Succinct LWE Sampling, Random Polynomials, and Obfuscation

Lalita Devadas[1], Willy Quach[2], Vinod Vaikuntanathan[1], Hoeteck Wee[3], and Daniel Wichs[2,3]

[1] Massachusetts Institute of Technology, Cambridge, MA 02139, USA
{lali,vinodv}@mit.edu
[2] Northeastern University, Boston, MA 02115, USA
quach.w@northeastern.edu, wichs@ccs.neu.edu
[3] NTT Research, Sunnyvale, CA 94085, USA
wee@di.ens.fr

Abstract. We present a construction of indistinguishability obfuscation (iO) that relies on the learning with errors (LWE) assumption together with a new notion of succinctly sampling pseudorandom LWE samples. We then present a candidate LWE sampler whose security is related to the hardness of solving systems of polynomial equations. Our construction improves on the recent iO candidate of Wee and Wichs (Eurocrypt 2021) in two ways: first, we show that a much weaker and simpler notion of LWE sampling suffices for iO; and secondly, our candidate LWE sampler is secure based on a compactly specified and falsifiable assumption about random polynomials, with a simple error distribution that facilitates cryptanalysis.

Keywords: Indistinguishability Obfuscation · Learning With Errors

1 Introduction

Indistinguishability obfuscation (iO) [BGI+01,GR07] is a probabilistic polynomial-time algorithm $\mathcal{O}$ that takes as input a circuit C and outputs an (obfuscated) circuit $C' = \mathcal{O}(C)$ satisfying two properties: (a) *functionality*: C and C' compute the same function; and (b) *security*: for any two circuits C_1 and C_2 that compute the same function (and have the same size), $\mathcal{O}(C_1)$ and $\mathcal{O}(C_2)$ are computationally indistinguishable. Since the first candidate for iO was introduced in [GGH+13], a series of works have shown that iO would have a huge impact on cryptography.

In this work, we build upon the recent line of works on lattice-inspired iO candidates [Agr19, CHVW19, AP20, BDGM20b, BDGM20a, WW21, GP21] that are plausibly post-quantum secure. The dream goal here is to ultimately base iO on the hardness of the learning with errors (LWE) problem together with an assumption about simple Boolean or integer pseudorandom generators (PRGs). Such a result would, in particular, eliminate pairings from the recent breakthrough result basing iO on well-founded assumptions [JLS21].

1.1 Our Contributions

We present a candidate construction of iO that relies on LWE together with a new notion of succinctly sampling pseudorandom LWE samples. In addition, we present a candidate sampler whose security is related to the hardness of solving systems of polynomial equations. Our construction improves on the recent iO candidate of Wee and Wichs [WW21] (henceforth referred to as the WW construction) in two ways:

- First, our new notion of succinct LWE sampling simplifies and relaxes the notion of oblivious LWE sampling from WW. Instead of a simulation-based definition as in WW, we have a simple indistinguishability-based definition, where the generated LWE sample can be used to drown out the differences between certain error distributions. Furthermore, we put forth two variants of succinct LWE sampling, and provide a general amplification from a weak (falsifiable) notion that refers to a specific error distribution to a strong (non-falsifiable) notion that refers to general error distributions.
- Next, our candidate succinct LWE sampler is easy to describe and is based on random polynomials. It yields an LWE sample with a simple error distribution that facilitates cryptanalysis. This is in contrast to WW, where the LWE sampler involved complex FHE evaluation, and the resulting error distribution in the samples was dependent on the concrete implementation of the circuit being evaluated. Indeed, a recent work of [HJL21] carefully crafted circuit implementations that would render the WW candidate as well as the related candidate in [GP21] insecure (see Section 1.3 for a more detailed discussion).

1.2 Technical Overview

The starting point of our construction is essentially the same as that of the Wee-Wichs (WW) iO candidate, which in turn builds on [BDGM20a]. We begin by describing a notion of succinct randomized encoding (SRE), which can be seen as a relaxation of the notions of split FHE and functional encodings used in prior works. It is also very related to the notion of exponentially efficient iO (XiO) from [LPST16], and is easily seen to imply it, but we find the SRE abstraction easier to work with in the context of our work. By leveraging prior results on XiO [LPST16], our notion of SRE implies iO under the LWE assumption.

Succinct Randomized Encodings. A succinct randomized encoding[4] [BGL+15, LPST16] of a function $f : \{0,1\}^\ell \to \{0,1\}^N$ is an efficient probabilistic algorithm Encode such that:

- functionality: we can efficiently recover $f(x)$ given f and $\mathsf{Encode}(f, x)$;

[4] Our notion of succinct randomized encodings is weaker than prior works: indeed, [BGL+15] required the encoder to run in time sublinear in N, whereas we allow the encoder run-time to be polynomial in N.

– security: for any x_0, x_1 such that $f(x_0) = f(x_1)$, we have $\mathsf{Encode}(f, x_0) \approx_c \mathsf{Encode}(f, x_1)$; and
– succinctness: $\mathsf{Encode}(f, x)$ is shorter than the output length of f. That is, $|\mathsf{Encode}(f, x)| = \widetilde{O}(N^\delta)$ for some constant $\delta < 1$, ignoring factors polynomial in ℓ and the security parameter.

Henceforth, we will focus on building SRE for circuits.

Base Scheme. We start with a base scheme for succinct randomized encodings implicit in WW, which is insecure, but serves as the basis of our eventual construction. The base scheme uses a variant of the homomorphic encryption/commitment schemes of [GSW13, GVW15], along with the "packing" techniques in [PVW08, MW16, BTVW17, PS19, GH19, BDGM19]. Given a commitment $\mathbf{C}$ to an input $x \in \{0,1\}^\ell$, along with a circuit $f : \{0,1\}^\ell \to \{0,1\}^N$, this scheme allows us to homomorphically compute a commitment $\mathbf{C}_f$ to the output $f(x)$. Moreover, the opening for the output commitment is shorter than the output size N. Concretely, we define $\mathbf{C}, \mathbf{C}_f$ as follows:

– We treat the function $f : \{0,1\}^\ell \to \{0,1\}^N$ as a function $f : \{0,1\}^\ell \to \{0,1\}^{M \times K}$, where M and K are parameters we shall specify shortly, such that $MK = N$.
– Given a public random matrix $\mathbf{A} \in \mathbb{Z}_q^{M \times w}$ where $M \gg w$, we define a commitment $\mathbf{C}$ to an input x as

$$\mathbf{C} := \mathbf{A}\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E}$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{M \times w}$, $\mathbf{R} \leftarrow \mathbb{Z}_q^{w \times \ell M \log q}$ are uniformly random, $\mathbf{E} \leftarrow \chi^{M \times \ell M \log q}$ has its entries chosen from an error distribution χ, $\mathbf{G} \in \mathbb{Z}_q^{M \times M \log q}$ is the gadget matrix [MP12], and we treat x as a row vector of length ℓ in $x \otimes \mathbf{G}$.
– Homomorphic evaluation of f on $\mathbf{C}$ yields $\mathbf{C}_f$ satisfying

$$\mathbf{C}_f = \mathbf{A}\mathbf{R}_{f,x} + \mathbf{E}_{f,x} + f(x) \cdot \tfrac{q}{2} \in \mathbb{Z}_q^{M \times K} \tag{1}$$

where $f(x) \in \{0,1\}^{M \times K}$, $\mathbf{R}_{f,x} \in \mathbb{Z}_q^{w \times K}$ and $\mathbf{E}_{f,x}$ has small entries.

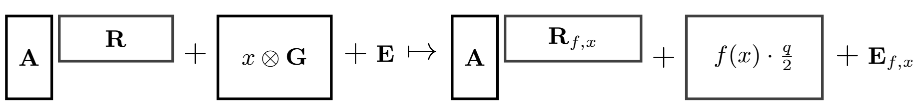

Our base scheme[5] simply outputs

$$\mathbf{A},\ \mathbf{C} := \mathbf{A}\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E},\ \mathbf{R}_{f,x}$$

as the encoding of x. Decoding computes $\mathbf{C}_f$ given $(\mathbf{C}, f)$, subtracts $\mathbf{A} \cdot \mathbf{R}_{f,x}$ to obtain $f(x) \cdot \frac{q}{2}$ plus error (following equation 1) and rounds to obtain $f(x)$.

[5] In the WW terminology, this would be a candidate K-sim functional encoding for $f_1, \ldots, f_K : \{0,1\}^\ell \to \{0,1\}^M$.

The encoding is also succinct: The total size of the encoding (in bits) is

$$O((Mw + M^2\ell + wK) \cdot \log q).$$

Setting $M = N^{1/3}, K = N^{2/3}, w = O(\lambda)$ yields encoding size $\widetilde{O}(N^{2/3})$, where $\widetilde{O}(\cdot)$ hides polynomial factors in λ, ℓ and the depth of the circuit computing f.

The scheme is, however, *completely insecure* as written because, given $\mathbf{C}, \mathbf{R}_{f,x}$ and a "guess" for x, we can recover $\mathbf{R}$ by solving a system of linear equations, and test if our guess was correct (see WW). This allows us to easily distinguish between encodings of any x_0 and x_1.

"Pseudorandom" LWE Sampling. Following [WW21], we fix the insecurity of the base scheme by masking $\mathbf{R}_{f,x}$ using a "pseudorandom" LWE sample; similar ideas were used in several prior works [BDGM20a, GP21, JLS21, AR17, Agr19, JLMS19, AJL+19] with "pseudorandom" noise. That is, we generate a "pseudorandom" LWE sample $\mathbf{B}^* = \mathbf{A}\mathbf{S}^* + \mathbf{E}^* \in \mathbb{Z}_q^{M \times K}$ and output

$$\mathsf{seed}_{\mathbf{B}^*},\ \mathbf{A},\ \mathbf{A}\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E},\ \mathbf{R}_{f,x} + \mathbf{S}^* \tag{2}$$

where $\mathsf{seed}_{\mathbf{B}^*}$ is a succinct description of $\mathbf{B}^*$, with $|\mathsf{seed}_{\mathbf{B}^*}| \leq (MK)^\delta$ for some $\delta < 1$. Correctness now relies on the fact that

$$\mathbf{A} \cdot (\mathbf{R}_{f,x} + \mathbf{S}^*) \approx \mathbf{B}^* + \mathbf{C}_f + f(x) \cdot \tfrac{q}{2}.$$

WW's security requirement for the pseudorandom LWE sample, "oblivious LWE sampling", was cumbersome to define, required a simulator, and only made sense in the common reference string model. The reliance on a simulator means the definition did not have an inherently falsifiable format that enables demonstrating insecurity by constructing an efficient attacker. Here, we reformulate a simpler and falsifiable variant that we call "succinct LWE sampling".[6]

Defining pseudorandom LWE sampling, in WW and in our work, is difficult because we want $\mathbf{B}^* = \mathbf{A}\mathbf{S}^* + \mathbf{E}^*$ to look like a random LWE sample, but this is impossible since it is succinctly described in $\mathsf{seed}_{\mathbf{B}^*}$. Instead, we essentially want $\mathbf{E}^*$ to drown out the difference between any two sufficiently small error distributions $\mathbf{Z}_0$ and $\mathbf{Z}_1$, in the sense that $\mathsf{seed}_{\mathbf{B}^*}, \mathbf{E}^* - \mathbf{Z}_b$ hides b. Unfortunately, this too is impossible, since $\mathsf{seed}_{\mathbf{B}^*}$ lets us get $\mathbf{B}^* = \mathbf{A}\mathbf{S}^* + \mathbf{E}^*$ from which we can then derive $\mathbf{A}\mathbf{S}^* + \mathbf{Z}_b$; this allows us to distinguish between (say) $\mathbf{Z}_0 = 0$ and $\mathbf{Z}_1$ being a small Gaussian by checking rank. Our main observation is that we don't need indistinguishability to hold for worst-case distributions $\mathbf{Z}_b$, but rather only for ones where an LWE sample $\mathbf{A}\mathbf{R} + \mathbf{Z}_b$ with the error $\mathbf{Z}_b$ and a truly random $\mathbf{R}$ would hide the bit b. Formally, the definition says that for any two distributions of $(\mathbf{Z}_b, \mathsf{aux}_b)$ where $\mathbf{Z}_b$ is sufficiently short:

$$\textbf{If} \quad (\mathsf{aux}_0, \mathbf{A}, \mathbf{A}\mathbf{R} + \mathbf{Z}_0) \approx_c (\mathsf{aux}_1, \mathbf{A}, \mathbf{A}\mathbf{R} + \mathbf{Z}_1), \tag{3}$$

$$\textbf{then} \quad (\mathsf{seed}_{\mathbf{B}^*}, \mathsf{aux}_0, \mathbf{A}, \mathbf{E}^* - \mathbf{Z}_0) \approx_c (\mathsf{seed}_{\mathbf{B}^*}, \mathsf{aux}_1, \mathbf{A}, \mathbf{E}^* - \mathbf{Z}_1). \tag{4}$$

[6] It is simpler in terms of syntax, since we do not refer to LWE trapdoors for $\mathbf{A}$, and in terms of the security requirement since we do not require a simulator, but instead have a simple indistinguishability criterion.

Note that, since $\mathsf{seed}_{\mathbf{B}^*}$ defines $\mathbf{AS}^* + \mathbf{E}^*$, giving $\mathbf{E}^* - \mathbf{Z}_b$ in (4) is equivalent to giving $\mathbf{AS}^* + \mathbf{Z}_b$, and hence we use these interchangeably in the definition.

The above definition is not falsifiable since it quantifies over all $(\mathsf{aux}_b, \mathbf{Z}_b)$ satisfying the pre-condition (3). However, we also consider a weaker, falsifiable definition, where we fix a specific $(\mathsf{aux}^*_b, \mathbf{Z}^*_b)$ that satisfies the pre-condition (3). We then show a generic transformation that lifts any scheme realizing the weak definition into one that realizes the general definition. Specifically, in the weak definition, we fix $\mathsf{aux}^*_b = (\widehat{\mathbf{B}}, \mathbf{C})$ to consist of a commitment $\widehat{\mathbf{B}}$ to 0, along with a commitment $\mathbf{C}$ to $-b$. We then homomorphically evaluate an AND operation (multiplication) on the commitments $\widehat{\mathbf{B}}, \mathbf{C}$, which results in a commitment to 0, and we define $\mathbf{Z}^*_b$ to be the error term for this commitment. Formally,

$$\mathsf{aux}^*_b = \left(\widehat{\mathbf{B}} = \mathbf{AS}_0 + \mathbf{F}, \quad \mathbf{C} = \mathbf{AR} + \mathbf{E} - b\mathbf{G}\right) \quad \text{and} \quad \mathbf{Z}^*_b = \mathbf{EG}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F},$$

where $\mathbf{E}$ and $\mathbf{F}$ are matrices with small entries. The transformation is inspired by a trick employed in WW to frame the security of their candidate oblivious LWE sampler construction as a falsifiable assumption. Here, we are able to abstract this trick out and formally prove that it amplifies a weak definition of security to a strong one. Therefore, we get a simple and falsifiable definition of succinct LWE sampling as our target. We refer to the full version for more details.

