := t_i^b \mathbf{A}_{i+1} + \mathbf{E}'_i^b$; otherwise, $\mathbf{Y}_i^b := s_i^b \mathbf{A}_{i+1} + \mathbf{E}_i^b$;
4. Samples $\{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{b \in \{0,1\}, i \in [v]}$ as the constrained-key below level v . Samples the rest of the \mathbf{D} matrices obliviously $\{\mathbf{D}_i^b \leftarrow \gamma\}_{b \in \{0,1\}, i \in [v+1, \ell]}$.
5. To simulate the outputs, the simulator maintains a list \mathcal{U} of \mathbf{U} matrices (to be specified) initiated empty. For $k \in [t]$, if the constraint is known to deviate in the path of $x^{(k)}$ from level $w^{(k)} \in [v+1, \ell]$, then compute $y^{(k)}$ as $\left[\prod_{i=1}^v s_i^{x_i^{(k)}} \mathbf{U}^{x_{[v+1, w]^{(k)}}} \prod_{j=w^{(k)}}^\ell \mathbf{D}_j^{x_j^{(k)}} \right]_p$ — here $\mathbf{U}^{x_{[v+1, w]^{(k)}}}$ is indexed by $x_{[v+1, w]^{(k)}}$; if it is not in the list \mathcal{U} , sample $\mathbf{U}^{x_{[v+1, w]^{(k)}}} \leftarrow U(R_q^m)$, include it in \mathcal{U} ; otherwise, reuse the one in \mathcal{U} . If $x^{(k)}$ has not deviated above level v , then $y^{(k)} = \left[\prod_{i=1}^v s_i^{x_i^{(k)}} \mathbf{A}_{v+1} \prod_{j=v+1}^\ell \mathbf{D}_j^{x_j^{(k)}} \right]_p$.

Lemma 12. $H_v \approx_c H_{v-1}$, for $v \in \{\ell, \dots, 1\}$.

Proof. The difference of H_v and H_{v-1} lies in the sampling of \mathbf{D}_v^0 , \mathbf{D}_v^1 and the outputs $\{y^{(k)}\}$. We first analyze the difference of the outputs between H_v and H_{v-1} by classifying the input queries into 3 cases:

1. For input $x^{(k)}$ that matches the partial constraint vector $\mathbf{c}_{[v, \ell]}$, observe that $\left[\prod_{i=1}^{v-1} s_i^{x_i^{(k)}} \mathbf{A}_v \prod_{j=v}^\ell \mathbf{D}_j^{x_j^{(k)}} \right]_p = \left[\prod_{i=1}^{v-1} s_i^{x_i^{(k)}} (s_v^{x_v^{(k)}} \mathbf{A}_{v+1} + \mathbf{E}_v^{x_v^{(k)}}) \prod_{j=v+1}^\ell \mathbf{D}_j^{x_j^{(k)}} \right]_p$
 $\approx_s \left[\prod_{i=1}^v s_i^{x_i^{(k)}} \mathbf{A}_{v+1} \prod_{j=v+1}^\ell \mathbf{D}_j^{x_j^{(k)}} \right]_p$, where \approx_s is due to the small norm of $\prod_{i=1}^{v-1} s_i^{x_i^{(k)}} \mathbf{E}_v^{x_v^{(k)}} \prod_{j=v+1}^\ell \mathbf{D}_j^{x_j^{(k)}}$. Hence the output is statistically close in H_{v-1} and H_v .
2. For the input $x^{(k)}$ that is preserving above level v but deviated at level v , the fresh LWE secret $t_v^{x_v^{(k)}}$ sampled in the constrained key is independent from the original key $s_v^{x_v^{(k)}}$. So $s_v^{x_v^{(k)}} \mathbf{A}_{v+1} + \mathbf{E}_v^{x_v^{(k)}}$ and $t_v^{x_v^{(k)}} \mathbf{A}_{v+1} + \mathbf{E}'_v^{x_v^{(k)}}$ are treated as independent LWE instances w.r.t. \mathbf{A}_{v+1} .

3. For $x^{(k)}$ that has deviated above level v , the output can be written as

$$\begin{aligned} y^{(k)} &= \left[\prod_{i=1}^v s_i^{x_i^{(k)}} \mathbf{U}^{x_{[v+1,w]}^{(k)}} \prod_{j=w^{(k)}}^{\ell} \mathbf{D}_j^{x_j^{(k)}} \right]_p \\ &\approx_s \left[\prod_{i=1}^{v-1} s_i^{x_i^{(k)}} (s_v^{x_v^{(k)}} \mathbf{U}^{x_{[v+1,w]}^{(k)}} + \mathbf{E}') \prod_{j=w^{(k)}}^{\ell} \mathbf{D}_j^{x_j^{(k)}} \right]_p, \end{aligned} \quad (2)$$

where $\mathbf{U}^{x_{[v+1,w]}^{(k)}}$ is uniform by induction.

To summarize, there are less than $3(|\mathcal{U}| + 1)$ matrices that are GLWE samples in \mathbf{H}_v while uniform in \mathbf{H}_{v-1} . The GLWE samples involves 3 independent secrets: s_v^0, s_v^1 and $t_v^{1-c_v}$ if $c_v \neq \star$. $t_v^{1-c_v}$ is only masked by \mathbf{A}_{v+1} ; $\{s_v^b\}_{b \in \{0,1\}}$ are masked by \mathbf{A}_{v+1} (in the constrained key and the outputs of cases (1) and (2)) and the uniform matrices in the list \mathcal{U} (the outputs of case (3)); all the samples are associated with independently sampled noises.

