

Lockable Obfuscation

Rishab Goyal
UT Austin
rgoyal@cs.utexas.edu

Venkata Koppula
UT Austin
kvenkata@cs.utexas.edu

Brent Waters
UT Austin
bwaters@cs.utexas.edu*

Abstract

In this paper we introduce the notion of lockable obfuscation. In a lockable obfuscation scheme there exists an obfuscation algorithm Obf that takes as input a security parameter λ , a program P , a message msg and “lock value” α and outputs an obfuscated program \tilde{P} . One can evaluate the obfuscated program \tilde{P} on any input x where the output of evaluation is the message msg if $P(x) = \alpha$ and otherwise receives a rejecting symbol \perp .

We proceed to provide a construction of lockable obfuscation and prove it secure under the Learning with Errors (LWE) assumption. Notably, our proof only requires LWE with polynomial hardness and does not require complexity leveraging.

We follow this by describing multiple applications of lockable obfuscation. First, we show how to transform any attribute-based encryption (ABE) scheme into one in which the attributes used to encrypt the message are hidden from any user that is not authorized to decrypt the message. (Such a system is also known as predicate encryption with one-sided security.) The only previous construction due to Gorbunov, Vaikuntanathan and Wee is based off of a specific ABE scheme of Boneh et al. By enabling the transformation of any ABE scheme we can inherit different forms and features of the underlying scheme such as: multi-authority, adaptive security from polynomial hardness, regular language policies, etc.

We also show applications of lockable obfuscation to separation and uninstantiability results. We first show how to create new separation results in circular encryption that were previously based on indistinguishability obfuscation. This results in new separation results from learning with error including a public key bit encryption scheme that is IND-CPA secure and not circular secure. The tool of lockable obfuscation allows these constructions to be almost immediately realized by translation from previous indistinguishability obfuscation based constructions.

In a similar vein we provide random oracle uninstantiability results of the Fujisaki-Okamoto transformation (and related transformations) from the lockable obfuscation combined with fully homomorphic encryption. Again, we take advantage that previous work used indistinguishability obfuscation that obfuscated programs in a form that could easily be translated to lockable obfuscation.

1 Introduction

The topic of indistinguishability obfuscation has received an tremendous amount of attention from the cryptographic community over the last several years. Initially, the concept was introduced by Barak et al. [BGI⁺01, BGI⁺12] as a possible alternative to the notion of virtual black box obfuscation which they showed to be impossible to achieve for some functionalities. However, the concept indistinguishability obfuscation did not receive much immediate attention since (1) there were no such obfuscation candidates at the time and (2) the perceived lack of applications due to the fact that it only guaranteed security between two functionally equivalent circuits.

*Supported by NSF CNS-1228599 and CNS-1414082. DARPA through the U.S. Office of Naval Research under Contract N00014-11-1-0382, Google Faculty Research award, the Alfred P. Sloan Fellowship, Microsoft Faculty Fellowship, and Packard Foundation Fellowship.

In 2013, two works in the literature addressed these questions. First, Garg et al. [GGH⁺13b] provided the first indistinguishability obfuscation candidate using the Garg, Gentry and Halevi [GGH13a] multilinear map candidate. Then Sahai and Waters [SW14] introduced the “punctured programming” methodology for building cryptographic primitives from indistinguishability obfuscation which was used in their work and several subsequent works to resolve many open problems in cryptography.

When the potential of indistinguishability obfuscation was exposed, attention naturally moved to establishing security of obfuscation candidates since the original work of Garg et al. [GGH⁺13b] only provided a heuristic argument of security. Initial work in this line attempted to prove security under certain multilinear map models or assumptions [BGK⁺14, BR14, PST14, GLW14, GLSW15]. However, the security guarantees delivered from such proofs could only be as strong as the underlying multilinear map candidates [GGH13a, CLT13, GGH15, CLT15] which have been under a steady stream of cryptanalysis (see e.g. [CLT14, CHL⁺15, CGH⁺15, BGH⁺15, CLLT16, CLLT17, BWZ14, HJ16, Hal15, CFL⁺16, MSZ16, CJL16, ADGM16] and the references therein). To combat this there have been new multilinear map candidates proposed as well as models meant to capture most existing attacks [BMSZ16, GMM⁺16]. While these techniques present progress in defense against currently known cryptanalysis, it is unclear whether they can be connected to standard assumptions. Another set of works [BV15, AJ15, AJS15, Lin16a, LV16, Lin16b, AS16] have shown connections between certain types of functional encryption schemes and indistinguishability obfuscation with results showing that a constant degree multilinear maps combined with a constant depth PRG give indistinguishability obfuscation.

In this paper we approach the problem of achieving provably secure obfuscation from a different direction. Our philosophy is to anchor ourselves to the Learning with Errors (LWE) assumption and explore what applications and forms of obfuscation are achievable. Here we propose and define a new form of obfuscation that we call lockable obfuscation for which we give a construction that is provably secure under the LWE assumption. In addition, we show several applications of lockable obfuscation that were only known to this point under indistinguishability obfuscation.

We begin by informally introducing the notion of lockable obfuscation. A lockable obfuscation scheme consists of an obfuscation and evaluation algorithms. The (randomized) obfuscation algorithm Obf takes as input a security parameter λ , a program P , a message msg and lock value α and outputs an obfuscated program \tilde{P} . The evaluation algorithm Eval takes as input an obfuscated program \tilde{P} and an input x . If $P(x) = \alpha$ then the evaluation algorithm outputs the message msg , where P , msg and α were the program, message and lock values input to create \tilde{P} . Otherwise, if $P(x) \neq \alpha$ the evaluation algorithm (with high probability) outputs the \perp symbol to indicate rejection.

Intuitively, security states that if the lock value is chosen at random and kept from the attacker, then the attacker cannot learn anything about the program and message other than their sizes. That is there exist a simulator Sim such that for all program message pairs P, msg

$$\{\text{Obf}(1^\lambda, P, \text{msg}, \alpha) : \alpha \leftarrow \{0, 1\}^{m(\lambda)}\} \approx_c \text{Sim}(1^\lambda, 1^{|P|}, 1^{|\text{msg}|}).$$

We show how to construct a lockable obfuscation scheme for any polynomial sized circuit of sufficient output length from the Learning with Errors assumption. Our construction relies only on the polynomial hardness of the assumption (i.e. unlike witness encryption/indistinguishability obfuscation constructions, there is no sub-exponential hardness or complexity leveraging involved).¹ For this reason lockable obfuscation could be a preferred abstraction for building certain primitives even in a possible future where indistinguishability obfuscation is realizable from the assumption of LWE with subexponential hardness. The reason we don’t require subexponential hardness of LWE is that the security of our construction is derived from the hidden lock value α . In particular, our proof of security does *not* step through each possible input like many reductions to witness encryption or indistinguishability obfuscation [GLW14, GLSW15, BV15, AJ15, AJS15]. Also, note that since the lock value is chosen at random (and independent of the program), therefore lockable obfuscation could also be interpreted as obfuscation for a family of evasive functions [BBC⁺14].

¹Note that we still require subexponential LWE modulus. This translates to a subexponential approximation factor in the worst case hardness of lattice problems.

We will defer the explanation of our construction and proof to the technical overview of Section 1.1 and move on to discussing applications.

Hiding Attributes in Attribute-Based Encryption In a (key-policy) Attribute-Based Encryption (ABE) system [SW05] a user will encrypt a message `msg` under an attribute string x . A private key, as issued by an authority, will be associated with a boolean function f . When a user holding a secret key for function f attempts to decrypt a ciphertext he will learn the message if and only if $f(x) = 1$. Otherwise, the message remains hidden. While an ABE ciphertext will securely hide a private message, its security definition provides no guarantees about hiding the attribute string x . This can be problematic in many practical scenarios. Suppose we use ABE to encrypt email messages and use header metadata such as the sender, recipients, time and subject as attributes, where access functions are written over these attributes. In many settings this metadata itself will be very sensitive and disclosing it as part of the ciphertext is undesirable.

Almost all expressive² ABE systems built from standard tools³ allow for an attacker to learn the attribute string associated with a ciphertext. One notable exception is the work of Gorbunov, Vaikuntanathan and Wee [GVW15] that achieves one-sided security in hiding attributes. In that their construction hides the attribute string x as long as $f(x) = 0$ for all functions f that the adversary has a secret key for. This gives a much improved security picture to before as in our example an attacker will not be able to learn the header metadata of the emails so long as they are not authorized to decrypt them. We will call such a scheme a predicate encryption scheme with one-sided security.

The GVW construction was built by adapting the ABE system of Boneh et al. [BGG⁺14] and required an intricate knowledge of the original system in order to utilize certain special mathematical properties. The resulting construction achieves the same functionality of bounded depth circuits as the original as well as maintains selective security under the Learning with Errors assumption.

In this paper we show how to use lockable obfuscations (sometimes in combination with fully homomorphic encryption) to generically transform *any* ABE scheme into a predicate encryption scheme with one-sided security with corresponding functionality. An advantage of using a generic transformation is that it can take advantages of the features or forms of ABE constructions that have been introduced over the last 10+ years. including key-policy [GPSW06], ciphertext-policy [BSW07], adaptive security (without complexity leveraging) [LOS⁺10], multi-authority [Cha07, CC09, LW11], and efficient expression of keys as deterministic finite automata [Wat12, Att14] and circuits [GVW13, BGG⁺14]. Currently, there is no single ABE construction from standard tools that simultaneously delivers all such features. The most desirable ABE form can vary significantly between applications. Our transformations will allow one to inherit the properties of the underlying ABE scheme and at the same obtain one-sided hiding of the ciphertext attributes.

The technique for hiding attributes is fairly simple. To perform a (key-policy) encryption to attributes x and message `msg` the encryption algorithm first chooses a random lock value α . Next it creates a (sub) ciphertext C which is an encryption of the message α under the attribute string x in the original ABE system. Now consider the program P which takes as input a secret key and then decrypts C and outputs the result. The final ciphertext `ct` is a lockable obfuscation of this program P , under the lock value α and message `msg`. Key generation is the same as in the original scheme and decryption is simply using the lockable obfuscation evaluation algorithm on the `ct` with the secret key as the input.

Correctness can be observed. Suppose a user has a secret key sk_f for function f and applies it to a ciphertext encryption with attributes x where $f(x) = 1$). The evaluation algorithm will output `msg` since $P(sk_f) = \alpha$, which is the lock value. On the other hand suppose that the attacker does not have any secret keys for a function f where $f(x) = 1$, then by the message hiding security of the underlying ABE scheme, he won't be able to distinguish an encryption of the lock value α from an encryption of the all 0's string. Now that the lock value is hidden, one can take an additional hybrid step to argue that this is indistinguishable from a simulated obfuscation to erase knowledge of the program P including C and the attribute string x .

²We note that there are constructions of more limited expressivity such as vector matching [BW07] or inner product testing [KSW08] that achieve such security.

³For the purposes of this discussion we roughly consider number theoretic constructions grounded on RSA, bilinear maps and (Ring) LWE to be standard tools and those based on multilinear maps or indistinguishability obfuscation not to be.

The above construction works for any scheme with an apriori bounded size decryption circuit. By adding a level of indirection one can leverage leveled homomorphic encryption (which is realizable under LWE [BV11, BGV12]) to upgrade this to any scheme with an apriori decryption circuit depth. Or leverage fully homomorphic encryption to work without any depth bounds. The details of the transformation are given in Section 5.

Separation and Uninstantiability Results We now demonstrate the power of lockable obfuscation for achieving negative results in cryptography by focusing on two families of separation and uninstantiability results.

In recent years there has been a significant interest on the problem of circular security [CL01, Lau02, ABHS09] perhaps in large part due to Gentry’s [Gen09] result showing that a leveled homomorphic encryption scheme that is circular secure implies unbounded fully homomorphic encryption. Roughly, an encryption system is circular secure if it maintains security in the setting where there are n pairs of public keys and ciphertexts arranged such that the i^{th} ciphertext encrypts the $(i + 1)^{th}$ secret key.

There have been several separation results [ABBC10, CGH12, Rot13, KRW15, MO14, BHW15, AP16, KW16, GKW17a, GKW17b] showing that such security does not come for free. In particular, they show in different contexts that there exist schemes that are IND-CPA secure, but not circular secure. (We discuss this prior work in detail in Section 6). A natural dichotomy of these results is between separations achieved using indistinguishability obfuscation [KRW15, GKW17a] and those built from standard assumptions [ABBC10, CGH12, BHW15, AP16, KW16, GKW17b] such as assumptions on bilinear groups or LWE.

Separations built on indistinguishability obfuscation have the advantage that they can be developed relatively rapidly and are also relatively simple as one can build constructions using “normal” programs without diving into number theory. The disadvantage is that indistinguishability is not currently known from any standard assumptions. On the flip side, the number theoretic constructions are based on much more standard assumptions. However, they developing and understanding such solutions is a much more arduous task.

In Section 6 we show how to translate existing two indistinguishability obfuscation separation results due to Koppula, Ramchen and Waters [KRW15] and one result due Goyal, Koppula and Waters [GKW17a] to rely on lockable obfuscation. The translations are extremely straightforward and our resulting solutions are almost identical with the exception that we use lockable obfuscation. The main insight is that the programs obfuscated in the above results come in a lockable friendly form. In particular, they perform a sequence of computations on the input that result in a value s . Then the program tests and reports if $\text{PRG}(s) = t$ for a pseudorandom generator PRG and hardwired value t . Using lockable obfuscation one simply uses s as the output (which is a possible lock value).

Our translations of these prior works to lockable obfuscation lead to separations that can be built on the Learning with Errors assumption. Concretely, we obtain new results from LWE that were previously known only from indistinguishability obfuscation. (1) We show how to build a *public key* bit encryption scheme that is not circular secure and (2) we show a separation for unbounded length cycles.

Our second family of negative results relate to a grouping of constructions related to the well known Fujisaki-Okamoto transformation that achieves chosen ciphertext security in the random oracle model from any scheme which is IND-CPA secure. Included in this grouping are: the Bellare et al. [BHSV98] transformation from an IND-CPA scheme to an injective trapdoor function, two transformations from IND-CPA to IND-CCA security due to Fujisaki and Okamoto [FO99a, FO99b] and the deterministic encryption construction of Bellare, Boldyreva and O’Neill [BBO07].

All of the constructions follow a similar paradigm where they encrypt a string x under random coins determined from $H(x)$. (How the string x is construed varies somewhat between the schemes.) The works above show that if H is presented as an oracle access to a random function, the transformation results in a secure scheme under the relevant definition.

We give a random oracle uninstantiability [CGH98] result where using lockable obfuscation there exists an encryption scheme where for any hash function of up to apriori bounded size the applying the above transformations will result in an insecure encryption scheme — the message will be easily discoverable. If

we add the assumption of fully homomorphic encryption we can remove the bounded size restriction.

Brzuska, Farshim and Mittelbach [BFM15] achieved these results using indistinguishability obfuscation.⁴ We realize our results by simply translating the BFM result to move from indistinguishability obfuscation to lockable obfuscation. Again, this is possible because the programs obfuscated in the BFM paper follow the same lockable friendly form.

Indistinguishability Obfuscation for Rejecting Programs For our final application we now consider a new notion of obfuscation that we call indistinguishability obfuscation for rejecting programs and show how to construct it from lockable obfuscation and witness encryption [GGSW13] for circuit satisfiability.

Obfuscators that meet this notion will be defined over boolean circuits. Like indistinguishability obfuscation our obfuscator will take in any (not necessarily rejecting) boolean circuit C in a class and output an obfuscated program that is functionally equivalent to C . However, the security guarantees given by such an obfuscator are limited to “rejecting” programs. Informally, they state that no PPT adversary can distinguish between circuits C_0 and C_1 so long as for all inputs x $C_0(x) = C_1(x) = 0$. In contrast, standard indistinguishability obfuscation security allows C_0, C_1 to have arbitrary (both 0 and 1) outputs so long as they are functionally equivalent.

Our construction is simple and follows along the same conceptual lines as our techniques for building predicate encryption with one sided security from Attribute-Based Encryption.

Concurrent Work In an independent and concurrent work, Wichs-Zirdelis proposed a similar notion called *Obfuscation for Compute-and-Compare Programs*. While the notions are very similar, the syntax is different. A compute-and-compare program $CC[f, y]$ is defined by a function f and a value y . Obfuscation of $CC[f, y]$ outputs a program P such that P , on any input x , outputs 1 if $f(x) = y$. For security, [WZ17] require that for a randomly chosen y (from a high-entropy distribution), the obfuscation of $CC[f, y]$ is indistinguishable from simulated obfuscation, where the simulator gets only size of f and y . Wichs-Zirdelis also extend this notion where the obfuscation algorithm also takes as input a message m together with f, y , and the obfuscated program P on input x , outputs m if $f(x) = y$. Security requires that the message m is also hidden. This message-based version of obfuscation of compute-and-compare programs is identical to our notion of lockable obfuscation, modulo the distribution of lock y , which is uniform in our case and can be any high-entropy distribution in their case. However, this gap can be simply bridged by using an PRG for pseudo-entropy seeds instead of a regular PRG.

Both works have a few applications in common. This includes the ABE to predicate encryption transformation, witness encryption to reject-iO and circular security counterexamples. In addition, we show new uninstantiability results in the random oracle model. Applications unique to [WZ17] are transformation from any secure sketch [DORS08] to private secure sketch [DS05], and obfuscation for conjunctions/affine testers.

1.1 Overview of our Lockable Obfuscation Construction

We will now describe our lockable obfuscation scheme for a family of poly-depth circuits. The construction is described in details in Section 4. At a high level, our scheme can be divided into three components — (1) A lockable obfuscation scheme for a family of low-depth circuits and 1-bit messages, (2) a bootstrapping mechanism to amplify to lockable obfuscation for a family of poly-depth circuits and 1-bit messages, and (3) extending to lockable obfuscation for a family of poly-depth circuits and multi-bit messages. (Note in our actual construction, we combine the first two components into one for technical reasons.)

Lockable Obfuscation for Low-Depth Circuits and 1-Bit Messages. The primary ingredients of our construction are low-depth pseudorandom generators (PRGs), lattice trapdoors [GPV08], telescoping products/cascading cancellations [KW16, AP16, GGH15] and oblivious sequence transformation [GKW17b].

⁴Technically, their result with no bounds on the hash function required indistinguishability obfuscation for Turing Machines with unbounded input. However, this could have been replaced with indistinguishability obfuscation for circuits and fully homomorphic encryption.

First, let us recall the notion of permutation branching programs, lattice trapdoors and oblivious sequence transformation. A permutation branching program of length L and width w can be represented using w states, $2L$ permutations $\sigma_{j,b}$ over states for each level $j \leq L$, an input-selector function $\text{inp}(\cdot)$ which determines the input read at each level, and an accepting and rejecting state. The program execution starts at state 1 of level 0. Suppose the branching program reads first input bit (say b) at level 1 (i.e., $\text{inp}(1) = 1$). Then, the state of the program changes to $\sigma_{1,b}(\text{st})$. Such a process can be carried out (iteratively) to compute the final program state at level L . Depending upon the final state, the program either accepts or rejects.

A lattice trapdoor generation algorithm can be used to sample a (uniformly looking) matrix \mathbf{A} together with a trapdoor $T_{\mathbf{A}}$. The trapdoor can be used to compute, for any matrix \mathbf{U} , a low norm matrix \mathbf{S} such that $\mathbf{A} \cdot \mathbf{S} = \mathbf{U}$.⁵ As a result, the matrix \mathbf{S} can be used to ‘transform’ any matrix $\tilde{\mathbf{A}} \approx \mathbf{A}$ to another matrix $\tilde{\mathbf{U}} \approx \mathbf{U}$. Oblivious sequence transformation is a technique that enables sampling a sequence of matrices $\mathbf{B}_1, \dots, \mathbf{B}_w$ along with some trapdoor such that for any sequence of matrices $\mathbf{U}_1, \dots, \mathbf{U}_w$, one can construct a short matrix \mathbf{X} such that $\mathbf{B}_i \cdot \mathbf{X} = \mathbf{U}_i$ for all i (and \mathbf{X} is oblivious of i).