Our final definition introduces additional relaxations. Instead of a uniformly random matrix $\mathbf{A}$, we allow the use of matrices $\mathbf{A}^*$, which may not be uniformly random and can have some additional structure, as long as LWE still holds w.r.t. $\mathbf{A}^*$. We also allow the succinct sampler to rely on a non-succinct common reference string (CRS) of length $\mathrm{poly}(N)$. This is analogous to the reliance on a CRS in WW (as well as [BDGM20a, GP21]) and suffices for iO.

Our Succinct Randomized Encoding. To go from succinct LWE sampling to SRE, we essentially follow WW, and replace $\mathbf{A}$ with $\mathbf{A}^*$ in (2). The SRE consists of:

$$\mathsf{seed}_{\mathbf{B}^*}, \ \mathbf{A}^*, \quad \mathbf{A}^*\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E}, \quad \mathbf{R}_{f,x} + \mathbf{S}^* \ . \tag{5}$$

Correctness and succinctness follow readily as before. To prove security, we need to argue as follows that $\mathsf{Encode}(f, x_b)$ hides b as long as $f(x_0) = f(x_1)$.

– As long as $\mathbf{A}^*$ is full-rank, $(\mathbf{R}_{f,x_b} + \mathbf{S}^*)$ can be computed from $\mathbf{A}^*$ and $\mathbf{A}^* \cdot (\mathbf{R}_{f,x_b} + \mathbf{S}^*)$, so it suffices to argue that:

$$\mathsf{seed}_{\mathbf{B}^*}, \quad \mathbf{A}^*, \quad \mathbf{A}^*\mathbf{R} + x_b \otimes \mathbf{G} + \mathbf{E}, \quad \mathbf{A}^* \cdot (\mathbf{R}_{f,x_b} + \mathbf{S}^*)$$

hides b.

– Using $\mathbf{C}_f = \mathbf{A}^*\mathbf{R}_{f,x_b} + \mathbf{E}_{f,x_b} + f(x_b) \cdot \frac{q}{2}$ and deriving $\mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$ from $\mathsf{seed}_{\mathbf{B}^*}$, we can write

$$\mathbf{A}^* \cdot (\mathbf{R}_{f,x_b} + \mathbf{S}^*) = \mathbf{C}_f - f(x_b) \cdot \tfrac{q}{2} + \mathbf{B}^* - \mathbf{E}^* - \mathbf{E}_{f,x_b},$$

so it suffices to argue that

$$\mathsf{seed}_{\mathbf{B}^*}, \quad \mathbf{A}^*, \quad \mathbf{A}^*\mathbf{R} + x_b \otimes \mathbf{G} + \mathbf{E}, \quad \mathbf{E}^* + \mathbf{E}_{f,x_b}$$

hides b.

– At this point, we will invoke security of our succinct LWE sampler with

$$\mathsf{aux}_b = \mathbf{A}^*\mathbf{R} + x_b \otimes \mathbf{G} + \mathbf{E}, \qquad \mathbf{Z}_b = \mathbf{E}_{f,x_b}$$

For this step, we need to show that the pre-condition (3) holds:

$$(\mathbf{A}^*\mathbf{R}+x_0\otimes\mathbf{G}+\mathbf{E},\ \mathbf{A}^*,\ \mathbf{A}^*\mathbf{S}'+\mathbf{E}_{f,x_0}) \approx_c (\mathbf{A}^*\mathbf{R}+x_1\otimes\mathbf{G}+\mathbf{E},\ \mathbf{A}^*,\ \mathbf{A}^*\mathbf{S}'+\mathbf{E}_{f,x_1}).$$

This follows from LWE w.r.t. $\mathbf{A}^*$ and the fact that $\mathbf{A}^*\mathbf{S}' + \mathbf{E}_{f,x_b} \equiv \mathbf{A}^*\mathbf{S}' + \mathbf{C}_f - f(x_b) \cdot \frac{q}{2}$, where $f(x_0) = f(x_1)$.

Note that, in the above, we only relied on the security of the LWE sampler for the special case where aux_b is an encryption of x_b and $\mathbf{Z}_b$ is the error in the ciphertext one gets by homomorphically computing $f(x_b)$ for some function f such that $f(x_0) = f(x_1)$. However, as mentioned previously, we can also rely on an even more restricted form of $(\mathsf{aux}_b, \mathbf{Z}_b)$, essentially corresponding to the extremely simple case where f just computes the AND of b and 0, and generically lift security to the completely general case.

Our Candidate Succinct LWE Sampler. We want to design a succinct LWE sampler generating $\mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$. The security requirement in Equation (4) implies that $\mathbf{E}^* - \mathbf{Z}_b$ hides b for any short matrices $\mathbf{Z}_0, \mathbf{Z}_1$ satisfying some additional properties which we shall ignore in the rest of this overview. In addition, we want $\mathbf{B}^*$ to admit a short description $\mathsf{seed}_{\mathbf{B}^*}$, which means that $\mathbf{E}^* \in \mathbb{Z}^{M\times K}$ should compute a "pseudorandom" noise-flooding distribution.

Following [JLMS19, AJL+19], a good candidate for $\mathbf{E}^*$ is to evaluate MK random degree-d polynomials in dmk variables drawn from independent Gaussian distributions, where $MK \ll (dmk)^{d/2}$ to avoid linearization and potential sum-of-squares-based attacks; the ensuing distribution is plausibly indistinguishable from MK independent samples from a "noise-flooding" distribution $\mathcal{D}$ for a suitable choice of parameters. Concretely, thinking of d as a small constant, we sample "secret" Gaussian matrices $\mathbf{E}_1, \ldots, \mathbf{E}_d \leftarrow \chi^{m\times k}$ and public Gaussian matrices $\mathbf{P} \leftarrow \chi^{M\times m^d}$ and $\mathbf{P}' \leftarrow \chi^{k^d\times K}$ and we define

$$\mathbf{E}^* := \mathbf{P}(\mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \cdots \otimes \mathbf{E}_d)\mathbf{P}' \in \mathbb{Z}^{M\times K}$$

where $\mathbf{P}, \mathbf{P}'$ are published in the CRS. In the special case of $m = M = 1$ and $\mathbf{P} = 1$, the distribution of $\mathbf{E}^* \in \mathbb{Z}^K$ corresponds roughly to the evaluation of K random (i.e. Gaussian) degree-d (multilinear) polynomials in dk variables (where the dk variables are the entries of the $\mathbf{E}_1, \ldots, \mathbf{E}_d$ and the coefficients of the polynomial are specified by $\mathbf{P}'$). In the general case, we have a collection of polynomials, where each one looks at a certain structured set of monomials. For more details, see Section 4.5.

Next, we specify $(\mathbf{B}^*, \mathbf{A}^*, \mathbf{S}^*, \mathsf{seed}_{\mathbf{B}^*})$, starting with $\mathsf{seed}_{\mathbf{B}^*}$. Following [JLMS19], we additionally sample $\mathbf{A}_i \leftarrow \mathbb{Z}_q^{m\times w}, \mathbf{S}_i \leftarrow \mathbb{Z}_q^{w\times k}$ for $i = 1, \ldots, d$ and some $w \ll m, k$, and we define:

$$\mathsf{seed}_{\mathbf{B}^*} := (\mathbf{B}_1 := \mathbf{A}_1\mathbf{S}_1 + \mathbf{E}_1\ ,\ \ldots\ ,\ \mathbf{B}_d := \mathbf{A}_d\mathbf{S}_d + \mathbf{E}_d) \in (\mathbb{Z}_q^{m\times k})^d.$$

Inspired by the homomorphic operations of the Brakerski-Vaikuntanathan FHE [BV11], we want to relate $\mathbf{E}^*$ to $\mathbf{B}_1 \otimes \cdots \otimes \mathbf{B}_d$ and from there, derive $\mathbf{B}^*, \mathbf{A}^*, \mathbf{S}^*$ such that $\mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$ (we will discuss succinctness after that). We start with $d = 2$ for simplicity. By the mixed product property:

$$\begin{aligned}\mathbf{B}_1 \otimes \mathbf{B}_2 &= \mathbf{A}_1\mathbf{S}_1 \otimes \mathbf{B}_2 \ + \ \mathbf{E}_1 \otimes \mathbf{A}_2\mathbf{S}_2 \ + \ \mathbf{E}_1 \otimes \mathbf{E}_2 \\ &= [\mathbf{A}_1 \otimes \mathbf{I}_m \mid \mathbf{I}_m \otimes \mathbf{A}_2]\begin{pmatrix}\mathbf{S}_1 \otimes \mathbf{B}_2 \\ \mathbf{E}_1 \otimes \mathbf{S}_2\end{pmatrix} \ + \ \mathbf{E}_1 \otimes \mathbf{E}_2.\end{aligned}$$

We start by defining $\mathbf{B}^*$ and "pre-cursor" values $\overline{\mathbf{A}}^*, \overline{\mathbf{S}}^*$, which we will use to derive the final $\mathbf{A}^*, \mathbf{S}^*$ later, via:

$$\overbrace{\mathbf{P} \cdot (\mathbf{B}_1 \otimes \mathbf{B}_2) \cdot \mathbf{P}'}^{\mathbf{B}^*} = \overbrace{\mathbf{P}[\mathbf{A}_1 \otimes \mathbf{I}_m \mid \mathbf{I}_m \otimes \mathbf{A}_2]}^{\overline{\mathbf{A}}^*} \cdot \overbrace{\begin{pmatrix}\mathbf{S}_1 \otimes \mathbf{B}_2 \\ \mathbf{E}_1 \otimes \mathbf{S}_2\end{pmatrix}\mathbf{P}'}^{\overline{\mathbf{S}}^*} + \overbrace{\mathbf{P}(\mathbf{E}_1 \otimes \mathbf{E}_2)\mathbf{P}'}^{\mathbf{E}^*}$$

For general d, we have:

$$\mathbf{B}^* = \mathbf{P} \cdot (\mathbf{B}_1 \otimes \cdots \otimes \mathbf{B}_d) \cdot \mathbf{P}' \in \mathbb{Z}_q^{M \times K}, \ \ \mathbf{E}^* = \mathbf{P}(\mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \cdots \otimes \mathbf{E}_d)\mathbf{P}' \in \mathbb{Z}^{M \times K},$$

$$\overline{\mathbf{A}}^* = \mathbf{P} \cdot (\mathbf{A}_1 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \| \cdots\cdots \| \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \otimes \mathbf{A}_d) \in \mathbb{Z}_q^{M \times dwm^{d-1}},$$

$$\overline{\mathbf{S}}^* = \begin{pmatrix}\mathbf{S}_1 \otimes \mathbf{B}_2 \otimes \cdots \otimes \mathbf{B}_d \\ \mathbf{E}_1 \otimes \mathbf{S}_2 \otimes \cdots \otimes \mathbf{B}_d \\ \vdots \\ \mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \cdots \otimes \mathbf{S}_d\end{pmatrix} \cdot \mathbf{P}' \in \mathbb{Z}_q^{dwm^{d-1} \times K}, \text{ which we show satisfy}$$

$$\mathbf{B}^* = \overline{\mathbf{A}}^* \cdot \overline{\mathbf{S}}^* + \mathbf{E}^*.$$

Note that while the width of $\mathbf{A}$ in both the base scheme and WW is $w = \mathrm{poly}(\lambda)$, the width of $\overline{\mathbf{A}}^*$ is much larger and will in fact grow with N.

As mentioned above, it seems reasonable to conjecture that $\mathbf{E}^*$ on its own is pseudo-iid. However, $\overline{\mathbf{S}}^*$ is structured and does not look random on its own, which is problematic since we want $\overline{\mathbf{S}}^* + \mathbf{R}_{f,x}$ to drown out differences in the distribution of $\mathbf{R}_{f,x}$. Therefore, we will rely on a variant of Kilian randomization [Kil88] to hide the structure of $\overline{\mathbf{A}}^*, \overline{\mathbf{S}}^*$. We compute a random basis $\mathbf{A}^*$ of the column span of $\overline{\mathbf{A}}^*$ and then solve for $\mathbf{S}^*$ subject to $\mathbf{A}^*\mathbf{S}^* = \overline{\mathbf{A}}^* \cdot \overline{\mathbf{S}}^*$. This ensures that $\mathbf{A}^*, \mathbf{S}^*$ essentially do not reveal more than the product $\overline{\mathbf{A}}^*\overline{\mathbf{S}}^*$.

Succinctness. With the above implementation of succinct LWE sampling, from (5), the encodings of the resulting SRE have size

$$|\mathsf{Encode}(f, x)| = \widetilde{O}\left(\underbrace{M^2}_{\mathbf{A}^*\mathbf{R}+x\otimes\mathbf{G}+\mathbf{E}} + \underbrace{dmk}_{\mathsf{seed}_{\mathbf{B}^*}} + \underbrace{Mdwm^{d-1}}_{\mathbf{A}^*} + \underbrace{Kdwm^{d-1}}_{\mathbf{S}^*+\mathbf{R}_{f,x}}\right)$$

where $\widetilde{O}(\cdot)$ hides $\mathrm{poly}(\lambda, \log q, \ell)$ factors, which is in turn polynomial in λ, ℓ and circuit depth of f. We set

$$w = \mathrm{poly}(\lambda), \qquad m = N^{\frac{1}{2d}}, \qquad k = m^5 = N^{\frac{5}{2d}},$$
$$M = m^{d-1/2} = N^{\frac{1}{2}-\frac{1}{4d}}, \qquad K = m^{d+1/2} = N^{\frac{1}{2}+\frac{1}{4d}}.$$

Then, $|\mathsf{Encode}(f, x)| = \widetilde{O}(m^{2d-1/6}) = \widetilde{O}(N^{1-\frac{1}{12d}})$, that is, our scheme achieves $(1-\frac{1}{12d})$-succinctness, which can then be lifted to iO using [AJ15,BV15,LPST16].

Our Final Assumption: Subspace Flooding. Combined with the transformation discussed earlier, we only need our sampler to satisfy weak security, which boils down to the following *subspace flooding* assumption: that

$$\mathbf{P}, \mathbf{P'}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \widehat{\mathbf{B}} = \mathbf{A}^*\mathbf{S}_0 + \mathbf{F}, \mathbf{C} = \mathbf{A}^*\mathbf{R} + \mathbf{E} - b\mathbf{G}, \mathbf{E}^* + \mathbf{E} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F} \quad (6)$$

hides b where $\mathbf{P} \in \mathbb{Z}^{M \times m^d}$, $\mathbf{P'} \in \mathbb{Z}^{k^d \times K}$, $\mathbf{E} \in \mathbb{Z}^{M \times M \log q}$, and $\mathbf{F} \in \mathbb{Z}^{M \times K}$ and $\{\mathbf{E}_i\}_{i \in [d]}$ are sampled from small distributions;

$$\mathbf{E}^* = \mathbf{P}(\mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \cdots \otimes \mathbf{E}_d)\mathbf{P'} \in \mathbb{Z}^{M \times K};$$

for $i = 1, \ldots, d$, $\mathbf{A}_i$ is sampled from $\mathbb{Z}_q^{m \times w}$ and $\mathbf{S}_i$ is sampled from $\mathbb{Z}_q^{w \times k}$;

$$\mathsf{seed}_{\mathbf{B}^*} = \{\mathbf{B}_i = \mathbf{A}_i\mathbf{S}_i + \mathbf{E}_i\}_{i \in [d]} \in (\mathbb{Z}_q^{m \times w})^d;$$

$\mathbf{S}_0$ is sampled from $\mathbb{Z}_q^{dwm^{d-1} \times K}$ and $\mathbf{R}$ is sampled from $\mathbb{Z}_q^{dwm^{d-1} \times M \log q}$ so $\widehat{\mathbf{B}} \in \mathbb{Z}_q^{M \times K}$ and $\mathbf{C} \in \mathbb{Z}_q^{M \times M \log q}$; and $\mathbf{A}^*$ is the result of the Kilian randomization process described above.