If there is an attacker A' that distinguishes \mathbf{H}_v and \mathbf{H}_{v-1} with non-negligible probability ζ , we can build an attacker A who distinguishes (a subset among the $3(|\mathcal{U}| + 1)$) GLWE samples

$$\{[\mathbf{A}_{v+1}, \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^{|\mathcal{U}|}], [s_v^0, s_v^1, t_v^{1-c_i}]^T \cdot [\mathbf{A}_{v+1}, \mathbf{U}^1, \mathbf{U}^2, \dots, \mathbf{U}^{|\mathcal{U}|}] + \tilde{\mathbf{E}}\},$$

where $\tilde{\mathbf{E}} \leftarrow \chi^{3 \times (|\mathcal{U}|+1)m}$ from

$$\{U(R_q^{(|\mathcal{U}|+1)m} \times R_q^{3 \times (|\mathcal{U}|+1)m})\}$$

To do so, once A obtains the samples, it places the samples under mask \mathbf{A}_{v+1} in the constrained key and the outputs of cases (1) and (2); places the samples under masks $\mathbf{U}^1, \dots, \mathbf{U}^{|\mathcal{U}|}$ in the outputs of cases (3). Then samples $\{\mathbf{A}_j\}_{j \in [v]}$ with trapdoors, GLWE secrets $\{s_i^b \leftarrow \eta\}_{b \in \{0,1\}, i \in [v]}$. Then samples $\{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{b \in \{0,1\}, i \in [v]}$ as the constrained-key below level v . Samples the rest of the \mathbf{D} matrices obviously $\{\mathbf{D}_i^b \leftarrow \gamma\}_{b \in \{0,1\}, i \in [v+1, \ell]}$.

With these matrices the attacker A is able to simulate the outputs, send the outputs and constrained key to A' . If the samples are from GLWE, then it corresponds to \mathbf{H}_v ; if the samples are uniform, then the matrices $\{\mathbf{D}_v^b\}_{b \in \{0,1\}}$ sampled via $\{\mathbf{D}_v^b \leftarrow \text{PreimgSam}(\mathbf{A}_v, \tau_v, \mathbf{Y}_v^b, \sigma)\}_{b \in \{0,1\}}$ are statistically close to the obviously sampled ones due to Lemma 4, so it is statistically close to \mathbf{H}_{v-1} . Hence A breaks GLWE with probability more than $\zeta/(3(t+1))$, which contradicts to Lemma 7. \square

Lemma 13. *If $C(x^{(k)}) = 0$, then the output $y^{(k)}$ in \mathbf{H}_0 is pseudorandom.*

Proof. A constrained output $y^{(k)}$ can be expressed as $\left[\mathbf{U}^{x_{[1,w]}^{(k)}} \prod_{j=w^{(k)}}^{\ell} \mathbf{D}_j^{x_j^{(k)}} \right]_p$,

where the secret $\mathbf{U}_{[1,w^{(k)}]}^{(k)}$ is uniform; the public \mathbf{D} matrices are sampled from discrete-Gaussian distribution γ . By Lemma 9 $y^{(k)}$ is pseudorandom. \square

The proof completes by combining Lemma 12 and Lemma 13. \square

5.2 Constraint-hiding for NC^1 circuits

Next we present the CHCPRF for NC^1 circuit constraints. For circuits of depth d , use Barrington's Theorem [8] to convert them into a set of oblivious branching program $\{\text{BP}\}$ with the same index-to-input map $\iota : [z] \rightarrow [\ell]$, the same w -cycle \mathbf{P} that represents the 0 output (by default $w = 5$). Let $\{\mathbf{B}_i^b \in \{0, 1\}^{w \times w}\}_{i \in [z], b \in \{0, 1\}}$ be the permutation matrices in each BP.

The master secret key for the CHCPRF consists of $2z$ secrets from GLWE-hard distributions η over R with small Euclidean norm $\{s_i^b \leftarrow \eta\}_{b \in \{0, 1\}, i \in [z]}$, together with a vector $\mathbf{J} \in R^{1 \times w}$. To generate an evaluation key, in the normal setting, let $\mathbf{S}_i^b := \mathbf{I}^{w \times w} \otimes s_i^b \in \{0, 1\}^{w \times w} \otimes_R R$; in the constrained setting for a constraint recognized by BP, let $\mathbf{S}_i^b := \mathbf{B}_i^b \otimes s_i^b \in \{0, 1\}^{w \times w} \otimes_R R$. For both settings, places $\{\mathbf{S}_i^b\}_{b \in \{0, 1\}, i \in [z]}$ in a chain of length z and width 2, places \mathbf{J} on the left end of the chain, and uses the GGH15 methodology to encode the chain. The encoding of \mathbf{J} is merged into \mathbf{A}_1 and denote the resultant matrix as \mathbf{A}_J . The evaluation key consists of \mathbf{A}_J and the \mathbf{D} matrices $\{\mathbf{D}_i^b\}_{b \in \{0, 1\}, i \in [z]}$.