Moving on to our construction, at a high level the obfuscator starts by first generating a sequence of permutation branching programs corresponding to the circuit C (where each branching program computes one output bit), and then encoding the state transition permutations for each level for every branching program using the technique of oblivious sequence transformation. Let ℓ be the output length of C . In other words, for obfuscating circuit C under lock α with message msg , the obfuscator first expresses C as a set of width 5 permutation branching programs $\{\text{BP}^{(i)}\}_i$ of polynomial length L , where for each $i \in [\ell]$ $\text{BP}^{(i)}$ computes the i^{th} output bit of circuit C .⁶ Without loss of generality, we can assume that all branching programs have a common input selector function $\text{inp}(\cdot)$ such that $\text{inp}(j)$ bit of the input is read at level j . The obfuscation algorithm continues by choosing $5(L+1)\ell$ matrices, one matrix for each (level, state) of each branching program.⁷ Let $\mathbf{B}_{j,w}^{(i)}$ be the matrix corresponding to state $w \in [5]$ at level $j \leq L$ for branching program $\text{BP}^{(i)}$, $i \in [\ell]$. For every $i \in [\ell]$, the matrices $\{\mathbf{B}_{j,1}^{(i)}, \dots, \mathbf{B}_{j,5}^{(i)}\}$ for the first L levels (i.e., all but top-level matrices) are sampled such that they have a common trapdoor $T_j^{(i)}$, i.e. using oblivious sequence transformation. The top-level matrices, however, are sampled uniformly at random without a trapdoor subject to the constraint that the top-level matrices corresponding to lock string α sum to a special fixed matrix depending upon the message msg . More formally, for each $i \in [\ell]$, let q be the LWE modulus, and $\text{acc}^{(i)}$ and $\text{rej}^{(i)}$ be the accepting and rejecting states for $\text{BP}^{(i)}$, then the obfuscator chooses the matrices $\mathbf{B}_{L,\text{acc}^{(i)}}^{(i)}, \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)}$ such that

$$\sum_{i : \alpha_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i : \alpha_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if msg} = 0. \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

For each level $j \in [L]$, the obfuscation algorithm also chooses two low-norm matrices $\mathbf{S}_j^{(0)}$ and $\mathbf{S}_j^{(1)}$ (these are shared across all branching programs), and computes 2ℓ low-norm matrices $\{\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}\}_{i,j}$ such that for every state $w \in [5]$, $\mathbf{B}_{j-1,w}^{(i)} \cdot \mathbf{C}_j^{(i,0)} \approx \mathbf{S}_j^{(0)} \cdot \mathbf{B}_{j,\sigma_{j,0}^{(i)}(w)}^{(i)}$ and $\mathbf{B}_{j-1,w}^{(i)} \cdot \mathbf{C}_j^{(i,1)} \approx \mathbf{S}_j^{(1)} \cdot \mathbf{B}_{j,\sigma_{j,1}^{(i)}(w)}^{(i)}$. That is, the matrices $\mathbf{C}_j^{(i,0)}$ and $\mathbf{C}_j^{(i,1)}$ represent the state transition from level $j-1$ to j when bit 0 or 1 is read at step j of branching program execution. For each $i \in [\ell], j \in [L]$, the $\mathbf{C}_j^{(i,b)}$ matrices can be generated using the lattice trapdoors $T_j^{(i)}$. The obfuscation algorithm outputs these matrices $\{\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}\}_{i,j}$ together with the base-level matrices $\{\mathbf{B}_{0,1}^{(i)}\}_i$ as the final obfuscated program.

At a high level, one could visualize the obfuscated program which consists of the base-level matrices

⁵For ease of exposition, we will use the notation $\mathbf{A}^{-1}(\cdot)$ to represent the pre-image operation. In the formal description of our algorithms later, we use the pre-image sampling algorithm `SamplePre`.

⁶From Barrington’s Theorem [Bar86], we know that for every NC^1 circuits there exists a width 5 permutation branching program of polynomial length.

⁷Note that if a branching program has length L , then it has $L+1$ levels.

$\{\mathbf{B}_{0,1}^{(i)}\}_i$ and matrices $\{\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}\}_{i,j}$ as “encodings” of the branching program starting states and state transition permutations, respectively. Therefore, evaluating an obfuscated program on some input x will be analogous to evaluating the branching programs $\text{BP}^{(i)}$ on input x directly. Fix some $i \in [\ell]$. Suppose the first input bit x_1 is read at level 1. Then evaluation of $\text{BP}^{(i)}$ at level 1 would map the state 1 at level 0 to state $\sigma_{1,x_1}^{(i)}$ at level 1. Analogously, the obfuscation evaluator can compute $\mathbf{B}_{0,1}^{(i)} \cdot \mathbf{C}_1^{(i,x_1)} \approx \mathbf{S}_1^{(x_1)} \cdot \mathbf{B}_{1,\sigma_{1,x_1}^{(i)}}^{(i)}$. In general, if the program state at level $j-1$ during execution is w , then the evaluator will accumulate the product of the form $\mathbf{\Gamma}_{j-1} \cdot \mathbf{B}_{j,w}^{(i)}$, where $\mathbf{\Gamma}_{j-1}$ is a product of $j-1$ low-norm matrices. This can be easily verified as follows. Suppose the next bit read is b , then the new state at level j will be $\sigma_{j,b}^{(i)}(w)$, thus the new accumulated product during obfuscation evaluation will be $\mathbf{\Gamma}_{j-1} \cdot \mathbf{B}_{j,w}^{(i)} \cdot \mathbf{C}_j^{(i,b)} \approx \mathbf{\Gamma}_j \cdot \mathbf{B}_{j+1,\sigma_{j,b}^{(i)}(w)}^{(i)}$, where $\mathbf{\Gamma}_j = \mathbf{\Gamma}_{j-1} \cdot \mathbf{S}_j^{(b)}$. Therefore, the invariant is maintained. Note that the matrix $\mathbf{\Gamma}_L$ will be same for all branching programs since the low-norm matrices $\mathbf{S}_j^{(0)}$ and $\mathbf{S}_j^{(1)}$ are shared across all branching programs.

Continuing this way, the evaluator can iteratively compute the matrix product at the top. Thus, for each branching program, the accumulated product at the top will either be $\approx \mathbf{\Gamma}_L \cdot \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)}$ or $\approx \mathbf{\Gamma}_L \cdot \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)}$, depending on whether $C(x)_i = 0$ or 1. Let $\mathbf{\Delta}^{(i)} = \mathbf{\Gamma}_L \cdot \mathbf{B}_{L,\text{st}^{(i)}}^{(i)}$, where $\text{st}^{(i)} = \text{acc}^{(i)}$ or $\text{rej}^{(i)}$ depending on $C(x)_i$. Finally, the evaluator simply sums the top-level accumulated products ($\approx \mathbf{\Delta}^{(i)}$) and checks whether the norm of the final summed matrix lies in appropriate range. More concretely, consider the case when $C(x) = \alpha$ and $\text{msg} = 0$, then $\sum_i \mathbf{\Delta}^{(i)} = \mathbf{\Gamma}_L \cdot \sum_i \mathbf{B}_{L,\text{st}^{(i)}}^{(i)} = \mathbf{0}^{n \times m}$. Since the final top-level matrix sum is close to $\sum_i \mathbf{\Delta}^{(i)}$, thus it will have norm close to 0, and hence the evaluator can simply test this and output 0 as the message.

Similarly we could argue correctness for the cases when $\text{msg} = 1$ or $C(x) \neq \alpha$. However, our current proof techniques do not seem sufficient for proving the security of above construction. The reason is that there is an inherent tension in setting scheme parameters while basing security on LWE. This is discussed in detail later in Section 4.3. We were able to bypass this problem by obfuscating an “expanded” circuit, which evaluates a low-depth pairwise independent hash function h and a low-depth pseudorandom generator PRG with a large enough polynomial stretch (in succession) on the output of circuit C , instead of directly obfuscating circuit C using the above matrix encoding procedure.

In other words, let Q be the circuit that on input x , outputs $\text{PRG}(h(C(x)))$. Observe that if h and PRG can be computed by low-depth circuits (\mathbf{NC}^1), then Q also can be computed by an \mathbf{NC}^1 circuit (since C is assumed to be a log-depth circuit). Therefore, Q can be expressed by a set of width 5 permutation branching programs of polynomial length L as well. Additionally, now the matrix component generation procedure will use $\beta = \text{PRG}(h(\alpha))$ as the lock instead of α . This modification is sufficient to avoid the tension between scheme parameters, and also allows us to prove security of our scheme under LWE with only polynomial hardness. This completes the description of our lockable obfuscation scheme for low-depth circuits and 1-bit messages.

Bootstrapping Lockable Obfuscation. Let $\mathcal{O}_{\mathbf{NC}^1}$ be a lockable obfuscator for log-depth circuits. We will use leveled homomorphic encryption (LHE) with an \mathbf{NC}^1 decryption circuit to bootstrap $\mathcal{O}_{\mathbf{NC}^1}$ to an obfuscator that works for any depth d . The obfuscator gets as input a circuit C of depth d , a string α and a message msg . It first chooses the LHE secret-evaluation keys and encrypts the circuit C . Let Q be the circuit which takes as input an LHE ciphertext and decrypts it using the hardwired LHE secret key and outputs the decrypted string. Note that the circuit Q is a logarithmic depth circuit. The obfuscator outputs $\mathcal{O} \leftarrow \mathcal{O}_{\mathbf{NC}^1}(Q, \alpha, \text{msg})$ together with the encryption of C and the LHE evaluation key.

Evaluating on input x . Let $U_x(\cdot)$ be the universal circuit with input x hardwired (that is, it takes a circuit C as input and outputs $C(x)$). The evaluation algorithm first homomorphically evaluates the circuit U_x on the encryption of C . This results in an LHE ciphertext ct , which is an encryption of $C(x)$. It then evaluates the obfuscation \mathcal{O} on input ct , and outputs the resulting string. Using the correctness of \mathcal{O} and the LHE scheme, we can argue that if $C(x) = \alpha$, then the evaluation outputs msg , and it outputs \perp otherwise.

The security proof here is fairly simple. Using the security of the underlying obfuscator $\mathcal{O}_{\text{NC}^1}$, we first switch the obfuscation \mathcal{O} to be a simulated obfuscation. Once the obfuscator \mathcal{O} is simulated, the obfuscator no longer needs the LHE secret key. Therefore, we can now replace the LHE encryption of C with encryption of zeros, thereby erasing all the information about circuit C except its size. Therefore, the final simulator simply outputs an encryption of zeros, together with a simulated obfuscation, and this is indistinguishable from the honestly computed obfuscation.

Extending Lockable Obfuscation for 1-Bit Messages to Multi-Bit Messages. We would like point out that the standard repetition method for extending message space does not work for lockable obfuscation schemes because the security is only guaranteed when the adversary does not know the lock string α . However, we observe there are still multiple ways to extend its message space. We briefly discuss two possible extensions. One option could be to encode a multi-bit message directly in the top-level matrices. Currently, the top-level matrices are set to sum to $\mathbf{0}^{n \times m}$ if $\text{msg} = 0$, otherwise to $\sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}]$. However, if we interpret the message msg as an integer $v < \sqrt{q}/2$, then we could simply set the sum to be $v \cdot \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}]$.

The second extension could be carried more generally without exploiting the mathematical structure of the underlying obfuscation scheme. The high level idea is to again “expand” the circuit C using a pairwise independent hash function and a pseudorandom generator before obfuscation. Suppose the lock α be a string of length k . Let $\beta = \text{PRG}(h(\alpha))$ and $|\beta| = \ell \cdot k$. To obfuscate circuit C under lock α for an ℓ -bit message msg , the multi-bit obfuscator (for each $i < \ell$) independently obfuscates the circuit $Q[i]$ under lock $\beta[i]$ for message msg_i using the 1-bit obfuscation scheme, where circuit $Q[i]$ denotes the circuit that outputs the $i \cdot k + 1, \dots, (i + 1) \cdot k$ output bits of circuit $\text{PRG}(h(C(\cdot)))$ and $\beta[i] = \beta_{i \cdot k + 1}, \dots, \beta_{(i + 1) \cdot k}$. The security proof follows from a simple hybrid argument. This transformation is described later in Appendix C.

This completes the technical overview of our lockable obfuscation scheme.

1.2 More on Related Encoding Works

The idea of using lattice trapdoors for constructing multilinear maps was first seen in the work of Gentry, Gorbunov and Halevi (GGH) where they proposed a candidate for graph-induced multilinear maps [GGH15]. Their work builds upon the homomorphic encryption scheme of Gentry, Sahai and Waters [GSW13] which could be considered as the starting point of cascading cancellations technique. In [GGH15], there is a fixed (directed acyclic) graph $G = (V, E)$, and plaintext messages are associated with edges. Each vertex u has an associated matrix \mathbf{A}_u and a secret parameter T_u which is the trapdoor for \mathbf{A}_u . The plaintext space consists of matrices, and the encoding of a matrix \mathbf{M} along an edge (u, v) is $\mathbf{A}_u^{-1}(\mathbf{M} \cdot \mathbf{A}_v + \text{noise})$. Note that given an encoding of matrix \mathbf{M}_1 for edge (u, v) and an encoding of matrix \mathbf{M}_2 for edge (v, w) , one can compute an encoding of $\mathbf{M}_1 \cdot \mathbf{M}_2$ along path (u, v, w) . Informally, their intuition was that given all the encodings it should be easy to compute an encoding for any path in graph G , however it should remain hard for anyone to generate encodings over any non-path combination of vertices. Using these multilinear maps, GGH gave candidate constructions for obfuscation and multipartite key-agreement. However, there were no security proofs for these candidates, and their key-agreement protocol was later shown to be broken [CLLT16]. In a later work, Brakerski et al. [BVWW16] gave a construction for obfuscating conjunctions based on an entropic variant of Ring-LWE. In that work, they observed that if the underlying graph was a straight-line graph (i.e., an ordered sequence of vertices), then the corresponding GGH multilinear maps provide some provable security properties.

In a more recent work, Koppula and Waters [KW16] used what they called cascading cancellations technique for constructing k -circular security separations. Their construction deviates from the GGH/BVWW paradigm in two crucial aspects. First, their construction involves multiple *strands* (say ℓ) of length k , instead of just a single strand. Unlike the previous works, which involved directly comparing two distinct encodings (or, components) along the same path for equality, they first combined (using matrix multiplications) all the

k components for each different strand, and then summed the ℓ final components (one for each strand) to perform an equality check with a fixed value. Second, what is mechanically close to the “plaintext” encoded values in GGH is viewed here as simply random matrices (and just part of the overall randomness) and not a value to be encoded. This differs from GGH in which the elements being encoded were labeled as “plaintext” values. For instance, the elements that were encoded in the GGH obfuscation candidate were the state transition matrices of the actual branching program being obfuscated and thus can reflect some stronger semantics. Concurrently, Alami and Peikert [AP16] also provided circular security separations, that had a cancellation type effect although with a different technical approach.

Most recently, Goyal, Koppula and Waters [GKW17b] further advanced the existing techniques for constructing bit-circular security separations. One of their most important contributions was an alternative mechanism to encode and *hide* (branching) programs using lattice trapdoors. To this end, they introduced a novel technique which they call oblivious sequence transformation. This is a significant departure from prior works as previously most works generated the matrix components (or encodings) independently, however it was essential in their construction that the components be jointly generated. At a high level, they provided new techniques to encode and hide a permutation between a sequence of nodes. Informally, they gave a mechanism to encode a permutation between nodes u_1, \dots, u_5 and w_1, \dots, w_5 such that given a node u_i one could obviously go to its corresponding node w_j . Another important aspect of their work was to encode a log-depth Pseudo Random Generator (PRG) using the oblivious sequence transformation technique such that the PRG could be publicly evaluated and if the output of computation is some fixed (but unknown) value, then it could be efficiently tested. This seems to be one of the most important technical aspect of their work since in order to prove security from the LWE assumption as well as guarantee efficient testing, encoding a log-depth PRG with a large polynomial stretch is essential. We finally remark that Canetti and Chen [CC17] recently used the LWE assumption to achieve 1-collusion secure constrained PRFs for **NC1** with constraint hiding. Their construction involves a single strand like BVWW and their embedding of a branching program follows the GGH style of putting branching programs into the encoded plaintext values much more closely.

2 Preliminaries

Notations. Let PPT denote probabilistic polynomial-time. We will use lowercase bold letters for vectors (e.g. \mathbf{v}) and uppercase bold letters for matrices (e.g. \mathbf{A}). For any integer $q \geq 2$, we let \mathbb{Z}_q denote the ring of integers modulo q . We represent \mathbb{Z}_q as integers in the range $(-q/2, q/2]$. For a vector \mathbf{v} , we let $\|\mathbf{v}\|$ denote its ℓ_2 norm and $\|\mathbf{v}\|_\infty$ denote its infinity norm. Similarly, for matrices $\|\cdot\|$ and $\|\cdot\|_\infty$ denote their ℓ_2 and infinity norms (respectively).

We denote the set of all positive integers upto n as $[n] := \{1, \dots, n\}$. For any finite set S , $x \leftarrow S$ denotes a uniformly random element x from the set S . Similarly, for any distribution \mathcal{D} , $x \leftarrow \mathcal{D}$ denotes an element x drawn from distribution \mathcal{D} . The distribution \mathcal{D}^n is used to represent a distribution over vectors of n components, where each component is drawn independently from the distribution \mathcal{D} . Two distributions \mathcal{D}_1 and \mathcal{D}_2 , parameterized by security parameter λ , are said to be computationally indistinguishable, represented by $\mathcal{D}_1 \approx_c \mathcal{D}_2$, if for all PPT adversaries \mathcal{A} , $|\Pr[1 \leftarrow \mathcal{A}(x) : x \leftarrow \mathcal{D}_1] - \Pr[1 \leftarrow \mathcal{A}(x) : x \leftarrow \mathcal{D}_2]| \leq \text{negl}(\lambda)$. For a family of distributions $\chi = \{\chi(\lambda)\}_\lambda$ over the integers, and integers bounds $B = \{B(\lambda)\}_\lambda$, we say that χ is B -bounded if $\Pr[|x| \leq B(\lambda) : x \leftarrow \chi(\lambda)] = 1$.

This section closely follows [GKW17b].

Min-Entropy and Randomness Extraction. The min-entropy of a random variable X is defined as $\mathbf{H}_\infty(X) \stackrel{\text{def}}{=} -\log_2(\max_x \Pr[X = x])$. Let $\text{SD}(X, Y)$ denote the statistical distance between two random variables X and Y . Below we state the Leftover Hash Lemma (LHL) from [HILL99, DRS04, DORS08].

Theorem 2.1. Let $\mathcal{H} = \{h : X \rightarrow Y\}_{h \in \mathcal{H}}$ be a universal hash family, then for any random variable W

taking values in X , the following holds

$$\text{SD}((h, h(W)), (h, U_Y)) \leq \frac{1}{2} \sqrt{2^{-\mathbf{H}_\infty(W)} \cdot |Y|}.$$

We will use the following corollary, which follows from the Leftover Hash Lemma.

Corollary 2.1. Let $\ell > m \cdot n \log_2 q + \omega(\log n)$ and q a prime. Let \mathbf{R} be an $k \times m$ matrix chosen as per distribution \mathcal{R} , where $k = k(n)$ is polynomial in n and $\mathbf{H}_\infty(\mathcal{R}) = \ell$. Let \mathbf{A} and \mathbf{B} be matrices chosen uniformly in $\mathbb{Z}_q^{n \times k}$ and $\mathbb{Z}_q^{n \times m}$, respectively. Then the statistical distance between the following distributions is negligible in n .

$$\{(\mathbf{A}, \mathbf{A} \cdot \mathbf{R})\} \approx_s \{(\mathbf{A}, \mathbf{B})\}$$

Proof. The proof of above corollary follows directly from the Leftover Hash Lemma. Note that for a prime q the family of hash functions $h_{\mathbf{A}} : \mathbb{Z}_q^{k \times m} \rightarrow \mathbb{Z}_q^{n \times m}$ for $\mathbf{A} \in \mathbb{Z}_q^{n \times k}$ defined by $h_{\mathbf{A}}(\mathbf{X}) = \mathbf{A} \cdot \mathbf{X}$ is universal. Therefore, if \mathcal{R} has sufficient min-entropy, i.e. $\ell > m \cdot n \log_2 q + \omega(\log n)$, then the Leftover Hash Lemma states that statistical distance between the distributions $(\mathbf{A}, \mathbf{A} \cdot \mathbf{R})$ and (\mathbf{A}, \mathbf{B}) is at most $2^{-\omega(\log n)}$ which is negligible in n as desired. \blacksquare

2.1 Lattice Preliminaries

An m -dimensional lattice \mathcal{L} is a discrete additive subgroup of \mathbb{R}^m . Given positive integers n, m, q and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we let $\Lambda_q^\perp(\mathbf{A})$ denote the lattice $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{0} \pmod{q}\}$. For $\mathbf{u} \in \mathbb{Z}_q^n$, we let $\Lambda_q^\mathbf{u}(\mathbf{A})$ denote the coset $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x} = \mathbf{u} \pmod{q}\}$.