Note that the columns of $\mathbf{E} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) \in \mathbb{Z}^{M \times K}$ live in a low-rank subspace defined by the columns of $\mathbf{E} \in \mathbb{Z}^{M \times M \log q}$ where $K \gg M \log q$ and $\mathbf{F}$ is sampled independently from a small distribution. Thus, the assumption states that $\mathbf{E}^*$ masks whether the error $\mathbf{E}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F} \in \mathbb{Z}^{M \times K}$ lives in this low-rank subspace, hence the name "subspace flooding".

A different, less syntactic, perspective on the subspace flooding assumption tells us that to protect arbitrary computations, it is sufficient to protect a single homomorphic multiplication. Indeed, consider $\mathbf{C}$ to be a GSW encryption of $-b$ and $\widehat{\mathbf{B}}$ to be a GSW encryption of 0. Their homomorphic multiplication gives us

$$\mathbf{C} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) = \mathbf{A}^*(\mathbf{R}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{S}_0) + (\mathbf{E} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F})$$

Subspace flooding says that adding $\mathbf{E}^*$ "protects" the error $\mathbf{E} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F}$ in the evaluated ciphertext in the sense of hiding b.

Theorem 1 (Informal). *Under the (subexponential hardness of the) learning with errors assumption and the subspace flooding assumption (Equation 6 above), there exists an indistinguishability obfuscation scheme.*

1.3 Discussion

Noise Distribution in Prior Works. The sampler in WW sampler works by homomorphically generating pseudorandom LWE samples using an encrypted (weak) pseudorandom function, such as that given by $k, u \mapsto \mathsf{round}(\langle k, u\rangle)$ for key k and random input u. Prior works used the GSW FHE for homomorphic evaluation, but did not specify the circuit implementation for the PRF. Hopkins, Jain and Lin (HJL) [HJL21] presented attacks on these prior LWE samplers that "exploit the flexibility to choose specific implementations of circuits and LWE error distributions in the Gay-Pass and Wee-Wichs assumptions." Specifically, they showed how to introduce redundancy into the circuit used in homomorphic evaluation following the GSW FHE so that the last two bits of $\mathbf{E}^* + \mathbf{Z}_b$ leak b.

Note that the above attack can be circumvented by fixing some natural choice of a concrete weak PRF, such as the aforementioned, which corresponds to FHE decryption; and a circuit evaluation of it, such as [AP14], which is in fact a read-once branching program with k hardwired. Unfortunately, writing down an explicit expression for the error distribution in the pseudorandom LWE sample is far from straightforward, which in turn impedes any cryptanalytic efforts. In this work, we avoid such considerations by directly considering succinct LWE samplers, as opposed to homomorphically evaluated weak PRFs.

Relation to the "LWE with Leakage" Assumption of [JLMS19]. Our assumption basically asserts that for small $\mathbf{Z}_0, \mathbf{Z}_1$ satisfying some precondition:

$$\mathbf{A}_1, \ldots, \mathbf{A}_d, \;\; (\mathbf{B}_i := \mathbf{A}_i\mathbf{S}_i + \mathbf{E}_i)_{i\in[d]}, \;\; \mathbf{P}, \mathbf{P}', \;\; \mathbf{P}(\mathbf{E}_1 \otimes \cdots \otimes \mathbf{E}_d)\mathbf{P}' - \mathbf{Z}_b$$

hides b. (In fact, we do not give away $\mathbf{A}_1, \ldots, \mathbf{A}_d$, rather a random basis for the column span of $\mathbf{A}^*$. We ignore this difference for the rest of the comparison.)

The LWE with leakage assumption of [JLMS19] basically asserts that for small $\mathbf{z}_0, \mathbf{z}_1$, and $\mathbf{A}_i \in \mathbb{Z}_q^{m\times w}$, $\mathbf{s}_i \in \mathbb{Z}_q^{w\times 1}$, $\mathbf{e}_i \in \chi^{m\times 1}$:

$$\mathbf{A}_1, \ldots, \mathbf{A}_{d-2}, \;\; (\mathbf{b}_i := \mathbf{A}_i\mathbf{s}_i + \mathbf{e}_i)_{i\in[d-2]}, \;\; \mathbf{P}, \;\; \mathbf{P}(\mathbf{e}_1 \otimes \cdots \otimes \mathbf{e}_d) + \mathbf{z}_b$$

hides b.

The LWE with leakage assumption of [JLMS19] can be viewed as a variant of our flooding assumption. Syntactically, their definition can be recovered from ours with three modifications:

1. Set $k = 1$ as opposed to our assumption where $k \gg m$;
2. Set $\mathbf{P}$ to be very compressing, namely, the output has length $M \ll m^{d/2}$, whereas in our case $M \approx m^{d-1/2}$; and
3. Do not release $\mathbf{A}_{d-1}, \mathbf{A}_d, \mathbf{B}_{d-1}, \mathbf{B}_d$ to the distinguisher, ensuring that the only leakage about $\mathbf{e}_{d-1}, \mathbf{e}_d$ comes from $\mathbf{E}^*$.

These syntactic differences have the following consequences:

- With $k = 1$ and $M \approx m^{d-1/2}$, the assumption can indeed be broken with sum-of-squares attacks (see, e.g., [BHJ+19].) Thus, our source of security comes from the fact that k is large. Semantically, this means that we take multiple, albeit correlated, instances of the [JLMS19] problem, defined by the k^d columns of our matrix $\mathbf{E}_1 \otimes \cdots \otimes \mathbf{E}_d$, and output a "few", namely, $K \ll k^{d/2}$ linear combinations of them.

– An adversary in our setting can check the rank of

$$\mathbf{P}(\mathbf{B}_1 \otimes \cdots \otimes \mathbf{B}_d)\mathbf{P}' - \mathbf{E}^* + \mathbf{Z}_b \bmod q$$

which is something that cannot be computed in the [JLMS19] assumption since $\mathbf{B}_{d-1}, \mathbf{B}_d$ are not given to the distinguisher. This allows the latter to plausibly handle *worst-case small* $\mathbf{z}_b$, whereas we require an additional pre-condition on $\mathbf{Z}_b$.

Their final iO scheme additionally assume bilinear groups (in addition to LWE), which we do not.

Cryptanalytic Challenges. A central open problem from this work is to design succinct LWE samplers based on weaker assumptions and to carry out cryptanalysis of our candidate succinct LWE sampler. To facilitate the latter, we describe concrete cryptanalytic challenges in Section 4.6. Thanks to our amplification theorem, in order to base iO on our candidate LWE sampler, it suffices for security to hold for a specific pair of distributions $(\mathbf{Z}_0, \mathbf{Z}_1)$. On the other hand, the heuristic underlying our candidate sampler (related to random polynomials being indistinguishable from independent copies of a noise-flooding distribution $\mathcal{D}$) does not refer to properties of the specific distribution. For this reason, our cryptanalytic challenges also refer to more general distributions $\mathbf{Z}_0, \mathbf{Z}_1$ that may not correspond to those which are sufficient for iO.

2 Preliminaries

2.1 Notations

We will denote by λ the security parameter. The notation $\mathrm{negl}(\lambda)$ denotes any function f such that $f(\lambda) = \lambda^{-\omega(1)}$, and $\mathrm{poly}(\lambda)$ denotes any function f such that $f(\lambda) = \mathcal{O}(\lambda^c)$ for some $c > 0$. For a probabilistic algorithm $\mathsf{alg}(\mathsf{inputs})$, we might explicitly refer to its random coins by writting $\mathsf{alg}(\mathsf{inputs};\mathsf{coins})$. We will denote vectors by bold lower case letters (e.g. $\mathbf{a}$) and matrices by bold upper cases letters (e.g. $\mathbf{A}$). We will denote by $\mathbf{a}^\top$ and $\mathbf{A}^\top$ the transposes of $\mathbf{a}$ and $\mathbf{A}$, respectively. We will denote by $\lfloor x \rceil$ the nearest integer to x, rounding towards 0 for half-integers. For matrices $\mathbf{A}, \mathbf{B}$ of appropriate dimensions, we will denote by $(\mathbf{A}\|\mathbf{B})$ their horizontal concatenation and $\binom{\mathbf{A}}{\mathbf{B}}$ their vertical concatenation. For an integer $n \geq 1$, we denote by $\mathbf{I}_n$ the identity matrix of dimension n. For integral vectors and matrices (i.e., those over $\mathbb{Z}$), we use the notation $\|\mathbf{r}\|, \|\mathbf{R}\|$ to denote the maximum absolute value over all the entries.

For matrices $\mathbf{A}, \mathbf{B}$, we denote by $\mathbf{A} \otimes \mathbf{B}$ their tensor (or Kronecker) product. We'll use the following mixed-product property: for matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ of appropriate dimensions, we have $(\mathbf{AB}) \otimes (\mathbf{CD}) = (\mathbf{A} \otimes \mathbf{C}) \cdot (\mathbf{B} \otimes \mathbf{D})$.

For $p \in \mathbb{Q}$, we write $\mathsf{Round}_p(x) = \lfloor x \cdot 1/p \rceil$. If $\mathbf{X}$ is a matrix, $\mathsf{Round}_p(\mathbf{X})$ denotes the rounded value applied component-wise. We denote by $\lceil x \rceil$ the smallest integer larger or equal to x.

For a finite set S, $s \leftarrow S$ denotes sampling uniformly in S. We define the statistical distance between two random variables X and Y over some domain Ω as: $\mathsf{SD}(X,Y) = \frac{1}{2}\sum_{w \in \Omega} |X(w) - Y(w)|$. We say that two ensembles of random variables $X = \{X_\lambda\}$, $Y = \{Y_\lambda\}$ are *statistically indistinguishable*, denoted $X \approx_s Y$, if $\mathsf{SD}(X_\lambda, Y_\lambda) \leq \mathrm{negl}(\lambda)$.

We say that two ensembles of random variables $X = \{X_\lambda\}$, and $Y = \{Y_\lambda\}$ are *computationally indistinguishable*, denoted $X \approx_c Y$, if, for all (non-uniform) PPT distinguishers $\mathcal{A}$, we have $|\Pr[\mathcal{A}(X_\lambda) = 1] - \Pr[\mathcal{A}(Y_\lambda) = 1]| \leq \mathrm{negl}(\lambda)$. We also refer to sub-exponential security, meaning that there exists some $\varepsilon > 0$ such that the distinguishing advantage is at most $2^{-\lambda^\varepsilon}$.

2.2 Learning With Errors

Definition 1 (B-bounded distribution). *We say that a distribution χ over $\mathbb{Z}$ is B-bounded if*

$$\Pr[\, \chi \in [-B, B]\,] = 1.$$

We recall the definition of the (decision) *Learning with Errors* problem, introduced by Regev [Reg05].

Definition 2 ((Decision) Learning with Errors ([Reg05])). *Let $n = n(\lambda)$ and $q = q(\lambda)$ be integer parameters and $\chi = \chi(\lambda)$ be a distribution over $\mathbb{Z}$. The Learning with Errors (LWE) assumption $LWE_{n,q,\chi}$ states that for all polynomials $m = \mathrm{poly}(\lambda)$ the following distributions are computationally indistinguishable:*

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{u})$$

where $\mathbf{A} \leftarrow \mathbb{Z}_q^{m \times n}, \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \chi^m, \mathbf{u} \leftarrow \mathbb{Z}_q^m$.

Just like many prior works, we rely on LWE security with the following range of parameters. We assume that for any polynomial $p = p(\lambda) = \mathrm{poly}(\lambda)$ there exists some polynomial $n = n(\lambda) = \mathrm{poly}(\lambda)$, some $q = q(\lambda) = 2^{\mathrm{poly}(\lambda)}$ and some $B = B(\lambda)$-bounded distribution $\chi = \chi(\lambda)$ such that $q/B \geq 2^p$ and the $LWE_{n,q,\chi}$ assumption holds. Throughout the paper, the *LWE assumption* without further specification refers to the above parameters. The *sub-exponentially secure LWE* assumption further assumes that $LWE_{n,q,\chi}$ with the above parameters is sub-exponentially secure, meaning that there exists some $\varepsilon > 0$ such that the distinguishing advantage of any polynomial-time distinguisher is $2^{-\lambda^\varepsilon}$.

The works of [Reg05,Pei09] showed that the (sub-exponentially secure) LWE assumption with the above parameters follows from the worst-case (sub-exponential) quantum hardness SIVP and classical hardness of GapSVP with sub-exponential approximation factors.

2.3 Lattice Tools

Noise Flooding. We will use the following fact.

Lemma 1 (Flooding Lemma (e.g., [AJL+12])). *Let $B = B(\lambda), B' = B'(\lambda) \in \mathbb{Z}$ be parameters and let $U([-B, B])$ be the uniform distribution over the integer interval $[-B, B]$. Then for any $e \in [-B', B']$, the statistical distance between $U([-B, B])$ and $U([-B, B]) + e$ is B'/B.*

Gadget Matrix [MP12]. For an integer $q \geq 2$, define: $\mathbf{g} = (1, 2, \cdots, 2^{\lceil \log q \rceil - 1}) \in \mathbb{Z}_q^{1 \times \lceil \log q \rceil}$. The *gadget matrix* $\mathbf{G}$ is defined as $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}_n \in \mathbb{Z}_q^{n \times m}$ where $n \in \mathbb{N}$ and $m = n \lceil \log q \rceil$. There exists an efficiently computable deterministic function $\mathbf{G}^{-1} : \mathbb{Z}_q^n \to \{0,1\}^m$ such that for all $\mathbf{u} \in \mathbb{Z}_q^n$ we have $\mathbf{G} \cdot \mathbf{G}^{-1}(\mathbf{u}) = \mathbf{u}$. We let $\mathbf{G}^{-1}(\$)$ denote the distribution obtained by sampling $\mathbf{u} \leftarrow \mathbb{Z}_q^n$ uniformly at random and outputting $\mathbf{t} = \mathbf{G}^{-1}(\mathbf{u})$. These extend directly to matrices: $\mathbf{G}^{-1} : \mathbb{Z}_q^{n \times k} \to \{0,1\}^{m \times k}$ by concatenating the outputs.