To evaluate on x , output $[\mathbf{A}_J \prod_{i=1}^z \mathbf{D}_i^{x_{\iota(i)}}]_p$. To elaborate the functionality, for x s.t. $C(x) = 1$, $[\mathbf{A}_J \prod_{i=1}^z \mathbf{D}_i^{x_{\iota(i)}}]_p \approx_s [\mathbf{J}(\mathbf{I}^{w \times w} \otimes \prod_{i=1}^z s_i^{x_{\iota(i)}})\mathbf{A}_{z+1}]_p$; for x s.t. $C(x) = 0$, $[\mathbf{A}_J \prod_{i=1}^z \mathbf{D}_i^{x_{\iota(i)}}]_p \approx_s [\mathbf{J}(\mathbf{P} \otimes \prod_{i=1}^z s_i^{x_{\iota(i)}})\mathbf{A}_{z+1}]_p$. As a reminder, the permutation matrix \mathbf{P} that represent the w -cycle is not a secret to the construction, so the use of the left-bookend \mathbf{J} is essential for security.

Construction 11 (CHCPRFs for NC^1 circuits) *We construct a function family $\mathcal{F} = \{f : \{0, 1\}^\ell \rightarrow R_p^{1 \times wm}\}$ equipped with 3 algorithms (Gen, Constrain, Eval), associated with a set of oblivious branching programs $\{\text{BP}\}$ of length z obtained by applying Lemma 3 on all the NC^1 circuits.*

- Gen(1^λ) samples parameters q, p, σ, m, z (the length of branching programs), $\{(\mathbf{A}_i, \tau_i) \leftarrow \text{TrapSam}(R^{w \times w}, 1^n, 1^m, q)\}_{i \in [z]}$, $\mathbf{A}_{z+1} \leftarrow U(R_q^{w \times wm})$. Samples $2z$ independent small secrets from GLWE-hard distributions $\{s_i^b \leftarrow \eta\}_{b \in \{0, 1\}, i \in [z]}$, sets the secret matrices to be $\mathbf{S}_i^b = \mathbf{I}^{w \times w} \otimes s_i^b$. Next, encode the secrets as follows: first compute $\{\mathbf{Y}_i^b = \mathbf{S}_i^b \mathbf{A}_{i+1} + \mathbf{E}_i^b, \mathbf{E}_i^b \leftarrow \chi^{w \times wm}\}_{i \in [z], b \in \{0, 1\}}$; then, sample $\{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{i \in [z], b \in \{0, 1\}}$. Additionally, sample a small secret $\mathbf{J} \leftarrow \eta^{1 \times w}$ as the left-bookend. Compute $\mathbf{A}_J := \mathbf{J}\mathbf{A}_1 + \mathbf{E}_J$ where $\mathbf{E}_J \leftarrow \chi^{1 \times wm}$.
Set $\text{MSK} := (\{\mathbf{A}_i\}_{i \in [1, z+1]}, \{\tau_i\}_{i \in [z]}, \mathbf{A}_J, \{s_i^b, \mathbf{D}_i^b\}_{i \in [z], b \in \{0, 1\}})$.
- Constrain(MSK, BP) takes MSK, and a matrix branching program $\text{BP} = \{\mathbf{B}_i^b \in R^{w \times w}\}_{i \in [z], b \in \{0, 1\}}$. For $i \in [z]$, $b \in \{0, 1\}$, compute $\mathbf{Y}_i^b = (\mathbf{B}_i^b \otimes s_i^b) \mathbf{A}_{i+1} + \mathbf{E}_i^b$, $\mathbf{E}_i^b \leftarrow \chi^{w \times wm}$, samples $\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)$.
Set the constrained key $\text{CK} := (\mathbf{A}_J, \{\mathbf{D}_i^b\}_{i \in [z], b \in \{0, 1\}})$.
- Eval(k, x) takes the input x and the key $k = (\mathbf{A}_J, \{\mathbf{D}_i^b\}_{i \in [z], b \in \{0, 1\}})$, outputs $[\mathbf{A}_J \prod_{i=1}^z \mathbf{D}_i^{x_{\iota(i)}}]_p$.

Setting of parameters. Settings of the distributions and their dimensions: For $R = \mathbb{Z}^{n \times n}$, set $\eta = \chi = D_{\mathbb{Z}, \sigma}^{n \times n}$, $\gamma = D_{\mathbb{Z}^{n \times wm}, \sigma}^{1 \times n \times wm}$. For $R = \mathbb{Z}[x]/(x^n + 1)$, n being a power of 2, set $\eta = \chi = D_{R, \sigma}$, $\gamma = D_{R^{wm}, \sigma}^{1 \times wm}$.

The restriction on the parameters are analogous to the settings in the bit-fixing construction.

Theorem 12. *Assuming $\text{GLWE}_{R, 1, 1, w^2 m, n, q, \eta, \chi}$, Construction 11 is a simulation-secure CHCPRF for NC¹ constraints.*

Proof overview. The simulation algorithm and the overall proof strategy is similar to the one for the bit-fixing constraints. Namely, we close the trapdoors for \mathbf{A} matrices from level z to level 1. Within each level v , there are several GLWE instance associated with \mathbf{A}_{v+1} whose trapdoor is closed in the previous hybrid. The additional complexity comes from dealing with secrets with permutation matrix structures. They are handled by the new GLWE packages from Section 3.