Discrete Gaussians. Let σ be any positive real number. The Gaussian distribution \mathcal{D}_σ with parameter σ is defined by the probability distribution function $\rho_\sigma(\mathbf{x}) = \exp(-\pi \|\mathbf{x}\|^2 / \sigma^2)$. For any set $\mathcal{L} \subset \mathbb{R}^m$, define $\rho_\sigma(\mathcal{L}) = \sum_{\mathbf{x} \in \mathcal{L}} \rho_\sigma(\mathbf{x})$. The discrete Gaussian distribution $\mathcal{D}_{\mathcal{L}, \sigma}$ over \mathcal{L} with parameter σ is defined by the probability distribution function $\rho_{\mathcal{L}, \sigma}(\mathbf{x}) = \rho_\sigma(\mathbf{x}) / \rho_\sigma(\mathcal{L})$ for all $\mathbf{x} \in \mathcal{L}$.

The following lemma (Lemma 4.4 of [MR07], [GPV08]) shows that if the parameter σ of a discrete Gaussian distribution is small, then any vector drawn from this distribution will be short (with high probability).

Lemma 2.1. Let m, n, q be positive integers with $m > n$, $q \geq 2$. Let $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ be a matrix of dimensions $n \times m$, $\sigma = \tilde{\Omega}(n)$ and $\mathcal{L} = \Lambda_q^\perp(\mathbf{A})$. Then

$$\Pr[\|\mathbf{x}\| > \sqrt{m} \cdot \sigma : \mathbf{x} \leftarrow \mathcal{D}_{\mathcal{L}, \sigma}] \leq \text{negl}(n).$$

Truncated Discrete Gaussians. The truncated discrete Gaussian distribution over \mathbb{Z}^m with parameter σ , denoted by $\tilde{\mathcal{D}}_{\mathbb{Z}^m, \sigma}$, is same as the discrete Gaussian distribution \mathcal{D}_σ except it outputs 0 whenever the ℓ_∞ norm exceeds $\sqrt{m} \cdot \sigma$. Note that, by definition, $\tilde{\mathcal{D}}_{\mathbb{Z}^m, \sigma}$ is $\sqrt{m} \cdot \sigma$ -bounded. Also, note that $\tilde{\mathcal{D}}_{\mathbb{Z}^m, \sigma} \approx_s \mathcal{D}_{\mathbb{Z}^m, \sigma}$.

Learning with Errors (LWE). The Learning with Errors (LWE) problem was introduced by Regev [Reg05]. The LWE problem has four parameters: the dimension of the lattice n , the number of samples m , the modulus q and the error distribution $\chi = \chi(n)$.

Assumption 1 (Learning with Errors). Let n, m and q be positive integers and χ a noise distribution over \mathbb{Z}_q . The Learning with Errors assumption (n, m, q, χ) -LWE, parameterized by n, m, q, χ , states that the following distributions are computationally indistinguishable:

$$\left\{ (\mathbf{A}, \mathbf{s}^\top \cdot \mathbf{A} + \mathbf{e}^\top) : \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \\ \mathbf{s} \leftarrow \mathbb{Z}_q^n, \mathbf{e} \leftarrow \chi^m \end{array} \right\} \approx_c \left\{ (\mathbf{A}, \mathbf{u}^\top) : \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \\ \mathbf{u} \leftarrow \mathbb{Z}_q^m \end{array} \right\}$$

Under a quantum reduction, Regev [Reg05] showed that for certain noise distributions, LWE is as hard as worst case lattice problems such as the decisional approximate shortest vector problem (GapSVP) and approximate shortest independent vectors problem (SIVP). Later works [Pei09, BLP⁺13] showed classical reductions from LWE to GapSVP_γ.

These works show that for B -bounded discretized Gaussian error distributions χ , solving (n, m, q, χ) -LWE is as hard as approximating GapSVP_γ and SIVP_γ to a factor of $\tilde{O}(n \cdot q/B)$. Given the current state of art in lattice algorithms, GapSVP_γ and SIVP_γ are believed to be hard for $\gamma = \tilde{O}(2^{n^\epsilon})$ (for fixed $\epsilon \in (0, 1/2)$), and therefore (n, m, q, χ) -LWE is believed to be hard for B -bounded discretized Gaussian error distributions χ with $B = 2^{-n^\epsilon} \cdot q \cdot \text{poly}(n)$.

LWE with Short Secrets. In this work, we will be using a variant of the LWE problem called *LWE with Short Secrets*. In this variant, introduced by Applebaum et al. [ACPS09], the secret vector is also chosen from the noise distribution χ . They showed that this variant is as hard as LWE for sufficiently large number of samples m .

Assumption 2 (LWE with Short Secrets). Let n, m and q be positive integers and χ a noise distribution on \mathbb{Z} . The LWE with Short Secrets assumption (n, m, q, χ) -LWE-ss, parameterized by n, m, q, χ , states that the following distributions are computationally indistinguishable⁸:

$$\left\{ (\mathbf{A}, \mathbf{S} \cdot \mathbf{A} + \mathbf{E}) : \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \\ \mathbf{S} \leftarrow \chi^{n \times n}, \mathbf{E} \leftarrow \chi^{n \times m} \end{array} \right\} \approx_c \left\{ (\mathbf{A}, \mathbf{U}) : \begin{array}{l} \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \\ \mathbf{U} \leftarrow \mathbb{Z}_q^{n \times m} \end{array} \right\}.$$

Lattices with Trapdoors. Lattices with trapdoors are lattices that are statistically indistinguishable from randomly chosen lattices, but have certain ‘trapdoors’ that allow efficient solutions to hard lattice problems.

Definition 2.1 ([Ajt99, GPV08]). A trapdoor lattice sampler consists of algorithms TrapGen and SamplePre with the following syntax and properties:

- **TrapGen**($1^n, 1^m, q$) $\rightarrow (\mathbf{A}, T_{\mathbf{A}})$: The lattice generation algorithm is a randomized algorithm that takes as input the matrix dimensions n, m , modulus q , and outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $T_{\mathbf{A}}$.
- **SamplePre**($\mathbf{A}, T_{\mathbf{A}}, \mathbf{u}, \sigma$) $\rightarrow \mathbf{s}$: The presampling algorithm takes as input a matrix \mathbf{A} , trapdoor $T_{\mathbf{A}}$, a vector $\mathbf{u} \in \mathbb{Z}_q^n$ and a parameter $\sigma \in \mathcal{R}$ (which determines the length of the output vectors). It outputs a vector $\mathbf{s} \in \mathbb{Z}_q^m$.

These algorithms must satisfy the following properties:

1. *Correct Presampling*: For all vectors \mathbf{u} , parameters σ , $(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$, and $\mathbf{s} \leftarrow \text{SamplePre}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{u}, \sigma)$, $\mathbf{A} \cdot \mathbf{s} = \mathbf{u}$ and $\|\mathbf{s}\|_\infty \leq \sqrt{m} \cdot \sigma$.

2. *Well Distributedness of Matrix*: The following distributions are statistically indistinguishable:

$$\{\mathbf{A} : (\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)\} \approx_s \{\mathbf{A} : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}\}.$$

3. *Well Distributedness of Preimage*: For all $(\mathbf{A}, T_{\mathbf{A}}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$, if $\sigma = \omega(\sqrt{n \cdot \log q \cdot \log m})$, then the following distributions are statistically indistinguishable:

$$\{\mathbf{s} : \mathbf{u} \leftarrow \mathbb{Z}_q^n, \mathbf{s} \leftarrow \text{SamplePre}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{u}, \sigma)\} \approx_s \mathcal{D}_{\mathbb{Z}^m, \sigma}.$$

These properties are satisfied by the gadget-based trapdoor lattice sampler of [MP12] for parameters m such that $m = \Omega(n \cdot \log q)$.

⁸Applebaum et al. showed that $\{(\mathbf{A}, \mathbf{s}^\top \cdot \mathbf{A} + \mathbf{e}) : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{s} \leftarrow \chi^n, \mathbf{e} \leftarrow \chi^m\} \approx_c \{(\mathbf{A}, \mathbf{u}) : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{u} \leftarrow \mathbb{Z}_q^m\}$, assuming LWE is hard. However, by a simple hybrid argument, we can replace vectors $\mathbf{s}, \mathbf{e}, \mathbf{u}$ with matrices $\mathbf{S}, \mathbf{E}, \mathbf{U}$ of appropriate dimensions.

2.2 Branching Programs

Branching programs are a model of computation used to capture space-bounded computations [BDFP86, Bar86]. In this work, we will be using a restricted notion called *permutation branching programs*.

Definition 2.2 (Permutation Branching Program). A permutation branching program of length L , width w and input space $\{0, 1\}^n$ consists of a sequence of $2L$ permutations $\sigma_{i,b} : [w] \rightarrow [w]$ for $1 \leq i \leq L, b \in \{0, 1\}$, an input selection function $\text{inp} : [L] \rightarrow [n]$, an accepting state $\text{acc} \in [w]$ and a rejection state $\text{rej} \in [w]$. The starting state st_0 is set to be 1 without loss of generality. The branching program evaluation on input $x \in \{0, 1\}^n$ proceeds as follows:

- For $i = 1$ to L ,
 - Let $\text{pos} = \text{inp}(i)$ and $b = x_{\text{pos}}$. Compute $\text{st}_i = \sigma_{i,b}(\text{st}_{i-1})$.
- If $\text{st}_L = \text{acc}$, output 1. If $\text{st}_L = \text{rej}$, output 0, else output \perp .

In a remarkable result, Barrington [Bar86] showed that any circuit of depth d can be simulated by a permutation branching program of width 5 and length 4^d .

Theorem 2.2 ([Bar86]). For any boolean circuit C with input space $\{0, 1\}^n$ and depth d , there exists a permutation branching program BP of width 5 and length 4^d such that for all inputs $x \in \{0, 1\}^n$, $C(x) = \text{BP}(x)$.

This permutation property will be useful for proving security of our main construction in Section 4. We will also require that the permutation branching program has a fixed input-selector function inp . In our construction, we will have multiple branching programs, and all of them must read the same input bit at any level $i \leq L$.

Definition 2.3. A permutation branching program with input space $\{0, 1\}^n$ is said to have a fixed input-selector $\text{inp}(\cdot)$ if for all $i \leq L$, $\text{inp}(i) = i \bmod n$.

Any permutation branching program of length L and input space $\{0, 1\}^n$ can be easily transformed to a fixed input-selector branching program of length $n \cdot L$. In this work, we only require that all branching programs share the same input selector function $\text{inp}(\cdot)$. The input selector which satisfies $\text{inp}(i) = i \bmod n$ is just one possibility, and we stick with it for simplicity. We will use the following corollary, which follows from Theorem 2.2.

Corollary 2.2. For any boolean circuit C with input space $\{0, 1\}^n$ and depth d , there exists a *fixed-input selector* permutation branching program BP of width 5 and length $n \cdot 4^d$ such that for all inputs $x \in \{0, 1\}^n$, $C(x) = \text{BP}(x)$.

2.3 Public Key Encryption

A public key encryption (PKE) scheme \mathcal{E} for message spaces $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ consists of the following polynomial-time algorithms.

- $\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$. The setup algorithm takes as input the security parameter λ , and outputs a public-secret key pair (pk, sk) .
- $\text{Enc}(\text{pk}, m \in \mathcal{M}_\lambda) \rightarrow \text{ct}$. The encryption algorithm takes as input a public key pk and a message m , and outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$. The decryption algorithm takes as input a secret key sk and a ciphertext ct , and outputs a message m .

Correctness. A PKE scheme \mathcal{E} for message spaces \mathcal{M} is said to be correct if for all λ , $m \in \mathcal{M}_\lambda$ and $(\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda)$, we have that $\text{Dec}(\text{sk}, \text{Enc}(\text{pk}, m)) = m$.

Security. The standard security notion for PKE schemes is IND-CPA security. Formally, it is defined as follows.

Definition 2.4. A public key encryption scheme $\mathcal{E} = (\text{Setup}, \text{Enc}, \text{Dec})$ is IND-CPA secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible functions $\text{negl}(\cdot)$, such that the following function of λ is bounded by $\text{negl}(\cdot)$

$$\left| \Pr \left[\mathcal{A}(\text{ct}) = b \mid \begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \text{Setup}(1^\lambda); b \leftarrow \{0, 1\} \\ (m_0, m_1) \leftarrow \mathcal{A}(\text{pk}); \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b) \end{array} \right] - \frac{1}{2} \right|$$

2.4 Homomorphic Encryption

Homomorphic encryption [RAD78, Gen09] is a powerful extension of public key encryption that allows one to evaluate functions on ciphertexts. Homomorphic encryption schemes can be classified as either leveled or fully homomorphic encryption schemes. A leveled homomorphic encryption (LHE) scheme allows bounded depth computation over the ciphertexts. The setup algorithm takes as input a ‘level bound’ ℓ together with the security parameter, and outputs a public-secret key pair. Given an ciphertext ct corresponding to message m , one can use the evaluation algorithm to evaluate a bounded depth circuit C on ct , and the resulting ciphertext ct' , when decrypted using the secret key, outputs $C(m)$ if the depth of C is less than ℓ . Fully homomorphic encryption, on the other hand, allows for arbitrary computation on the ciphertext.

2.4.1 Leveled Homomorphic Encryption

A secret key leveled homomorphic encryption scheme \mathcal{HE} with message space $\{0, 1\}$ consists of four algorithms $\text{Setup}, \text{Enc}, \text{Dec}, \text{Eval}$ with the following syntax:

1. $\text{Setup}(1^\lambda, 1^\ell) \rightarrow (\text{sk}, \text{ek})$ The setup algorithm takes as input the security parameter λ , bound on circuit depth ℓ and outputs a secret key sk and evaluation key ek .
2. $\text{Enc}(\text{sk}, m \in \{0, 1\}) \rightarrow \text{ct}$ The encryption algorithm takes as input a secret key sk , message $m \in \{0, 1\}$ and outputs a ciphertext ct .
3. $\text{Eval}(\text{ek}, C \in \mathcal{C}_\ell, \text{ct}) \rightarrow \text{ct}'$ The evaluation algorithm takes as input an evaluation key ek , a circuit $C \in \mathcal{C}_\ell$, a ciphertext ct and outputs a ciphertext ct' .
4. $\text{Dec}(\text{sk}, \text{ct}) \rightarrow x$ The decryption algorithm takes as input a secret key sk and ciphertext ct and outputs $x \in \{0, 1\} \cup \{\perp\}$.

We will now define some properties of leveled homomorphic encryption schemes. Let \mathcal{HE} be any homomorphic encryption scheme with message space $\{0, 1\}$. First, we have the correctness property, which states that the decryption of a homomorphic evaluation on a ciphertext must be equal to the evaluation on the underlying message.

Definition 2.5 (Correctness). The scheme \mathcal{HE} is said to be (perfectly) correct if for all security parameter λ , circuit-depth bound ℓ , $(\text{sk}, \text{ek}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, circuit $C \in \mathcal{C}_\ell$ and message $m \in \{0, 1\}$,

$$\text{Dec}(\text{sk}, \text{Eval}(\text{ek}, C, \text{Enc}(\text{sk}, m))) = C(m).$$

Next, we have the compactness property which requires that the size of the output of an evaluation on a ciphertext must not depend upon the evaluation circuit. In particular, we require that there exists one decryption circuit such that this circuit can decrypt any bounded-depth evaluations on ciphertexts.

Definition 2.6 (Compactness). A homomorphic encryption scheme \mathcal{HE} is said to be compact if for all λ , ℓ there is a decryption circuit $C_{\lambda, \ell}^{\text{Dec}}$ such that for all $(\text{sk}, \text{ek}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, $m \in \{0, 1\}$, $C \in \mathcal{C}_\ell$, $C_{\lambda, \ell}^{\text{Dec}}(\text{sk}, \text{Eval}(\text{ek}, C, \text{Enc}(\text{sk}, m))) = C(m)$.

Finally, we require that the depth of the decryption circuit is bounded by a logarithmic function in the security parameter λ .

Definition 2.7. A compact homomorphic encryption scheme \mathcal{HE} is said to have log-depth decryption circuit if for all λ, ℓ , $\text{depth}(C_{\lambda, \ell}^{\text{Dec}}) = O(\log \lambda)$.

For security, we require that the underlying scheme is IND-CPA secure.

Definition 2.8 (Security). A homomorphic encryption scheme \mathcal{HE} is secure if $\Gamma = (\text{Setup}, \text{Enc}, \text{Dec})$ is IND-CPA secure (Definition 2.4).

Starting with the work of Gentry [Gen09], there has been a long line of interesting works seeking to improve the efficiency/security of homomorphic encryption schemes. Today, we have LHE schemes [BV11, BGV12, GSW13] with log-depth decryption circuits that can be proven secure under the Learning with Errors assumption.

2.4.2 Fully Homomorphic Encryption

Fully homomorphic encryption is a stronger notion of homomorphic encryption where the evaluator can perform unbounded computations on the ciphertext. As a result, the algorithm does not need to take the depth bound as input.

A secret key fully homomorphic encryption scheme \mathcal{HE} with message space $\{0, 1\}$ consists of four algorithms $\text{Setup}, \text{Enc}, \text{Dec}, \text{Eval}$ with the following syntax:

1. $\text{Setup}(1^\lambda) \rightarrow (\text{sk}, \text{ek})$ The setup algorithm takes as input the security parameter λ and outputs a secret key sk and evaluation key ek .
2. $\text{Enc}(\text{sk}, m \in \{0, 1\}) \rightarrow \text{ct}$ The encryption algorithm takes as input a secret key sk , message $m \in \{0, 1\}$ and outputs a ciphertext ct .
3. $\text{Eval}(\text{ek}, C, \text{ct}) \rightarrow \text{ct}'$ The evaluation algorithm takes as input an evaluation key ek , a circuit C , a ciphertext ct and outputs a ciphertext ct' .
4. $\text{Dec}(\text{sk}, \text{ct}) \rightarrow x$ The decryption algorithm takes as input a secret key sk and ciphertext ct and outputs $x \in \{0, 1\} \cup \{\perp\}$.

Definition 2.9 (Correctness). The scheme \mathcal{HE} is said to be (perfectly) correct if for all security parameter λ , $(\text{sk}, \text{ek}) \leftarrow \text{Setup}(1^\lambda, 1^\ell)$, circuit C and message $m \in \{0, 1\}$,

$$\text{Dec}(\text{sk}, \text{Eval}(\text{ek}, C, \text{Enc}(\text{sk}, m))) = C(m).$$

The security and compactness requirement is same as that for leveled homomorphic encryption.

2.5 Pairwise Independent Hash Functions

Definition 2.10. Let \mathcal{H} be a family of hash functions where each $h \in \mathcal{H}$ has domain $\{0, 1\}^\ell$ and co-domain $\{0, 1\}^{\ell_{\mathcal{H}}}$. The family \mathcal{H} is said to be a pairwise independent hash function family if for all $x, y \in \{0, 1\}^\ell$ and $a, b \in \{0, 1\}^{\ell_{\mathcal{H}}}$ s.t. $x \neq y$,

$$\Pr_{h \leftarrow \mathcal{H}} [h(x) = a \wedge h(y) = b] = \frac{1}{(2^{\ell_{\mathcal{H}}})^2}.$$

A simple family of pairwise independent hash functions is the following set of functions:

$$\{h_{\mathbf{A}, \mathbf{b}}(\mathbf{x}) = \mathbf{Ax} + \mathbf{b} : \mathbf{A} \in \mathbb{Z}_2^{\ell_{\mathcal{H}} \times \ell}, \mathbf{b} \in \mathbb{Z}_2^{\ell_{\mathcal{H}}}\}.$$

This family can be implemented using log-depth circuits.

2.6 NC¹ Pseudorandom Generators with Polynomial Stretch

In this work, we require polynomial stretch pseudorandom generators which can be implemented by polynomial length branching programs.