2.4 Homomorphic Operations

In this section, we describe how to perform homomorphic operations over certain encodings of inputs. For readers familiar with lattice-based primitives, these essentially are packed versions of the GSW homomorphism.

Our operations follow readily from [WW21] (building on [GSW13, GVW15], along with the "packing" techniques in [PVW08, MW16, BTVW17, PS19, GH19, BDGM19]), who build homomorphic operations for $f : \{0,1\}^\ell \to \{0,1\}^M$, producing some vector $\mathbf{c}_f \in \mathbb{Z}_q^M$. We extend these operations to functions $f : \{0,1\}^\ell \to \{0,1\}^{M \times K}$ to produce some matrix $\mathbf{C}_f \in \mathbb{Z}_q^{M \times K}$, obtained by concatenating K vectors $\mathbf{c}_{f_i}$. This yields the following.

Definition 3 (Homomorphic operations). *Let M, W, q, ℓ, K, t be parameters. We define the following efficient algorithms:*

- $\mathsf{Eval}(f : \{0,1\}^\ell \to \{0,1\}^{M \times K}, \mathbf{C} \in \mathbb{Z}_q^{M \times \ell M \log q})$*: deterministically outputs a matrix $\mathbf{C}_f \in \mathbb{Z}_q^{M \times Q}$.*
- $\mathsf{Eval}_{\mathsf{open}}(f, \mathbf{A} \in \mathbb{Z}_q^{M \times W}, x \in \{0,1\}^\ell, \mathbf{R} \in \mathbb{Z}_q^{W \times \ell M \log q}, \mathbf{E} \in \mathbb{Z}^{M \times \ell M \log q})$*: deterministically outputs two matrices $(\mathbf{R}_{f,x} \in \mathbb{Z}_q^{W \times Q}, \mathbf{E}_{f,x} \in \mathbb{Z}^{M \times Q})$.*

These operations have the following property. For all $f : \{0,1\}^\ell \to \{0,1\}^{M \times K}$ of depth t, $\mathbf{x} \in \{0,1\}^\ell$, $\mathbf{A} \in \mathbb{Z}_q^{M \times W}$, $\mathbf{R} \in \mathbb{Z}_q^{W \times \ell M \log q}$ and $\mathbf{E} \in \mathbb{Z}^{M \times \ell M \log q}$, if

$$\mathbf{C} = \mathbf{A}\mathbf{R} + x^\top \otimes \mathbf{G} + \mathbf{E} \in \mathbb{Z}_q^{M \times \ell M \log q},$$

$$\mathbf{C}_f = \mathsf{Eval}(f, \mathbf{C}),$$

$$(\mathbf{R}_{f,x}, \mathbf{E}_{f,x}) = \mathsf{Eval}_{\mathsf{open}}(f, \mathbf{A}, x, \mathbf{R}, \mathbf{E}),$$

where we view x as a row vector $x \in \{0,1\}^{1 \times \ell}$, then

$$\mathbf{C}_f = \mathbf{A}\mathbf{R}_{f,x} + q/2 \cdot f(x) + \mathbf{E}_{f,x} \in \mathbb{Z}_q^{M \times K},$$

where $f(x) \in \{0,1\}^{M \times K}$. Furthermore $\|\mathbf{E}_{f,x}\| = \|\mathbf{E}\| \cdot M^{g(t)}$ for some efficiently computable g such that $g(t) = \mathcal{O}(t)$.

Similarly to [WW21], these algorithms extend to functions f with outputs in $\mathbb{Z}_q$.

- $\mathsf{Eval}^q(f : \{0,1\}^\ell \to \mathbb{Z}_q^{M\times K}, \mathbf{C} \in \mathbb{Z}_q^{M\times \ell M \log q})$: deterministically outputs a matrix $\mathbf{C}_f \in \mathbb{Z}_q^{M\times Q}$.
- $\mathsf{Eval}^q_{\mathsf{open}}(f, \mathbf{A} \in \mathbb{Z}_q^{M\times W}, x \in \{0,1\}^\ell, \mathbf{R} \in \mathbb{Z}_q^{W\times \ell M \log q}, \mathbf{E} \in \mathbb{Z}^{M\times \ell M \log q})$: deterministically outputs two matrices $(\mathbf{R}_f \in \mathbb{Z}_q^{W\times Q}, \mathbf{E}_f \in \mathbb{Z}^{M\times Q})$.

The correctness requirement becomes:

$$\mathbf{C}_f = \mathbf{A}\mathbf{R}_{f,x} + f(x) + \mathbf{E}_{f,x} \in \mathbb{Z}_q^{M\times K},$$

where $\mathbf{C} = \mathbf{A}\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E} \in \mathbb{Z}_q^{M\times \ell M \log q}$, x being again seen as a row vector, $\mathbf{C}_f = \mathsf{Eval}^q(f, \mathbf{C})$ and $(\mathbf{R}_{f,x}, \mathbf{E}_{f,x}) = \mathsf{Eval}^q_{\mathsf{open}}(f, \mathbf{A}, x, \mathbf{R}, \mathbf{E})$, and $f(x) \in \mathbb{Z}_q^{M\times K}$. Again, $\|\mathbf{E}_{f,x}\| = \|\mathbf{E}\| \cdot M^{g(t)}$.

2.5 Succinct Randomized Encodings

Next, we define succinct randomized encodings [BGL+15, BCG+18, LPST16].

Definition 4. *A* succinct randomized encoding scheme *(SRE) for the function family $\mathcal{F}_{\ell,N,t} = \{f : \{0,1\}^\ell \to \{0,1\}^N\}$ of circuits of depth at most t, is a tuple of PPT algorithms* $(\mathsf{CRSGen}, \mathsf{Encode}, \mathsf{Decode})$ *with the following syntax:*

- $\mathsf{CRSGen}(1^\lambda, \mathcal{F}_{\ell,N,t}) \to \mathsf{crs}$*: on input the security parameter and a function family, outputs* crs.
- $\mathsf{Encode}(\mathsf{crs}, f, x) \to C$*: on input* crs*, a function $f \in \mathcal{F}_{\ell,N,t}$ and $x \in \{0,1\}^\ell$, outputs an encoding C.*
- $\mathsf{Decode}(\mathsf{crs}, C, f) \to y$*: a deterministic algorithm which, on input* crs*, an encoding C, and a function $f \in \mathcal{F}_{\ell,N,t}$, outputs a value $y \in \{0,1\}^N$.*

We require the following properties:

Correctness: For $f \in \mathcal{F}_{\ell,N,t}$ and any $x \in \{0,1\}^\ell$:

$$\Pr\left[\mathsf{Decode}(\mathsf{crs}, \mathsf{Encode}(\mathsf{crs}, f, x), f) = f(x)\right] \geq 1 - \mathrm{negl}(\lambda),$$

where $\mathsf{crs} \leftarrow \mathsf{CRSGen}(1^\lambda, \mathcal{F}_{\ell,N,t})$ *(over the randomness of* $\mathsf{CRSGen}, \mathsf{Encode}$*).*

δ-Succinctness: There exists a constant $\delta < 1$ such that, for all $\mathsf{crs} \leftarrow \mathsf{CRSGen}(1^\lambda, \mathcal{F}_{\ell,N,t})$*, $C \leftarrow \mathsf{Encode}(\mathsf{crs}, f, x)$, we have:*

$$|C| = N^\delta \cdot \mathrm{poly}(\lambda, \ell, t).$$

Indistinguishability-based Security: For all PPT $\mathcal{A}$, all $x_0, x_1 \in \ell$, and all $f \in \mathcal{F}_{t,\ell,N}$ such that $f(x_0) = f(x_1)$, the following distributions are indistinguishable for $b = 0$ and $b = 1$:

- $\mathcal{D}_b$*: Sample* $\mathsf{crs} \leftarrow \mathsf{CRSGen}(1^\lambda, \mathcal{F}_{t,\ell,N})$*, $C_b \leftarrow \mathsf{Encode}(\mathsf{crs}, f, x_b)$. Output* (crs, C_b)*.*

Relation to XiO. Our notion of SRE is also very related to the notion of exponentially efficient iO (XiO) from [LPST16]. An XiO scheme obfuscates a circuit $C : \{0,1\}^{\log N} \to \{0,1\}$ with the same security guarantee as iO, but the runtime of the obfuscator can be as high as $\mathrm{poly}(\lambda, |C|, N)$ and the only constraint that makes the problem non-trivial is that the obfuscated circuit is succinct, of size at most $N^\delta \mathrm{poly}(\lambda, |C|)$ for $\delta < 1$. An SRE scheme immediately yields an XiO scheme by thinking of f as the universal circuit that takes as input a circuit $x = C$ an evaluates it on all N inputs in $\{0,1\}^{\log N}$. The output size of f is N and the depth of f can be bounded by $t = \mathrm{poly}(|C|)$, so the succinctness of the SRE yields the corresponding succinctness of the XiO. Therefore, by leveraging the prior work of [LPST16] that shows how to go from XiO (in the CRS model) to iO via LWE, we get the following theorem.

Theorem 2. [AJ15, BV15, LPST16] *Assuming sub-exponentially secure SRE exist and sub-exponentially secure LWE, there exists an iO scheme.*

3 Succinct LWE Sampler: Definition and Amplification

In Section 3.1, we define the notion of succinct LWE samplers. In Section 3.2, we describe a seemingly weaker notion of LWE sampler, and prove that it implies the first (and stronger) notion.

3.1 Definition and Discussion

Definition 5 (Succinct LWE Sampler). *A* succinct LWE sampler *is a tuple of PPT algorithms* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *with the following syntax:*

- $\mathsf{SampCRSGen}(1^\lambda, 1^N, \alpha)$*: on input the security parameter* λ*, a size parameter* N *and a blow-up factor* α*, samples a common reference string* crs*, which include parameters* $\mathsf{params} = (q, M, K, \overline{\chi}, \overline{B})$*.*
- $\mathsf{LWEGen}(\mathsf{crs})$*: samples* $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*)$*.*
- $\mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$ *is a deterministic algorithm that outputs a matrix* $\mathbf{B}^*$*.*

Domains and Parameters. The outputs of LWEGen *and* Expand *satisfy:*

$$\mathbf{A}^* \in \mathbb{Z}_q^{M\times W}, \quad \mathbf{S}^* \in \mathbb{Z}_q^{W\times K}, \quad \mathbf{B}^* \in \mathbb{Z}_q^{M\times K},$$

for some integer W*. We require that:*

- $N = MK$*;*
- $\overline{B} = \mathrm{poly}(N)$*;*
- $\overline{\chi}$ *is a* $\overline{B}$*-bounded noise distribution; and*
- $q \geq 8 \cdot 2^\lambda \cdot \alpha \cdot \overline{B}$*.*

Correctness. We require that

$$||\mathbf{B}^* - \mathbf{A}^*\mathbf{S}^*|| := \beta \leq q/8$$

where $\mathsf{crs} \leftarrow \mathsf{SampCRSGen}(1^\lambda, 1^N, \alpha)$*,* $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) \leftarrow \mathsf{LWEGen}(\mathsf{crs})$ *and* $\mathbf{B}^* := \mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$*. Furthermore, we require that* $\mathbf{A}^*$ *is full-rank with overwhelming probability over the randomness of* $\mathsf{SampCRSGen}$ *and* LWEGen*.*

δ-Succinctness. We require the total bit length of the output of LWEGen *is small. That is,*

$$\mathsf{bitlength}(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) \leq N^\delta \cdot \mathrm{poly}(\lambda, \log q) = (MK)^\delta \cdot \mathrm{poly}(\lambda, \log q) \ ,$$

where $\delta < 1$ is a constant. When we omit δ, it means succinctness holds for some constant $\delta < 1$.

LWE with respect to $\mathbf{A}^*$. We require that

$$(\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{A}^*\mathbf{s}' + \mathbf{e}') \approx_c (\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{b}),$$

where $\mathsf{crs} = \mathsf{SampCRSGen}(1^\lambda, 1^N, \alpha; \mathsf{coins}_{\mathsf{crs}})$*,* $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) \leftarrow \mathsf{LWEGen}(\mathsf{crs}; \mathsf{coins}_{\mathsf{seed}})$*, $\mathbf{s}' \leftarrow \mathbf{Z}_q^W$, and $\mathbf{e}' \leftarrow \overline{\chi}^M$.*

Security (or β_0-Flooding). Let D_0, D_1 be any two polynomial-time samplable distributions such that $(\mathsf{aux}_b, \mathbf{Z}_b) \leftarrow D_b(\mathbf{A}^*)$ satisfies $\mathbf{Z}_b \in \mathbb{Z}^{M \times K}$, $\|\mathbf{Z}_b\| \leq \beta_0$ where $\beta_0 \cdot 2^\lambda \leq \beta$ and

$$(\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_1, \mathsf{aux}_1)$$

where $\mathsf{crs} = \mathsf{SampCRSGen}(1^\lambda, 1^N, \alpha; \mathsf{coins}_{\mathsf{crs}})$*,* $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) = \mathsf{LWEGen}(\mathsf{crs}; \mathsf{coins}_{\mathsf{seed}})$ *and $\mathbf{S}' \leftarrow \mathbb{Z}_q^{W \times K}$. Then,*

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_1, \mathsf{aux}_1).$$

We will refer to the assumption on D_0, D_1 as the pre-condition *for security, and the resulting indistinguishability the* post-condition*.*
Furthermore, as we will later describe a relaxed notion of security, we will sometimes refer to the notion above as ***strong security*** *to avoid ambiguity.*

Remark 1 (Alternate formulation). Since the sampler allows us to compute $\mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}) = \mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$, the security post-condition can be equivalently stated as:

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{E}^* - \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{E}^* - \mathbf{Z}_1, \mathsf{aux}_1).$$

Remark 2 (Implied Statements). The randomness $\mathsf{coins}_{\mathsf{crs}}$ and $\mathsf{coins}_{\mathsf{seed}}$ respectively used by $\mathsf{SampCRSGen}$ and LWEGen allow us to compute $\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*$. In particular, LWE with respect to $\mathbf{A}^*$ implies that

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*, \mathbf{A}^*\mathbf{s}' + \mathbf{e}) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*, \mathbf{b}),$$

and the pre-condition on D_0, D_1 for security implies that

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*, \mathsf{aux}_0, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_0) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*, \mathsf{aux}_1, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_1).$$

Remark 3 (Restrictions on $\mathbf{Z}_0, \mathbf{Z}_1$). We note that security (namely, the post-conditionition) cannot hold for arbitrary $\mathbf{Z}_0, \mathbf{Z}_1$, for which the pre-condition does not hold. Even if one only required that $\mathbf{Z}_0$ and $\mathbf{Z}_1$ had small entries, one can efficiently distinguish $\mathbf{Z}_0 = \mathbf{0}$ from any $\mathbf{Z}_1$ not in the column span of $\mathbf{A}^*$. In particular, the rank of $\mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_b$ would leak b: this is because $\mathbf{A}^*\mathbf{S}^*$ is rank-deficient by succinctness. We can rule out such distinguishers simply by requiring that $\mathbf{Z}_0 - \mathbf{Z}_1$ lies in the column span of $\mathbf{A}^*$; our pre-condition is in some sense a "distributional" or "computational" relaxation of such a requirement.

