Candidate Witness Encryption from Lattice Techniques

Rotem Tsabary[⋆]

Google Research, Israel

Abstract. Witness encryption (WE), first introduced by Garg, Gentry, Sahai and Waters in [GGSW13], is an encryption scheme where messages are encrypted with respect to instances of an **NP** relation, such that in order to decrypt one needs to know a valid witness for the instance that is associated with the ciphertext.

Despite of significant efforts in the past decade to construct WE from standard assumptions, to the best of our knowledge all of the existing WE candidates either rely directly on iO or use techniques that also seem to imply iO in the same way that they seem to imply WE.

In this work we propose a new hardness assumption with regard to lattice trapdoors and show a witness encryption candidate which is secure under it. Contrary to previous WE candidates, our technique is trivially broken when one tries to convert it to iO, which suggests that the security relies on a different mechanism. We view the gap between WE and iO as an analogue to the gap between ABE and FE and thus potentially significant.

Intuitively, the assumption says that *"the best an attacker can do with a trapdoor sample is to use it semi-honestly"* – i.e. that LWE with respect to a public matrix $\mathbf{A}$, given as auxiliary information a trapdoor sample $\mathbf{K} \leftarrow \mathbf{A}^{\mathsf{TD}}(\mathbf{B})$, is as hard as LWE with respect to the public matrix $[\mathbf{A}|\mathbf{B}]$ and no auxiliary information.

In order to formally utilize the assumption we define a notion of LWE oracles with generic distributions of public matrices and auxiliary information. This model allows to bound the hardness of LWE with respect to one distribution as a function of the hardness of LWE with respect to another distribution. Repeated arguments of this flavor can be used as a sequence of hybrids in order to gradually change the challenge that an adversary is facing while keeping track on the security loss in each step of the proof. Typically security proofs of LWE-based systems implicitly make arguments of this flavor for distributions that are indistinguishable, while our model allows to make relaxed arguments that in some cases suffice for the proof requirements.

[⋆] `rotem.ts0@gmail.com`. Research done while at Reichman University (IDC Herzliya), Israel, supported by ISF grant No. 1399/17.

1 Introduction

Witness Encryption. Witness encryption (WE), first introduced by Garg, Gentry Sahai and Waters in [GGSW13], is an encryption scheme where messages are encrypted with respect to instances of an **NP** relation, such that in order to decrypt a ciphertext one needs to know a valid witness for the instance that is associated with the ciphertext. Despite of significant efforts in the past decade to construct witness encryption from a standard hardness assumption, to the best of our knowledge the only construction that was proven secure under an explicit assumption relies on multi-linear maps [GLW14].

WE via iO. Indisitinguishability obfuscation (iO) is known to imply WE. In the past few years there have been major breakthroughs that lead to an iO scheme with provable security from standard hardness assumptions [Lin16, LV16, Lin17, LT17, AJL+19, Agr19, JLMS19, GJLS21, JLS21]. While the result serves as a proof of feasibility, it is challenging to grasp the intuition behind the resulting construction because it consists of a long sequence of reduction steps between various notions. It is therefore natural to ask whether there exist "simple" WE and iO candidates with provable security from standard assumptions. Recently Wee and Wichs proposed in [WW21] an iO candidate from lattices, based on the framework of [BDGM20], and proved its security assuming the existence of an oblivious LWE sampler. Other LWE-based iO candidates were suggested by Wichs and Zirdelis in [WZ17] and by Chen et al. in [CVW18], however without any security proof.

1.1 Our Contribution

We propose a new hardness assumption with regard to lattice trapdoors and show a witness encryption candidate which is secure under it. Intuitively, the assumption says that *"the best an attacker can do with a trapdoor sample is to use it semi-honestly"* – i.e. that LWE with respect to a public matrix $\mathbf{A}$, given as auxiliary information a trapdoor sample $\mathbf{K} \leftarrow \mathbf{A}^{\mathsf{TD}}(\mathbf{B})$, is as hard as LWE with respect to the public matrix $[\mathbf{A}|\mathbf{B}]$ and no auxiliary information.

The assumption can be viewed as a generalization of provable cases that are widely used in the literature, where the two extreme cases are when $\mathbf{B}$ is trivially LWE-hard (e.g. uniformly random) and when $\mathbf{B}$ is trivially LWE-broken (e.g. the gadget matrix). We believe that it also captures the intuition behind other structured LWE scenarios that seem to be secure but fail to pass traditional analysis techniques, so any counter example would greatly improve our understanding of the lattice toolbox.

In order to formally state the assumption and use it in the security proof of the scheme we define a notion of LWE challenges with general distributions of public matrices and auxiliary information. This model allows us to bound the hardness of LWE with respect to one distribution as a function of the hardness of LWE with respect to another distribution. Repeated arguments of this flavor can be used as a sequence of hybrids in order to gradually change the challenge that an adversary is facing while keeping track on the security loss in each step of the proof. Typically security proofs of LWE-based systems implicitly make arguments of this flavor for distributions that are indistinguishable, while our model allows to make relaxed arguments that in some cases suffice for the proof requirements.

We note that under some circumstances a polynomial number of trapdoor samples can lead to exponentially many public matrices that are *potentially* accessible, so the model is defined such that the number of public matrices for which the adversary gets to see an LWE expression is bounded by its running time. This is enforced by modeling the LWE experiment as a game between the adversary and an oracle that has all of the accessible matrices hard-wired in it, where the oracle provides upon request LWE samples with respect to any of these matrices.

Comparison with Previous WE Candidates

WE vs. iO. To the best of our knowledge all of the existing WE candidates either rely directly on iO or use techniques that also seem to imply iO in the same way that they seem to imply WE (see the discussion above). Contrary to that, our technique is trivially broken when one tries to convert it to iO. In particular it is easy to break the underlying LWE assumption given a valid witness, which is not the case in a generic iO-based WE construction. This suggests that the differences between our technique and existing candidates are not merely technical, rather, the security is based on a different mechanism and our assumption is possibly tighter. We view the gap between WE and iO as an analogue to the gap between ABE and FE and thus potentially significant.

Number of Hybrids. The security analysis of our candidate requires $2^{\mathrm{poly}(\ell)}$ hybrids where ℓ is the witness length. This is compatible with the barriers that were discussed in [GGSW13,GLW14]. We note that the only other explicit candidate with a security proof (the MMaps-based construction of [GLW14]) requires $\mathrm{poly}(2^{\ell})$ hybrids.

2 Overview of Techniques

Our candidate conceptually consists of two layers. First, we define and show the existence of *resilient* branching programs that behave in a predicted manner when they are executed with a sequence of bits that violate the index-to-input map. This part of the work is information-theoretic in nature and relies on an iO candidate suggested by [CHVW19]. The second layer takes a branching program of this form and generates a collection of matrices and [GGH15] encodings that correspond to the nodes and edges respectively of the branching program. The matrices of the end layer are generated according to a secret-sharing access structure that is provided as part of the resilient branching program, and the ciphertext is provided as an LWE challenge with respect to the start layer.

The security analysis of the scheme takes two main steps. First, we use the security assumption about lattice trapdoors to gradually replace the GGH15 encodings with "accessible matrices" that the adversary could derive by using any sequence of trapdoors of his choice. This is done in the direction of the computation of the branching program (it is essentially a BFS over the branching program nodes). This can be interpreted as a reduction from a security game of the scheme where the trapdoor samples can be used maliciously to a security game of the scheme where the trapdoor samples can only be used as intended.

In the second part of the security analysis we prove from standard assumptions the security of the scheme in the latter game. To formalize this game, we consider an LWE experiment where the adversary can receive upon request LWE samples with respect to any accessible node in the branching program, and bound the success probability of all PPT adversaries in this experiment. In this part of the proof we use the fact that $f(x) = 0$ for all $x \in \{0,1\}^\ell$ (since security should hold only when there are no valid witnesses) and the information-theoretic properties of the resilient branching program when it is executed with corrupted sequences. This part of the proof takes $2^{\mathrm{poly}(\ell)}$ hybrids since it handles each evaluation sequence individually.

2.1 Notations

For ease of exposition we treat vectors as row vectors by default. We let $\mathbf{v}[i]$ denote the ith entry of the vector $\mathbf{v}$ and we let $\mathbf{M}[i,j]$ denote the entry in the ith row and jth column of the matrix $\mathbf{M}$. We let $\equiv_\lambda$ denote computational/statistical indistinguishability with respect to a security parameter λ.

2.2 Branching Programs with Resiliency to Corrupted Inputs

This section is based on an iO candidate suggested by [CHVW19]. The computational model that is used in this work is a generalized form of matrix branching programs. Recall that a width-w length-t matrix BP that computes a predicate $f : \{0,1\}^\ell \to \{0,1\}$ is described by a start state vector $\mathbf{v}_{start} \in \{0,1\}^w$ where $\|\mathbf{v}_{start}\|_1 = 1$, an index-to-input map $\rho : [t] \to [\ell]$ and a tuple of $2t$ evaluation matrices $\{\mathbf{M}_{j,b}\}_{j\in[t],b\in\{0,1\}}$, where the guarantee is that for all x the first entry of $\mathbf{v}_x \in \{0,1\}^w$ equals $f(x)$, where

$$\mathbf{v}_x = \mathbf{v}_{start} \prod_{j=1}^{t} \mathbf{M}_{j,x_{\rho(j)}} \ .$$

By Barrington's Theorem [Bar89], each $f \in \mathrm{NC}^1$ can be computed by a polynomial-sized BP whose evaluation matrices are permutations. A central problem that arises when one attempts to use branching programs in the context of witness encryption is handling evaluation sequences $z \in \{0,1\}^t$ that are inconsistent with any input string $x \in \{0,1\}^\ell$. This happens whenever there exist indices $i, i' \in [t]$ for which $\rho(i) = \rho(i')$ but $z[i] \neq z[i']$. The problem is that in a BP-based solution, for correctness we typically would provide in the ciphertext some information that allows to evaluate the BP with respect to any sequence $z \in \{0,1\}^t$, while in the security proof we can only make assumptions about the BP output when it is executed with non-corrupted sequences z.