Definition 2.11. Let ℓ, ℓ_{PRG} be polynomials. A circuit class $\text{PRG}_{\ell, \ell_{\text{PRG}}} = \{\text{PRG}_\lambda\}_\lambda$ where $\text{PRG}_\lambda : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{PRG}}(\lambda)}$, is said to be a pseudorandom generator if for all PPT adversaries \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that the following function of λ is bounded by $\text{negl}(\cdot)$:

$$\left| \Pr \left[1 \leftarrow \mathcal{A}(\text{PRG}_\lambda(s)) : s \leftarrow \{0, 1\}^{\ell(\lambda)} \right] - \Pr \left[1 \leftarrow \mathcal{A}(u) : u \leftarrow \{0, 1\}^{\ell_{\text{PRG}}(\lambda)} \right] \right|$$

Using PRFs in NC¹, we can show that there exist PRGs of logarithmic depth which can achieve any polynomial stretch. Low depth PRFs, in turn, can be constructed from the Learning with Errors assumption.

Theorem 2.3 (PRFs in NC¹ [BPR12]). For some $\sigma > 0$, suitable universal constant $C > 0$, modulus $p \geq 2$, any $m = \text{poly}(n)$, let $\chi = \mathcal{D}_{\mathbb{Z}, \sigma}$ and $q \geq p \cdot k(C\sigma\sqrt{n})^k \cdot n^{\omega(1)}$, assuming hardness of (n, m, q, χ) -LWE, there exists a function family PRF consisting of functions from $\{0, 1\}^k$ to $\mathbb{Z}_p^{m \times n}$ that satisfies pseudorandomness property and the entire function can be computed in $\text{TC}^0 \subseteq \text{NC}^1$.

From Theorems 2.2 and 2.3, the following corollary is immediate.

Corollary 2.3. There exists a polynomial p such that for any polynomial ℓ_{PRG} , there exists a family of pseudorandom generators $\{\text{PRG}_\lambda\}$ where $\text{PRG}_\lambda : \{0, 1\}^{\ell(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{PRG}}(\lambda)}$, and for all $i \in [\ell_{\text{PRG}}(\lambda)]$, the i^{th} bit of PRG_λ can be computed by a branching program of length $p(\lambda)$. This family of pseudorandom generators can be proven secure assuming hardness of (n, m, q, χ) -LWE with parameters as in Theorem 2.3.

3 Lockable Obfuscation

In this section, we define the notion of lockable obfuscation, and discuss how it is related to other cryptographic primitives. An lockable obfuscator takes as input a program P , a message msg , and a ‘lock’ α . It outputs an obfuscated program \tilde{P} which has the same domain as the program P . The program \tilde{P} , on input x , outputs the message msg if $P(x) = \alpha$. If not, then with overwhelming probability, it will output \perp . For security, we require that for all programs P and messages msg , \tilde{P} for a uniformly random α can be efficiently simulated given only the sizes of P and msg . We will now present the syntax, correctness and security definition.

Let n, m, d be polynomials, and $\mathcal{C}_{n, m, d}(\lambda)$ be the class of depth $d(\lambda)$ circuits with $n(\lambda)$ bit input and $m(\lambda)$ bit output. A lockable obfuscator for $\mathcal{C}_{n, m, d}$ consists of algorithms Obf and Eval with the following syntax. Let \mathcal{M} be the message space.

- $\text{Obf}(1^\lambda, P, \text{msg}, \alpha) \rightarrow \tilde{P}$. The obfuscation algorithm is a randomized algorithm that takes as input the security parameter λ , a program $P \in \mathcal{C}_{n, m, d}$, message $\text{msg} \in \mathcal{M}$ and ‘lock string’ $\alpha \in \{0, 1\}^{m(\lambda)}$. It outputs a program \tilde{P} .
- $\text{Eval}(\tilde{P}, x) \rightarrow y \in \mathcal{M} \cup \{\perp\}$. The evaluator is a deterministic algorithm that takes as input a program \tilde{P} and a string $x \in \{0, 1\}^{n(\lambda)}$. It outputs $y \in \mathcal{M} \cup \{\perp\}$.

3.1 Correctness

For correctness, we require that if $P(x) = \alpha$, then the obfuscated program $\tilde{P} \leftarrow \text{Obf}(1^\lambda, P, \text{msg}, \alpha)$, evaluated on input x , outputs msg , and if $P(x) \neq \alpha$, then \tilde{P} outputs \perp on input x . Here, we give three correctness definitions, which differ in the case where $P(x) \neq \alpha$.

Definition 3.1 (Perfect Correctness). Let n, m, d be polynomials. A lockable obfuscation scheme for $\mathcal{C}_{n, m, d}$ and message space \mathcal{M} is said to be perfectly correct if it satisfies the following properties:

1. For all security parameters λ , inputs $x \in \{0, 1\}^{n(\lambda)}$, programs $P \in \mathcal{C}_{n,m,d}$ and messages $\text{msg} \in \mathcal{M}$, if $P(x) = \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}.$$

2. For all security parameters λ , inputs $x \in \{0, 1\}^{n(\lambda)}$, programs $P \in \mathcal{C}_{n,m,d}$ and messages $\text{msg} \in \mathcal{M}$, if $P(x) \neq \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \perp.$$

Definition 3.2 (Statistical Correctness). An lockable obfuscation scheme is said to be statistically correct if it satisfies the following properties:

1. For all security parameters λ , inputs $x \in \{0, 1\}^{n(\lambda)}$, programs $P \in \mathcal{C}_{n,m,d}$ and messages $\text{msg} \in \mathcal{M}$, if $P(x) = \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}.$$

2. For all security parameters λ , programs $P \in \mathcal{C}_{n,m,d}$, $\alpha \in \{0, 1\}^{m(\lambda)}$ and $\text{msg} \in \mathcal{M}$,

$$\Pr[\exists x \text{ s.t. } P(x) \neq \alpha \text{ and } \text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}] \leq \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Definition 3.3 (Semi-statistical Correctness). Let n, m, d be polynomials. A lockable obfuscation scheme for $\mathcal{C}_{n,m,d}$ and message space \mathcal{M} is said to be semi-statistical correct if it satisfies the following properties:

1. For all security parameters λ , inputs $x \in \{0, 1\}^{n(\lambda)}$, programs $P \in \mathcal{C}_{n,m,d}$ and messages $\text{msg} \in \mathcal{M}$, if $P(x) = \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}.$$

2. For all security parameters λ , inputs $x \in \{0, 1\}^{n(\lambda)}$, programs $P \in \mathcal{C}_{n,m,d}$, $\alpha \in \{0, 1\}^{m(\lambda)}$ such that $P(x) \neq \alpha$ and $\text{msg} \in \mathcal{M}$,

$$\Pr[\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}] \leq \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Looking ahead, our construction satisfies the semi-statistical correctness condition, and will be sufficient for our applications.

3.2 Security

Definition 3.4. Let n, m, d be polynomials. A lockable obfuscation scheme $(\text{Obf}, \text{Eval})$ for $\mathcal{C}_{n,m,d}$ and message space \mathcal{M} is said to be secure if there exists a PPT simulator Sim such that for all PPT adversaries $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, there exists a negligible function $\text{negl}(\cdot)$ such that the following function is bounded by $\text{negl}(\cdot)$:

$$\left| \Pr \left[\mathcal{A}_1(\tilde{P}_b, \text{st}) = b : \begin{array}{l} (P \in \mathcal{C}_{n,m,d}, \text{msg} \in \mathcal{M}, \text{st}) \leftarrow \mathcal{A}_0(1^\lambda) \\ b \leftarrow \{0, 1\}, \alpha \leftarrow \{0, 1\}^{m(\lambda)} \\ \tilde{P}_0 \leftarrow \text{Obf}(1^\lambda, P, \text{msg}, \alpha) \\ \tilde{P}_1 \leftarrow \text{Sim}(1^\lambda, 1^{|P|}, 1^{|\text{msg}|}) \end{array} \right] - \frac{1}{2} \right|$$

4 Our Construction

In this section, we present our lockable obfuscation scheme. For any polynomials $\ell_{\text{in}}, \ell_{\text{out}}, d$ such that $\ell_{\text{out}} = \omega(\log \lambda)$, we construct a lockable obfuscation scheme $\mathcal{O} = (\text{Obf}, \text{Eval})$ for the circuit class $\mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$. The message space for our construction will be $\{0, 1\}$, although one can trivially extend it to $\{0, 1\}^{\ell(\lambda)}$ for any polynomial ℓ . More details are provided in Appendix C.

The tools required for our construction are as follows:

- A compact leveled homomorphic bit encryption scheme ($\text{LHE.Setup}, \text{LHE.Enc}, \text{LHE.Eval}, \text{LHE.Dec}$) with decryption circuit of depth $d_{\text{Dec}}(\lambda)$ and ciphertexts of length $\ell_{\text{ct}}(\lambda)$.
- A pairwise independent hash family $\mathcal{H} : \{0, 1\}^{\ell_{\text{out}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\mathcal{H}}(\lambda)}$ which can be implemented by a circuit of logarithmic depth $d_{\mathcal{H}}(\lambda)$.
- A pseudorandom generator $\text{PRG} : \{0, 1\}^{\ell_{\mathcal{H}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{PRG}}(\lambda)}$ of depth $d_{\text{PRG}}(\lambda)$.

For notational convenience, let $\ell_{\text{in}} = \ell_{\text{in}}(\lambda)$, $\ell_{\text{out}} = \ell_{\text{out}}(\lambda)$, $\ell_{\mathcal{H}} = \ell_{\mathcal{H}}(\lambda)$, $\ell_{\text{PRG}} = \ell_{\text{PRG}}(\lambda)$, $d_{\text{Dec}} = d_{\text{Dec}}(\lambda)$, $d_{\mathcal{H}} = d_{\mathcal{H}}(\lambda)$, $d_{\text{PRG}} = d_{\text{PRG}}(\lambda)$ and $d = d(\lambda)$.

Fix any $\epsilon < 1/2$. Let χ be a B -bounded discrete Gaussian distribution with parameter σ such that $B = \sqrt{m} \cdot \sigma$. Let $n, m, \ell, \sigma, q, \text{Bd}$ be parameters with the following constraints:

- $n = \text{poly}(\lambda)$ and $q \leq 2^{n^\epsilon}$ (for LWE security)
- $\ell_{\mathcal{H}} = \omega(\log \lambda)$ (for PRG security)
- $\ell_{\text{out}} - \ell_{\mathcal{H}} = \omega(\log \lambda)$ (for strong extraction)
- $m = \Omega(n \cdot \log q)$ (for SamplePre)
- $\sigma = \omega(\sqrt{n \cdot \log q \cdot \log m})$ (for Preimage Well Distributedness)
- $\ell_{\text{PRG}} = n \cdot m \cdot \log q + \omega(\log n)$ (for applying Leftover Hash Lemma)
- $\ell_{\text{PRG}} - \ell_{\mathcal{H}} = \omega(\log \lambda)$ (for correctness of scheme)
- $\text{Bd} = \ell_{\text{PRG}} \cdot L \cdot (m^2 \cdot \sigma)^L < q^{1/4}$ (for correctness of scheme)
(where $L = \ell_{\text{out}} \cdot \ell_{\text{ct}} \cdot 4^{d_{\text{Dec}} + d_{\mathcal{H}} + d_{\text{PRG}}}$)

First, note that it is important that $L = \lambda^c$ for some constant c and $\text{Bd} = \ell_{\text{PRG}} \cdot L \cdot (m \cdot \sigma)^L < q^{1/4}$. This crucially relies on the fact that the LHE scheme is compact (so that ℓ_{ct} is bounded by a polynomial) and that the LHE decryption, hash function evaluation and PRG computation can be performed by a log depth circuit (i.e, have poly length branching programs). The constant c depends on the LHE scheme, pairwise independent hash function and PRG.

One possible setting of parameters is as follows: $n = \lambda^{4c/\epsilon}$, $m = n^{1+\epsilon}$, $q = 2^{n^\epsilon}$, $\sigma = n$, $\ell_{\text{PRG}} = n^{2\epsilon+2}$ and $\ell_{\mathcal{H}} = \ell_{\text{out}}/2$.

We will now describe the obfuscation and evaluation algorithms.

- $\text{Obf}(1^\lambda, P, \text{msg}, \alpha)$: The obfuscation algorithm takes as input a program $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$, message $\text{msg} \in \{0, 1\}$ and $\alpha \in \{0, 1\}^{\ell_{\text{out}}}$. The obfuscator proceeds as follows:
 1. First, it chooses the LHE key pair as $(\text{lhe.sk}, \text{lhe.ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$.⁹
 2. Next, it encrypts the program P . It sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{lhe.sk}, P)$.¹⁰
 3. It then chooses a pairwise independent hash function $h \leftarrow \mathcal{H}$. Let $\beta = \text{PRG}(h(\alpha))$.

⁹We set the LHE depth bound to be $d \log d$, where the extra log factor is to account for the constant blowup involved in using a universal circuit. In particular, we can set the LHE depth bound to be $c \cdot d$ where c is some fixed constant depending on the universal circuit.

¹⁰Note that LHE scheme supports bit encryption. Therefore, to encrypt P , a multi-bit message, the FHE.Enc algorithm will be run independently on each bit of P . However, for notational convenience throughout this section we overload the notation and use FHE.Enc and FHE.Dec algorithms to encrypt and decrypt multi-bit messages respectively.

4. Next, consider the following circuit Q which takes as input $\ell_{\text{out}} \cdot \ell_{\text{ct}}$ bits of input and outputs ℓ_{PRG} bits. Q takes as input ℓ_{out} LHE ciphertexts $\{\text{ct}_i\}_{i \leq \ell_{\text{out}}}$, has LHE secret key lhe.sk hardwired and computes the following — (1) it decrypts each input ciphertext ct_i (in parallel) to get string x of length ℓ_{out} bits, (2) it sequentially applies the hash function and PRG on x and outputs $\text{PRG}(h(x))$. Concretely, $Q(\text{ct}_1, \dots, \text{ct}_{\ell_{\text{out}}}) = \text{PRG}(h(\text{LHE.Dec}(\text{lhe.sk}, \text{ct}_1) \parallel \dots \parallel \text{LHE.Dec}(\text{lhe.sk}, \text{ct}_{\ell_{\text{out}}}))$.

For $i \leq \ell_{\text{PRG}}$, we use $\text{BP}^{(i)}$ to denote the fixed-input selector permutation branching program that outputs the i^{th} bit of output of circuit Q . Note that Q has depth $d_{\text{tot}} = d_{\text{Dec}} + d_{\mathcal{H}} + d_{\text{PRG}}$. By Corollary 2.2, we know that each branching program $\text{BP}^{(i)}$ has length $L = \ell_{\text{out}} \cdot \ell_{\text{ct}} \cdot 4^{d_{\text{tot}}}$ and width 5.

5. Finally, the obfuscator creates matrix components which enable the evaluator to compute msg if it has an input strings (ciphertexts) $\text{ct}_1, \dots, \text{ct}_{\ell_{\text{out}}}$ such that $Q(\text{ct}_1, \dots, \text{ct}_{\ell_{\text{out}}}) = \beta$. Concretely, it runs the (randomized) routine **Comp-Gen** (defined in Figure 1). This routine takes as input the circuit Q in the form of ℓ_{PRG} branching programs $\{\text{BP}^{(i)}\}_i$, string β and message msg . Let $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ \mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)} \right\}_{i,j} \right) \leftarrow \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})$.

6. The final obfuscated program consists of the LHE evaluation key $\text{ek} = \text{lhe.ek}$, LHE ciphertexts \mathbf{ct} , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

• $\text{Eval}(\tilde{P}, x)$: The evaluation algorithm takes as input $\tilde{P} = \left(\text{ek}, \mathbf{ct}, \left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$ and input $x \in \{0, 1\}^{\ell_{\text{in}}}$. It performs the following steps.

1. The evaluator first constructs a universal circuit $U_x(\cdot)$ with x hardwired as input. This universal circuit takes a circuit C as input and outputs $U_x(C) = C(x)$. Using the universal circuit of Cook and Hoover [CH85], it follows that $U_x(\cdot)$ has depth $O(d)$.
2. Next, it performs homomorphic evaluation on \mathbf{ct} using circuit $U_x(\cdot)$. It computes $\tilde{\mathbf{ct}} = \text{LHE.Eval}(\text{ek}, U_x(\cdot), \mathbf{ct})$. Note that $\ell_{\text{ct}} \cdot \ell_{\text{out}}$ denotes the length of \mathbf{ct} (as a bitstring), and let $\tilde{\mathbf{ct}}_i$ denote the i^{th} bit of $\tilde{\mathbf{ct}}$.
3. The evaluator then obviously evaluates the ℓ_{PRG} branching programs on input $\tilde{\mathbf{ct}}$ using the matrix components. It calls the component evaluation algorithm **Comp-Eval** (defined in Figure 2). Let $y = \text{Comp-Eval} \left(\tilde{\mathbf{ct}}, \left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right) \right)$. The evaluator outputs y .

4.1 Correctness

We will prove that the lockable obfuscation scheme described above satisfies the semi-statistical correctness property (see Definition 3.3). To prove this, we need to prove that if $P(x) = \alpha$, then the evaluation algorithm always outputs the message, and if $P(x) \neq \alpha$, then with high probability, it outputs \perp .

First, we will prove the following lemma about the **Comp-Gen** and **Comp-Eval** routines. Intuitively, this lemma states that for all fixed input branching programs $\{\text{BP}^{(i)}\}_i$, strings β and messages msg , if $\text{BP}^{(i)}(z) = \beta_i$ for all $i \leq \ell_{\text{PRG}}$, then the component evaluator outputs msg .

Lemma 4.1. For any set of branching programs $\{\text{BP}^{(i)}\}_{i \leq \ell_{\text{PRG}}}$, string $\beta \in \{0, 1\}^{\ell_{\text{PRG}}}$, message $\text{msg} \in \{0, 1\}$ and input z ,

1. if $\text{BP}^{(i)}(z) = \beta_i$ for all $i \leq \ell_{\text{PRG}}$, then $\text{Comp-Eval}(z, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) = \text{msg}$.
2. if $\text{BP}^{(i)}(z) \neq \beta_i$ for some $i \leq \ell_{\text{PRG}}$, then

$$\Pr[\text{Comp-Eval}(z, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) = \perp] \geq 1 - \text{negl}(\lambda).$$

Comp-Gen

Input: $\{\text{BP}^{(i)}\}_i, \beta \in \{0, 1\}^{\ell_{\text{PRG}}}, \text{msg} \in \{0, 1\}$

Output: Components $\left(\left\{\mathbf{B}_{0,1}^{(i)}\right\}_i, \left\{(\mathbf{C}_{\text{level}}^{(i,0)}, \mathbf{C}_{\text{level}}^{(i,1)})\right\}_{i \leq \ell_{\text{PRG}}, \text{level} \leq L}\right)$.

(a) Let $\text{BP}^{(i)} = \left(\left\{\sigma_{j,b}^{(i)} : [5] \rightarrow [5]\right\}_{j \in [L], b \in \{0,1\}}, \text{acc}^{(i)} \in [5], \text{rej}^{(i)} \in [5]\right)$ for all $i \leq \ell_{\text{PRG}}$.

(b) First, it chooses a matrix for each state of each branching program. Recall, there are ℓ_{PRG} branching programs, and each branching program has L levels, and each level has 5 states. For each $i \leq \ell_{\text{PRG}}$, $j \in [0, L-1]$, it chooses a matrix of dimensions $5n \times m$ along with its trapdoors (independently) as $(\mathbf{B}_j^{(i)}, T_j^{(i)}) \leftarrow \text{TrapGen}(1^{5n}, 1^m, q)$. The matrix $\mathbf{B}_j^{(i)}$ can be parsed as follows

$$\mathbf{B}_j^{(i)} = \begin{bmatrix} \mathbf{B}_{j,1}^{(i)} \\ \vdots \\ \mathbf{B}_{j,5}^{(i)} \end{bmatrix}$$

where matrices $\mathbf{B}_{j,k}^{(i)} \in \mathbb{Z}_q^{n \times m}$ for $k \leq 5$. The matrix $\mathbf{B}_{j,k}^{(i)}$ corresponds to state k at level j of branching program $\text{BP}^{(i)}$.