Remark 4 (Triviality without succinctness). We remark that it is easy to build a succinct LWE sampler if there are no restrictions on the bit-length of $\mathsf{seed}_{\mathbf{B}^*}$ (looking ahead, such a sampler would not be sufficient to build iO). Indeed, without any succinctness requirement, we could set:

$$\mathsf{crs} = \emptyset, \quad \mathsf{seed}_{\mathbf{B}^*} = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^* \in \mathbb{Z}_q^{M \times K}$$

where $\mathbf{S}^*$ is random and $\mathbf{E}^*$ has small entries, but large enough to "noise-flood" $\mathbf{Z}_b$ (namely, $\beta_0/\beta = 2^{-\lambda}$).

For convenience, we consider the equivalent notion of security from Remark 1. We claim that this construction (unconditionally) satisfies security. To see this, first note that for all $b \in \{0,1\}$:

$$(\mathsf{seed}_{\mathbf{B}^*},\ \mathbf{A}^*,\ \mathbf{E}^* - \mathbf{Z}_b,\ \mathsf{aux}_b) \approx_s (\mathbf{A}^*\mathbf{S}^* + (\mathbf{E}^* + \mathbf{Z}_b),\ \mathbf{A}^*,\ \mathbf{E}^*,\ \mathsf{aux}_b)$$

by noise flooding, where we use that $\mathbf{E}^*$ is sampled independently of $\mathsf{aux}_b, \mathbf{Z}_b$. The pre-condition then implies that

$$(\mathbf{A}^*, (\mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_0) + \mathbf{E}^*, \mathbf{E}^*, \mathsf{aux}_0) \approx_c (\mathbf{A}^*, (\mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_1) + \mathbf{E}^*, \mathbf{E}^*, \mathsf{aux}_1),$$

where we again use that $\mathbf{E}^*$ is sampled independently of $\mathsf{aux}_b, \mathbf{Z}_b, \mathbf{S}^*$, and that $\mathbf{S}^*$ is sampled uniformly at random independently of the other components (and takes the role of $\mathbf{S}'$ in the pre-condition).

Remark 5 (Heuristic necessity of a CRS). We heuristically show that security requires a (long) CRS if $\mathsf{seed}_{\mathbf{B}^*}$ is required to be short, namely the CRS needs to be of length $\approx N$ for any δ-succinct scheme with $\delta < 1$.

Suppose for contradiction that there is such a sampler that expands some short input $(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$ of length at most $N^\delta \cdot \mathrm{poly}(\lambda, \log q)$ to some $\mathsf{Expand}(\mathsf{seed}_{\mathbf{B}^*}) = \mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$ of bit-length $N \log q$. Let $\mathbf{Z}_b$ be a random LWE error and let aux_b be an obfuscation of the following program:

$P_{b,\mathbf{A}^*,\mathbf{Z}_b}$: on input $(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$ of bit-length $N^\delta \cdot \mathrm{poly}(\lambda, \log q)$, and $\widetilde{\mathbf{B}}$ of bit-length $N \log q$,
- Check that $\widetilde{\mathbf{B}} - \mathbf{Z}_b$ is in the column span of $\mathbf{A}^*$, and output $\perp$ if not.
- Compute $\mathbf{B}^* = \mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}) = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$. Output b if $\|\mathbf{B}^* - \widetilde{\mathbf{B}} + \mathbf{Z}_b\| \leq \beta$, and output $\perp$ otherwise.

Then $(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \widetilde{\mathbf{B}} = \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_0, \mathsf{aux}_0)$ is efficiently distinguishable from $(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \widetilde{\mathbf{B}} = \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_1, \mathsf{aux}_1)$, by running aux_b on input $((\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}), \widetilde{\mathbf{B}})$ and using the fact that $(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$ has bit-length at most $\widetilde{O}(N^\delta)$ by assumption, that $\|\mathbf{E}^*\| \leq \beta$, that $\mathbf{A}^*\mathbf{S}^*$ has low rank by succinctness, and that $\mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_0 - \mathbf{Z}_1$ has high rank w.h.p.

Furthermore, suppose heuristically that aux_b acts like an ideal obfuscation of $P_{b,\mathbf{Z}_b}$, meaning that it does not reveal more than black-box access to the program. Then, the pre-condition would hold since given $(\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{B}_b = \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_b)$ and black-box access to $P_{b,\mathbf{Z}_b}$, one cannot distinguish $b = 0$ vs $b = 1$. The idea is that the only way to learn anything about b is to provide a "good" input to $P_{b,\mathbf{Z}_b}$ that makes it output something other than $\perp$. Any good input must be of the form $((\mathsf{crs}', \mathsf{seed}'_{\mathbf{B}^*}), \mathbf{B}_b + \mathbf{A}^*\mathbf{S})$ for some $\mathbf{S} \in \mathbb{Z}_q^{W \times K}$. But if $\mathbf{B}_b$ was uniform, there would be no inputs of this form, where $(\mathsf{crs}', \mathsf{seed}'_{\mathbf{B}^*})$ is short, such that $\|\mathsf{Expand}(\mathsf{crs}', \mathsf{seed}'_{\mathbf{B}^*}) - \mathbf{B}_b + \mathbf{A}^*\mathbf{S}\|$ is also small, meaning that $P_{b,\mathbf{Z}_b}$ would always output $\perp$ in this case. This follows by a counting argument, where the sizes of crs', $\mathsf{seed}'_{\mathbf{B}^*}$ and $\mathbf{S}$ are much smaller than the size of $\mathbf{B}_b$ whenever δ is sufficiently small, and β is relatively small compared to q. Therefore finding a good input to $P_{b,\mathbf{Z}_b}$ would require breaking LWE with respect to $\mathbf{A}^*$.

3.2 Weak Succinct LWE Samplers

We now present a weaker security notion for succinct LWE samplers. Instead of quantifying over all $(\mathbf{Z}_b, \mathsf{aux}_b)$ that satisfy the specified pre-condition as we did previously, we now fix one particular and simple choice of $(\mathbf{Z}_b, \mathsf{aux}_b)$. In particular, this makes the definition falsifiable. We then show in Theorem 3 that there is a generic compiler that upgrades this type of weak security to the previous definition of strong security (Definition 5).

Definition 6. *Weak Security (or Weak β_0-Flooding). Define D_0, D_1 as follows.*

$$D_b : \quad \mathsf{aux}_b = \left(\widehat{\mathbf{B}} := \mathbf{A}^*\widehat{\mathbf{S}} + \widehat{\mathbf{E}}, \quad \mathbf{C} = \mathbf{A}^*\mathbf{R} + \mathbf{E} - b \cdot \mathbf{G}\right)$$
$$\mathbf{Z}_b = \mathbf{E}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\widehat{\mathbf{E}},$$

where

- $\mathsf{SampCRSGen}$ *defines* $(q, M, K, \overline{\chi}, \overline{B}) = \mathsf{params}$*;*
- LWEGen *defines* $\mathbf{A}^* \in \mathbb{Z}_q^{M \times W}$*;*
- $\widehat{\mathbf{B}} \in \mathbb{Z}_q^{M \times K}$*,* $\widehat{\mathbf{S}} \leftarrow \mathbb{Z}_q^{W \times K}$*, and* $\widehat{\mathbf{E}} \leftarrow [-B_{\mathsf{flood}}, B_{\mathsf{flood}}]^{M \times K}$*, where* $B_{\mathsf{flood}} = (\beta_0 + \overline{B}) \cdot 2^\lambda$*;*
- $\mathbf{C} \in \mathbb{Z}_q^{M \times M \log q}$*,* $\mathbf{R} \leftarrow \mathbb{Z}_q^{W \times M \log q}$*, and* $\mathbf{E} \leftarrow \overline{\chi}^{M \times M \log q}$*.*

We say that the sampler $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *is weakly secure if*

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_1, \mathsf{aux}_1).$$

Remark 6 (Alternate formulation of security). Similar to Remark 1, as the sampler allows us to compute $\mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}) = \mathbf{B}^* = \mathbf{A}^*\mathbf{S}^* + \mathbf{E}^*$, weak security equivalently states that:

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{E}^* - \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{E}^* - \mathbf{Z}_1, \mathsf{aux}_1).$$

Remark 7 (Pre-condition from LWE). We note that the distributions D_0, D_1 satisfy the pre-condition for security of Definition 5, assuming LWE, namely:

$$(\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_0, \mathsf{aux}_0) \approx_c (\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathbf{A}^*\mathbf{S}' + \mathbf{Z}_1, \mathsf{aux}_0), \quad (7)$$

where $(\mathsf{aux}_b, \mathbf{Z}_b) \leftarrow D_b$ and $\mathbf{S}' \leftarrow \mathbb{Z}_q^{W \times K}$.

This is true because one can efficiently sample $\mathbf{A}^*\mathbf{S}' + \mathbf{Z}_b$ given only $(\mathbf{A}^*, \mathsf{aux}_b)$, as follows:

- Compute $\mathbf{C}_{\widehat{\mathbf{B}}} = \mathbf{C}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) \in \mathbb{Z}_q^{M \times K}$; and
- Output $\mathbf{C}_{\widehat{\mathbf{B}}} + \mathbf{A}^*\mathbf{S}$ for some random $\mathbf{S} \leftarrow \mathbb{Z}_q^{W \times K}$.

Indeed,

$$\begin{aligned}\mathbf{C}_{\widehat{\mathbf{B}}} + \mathbf{A}^*\mathbf{S} &= (\mathbf{A}^*\mathbf{R} + \mathbf{E} - b\mathbf{G})\mathbf{G}^{-1}(\widehat{\mathbf{B}}) + \mathbf{A}^*\mathbf{S} \\ &= \mathbf{A}^*(\mathbf{R}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\widehat{\mathbf{S}} + \mathbf{S}) + (\mathbf{E}\mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\widehat{\mathbf{E}})\end{aligned}$$

and the latter term is distributed identically to $\mathbf{A}^*\mathbf{S}' + \mathbf{Z}_b$ with a random $\mathbf{S}'$.

Therefore, to show the precondition equation (7), it suffices to prove that $(\mathsf{coins}_{\mathsf{crs}}, \mathsf{coins}_{\mathsf{seed}}, \mathsf{aux}_b)$ hides b. But this follows from LWE with respect to $\mathbf{A}^*$ (Definition 5) with noise distribution $\overline{\chi}$.

3.3 Amplification

The following theorem allows to lift weak security (Definition 6) to strong security (Definition 5).

Theorem 3. *Suppose there exists a weakly secure, δ-succinct LWE sampler (Definition 6). Suppose furthermore that it satisfies $M^2 \leq N^\delta \cdot \mathrm{poly}(\lambda, \log q)$. Then, assuming LWE, there exists a secure δ-succinct LWE sampler, satisfying strong security (Definition 5). Moreover, with the parameters of Definition 6, there exists such a sampler that is (strongly) β_0-flooding.*

We refer to the full version for a construction and a proof.

4 Candidate Succinct LWE Sampler

In Section 4.1, we present the template of our main candidate. In Section 4.2, we state correctness and succinctness (and refer to the full version for proofs). In Section 4.3, we explain how to setup parameters, and state our conjectured security. Last, we discuss the plausibility of our conjecture in Section 4.5.

4.1 A Basic Framework

We describe a basic template to build succinct LWE samplers. Looking ahead, the SRE construction in Section 5 requires an additional succinctness requirement, namely, that additional encodings produced by the SRE are succinct. We make sure that our template and the parameters we propose are compatible with that constraint.

We now describe our framework. It uses a set of parameters:

$$\mathsf{parameters} := (d, m, k, w, M, K, \overline{\chi}, \chi, \beta, q)$$

which in particular includes $\mathsf{params} = (q, M, K, \overline{\chi}, \overline{B}, \chi)$ directly output by $\mathsf{SampCRSGen}$. Informally,

- the security of our sampler is related to the hardness of solving systems of random degree d polynomials;
- q is the underlying LWE modulus;
- m, k, w define the dimensions of the "seed" LWE samples $\mathbf{A}_i, \mathbf{S}_i, \mathbf{E}_i$, which together with d, determine M, K, which are the dimensions for "expanded" sample $\mathbf{B}^*$;
- χ is the noise distribution for $\mathbf{E}_i$; it is B-bounded over $\mathbb{Z}$;
- $\overline{\chi}$ is the noise distribution used for LWE w.r.t $\mathbf{A}^*$; it is $\overline{B}$-bounded over $\mathbb{Z}$;
- D_P a σ-bounded distribution over $\mathbb{Z}$. We will take $D_P = \chi$ for simplicity.

We now describe our candidate $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$.

- $\mathsf{SampCRSGen}(1^\lambda, 1^N, \alpha)$: Derive $\mathsf{parameters} = (d, m, k, w, M, K, \overline{\chi}, \overline{B}, \chi, \beta, q)$ from $(1^\lambda, 1^N, \alpha)$ as described later in Section 4.3. Set $\mathsf{params} = (q, M, K, \overline{\chi}, \overline{B}, \chi)$.
  Sample $\mathbf{P'} \leftarrow \chi^{k^d \times K}$ and $\mathbf{P} \leftarrow \chi^{M \times m^d}$. Output
  $$\mathsf{crs} = (\mathsf{params}, \mathbf{P}, \mathbf{P'}).$$
- $\mathsf{LWEGen}(\mathsf{crs})$: On input $\mathsf{crs} = (\mathsf{params}, \mathbf{P}, \mathbf{P'})$, sample, for $i \in [d]$, $\mathbf{A}_i \leftarrow \mathbb{Z}_q^{m \times w}$, $\mathbf{S}_i \leftarrow \mathbb{Z}_q^{w \times k}$, $\mathbf{E}_i \leftarrow \chi^{m \times k}$ where χ is specified in params. Compute:
  $$\mathbf{B}_i = \mathbf{A}_i \mathbf{S}_i + \mathbf{E}_i \in \mathbb{Z}_q^{m \times k}.$$
  Set:
  $$\overline{\mathbf{A}}^* = \mathbf{P} \cdot \Big(\mathbf{A}_1 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \mathbf{I}_m \otimes \mathbf{A}_2 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \cdots \\ \cdots \,\|\, \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \otimes \mathbf{A}_d \Big) \in \mathbb{Z}_q^{M \times dwm^{d-1}}$$
  $$\overline{\mathbf{S}}^* = \begin{pmatrix} \mathbf{S}_1 \otimes \mathbf{B}_2 \otimes \cdots \otimes \mathbf{B}_d \\ \mathbf{E}_1 \otimes \mathbf{S}_2 \otimes \cdots \otimes \mathbf{B}_d \\ \vdots \\ \mathbf{E}_1 \otimes \mathbf{E}_2 \otimes \cdots \otimes \mathbf{S}_d \end{pmatrix} \cdot \mathbf{P'} \in \mathbb{Z}_q^{dwm^{d-1} \times K}.$$

Sample a random basis $\mathbf{A}^* \in \mathbb{Z}_q^{M \times W}$ of the column space of $\overline{\mathbf{A}}^*$, and solve for $\mathbf{S}^* \in \mathbb{Z}_q^{W \times K}$ such that $\mathbf{A}^*\mathbf{S}^* = \overline{\mathbf{A}}^* \cdot \overline{\mathbf{S}}^*$. Output:

$$\mathsf{seed}_{\mathbf{B}^*} = \{\mathbf{B}_i\}_{i \in [d]}, \quad \mathbf{A}^*, \mathbf{S}^*.$$

– $\mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$: On input $\mathsf{crs} = (\mathsf{params}, \mathbf{P}, \mathbf{P}')$ and $\mathsf{seed}_{\mathbf{B}^*} = \{\mathbf{B}_i\}_{i \in [d]}$, output:

$$\mathbf{B}^* = \mathbf{P} \cdot (\mathbf{B}_1 \otimes \cdots \otimes \mathbf{B}_d) \cdot \mathbf{P}' \in \mathbb{Z}_q^{M \times K}.$$

4.2 Correctness, Succinctness, and LWE with respect to $\mathbf{A}^*$

We show that for appropriate parameters, the sampler described above is correct and succinct.