Proof. The simulator samples $\{\mathbf{A}_i \leftarrow U(R_q^{w \times wm})\}_{i \in [1, z+1]}$, and $\{\mathbf{D}_i^b \leftarrow \gamma\}_{b \in \{0, 1\}, i \in [z]}$. It also samples $\mathbf{J} \leftarrow \eta^{1 \times w}$, computes $\mathbf{A}_J := \mathbf{J}\mathbf{A}_1 + \mathbf{E}_J$ where $\mathbf{E}_J \leftarrow \chi^{1 \times wm}$. Outputs the constrained key $(\mathbf{A}_J, \{\mathbf{D}_i^b\}_{i \in [z], b \in \{0, 1\}})$. The simulator responds the input queries by picking $\{y^{(k)}\}_{k \in [t]}$ according to $\{d^{(k)}\}_{k \in [t]}$: if $d^{(k)} = 1$, then outputs $y^{(k)} = \left[\mathbf{A}_J \prod_{i=1}^z \mathbf{D}_i^{x^{(k)}} \right]_p$; otherwise $y^{(k)} \leftarrow U(R_p^{1 \times wm})$.

The proof consists of two parts. The first part (Lemma 14) shows that the real distribution is indistinguishable from a semi-simulated one, where all the \mathbf{D} matrices on the constrained key are sampled without knowing the constraint and trapdoors of \mathbf{A} matrices, and all the outputs are expressed by these obliviously sampled \mathbf{A} and \mathbf{D} matrices. The second part (Lemma 15) argues that the outputs are pseudorandom if they are in the constrained area.

In the first part, we define intermediate hybrid distributions $\{H_v\}_{v \in [0, z]}$. H_z corresponds to the real constrained key and output distributions, H_0 corresponds to the simulated constrained key and the semi-simulated outputs. The simulators in H_z, H_{z-1}, \dots, H_1 know the full description of the constraint $\text{BP} = \{\mathbf{B}_i^b\}_{i \in [z], b \in \{0, 1\}}$; the simulator in H_0 only knows the indicators $\{d^{(k)}\}_{k \in [t]}$.

Descriptions of $H_v, v \in [0, z]$: The simulator in H_v

1. Samples $\{(\mathbf{A}_j, \tau_j) \leftarrow \text{TrapSam}(R^{w \times w}, 1^n, 1^m, q)\}_{j \in [v]}$ with trapdoors; samples $\{\mathbf{A}_{j'} \leftarrow U(R_q^{w \times wm})\}_{j' \in [v+1, z+1]}$ uniformly random;
2. Samples the GLWE secrets $\{s_i^b \leftarrow \eta\}_{b \in \{0, 1\}, i \in [v]}$ below level v ; and a bookend vector $\mathbf{J} \leftarrow \eta^{1 \times w}$;
3. Samples $\mathbf{Y}_i^b := (\mathbf{B}_i^b \otimes s_i^b) \mathbf{A}_{i+1} + \mathbf{E}'_i^b$; computes $\mathbf{A}_J := \mathbf{J}\mathbf{A}_1 + \mathbf{E}_J$;
4. Simulates $\{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{b \in \{0, 1\}, i \in [v]}$ as the constrained-key below level v . Samples the rest of the \mathbf{D} matrices obliviously $\{\mathbf{D}_i^b \leftarrow \gamma\}_{b \in \{0, 1\}, i \in [v+1, z]}$.
5. Simulates the outputs. For $k \in [t]$, computes $y^{(k)}$ as

$$y^{(k)} = \left[\mathbf{J} \times \left(\left(\prod_{j=v+1}^z \mathbf{B}_j^{x^{(k)}} \right)^{-1} \otimes \prod_{i=1}^v s_i^{x^{(k)}} \right) \times \mathbf{A}_{v+1} \prod_{j=v+1}^z \mathbf{D}_j^{x^{(k)}} \right]_p \quad (3)$$

Lemma 14. $H_v \approx_c H_{v-1}$, for $v \in [z]$.

Proof. The difference of H_v and H_{v-1} lies in the sampling of $\mathbf{D}_v^0, \mathbf{D}_v^1$ and the outputs $\{y^{(k)}\}$. We first examine the outputs. For $k \in [t]$, we express the output $y^{(k)}$, starting from the expression in H_v to the one in H_{v-1} :