To overcome this problem we define a generalized form of matrix BP with a related security notion, where it is required that the BP outputs 0 for any x for which $f(x) = 0$ *and also for any corrupted sequence* z. In this model we allow start vectors $\mathbf{v}_{start} \in \{0,1\}^w$ with multiple 1 entries, i.e. $\|\mathbf{v}_{start}\|_1 \geq 1$, which can be thought of as if the BP holds multiple active nodes that are being updated simultaneously according to the same sequence of input bits z. We require an additional map $F : \{0,1\}^w \to \{0,1\}$ that dictates the final output of the BP according to the active nodes in the end layer. That is, for every sequence $z \in \{0,1\}^t$ we define

$$\mathbf{v}_z = \mathbf{v}_{start} \prod_{j=1}^{t} \mathbf{M}_{j,z_j}$$

and we require that if z is consistent with some input x (according to ρ) then $F(\mathbf{v}_z) = f(x)$, and if z is corrupted then $F(\mathbf{v}_z) = 0$. This allows us in the security proof to rely on the assumption that for all $z \in \{0,1\}^t$ it holds that $F(\mathbf{v}_z) = 0$. We note that for any $f \in \mathrm{NC}^1$ there is a degenerate width-2ℓ BP of this form in which the evaluation steps simply read each input-bit once and record it in one of two possible states, and the complexity of f is pushed to F. However, in this work we impose additional restrictions about F – we require that it will be a read-once monotone CNF formula, i.e.

$$F(\mathbf{v}_z) = \bigwedge_{j \in [k]} \left(\bigvee_{i \in S_j} \mathbf{v}_z[i] \right) \tag{1}$$

where $\{S_j\}_{j\in[k]}$ are disjoint subsets of $[w]$. We also require that the evaluation matrices will be *almost* injective, where by that we mean that each evaluation matrix induces a map $[w] \to [w] \cup \{\perp\}$ where $\perp$ represents a “dead end” sink state that can have multiple pre-images, but each of the other values in $[w]$ has at most a single pre-image. Formally this translates to the requirement that every row vector and every column vector of every evaluation matrix has at most a singe entry that equals 1. We note that a Barrington's BP satisfies these structural properties and show how to compile it into a BP that is resilient to corrupted inputs.

Our compiler adds 3ℓ nodes to the width of the BP, where each triplet works as a memory cell for an input index $i \in [\ell]$. The nodes of the ith triplet represent the following possible states:

1. x_i *has not been read yet.*
2. x_i *was* 0 *every time it was read.*
3. x_i *was* 1 *every time it was read.*

The evaluation matrices can be defined naturally to maintain this information with respect to each index $i \in [\ell]$ after each evaluation step $j \in [t]$ by using a diagonal concatenation of sub-matrices that work on each index individually. In every step where the i'th bit should be read, if the received bit is inconsistent with the current state of the memory cell that corresponds to i, the BP moves to the sink state. The vector $\mathbf{v}_{start}$ is defined such that each memory cell starts at the state *"x_i has not been read yet"*. The state-to-output map F consists of $k = \ell + 1$ clauses that check that each of the ℓ memory cells is in either of the three states mentioned above (i.e. that neither of the memory cells is in the sink state), and that the underlying Barrington's BP is in its accepting state.

Secret Sharing for F Consider a simple perfect secret-sharing according to F as follows. Let F be as in Eq (1) with respect to the subsets $\{S_j\}_{j\in[k]}$. Sample k values $\{r_j\}_{j\in[k]}$ such that each subset of size $k-1$ is uniformly random, but the sum of all of them is 0. Let each index $i \in [w]$ hold the share r_j iff $i \in S_j$. Note that since $\{S_j\}$ are disjoint each index $i \in [w]$ holds at most a single share. Each index $i \in [w]$ that does not appear in any of the sets S_j receives an independent uniformly random share $r \xleftarrow{\$} \{0,1\}$.

The properties of this secret sharing procedure that we would like to highlight are as follows. For every sequence $z \in \{0,1\}^t$ consider the subset of active indices of $\mathbf{v}_z$, i.e. $S_z = \{i \,:\, \mathbf{v}_z[i] = 1\}_{i\in[w]}$. Then according to the way we defined $\mathbf{v}_{start}$ and the evaluation matrices of the BP, it holds that:

- For $j = 1, \ldots, \ell$, if z is consistent with respect to the jth bit of the input then $|S_j \cap S_z| = 1$, otherwise $|S_j \cap S_z| = 0$.
- For $j = \ell + 1$, if the underlying Barrington's BP ends in the accepting state when it is executed with the sequence z then $|S_j \cap S_z| = 1$, otherwise $|S_j \cap S_z| = 0$.

It follows that:

- If $F(\mathbf{v}_z) = 1$ then $|S_j \cap S_z| = 1$ for all $j \in [\ell + 1]$ and therefore the sum of the shares that correspond to indices in S_z is 0.
- I $F(\mathbf{v}_z) = 0$ then there is at least one set S_{j^*} for which $|S_{j^*} \cap S_z| = 0$. This implies that all of the shares that correspond to the indices in S_z are uniformly random when observed together (since there are at most $k-1$ or them). Moreover, each of the other shares that are not in S_z can be simulated as a sum of the shares that are in S_z and possibly additional random samples.

2.3 The Witness Encryption Construction

To encrypt a message $\mu \in \{0,1\}$ with respect to an **NP** relation we first generate a branching program as was described above for the verification algorithm of the relation. The encryption then proceeds as follows:

1. Sample w matrices according to the secret-sharing algorithm with respect to F that was described above. Denote these matrices by $\{\mathbf{A}_{t+1,i}\}_{i\in[w]}$. We sometimes refer to them as *the end layer*.

2. For every node of the BP, i.e. for every $i \in [w]$ and $j \in [t]$, sample a matrix $\mathbf{A}_{j,i}$ with a trapdoor.
3. For every step $j \in [t]$ sample a pair of short matrices $\mathbf{S}_{j,0}$, $\mathbf{S}_{j,1}$.
4. For every edge of the BP $(j, i) \to (j + 1, i')$ with respect to an input bit $b \in \{0, 1\}$, if i' is not the sink state then use the trapdoor of $\mathbf{A}_{j,i}$ to sample
$$\mathbf{K} \leftarrow \mathbf{A}_{j,i}^{\mathsf{TD}}(\mathbf{S}_{j,b}\mathbf{A}_{j+1,i'}) \ .$$
5. Sample a secret vector $\mathbf{s}$. For each $i \in [w]$ for which $\mathbf{v}_{start}[i] = 1$:
   - If $\mu = 0$ then sample ct_i as a uniformly random vector.
   - If $\mu = 1$ then compute $\mathsf{ct}_i = \mathbf{s}\mathbf{A}_{1,i} + \mathbf{e}_i$ for some error vector $\mathbf{e}_i$.

   Output all of $\{\mathsf{ct}_i\}_i$ along with the trapdoor-samples from Step 4.

Note that in Step 4 the targets of the trapdoor samples do not have an error term. We *will* use an LWE assumption with respect to the secrets $\{\mathbf{S}_{j,b}\}_{j\in[t],b\in\{0,1\}}$ in intermediate hybrids during the security analysis by adding error terms $\mathbf{E}$. Intuitively, the trapdoor-targets are only accessible to the adversary as part of LWE expressions that include the error vectors $\{\mathbf{e}_i\}_i$. In the security proof we will add to the trapdoor-targets error terms $\mathbf{E}$ that are swallowed by $\{\mathbf{e}_i\}_i$ and so they do not noticeably affect the view of the adversary. We note that adding error terms $\mathbf{E}$ in the real construction would not detract from the correctness nor the security of the scheme (up to some polynomial changes in the parameters), but since it is not required we avoid it for simplicity.

Decryption. In order to decrypt a ciphertext using a valid witness $x \in \{0, 1\}^\ell$, apply trapdoor samples to the vectors $\{\mathsf{ct}_i\}$ according to the steps that the BP takes when it is evaluated with the input x. This should result in LWE expressions with respect to the matrices
$$\{\mathbf{S}_x\mathbf{A}_{t+1,i} \ : \ \mathbf{v}_x[i] = 1\}_{i\in[w]} \tag{2}$$
where $\mathbf{S}_x = \prod_{j=1}^{t} \mathbf{S}_{j,x_{\rho(j)}}$. Since x is a valid witness then $F(\mathbf{v}_x) = 1$ and then according to the secret sharing algorithm it should hold that
$$\sum_{i:\mathbf{v}_x[i]=1} \mathbf{A}_{t+1,i} = \mathbf{0} \ .$$
Therefore, if the sum of the LWE expressions with respect to the matrices in (2) is close to zero then output $\mu = 1$ and otherwise output $\mu = 0$.

2.4 Analysis Model

We now discuss the model that will be used in the security analysis.

Standard Decisional LWE. A standard (decisional) LWE experiment considers a distinguisher D that needs to distinguish whether it interacts with a random oracle or with an LWE oracle $\mathcal{O}^{LWE}$, where $\mathcal{O}^{LWE}$ has some secret vector $\mathbf{s}$ and error distribution χ hard-wired in it, and upon receiving a request from D the oracle samples a uniformly random matrix $\mathbf{B}$ along with an error term $\mathbf{e} \leftarrow \chi$ and replies with $(\mathbf{B}, \mathbf{sB} + \mathbf{e})$. Let us denote by $\mathsf{Adv}_{\mathsf{D}}^{LWE}$ the advantage of D in the standard LWE experiment. The assumption that LWE is hard then can be stated as

$$\mathsf{Adv}_{\mathsf{D}}^{LWE} \leq \text{negl}(\lambda) \textit{ for all } \text{PPT} \textit{ distinguishers } \mathsf{D}.$$

We now generalize this model so that we can make arguments about the hardness of LWE with respect to matrices $\mathbf{B}$ that are not necessarily uniform, where D might be exposed to some auxiliary information. We let aux be correlated to $\mathbf{B}$, e.g. it can be a trapdoor-sample with respect to it. We will also need to capture scenarios where D gets to *choose* matrices from an indexed domain $\{\mathbf{B}_i\}_i$.