(c) For the top level, it chooses top level matrices $\mathbf{B}_{L,k}^{(i)}$ for each $i \leq \ell_{\text{PRG}}, k \leq 5$, uniformly at random, subject to the following constraint:

$$\sum_{i : \beta_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i : \beta_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \mathbf{0}^{n \times m} \text{ if msg} = 0.$$

$$\sum_{i : \beta_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i : \beta_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] \text{ if msg} = 1.$$

(d) Next, it generates the components for each level. For each level $\text{level} \in [1, L]$, do the following:

- i. Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$. If either $\mathbf{S}_{\text{level}}^{(0)}$ or $\mathbf{S}_{\text{level}}^{(1)}$ has determinant zero, then set it to be \mathbf{I}_n .
- ii. For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$. More formally, for $i \leq \ell_{\text{PRG}}$, set

$$\mathbf{D}_{\text{level}}^{(i,b)} = \begin{bmatrix} \mathbf{B}_{\text{level},\sigma_{\text{level},b}^{(i)}(1)}^{(i)} \\ \vdots \\ \mathbf{B}_{\text{level},\sigma_{\text{level},b}^{(i)}(5)}^{(i)} \end{bmatrix}.$$

iii. Set $\mathbf{M}_{\text{level}}^{(i,b)} = (\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)}) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.

iv. Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

(e) Output $\left(\left\{\mathbf{B}_{0,1}^{(i)}\right\}_i, \left\{(\mathbf{C}_{\text{level}}^{(i,0)}, \mathbf{C}_{\text{level}}^{(i,1)})\right\}_{i \leq \ell_{\text{PRG}}, \text{level} \leq L}\right)$.

Figure 1: Routine Comp-Gen

Proof. The proof of this lemma will be similar to the proof of correctness of the testing algorithm in [GKW17b, Section 4.2].

Recall that the component generation algorithm chooses matrices $\mathbf{B}_j^{(i)}$ for each $i \leq \ell_{\text{PRG}}, j \leq L$, $\mathbf{S}_j^{(0)}, \mathbf{S}_j^{(1)}$ for each $j \leq L$ and $\mathbf{E}_j^{(i,0)}, \mathbf{E}_j^{(i,1)}$ for each $i \leq \ell_{\text{PRG}}, j \leq L$. Note that the $\mathbf{S}_j^{(b)}$ and $\mathbf{E}_j^{(i,b)}$ matrices have l_∞ norm bounded by $\sigma \cdot m$ since they are chosen from truncated Gaussian distribution with parameter σ .

We start by introducing some notations for this proof.

- $\text{st}_j^{(i)}$: the state of $\text{BP}^{(i)}$ after j steps when evaluated on z

Comp-Eval

Input: Input string z , Components $\left(\left\{\mathbf{B}_{0,1}^{(i)}\right\}_i, \left\{\left(\mathbf{C}_{\text{level}}^{(i,0)}, \mathbf{C}_{\text{level}}^{(i,1)}\right)\right\}_{i \leq \ell_{\text{PRG}}, \text{level} \leq L}\right)$.

Output: $y \in \{0, 1, \perp\}$.

- (a) For each $i \in [1, \ell_{\text{PRG}}]$, do the following
 - i. Set $\mathbf{M}^{(i)} = \mathbf{B}_{0,1}^{(i)}$.
 - ii. For $j = 1$ to L , do the following
 - If $z_{\text{inp}(j)} = 0$, set $\mathbf{M}^{(i)} = \mathbf{M}^{(i)} \cdot \mathbf{C}_j^{(i,0)}$. Else, set $\mathbf{M}^{(i)} = \mathbf{M}^{(i)} \cdot \mathbf{C}_j^{(i,1)}$.
- (b) Compute $\mathbf{M} = \sum_i \mathbf{M}^{(i)}$. Let $\mathbf{M} = \left[\mathbf{M}^{(1)} \parallel \mathbf{M}^{(2)}\right]$ where $\mathbf{M}^{(1)} \in \mathbb{Z}_q^{n \times n}$ (i.e. a square matrix). Next, do the following
 - If $\|\mathbf{M}\|_\infty \leq \text{Bd}$, output 0.
 - Otherwise, if $\left\|\mathbf{M}^{(2)}\right\|_\infty \leq \text{Bd}$ and $\|\mathbf{M}\|_\infty \in [\sqrt{q} - \text{Bd}, (\sqrt{q} + 1) \cdot \text{Bd}]$, output 1.
 - Else output \perp .

Figure 2: Routine Comp-Eval

- $\mathbf{S}_j = \mathbf{S}_j^{(z_{\text{inp}(j)})}$, $\mathbf{E}_j^{(i)} = \mathbf{E}_j^{(i, z_{\text{inp}(j)})}$, $\mathbf{C}_j^{(i)} = \mathbf{C}_j^{(i, z_{\text{inp}(j)})}$ for all $j \leq L$
- $\mathbf{\Gamma}_{j^*} = \prod_{j=1}^{j^*} \mathbf{S}_j$ for all $j^* \leq L$
- $\mathbf{\Delta}_{j^*}^{(i)} = \mathbf{B}_{0,1}^{(i)} \cdot \left(\prod_{j=1}^{j^*} \mathbf{C}_j^{(i)}\right)$, $\tilde{\mathbf{\Delta}}_{j^*}^{(i)} = \mathbf{\Gamma}_{j^*} \cdot \mathbf{B}_{j^*, \text{st}_{j^*}}^{(i)}$, $\mathbf{Err}_{j^*}^{(i)} = \mathbf{\Delta}_{j^*}^{(i)} - \tilde{\mathbf{\Delta}}_{j^*}^{(i)}$ for all $j^* \leq L$

Observe that the Comp-Eval algorithm computes matrix $\mathbf{M} = \sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{\Delta}_L^{(i)}$. First, we show that for all $i \leq \ell_{\text{PRG}}$, $j^* \leq L$, $\mathbf{Err}_{j^*}^{(i)}$ is small and bounded. This would help us in arguing that matrices $\sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{\Delta}_L^{(i)}$ and $\sum_{i=1}^{\ell_{\text{PRG}}} \tilde{\mathbf{\Delta}}_L^{(i)}$ are very close to each other.

Now, since we already know that $\sum_{i=1}^{\ell_{\text{PRG}}} \tilde{\mathbf{\Delta}}_L^{(i)}$ is close to all zeros matrix if $\text{msg} = 0$ and is close to $\sqrt{q} \cdot \left[\mathbf{\Gamma}_L \parallel \mathbf{0}^{n \times (m-n)}\right]$ if $\text{msg} = 1$, therefore combining this with the fact that $\mathbf{Err}_{j^*}^{(i)}$ is small and bounded, we can conclude by arguing that matrix \mathbf{M} will have appropriately bounded entries depending on the value of msg .

First, we show that $\mathbf{Err}_{j^*}^{(i)}$ is bounded.

Claim 4.1. $\forall i \in \{1, \dots, \ell_{\text{PRG}}\}, j^* \in \{1, \dots, L\}$, $\left\|\mathbf{Err}_{j^*}^{(i)}\right\|_\infty \leq j^* \cdot (m^2 \cdot \sigma)^{j^*}$.

Proof. The above claim is proven by induction over j^* , and all arguments hold irrespective of the value of i . Therefore, for simplicity of notation, we will drop the dependence on i .

Base case ($j^* = 1$). We know that $\mathbf{\Delta}_1 = \mathbf{B}_{0,1} \cdot \mathbf{C}_1 = \mathbf{S}_1 \cdot \mathbf{B}_{1, \text{st}_1} + \mathbf{E}_{1, \text{st}_1}$ where $\mathbf{E}_{1, \text{st}_1}$ is a submatrix of \mathbf{E}_1 . Thus, we could write the following

$$\|\mathbf{Err}_1\|_\infty = \left\|\mathbf{\Delta}_1 - \tilde{\mathbf{\Delta}}_1\right\|_\infty = \|\mathbf{E}_{1, \text{st}_1}\|_\infty \leq m \cdot \sigma < m^2 \cdot \sigma.$$

This completes the proof of base case. For the induction step, we assume that the above lemma holds for $j^* - 1$, and show that it holds for j^* as well.

Induction Step. We know that $\Delta_{j^*} = \Delta_{j^*-1} \cdot \mathbf{C}_{j^*}$. Also, $\Delta_{j^*-1} = \tilde{\Delta}_{j^*-1} + \mathbf{Err}_{j^*-1}$. So, we could write the following

$$\begin{aligned} \Delta_{j^*} &= \tilde{\Delta}_{j^*-1} \cdot \mathbf{C}_{j^*} + \mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*} \\ &= \Gamma_{j^*-1} \cdot (\mathbf{B}_{j^*-1, \text{st}_{j^*-1}} \cdot \mathbf{C}_{j^*}) + \mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*} \\ &= \Gamma_{j^*-1} \cdot (\mathbf{S}_{j^*} \cdot \mathbf{B}_{j^*, \text{st}_{j^*}} + \mathbf{E}_{j^*, \text{st}_{j^*-1}}) + \mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*} \\ &= \tilde{\Delta}_{j^*} + \Gamma_{j^*-1} \cdot \mathbf{E}_{j^*, \text{st}_{j^*-1}} + \mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*} \end{aligned}$$

Here, $\mathbf{E}_{j^*, \text{st}_{j^*-1}}$ is an $n \times m$ submatrix of \mathbf{E}_{j^*} . Finally, we can bound \mathbf{Err}_{j^*} as follows

$$\begin{aligned} \|\mathbf{Err}_{j^*}\|_\infty &= \|\Delta_{j^*} - \tilde{\Delta}_{j^*}\|_\infty = \|\Gamma_{j^*-1} \cdot \mathbf{E}_{j^*, \text{st}_{j^*-1}} + \mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*}\|_\infty \\ &\leq \|\Gamma_{j^*-1} \cdot \mathbf{E}_{j^*, \text{st}_{j^*-1}}\|_\infty + \|\mathbf{Err}_{j^*-1} \cdot \mathbf{C}_{j^*}\|_\infty \\ &\leq (n^2 \cdot \sigma)^{j^*-1} \cdot m \cdot \sigma + (j^* - 1) \cdot (m^2 \cdot \sigma)^{j^*-1} \cdot m^2 \cdot \sigma \leq j^* \cdot (m^2 \cdot \sigma)^{j^*} \end{aligned}$$

This completes the proof of above claim. ■

The remaining proof of the lemma will have two parts, the first part dealing with the case when $\text{BP}^{(i)}(z) = \beta_i$ for all i , and the second one dealing with when $\text{BP}^{(i)}(z) \neq \beta_i$ for some i .

Recall that the Comp-Eval algorithm computes matrix $\mathbf{M} = \sum_{i=1}^{\ell_{\text{PRG}}} \Delta_L^{(i)}$. Let $\tilde{\mathbf{M}} = \sum_{i=1}^{\ell_{\text{PRG}}} \tilde{\Delta}_L^{(i)}$ and $\mathbf{Err} = \sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{Err}_L^{(i)}$. Also, we parse these matrices as $\mathbf{M} = [\mathbf{M}^{(1)} \parallel \mathbf{M}^{(2)}]$, $\tilde{\mathbf{M}} = [\tilde{\mathbf{M}}^{(1)} \parallel \tilde{\mathbf{M}}^{(2)}]$ and $\mathbf{Err} = [\mathbf{Err}^{(1)} \parallel \mathbf{Err}^{(2)}]$, where $\mathbf{M}^{(1)}$, $\tilde{\mathbf{M}}^{(1)}$ and $\mathbf{Err}^{(1)}$ are $n \times n$ (square) matrices.

First, note that $\mathbf{M} = \tilde{\mathbf{M}} + \mathbf{Err}$. Also, using Claim 4.1, we can write that

$$\|\mathbf{Err}\|_\infty = \left\| \sum_{i=1}^{\ell_{\text{PRG}}} \left(\Delta_L^{(i)} - \tilde{\Delta}_L^{(i)} \right) \right\|_\infty \leq \sum_{i=1}^{\ell_{\text{PRG}}} \left\| \Delta_L^{(i)} - \tilde{\Delta}_L^{(i)} \right\|_\infty \leq \ell_{\text{PRG}} \cdot L \cdot (m^2 \cdot \sigma)^L = \text{Bd}.$$

Next, consider the following scenarios.

Part 1: $\text{BP}^{(i)}(z) = \beta_i$ for all $i \leq \ell_{\text{PRG}}$. First, recall that the top level matrices always satisfy the following constraints during honest obfuscation:

$$\sum_{i: \beta_i=0} \mathbf{B}_{L, \text{rej}^{(i)}}^{(i)} + \sum_{i: \beta_i=1} \mathbf{B}_{L, \text{acc}^{(i)}}^{(i)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if msg} = 0 \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

Since $\text{BP}^{(i)}(z) = \beta_i$ for all $i \leq \ell_{\text{PRG}}$, we can alternatively write it as

$$\sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{B}_{L, \text{st}_L^{(i)}}^{(i)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if msg} = 0 \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

Note that

$$\tilde{\mathbf{M}} = \sum_{i=1}^{\ell_{\text{PRG}}} \tilde{\Delta}_L^{(i)} = \sum_{i=1}^{\ell_{\text{PRG}}} \Gamma_L \cdot \mathbf{B}_{L, \text{st}_L^{(i)}}^{(i)} = \Gamma_L \cdot \sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{B}_{L, \text{st}_L^{(i)}}^{(i)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if msg} = 0 \\ \sqrt{q} \cdot [\Gamma_L \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

Next, we consider the following two cases depending upon the message being obfuscated — (1) $\text{msg} = 0$, (2) $\text{msg} = 1$.

Case 1 ($\text{msg} = 0$). We already know that $\|\text{Err}\|_\infty \leq \text{Bd}$. Now since in this case $\widetilde{\mathbf{M}} = \mathbf{0}^{n \times m}$, we can say that

$$\|\mathbf{M}\|_\infty = \left\| \widetilde{\mathbf{M}} + \text{Err} \right\|_\infty = \|\text{Err}\|_\infty \leq \text{Bd}.$$

Thus, matrix \mathbf{M} (computed during evaluation) always satisfies the condition that $\|\mathbf{M}\|_\infty \leq \text{Bd}$ if $\text{msg} = 0$.

Case 2 ($\text{msg} = 1$). First, note that during obfuscation we sample secret matrices $\mathbf{S}_{\text{level}}^{(b)}$ (for each level and bit b) such that they are short and *always* invertible. Therefore, matrix $\mathbf{\Gamma}_L$ (which is product of L secret matrices) is also invertible. Thus, we can write that

$$\|\mathbf{\Gamma}_L\|_\infty \geq 1, \quad \|\mathbf{\Gamma}_L\|_\infty = \left\| \prod_{j=1}^L \mathbf{S}_j \right\|_\infty \leq n^L \prod_{j=1}^L \|\mathbf{S}_j\|_\infty \leq (m^2 \cdot \sigma)^L < \text{Bd}.$$

The lower bound of 1 follows from the fact that $\mathbf{\Gamma}_L$ is non-singular (and integral) matrix, and upper bound follows from the fact that matrices $\mathbf{S}_{\text{level}}^{(b)}$ are sampled from bounded Gaussian distribution.

Since $\text{msg} = 1$, we know that $\widetilde{\mathbf{M}}^{(1)} = \sqrt{q} \cdot \mathbf{\Gamma}_L$ and $\widetilde{\mathbf{M}}^{(2)} = \mathbf{0}^{n \times (m-n)}$. Using all these facts, we can show that if $\text{msg} = 1$, then the checks during component evaluation always succeed. More formally, first we show that $\mathbf{M}^{(2)}$ is a short matrix.

$$\left\| \mathbf{M}^{(2)} \right\|_\infty = \left\| \widetilde{\mathbf{M}}^{(2)} + \text{Err}^{(2)} \right\|_\infty = \left\| \text{Err}^{(2)} \right\|_\infty \leq \|\text{Err}\|_\infty \leq \text{Bd}.$$

Next, we show that matrix $\mathbf{M}^{(1)}$ has large entries. In other words, matrix \mathbf{M} has high l_∞ norm. Concretely,

$$\|\mathbf{M}\|_\infty = \left\| \widetilde{\mathbf{M}} + \text{Err} \right\|_\infty \leq \left\| \widetilde{\mathbf{M}} \right\|_\infty + \|\text{Err}\|_\infty = \left\| \widetilde{\mathbf{M}}^{(1)} \right\|_\infty + \|\text{Err}\|_\infty = \|\sqrt{q} \cdot \mathbf{\Gamma}_L\|_\infty + \|\text{Err}\|_\infty \leq \sqrt{q} \cdot \text{Bd} + \text{Bd}.$$

$$\|\mathbf{M}\|_\infty = \left\| \widetilde{\mathbf{M}} + \text{Err} \right\|_\infty \geq \left\| \widetilde{\mathbf{M}} \right\|_\infty - \|\text{Err}\|_\infty \geq \left\| \widetilde{\mathbf{M}}^{(1)} \right\|_\infty - \|\text{Err}\|_\infty = \|\sqrt{q} \cdot \mathbf{\Gamma}_L\|_\infty - \|\text{Err}\|_\infty \geq \sqrt{q} - \text{Bd}.$$

Therefore, $\|\mathbf{M}\|_\infty \in [\sqrt{q} - \text{Bd}, (\sqrt{q} + 1) \cdot \text{Bd}]$. Thus, if $\text{msg} = 1$ then the evaluation always outputs 1.

Part 2: $\text{BP}^{(i)}(z) \neq \beta_i$ for some $i \leq \ell_{\text{PRG}}$. Let k be the first such index. Then $\mathbf{B}_{L, \text{st}_L}^{(k)}$ is a uniformly random $n \times m$ matrix, independent of the other top level matrices. As a result, $\sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{B}_{L, \text{st}_L}^{(i)}$ is a uniformly random matrix. As we noted before $\mathbf{\Gamma}_L$ is an invertible matrix, therefore $\widetilde{\mathbf{M}} = \sum_{i=1}^{\ell_{\text{PRG}}} \widetilde{\mathbf{\Delta}}_L^{(i)} = \mathbf{\Gamma}_L \cdot \sum_{i=1}^{\ell_{\text{PRG}}} \mathbf{B}_{L, \text{st}_L}^{(i)}$ is a uniformly random matrix. This also implies that sub-matrices $\widetilde{\mathbf{M}}^{(1)}$ and $\widetilde{\mathbf{M}}^{(2)}$ are random matrices as well.

Therefore, with all but negligible probability, both $\|\mathbf{M}\|_\infty$ as well as $\left\| \mathbf{M}^{(2)} \right\|_\infty$ will not be small (i.e., less than Bd). Thus, in this scenario, the evaluator will output \perp (with all but negligible probability).

This concludes the proof of Lemma 4.1. ■

Using the above lemma, we can now argue the correctness of our scheme. First, we need to show correctness for the case when $P(x) = \alpha$.

Claim 4.2. For all security parameters λ , inputs $x \in \{0, 1\}^{\ell_{\text{in}}}$, programs $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$ and messages $\text{msg} \in \{0, 1\}$, if $P(x) = \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}.$$

Proof. First, the obfuscator encrypts the program P using an LHE secret key lhe.sk , and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{lhe.sk}, P)$. The evaluator evaluates the LHE ciphertext on universal circuit $U_x(\cdot)$, which results in an evaluated ciphertext $\tilde{\text{ct}}$. Now, by the correctness of the LHE scheme, decryption of $\tilde{\text{ct}}$ using lhe.sk outputs α . Therefore, $\text{PRG}(h(\text{LHE.Dec}(\text{lhe.sk}, \tilde{\text{ct}}))) = \beta$.¹¹ Then, using Lemma 4.1, we can argue that Comp-Eval outputs msg , and thus Eval outputs msg . \blacksquare

Claim 4.3. For all security parameters λ , inputs $x \in \{0, 1\}^{\ell_{\text{in}}}$, programs $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$, $\alpha \in \{0, 1\}^{\ell_{\text{out}}}$ such that $P(x) \neq \alpha$ and $\text{msg} \in \{0, 1\}$,

$$\Pr[\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \perp] \geq 1 - \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Proof. Fix any security parameter λ , program P , α , x such that $P(x) \neq \alpha$ and message msg . Let us now consider the events $\text{Error} : \text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) \neq \perp$ and $\text{Bad} : \text{PRG}(h(P(x))) = \text{PRG}(h(\alpha))$. Clearly, $\Pr[\text{Error}] \leq \Pr[\text{Error} \mid \overline{\text{Bad}}] + \Pr[\text{Bad}]$. We will prove that both terms $\Pr[\text{Bad}]$ and $\Pr[\text{Error} \mid \overline{\text{Bad}}]$ are both negligible in the security parameter.