$$\begin{aligned}
y^{(k)} &= \left[\mathbf{J} \times \left(\left(\prod_{j=v+1}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^v s_i^{x_{\iota(i)}^{(k)}} \right) \times \mathbf{A}_{v+1} \prod_{j=v+1}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p \\
&= \left[\mathbf{J} \times \left(\left(\prod_{j=v+1}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^{v-1} s_i^{x_{\iota(i)}^{(k)}} \right) \times \left(\mathbf{I}^{w \times w} \otimes s_v^{x_{\iota(v)}^{(k)}} \right) \mathbf{A}_{v+1} \prod_{j=v+1}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p \\
&= \left[\mathbf{J} \times \left(\left(\prod_{j=v+1}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^{v-1} s_i^{x_{\iota(i)}^{(k)}} \right) \right. \\
&\quad \left. \times \left(\mathbf{B}_v^{x_{\iota(v)}^{(k)}} \otimes 1_R \right)^{-1} \times \left(\mathbf{B}_v^{x_{\iota(v)}^{(k)}} \otimes s_v^{x_{\iota(v)}^{(k)}} \right) \mathbf{A}_{v+1} \prod_{j=v+1}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p \\
&\approx_s \left[\mathbf{J} \times \left(\left(\prod_{j=v}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^{v-1} s_i^{x_{\iota(i)}^{(k)}} \right) \right. \\
&\quad \left. \times \left[\left(\mathbf{B}_v^{x_{\iota(v)}^{(k)}} \otimes s_v^{x_{\iota(v)}^{(k)}} \right) \mathbf{A}_{v+1} + \mathbf{E}' \right] \prod_{j=v+1}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p \\
&= \left[\mathbf{J} \times \left(\left(\prod_{j=v}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^{v-1} s_i^{x_{\iota(i)}^{(k)}} \right) \times \mathbf{Y}_v^{x_{\iota(v)}^{(k)}} \prod_{j=v+1}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p \\
&= \left[\mathbf{J} \times \left(\left(\prod_{j=v}^z \mathbf{B}_j^{x_{\iota(j)}^{(k)}} \right)^{-1} \otimes \prod_{i=1}^{v-1} s_i^{x_{\iota(i)}^{(k)}} \right) \times \mathbf{A}_v \prod_{j=v}^z \mathbf{D}_j^{x_{\iota(j)}^{(k)}} \right]_p
\end{aligned} \tag{4}$$

where $\mathbf{Y}_v^{x_{\iota(v)}^{(k)}} = \mathbf{A}_v \mathbf{D}_v^{x_{\iota(v)}^{(k)}}$. The correctness of this equation is a routine check. The implication is that the difference of H_v and H_{v-1} fully lies in the sampling of $\mathbf{Y}_v^0, \mathbf{Y}_v^1$ (being GLWE samples in H_v or uniform in H_{v-1}) and their preimages $\mathbf{D}_v^0, \mathbf{D}_v^1$ sampled by the trapdoor of \mathbf{A}_v .

Formally, suppose there is an attacker A' that distinguishes H_v and H_{v-1} with non-negligible probability ζ , we can build an attacker A who distinguishes:

$$\mathbf{A}_{v+1}, \{\mathbf{Y}_v^b = (\mathbf{B}_v^b \otimes s_v^b) \mathbf{A}_{v+1} + \mathbf{E}_v^b\}_{b \in \{0,1\}}$$

from

$$\{U(R_q^{w \times wm} \times R_q^{w \times wm} \times R_q^{w \times wm})\}$$

To do so, once A obtains the samples, it samples $\{\mathbf{A}_j\}_{j \in [v]}$ with trapdoors, and produce the preimages $\{\mathbf{D}_v^b \leftarrow \text{PreimgSam}(\mathbf{A}_v, \tau_v, \mathbf{Y}_v^b, \sigma)\}_{b \in \{0,1\}}$. Then places $\mathbf{A}_{v+1}, \mathbf{Y}_v^0, \mathbf{Y}_v^1, \mathbf{D}_v^0, \mathbf{D}_v^1$ in the constrained key and the outputs. It further samples GLWE secrets $\{s_i^b \leftarrow \eta\}_{b \in \{0,1\}, i \in [v]}, \{\mathbf{D}_i^b \leftarrow \text{PreimgSam}(\mathbf{A}_i, \tau_i, \mathbf{Y}_i^b, \sigma)\}_{b \in \{0,1\}, i \in [v]}$ as the constrained-key below level v . Samples the rest of the \mathbf{D} matrices obliviously $\{\mathbf{D}_i^b \leftarrow \gamma\}_{b \in \{0,1\}, i \in [v+1, z]}$.

With these matrices the attacker A is able to simulate the rest of the outputs, send the outputs and constrained key to A' . If the samples are from GLWE, then it corresponds to \mathbf{H}_v ; if the samples are uniform, then the matrices $\{\mathbf{D}_v^b\}_{b \in \{0,1\}}$ sampled via $\{\mathbf{D}_v^b \leftarrow \text{PreimgSam}(\mathbf{A}_v, \tau_v, \mathbf{Y}_v^b, \sigma)\}_{b \in \{0,1\}}$ are statistically close to the obliviously sampled ones due to Lemma 4, so it is statistically close to \mathbf{H}_{v-1} . Hence A breaks GLWE with probability more than $\zeta/2$, which contradicts to Lemma 10. \square

Lemma 15. *If $C(x^{(k)}) = 0$, then the output $y^{(k)}$ in \mathbf{H}_0 is pseudorandom.*

Proof. Following Eqn. 3, a constrained output $y^{(k)}$ in \mathbf{H}_0 can be expressed as:

$$y^{(k)} = \left[\mathbf{J} \times (\mathbf{P}^{-1} \otimes 1_R) \times \mathbf{A}_1 \prod_{j=1}^z \mathbf{D}_j^{x_{i(j)}^{(k)}} \right]_p \approx_s \left[(\mathbf{J} \times (\mathbf{P}^{-1} \otimes 1_R) \times \mathbf{A}_1 + \mathbf{E}) \prod_{j=1}^z \mathbf{D}_j^{x_{i(j)}^{(k)}} \right]_p \quad (5)$$