Claim 1. *Assume $\beta \geq B^2(mkB)^d$. Then the sampler* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *described above satisfies correctness (Definition 5).*

Claim 2. *Suppose there exists $\delta < 1$ such that*

$$(dmk + MW + WK) \leq N^\delta \cdot \mathrm{poly}(\lambda, \log q),$$

where W is the width of $\mathbf{A}^*$. Then $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *described above is δ-succinct.*

Proof. This follows as $\mathsf{bitlength}(\{\mathbf{B}_i\}_{i\in[d]}, \mathbf{A}^*, \mathbf{S}^*) = (dmk + MW + WK) \cdot \log q$.

Next, we show that LWE holds with respect to $\mathbf{A}^*$ (assuming standard LWE), for our candidate sampler. We first show that it holds with respect to $\overline{\mathbf{A}}^*$.

Lemma 2 (LWE with respect to $\overline{\mathbf{A}}^*$). *Let $\chi(\lambda)$ be a $B(\lambda)$-bounded distribution. Let D_P be a σ-bounded distribution over $\mathbb{Z}$ such that if $\mathbf{P} = D_P^{M \times m^d}(\mathsf{coins}_P)$ is sampled using randomness coins_P, then with overwhelming probability over coins_P, $\mathbf{P}$ is full-rank. Suppose furthermore that $M \leq m^d$.*

Suppose $LWE_{w,q,\chi}$ holds. Let $\overline{\chi} = \mathcal{U}([-\overline{B}, \overline{B}])$ be the uniform distribution in $[-\overline{B}, \overline{B}]$, where $\overline{B} \geq \sigma m^d B \cdot 2^\lambda$. Then:

$$\left(\mathsf{coins}_P, \mathbf{P}, \{\mathbf{A}_i\}_{i\in[d]}, \overline{\mathbf{A}}^*, \overline{\mathbf{A}}^* \cdot \mathbf{s} + \mathbf{e}\right) \approx_c \left(\mathsf{coins}_P, \mathbf{P}, \{\mathbf{A}_i\}_{i\in[d]}, \overline{\mathbf{A}}^*, \mathbf{b}\right),$$

where $\mathbf{P} = D_P^{M \times m^d}(\mathsf{coins}_P)$, $\mathbf{b} \leftarrow \mathbb{Z}_q^M$, $\mathbf{s} \leftarrow \mathbb{Z}_q^{dwm^{d-1}}$, $\mathbf{e} \leftarrow \overline{\chi}^M$.

Corollary 1 (LWE with respect to $\mathbf{A}^*$). *Let $\chi(\lambda)$ be a $B(\lambda)$-bounded distribution. Suppose furthermore that $M \leq m^d$. Then, assuming $LWE_{w,\chi,q}$,* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *satisfies LWE with respect to $\mathbf{A}^*$ with noise distribution $\overline{\chi} = \mathcal{U}([-\overline{B}, \overline{B}])$ where $\overline{B} = B^2 \cdot m^d \cdot 2^\lambda$.*

We refer to the full version for proofs of Claim 1, Lemma 2, and Corollary 1.

4.3 Instantiating the Parameters

Parameters. We first go through our parameters, and show that they satisfy the constraints of Definition 5.

Our candidate is a "degree-d" sampler, where $d \geq 2$ is a fixed constant integer. It expands LWE samples $\mathbf{B}_i \in \mathbb{Z}_q^{m \times k}$ to a matrix $\mathbf{B}^* \in \mathbb{Z}_q^{M \times K}$, using matrices $\mathbf{P} \leftarrow \chi^{M \times m^d}$ and $\mathbf{P'} \leftarrow \chi^{k^d \times K}$.[7] This expansion has stretch γ, in the sense that $MK = (mk)^\gamma$. w and W are the respective widths of the underlying matrices $\mathbf{A}_i \in \mathbb{Z}_q^{m \times w}$ and $\mathbf{A}^* \in \mathbb{Z}_q^{M \times W}$. δ is the succinctness parameter of our sampler.

χ denotes a B-bounded distribution used to sample $\mathsf{seed}_{\mathbf{B}^*}$, namely the matrices $\{\mathbf{E}_i\}_{i \in [d]}$, and we assume that $\mathrm{LWE}_{w,q,\chi}$ holds. β is a bound on $\|\mathbf{E}^*\|$ which depends on B.

$\overline{\chi}$ denotes a $\overline{B}$-bounded distribution such that LWE with respect to $\mathbf{A}^*$ holds (assuming LWE holding for some fixed parameters only dependent on the security parameter λ). α denotes a blow-up factor that defines the noise bound β_0 that the sampler is masking in the security property, namely $\beta_0 = \alpha\overline{B}$.

We gather the constraints on our parameters below:

- $N = MK$ //constraint of the sampler
- $(dmk + MW + WK) \leq N^\delta \cdot \mathrm{poly}(\lambda, \log q)$ for some $\delta < 1$ //δ-succinctness
- $M^2 \leq N^\delta \cdot \mathrm{poly}(\lambda, \log q)$ //for SRE succinctness
- $M \leq m^d$ //LWE with respect to $\mathbf{A}^*$ (Corollary 1)
- χ is a B-bounded distribution s.t. $\mathrm{LWE}_{w,q,\chi}$ holds. //base LWE assumption
- $\overline{B} = B^2 m^d \cdot 2^\lambda$ //LWE with respect to $\mathbf{A}^*$ (Corollary 1)
- $\beta = B^2(mkB)^d$ //bound on $\|\mathbf{E}^*\|$
- B large enough s.t. $\beta \geq \beta_0 \cdot 2^\lambda$ where $\beta_0 = \alpha\overline{B}$. //constraint of the sampler
- $q \geq 8\beta$. //constraint of the sampler

We additionally add the following constraints to ensure security:

- $\gamma < d/2$ //to avoid SOS attacks (Section 4.5).
- $M \leq m^d, K \leq k^d$ //to avoid rank attacks[8] (Section 4.5).

Next, we show our candidate sampler satisfies these constraints. Given the security parameter λ, fix a degree $d = \mathcal{O}(1)$, a dimension $w = w(\lambda)$, and a bound $B = B(\lambda)$. Given additional parameters $N \geq w^{6d}$ and α as input, our candidate sets the following parameters.

It fixes a stretch parameter $\gamma \in \left[\frac{2d}{2d-1/6}, d/2\right)$.

Set $m = N^{1/2d} \geq w^3$. It then defines the following "dimension" parameters k, M, K:

$$k = m^{\frac{2d}{\gamma}-1}, \quad M = m^{d-1/2}, \quad K = m^{d+1/2}$$

[7] In general, we can use a different (small) distributions D_P and $D_{P'}$ for $\mathbf{P}$, $\mathbf{P'}$. We only set $D_P = D'_P = \chi$ to minimize the number of distributions and parameters.

[8] The first constraint is redundant with the constraints of Corollary 1.

and $wm^{d-1} \leq W = \text{rank}\left(\overline{\mathbf{A}}^*\right) \leq m^d - (m-w)^d < dwm^{d-1} = \mathsf{width}\left(\overline{\mathbf{A}}^*\right)$ by construction of $\mathbf{A}^*$.[9] Note that the second inequality is strict as $m > w$,[10] that is, $\overline{\mathbf{A}}^*$ is rank deficient.

It then defines the following "bound" parameters $\overline{B}, \beta$:

$$\overline{B} = B^2 m^d \cdot 2^\lambda, \quad \beta = B^2 (mkB)^d,$$

where we assume that χ is B-bounded with $B \geq \frac{(\alpha \cdot 2^{2\lambda})^{1/d}}{k}$ such that $\mathsf{LWE}_{w,q,\chi}$ holds.[11]

Let $\overline{\chi} = \mathcal{U}([-\overline{B}, \overline{B}])$ be the uniform distribution over $[-\overline{B}, \overline{B}]$. It finally sets the modulus q as

$$q = 8\beta.$$

We show that the setting of parameters satisfy all the constraints described above. First, by definition, $N = m^{2d} = MK$. Furthermore:

$$\begin{aligned}
\mathsf{bitlength}(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) &= dmk \log q \;+\; M \cdot W \log q \;+\; W \cdot K \log q \\
&< \left(dm^{2d/\gamma} + dwm^{2d-3/2} + dwm^{2d-1/2}\right) \cdot \log q \\
&= \left(m^{2d-\frac{1}{6}} \;+\; dm^{2d/\gamma}\right) \cdot \log q \\
&= N^\delta \cdot \text{poly}(\lambda, \log q)
\end{aligned}$$

with $\delta = 1 - \frac{1}{12d} = \frac{2d-1/6}{2d}$, where we used $W < dwm^{d-1}$, $w \leq m^{1/3}$, which follows as $N \geq w^{6d}$ and $m = N^{1/2d}$, and $1/\gamma \leq \delta$.

We furthermore have $M^2 = m^{2d-1} \leq N^\delta$.

We have by construction: $\overline{B} = B^2 m^d \cdot 2^\lambda$, $\beta_0 = \alpha\overline{B}$, $\beta = \beta_0 \cdot 2^\lambda$, $\beta \geq B^2 (mkB)^d$ and $q = 8\beta$, so that the constraint $\beta \geq \beta_0 \cdot 2^\lambda$ can be rewritten as:

$$B^2(mkB)^d \geq \alpha \cdot 2^\lambda \cdot (B^2 m^d 2^\lambda),$$

which is exactly our constraint on B.

Last, we have $\gamma < d/2$ by definition, $M = m^{d-1/2} \leq m^d$, and $K = m^{d+1/2} \leq (m^3)^d$.

Remark 8 (Length of the CRS). As noted in Remark 5, a long CRS is required for security to hold if we allow arbitrary auxiliary information aux. We note this is the case for the parameters of Conjecture 1. Indeed: $\mathsf{bitlength}(\mathbf{P}') = k^d K \;\log q \geq m^{4d+1/2} \;\log q \geq N \;\text{poly}(\lambda, \log q) = m^{2d} \;\text{poly}(\lambda, \log q)$.

---

[9] We prove that $\text{rank}\left(\overline{\mathbf{A}}^*\right) \leq m^d - (m-w)^d$ in Section 4.5, paragraph Rank of $\mathbf{A}^*\mathbf{S}^*$.

[10] Writing $m = m' + w$ where $m' > 0$, the difference $(m'+w)^d - (m'^d + dw(m'+w)^{d-1})$ is the sum of monomials in m', w with positive coefficients.

[11] This is without loss of generality by defining for instance $\chi' = \chi + [-B, B]$ where B' is large enough to satisfy the previous constraint. A direct reduction ensures that if LWE holds with χ, then it holds with χ'.

Remark 9 (Parameters as a function of γ.). Our construction induces different parameters, according the choice of γ. The main affected parameter is k, which goes from $k = m^{3+o(1)}$ to $k \approx m^{2d}$. We note here that it also makes sense to use a constant $\gamma \in \left(1, \frac{2d}{2d-1/6}\right]$ for our construction. The only difference is that the succinctness of the scheme then becomes $1/\gamma$ as opposed to $1 - \mathcal{O}(1/d)$.

We gather some example parameters in the table below. In all cases, we set $d \geq 4$ be a constant, $m \geq w^3$ so that $N = m^{2d}$, $M = m^{d-1/2}$ and $K = m^{d+1/2}$. The third column represent the components that should have size bounded by N^δ to satisfy δ-succinctness.

Stretch γ	Dimension k	$M^2 + \mathsf{bitlength}(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*)$	Succinctness δ
$\gamma = d/3$	$k = m^5$	$\mathcal{O}(m^{2d-1/6})$	$\delta = 1 - \frac{1}{12d}$
$\gamma = \frac{2d}{2d-1/6}$	$k = m^{2d-7/6}$	$\mathcal{O}(m^{2d-1/6})$	$\delta = 1 - \frac{1}{12d} = 1/\gamma$
$\gamma = \frac{2d}{2d-\epsilon}$	$k = m^{2d-\epsilon-1}$	$\mathcal{O}(m^{2d-\epsilon})$	$\delta = 1/\gamma$

Fig. 1. Example parameters. In the above, we fix a constant $d \geq 4$ and $w = w(\lambda)$. The output size is $N = m^{2d}$ where $N \geq w^{6d}$.

Next, we state our main conjecture for our candidate, namely that it satisfies the weak notion of security of Definition 6. Looking ahead, thanks to Theorem 3, this suffices to imply iO.

Conjecture 1 (Conjectured security). Let χ be a B-bounded distribution, and assume $\mathrm{LWE}_{w,q,\chi}$ holds. Then ($\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand}$) with any of the parameters above satisfies weak β_0-flooding (Definition 6), where $\beta_0 = \alpha\overline{B}$.

Remark 10 (Security as a function of d). Our constructions decouples the stretch γ, defined as $(\mathsf{bitlength}(\{\mathbf{B}_i\}_{i\in[d]}))^\gamma = \mathsf{bitlength}(\mathbf{B}^*)$ (up to polynomial factors in $\lambda, \log q$), from the degree d. In particular, for a fixed (constant) stretch $\gamma \geq \frac{2d}{2d-1/6}$, we expect Conjecture 1 to be weaker as d increases.

Next, combining the above with Theorem 3, we describe two distributions whose indistinguishability would imply the existence of succinct LWE sampler with θ-flooding (Definition 5) for some parameter θ. Looking ahead, combined with Theorem 4, this suffices to imply an iO scheme.