General Public Matrices and Auxiliary Information. For any pair of (possibly correlated) distributions $\mathcal{B}$ and aux consider a distinguisher D that is provided with aux as input and needs to distinguish whether it interacts with a random oracle or with an oracle $\mathcal{O}^{LWE(\mathcal{B})}$, where $\mathcal{O}^{LWE(\mathcal{B})}$ has some secret vector $\mathbf{s}$ and error distribution χ hard-wired in it, and upon receiving a request from D the oracle samples a matrix $\mathbf{B} \leftarrow \mathcal{B}$ along with an error term $\mathbf{e} \leftarrow \chi$ and replies with $(\mathbf{B}, \mathbf{sB} + \mathbf{e})$. We denote by $\mathsf{Adv}_{\mathsf{D}(\mathsf{aux})}^{LWE(\mathcal{B})}$ the advantage of D in an experiment of this form.

This notation allows us to make relative statements about the hardness of LWE with respect to different distributions $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}', \mathsf{aux}')$ before stating anything about the hardness of either of them individually. We can argue that LWE with respect to $(\mathcal{B}, \mathsf{aux})$ is at least as hard as LWE with respect to $(\mathcal{B}', \mathsf{aux}')$ as follows:

$$\textit{For any } \text{PPT} \textit{ distinguisher } \mathsf{D} \textit{ there exists a } \text{PPT} \textit{ distinguisher } \mathsf{D}' \textit{ such that} \\ \mathsf{Adv}_{\mathsf{D}(\mathsf{aux})}^{LWE(\mathcal{B})} \leq \mathsf{Adv}_{\mathsf{D}'(\mathsf{aux}')}^{LWE(\mathcal{B}')} + \text{negl}(\lambda). \tag{3}$$

Standard proof strategies oftentimes implicitly make arguments similar to (3) for distributions $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}', \mathsf{aux}')$ that are indistinguishable. In fact it easy to show that whenever $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}', \mathsf{aux}')$ are indistinguishable then (3) indeed holds by setting $\mathsf{D}' = \mathsf{D}$, because the view of D remains negl(λ)-close in both of the experiments.

There are also scenarios where Argument (3) holds for distributions $\mathcal{B}$ and $\mathcal{B}'$ that are clearly distinguishable. As an example, for any $\mathcal{B}$ consider $\mathcal{B}' = \mathcal{B} + \mathcal{E}$ where $\mathcal{E}$ is a distribution such that for all $\mathbf{E} \leftarrow \mathcal{E}$ the term $\mathbf{sE}$ is swallowed by the error distribution χ of the oracle $\mathcal{O}^{LWE(\mathcal{B})}$. Then $\mathcal{B}$ and $\mathcal{B}'$ might be distinguishable (e.g. if $\mathcal{B}$ is a constant), but it can be shown that (3) still holds by setting $\mathsf{D}' = \mathsf{D}$ because the view of D remains statistically-close in the experiments against $\mathcal{O}^{LWE(\mathcal{B})}$ and $\mathcal{O}^{LWE(\mathcal{B}+\mathcal{E})}$.

Another common scenario is when aux is efficiently sampleable and not correlated to $\mathcal{B}$. In that case Argument (3) can be proved with respect to $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}, 1^\lambda)$ by considering a distinguisher $\mathsf{D}'(1^\lambda)$ that locally samples aux and then works the same as $\mathsf{D}(\mathsf{aux})$.

Indexed Public Matrices. Up until now we considered an oracle $\mathcal{O}^{LWE(\mathcal{B})}$ that upon request samples a fresh instance of the form $\mathbf{B} \leftarrow \mathcal{B}$. This model does not capture cases where the distinguisher gets to choose the matrices from the domain of $\mathcal{B}$ for which he sees LWE samples. We change the definition of the oracle to support scenarios of this form. Let $\mathcal{B} = \{\mathbf{B}_i\}_{i \in [p]}$ be a collection of p matrices. The random oracle $\mathcal{O}^{LWE(\mathcal{B})}$ has a secret vector $\mathbf{s}$ and error distribution χ hard-wired in it, and upon receiving a request $i \in [p]$ it samples an error term $\mathbf{e}_i \leftarrow \chi$ and replies with $\mathbf{s}\mathbf{B}_i + \mathbf{e}_i$. Repeated queries for the same i are replied with the same error term $\mathbf{e}_i$.

Notes.

1. p can be exponentially large, but any PPT distinguisher D that interacts with $\mathcal{O}^{LWE(\mathcal{B})}$ only gets to see samples with respect to a polynomial-sized subset of $\mathcal{B}$ that it can choose adaptively upon seeing previous replies.
2. aux must be of polynomial size because it is given explicitly to D. It is also reasonable to define a model where aux can be of exponential size, such that it is hard-wired to the LWE oracle and parts of it are provided to D upon request via auxiliary queries. In this work we use the version that was discussed above because it suffices for our needs.
3. The model can capture scenarios where the oracle should sample a fresh matrix $\mathbf{B} \leftarrow \mathcal{B}'$ for each query by considering a collection $\mathcal{B} = \{\mathbf{B}_i\}_{i \in [p]}$ that consists of samples $\mathbf{B}_i \leftarrow \mathcal{B}'$ for $p \in \exp(\lambda)$.
4. The term $\mathbf{B}_i$ "in the clear" does not appear anymore in the reply of the oracle, so the model can capture cases where the matrices in $\mathcal{B}$ are not necessarily known to the attacker. Additional information about $\mathcal{B}$ can be revealed in aux if required.

Example Consider $2t$ matrices $\{\mathbf{S}_{j,b}\}_{j \in [t], b \in \{0,1\}}$ from a Gaussian distribution and w uniformly random matrices $\{\mathbf{A}_i\}_{i \in [w]}$. Define the collection

$$\mathcal{B} = \left\{ \mathcal{B}_z = \left\{ \mathbf{B}_{z,i} = \prod_{j=1}^{t} \mathbf{S}_{j,z_j} \mathbf{A}_i \right\}_{i \in [w]} \right\}_{z \in \{0,1\}^t} .$$

It can be shown via a sequence of $O(t)$ hybrids that LWE with respect to $\mathcal{B}$ is hard. Consider the sequence of hybrids $\mathcal{B}^t, (\mathcal{B}^{t-1}_{temp}, \mathcal{B}^{t-1}), \ldots, (\mathcal{B}^0_{temp}, \mathcal{B}^0)$ where in each hybrid we change the way that matrices in $\mathcal{B}$ are sampled. Let $\mathcal{B}^t = \mathcal{B}$ as above and make the following inductive assumption about $\mathcal{B}^j$:

There is a collection of uniformly random matrices

$$\mathcal{A}^j = \left\{ \mathcal{A}^j_{\hat{z}} = \{\mathbf{A}_{\hat{z},i}\}_{i\in[w]} \right\}_{\hat{z}\in\{0,1\}^{t-j}}$$

such that each $\mathbf{B}^j_{z,i} \in \mathcal{B}^j$ is of the form

$$\mathbf{B}^j_{z,i} = \prod_{j'=1}^{j} \mathbf{S}_{j',z_{j'}} \mathbf{A}_{\hat{z},i}$$

where $\hat{z}$ is the length-$(t-j)$ suffix of z.

Note that the assumption holds for $j = t$ by definition for the set $\mathcal{A}^t = \{\mathbf{A}_i\}_{i\in[w]}$. For $j = t-1, \ldots, 0$, under the inductive assumption with respect to $\mathcal{B}^{j+1}$, define hybrid $\mathcal{B}^j_{temp}$ as follows:

Hybrid $\mathcal{B}^j_{temp}$. Consider the set $\mathcal{A}^{j+1}$ that satisfies the assumption for $\mathcal{B}^{j+1}$. For each $\mathbf{A}_{\hat{z},i} \in \mathcal{A}^{j+1}$ sample a pair of Gaussian matrices $\mathbf{E}^0_{\hat{z},i}, \mathbf{E}^1_{\hat{z},i}$ and define

$$\mathbf{M}^{0}_{\hat{z},i} := \mathbf{S}_{j,0}\mathbf{A}_{\hat{z},i} + \mathbf{E}^{0}_{\hat{z},i}, \qquad \mathbf{M}^{1}_{\hat{z},i} := \mathbf{S}_{j,1}\mathbf{A}_{\hat{z},i} + \mathbf{E}^{1}_{\hat{z},i} \ .$$

Define each $\mathbf{B}^j_{z,i} \in \mathcal{B}^j_{temp}$ as

$$\mathbf{B}^j_{z,i} = \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{M}^{z[j]}_{\hat{z},i}$$

where $z[j]$ is the jth bit of z and $\hat{z}$ is the length-$(t-j-1)$ suffix of z. Note that

$$\begin{aligned}
\mathbf{B}^j_{z,i} &= \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{M}^{z[j]}_{\hat{z},i} \\
&= \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \left(\mathbf{S}_{j,z[j]} \mathbf{A}_{\hat{z},i} + \mathbf{E}^{z[j]}_{\hat{z},i} \right) \\
&= \underbrace{\prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{S}_{j,z[j]} \mathbf{A}_{\hat{z},i}}_{\text{same as } \mathbf{B}^{j+1}_{z,i} \text{ in } \mathcal{B}^{j+1}} + \underbrace{\prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{E}^{z[j]}_{\hat{z},i}}_{\mathbf{E}'} \ .
\end{aligned}$$

If we define the error distribution χ of the LWE oracle such that it swallows $\mathbf{E}'$ then $\mathbf{B}^j_{z,i}$ and $\mathbf{B}^{j+1}_{z,i}$ are statistically close and therefore LWE with respect to $\mathcal{B}^{j+1}$ and $\mathcal{B}^j_{temp}$ is equivalently hard.

Hybrid $\mathcal{B}^j$. The same as $\mathcal{B}^j_{temp}$, except that each matrix $\mathbf{M}^b_{\hat{z},i}$ is sampled uniformly at random. Under the standard LWE assumption, the distributions $\mathcal{B}^j_{temp}$ and $\mathcal{B}^j$ are indistinguishable and therefore LWE with respect to $\mathcal{B}^j_{temp}$ and $\mathcal{B}^j$ is equivalently hard. Moreover, the inductive assumption holds for $\mathcal{B}^j$ with respect to the set

$$\mathcal{A}^j = \left\{ \mathbf{A}_{(b,\hat{z}),i} = \mathbf{M}^b_{\hat{z},i} \right\}_{b \in \{0,1\}, \hat{z} \in \{0,1\}^{t-j-1}, i \in [w]} .$$

Lastly, from the inductive assumption Hybrid $\mathcal{B}^0$ is the uniform distribution and therefore LWE with respect to $\mathcal{B}^0$ is identical to the standard LWE experiment.