For $\mathbf{J}\mathbf{A}_1 + \mathbf{E}_J$ as part of the constrained key, $\mathbf{J} \times (\mathbf{P}^{-1} \otimes 1_R) \times \mathbf{A}_1 + \mathbf{E}$ as part of the constrained output $y^{(k)}$, $(\mathbf{J}\mathbf{A}_1 + \mathbf{E}_J, \mathbf{J} \times (\mathbf{P}^{-1} \otimes 1_R) \times \mathbf{A}_1 + \mathbf{E})$ is indistinguishable from $U(R_q^{1 \times wm}, R_q^{1 \times wm})$ due to Lemma 11. This means each constrained output $y^{(k)}$ is indistinguishable from $\left[\mathbf{U} \prod_{j=1}^z \mathbf{D}_j^{x_{i(j)}^{(k)}} \right]_p$ where $\mathbf{U} \leftarrow U(R_q^{1 \times wm})$. Hence $y^{(k)}$ is pseudorandom if $C(x^{(k)}) = 0$ due to Lemma 9. \square

The proof completes by combining the Lemmas 14 and 15. \square

6 Private-key functional encryption from CHCPRF

We construct private-key function-hiding functional encryptions for NC^1 circuits from (1) CHCPRFs for NC^1 ; (2) semantic secure private-key encryption schemes with decryption in NC^1 . The scheme satisfies 1-key simulation-based security.

6.1 The definition of functional encryption

Definition 8 (Function-hiding private-key functional encryption [34]).

A functional encryption scheme for a class of functions $\mathcal{C}_\mu = \{C : \{0,1\}^\mu \rightarrow \{0,1\}\}$ is a tuple of p.p.t. algorithms (Setup, FSKGen, Enc, Dec) such that:

- $\text{Setup}(1^\lambda)$ takes as input the security parameter 1^λ , outputs the master secret key MSK.
- $\text{FSKGen}(\text{MSK}, C)$ takes MSK and a function $C \in \mathcal{C}_\mu$, outputs a functional decryption key FSK_C .
- $\text{Enc}(\text{MSK}, m)$ takes MSK and a message $m \in \{0,1\}^\mu$, outputs a ciphertext CT_m .
- $\text{Dec}(\text{FSK}_C, \text{CT}_m)$ takes as input a ciphertext CT_m and a functional decryption key FSK_C , outputs (in the clear) the result $C(m)$ of applying the function on the message.

We require that:

Correctness. For every message $m \in \{0, 1\}^\mu$ and function $C \in \mathcal{C}_\mu$ we have:

$$\Pr \left[b = C(m) \left| \begin{array}{l} \text{MSK} \leftarrow \text{Setup}(1^\lambda) \\ \text{FSK}_C \leftarrow \text{FSKGen}(\text{MSK}, C) \\ \text{CT}_m \leftarrow \text{Enc}(\text{MSK}, m) \\ b \leftarrow \text{Dec}(\text{FSK}_C, \text{CT}_m) \end{array} \right. \right] = 1 - \text{negl}(\lambda)$$

Security. We require that for all polytime, stateful algorithm Adv , there is a polytime, stateful algorithm Sim such that:

$$\{\text{Experiment } \text{REAL}_{\text{Adv}}(1^\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Experiment } \text{IDEAL}_{\text{Adv}, \text{Sim}}(1^\lambda)\}_{\lambda \in \mathbb{N}}$$

The real and ideal experiments of stateful algorithms Adv, Sim are as follow:

Experiment $\text{REAL}_{\text{Adv}}(1^\lambda)$

$\text{MSK} \leftarrow \text{Gen}(1^\lambda),$

Repeat :

$\text{Adv} \rightarrow (m, d_m); \text{Adv} \leftarrow \text{Enc}(\text{MSK}, m);$

$\text{Adv} \rightarrow C;$

 if $d_m \neq C(m)$ for some m then Output \perp

 else $\text{Adv} \leftarrow \text{FSK}_C = \text{FSKGen}(\text{MSK}, C);$

Repeat :

$\text{Adv} \rightarrow m; \text{Adv} \leftarrow \text{Enc}(\text{MSK}, m)$

$\text{Adv} \rightarrow b;$ Output b

Experiment $\text{IDEAL}_{\text{Adv}, \text{Sim}}(1^\lambda)$

$\text{Sim} \leftarrow 1^\lambda$

Repeat :

$\text{Adv} \rightarrow (m, d_m); \text{Adv} \leftarrow \text{Sim}(1^{|m|}, d_m);$

$\text{Adv} \rightarrow C;$

 if $d_m \neq C(m)$ for some m then Output \perp

 else $\text{Adv} \leftarrow \text{FSK}_S = \text{Sim}(1^{|C|});$

Repeat :

$\text{Adv} \rightarrow m; \text{Adv} \leftarrow \text{Sim}(1^{|m|}, C(m))$

$\text{Adv} \rightarrow b;$ Output b

That is, in the experiments Adv can ask for a single functional decryption key and polynomially many input queries, in any order. For encryption queries m made before the decryption key query, Adv is expected to provide a bit d_x indicating whether $C(m) = 1$. In the real experiment Adv obtains the encryption of m . In the ideal experiment Adv obtains a value generated by Sim , whereas Sim is given only $1^{|m|}$ and d_m . Once Adv makes the functional key query for circuit $C \in \mathcal{C}_\lambda$, both experiments verify the consistency of the indicator bits d_m for all the encryption queries m made by Adv so far. If any inconsistency is found then the experiment halts. Next, in the real experiment Adv obtains the constrained key generated by the constraining algorithm; in the ideal experiment Adv obtains a key generated by Sim , whereas Sim is given only the size of C . The handling of encryption queries made by Adv after the circuit query is similar to the ones before, with the exception that the indicator bit d_m is no longer needed and Sim obtains the value of $C(m)$ instead. The output of the experiment is the final output bit of Adv .