$\Pr[\text{Error} \mid \overline{\text{Bad}}]$: Since we are conditioning on $\overline{\text{Bad}}$, it follows that $\text{PRG}(h(P(x))) \neq \text{PRG}(h(\alpha))$, and note that this probability is only over the coins chosen during Comp-Gen and Comp-Eval . Let ct be the LHE encryption of P , and $\tilde{\text{ct}} = \text{LHE.Eval}(\text{ek}, U_x(\cdot), \text{ct})$. As described in the construction, let Q be the program that, on input a LHE ciphertext, first decrypts it using sk , then computes $\text{PRG}(h(y))$ (where y is the decrypted string). Let $\text{BP}^{(i)}$ denote the branching program computing the i^{th} output bit of Q . Using the correctness of LHE decryption, it follows that $Q(\tilde{\text{ct}}) = \text{PRG}(h(P(x))) \neq \text{PRG}(h(\alpha))$, and therefore there exists some $i \leq \ell_{\text{PRG}}$ such that $\text{BP}^{(i)}(\tilde{\text{ct}}) \neq \beta_i$ (recall $\beta = \text{PRG}(h(\alpha))$). As a result, using Lemma 4.1, we can conclude that for all $\text{msg} \in \{0, 1\}$, $\Pr[\text{Comp-Eval}(\tilde{\text{ct}}, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) \neq \perp] \leq \text{negl}(\lambda)$.

$\Pr[\text{Bad}]$: For any $y \in \{0, 1\}^{\ell_{\text{PRG}}}$, let $\text{PRG}^{-1}(y) = \{z \in \{0, 1\}^{\ell_{\mathcal{H}}} \text{ s.t. } \text{PRG}(z) = y\}$. Using the PRG security, we can argue that, for any $y \in \{0, 1\}^{\ell_{\text{PRG}}}$, $|\text{PRG}^{-1}(y)|/2^{\ell_{\mathcal{H}}} \leq \text{negl}(\lambda)$. Suppose, on the contrary, there exists some y^* such that $|\text{PRG}^{-1}(y^*)|/2^{\ell_{\mathcal{H}}} = \epsilon$, where ϵ is non-negligible, then we can construct a PPT adversary that can break the PRG security with advantage $\epsilon^2 - \text{negl}(\lambda)$. Concretely, the adversary receives a challenge string t . It chooses a uniformly random string $s \leftarrow \{0, 1\}^{\ell_{\mathcal{H}}}$ and checks if $\text{PRG}(s) = t$. If so, it outputs that t is a pseudorandom string, else it guesses randomly. Clearly, the advantage of this attacker is at least $\epsilon^2 - 2^{-(\ell_{\text{PRG}} - \ell_{\mathcal{H}})}$, where $2^{-(\ell_{\text{PRG}} - \ell_{\mathcal{H}})}$ loss is in the case when the challenge string t lies in the image space of PRG. Since, $\ell_{\text{PRG}} - \ell_{\mathcal{H}} = \omega(\log \lambda)$. Therefore, if PRG is secure then for every $y \in \{0, 1\}^{\ell_{\text{PRG}}}$, $\frac{|\text{PRG}^{-1}(y)|}{2^{\ell_{\mathcal{H}}}} \leq \text{negl}(\lambda)$.

Let $\phi_\gamma = \text{PRG}^{-1}(\text{PRG}(\gamma))$. Now we consider the probability $\Pr[\text{Bad}]$, where the probability is over the choice of $h \leftarrow \mathcal{H}$.

$$\begin{aligned} \Pr[\text{Bad}] &= \Pr[\text{PRG}(h(P(x))) = \text{PRG}(h(\alpha))] \\ &= \sum_{\gamma \in \{0, 1\}^{\ell_{\mathcal{H}}}} \Pr[h(\alpha) = \gamma \wedge h(P(x)) \in \phi_\gamma] \\ &= \sum_{\gamma \in \{0, 1\}^{\ell_{\mathcal{H}}}} \sum_{\delta \in \phi_\gamma} \Pr[h(\alpha) = \gamma \wedge h(P(x)) = \delta] \\ &\leq \sum_{\gamma \in \{0, 1\}^{\ell_{\mathcal{H}}}} \sum_{\delta \in \phi_\gamma} \frac{1}{(2^{\ell_{\mathcal{H}}})^2} = \sum_{\gamma \in \{0, 1\}^{\ell_{\mathcal{H}}}} \frac{|\phi_\gamma|}{(2^{\ell_{\mathcal{H}}})^2} \\ &\leq \max_{\gamma} \frac{|\phi_\gamma|}{2^{\ell_{\mathcal{H}}}} \leq \text{negl}(\lambda). \end{aligned}$$

¹¹As before, we are overloading the notation and using LHE.Dec to decrypt multiple ciphertexts.

Therefore, $\Pr[\text{Error}]$ is bounded by a negligible function in λ . This concludes the proof. ■

4.2 Security

We will now prove the scheme secure as per Definition 3.4. For proving security, we will first need to define a PPT simulator. The simulator Sim gets as input the security parameter λ and size of program $1^{|P|}$, and it must output a simulated program.

We emphasize that our proof will use a polynomial number of hybrids and each (computational) reduction will only rely on a primitive being secure against poly-time attackers. Thus we avoid making any subexponential hardness assumptions.

4.2.1 Simulator Sim

The simulator first chooses the parameters $n, m, q, \sigma, \ell_{\text{PRG}}$ as in the original scheme. Next, it chooses LHE secret/evaluation keys. It computes $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ (note that the depth d of the circuit class is fixed for the scheme). It then computes an encryption of $\mathbf{0}^{|P|}$. Let $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, \mathbf{0}^{|P|})$. Finally, the simulator chooses ℓ_{PRG} matrices $\mathbf{B}_{0,1}^{(i)} \leftarrow \mathbb{Z}_q^{n \times m}$ and low norm matrices $\{(\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)})\}_{i,j}$ for $i \leq \ell_{\text{PRG}}$, $j \leq L$ where $\mathbf{C}_j^{(i,b)} \leftarrow \chi^{m \times m}$. The obfuscation consists of the LHE evaluation key ek , ciphertext ct , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

To prove security, we will define a sequence of hybrids, and show that the hybrids are computationally indistinguishable. The first hybrid corresponds to the case when the challenger outputs an obfuscation of the program, while the final hybrid corresponds to one where the challenger outputs an obfuscation of an ‘all-reject’ program. At a high level, our aim is to replace the LHE encryption of the program P with an encryption of all-zeroes string, and to remove the message msg from the matrix components. To achieve this, we will first replace the matrix components with random low-norm matrices. This switch ensures that the matrix components contain no information about the msg or the LHE secret key. We can then switch the encryption of P with encryption of all-zeroes string.

4.2.2 Sequence of Hybrid Games

In the security game, the adversary \mathcal{A} starts by choosing a program $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$ and message $\text{msg} \in \{0, 1\}$. It sends the P, msg to the challenger and the challenger chooses a random bit ζ . If $\zeta = 0$, the challenger computes program \tilde{P} as an obfuscation of P for message msg with random lock using the obfuscation algorithm. Otherwise, the challenger computes \tilde{P} using the simulator. Now, in the sequence of hybrids described below, **Game 0** will correspond to the case $\zeta = 0$ (i.e., the challenger computes the obfuscated program honestly) and (final) **Game 5** will correspond to $\zeta = 1$ (i.e., the challenger simulates the obfuscated program).

In the following hybrid description, we only describe the challenger’s execution and in the later games, we only highlight the changes made in each consecutive hybrid. Note that the program P and message msg are provided by \mathcal{A} to the challenger.

Game 0: As mentioned above, in this game the challenger computes the obfuscated program honestly.

1. The challenger first chooses the LWE parameters n, m, q, σ, χ and ℓ_{PRG} . Recall L denotes the length of the branching programs.
2. It chooses $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, P)$.
3. Next, it chooses a uniformly random string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$, hash function $h \leftarrow \mathcal{H}$ and sets $\beta = \text{PRG}(h(\alpha))$.

4. Next, consider the following program Q . It takes as input an LHE ciphertext ct , has sk hardwired and does the following: it decrypts the input ciphertext ct to get string x and outputs $\text{PRG}(h(x))$. For $i \leq \ell_{\text{PRG}}(\lambda)$, let $\text{BP}^{(i)}$ denote the branching program that outputs the i^{th} bit of $\text{PRG}(h(x))$.
5. It chooses ℓ_{PRG} uniformly random matrices $\mathbf{B}_L^{(i)}$ of dimensions $5n \times m$, such that the following constraint is satisfied (recall $\mathbf{B}_{L,1}^{(i)}$ represents the first n rows of $\mathbf{B}_L^{(i)}$, $\mathbf{B}_{L,2}^{(i)}$ represents the next n rows of $\mathbf{B}_L^{(i)}$, etc)

$$\sum_{i: \beta_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i: \beta_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \begin{cases} \mathbf{0} & \text{if msg} = 0. \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

6. For $i = 1$ to ℓ_{PRG} and $j = 0$ to $L - 1$, it chooses $(\mathbf{B}_j^{(i)}, T_j^{(i)}) \leftarrow \text{TrapGen}(1^{5n}, 1^m, q)$.
7. Next, it generates the components for each level. For each $i \in [1, \ell_{\text{PRG}}]$ and each level $\text{level} \in [1, L]$, do the following:
 - (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$. If either $\mathbf{S}_{\text{level}}^{(0)}$ or $\mathbf{S}_{\text{level}}^{(1)}$ has determinant zero, then set it to be \mathbf{I}_n .
 - (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
 - (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = (\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)}) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
 - (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$
8. The final obfuscated program consists of the LHE evaluation key ek , LHE encryption ct , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

Game 1: In this experiment, the challenger replaces $h(\alpha)$ with a uniformly random string α' . Also, it chooses the matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)}$ without checking if their determinant is non-zero.

3. Next, it chooses a hash function $h \leftarrow \mathcal{H}$, a uniformly random string $\alpha' \leftarrow \{0, 1\}^{\ell_{\text{PRG}}}$ and sets $\beta = \text{PRG}(\alpha')$.
7. Next, it generates the components for each level. For each level $\text{level} \in [1, L]$, do the following:
 - (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
 - (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
 - (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = (\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)}) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
 - (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

Game 2: In this experiment, the challenger replaces the pseudorandom string $\beta = \text{PRG}(\alpha')$ with a truly random string.

3. Next, it chooses a hash function $h \leftarrow \mathcal{H}$, a uniformly random string $\beta \leftarrow \{0, 1\}^{\ell_{\text{PRG}}}$.

Game 3: In this experiment, the challenger chooses the top level matrices uniformly at random.

5. It chooses ℓ_{PRG} uniformly random matrices $\mathbf{B}_L^{(i)}$ of dimensions $5n \times m$, without any constraints.

We will now define $4L+1$ intermediate game $\text{Game}(4, j^*, 0)$, $\text{Game}(4, j^*, 1)$, $\text{Game}(4, j^*, 2)$, $\text{Game}(4, j^*, 3)$ for $j^* \in \{0, 1, \dots, L-1\}$ and $\text{Game}(4, L, 0)$. The experiment $\text{Game}(4, 0, 0)$ will be same as Game 3 .

Game (4, j^* , 0): In this experiment, the challenger chooses the top $j^* + 1$ matrices $\mathbf{B}_j^{(i)}$ uniformly at random (without a trapdoor). Also, the top j^* $\mathbf{C}_j^{(i,b)}$ matrices, for $j > L - j^*$, are chosen randomly from the noise distribution $\chi^{m \times m}$.

6. For $i = 1$ to ℓ_{PRG} and $j = 0$ to $L - 1$,

- if $j < L - j^*$ it chooses $(\mathbf{B}_j^{(i)}, T_j^{(i)}) \leftarrow \text{TrapGen}(1^{5n}, 1^m, q)$.
- if $j \geq L - j^*$, it chooses $\mathbf{B}_j^{(i)}$ uniformly at random from $\mathbb{Z}_q^{5n \times m}$.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
- (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game (4, j^* , 1): In this experiment, the challenger chooses $\mathbf{M}_{\text{level}}^{(i,0)}$ uniformly at random for $\text{level} = L - j^*$ ($\mathbf{M}_{\text{level}}^{(i,1)}$ is same as before).

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
- (c) If $\text{level} = L - j^*$, choose $\mathbf{M}_{\text{level}}^{(i,0)} \leftarrow \mathbb{Z}_q^{5n \times m}$ and $\mathbf{M}_{\text{level}}^{(i,1)}$ same as previous hybrid.
Else set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game (4, j^* , 2): In this experiment, the challenger chooses $\mathbf{M}_{\text{level}}^{(i,b)}$ uniformly at random for $\text{level} = L - j^*$ and $b \in \{0, 1\}$.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
- (c) If $\text{level} = L - j^*$, choose $\mathbf{M}_{\text{level}}^{(i,b)} \leftarrow \mathbb{Z}_q^{5n \times m}$.
Else set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game $(4, j^*, 3)$: In this experiment, the challenger chooses $\mathbf{C}_{L-j^*}^{(i,b)}$ from the noise distribution.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level}, b}^{(i)}(\cdot)$.
- (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = (\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)}) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) **If $\text{level} = L - j^*$, then choose $\mathbf{C}_{L-j^*}^{(i,b)} \leftarrow \chi^{m \times m}$ for $b \in \{0, 1\}$.**
Else compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game 5: In this experiment, the challenger computes an encryption of $\mathbf{0}^{|P|}$ instead of an encryption of P . Note that the LHE secret key sk is not required at any other step in this experiment. This corresponds to the challenger simulating the program.

- 2. It chooses $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, \mathbf{0}^{|P|})$.

4.2.3 Analysis

We will now show that the above hybrid experiments are computationally indistinguishable. For any adversary \mathcal{A} , let $p_i^{\mathcal{A}}$ denote the probability that the adversary outputs 1 in Game i .

Claim 4.4. For any adversary \mathcal{A} , $|p_0^{\mathcal{A}} - p_1^{\mathcal{A}}| \leq \text{negl}(\lambda)$.

Proof. The proof of this claim relies on the fact that a pairwise independent hash function (with appropriate domain and co-domain) is a strong extractor. There are two differences between these two experiments. First, in one case, the challenger chooses $h \leftarrow \mathcal{H}$, $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ and sets $\beta = \text{PRG}(h(\alpha))$, while in the other case, the challenger chooses $h \leftarrow \mathcal{H}$, $\alpha' \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ and sets $\beta = \text{PRG}(\alpha')$. Secondly, in Game 0, the challenger ensures that the matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)}$ are full-rank, while in Game 1, it chooses these matrices from χ without the rank check. Note that any $n \times n$ matrix sampled from χ is full rank with all but negligible probability. Therefore, this switch is indistinguishable. We will now show that switching $h(\alpha)$ to a uniformly random string is also statistically indistinguishable.

Suppose there exists an adversary \mathcal{A} such that $|p_0^{\mathcal{A}} - p_1^{\mathcal{A}}| = \epsilon$. We can construct a reduction algorithm \mathcal{B} that can break the strong extraction property of \mathcal{H} with probability ϵ . The reduction algorithm receives (h, γ) from the strong extractor challenger, and P, msg from the adversary \mathcal{A} . It chooses LHE keys $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$, computes an encryption of P as $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, P)$ and sets $\beta = \text{PRG}(\gamma)$. Next, it defines the program Q which takes as input an LHE ciphertext, and has sk hardwired. Q first decrypts the input ciphertext to compute y , then computes $\text{PRG}(h(y))$. Let $\{\text{BP}^{(i)}\}_i$ be the branching programs corresponding to Q . The reduction algorithm computes the components using $\text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})$, which is identical in both experiments. It sends the components, together with ek and ct , to \mathcal{A} and the adversary outputs a bit b' . If $b' = 1$, the reduction algorithm guesses that γ is an evaluation of the hash function. The reduction algorithm has advantage $|p_0^{\mathcal{A}} - p_1^{\mathcal{A}}| = \epsilon$ in the strong extractor experiment. \blacksquare

Claim 4.5. Assuming the security of PRG, for any PPT adversary \mathcal{A} , $|p_1^{\mathcal{A}} - p_2^{\mathcal{A}}| \leq \text{negl}(\lambda)$.

Proof. We describe a reduction algorithm \mathcal{B} which plays the indistinguishability based game with PRG challenger.¹² \mathcal{B} receives β from the PRG challenger, and P, msg from the adversary \mathcal{A} . It then chooses

¹²Note that if \mathcal{A} is a non-uniform attacker, then our reduction algorithm will be non-uniform as well.

$(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ and computes $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, P)$. Next, it defines the program Q which takes as input an LHE ciphertext, and has sk hardwired. Q first decrypts the input ciphertext to compute y , then computes $\text{PRG}(h(y))$. Let $\{\text{BP}^{(i)}\}_i$ be the branching programs corresponding to Q . The reduction algorithm computes the components using $\text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})$, which is identical in both experiments. It sends the components, together with ek and ct , and the adversary outputs a bit b' . If $b' = 1$, the reduction algorithm guesses that β is an evaluation of the hash function.

Note that when the PRG challenger sends a pseudorandom string $\beta = \text{PRG}(\alpha')$ for a uniformly random α' , then \mathcal{B} exactly simulates the view of Game 1 for \mathcal{A} . Otherwise if the PRG challenger outputs a uniformly random string, then \mathcal{B} exactly simulates the view of Game 2. Therefore, if $|p_1^A - p_2^A|$ is non-negligible, then PRG is not a secure pseudorandom generator. \blacksquare

Claim 4.6. For any adversary \mathcal{A} , $|p_2^A - p_3^A| \leq \text{negl}(\lambda)$.

Proof. This step is information theoretic, and uses the Leftover Hash Lemma (matrix version, Corollary 2.1).

Note that the difference between Game 2 and 3 is the way top level matrices $\mathbf{B}_{L, \text{st}_L}^{(i)}$ are sampled during the component generation phase. In Game 2, matrix $\mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})}$ is chosen as

$$\mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})} = \begin{cases} - \left(\sum_{i \leq \ell_{\text{PRG}} - 1 : \beta_i = 0} \mathbf{B}_{L, \text{rej}^{(i)}}^{(i)} + \sum_{i \leq \ell_{\text{PRG}} - 1 : \beta_i = 1} \mathbf{B}_{L, \text{acc}^{(i)}}^{(i)} \right) & \text{if msg} = 0 \\ - \left(\sum_{i \leq \ell_{\text{PRG}} - 1 : \beta_i = 0} \mathbf{B}_{L, \text{rej}^{(i)}}^{(i)} + \sum_{i \leq \ell_{\text{PRG}} - 1 : \beta_i = 1} \mathbf{B}_{L, \text{acc}^{(i)}}^{(i)} \right) + \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1 \end{cases}$$

where $\text{st}_L^{(\ell_{\text{PRG}})}$ is $\text{acc}^{(\ell_{\text{PRG}})}$ if $\beta_{\ell_{\text{PRG}}} = 1$, and $\text{rej}^{(\ell_{\text{PRG}})}$ otherwise. It can be equivalently written as follows

$$\mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})} = -\mathbf{A} \cdot \mathbf{R}, \quad \mathbf{A} = \left[\mathbf{B}_{L, \text{rej}^{(1)}}^{(1)} \parallel \mathbf{B}_{L, \text{acc}^{(1)}}^{(1)} \parallel \dots \parallel \mathbf{B}_{L, \text{rej}^{(\ell_{\text{PRG}}-1)}}^{(\ell_{\text{PRG}}-1)} \parallel \mathbf{B}_{L, \text{acc}^{(\ell_{\text{PRG}}-1)}}^{(\ell_{\text{PRG}}-1)} \right]$$

where $\mathbf{R} = \mathbf{u} \otimes \mathbf{I}_m \in \mathbb{Z}_q^{2m(\ell_{\text{PRG}}-1) \times m}$, $\mathbf{u} = (u_1, \dots, u_{2\ell_{\text{PRG}}-2})^\top \in \{0, 1\}^{2\ell_{\text{PRG}}-2}$ and for all $i \leq \ell_{\text{PRG}} - 1$, $u_{2i} = \beta_i$ and $u_{2i-1} = 1 - \beta_i$. That is, matrix \mathbf{R} consists of $2\ell_{\text{PRG}} - 2$ submatrices where if $\beta_i = 1$, then its $2i^{\text{th}}$ submatrix is identity and $(2i - 1)^{\text{th}}$ submatrix is zero, otherwise it is the opposite. Let \mathcal{R} denote the distribution of matrix \mathbf{R} as described above with β drawn uniformly from $\{0, 1\}^{\ell_{\text{PRG}}}$. Note that $\mathbf{H}_\infty(\mathcal{R}) = \ell_{\text{PRG}} - 1$ (min-entropy of \mathcal{R}), and $\ell_{\text{PRG}} > m \cdot n \log_2 q + \omega(\log n)$. Therefore, it follows (from Corollary 2.1) that

$$\begin{aligned} & \left\{ \left(\mathbf{A}, \mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})} = -\mathbf{A} \cdot \mathbf{R} \right) : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times 2m(\ell_{\text{PRG}}-1)}, \mathbf{R} \leftarrow \mathcal{R} \right\} \\ & \quad \approx_s \\ & \left\{ \left(\mathbf{A}, \mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})} \right) : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times 2m(\ell_{\text{PRG}}-1)}, \mathbf{B}_{L, \text{st}_L}^{(\ell_{\text{PRG}})} \leftarrow \mathbb{Z}_q^{n \times m} \right\} \end{aligned}$$

Thus, $|p_2^A - p_3^A|$ is negligible in the security parameter for all adversaries \mathcal{A} . \blacksquare

Claim 4.7. For all $j^* \in [1, L]$, for all \mathcal{A} , $|p_{(4, j^* - 1, 3)}^A - p_{(4, j^*, 0)}^A| \leq \text{negl}(\lambda)$.