2.5 Security of the Construction

Semi-Honest Security Assume that all that the adversary can do is to evaluate the encoded BP with arbitrary sequences $z \in \{0,1\}$ of his choice. This would result in LWE challenges with respect to matrices of the form

$$\mathcal{B} = \left\{ \mathcal{B}_z = \left\{ \mathbf{B}_{z,i} = \prod_{j=1}^{t} \mathbf{S}_{j,z_j} \mathbf{A}_{t+1,i} \ : \ \mathbf{v}_z[i] = 1 \right\}_{i \in [w]} \right\}_{z \in \{0,1\}^t} . \qquad (4)$$

(For simplicity of the overview we ignore here the accessible matrices in intermediate steps). In this part of the proof we show that an LWE experiment against the oracle $\mathcal{O}^{LWE(\mathcal{B})}$ is hard. Note the similarities between Equation (4) and the example in Section 2.4. If all of the output matrices $\{\mathbf{A}_{t+1,i}\}_{i \in [w]}$ were uniformly random we could apply the same proof. However, here the matrices $\{\mathbf{A}_{t+1,i}\}_{i \in [w]}$ are sampled according to the secret sharing of F and therefore some subsets of these matrices are correlated. We will use the relaxed guarantee that for all sequences $z \in \{0,1\}$ it holds that $F(\mathbf{v}_z) = 0$ and therefore the output matrices $\{\mathbf{A}_{t+1,i} \ : \ \mathbf{v}_z[i] = 1\}_{i \in [w]}$ are not correlated to each other.

The proof goes via a sequence of 2^t steps iterating over $z^* \in \{0,1\}^t$, where in each step the matrices in $\mathcal{B}_{z^*}$ are replaced with independent uniformly random matrices. Each such step takes $2t$ sub-hybrids that resemble the proof in Section 2.4. It begins at the end layer of the BP by sampling $\{\mathbf{A}_{t+1,i} \ : \ \mathbf{v}_z[i]{=}1\}_{i \in [w]}$ uniformly at random and simulating $\{\mathbf{A}_{t+1,i} \ : \ \mathbf{v}_z[i]{\neq}1\}_{i \in [w]}$ according to the secret-sharing simulation properties that were discussed in Section 2.2. It then moves in t hybrids towards the start layer, in each step replacing with uniform the matrices that correspond to the length-j suffix of z^*, until all of the matrices in $\mathcal{B}_{z^*}$ are sampled uniformly at random. Lastly, it moves again in the forward direction to undo changes that affect matrices in $\mathcal{B} \backslash \mathcal{B}_{z^*}$.

In more detail, consider the sequence $\mathcal{B}^{0^t}, \ldots, \mathcal{B}^{1^t}$ of length 2^t, where for $z^* \in \{0,1\}^t$ in hybrid $\mathcal{B}^{z^*} = \{\mathcal{B}^{z^*}_z\}_{z \in \{0,1\}^t}$ the matrices in $\{\mathcal{B}^{z^*}_z\}_{z < z^*}$ are sampled uniformly at random and the matrices in $\{\mathcal{B}^{z^*}_z\}_{z \geq z^*}$ are sampled as in Equation (4). In particular $\mathcal{B}^{0^t}$ is identical to the distribution in (4) and $\mathcal{B}^{1^t}$ is identical to the uniform distribution. For every $z^* \in \{0,1\}^t$ we show that LWE with respect to $\mathcal{B}^{z^*}$ and $\mathcal{B}^{z^*+1}$ is (almost) identically hard. The proof goes via a sequence of hybrids

$$\mathcal{B}^{z^*} = \mathcal{B}^t, (\mathcal{B}^{t-1}_{temp}, \mathcal{B}^{t-1}), \ldots, (\mathcal{B}^0_{temp}, \mathcal{B}^0),\ \ (\widetilde{\mathcal{B}}^0, \widetilde{\mathcal{B}}^0_{temp}), \ldots, (\widetilde{\mathcal{B}}^{t-1}, \widetilde{\mathcal{B}}^{t-1}_{temp}), \widetilde{\mathcal{B}}^t = \mathcal{B}^{z^*+1}.$$

Make the following inductive assumption about $\mathcal{B}^j$:

There is a collection of uniformly random matrices

$$\mathcal{A}^j = \left\{ \mathcal{A}^j_{\hat{z}} = \{\mathbf{A}_{\hat{z},i}\ :\ \mathbf{v}_{z^*}[i] = 1\}_{i\in[w]} \right\}_{\hat{z}\in\{0,1\}^{t-j}}$$

such that each $\mathbf{B}^j_{z,i} \in \mathcal{B}^j$ (for $z \geq z^*$) is of the form

$$\mathbf{B}^j_{z,i} = \prod_{j'=1}^{j} \mathbf{S}_{j',z_{j'}} \mathbf{A}_{\hat{z},i}$$

where $\hat{z}$ is the length-$(t-j)$ suffix of z, and $\mathbf{A}_{\hat{z},i}$ is either a matrix in $\mathcal{A}^j_{\hat{z}}$ (whenever $\mathbf{v}_{z^*}[i] = 1$), and otherwise it is a sum of up to w matrices in $\mathcal{A}^j_{\hat{z}}$.

Note that the assumption holds for $j = t$ for the set $\mathcal{A}^t = \{\mathbf{A}_{t+1,i}\ :\ \mathbf{v}_{z^*}[i] = 1\}_{i\in[w]}$ according to the secret sharing properties in Section 2.2 since $F(\mathbf{v}_z) = 0$. For $j = t-1, \ldots, 0$, under the inductive assumption with respect to $\mathcal{B}^{j+1}$, define hybrid $\mathcal{B}^j_{temp}$ as follows:

Hybrid $\mathcal{B}^j_{temp}$. Consider the set $\mathcal{A}^{j+1}$ that satisfies the assumption for $\mathcal{B}^{j+1}$. For each $\mathbf{A}_{\hat{z},i} \in \mathcal{A}^{j+1}$ sample a pair of Gaussian matrices $\mathbf{E}^0_{\hat{z},i}, \mathbf{E}^1_{\hat{z},i}$ and define

$$\mathbf{M}^{0}_{\hat{z},i} := \mathbf{S}_{j,0}\mathbf{A}_{\hat{z},i} + \mathbf{E}^{0}_{\hat{z},i}, \qquad \mathbf{M}^{1}_{\hat{z},i} := \mathbf{S}_{j,1}\mathbf{A}_{\hat{z},i} + \mathbf{E}^{1}_{\hat{z},i}\ . \tag{5}$$

For each $i \in [w]$ for which $\mathbf{v}_{z^*}[i] \neq 1$ recall that $\mathbf{A}_{\hat{z},i}$ is a sum of matrices in $\mathcal{A}^j_{\hat{z}}$ and define for $b \in \{0,1\}$ the matrix $\mathbf{M}^b_{\hat{z},i}$ as a sum of the corresponding matrices in $\{\mathbf{M}^b_{\hat{z},i'}\}_{\mathbf{v}_{z^*}[i']=1}$.

Define each $\mathbf{B}^j_{z,i} \in \mathcal{B}^j_{temp}$ as

$$\mathbf{B}^j_{z,i} = \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{M}^{z[j]}_{\hat{z},i}$$

where $z[j]$ is the jth bit of z and $\hat{z}$ is the length-$(t-j-1)$ suffix of z. Note that

$$\begin{aligned}\mathbf{B}^j_{z,i} &= \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{M}^{z[j]}_{\hat{z},i} \\ &= \prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \left(\mathbf{S}_{j,z[j]} \mathbf{A}_{\hat{z},i} + \mathbf{E}^{z[j]}_{\hat{z},i}\right) \\ &= \underbrace{\prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{S}_{j,z[j]} \mathbf{A}_{\hat{z},i}}_{\text{same as } \mathbf{B}^{j+1}_{z,i} \text{ in } \mathcal{B}^{j+1}} + \underbrace{\prod_{j'=1}^{j-1} \mathbf{S}_{j',z_{j'}} \mathbf{E}^{z[j]}_{\hat{z},i}}_{\mathbf{E}'} \ .\end{aligned}$$

If we define the error distribution χ of the LWE oracle such that it swallows $\mathbf{E}'$ then $\mathbf{B}^j_{z,i}$ and $\mathbf{B}^{j+1}_{z,i}$ are statistically close and therefore LWE with respect to $\mathcal{B}^{j+1}$ and $\mathcal{B}^j_{temp}$ is equivalently hard.

Hybrid $\mathcal{B}^j$. The same as $\mathcal{B}^j_{temp}$, except that each matrix $\mathbf{M}^b_{\hat{z},i}$ is sampled uniformly at random. Under the standard LWE assumption, the distributions $\mathcal{B}^j_{temp}$ and $\mathcal{B}^j$ are indistinguishable and therefore LWE with respect to $\mathcal{B}^j_{temp}$ and $\mathcal{B}^j$ is equivalently hard. Moreover, the inductive assumption holds for $\mathcal{B}^j$ with respect to the set

$$\mathcal{A}^j = \left\{\mathbf{A}_{(b,\hat{z}),i} = \mathbf{M}^b_{\hat{z},i} \ : \ \mathbf{v}_z[i] = 1\right\}_{b\in\{0,1\},\hat{z}\in\{0,1\}^{t-j-1},i\in[w]} \ .$$

From the inductive assumption in Hybrid $\mathcal{B}^0$ all of the matrices in $\mathcal{B}_{z^*}$ are sampled from the uniform distribution. However, matrices in $\mathcal{B}_z$ for $z > z^*$ are still simulated as sums of intermediate matrices that were sampled in previous hybrids according to z^*. In hybrids $(\widetilde{\mathcal{B}}^0, \widetilde{\mathcal{B}}^0_{temp}), \ldots, (\widetilde{\mathcal{B}}^{t-1}, \widetilde{\mathcal{B}}^{t-1}_{temp}), \widetilde{\mathcal{B}}^t$ we undo these changes and end up with the distribution $\widetilde{\mathcal{B}}^t$ that is identical to $\mathcal{B}^{z^*+1}$.