6.2 The construction

Theorem 13. *If there are 1-key secure constraint-hiding constraint PRFs for constraint class \mathcal{C} , and symmetric-key encryption schemes with decryption in*

the class \mathcal{C} , then there are 1-key secure private-key function-hiding functional encryptions for function class \mathcal{C} .

Corollary 2. *Assuming the intractability of GLWE, there are 1-key secure private-key function-hiding functional encryptions for NC^1 .*

Construction 14 *Given a CHCPRF $(\text{F.Gen}, \text{F.Constrain}, \text{F.Eval})$, a semantic secure symmetric-key encryption scheme $(\text{Sym.Gen}, \text{Sym.Enc}, \text{Sym.Dec})$, we build a private-key functional encryption FE as follows:*

- $\text{FE.Setup}(1^\lambda)$ takes as input the security parameter 1^λ , runs $\text{Sym.Gen}(1^\lambda) \rightarrow \text{Sym.SK}$, $\text{F.Gen}(1^\lambda) \rightarrow \text{F.MSK}$, outputs the master secret key $\text{FE.MSK} = (\text{Sym.SK}, \text{F.MSK})$.
- $\text{FE.Enc}(\text{FE.MSK}, m)$ parses $\text{FE.MSK} = (\text{Sym.SK}, \text{F.MSK})$, computes $\text{Sym.CT} = \text{Sym.Enc}(m)$, $\text{Tag} = \text{F.Eval}(\text{F.MSK}, \text{Sym.CT})$. Outputs $\text{FE.CT} = (\text{Sym.CT}, \text{Tag})$.
- $\text{FE.FSKGen}(\text{FE.MSK}, C)$ parses $\text{FE.MSK} = (\text{Sym.SK}, \text{F.MSK})$, outputs the functional decryption key $\text{FE.FSK}_C = \text{F.Constrain}(\text{F.MSK}, F[\text{Sym.SK}, C])$, where the functionality of $F[\text{Sym.SK}, C](\cdot)$ is:
 - On input x , computes $\text{Sym.Dec}(\text{Sym.SK}, x) \rightarrow m \in \{0, 1\}^\mu \cap \perp$;
 - if $m = \perp$, return 0; else, return $C(m)$.
- $\text{FE.Dec}(\text{FE.FSK}_C, \text{FE.CT})$ parses $\text{FE.FSK}_C = \text{F.CK}_F$, $\text{FE.CT} = (\text{Sym.CT}, \text{Tag})$, computes $T = \text{F.Eval}(\text{F.CK}_F, \text{Sym.CT})$. Outputs 1 if $T = \text{Tag}$, 0 if not.

Correctness. Correctness follows the correctness of Sym and F.

Proof. We build the FE simulator FE.Sim from the symmetric-key encryption simulator Sym.Sim and CHCPRF simulator F.Sim :

1. Generates the simulated master secret-keys Sym.SK^S and F.MSK^S
2. Given a function-decryption key query (for function C), FE.Sim runs $\text{CK}^S \leftarrow \text{F.Sim}_1(1^\lambda, 1^{|F[\text{Sym.SK}, C]|}, \text{F.MSK}^S)$, outputs CK^S as FE.FSK^S .
3. Given a ciphertext query and the output bit $C(m)$, FE.Sim runs $\text{Sym.CT}^S \leftarrow \text{Sym.Sim}(1^\lambda, 1^{|m|}, \text{Sym.SK}^S)$ and $\text{Tag}^S \leftarrow \text{FSim}_2(\text{F.MSK}^S, \text{CK}^S, \text{Sym.CT}^S, C(m))$, outputs $(\text{Sym.CT}^S, \text{Tag}^S)$ as FE.CT^S .

To show that the simulated outputs are indistinguishable from the real outputs, consider an intermediate simulator $\text{FE.Sim}'$ which is the same to FE.Sim , except that it uses the real Sym ciphertexts in the ciphertext queries. Observe that the secret-key of Sym is not exposed in $\text{FE.Sim}'$ or FE.Sim , the output distributions of $\text{FE.Sim}'$ and FE.Sim are indistinguishable following the security of Sym.

Next, assume there is a distinguisher D for the outputs of the real FE scheme and $\text{FE.Sim}'$, we build an attacker A for the CHCPRF F. A samples a secret key for Sym, sends a constrained circuit query, obtains the real CK if it is the real distribution, or the simulated CK^S if it is the simulated distribution; then creates symmetric-key ciphertexts, sends as the input queries to the CHCPRF. It obtains the real outputs if it is the real case, or the simulated outputs if it is the simulated case. A treats the outputs as tags. A forwards the ciphertexts, tags and FSK to D . D 's success probability transfers to the one for A . \square

Acknowledgments

We thank Leonid Reyzin for helpful discussions and the formalization of LWE-hardness. The first author is a member of the Check Point Institute for Information Security. This work is supported by the NSF MACS project, NSF grants CNS1012798, CNS1012910, and ISF grant 1523/14.