Proof. This proof relies on the Matrix Well Distributedness property of TrapGen (see Definition 2.1). The only difference in Game $(4, j^* - 1, 3)$ and Game $(4, j^*, 0)$ is the choice of $\{\mathbf{B}_{L-j^*}^{(i)}\}$. In Game $(4, j^* - 1, 3)$, $\{\mathbf{B}_{L-j^*}^{(i)}\}$ is chosen with a trapdoor for all $i \leq \ell_{\text{PRG}}$. In Game $(4, j^*, 0)$, these matrices are chosen uniformly at random. Note that in both the experiments, the trapdoors corresponding to $\mathbf{B}_{L-j^*}^{(i)}$ are not used, since the matrix components $\{\mathbf{C}_{L-j^*+1}^{(i,b)}\}_{i,b}$ are chosen from the noise distribution χ . As a result, using the Matrix Well Distributedness property of TrapGen, we can argue that these two games are statistically indistinguishable, and therefore $|p_{(4, j^* - 1, 3)}^A - p_{(4, j^*, 0)}^A| \leq \text{negl}(\lambda)$. \blacksquare

Claim 4.8. Assuming the LWE with short secrets assumption (Assumption 2) For all $j^* \in [0, L - 1]$, for all PPT adversaries \mathcal{A} , $|p_{(4,j^*,0)}^{\mathcal{A}} - p_{(4,j^*,1)}^{\mathcal{A}}| \leq \text{negl}(\lambda)$

Proof. In Game $(4, j^*, 0)$, the challenger computes $\mathbf{M}_{L-j^*}^{(i,0)}$ as $(\mathbf{I}_5 \otimes \mathbf{S}_{L-j^*}^{(0)}) \cdot \mathbf{D}_{L-j^*}^{(i,0)} + \mathbf{E}_{L-j^*}^{(i,0)}$. In Game $(4, j^*, 1)$, this matrix is chosen uniformly at random. We will use the LWE assumption (short secrets version) to prove that these two hybrids are indistinguishable.

Suppose there exists a PPT adversary \mathcal{A} that can distinguish between these two games with advantage ϵ . Then we can construct a reduction algorithm \mathcal{B} that can break the LWE-ss assumption. First, \mathcal{B} receives as LWE-ss challenge two matrices (\mathbf{F}, \mathbf{G}) of dimensions $n \times 5\ell_{\text{PRG}} \cdot m$, and P, msg from the adversary \mathcal{A} . It then partitions \mathbf{F} into ℓ_{PRG} submatrices of dimensions $n \times 5m$ as $[\mathbf{F}_1^{(1)} \parallel \dots \parallel \mathbf{F}_{\ell_{\text{PRG}}}^{(\ell_{\text{PRG}})}]$. Further, each matrix $\mathbf{F}^{(i)}$ is parsed into 5 matrices of dimensions $n \times m$ as $[\mathbf{F}_1^{(i)} \parallel \dots \parallel \mathbf{F}_5^{(i)}]$. Similarly, the matrix \mathbf{G} is partitioned into $\{\mathbf{G}_k^{(i)}\}_{i \leq \ell_{\text{PRG}}, k \leq 5}$.

The reduction algorithm chooses the LHE keys and computes the ciphertext as in the two games. Next, it needs to choose the matrix components. First, it chooses the top $j^* - 1$ matrices $\{\mathbf{B}_j^{(i)}\}_{i \leq \ell_{\text{PRG}}, j > L-j^*}$ uniformly at random, and the top $j^* - 1$ components $\{(\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)})\}_{i \leq \ell_{\text{PRG}}, j > L-j^*}$ from the noise distribution. It then sets $\mathbf{B}_{L-j^*,k}^{(i)} = \mathbf{F}_k^{(i)}$ for all $i \leq \ell_{\text{PRG}}, k \leq 5$.

The reduction algorithm then uses $\mathbf{G}_k^{(i)}$ to define $\mathbf{M}_{L-j^*}^{(i,0)}$ as follows:

$$\mathbf{M}_{L-j^*}^{(i,0)} = \begin{bmatrix} \mathbf{G}_{\sigma_{L-j^*,0}^{(i)}(1)}^{(i)} \\ \vdots \\ \mathbf{G}_{\sigma_{L-j^*,0}^{(i)}(5)}^{(i)} \end{bmatrix}.$$

The matrix $\mathbf{M}_{L-j^*}^{(i,1)}$ is computed identically in both the games. Using $\mathbf{B}_{L-j^*}^{(i)}$, the reduction algorithm can compute $\mathbf{M}_{L-j^*}^{(i,1)}$. Finally, it computes $\mathbf{C}_{L-j^*}^{(i,b)}$ for all $i \leq \ell_{\text{PRG}}, b \in \{0, 1\}$.

The remaining components are computed identically in both games, and the reduction algorithm can perfectly simulate them. Finally, it sends all the components, together with the LHE evaluation key and ciphertext. The adversary sends a bit b' . If $b' = 1$, the reduction algorithm guesses that $\mathbf{G} = \mathbf{S} \cdot \mathbf{F} + \mathbf{E}$ for some matrices \mathbf{S}, \mathbf{E} drawn from the noise distribution. If $b' = 0$, then the reduction algorithm guesses that \mathbf{G} is random.

If \mathbf{G} is an LWE sample, then \mathcal{B} simulates Game $(4, j^*, 0)$, else it simulates Game $(4, j^*, 1)$. The advantage of \mathcal{B} is $p_{(4,j^*,0)}^{\mathcal{A}} - p_{(4,j^*,1)}^{\mathcal{A}}$. ■

Claim 4.9. Assuming the LWE with short secrets assumption (Assumption 2) For all $j^* \in [0, L - 1]$, for all PPT adversaries \mathcal{A} , $|p_{(4,j^*,1)}^{\mathcal{A}} - p_{(4,j^*,2)}^{\mathcal{A}}| \leq \text{negl}(\lambda)$

Proof. In Game $(4, j^*, 1)$, the challenger computes $\mathbf{M}_{L-j^*}^{(i,1)}$ as $(\mathbf{I}_5 \otimes \mathbf{S}_{L-j^*}^{(1)}) \cdot \mathbf{D}_{L-j^*}^{(i,1)} + \mathbf{E}_{L-j^*}^{(i,1)}$. In Game $(4, j^*, 2)$, this matrix is chosen uniformly at random. This proof is identical to the previous proof. ■

Claim 4.10. For all $j^* \in [0, L - 1]$, for all \mathcal{A} , $|p_{(4,j^*,2)}^{\mathcal{A}} - p_{(4,j^*,3)}^{\mathcal{A}}| \leq \text{negl}(\lambda)$.

Proof. The proof of this claim relies on the *preimage well-distributedness property* of lattice trapdoor sampler (TrapGen, SamplePre) (Definition 2.1). Note that the only difference between these two games is with regard to $\mathbf{C}_{L-j^*}^{(i,b)}$. In Game $(4, j^*, 2)$ these matrices are computed using the SamplePre algorithm. In Game $(4, j^*, 3)$, they're chosen from the noise distribution χ . Since $\mathbf{M}_{L-j^*}^{(i,b)}$ is chosen uniformly at random, we can use the preimage well-distributedness property to argue that the two games are statistically indistinguishable. ■

Claim 4.11. Assuming the LHE scheme is IND-CPA secure, for all PPT adversaries \mathcal{A} , $|p_{4,L,0}^{\mathcal{A}} - p_5^{\mathcal{A}}| \leq \text{negl}(\lambda)$.

Proof. The only difference between $p_{4,L,0}^{\mathcal{A}}$ and $p_5^{\mathcal{A}}$ is that the LHE ciphertext is an encryption of P in one case, and an encryption of $\mathbf{0}^{|P|}$ in the other. Note that in both games, the LHE secret key is only used for computing the ciphertext (it is not used by the branching programs $\text{BP}^{(i)}$). Suppose there exists an adversary \mathcal{A} that can distinguish between these two games. Then we can construct a reduction algorithm \mathcal{B} that breaks the IND-CPA security of the LHE scheme.

The reduction algorithm receives P, msg from the adversary \mathcal{A} , and sends $P, \mathbf{0}^{|P|}$ to the LHE challenger. It receives an evaluation key ek and a ciphertext ct . It then chooses the matrix components $\mathbf{B}_0^{(i)}$ uniformly at random, and the components $\{\mathbf{C}_j^{(i,b)}\}_{i,j,b}$ from the noise distribution. Finally, it sends ek, ct and the matrix components to the adversary. The adversary sends a bit b' . If $b' = 1$, then the reduction algorithm guesses that ct is an encryption of P , else it guesses that ct is an encryption of $\mathbf{0}^{|P|}$. Note that the advantage of \mathcal{B} in the IND-CPA game is $|p_{(4,L,0)}^{\mathcal{A}} - p_5^{\mathcal{A}}|$. \blacksquare

4.3 On the Need for a Pseudo Random Generator

An interesting question is whether the application of a Pseudo Random Generator is actually necessary for the security of our construction. A natural alternative construction would be to use the same overall structure with the modification that the circuit Q computes $Q(\text{ct}_1, \dots, \text{ct}_{\ell_{\text{out}}}) = \text{LHE.Dec}(\text{lhe.sk}, \text{ct}_1) \parallel \dots \parallel \text{LHE.Dec}(\text{lhe.sk}, \text{ct}_{\ell_{\text{out}}})$. In particular, Q just performs an LHE decryption and does not compute a pairwise independent hash and pseudo random generator to the result. In this case we set $\beta = \alpha$ instead of $\beta = \text{PRG}(h(\alpha))$.

While we do not rule out that such a construction might work, there appear to be some inherent tensions in setting parameters to prove such a system secure. Let's return to the proof of the scheme as it is. A critical proof step is to move from Game 2 to Game 3 where in Game 2 it chooses the $\mathbf{B}_{L,\text{acc}^{(i)}}^{(i)}, \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)}$ values such that

$$\sum_{i : \beta_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i : \beta_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if msg} = 0. \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

Whereas in Game 3 all are chosen uniformly at random. By our choice of parameters we have that $\ell_{\text{PRG}} \gg n \cdot m \cdot \log q$. Therefore, with high probability, there will be many strings w such that $\sum_{i : w_i=0} \mathbf{B}_{L,w_i}^{(i)} = \mathbf{0}$. It follows that (in the case where the message bit is 0) programming the matrices so that on a random β $\sum_{i : \beta_i=0} \mathbf{B}_{L,\beta_i}^{(i)} = \mathbf{0}$ is statistically close to just choosing them uniformly at random.

Now if we instead use the alternative construction we have that there are ℓ_{out} pairs of matrices instead at the last level instead of ℓ_{PRG} . In our construction $\ell_{\text{out}} < \ell_{\text{in}}$ since the input HE ciphertext must be big enough to encrypt α . And $\ell_{\text{in}} < \log q$ due to the noise accumulation in our construction. So we have that $\ell_{\text{out}} \ll \log q$ and the statistical argument we used earlier for going from Game 2 to Game 3 no longer holds. We might instead consider a computational argument for moving between these games where the LWE assumption with binary secrets is a natural candidate. Unfortunately, known conversions [BLP⁺13, MP13] between standard LWE and binary LWE require the LWE challenge vector to increase by a factor of $\log q$ which would again violate our parameter constraints.

Another interesting question is whether the pairwise independent hash function could be removed from our construction where the circuit Q would go straight from decrypting the HE ciphertext to evaluating a pseudorandom generator. Currently, the pairwise independent hash function is in place to assure semi-statistical correctness of Definition 3.3. Consider a pseudorandom generator PRG that always ignored its first input bit. Furthermore consider a program P , input x , lock value α such that $P(x) = \tilde{\alpha}$, where $\tilde{\alpha}$ is a bitstring that equals α at all bit positions except the first. Then the triple P , input x , lock value α would violate semi-statistical correctness under the construction where the pairwise independent hash function is removed.

At the same time if we were to use an *injective* pseudo random generator such issues are not a factor and the construction will actually achieve statistical correctness of Definition 3.2. We show this formally in Appendix D.

5 Predicate Encryption: Achieving 1-Sided Security

Below we first give a formal description of our transformation for the case of a key-policy ABE (KP-ABE) system where the decryption algorithms are of a bounded depth and size. In Appendix E we show a slight generalization of the first transformation where we add a level of indirection using fully homomorphic encryption to remove the depth and size constraints. We follow (in this section) by enumerating different forms of Attribute-Based Encryption and informally discussing how our transformation could be adapted to each case. We finally sketch a derivative application where we can encrypt to an arbitrarily formed public key and hide from everyone except the receiver which public key we encrypted to.

5.1 Key-Policy ABE with Bounded Decryption Depth and Size

In this section, we construct a PE scheme which achieves 1-sided security from any KP-ABE scheme and a lockable obfuscator. Our construction inherits the attribute space and predicate class of the underlying KP-ABE scheme.

Outline. The central idea is to hide the ciphertext inside an obfuscated circuit. During encryption, instead of encrypting the message m (using the underlying encryption algorithm), the encryptor chooses a lock value α for obfuscation and encrypts the lock value α under the input attribute x to get ciphertext ct_α . Next, it obfuscates (under the same lock value α) the decryption circuit of the underlying ABE scheme with the ciphertext ct_α (encrypting lock value) hardwired. Now, output of this obfuscated circuit \tilde{P} would be set to be the actual message m and the obfuscated circuit \tilde{P} will be the final ciphertext. Observe that given any valid secret key sk_C (such that $C(x) = 1$), one could simply evaluate the obfuscated circuit using the secret key and its output would guaranteed to be message m by correctness of the underlying decryption and evaluation algorithms. At a high level, this hides the attribute x because if the adversary can not decrypt the ciphertext ct_α containing the lock value α (i.e, it does not know the secret key for the corresponding attribute), then it has not no information about the lock value α . Therefore, by the obfuscation security, we can guarantee that the circuit \tilde{P} (i.e., challenge ciphertext) could be simulated instead. Below we describe our construction $\mathcal{PE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ for attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and 1-bit messages.

Construction. Let $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ be a key-policy attribute based encryption scheme for set of attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\mathcal{M}_\lambda\}_\lambda$ with decryption circuit of depth $d(\lambda)$ and secret key space $\{0, 1\}^{\ell(\lambda)}$. Also, let $\mathcal{O} = (\text{Obf}, \text{Eval})$ be a lockable obfuscator for circuit class $\mathcal{C}_{\ell, k, d}$ (i.e., the class of depth $d(\lambda)$ circuits with $\ell(\lambda)$ bit input and $k(\lambda)$ bit output). For simplicity, assume that $\mathcal{M}_\lambda = \{0, 1\}^{k(\lambda)}$. Below we describe our construction. For notational convenience, let $\ell = \ell(\lambda), k = k(\lambda)$ and $d = d(\lambda)$.

- $\text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm runs ABE.Setup to generate public parameters and master secret key as $(\text{pp}, \text{msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, 1^k)$.
- $\text{Enc}(\text{pp}, x, m) \rightarrow \text{ct}$. The encryption algorithm chooses a random string $\alpha \leftarrow \{0, 1\}^k$ and encrypts it as $\text{ct}_\alpha \leftarrow \text{ABE.Enc}(\text{pp}, x, \alpha)$. Next, it obfuscates the decryption circuit with message m and lock α as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{ABE.Dec}(\cdot, \text{ct}_\alpha), m, \alpha)$. Finally, it outputs the ciphertext as $\text{ct} = \tilde{P}$.
- $\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm runs ABE.KeyGen to generate the secret key as $\text{sk}_C \leftarrow \text{ABE.KeyGen}(\text{msk}, C)$.

- $\text{Dec}(\text{sk}_C, \text{ct}) \rightarrow m$ or \perp . Let $\text{ct} = \tilde{P}$. The decryption algorithm evaluates the obfuscated program on input sk_C , and outputs $\text{Eval}(\tilde{P}, \text{sk}_C)$.

Remark 5.1 (Preservation of Algorithms and Retroactive Application). We would like to point out that the above transformation preserves the original ABE setup and key generation algorithms as well as the private key structure. This could allow some flexibility in encryption. If the user encrypting data feels that hiding the attributes is important, he can apply our transformation. If not, he can default to the original scheme. Also, it allows for one-sided predicate encryption to be retroactively applied to an ABE scheme that has already published its public parameters and distributed its keys. Of course, only newly generated ciphertexts generated with using the transformation would have the attributes hidden.

Correctness. For all $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}$, attribute $x \in \mathcal{X}_\lambda$, public parameters and master secret key $(\text{pp}, \text{msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda)$, the ciphertext corresponding to message m under attribute x in our scheme is an obfuscated circuit \tilde{P} , where $\alpha \leftarrow \{0, 1\}^k$, $\text{ct}_\alpha \leftarrow \text{ABE.Enc}(\text{pp}, x, \alpha)$ and $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{ABE.Dec}(\cdot, \text{ct}_\alpha), m, \alpha)$.

For any predicate $C \in \mathcal{C}_\lambda$, the corresponding secret key in our scheme is simply $\text{sk}_C \leftarrow \text{ABE.KeyGen}(\text{msk}, C)$. Consider the following two cases:

1. $C(x) = 1$: We know that if $C(x) = 1$, then with all but negligible probability $\text{ABE.Dec}(\text{sk}_C, \text{ct}_\alpha) = \alpha$. This follows from correctness of the \mathcal{ABE} scheme. Next, using correctness of obfuscation, we can conclude that $\text{Eval}(\tilde{P}, \text{sk}_C) = m$ as $\text{ABE.Dec}(\text{sk}_C, \text{ct}_\alpha) = \alpha$. Therefore, if $C(x) = 1$, then $\Pr[\text{Dec}(\text{sk}_C, \text{ct} = \tilde{P}) = m] \geq 1 - \text{negl}(\lambda)$.
2. $C(x) = 0$: Similarly, we know that if $C(x) = 0$, then with all but negligible probability $\text{ABE.Dec}(\text{sk}_C, \text{ct}_\alpha) = \perp$. Also by correctness of obfuscation, we can conclude that if $\text{ABE.Dec}(\text{sk}_C, \text{ct}_\alpha) \neq \alpha$, then $\Pr[\text{Eval}(\tilde{P}, \text{sk}_C) = \perp] \geq 1 - \text{negl}(\lambda)$. Since $\perp \neq \alpha$, therefore combining these two facts we know that if $C(x) = 0$, then $\Pr[\text{Dec}(\text{sk}_C, \text{ct} = \tilde{P}) = \perp] \geq 1 - \text{negl}'(\lambda)$ for some negligible function $\text{negl}'(\cdot)$.

Therefore, \mathcal{PE} satisfies the predicate encryption correctness condition.

Security. We will now show that the scheme described above achieves 1-sided security as per Definition A.2. Formally, we prove the following.

Theorem 5.1. If $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a fully secure attribute based encryption for set of attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\mathcal{M}_\lambda\}_\lambda$ satisfying Definition A.1, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for 1-bit messages and circuit class $\{\mathcal{C}_{\ell, k, d}\}_\lambda$ satisfying Definition 3.4, then \mathcal{PE} is a secure predicate encryption scheme satisfying 1-sided security as per Definition A.2 for 1-bit messages and same attribute space and predicate class as the ABE scheme.

Our proof proceeds via a sequence of hybrid games. Each game is played between the challenger and attacker \mathcal{A} . Let \mathcal{A} be any PPT adversary that wins the 1-sided security game with non-negligible advantage. We argue that such an adversary must break security of at least one underlying primitive. The first game corresponds to the 1-sided security game as described in Definition A.2. In the final game, the challenge ciphertext is simulated and does not contain any information about the challenge messages.

We will first define the sequence of hybrid games, and then show that they are computationally indistinguishable.

Game 1: In this game, the challenge ciphertext ct is honestly generated, i.e. by first choosing a uniformly random lock α , then encrypting it under attribute x to get ciphertext ct_α and finally obfuscating the decryption circuit with lock α and message m_b .