Full Security The security game of the construction is identical to an LWE experiment where D receives as auxiliary information all of the trapdoor samples that were generated in Step 4 of the construction, and is playing against an oracle $\mathcal{O}^{LWE(\mathcal{B}_0)}$ where

$$\mathcal{B}_0 = \{\mathbf{A}_{1,i} \ : \ \mathbf{v}_{start}[i] = 1\}_{i\in[w]} \ .$$

For all $j = 1, \ldots, t$ we let aux_j denote the trapdoor samples that go from the jth layer to the $(j+1)$th layer of the BP, and for all $j = 0, \ldots, t$ we let $\mathcal{B}_j$ denote the matrices that can be accessed via any length-j evaluation sequence:

$$\mathcal{B}_j = \left\{\prod_{j'=1}^{j} \mathbf{S}_{j',z'_j} \mathbf{A}_{j+1,i} \ : \ \mathbf{v}_z[i] = 1\right\}_{z\in\{0,1\}^j,i\in[w]} \ .$$

Using this notation, we prove that for any PPT D there exists a PPT D' such that:

$$\underbrace{\mathsf{Adv}^{LWE(\mathcal{B}_0)}_{\mathsf{D}(\mathsf{aux}_1,\ldots,\mathsf{aux}_t)}}_{\text{real security game}} \leq \underbrace{\mathsf{Adv}^{LWE(\mathcal{B}_0,\cdots,\mathcal{B}_t)}_{\mathsf{D}'(1^\lambda)}}_{\text{semi-honest security game}} + \mathrm{negl}(\lambda) \ . \tag{6}$$

Note that $\mathcal{B}_t$ is identical to the distribution $\mathcal{B}$ that was discussed in the previous part of the overview (Section 2.4), i.e. we already showed that LWE with respect to $\mathcal{B}_t$ and $\mathsf{aux} = 1^\lambda$ is hard. In the body of the paper we show that LWE with respect to $(\mathcal{B}_0, \ldots, \mathcal{B}_t)$ and $\mathsf{aux} = 1^\lambda$ is hard similarly to the proof overview that was provided here for $\mathcal{B}_t$.

The proof of (6) takes t steps, where for $j = 1, \ldots, t$ we show that

$$\mathsf{Adv}^{LWE(\mathcal{B}_0,\ldots,\mathcal{B}_{j-1})}_{\mathsf{D}(\mathsf{aux}_j,\ldots,\mathsf{aux}_t)} \leq \mathsf{Adv}^{LWE(\mathcal{B}_0,\cdots,\mathcal{B}_j)}_{\mathsf{D}'(\mathsf{aux}_{j+1},\ldots,\mathsf{aux}_t)} + \mathrm{negl}(\lambda) \tag{7}$$

via $O(w)$ reductions to our assumption.

2.6 Phrasing the Assumption

The essence of the assumption is that if there exists a successful attack on LWE with respect to $\mathbf{A}$ given a trapdoor sample $\mathbf{K} \overset{\$}{\leftarrow} \mathbf{A}^{\mathsf{TD}}(\mathbf{T})$, then there also exists a successful attack on LWE with respect to $[\mathbf{A}|\mathbf{T}]$ without $\mathbf{K}$. We make a few generalizations to this base case. First, we consider a case where the challenge is with respect to $\mathbf{CA}$ for an arbitrary matrix $\mathbf{C}$ and not necessarily $\mathbf{C} = \mathbf{I}$. More generally, we consider a challenge with respect to multiple matrices $\{\mathbf{C}_j\mathbf{A}\}_{j\in[p]}$ and a trapdoor sample $\mathbf{K} \overset{\$}{\leftarrow} \mathbf{A}^{\mathsf{TD}}(\mathbf{T})$, and assume that is not easier than a challenge with respect to $\{[\mathbf{C}_j\mathbf{A} \mid \mathbf{C}_j\mathbf{T}]\}_{j\in[p]}$ without $\mathbf{K}$. Lastly we allow additional auxiliary information and public matrices that are available to the adversary. This sums up to the following:

Main Assumption (Informal). Consider

$$\left(\mathbf{A}, \mathbf{A}^{\mathsf{TD}}\right) \leftarrow \mathsf{TrapGen}(1^n, q, m) \ .$$

Let $\mathbf{T}$ be a $n \times m'$ *target matrix*, let $\mathcal{C}$ be a collection of $n \times n$ *prefix matrices*, let aux be a poly-sized auxiliary information and let $\mathcal{B}$ be a collection of $n \times m$ matrices (where $\mathcal{B}$ and $\mathcal{C}$ are possibly of exponential size), where $\mathbf{T}, \mathcal{C}, \mathsf{aux}, \mathcal{B}$ can be correlated to each other and to $\mathbf{A}, \mathbf{A}^{\mathsf{TD}}$. Consider

$$\mathbf{K} \leftarrow \mathbf{A}^{\mathsf{TD}}(\mathbf{T}) \ .$$

Then for any PPT distinguisher D there exists a PPT distinguisher D' such that

$$\mathsf{Adv}^{LWE(\mathcal{B},\{\mathbf{CA}\}_{\mathbf{C}\in\mathcal{C}})}_{\mathsf{D}(\mathsf{aux},\mathbf{K})} \leq \mathsf{Adv}^{LWE(\mathcal{B},\{\mathbf{CA}\}_{\mathbf{C}\in\mathcal{C}},\{\mathbf{CT}\}_{\mathbf{C}\in\mathcal{C}})}_{\mathsf{D}'(\mathsf{aux})} + \mathrm{negl}(\lambda) \ .$$

Note that all of the distributions are sampled before $\mathbf{K}$ and in particular aux does not include a copy of $\mathbf{K}$ or the randomness that was used during its sampling. aux might contain another trapdoor-sample with respect to the same target $\mathbf{T}$ (or even contain $\mathbf{A}^{\mathsf{TD}}$), but in that case the assumption holds trivially because there is always a PPT D' with large advantage in the latter experiment.

Proof of (7) via the Main Assumption Equation (7) considers all of the encodings that go from the j^*th layer to the (j^*+1)th layer of the branching program. These are all of the encodings that are sampled with trapdoors of the nodes $\{\mathbf{A}_{j^*,i}\}_{i\in[w]}$, where each node contributes at most two encodings (one for the 0 transition and one for the 1 transition). We replace each of these encodings via a reduction to the main assumption. Let $\mathbf{K}$ be an encoding from $\mathbf{A}_{j^*,i^*}$ to $\mathbf{S}_{j^*,b}\mathbf{A}_{j^*+1,i'}$, then apply the reduction with respect to the following distributions:

- The target matrix $\mathbf{T}$ is $\mathbf{S}_{j^*,b}\mathbf{A}_{j^*+1,i'}$.
- The prefix matrices $\mathcal{C}$ are all of the possible length-(j^*-1) evaluation sequences that lead to the node $\mathbf{A}_{j^*,i^*}$:

$$\mathcal{C} = \left\{ \prod_{j=1}^{j^*} \mathbf{S}_{j,z_j} \ : \ \mathbf{v}_z[i^*] = 1 \right\}_{z\in\{0,1\}^{j^*-1}} .$$

- The auxiliary information aux consists of all of the other trapdoor samples that are available to D, i.e. $(\mathsf{aux}_{j^*+1},\ldots,\mathsf{aux}_t)$ and the encodings in aux_{j^*} that were not replaced yet.
- The collection $\mathcal{B}$ consists of all of the matrices for which D can see an LWE sample that do not correspond to the node $\mathbf{A}_{j^*,i^*}$, i.e. $(\mathcal{B},\ldots,\mathcal{B}_{j^*-2})$, the matrices in $\mathcal{B}_{j^*-1}$ that do not correspond to $\mathbf{A}_{j^*,i^*}$ and the matrices in $\mathcal{B}_{j^*}$ that were already added in the previous hybrids.

This sums up the technical overview of the construction and its security analysis.

2.7 Paper Structure

In Appendix A we discuss a simplified LWE scenario that resembles the security properties of our WE candidate. Due to page limit constraints we moved the information-theoretic part of the work to Appendix B. In Section 3 we present the generalized LWE model and our assumption. In Section 4 we present the witness-encryption candidate and in Appendix C we analyze its security. Appendices D and E contain supplementary materials for Sections 3 and 4 respectively.

3 Lattice Tools

Some lattice preliminaries can be found in Appendix D.1.

3.1 LWE for General Matrices and Auxiliary Information

We now present the notation that will be used throughout this text to analyze the security of LWE with auxiliary information and public matrices from arbitrary distributions.

Definition 1 (Generalized Decisional LWE).
Let $\lambda \in \mathbb{N}$ be the security parameter, let $n, q, m \in \text{poly}(\lambda)$ be integers and let χ_e, χ_s be probability distributions over $\mathbb{Z}_q$. For every (possibly exponential) integer $p \in \mathbb{Z}$ and every distribution $\mathcal{B} = \{\mathbf{B}_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[p]}$ define an LWE oracle $\mathcal{O}^{\text{LWE}(\mathcal{B})} : [p] \to \mathbb{Z}_q^m$ as follows:

1. *$\mathcal{O}^{\text{LWE}(\mathcal{B})}$ samples a secret $\mathbf{s} \leftarrow \chi_s^n$.*
2. *$\mathcal{O}^{\text{LWE}(\mathcal{B})}$ associates each $i \in [p]$ with an error vector $\mathbf{e}_i \leftarrow \chi_e^m$.*
3. *Upon receiving a query $i \in [p]$, $\mathcal{O}^{\text{LWE}(\mathcal{B})}$ returns*

$$\mathbf{u}_i := \mathbf{s}\mathbf{B}_i + \mathbf{e}_i \ .$$

In addition consider the random oracle $\mathcal{O}^{\mathsf{rand}(\mathcal{B})} : [p] \to \mathbb{Z}_q^m$ that associates with every $i \in [p]$ a random vector $\mathbf{u}_i \overset{\$}{\leftarrow} \mathbb{Z}_q^m$.