Conjecture 2 (Stand-alone θ-flooding). Let $\beta_0 = \theta \cdot 2^\lambda$. With any of the parameters params described above, the following distributions Δ_b are indistinguishable:

$$\Delta_b = \Big(\mathbf{P}, \mathbf{P}', \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \quad \widehat{\mathbf{B}} = \mathbf{A}^*\mathbf{S}_0 + \mathbf{F},$$
$$\mathbf{C} = \mathbf{A}^*\mathbf{R} + \mathbf{E} - b\mathbf{G}, \quad \mathbf{E}^* + \mathbf{E} \cdot \mathbf{G}^{-1}(\widehat{\mathbf{B}}) - b\mathbf{F}\Big)$$

where

$$\mathbf{P} \leftarrow \chi^{M \times m^d}, \quad \mathbf{P}' \leftarrow \chi^{k^d \times K},$$

$$\mathsf{seed}_{\mathbf{B}^*} = \{\mathbf{B}_i\}_{i \in [d]} \in (\mathbb{Z}_q^{m \times w})^d$$

$$\widehat{\mathbf{B}} \in \mathbb{Z}_q^{M \times K}, \quad \text{where} \quad \mathbf{S}_0 \leftarrow \mathbb{Z}_q^{W \times K}, \quad \mathbf{F} \leftarrow \chi_{\mathsf{flood}}^{M \times K}$$

$$\mathbf{C} \in \mathbb{Z}_q^{M \times M \log q}, \quad \text{where} \quad \mathbf{R} \leftarrow \mathbb{Z}_q^{W \times M \log q}, \quad \mathbf{E} \leftarrow \overline{\chi}^{M \times M \log q}$$

where $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) \leftarrow \mathsf{LWEGen}(\mathsf{params}, \mathbf{P}, \mathbf{P}')$, $\mathbf{B}^* = \mathsf{Expand}(\mathsf{params}, \mathbf{P}, \mathbf{P}', \mathsf{seed}_{\mathbf{B}^*})$, and $\mathbf{E}^* = \mathbf{B}^* - \mathbf{A}^*\mathbf{S}^*$. Furthermore, $\overline{\chi}$ is a noise distribution such that LWE with respect to $\mathbf{A}^*$ holds, and χ_{flood} is a β_0-bounded distribution that floods θ-bounded distributions.

4.4 Alternate Candidate Construction

In the full version, we present a variant of the construction in Section 4.1. The main intuition is that this new variant sums T copies of the candidate of Section 4.1, but reusing the same matrices $\mathbf{A}_i$ across all copies. We refer to the full version for a complete description of that candidate.

4.5 Cryptanalysis

Recall that security of a succinct LWE sampler requires

$$(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{E}^* - \mathbf{Z}_b, \mathsf{aux}_b)$$

to hide b for appropriate aux_b and small $\mathbf{Z}_b$.

Ignoring the auxiliary information related to the sampler for now, the crucial requirement is that $\mathbf{E}^* - \mathbf{Z}_b$ (or, equivalently, $\mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_b$) hides b for sufficiently small $\mathbf{Z}_b$. As noted in the technical overview, pseudorandomness of $\mathbf{E}^*$ cannot hold given $\mathsf{seed}_{\mathbf{B}^*}$: one can compute $\mathbf{B}^* - \mathbf{E}^*$ and check that it is low rank. Nonetheless, as a sanity check, we would like to ensure that the marginal distribution of $\mathbf{E}^*$ is pseudorandom *by itself*, i.e. in the absence of $\mathsf{seed}_{\mathbf{B}^*}$. We first describe some attacks on the pseudorandomness of $\mathbf{E}^*$, and their influence on our parameters in Section 4.3.

Linearization Attacks. A strong break for the pseudorandomness of $\mathbf{E}^*$ is to recover the initial errors $\mathbf{E}_i \in \mathbb{Z}^{m \times k}$ such that $\mathbf{P}\left(\bigotimes_{i=1}^d \mathbf{E}_i\right)\mathbf{P}' = \mathbf{E}^*$. This would be enough to break pseudorandomness: only a small fraction of small $\mathbf{E}^* \in \mathbb{Z}^{M \times K}$ have such a succinct description as long as $N = MK$ is large enough compared to m and k (say $MK = (mk)^\gamma$ for some constant $\gamma > 1$).

One way of recovering the $\mathbf{E}_i$'s given $\mathbf{E}^*$, $\mathbf{P}$ and $\mathbf{P}'$ is to view the equation

$$\mathbf{P}\left(\bigotimes_{i=1}^{d} \mathbf{E}_i\right)\mathbf{P}' = \mathbf{E}^*$$

as a set of *linear* equations with the $(mk)^d$ variables

$$X_{i_1,j_1,\cdots,i_d,j_d} = \mathbf{E}_1^{i_1,j_1} \times \cdots \times \mathbf{E}_d^{i_d,j_d}$$

where $i_1, \cdots, i_d \in [m]$ and $j_1, \cdots, j_d \in [k]$, and where $\mathbf{E}^{i,j}$ denotes the (i,j)th component of $\mathbf{E}$. In particular, this is solvable as long as the number of equations is no smaller than the number of variables, that is:

$$MK \geq (mk)^d.$$

Our choice of parameters reflects security against linearization attacks. We also note that the linearization attack (in contrast to the sum of squares attack) works just as well over any finite field as it does over the integers.

Low-Degree Polynomials and Sum of Squares. The recovery attack described above can be generically improved using the more refined *sum of squares* (SOS) attacks. These ensure that pseudorandomness of $\mathbf{E}^*$ cannot hold whenever

$$MK \geq (mk)^{d/2}.$$

We refer the reader to [BHJ+19] for more details on sum of squares attacks. In our scheme, we explicitly require that the stretch of our sampler, namely γ such that $MK = (mk)^\gamma$, is smaller than $d/2$.

Security when $m = 1$. When $m = 1$, $\mathbf{P}$ is a scalar that we will ignore. We are given

$$\mathbf{e}^* = \left(\bigotimes_{i=1}^{d} \mathbf{e}_i \right) \mathbf{P}'$$

which is a vector of length K. Since $\bigotimes_{i=1}^{d} \mathbf{e}_i$ is simply the set of all degree-d multilinear monomials with a variable from each of the $\mathbf{e}_i$, this can be interpreted as evaluating K degree-d polynomials with Gaussian coefficients on the dk variables in $\mathbf{e}_1, \ldots, \mathbf{e}_d$. Since $K \ll k^{d/2}$, neither linearization nor sum of squares seems to apply [BHJ+19].

The work of Kosov [Kos20] tells us each entry in $\mathbf{E}^*$ by itself, namely a polynomial with Gaussian coefficients evaluated on Gaussian inputs, comes from a noise-flooding distribution (for mild choices of parameters).

This analysis also points to the qualitative distinction between our assumption and the analysis above for $m = 1$. When $m = 2$, for example, we obtain MK polynomials evaluated on a number of correlated random variables. That is, setting the two rows of $\mathbf{E}_i$ to be $\mathbf{e}_{i1}$ and $\mathbf{e}_{i2}$,

$$\mathbf{E}^* = \mathbf{P} \begin{bmatrix} \mathbf{e}_{11} \otimes \mathbf{e}_{21} \otimes \cdots \otimes \mathbf{e}_{d1} \\ \mathbf{e}_{12} \otimes \mathbf{e}_{21} \otimes \cdots \otimes \mathbf{e}_{d1} \\ \vdots \\ \mathbf{e}_{12} \otimes \mathbf{e}_{22} \otimes \cdots \otimes \mathbf{e}_{d2} \end{bmatrix} \mathbf{P}'$$

To the best of our knowledge, all attacks described above still fail. In fact, we don't even have an attack if $\mathbf{P} = \mathbf{I}_{2^d}$ was the identity and $M = 2^d$. However, this is certainly a cryptanalytic avenue worth pursuing in the future.

Rank Attacks. Towards analyzing the case of larger m, we attempt another class of attacks which consist of looking at the *rank* of the various matrices that arise in the assumption.

Rank Attack on $\mathbf{E}^*$. Note that a random (e.g. Gaussian) $\mathbf{E}^*$ would be full-rank with overwhelming probability. In particular, as

$$\mathbf{E}^* = \mathbf{P}\left(\bigotimes_{i=1}^{d} \mathbf{E}_i\right)\mathbf{P}',$$

where $\mathbf{P} \in \mathbb{Z}^{M\times m^d}$ and $\mathbf{P}' \in \mathbb{Z}^{k^d\times K}$, the rank of $\mathbf{E}^*$ is at most the rank of $\mathbf{P}, \mathbf{P}'$. In particular, $\mathbf{P}$ and $\mathbf{P}'$ need to be full-rank and compressing, meaning that $M \leq m^d$ and $K \leq k^d$, respectively. Our setting of parameters (see Section 4.3) ensure these restrictions hold.

The rank of $\bigotimes_{i=1}^{d} \mathbf{E}_i$ is the product of the ranks of $\mathbf{E}_i$, and is therefore, $\mathsf{min}(m^d, k^d)$ with high probability. Heuristically, then, the rank of $\mathbf{E}^*$ is exactly $\mathsf{min}(K, M)$ with high probability, as long as the Gaussians have sufficiently large width, a statement that we verified experimentally.

Rank Attack on $\mathbf{A}^*\mathbf{S}^*$. Note that if $\mathbf{A}^*\mathbf{S}^*$ is computationally indistinguishable from $\mathbf{A}^*\mathbf{S}'$ for a uniformly random $\mathbf{S}'$ given $\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathsf{aux}_b$, then the pre-condition implies the post-condition in Definition 5, guaranteeing security. Thus, we evaluate possible distinguishers between $\mathbf{A}^*\mathbf{S}^*$ and $\mathbf{A}^*\mathbf{S}'$.

One such class of attacks consist in comparing the rank of $\mathbf{A}^*\mathbf{S}^*$ to the rank of $\mathbf{A}^*$. We heuristically and experimentally analyzed the ranks of $\mathbf{A}^*$ and $\mathbf{A}^*\mathbf{S}^*$ to reason about these attacks.

First, note that $\overline{\mathbf{A}}^*\overline{\mathbf{S}}^* = \mathbf{A}^*\mathbf{S}^*$. Recall that the matrices $\mathbf{A}_i \in \mathbb{Z}_q^{m\times w}$ are random and therefore w.h.p. full-rank (i.e., rank w). Let $\mathbf{A}_i^\perp \in \mathbb{Z}_q^{(m-w)\times m}$ be a basis for the left-kernel of $\mathbf{A}_i$, that is, they are rank-$(m-w)$ matrices such that

$$\mathbf{A}_i^\perp \mathbf{A}_i = 0 \pmod q$$

We note that w.h.p. the rank of the matrix

$$(\,\mathbf{A}_1 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \mathbf{I}_m \otimes \mathbf{A}_2 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \cdots \,\|\, \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \otimes \mathbf{A}_d)$$

is at most $m^d - (m-w)^d \approx dwm^{d-1} - d^2w^2m^{d-2}/2$ (the approximation assumes that $m \gg w$ which is the case for us) since the row-span of $\mathbf{A}^*$ is contained in the right kernel of $(\mathbf{A}_1^\perp \otimes \cdots \otimes \mathbf{A}_d^\perp)$, and the latter has rank $m^d - (m-w)^d$. Our experiments indicate that the rank is indeed $m^d - (m-w)^d$ w.h.p. In other words, this matrix is *rank-deficient* by approximately $d^2w^2m^{d-2}/2$.

Heuristically,

$$\overline{\mathbf{A}}^* = \mathbf{P}\cdot(\,\mathbf{A}_1 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \mathbf{I}_m \otimes \mathbf{A}_2 \otimes \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \,\|\, \mathbf{I}_m \otimes \cdots \otimes \mathbf{I}_m \otimes \mathbf{A}_d)$$

has the same rank since $\mathbf{P} \in \mathbb{Z}^{M\times m^d}$ is Gaussian and nearly full-rank, i.e., rank $M \approx m^{d-1/2}$. That is, w.h.p., (heuristically)

$$\mathsf{rank}(\overline{\mathbf{A}}^*) = m^d - (m-w)^d$$

Also, heuristically, $\overline{\mathbf{A}}^*\overline{\mathbf{S}}^*$ has this rank as long as $\overline{\mathbf{S}}^*$ has sufficiently many columns, i.e. as long as K is large enough compared to $\mathsf{rank}(\overline{\mathbf{A}}^*)$. (Note that the entries of $\overline{\mathbf{A}}^*$ and $\overline{\mathbf{S}}^*$ are correlated.)

To test these heuristic statements, we ran experiments for $d = 3$ and a range of values of m, k and q. We found that $\overline{\mathbf{A}}^*$ had rank $m^d-(m-w)^d$ as expected (in all the runs of our experiment, suggesting a high probability statement). We also found that when $k \geq m$ and K is large enough so that $\mathbf{S}^*$ is wide, $\overline{\mathbf{A}}^*\overline{\mathbf{S}}^* = \mathbf{A}^*\mathbf{S}^*$ also had rank $m^d - (m-w)^d$ with high probability. This is the same as one would expect from $\mathbf{A}^*\mathbf{S}'$ for a random $\mathbf{S}'$, suggesting that rank attacks fail.

4.6 Cryptanalytic Challenges

We describe a few cryptanalytic challenges and how they relate to our candidate and our assumptions. For each of these problems, we can also consider easier challenges where (a) the challenger also gets $\mathbf{A}^*$; and (b) we replace $\mathbf{P}$ with the identity matrix.

Pseudo-flooding in the Absence of $\mathsf{seed}_{\mathbf{B}^*}$. Our intuition says that for any two low-norm matrices $\mathbf{Z}_0$ and $\mathbf{Z}_1$, $\mathbf{E}^* + \mathbf{Z}_b$ hides b. Concretely, let χ be a discrete Gaussian of sufficiently large parameter σ. A challenge is to come up with matrices $\mathbf{Z}_0$ and $\mathbf{Z}_1$ where $||\mathbf{Z}_b|| < \sigma/2^\lambda$ such that the bit b can be recovered given

$$\mathbf{P}\left(\bigotimes_{i=1}^{d}\mathbf{E}_i\right)\mathbf{P}' + \mathbf{Z}_b \ .$$

We note that when $m = 1$ and $\mathbf{P} = 1$, as argued above, this seems to follow from the noise-flooding properties of random (e.g. Gaussian) polynomials [BHJ+19].

Pseudo-flooding in the Presence of $\mathsf{seed}_{\mathbf{B}^*}$. Our stronger notion of security (Definition 5) would imply that it would be hard to recover b from

$$(\mathsf{seed}_{\mathbf{B}^*},\ \mathbf{A}^*\mathbf{S}^* + \mathbf{Z}_b,\ \mathbf{E}^* - \mathbf{Z}_b), \quad b \leftarrow \{0,1\}$$

for the following concrete distributions of $\mathbf{Z}_0, \mathbf{Z}_1$:

– (norm and ideal membership) $\mathbf{Z}_0$ is drawn from a Gaussian, and $\mathbf{Z}_1 = 2\mathbf{Z}_0$, and q is odd. In particular, an attacker that manages to learn the parity of $\mathbf{Z}_b$ or accurately approximate the norm of $\mathbf{Z}_b$ will be able to learn b.
– (subspace membership) $\mathbf{Z}_b = \mathbf{E}_0\mathbf{M} + b\hat{\mathbf{E}}$ where $\|\mathbf{E}_0\| \gg \|\hat{\mathbf{E}}\|$ and $\mathbf{M}$ is a public low-norm matrix. The distribution here is closely related to that for weak flooding. Here, $\|\mathbf{Z}_0\| \approx \|\mathbf{Z}_1\|$, but an attacker that manages to learn whether $\mathbf{Z}_b$ lies in the row span of $\mathbf{M}$ will be able to learn b.