References

1. Shweta Agrawal, Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Functional encryption: New perspectives and lower bounds. In *CRYPTO (2)*, volume 8043 of *LNCS*, pages 500–518. Springer, 2013.
2. Miklós Ajtai. Generating hard instances of the short basis problem. In Jirí Wiedermann, Peter van Emde Boas, and Mogens Nielsen, editors, *Automata, Languages and Programming, 26th International Colloquium, ICALP'99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, volume 1644 of *LNCS*, pages 1–9. Springer, 1999.
3. Joël Alwen and Chris Peikert. Generating shorter bases for hard random lattices. *Theory of Computing Systems*, 48(3):535–553, 2011.
4. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Advances in Cryptology-CRYPTO 2009*, pages 595–618. Springer, 2009.
5. Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 353–370, 2014.
6. Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT*, volume 7237 of *LNCS*, pages 719–737. Springer, 2012.
7. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6, 2012.
8. David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in nc^1 . In Juris Hartmanis, editor, *STOC*, pages 1–5. ACM, 1986.
9. Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 893–902. SIAM, 2016.
10. Dan Boneh, Sam Kim, and Hart Montgomery. Private puncturable prfs from standard lattice assumptions. To be appear at Eurocrypt 2017. Cryptology ePrint Archive, Report 2017/100, 2017.
11. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In Canetti and Garay [22], pages 410–428.
12. Dan Boneh, Kevin Lewi, and David J. Wu. Constraining pseudorandom functions privately. To be appear at PKC 2017, Cryptology ePrint Archive, Report 2015/1167, 2017.

13. Dan Boneh, Amit Sahai, and Brent Waters. Functional encryption: Definitions and challenges. In *Theory of Cryptography Conference*, pages 253–273. Springer, 2011.
14. Dan Boneh and Alice Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324(1):71–90, 2003.
15. Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In *Advances in Cryptology-ASIACRYPT 2013*, pages 280–300. Springer, 2013.
16. Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *Public Key Cryptography*, volume 8383 of *LNCS*, pages 501–519. Springer, 2014.
17. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325. ACM, 2012.
18. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 575–584. ACM, 2013.
19. Zvika Brakerski and Vinod Vaikuntanathan. Constrained key-homomorphic prfs from standard lattice assumptions. In *Theory of Cryptography*, pages 1–30. Springer, 2015.
20. Zvika Brakerski, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Obfuscating conjunctions under entropic ring LWE. In *ITCS*, pages 147–156. ACM, 2016.
21. Ran Canetti and Yilei Chen. Constraint-hiding constrained prfs for nc^1 from lwe. To be appear at Eurocrypt 2017. Cryptology ePrint Archive, 2017.
22. Ran Canetti and Juan A. Garay, editors. *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, volume 8042 of *LNCS*. Springer, 2013.
23. Yilei Chen, Craig Gentry, and Shai Halevi. Cryptanalyses of candidate branching program obfuscators. To be appear at Eurocrypt 2017. Cryptology ePrint Archive, Report 2016/998, 2017.
24. Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In *EUROCRYPT (1)*, volume 9056 of *LNCS*, pages 3–12. Springer, 2015.
25. Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. Cryptanalysis of GGH15 multilinear maps. In *CRYPTO (2)*, volume 9815 of *LNCS*, pages 607–628. Springer, 2016.
26. Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Canetti and Garay [22], pages 476–493.
27. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT (2)*, volume 9666 of *LNCS*, pages 559–585. Springer, 2016.
28. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-svp. To be appear at Eurocrypt 2017. Cryptology ePrint Archive, Report 2016/885, 2017.
29. Léo Ducas and Alain Durmus. Ring-lwe in polynomial rings. In *International Workshop on Public Key Cryptography*, pages 34–51. Springer, 2012.
30. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Johansson and Nguyen [38], pages 1–17.
31. Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *Theory of Cryptography*, pages 498–527. Springer, 2015.

32. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
33. Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986.
34. Shafi Goldwasser, Yael Kalai, Raluca Ada Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 555–564. ACM, 2013.
35. Satoshi Hada. Zero-knowledge and code obfuscation. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 443–457. Springer, 2000.
36. Shai Halevi. Graded encoding, variations on a scheme. Cryptology ePrint Archive, Report 2015/866, 2015.
37. Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. In *EUROCRYPT (1)*, volume 9665 of *LNCS*, pages 537–565. Springer, 2016.
38. Thomas Johansson and Phong Q. Nguyen, editors. *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, volume 7881 of *LNCS*. Springer, 2013.
39. Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *ACM Conference on Computer and Communications Security*, pages 669–684. ACM, 2013.
40. Philip N. Klein. Finding the closest lattice vector when it’s unusually close. In *SODA*, pages 937–941. ACM/SIAM, 2000.
41. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.
42. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.
43. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. In Johansson and Nguyen [38], pages 35–54.
44. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology–EUROCRYPT 2012*, pages 700–718. Springer, 2012.
45. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measure. *SIAM Journal on Computing*, 37(1):267–302, 2007.
46. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342. ACM, 2009.
47. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Theory of Cryptography*, pages 145–166. Springer, 2006.
48. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.