For any distribution $\mathsf{aux} \in \{0,1\}^{\text{poly}(\lambda)}$ we let $\text{LWE}_{n,m,q,\chi_s,\chi_e}[\mathcal{B}, \mathsf{aux}]$ denote the challenge of distinguishing between $\mathcal{O}^{\text{LWE}(\mathcal{B})}$ and $\mathcal{O}^{\mathsf{rand}(\mathcal{B})}$, given the auxiliary information aux.

For every distinguisher D we say that the advantage of D in the challenge $\text{LWE}_{n,m,q,\chi_s,\chi_e}[\mathcal{B}, \mathsf{aux}]$ is

$$\begin{aligned}&\mathsf{Adv}^{\mathsf{D}}_{n,m,q,\chi_s,\chi_e}[\mathcal{B}, \mathsf{aux}] \\ &\qquad := \left| Pr\left[\mathsf{D}^{\mathcal{O}^{\text{LWE}(\mathcal{B})}}(1^\lambda, \mathsf{aux}) = 1\right] - Pr\left[\mathsf{D}^{\mathcal{O}^{\mathsf{rand}(\mathcal{B})}}(1^\lambda, \mathsf{aux}) = 1\right]\right| \ .\end{aligned}$$

We sometimes omit the parameters n, m, q, χ_s, χ_e when they are clear from the context.

We also define a succinct notation for relative statements about the hardness of LWE with respect to different distributions of $\mathcal{B}, \mathsf{aux}$:

Definition 2 (Relative Hardness of LWE).
Let $\lambda \in \mathbb{N}$ be the security parameter, let $n, q, m \in \text{poly}(\lambda)$ be integers and let χ_e, χ_s be probability distributions over $\mathbb{Z}_q$. For any pair of distributions $\mathcal{B} = \{\mathbf{B}_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[p]}$, $\mathcal{B}' = \{\mathbf{B}'_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[p']}$ and

$\mathsf{aux}, \mathsf{aux}' \in \{0,1\}^{\mathrm{poly}(\lambda)}$ *and any function* ϵ*, let*

$$\mathrm{LWE}[\mathcal{B}, \mathsf{aux}] \quad \preccurlyeq \quad \mathrm{LWE}[\mathcal{B}', \mathsf{aux}'] + \epsilon(\lambda)$$

denote that for any PPT *distinguisher* D *there exists a* PPT *distinguisher* D' *such that*

$$\mathsf{Adv}^{\mathsf{D}}[\mathcal{B}, \mathsf{aux}] \quad \leq \quad \mathsf{Adv}^{\mathsf{D}'}[\mathcal{B}', \mathsf{aux}'] + \epsilon(\lambda)$$

Informally, this means that $\mathrm{LWE}[\mathcal{B}, \mathsf{aux}]$ *is at least (*$\epsilon(\lambda)$*-almost) as hard as* $\mathrm{LWE}[\mathcal{B}', \mathsf{aux}']$.

We recall the standard notion of Decisional LWE as was introduced by Regev in [Reg05], where $\mathcal{B}$ is the uniform distribution over $\mathbb{Z}_q^{n\times m}$, there is no auxiliary information and the LWE oracle sends $\mathbf{B}_i$ as a part of its response:

Definition 3 (Decisional LWE).
Let $\lambda \in \mathbb{N}$ *be the security parameter, let* $n, q, m \in \mathrm{poly}(\lambda)$ *be integers and let* χ_e, χ_s *be probability distributions over* $\mathbb{Z}_q$*. Let* $\mathcal{O}^{\mathrm{LWE}}$ *denote the oracle that works as follows:*

1. $\mathcal{O}^{\mathrm{LWE}}$ *samples a secret* $\mathbf{s} \leftarrow \chi_s^n$.
2. *Upon receiving a query,* $\mathcal{O}^{\mathrm{LWE}}$ *samples* $\mathbf{B} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $\mathbf{e} \leftarrow \chi_e^m$ *and returns* $(\mathbf{B}, \mathbf{sB} + \mathbf{e})$.

Let $\mathcal{O}^{\mathsf{rand}}$ *denote the oracle that upon receiving a query, samples* $\mathbf{B} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $\mathbf{u} \leftarrow \mathbb{Z}_q^m$ *and returns* $(\mathbf{B}, \mathbf{u})$.

We let $\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}$ *denote the challenge of distinguishing between* $\mathcal{O}^{\mathrm{LWE}}$ *and* $\mathcal{O}^{\mathsf{rand}}$*. For every distinguisher* D *we say that the advantage of* D *in the challenge* $\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}$ *is*

$$\mathsf{Adv}^{\mathsf{D}}_{n,m,q,\chi_s,\chi_e} := \left| Pr\left[\mathsf{D}^{\mathcal{O}^{\mathrm{LWE}}}(1^\lambda) = 1\right] - Pr\left[\mathsf{D}^{\mathcal{O}^{\mathsf{rand}}}(1^\lambda) = 1\right]\right| \ .$$

We sometimes omit the parameters n, m, q, χ_s, χ_e *when they are clear from the context.*

As was shown in a long sequence of works, LWE is at least as hard as the lattice problems GapSVP and SIVP under various choices of parameters. In this work we sample the LWE secret from the same distribution as the error term. [ACPS09] showed that this setting is at least as hard as when the secret is sampled uniformly form $\mathbb{Z}_q^n$.

Theorem 1 ([Reg05, Pei09, MM11, MP12, BLP+13, ACPS09]). *For all $\epsilon > 0$ there exist functions $q = q(n) \leq 2^n$, $\chi = \chi(n)$ such that χ is (B, ϵ)-bounded for some $B = B(n)$, $q/B \geq 2^{n^\epsilon}$ and such that for all $m \in \mathrm{poly}(m)$ $\mathrm{LWE}_{n,m,q,\chi,\chi}$ is at least as hard as the classical hardness of GapSVP_γ and the quantum hardness of SIVP_γ for $\gamma = 2^{\Omega(n^\epsilon)}$.*

We now consider a number of special cases for $\mathcal{B}, \mathsf{aux}$:

1. **Lemma ?? (Informal)** – If $\mathcal{B}$ consists of uniformly random matrices in $\mathbb{Z}_q^{n\times m}$ then $\mathrm{LWE}[\mathcal{B}, \mathsf{aux}]$ is at least as hard as standard decisional LWE for any efficiently samplable aux.
2. **Lemma ?? (Informal)** – For any pair of distributions $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}', \mathsf{aux}')$, if $(\mathcal{B}, \mathsf{aux})$ and $(\mathcal{B}', \mathsf{aux}')$ are indistinguishable then $\mathrm{LWE}[\mathcal{B}, \mathsf{aux}]$ and $\mathrm{LWE}[\mathcal{B}', \mathsf{aux}']$ are equivalently hard.
3. **Lemma ?? (Informal)** – If the error distribution χ_e swallows an error distribution $\mathcal{E}$, then $\mathrm{LWE}[\mathcal{B}, \mathsf{aux}]$ and $\mathrm{LWE}[\mathcal{B} + \mathcal{E}, \mathsf{aux}]$ are equivalently hard for any $(\mathcal{B}, \mathsf{aux})$.

The full Lemmas and proofs can be found in Appendix D.2.

3.2 Main Assumption

Assumption 31 *Let $\lambda \in \mathbb{N}$ be the security parameter, let $n, q, m \in \mathrm{poly}(\lambda)$ be integers and let χ_e, χ_s be probability distributions over $\mathbb{Z}_q$. Consider the following experiment:*

1. *Sample a matrix in $\mathbb{Z}_q^{n\times m}$ with a trapdoor:*
$$(\mathbf{A}, \mathbf{A}^{\mathsf{TD}}) \leftarrow \mathsf{TrapGen}(1^n, q, m) \ .$$
2. *Sample the following terms from arbitrary distributions that are possibly correlated to each other and to $\mathbf{A}, \mathbf{A}^{\mathsf{TD}}$:*
   - *Auxiliary information $\mathsf{aux} \in \{0,1\}^{\mathrm{poly}(\lambda)}$.*
   - *A target matrix $\mathbf{T} \in \mathbb{Z}_q^{n\times m}$.*
   - *A collection of prefix matrices $\mathcal{S} \subseteq \mathbb{Z}_q^{n\times n}$ (possibly of exponential size).*
   - *A collection of additional public matrices $\mathcal{B} \subseteq \mathbb{Z}_q^{n\times m}$ (possibly of exponential size).*
3. *Sample*
$$\mathbf{K} \leftarrow \mathbf{A}^{\mathsf{TD}}(\mathbf{T}) \ .$$

Then

$$\begin{aligned}&\mathrm{LWE}\left[\mathcal{B} \cup \{\mathbf{SA}\}_{\mathbf{S}\in\mathcal{S}} \ , \ \ (\mathbf{K}, \mathsf{aux})\right] \preccurlyeq \\ &\mathrm{LWE}\left[\mathcal{B} \cup \{\mathbf{SA}, \mathbf{ST}\}_{\mathbf{S}\in\mathcal{S}} \ , \ \mathsf{aux}\right] + \mathrm{negl}(\lambda) \ .\end{aligned}$$

4 Candidate Witness Encryption

The definition of witness encryption can be found in Appendix E.1.

4.1 Encodings of Matrix Branching Programs

Before proceeding to the full construction we define algorithms that create a collection of matrices and trapdoor-samples according to the nodes and edges respectively of a matrix branching program.