In both cases, an attacker could try to exploit the leakage on b from $\mathbf{A}^*\mathbf{S}^*+\mathbf{Z}_b$ or from $\mathbf{E}^* - \mathbf{Z}_b$. For instance, an efficient algorithm that recovers $\mathbf{E}^*$ from $\mathsf{seed}_{\mathbf{B}^*}$ or one that recovers b from $\mathbf{E}^* - \mathbf{Z}_b$ solves this problem.

Distinguishing $\mathbf{A}^*\mathbf{S}^*$ *from* $\mathbf{A}^*\mathbf{S}'$. As described above, we think the following claim is plausible:

$$\mathbf{A}^*\mathbf{S}^* \approx_c \mathbf{A}^*\mathbf{S}'$$

where $\mathbf{S}' \leftarrow \mathbb{Z}_q^{W\times K}$. As $\mathbf{A}^*\mathbf{S}^* = \overline{\mathbf{A}}^* \cdot \overline{\mathbf{S}}^*$ (where $\overline{\mathbf{A}}^*, \overline{\mathbf{S}}^*$ are defined in Section 4.1), and given that $\mathbf{A}^*$ and $\overline{\mathbf{A}}^*$ have the same column span, this is equivalent to

$$\overline{\mathbf{A}}^* \cdot \overline{\mathbf{S}}^* \approx_c \overline{\mathbf{A}}^* \cdot \mathbf{S}''$$

where $\mathbf{S}'' \leftarrow \mathbb{Z}_q^{dwm^{d-1}\times K}$, and $\overline{\mathbf{A}}^*, \overline{\mathbf{S}}^*$ have closed form expressions described in Section 4.1.

A distinguisher here does not immediately break strong or weak-flooding, but we believe it constitutes strong evidence that strong-flooding is false.

5 Our Succinct Randomized Encoding Construction

Let $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ be a succinct LWE sampler (Definition 5) with parameters to be determined later.

We now describe our SRE for the family $\mathcal{F}_{\ell,N,t} = \{f : \{0,1\}^\ell \to \{0,1\}^N\}$ of depth-t circuits. Let q be a modulus and $\overline{\chi}$ be a $\overline{B}$-bounded distribution to be determined later.

Let $g(t) = \mathcal{O}(t)$ be the function defined in Definition 3.

- $\mathsf{CRSGen}(1^\lambda, \mathcal{F}_{\ell,N,t})$: Output $\mathsf{crs} \leftarrow \mathsf{SampCRSGen}(1^\lambda, 1^N, N^{g(t)})$. It in particular includes parameters $\mathsf{params} = (q, M, K, \overline{\chi}, \overline{B})$.
- $\mathsf{Encode}(\mathsf{crs}, f, x)$: Compute $(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{A}^*, \mathbf{S}^*) \leftarrow \mathsf{LWEGen}(\mathsf{crs})$, where $\mathbf{A}^* \in \mathbb{Z}_q^{M\times W}$, $\mathbf{S}^* \in \mathbb{Z}_q^{W\times K}$.
  Sample $\mathbf{R} \leftarrow \{0,1\}^{W\times \ell M\log q}$, and $\mathbf{E} \leftarrow \overline{\chi}^{M\times \ell M\log q}$. Compute
  $$\mathbf{C} = \mathbf{A}^*\mathbf{R} + x \otimes \mathbf{G} + \mathbf{E} \in \mathbb{Z}_q^{M\times \ell M\log q},$$
  where we view $x \in \{0,1\}^{1\times\ell}$ as a row vector, and compute $(\mathbf{R}_{f,x}, \mathbf{E}_{f,x}) = \mathsf{Eval}_{\mathsf{open}}(f, \mathbf{A}^*, x, \mathbf{R}, \mathbf{E})$.
  Output:
  $$C = (\mathsf{seed}_{\mathbf{B}^*},\ \mathbf{C},\ \mathbf{A}^*,\ (\mathbf{R}_{f,x} + \mathbf{S}^*)).$$
- $\mathsf{Decode}(\mathsf{crs}, C, f))$: On input $C = (\mathsf{seed}_{\mathbf{B}^*}, \mathbf{C}, \mathbf{A}^*, \mathbf{V})$, compute $\mathbf{C}_f = \mathsf{Eval}(f, \mathbf{C})$, and $\mathbf{B}^* = \mathsf{Expand}(\mathsf{crs}, \mathsf{seed}_{\mathbf{B}^*})$. Output
  $$f(x) = \mathsf{Round}_{q/2}\left(\mathbf{C}_f + \mathbf{B}^* - \mathbf{A}^* \cdot \mathbf{V}\right) \in \{0,1\}^{M\times K}.$$

Theorem 4. *Suppose* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *is a succinct LWE sampler satisfying* δ*-succinctness and* β_0*-flooding (Definition 5) with* $\beta_0 = \overline{B} \cdot N^{g(t)}$*. Suppose furthermore that:*

$$M^2 = N^\delta \cdot \mathrm{poly}(\lambda, \ell, t).$$

Then $(\mathsf{CRSGen}, \mathsf{Encode}, \mathsf{Decode})$ *is an SRE for* $\mathcal{F}_{\ell,N,t}$ *satisfying* δ*-succinctness.*

Next, we show that the construction above satisfies correctness and succinctness.

Claim 3 (Correctness). *Suppose* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *satisfy the parameters constraints and correctness Definition 5. Then* $(\mathsf{CRSGen}, \mathsf{Encode}, \mathsf{Decode})$ *is correct.*

Proof. Define $\mathbf{V} = (\mathbf{R}_{f,x} + \mathbf{S}^*)$. By Definition 3, we have

$$\mathbf{C}_f + \mathbf{B}^* - \mathbf{A}^* \cdot (\mathbf{R}_{f,x} + \mathbf{S}^*) = f(x) \cdot q/2 + \mathbf{E}_{f,x} + \mathbf{E}^*.$$

Let $\beta_0 = \overline{B} \cdot N^{g(t)}$. The setting of parameters β, $\overline{B}$ and q from $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ imply $\|\mathbf{E}\| \leq \overline{B}$ and therefore $\|\mathbf{E}_{f,x}\| \leq \overline{B} M^{g(t)} \leq N^{g(t)} = \beta_0$ by definition of g (Definition 3), and using $M \leq N$. Furthermore $\beta \geq \beta_0 \cdot 2^\lambda$ and $q \geq 8\beta$ so that $\|\mathbf{E}_{f,x} + \mathbf{E}^*\| < q/4$, and therefore

$$\mathsf{Round}_{q/2}\left(\mathbf{C}_f + \mathbf{B}^* - \mathbf{A}^* \cdot \mathbf{V}\right) = \mathsf{Round}_{q/2}\left(f(x) \cdot q/2 + \mathbf{E}_{f,x} + \mathbf{E}^*\right) = f(x).$$

Claim 4. *Suppose the sampler* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *is* δ*-succinct (Definition 5), and suppose that the sampler furthermore satisfies*

$$M^2 = N^\delta \cdot \mathrm{poly}(\lambda, \ell, t).$$

Then $(\mathsf{CRSGen}, \mathsf{Encode}, \mathsf{Decode})$ *is* δ*-succinct.*

Proof. The setting of the parameters implies $\log q = \mathrm{poly}(\lambda, t)$. Then $\ell M^2 \log^2 q = N^\delta \cdot \mathrm{poly}(\lambda, \ell, t)$.

Furthermore $\mathbf{V} = (\mathbf{R}_{f,x} + \mathbf{S}^*) \in \mathbb{Z}_q^{W \times K}$ and therefore $\mathsf{bitlength}(\mathsf{seed}_{\mathbf{B}^*}, \mathbf{C}, \mathbf{A}^*, \mathbf{V}) \leq N^\delta \cdot \mathrm{poly}(\lambda, \ell, t)$ by δ-succinctness of $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$. Therefore the SRE is δ-succinct.

5.1 Security

Claim 5 (Indistinguishability-based security.). *Let* $f : \{0,1\}^\ell \to \{0,1\}^N$ *of depth* t*, and* $x_0, x_1 \in \{0,1\}^\ell$ *such that* $f(x_0) = f(x_1)$*. Suppose* $(\mathsf{SampCRSGen}, \mathsf{LWEGen}, \mathsf{Expand})$ *is secure (Definition 5), and LWE hold. Then:*

$$(\mathsf{crs}, \mathsf{Encode}(\mathsf{crs}, f, x_0)) \approx_c (\mathsf{crs}, \mathsf{Encode}(\mathsf{crs}, f, x_1)),$$

where $\mathsf{crs} \leftarrow \mathsf{CRSGen}(1^\lambda, \mathcal{F}_{\ell,N,t})$.

We refer to the full version for a proof of Claim 5. Combining Theorem 4 with our candidate succinct LWE sampler (Sections 4.1 and 4.3), noting that our proposed parameters in Section 4.3 satisfy $M^2 = N^\delta \cdot \mathrm{poly}(\lambda, \ell, t)$, gives a candidate SRE. Invoking Theorem 2, we obtain the following.

Corollary 2. *Assuming Conjecture 1 and sub-exponential LWE, there exists an iO scheme.*

We can furthermore use Theorem 3 to relax the requirement on our candidate succinct LWE sampler (Section 4.1), and only rely on weak security (Definition 6), thus obtaining the following.

Corollary 3. *Assuming Conjecture 2 and sub-exponential LWE, there exists an iO scheme.*

Acknowledgements

We thank Pravesh Kothari for his pointers to and conversations about the literature on SOS and low-degree polynomial attacks. LD and VV were supported by DARPA under Agreement No. HR00112020023, a grant from the MIT-IBM Watson AI, a grant from Analog Devices, a Microsoft Trustworthy AI grant, and a DARPA Young Faculty Award. WQ completed part of this work during an internship at NTT Research. DW was supported by NSF grant CNS-1750795, CNS-2055510, and the Alfred P. Sloan Research Fellowship.

References

Agr19. Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 191–225. Springer, Heidelberg, May 2019. 1, 4

AJ15. Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 308–326. Springer, Heidelberg, August 2015. 8, 14

AJL+12. Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Heidelberg, April 2012. 12

AJL+19. Prabhanjan Ananth, Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. Indistinguishability obfuscation without multilinear maps: New paradigms via low degree weak pseudorandomness and security amplification. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 284–332. Springer, Heidelberg, August 2019. 4, 6

AP14. Jacob Alperin-Sheriff and Chris Peikert. Faster bootstrapping with polynomial error. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 297–314. Springer, Heidelberg, August 2014. 9

AP20. Shweta Agrawal and Alice Pellet-Mary. Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 110–140. Springer, Heidelberg, May 2020. 1

AR17. Shweta Agrawal and Alon Rosen. Functional encryption for bounded collusions, revisited. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 173–205. Springer, Heidelberg, November 2017. 4

BCG+18. Nir Bitansky, Ran Canetti, Sanjam Garg, Justin Holmgren, Abhishek Jain, Huijia Lin, Rafael Pass, Sidharth Telang, and Vinod Vaikuntanathan. Indistinguishability obfuscation for RAM programs and succinct randomized encodings. *SIAM J. Comput.*, 47(3):1123–1210, 2018. 13

BDGM19. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 407–437. Springer, Heidelberg, December 2019. 3, 12

BDGM20a. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate iO from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 79–109. Springer, Heidelberg, May 2020. 1, 2, 4, 5

BDGM20b. Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Factoring and pairings are not necessary for iO: Circular-secure LWE suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. 1

BGI+01. Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18. Springer, Heidelberg, August 2001. 1

BGL+15. Nir Bitansky, Sanjam Garg, Huijia Lin, Rafael Pass, and Sidharth Telang. Succinct randomized encodings and their applications. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 439–448. ACM Press, June 2015. 2, 13

BHJ+19. Boaz Barak, Samuel B. Hopkins, Aayush Jain, Pravesh Kothari, and Amit Sahai. Sum-of-squares meets program obfuscation, revisited. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 226–250. Springer, Heidelberg, May 2019. 9, 25, 27

BTVW17. Zvika Brakerski, Rotem Tsabary, Vinod Vaikuntanathan, and Hoeteck Wee. Private constrained PRFs (and more) from LWE. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 264–302. Springer, Heidelberg, November 2017. 3, 12

BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011. 7

BV15. Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *56th FOCS*, pages 171–190. IEEE Computer Society Press, October 2015. 8, 14

CHVW19. Yilei Chen, Minki Hhan, Vinod Vaikuntanathan, and Hoeteck Wee. Matrix PRFs: Constructions, attacks, and applications to obfuscation. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part I*, volume 11891 of *LNCS*, pages 55–80. Springer, Heidelberg, December 2019. 1

GGH+13. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013. 1

GH19. Craig Gentry and Shai Halevi. Compressible FHE with applications to PIR. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 438–464. Springer, Heidelberg, December 2019. 3, 12

GP21. Romain Gay and Rafael Pass. Indistinguishability obfuscation from circular security. In *STOC*, 2021. 1, 2, 4, 5

GR07. Shafi Goldwasser and Guy N. Rothblum. On best-possible obfuscation. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 194–213. Springer, Heidelberg, February 2007. 1

GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013. 3, 12

GVW15. Sergey Gorbunov, Vinod Vaikuntanathan, and Daniel Wichs. Leveled fully homomorphic signatures from standard lattices. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th ACM STOC*, pages 469–477. ACM Press, June 2015. 3, 12

HJL21. Sam Hopkins, Aayush Jain, and Huijia Lin. Counterexamples to new circular security assumptions underlying iO, 2021. 2, 9

JLMS19. Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials overa $\mathbb{R}$ to build $i\mathcal{O}$. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 251–281. Springer, Heidelberg, May 2019. 4, 6, 9, 10

JLS21. Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In *STOC*, 2021. 1, 4

Kil88. Joe Kilian. Founding cryptography on oblivious transfer. In *20th ACM STOC*, pages 20–31. ACM Press, May 1988. 7

Kos20. Egor Kosov. Distributions of polynomials in gaussian random variables under structural constraints, 2020. 25

LPST16. Huijia Lin, Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation with non-trivial efficiency. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 447–462. Springer, Heidelberg, March 2016. 2, 8, 13, 14

MP12. Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012. 3, 12

MW16. Pratyay Mukherjee and Daniel Wichs. Two round multiparty computation via multi-key FHE. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 735–763. Springer, Heidelberg, May 2016. 3, 12

Pei09. Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 333–342. ACM Press, May / June 2009. 11

PS19. Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114. Springer, Heidelberg, August 2019. 3, 12

PVW08. Chris Peikert, Vinod Vaikuntanathan, and Brent Waters. A framework for efficient and composable oblivious transfer. In David Wagner, editor, *CRYPTO 2008*, volume 5157 of *LNCS*, pages 554–571. Springer, Heidelberg, August 2008. 3, 12

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. 11

WW21. Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In *EUROCRYPT*, 2021. 1, 2, 4, 12, 13