- $\mathsf{EncodeEdge}\left(\mathbf{A}, \mathbf{A}^{\mathsf{TD}}, \mathbf{S}, \mathbf{A}'\right)$ takes as input a source matrix $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$ along with its trapdoor $\mathbf{A}^{\mathsf{TD}}$, a secret $\mathbf{S} \in \mathbb{Z}_q^{n\times n}$ and a target matrix $\mathbf{A}' \in \mathbb{Z}_q^{n\times m}$. It computes and outputs
$$\mathbf{K} \leftarrow \mathbf{A}^{\mathsf{TD}}\left(\mathbf{S}\mathbf{A}'\right) \ .$$
We sometimes use the notations
$$\mathbf{K}[\mathsf{Source}] := \mathbf{A}, \qquad \mathbf{K}[\mathsf{Target}] := \mathbf{S}\mathbf{A}' \ .$$
- $\mathsf{EncodeMatrix}\left(\mathbf{M}, \left\{\mathbf{A}_i, \mathbf{A}_i^{\mathsf{TD}}\right\}_{i\in[w]}, \mathbf{S}, \{\mathbf{A}'_i\}_{i\in[w]}\right)$ takes as input an evaluation matrix $\mathbf{M} \in \{0,1\}^{w\times w}$, a collection of source matrices $\{\mathbf{A}_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[w]}$ along with their trapdoors $\{\mathbf{A}_i^{\mathsf{TD}} \in \mathbb{Z}_q^{n\times m}\}_{i\in[w]}$, a secret $\mathbf{S} \in \mathbb{Z}_q^{n\times n}$ and a collection of target matrices $\{\mathbf{A}'_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[w]}$. It computes and outputs
$$\mathcal{K} = \left\{\mathbf{K}^{i,i'} \leftarrow \mathsf{EncodeEdge}\left(\mathbf{A}_i, \mathbf{A}_i^{\mathsf{TD}}, \mathbf{S}, \mathbf{A}'_i\right) \ : \ \mathbf{M}[i,i'] = 1\right\}_{i\in[w],i'\in[w]}$$
where $\mathbf{M}[i,i']$ denote the entry of $\mathbf{M}$ in the i'th row and i''th column.
- $\mathsf{EncodeBP}_{\chi_S}\left(\{\mathbf{M}_{j,b}\}_{j\in[t],b\in\{0,1\}}, \{\mathbf{A}_i, \mathbf{A}_i^{\mathsf{TD}}\}_{i\in[w]}, \{\mathbf{A}'_i\}_{i\in[w]}\right)$ is defined with respect to a distribution of secrets χ_S. It takes as input a collection of evaluation matrices $\{\mathbf{M}_{j,b} \in \{0,1\}^{w\times w}\}_{j\in[t],b\in\{0,1\}}$, a collection of source matrices $\{\mathbf{A}_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[w]}$ along with their trapdoors $\{\mathbf{A}_i^{\mathsf{TD}}\}_{i\in[w]}$ and a collection of target matrices $\{\mathbf{A}'_i \in \mathbb{Z}_q^{n\times m}\}_{i\in[w]}$. It works as follows:
  1. For $i \in [w]$ denote $\left(\mathbf{A}_{i,1}, \mathbf{A}_{i,1}^{\mathsf{TD}}\right) := \left(\mathbf{A}_i, \mathbf{A}_i^{\mathsf{TD}}\right)$ and $\mathbf{A}_{i,t+1} := \mathbf{A}'_i$.
  2. For $j = 2, \ldots, t$ and $i \in [w]$ sample a matrix with a trapdoor
  $$\left(\mathbf{A}_{i,j}, \mathbf{A}_{i,j}^{\mathsf{TD}}\right) \xleftarrow{\$} \mathsf{TrapGen}(1^n, q, m) \ .$$
  3. For $j \in [t]$ and $b \in \{0,1\}$ sample a secret matrix $\mathbf{S}_{j,b} \leftarrow \chi_S^{n\times n}$ and compute
  $$\mathcal{K}_{j,b} \leftarrow \mathsf{EncodeMatrix}\left(\mathbf{M}_{j,b}, \{\mathbf{A}_{j,i}, \mathbf{A}_{j,i}^{\mathsf{TD}}\}_{i\in[w]}, \mathbf{S}_{j,b}, \{\mathbf{A}_{j+1,i}\}_{i\in[w]}\right) \ .$$
  4. Output $\{\mathcal{K}_{j,b}\}_{j\in[t],b\in\{0,1\}}$.

4.2 The Construction

Construction 4
Fix the parameters $n, m, q, \tau, B' \in \mathbb{N}$ and the distributions $\chi_e, \chi_s, \chi_S, \chi_E$ as will be described later.

- $\mathsf{Enc}(1^\lambda, f, \mu)$:
  1. *Let*
  $$\mathsf{BP} = \big(\ell, w, t, \rho, \mathbf{v}_{\mathrm{start}}, \{\mathbf{M}_{i,b}\}_{i\in[t],b\in\{0,1\}}, F\big)$$
  be a BP for f as in Construction **??**.
  2. *For $i \in [w]$ sample a matrix with a trapdoor*
  $$\big(\mathbf{A}_{i,1}, \mathbf{A}_{i,1}^{\mathsf{TD}}\big) \xleftarrow{\$} \mathsf{TrapGen}(1^n, q, m) \ .$$
  3. *Perform a secret sharing as in Theorem* **??** *for the share $\mathbf{0}^{n\times m}$ according to F:*
  $$\mathsf{Share}(F) \to (\mathbf{A}_{1,t+1}, \ldots, \mathbf{A}_{w,t+1})$$
  where $\mathbf{A}_{i,t+1} \in \mathbb{Z}_q^{n\times m}$ for $i \in [w]$.
  4. *Encode* BP *with the source matrices $\{\mathbf{A}_{i,1}\}_i$ and target matrices $\{\mathbf{A}_{i,t+1}\}_i$:*
  $$\{\mathcal{K}_{j,b}\}_{j\in[t],b\in\{0,1\}} \leftarrow \mathsf{EncodeBP}_{\chi_S}\left(\{\mathbf{M}_{j,b}\}_{j\in[t],b\in\{0,1\}}, \{\mathbf{A}_{i,1}, \mathbf{A}_{i,1}^{\mathsf{TD}}\}_{i\in[w]}, \{\mathbf{A}_{i,t+1}\}_{i\in[w]}\right) \ .$$
  5. *Let $I_{start} \subseteq [w]$ denote the indices in which the i'th bit of $\mathbf{v}_{start}$ is 1. That is,*
  $$I_{start} := \{i \ : \ \mathbf{v}_{start}[i] = 1\}_{i\in[w]}$$
  (where $\mathbf{v}_{start}[i]$ denotes the i'th bit of $\mathbf{v}_{start}$).
     - *If $\mu = 1$ then sample a secret vector $\mathbf{s} \leftarrow \chi_s^n$ and for $i \in I_{start}$ sample an error vector $\mathbf{e}_i \leftarrow \chi_e^m$, and compute*
     $$\mathsf{ct}_i := \mathbf{s}\mathbf{A}_{i,1} + \mathbf{e}_i \ .$$
     - *If $\mu = 0$ then for $i \in I_{start}$ sample*
     $$\mathsf{ct}_i \xleftarrow{\$} \mathbb{Z}_q^m \ .$$
  6. *Output* $\mathsf{ct} := \big(\{\mathsf{ct}_i\}_{i\in I_{start}}, \mathsf{BP}, \{\mathcal{K}_{j,b}\}_{j\in[t],b\in\{0,1\}}\big)$.
- $\mathsf{Dec}(\mathsf{ct}, x)$: *Parse*
$$\mathsf{ct} = \big(\{\mathsf{ct}_i\}_{i\in I_{start}}, \mathsf{BP}, \{\mathcal{K}_{j,b}\}_{j\in[t],b\in\{0,1\}}\big) \ .$$
For every $i \notin I_{start}$ denote $\mathsf{ct}_i = \mathbf{0}^m$. Consider the vector
$$[\mathsf{ct}_1|\ldots|\mathsf{ct}_w] \in \{\mathbb{Z}_q^m\}^w \ .$$
For every $j \in [t]$ and $b \in \{0,1\}$ parse
$$\mathcal{K}_{j,b} = \left\{\mathbf{K}_{j,b}^{i,i'} \ : \ \mathbf{M}_{j,b}[i,i'] = 1\right\}_{i\in[w],i'\in[w]}$$

and define the blocks matrix $\mathbf{K}_{j,b} \in \{\mathbb{Z}_q^{m\times m}\}^{w\times w}$ *where for* $i, i' \in [w]$*, the block in the* i *th row and* i' *th column of* $\mathbf{K}_{j,b}$ *is* $\mathbf{0}^{m\times m}$ *if* $\mathbf{M}_{j,b}[i,i'] = 0$ *and otherwise it is* $\mathbf{K}_{j,b}^{i,i'}$*:*

$$\mathbf{K}_{j,b} := \mathbf{M}_{j,b} \otimes \begin{bmatrix} \mathbf{K}_{j,b}^{1,1} & \dots & \mathbf{K}_{j,b}^{1,w} \\ \dots & \ddots & \dots \\ \mathbf{K}_{j,b}^{w,1} & \dots & \mathbf{K}_{j,b}^{w,w} \end{bmatrix} .$$

Compute

$$[\mathsf{ct}'_1 | \dots | \mathsf{ct}'_w] \quad := \quad [\mathsf{ct}_1 | \dots | \mathsf{ct}_w] \prod_{j=1}^{t} \mathbf{K}_{j,x_{\rho(j)}}$$

and

$$\mathbf{v}_x \quad := \quad \mathbf{v}_{start}^T \cdot \prod_{j=1}^{t} \mathbf{M}_{j,x_{\rho(j)}} .$$

Let $I_x \subseteq [w]$ *be the indices in which the i'th bit of* $\mathbf{v}_x$ *is* 1*, i.e.*

$$I_x := \{i \ : \ \mathbf{v}_x[i] = 1\}_{i\in[w]} .$$

Compute

$$\mathsf{ct}' := \sum_{i\in I_x} \mathsf{ct}'_i$$

and output 1 *iff* $\|\mathsf{ct}'\|_\infty \le B'$*.*

The correctness analysis can be found in Appendix E.2.

4.3 Security

Notations The following notations will be used throughout the proof.

- For every $j = 0, \dots, t$ and $z \in \{0,1\}^j$ let $\mathbf{v}_z \in \{0,1\}^w$ denote the state vector of BP after j steps when it is executed with the evaluation sequence z, i.e.

$$\mathbf{v}_z := \mathbf{v}_{start}{}^T \prod_{j'=1}^{j} \mathbf{M}_{j',z_{j'}} .$$

  For $i \in [w]$ let $\mathbf{v}_z[i] \in \{0,1\}$ denote the i'th bit of $\mathbf{v}_z$.

Consider the witness encryption security experiment as in Definition **??** and let $\{\mathbf{S}_{j,b}\}_{b\in\{0,1\},j\in[t]}$ and $\{\mathbf{A}_{i,j}\}_{i\in[w],j\in[t+1]}$ be the matrices that are sampled during Enc.

- For every $j = 0, \dots, t$ and $z \in \{0,1\}^j$ denote

$$\mathbf{S}_z := \mathbf{I} \cdot \prod_{j'=1}^{j} \mathbf{S}_{j',z_{j'}} .$$

- For $j \in [t+1]$ define the collection:

$$\mathcal{B}^j := \{\mathbf{S}_z \mathbf{A}_{i,j} \quad : \quad \mathbf{v}_z[i] = 1\}_{i \in [w], z \in \{0,1\}^{j-1}} \ . \tag{1}$$

  Intuitively, $\mathcal{B}^j$ consists of all of the possible states of the BP after $j-1$ evaluation steps, where $\mathbf{A}_{i,j}$ corresponds to an active BP node in the jth level and $\mathbf{S}_z$ corresponds to the evaluation sequence $z \in \{0,1\}^{j-1}$ that lead to this node.
- For $j = 1, \ldots, t$ denote $\mathsf{aux}^j := \{\mathcal{K}_{j,0}, \mathcal{K}_{j,1}\}$.

Proof Overview Consider the security experiment and let $\mathsf{ct} = \left(\{\mathsf{ct}_i\}_{i \in I_{start}}, \mathsf{BP}, \{\mathcal{K}_{j,b}\}_{j \in [t], b \in \{0,1\}}\right)$ be the challenge ciphertext with respect to a secret message $\mu \in \{0,1\}$. Recall that if $\mu = 1$ then $\{\mathsf{ct}_i\}_{i \in I_{start}}$ are of the form $\{\mathbf{s}\mathbf{A}_{i,1} + \mathbf{e}_i\}_{i \in I_{start}}$ and if $\mu = 0$ then $\{\mathsf{ct}_i\}_{i \in I_{start}}$ are uniformly random vectors in $\mathbb{Z}_q^m$. Therefore, the security experiment is identical to an LWE challenge with respect to the public matrices $\{\mathbf{A}_{i,1}\}_{i \in I_{start}}$ and auxiliary information $\left(\mathsf{BP}, \{\mathcal{K}_{j,b}\}_{j \in [t], b \in \{0,1\}}\right)$. Using the notation that was presented above, we note that

$$\mathcal{B}^1 = \{\mathbf{A}_{i,1}\}_{i \in I_{start}}, \qquad \{\mathsf{aux}^j\}_{j \in [t]} = \{\mathcal{K}_{j,b}\}_{j \in [t], b \in \{0,1\}}$$

and therefore we can say that the security experiment is identical to the LWE challenge

$$\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}\left[\mathcal{B}^1, \left(\mathsf{BP}, \{\mathsf{aux}^j\}_{j \in [t]}\right)\right] \ . \tag{2}$$

The security proof goes in two main steps. Fist, we show that under Assumption 31, the challenge in Equation (2) is at least as hard as $\mathrm{LWE}[\{\mathcal{B}^j\}_{j \in [t+1]}, \mathsf{BP}]$. In the second step we show that $\mathrm{LWE}[\{\mathcal{B}^j\}_{j \in [t+1]}, \mathsf{BP}]$ is at least as hard as standard decisional LWE by showing that $\{\mathcal{B}^j\}_{j \in [t+1]}$ are indistinguishable from uniform random matrices. The latter step requires $O(t \cdot 2^t)$ hybrids since it handles every evaluation sequence $z \in \{0,1\}^t$ individually. We now state the two main lemmas.

Lemma 1 (First Step). *Let $\{\mathcal{B}^j\}_{j \in [t+1]}, \{\mathsf{aux}^j\}_{j \in [t]}$ be as defined above, then under Assumption 31,*

$$\begin{aligned} &\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}\left[\mathcal{B}^1 \ , \ \left(\mathsf{BP}, \{\mathsf{aux}^j\}_{j \in [t]}\right)\right] \preccurlyeq \\ &\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}\left[\{\mathcal{B}^j\}_{j \in [t+1]} \ , \ \mathsf{BP}\right] + \mathrm{negl}(\lambda) \ . \end{aligned}$$

Lemma 2 (Second Step). *Let $\{\mathcal{B}^j\}_{j \in [t+1]}$ be as defined above and let $\mathcal{U}$ be a collection of uniformly random matrices in $\mathbb{Z}_q^{n \times m}$ such that $|\mathcal{U}| = \left|\{\mathcal{B}^j\}_{j \in [t+1]}\right|$, then under the standard decisional LWE assumption, if $f(x) = 0$ for all $x \in \{0,1\}^\ell$ then*

$$\begin{aligned} &\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}\left[\{\mathcal{B}^j\}_{j \in [t+1]} \ , \ \mathsf{BP}\right] \preccurlyeq \\ &\mathrm{LWE}_{n,m,q,\chi_s,\chi_e}\left[\mathcal{U} \ , \ \mathsf{BP}\right] + 2^t \cdot \mathrm{negl}(\lambda) \ . \end{aligned}$$

This pair of Lemmas along with Lemma **??** directly imply the security of the scheme:

Corollary 1 (Security of Construction 4). *Under Assumption 31 and the standard decisional LWE assumption, for any sufficiently large* λ*, for any* PPT *adversary* A *there is a negligible function* $\mathrm{negl}(\cdot)$ *such that for any* $f \in \mathcal{F}$ *where* $f : \{0,1\}^\ell \to \{0,1\}$*, if* $f(x) = 0$ *for all* $x \in \{0,1\}^\ell$ *then*

$$\left|Pr\left[A\left(\mathsf{Enc}(1^\lambda, f, 0)\right) = 1\right] - Pr\left[A\left(\mathsf{Enc}(1^\lambda, f, 1)\right) = 1\right]\right| \leq 2^t \cdot \mathrm{negl}(\lambda) \ .$$

In the full version of this paper we prove Lemmas 1 and 2.

References

[ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *CRYPTO*, pages 98–115, 2010.

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618, 2009.

[Agr19] Shweta Agrawal. Indistinguishability obfuscation without multilinear maps: New methods for bootstrapping and instantiation. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 191–225. Springer, 2019.

[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, *Advances in Cryptology - CRYPTO 2019 - 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2019, Proceedings, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 284–332. Springer, 2019.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.

[Bar89] David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in nc^1. *J. Comput. Syst. Sci.*, 38(1):150–164, 1989.

[BDGM20] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Candidate io from homomorphic encryption schemes. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology - EUROCRYPT 2020 - 39th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, May 10-14, 2020, Proceedings, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 79–109. Springer, 2020.

[BL88] Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shafi Goldwasser, editor, *Advances in Cryptology*

- CRYPTO '88, 8th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1988, Proceedings, volume 403 of *Lecture Notes in Computer Science*, pages 27–35. Springer, 1988.

[BLP^{+}13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584. ACM, 2013.

[BV16] Zvika Brakerski and Vinod Vaikuntanathan. Circuit-abe from LWE: unbounded attributes and semi-adaptive security. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part III*, pages 363–384, 2016.

[CHKP12] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai trees, or how to delegate a lattice basis. *J. Cryptology*, 25(4):601–639, 2012.

[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, *Theory of Cryptography - 17th International Conference, TCC 2019, Nuremberg, Germany, December 1-5, 2019, Proceedings, Part I*, volume 11891 of *Lecture Notes in Computer Science*, pages 55–80. Springer, 2019.

[CVW18] Yilei Chen, Vinod Vaikuntanathan, and Hoeteck Wee. GGH15 beyond permutation branching programs: Proofs, attacks, and candidates. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology - CRYPTO 2018 - 38th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2018, Proceedings, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 577–607. Springer, 2018.

[GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, volume 9015 of *Lecture Notes in Computer Science*, pages 498–527. Springer, 2015.

[GGSW13] S. Garg, C. Gentry, A. Sahai, and B. Waters. Witness encryption and its applications. In *STOC*, 2013.

[GJLS21] Romain Gay, Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from simple-to-state hard problems: New assumptions, new techniques, and simplification. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 97–126. Springer, 2021.

[GLW14] Craig Gentry, Allison B. Lewko, and Brent Waters. Witness encryption from instance independent assumptions. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, volume 8616 of *Lecture Notes in Computer Science*, pages 426–443. Springer, 2014.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 197–206. ACM, 2008.

[JLMS19] Aayush Jain, Huijia Lin, Christian Matt, and Amit Sahai. How to leverage hardness of constant-degree expanding polynomials over r to build io. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology - EUROCRYPT 2019 - 38th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Darmstadt, Germany, May 19-23, 2019, Proceedings, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 251–281. Springer, 2019.

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 60–73. ACM, 2021.

[Lin16] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. *IACR Cryptology ePrint Archive*, 2016:257, 2016.

[Lin17] Huijia Lin. Indistinguishability obfuscation from SXDH on 5-linear maps and locality-5 prgs. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 599–629. Springer, 2017.

[LT17] Huijia Lin and Stefano Tessaro. Indistinguishability obfuscation from trilinear maps and block-wise local prgs. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 630–660. Springer, 2017.

[LV16] Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from ddh-like assumptions on constant-degree graded encodings. In Irit Dinur, editor, *IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9-11 October 2016, Hyatt Regency, New Brunswick, New Jersey, USA*, pages 11–20. IEEE Computer Society, 2016.

[MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In *Advances in Cryptology - CRYPTO 2011*, pages 465–484, 2011.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Advances in Cryptology - EUROCRYPT 2012 - 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cambridge, UK, April 15-19, 2012. Proceedings*, pages 700–718, 2012.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 333–342, 2009.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on*

Theory of Computing, Baltimore, MD, USA, May 22-24, 2005, pages 84–93, 2005.

[WW21] Hoeteck Wee and Daniel Wichs. Candidate obfuscation via oblivious LWE sampling. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology - EUROCRYPT 2021 - 40th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zagreb, Croatia, October 17-21, 2021, Proceedings, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 127–156. Springer, 2021.

[WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 600–611. IEEE Computer Society, 2017.