1. **Setup Phase.** The challenger sets up by generating public parameters and master secret key as $(\text{pp}, \text{msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, 1^k)$. It sends pp to \mathcal{A} .

2. **Pre-Challenge Query Phase.** \mathcal{A} queries the challenger on polynomially many predicate circuits C_i to receive the corresponding secret keys as $\text{sk}_{C_i} \leftarrow \text{ABE.KeyGen}(\text{msk}, C_i)$.
3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$ and a uniformly random lock string $\alpha \leftarrow \{0, 1\}^k$. It computes ciphertext ct as $\text{ct} \leftarrow \text{ABE.Enc}(\text{pp}, x_b, \alpha)$. It also generates obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{ABE.Dec}(\cdot, \text{ct}), m_b, \alpha)$, and sends \tilde{P} to \mathcal{A} .
4. **Post-Challenge Query Phase.** \mathcal{A} queries the challenger on polynomially many predicate circuits C_i as before. The challenger handles these as in pre-challenge query phase.
5. **Guess.** \mathcal{A} outputs its guess b' and wins if $b = b'$.

Game 2: This is same as Game 1, except the challenger computes ciphertext ct as an encryption of an all zero string instead of α .

3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$ and a uniformly random lock string $\alpha \leftarrow \{0, 1\}^k$. It computes ciphertext ct as $\text{ct} \leftarrow \text{ABE.Enc}(\text{pp}, x_b, 0^k)$. It also generates obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{ABE.Dec}(\cdot, \text{ct}), m_b, \alpha)$, and sends \tilde{P} to \mathcal{A} .

Game 3: This is same as Game 2, except the challenger does not choose the lock α anymore and it simulates the obfuscated program \tilde{P} instead of generating it honestly as an obfuscation of the ABE decryption circuit.

3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$. It generates obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$ (where $s = |\text{ABE.Dec}(\cdot, \text{ct})|$), and sends \tilde{P} to \mathcal{A} .

Let $\text{Adv}_{\mathcal{A}}^i = |\Pr[b' = b] - 1/2|$ denote the advantage of adversary \mathcal{A} in guessing the bit b in Game i . First, note that $\text{Adv}_{\mathcal{A}}^3 = 0$. In other words, \mathcal{A} 's advantage in Game 3 is 0. This is because the challenge ciphertext \tilde{P} does not contain any information about bit b .

To complete the proof, we establish via a sequence of lemmas that no PPT adversary \mathcal{A} can distinguish between each adjacent game with non-negligible probability. We show via a sequence of lemmas that $|\text{Adv}_{\mathcal{A}}^i - \text{Adv}_{\mathcal{A}}^{i+1}|$ is negligible for all $i = 1, 2$. Below we discuss our lemmas in detail.

Lemma 5.1. If $\text{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a fully secure attribute based encryption, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish encryptions of lock α and an all zeros string, therefore break security of the ABE scheme.

The ABE challenger generates a key pair (pp, msk) and sends pp to \mathcal{B} . \mathcal{B} simply forwards the public parameters pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} forwards each predicate query C_i to the ABE challenger and receives back secret key sk_{C_i} , which it then sends to \mathcal{A} as its response. Next, \mathcal{A} sends the challenge message-attribute pairs $((m_0, x_0), (m_1, x_1))$ to \mathcal{B} such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. \mathcal{B} chooses a random string $\alpha \leftarrow \{0, 1\}^k$ and bit $b \leftarrow \{0, 1\}$, and sends $t_0 = \alpha$ and $t_1 = 0^k$ as its challenge messages and x_b as the challenge attribute to the ABE challenger. The ABE challenger chooses a random bit $\beta \leftarrow \{0, 1\}$, computes the challenge ciphertext $\text{ct}^* \leftarrow \text{ABE.Enc}(\text{pp}, x_b, t_\beta)$ and sends ct^* to \mathcal{B} . After receiving ct^* from the challenger, \mathcal{B} runs Step 3 as in Game 1. That is, \mathcal{B} generates obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{ABE.Dec}(\cdot, \text{ct}^*), m_b, \alpha)$,

and sends \tilde{P} to \mathcal{A} . Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before by forwarding those to ABE challenger. Finally, \mathcal{A} outputs its guess b' . If $b = b'$, then \mathcal{B} sends 0 as its guess (i.e., α was encrypted), otherwise it sends 1 as its guess (i.e., 0^k was encrypted) to the ABE challenger.

Note that if the ABE challenger encrypted α (i.e., $\beta = 0$), then \mathcal{B} perfectly simulates Game 1 for adversary \mathcal{A} . Otherwise it simulates Game 2 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible, then \mathcal{B} breaks the ABE scheme's security with non-negligible advantage. \blacksquare

Lemma 5.2. If $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a secure lockable obfuscator, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish an obfuscation of the ABE decryption circuit with a random lock from a simulated obfuscated program, therefore break security of the obfuscation scheme.

\mathcal{B} samples an ABE key pair (pp, msk) and sends pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} answers each query with secret key sk_{C_i} , with $\text{sk}_{C_i} \leftarrow \text{ABE.KeyGen}(\text{msk}, C_i)$. Next, \mathcal{A} sends the challenge message-attribute pairs $((m_0, x_0), (m_1, x_1))$ to \mathcal{B} such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. \mathcal{B} chooses a random bit $b \leftarrow \{0, 1\}$ and computes ciphertext ct as $\text{ct} \leftarrow \text{ABE.Enc}(\text{pp}, x_b, 0^k)$. \mathcal{B} sends the $\text{ABE.Dec}(\cdot, \text{ct})$ circuit along with message m_b to the obfuscation challenger, and receives an obfuscated program \tilde{P} . \mathcal{B} finally sends \tilde{P} to \mathcal{A} as the challenge ciphertext. Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before. Finally, \mathcal{A} outputs its guess b' . If $b = b'$, then \mathcal{B} sends 0 as its guess (i.e., ABE decryption circuit was obfuscated), otherwise it sends 1 as its guess (i.e., program \tilde{P} was simulated) as its guess to the obfuscation challenger.

Note that if the obfuscation challenger obfuscated $\text{ABE.Dec}(\cdot, \text{ct})$ for some lock α , then \mathcal{B} perfectly simulates Game 2 for adversary \mathcal{A} . Otherwise it simulates Game 3 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible, then \mathcal{B} breaks the obfuscation scheme's security with non-negligible advantage. \blacksquare

5.2 Generalizing to Other Types

We now briefly discuss other forms of Attribute-Based Encryption and informally describe how our transformation could be adapted to them.

Key-Policy ABE with Unbounded Decryption Depth. The above construction above applies to ABE systems where there is a system wide bound on the depth of the decryption circuit and bounded key size. This can account for both for ABE schemes where there is already such a bound in place such as circuit ABE schemes from LWE or be applied to ABE schemes where no such bound is in place, but one is created for purposes of the transformation. The advantage of this is that the only added assumption is the Learning with Error from lockable obfuscation.

However, in many cases we will want to use systems that don't have such a bound. For example, most key-policy ABE systems that support keys as arbitrary boolean formulas could have decryption circuits of arbitrary depth. Note that bilinear group ABE systems are the only ones that currently support dual system encryption and adaptive security using only polynomial-hardness assumptions [Wat09, LW10, LOS⁺10].

To accommodate this we modify the transformation by adding an extra FHE layer as follows. The encryption algorithm starts as before by choosing a random lock value α and encrypting the lock value under attribute string x using ABE encryption algorithm. However, it next creates a fresh FHE public-private key pair and encrypts the ABE ciphertext under the FHE public key. Next, it creates a lockable obfuscation that is created with lock value α , the message we want to encrypt and a program that takes as input an FHE ciphertext and decrypts using the FHE secret key. The ciphertext consists of the FHE public key, ciphertext and (lockable) obfuscated program. To decrypt simply use the fully homomorphic decryption algorithm to decrypt the ABE ciphertext inside the FHE ciphertext and then submit the resulting ciphertext as input to the obfuscated program.

The security argument is almost the same as before. In the first step ABE security is used to change the ciphertext encrypting α to encrypt the all 0's string. Next the obfuscated program is changed to a simulated program which loses knowledge of the FHE secret key. Then we can finally change the FHE ciphertext encrypting the ABE ciphertext to encrypt the all 0's string. The transformation is provided in Appendix E.

Such a technique will be applicable to many forms of key-policy Attribute-Based encryption including those that use dual system encryption for adaptive security [Wat09, LOS⁺10], ABE for boolean formulas [GPSW06] and ABE for deterministic finite automata [Wat12, Att14].

Unbounded Size, but Bounded Depth. An in-between variant of our first transformation and the methods described above would be to accommodate ABE systems that inherently have decryption circuits that have bounded depth, but not bounded size. A prototypical example is an ABE for circuits scheme from LWE [GVW13, BGG⁺14]. In this case we can do the same construction as above, except use a leveled homomorphic encryption scheme where the level is chosen to match the maximum ABE circuit decryption depth. The advantage of this approach is that realizations of leveled homomorphic encryption are known from the LWE assumption [BV11, BGV12].

Ciphertext Policy Attribute-Based Encryption. An alternate form of Attribute-Based Encryption is ciphertext-policy ABE. In this setting the semantics associated with the keys and ciphertexts get flipped. Here a function f is associated with a ciphertext and an attribute string x with a private key. A user with a private key for x should be able to decrypt a ciphertext associated with f if and only if $f(x) = 1$. In the setting of one-sided predicate encryption we want to hide the function descriptor f associated with the ciphertext from all those users who cannot decrypt it. Using universal circuits (or formulas) one can build [GJPS08] a ciphertext policy-ABE schemes from key-policy ABE scheme, however, this only works for circuits of a priori bounded size.

Using a variant of our approach from Section 5.1 we can handle all ABE schemes that have ciphertexts of bounded size. If we add a layer of fully homomorphic encryption the bounded size restriction goes away and we can handle additional schemes including those that are in the multi-authority setting [Cha07, CC09, LW11] or have ciphertext functions associated with deterministic finite automatas [Wat12, Att14].

5.3 Encrypting to a Hidden Public Key

Our final application of this section is that using lockable encryption one can create a ciphertext that encrypts to a public key, but hides the identity of the public key from anyone who does not have the secret key. Such functionality was considered previously, see for example [WFS03, BBDP01, ABN08] and the references therein. However, in previous cases the hiding would only be among all public keys that shared some common parameter such as all using the same elliptic curve group. Suppose instead that we wanted to encrypt in a way that it is not clear whether it was to someone with a 4096 bit RSA key, a 2048 bit RSA key or using an ECC key.

To do this the encryption algorithm first creates a ciphertext ct that encrypts a randomly chosen lock value α using encryption under the recipient's public key cryptosystem. Then the encryption algorithm creates a program P that takes as input a decryption circuit description and applies it to the ciphertext ct . The final ciphertext is a lockable obfuscation of this program P under the lock α and the message to be transmitted. The decryption circuit description and program size should be padded to include all decryption circuits of a chosen class. To attempt to decrypt such a ciphertext, one must first translate their secret key and decryption algorithm into a decryption circuit and then input this circuit description into the obfuscated program.

6 Separating IND-CPA security and Circular Security

A group of n public keys and ciphertexts is said to form a key cycle [CL01] if the i^{th} ciphertext encrypts the $(i + 1)^{th}$ secret key. A system is said to be n -circular secure if key cycles of length n are indistinguishable

from the case where the ciphertexts consists of encryptions to the all zeros message. Achieving circular security is very desirable in some cases. For instance, a leveled homomorphic encryption scheme that is circular secure implies the existence of a fully homomorphic encryption scheme via Gentry’s [Gen09] bootstrapping technique. As long as the levels of the homomorphic decryption circuit fit within the levels given by evaluation.

There have been several recent works that have shown that circular security, unfortunately, does not come for free. In particular, there exists IND-CPA schemes (in various contexts) that are not circular secure. Achieving such results initially depended on the existence of indistinguishability obfuscation. Subsequent works showed how in some scenarios one could make the desirable trade of assuming learning with errors instead of indistinguishability obfuscation. However, each of these constructions required the introduction of specialized LWE oriented techniques. Here we show that these separation results can instead be achieved by lockable obfuscation. The insight is simply that the original results used iO on programs that were “lockable obfuscation compatible” to begin with. So all we need to do is make very minor tweaks to the programs and then essentially copy the results. (We work things out fully below for completeness.)

The benefits are two fold. First, we are able to achieve some separation results that were not previously known under standard assumptions. And second, the solutions provided via lockable obfuscation are simpler to describe and understand in comparison to direct number theoretic constructions. Of course, direct constructions will likely be more practically efficient, although the importance of this is not clear in the setting of separation results. We now overview our three solutions along with the preceding work in each case. We present our results using the language of cycle testers as introduced by Bishop, Hohenberger and Waters [BHW15].

Cycle testers of length n . Acar et al. [ABBC10] using the SXDH assumption in bilinear groups provided an IND-CPA scheme that was not two circular secure. Subsequently, Koppula, Ramchen and Waters [KRW15] showed separation results for any n cycle length assuming indistinguishability obfuscation. (Concurrently, Marcedone and Orlandi [MO14] did this from VBB.) After these works Bishop, Hohenberger and Waters [BHW15] gave a cycle tester under the LWE assumption for $n = 2$. Finally, Koppula and Waters [KW16] and Alapati and Peikert [AP16] gave cycle testers for any a priori bounded length under the learning with errors and ring learning with errors problem.¹³ Below we will provide a translation of the KRW n cycle testers.

Cycle tester for bit encryption schemes. The above separation results do not apply for when the underlying encryption scheme encrypts single bit messages. Rothblum [Rot13] first showed a separation for bit encryption cycles of length one using a high degree and non-noisy multilinear multilinear paper. Koppula, Ramchen and Waters [KRW15] (in a different solution from above) showed the separation using indistinguishability obfuscation. Very recently, Goyal, Koppula and Waters [GKW17b] gave a bit encryption cycle tester secure under the Learning with Errors assumption. Their cycle tester, however, only worked for secret key encryption while the others were public key. Below we provide a translation of the KRW bit encryption cycle tester.

Cycle testers of unbounded length. Returning to the case of encryption schemes with multi-bit messages Goyal, Koppula and Waters [GKW17a] recently provided a cycle tester that had no a priori bound on cycle length. Their solution requires both indistinguishability obfuscation and leveled (bootstrappable) homomorphic encryption. We show how we can swap out indistinguishability obfuscation for lockable obfuscation.

¹³In their ring LWE solution Alapati and Peikert [AP16] built testers for any cycle length that could depend on the security parameter. In their LWE solution the cycle length had to be constant.

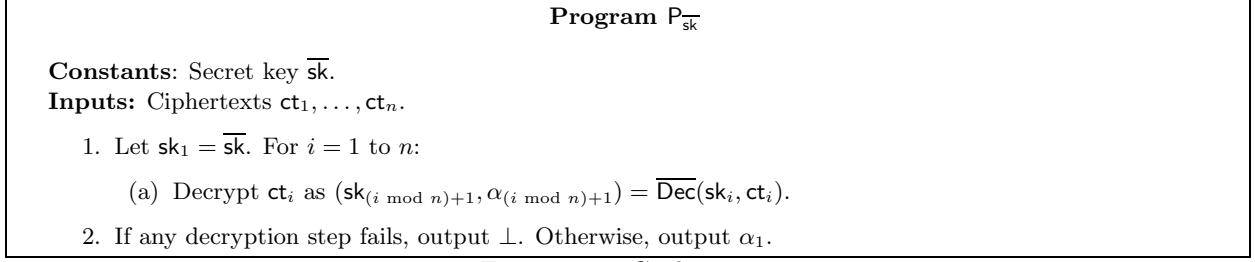


Figure 3: n -Cycle Tester

6.1 Separating IND-CPA Security from n -Circular Security

In this section, we construct a key cycle tester $\Gamma = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Test})$ from any public key encryption scheme and a lockable obfuscator.

Let $\overline{E} = (\overline{Setup}, \overline{Enc}, \overline{Dec})$ be a public key encryption scheme for message space $\{0, 1\}^{2k(\lambda)}$ with decryption circuit of depth $d(\lambda)$, secret key space $\{0, 1\}^{k(\lambda)}$ and ciphertext space $\{0, 1\}^{\ell(\lambda)}$. Also, let $\mathcal{O} = (\text{Obf}, \text{Eval})$ be a lockable obfuscator for 1-bit messages and circuit class $\mathcal{C}_{n \cdot \ell, k, n \cdot d}$ (i.e., the class of depth $n \cdot d(\lambda)$ circuits with $n \cdot \ell(\lambda)$ bit input and $k(\lambda)$ bit output). Below we describe our construction. For notational convenience, let $k = k(\lambda)$.

- **Setup**($1^\lambda, 1^n$) : The setup algorithm outputs public parameters as $pp = (\lambda, n)$.
- **KeyGen**(pp) : Let $pp = (\lambda, n)$. The key generation algorithm runs \overline{Setup} to obtain a public-secret key pair as $(\overline{pk}, \overline{sk}) \leftarrow \overline{Setup}(1^\lambda)$. It also chooses a string α uniformly at random as $\alpha \leftarrow \{0, 1\}^k$.
Next, consider the program $P_{\overline{sk}}$ described in Figure 3. It obfuscates program $P_{\overline{sk}}$ with message 1 and lock α as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, P_{\overline{sk}}, 1, \alpha)$. Finally, it outputs the key pair as $pk = (\overline{pk}, \tilde{P}), sk = (\overline{sk}, \alpha)$.
- **Enc**(pk, m) : Let $pk = (\overline{pk}, \tilde{P})$. The encryption algorithm computes ciphertext as $ct \leftarrow \overline{Enc}(\overline{pk}, m)$.
- **Test**(pk, ct) : Let $pk = (pk_1, \dots, pk_n)$ and $pk_1 = (\overline{pk}_1, \tilde{P}_1)$. The testing algorithm evaluates the obfuscated program \tilde{P}_1 on input ct as $b = \text{Eval}(\tilde{P}_1, ct)$. If $b = \perp$, it outputs 0. Otherwise, it outputs b .
- **Dec**(sk, ct) : Let $sk = (\overline{sk}, \alpha)$. The decryption algorithm outputs $\overline{Dec}(\overline{sk}, ct)$.¹⁴

Correctness. For all $\lambda, n \in \mathbb{N}$, public-secret key pairs $(\overline{pk}_i, \overline{sk}_i) \leftarrow \overline{Setup}(1^\lambda)$ and strings $\alpha_i \leftarrow \{0, 1\}^k$ (for $i \leq n$), a key cycle of length n consists of a sequence of n ciphertexts $ct = (ct_1, \dots, ct_n)$, where each ciphertext ct_i is computed as $ct_i \leftarrow \overline{Enc}(\overline{pk}_i, (\overline{sk}_{(i \bmod n)+1}, \alpha_{(i \bmod n)+1}))$. Also, the program \tilde{P}_1 (which is part of public key pk_1) is computed as $\tilde{P}_1 \leftarrow \text{Obf}(1^\lambda, P_{\overline{sk}_1}, 1, \alpha_1)$.

Let $\tilde{ct} = (\tilde{ct}_1, \dots, \tilde{ct}_n)$ be encryption of zero strings, where each ciphertext \tilde{ct}_i is computed as $\tilde{ct}_i \leftarrow \overline{Enc}(\overline{pk}_i, 0^{2k})$. For correctness of **Test** algorithm, we want to argue that it's advantage in distinguishing a sequence of encryptions of secret keys from encryptions of zeros is non-negligible.

First, note that $P_{\overline{sk}_1}(ct) = \alpha_1$ (by construction). Therefore, by correctness of obfuscation $\text{Eval}(\tilde{P}_1, ct) = 1$. Similarly, $P_{\overline{sk}_1}(\tilde{ct}) \neq \alpha_1$ (with all but negligible probability) as α_1 is chosen uniformly at random and $P_{\overline{sk}_1}$ is independent of string α_1 . This suggests that $\Pr[\text{Eval}(\tilde{P}_1, \tilde{ct}) = \perp] \geq 1 - \text{negl}(\lambda)$.

Therefore, we can conclude that Γ satisfies cycle tester correctness condition as it can distinguish key cycles and all zero encryption with non-negligible probability.

¹⁴We would like to point out that cycle tester framework does not contain a decryption algorithm. However, since it is naturally available in our scheme, therefore we also specify it.

Security. We will now show that the scheme described above achieves IND-CPA security as per Definition 2.4. Formally, we prove the following.

Theorem 6.1. If $\bar{E} = (\bar{\text{Setup}}, \bar{\text{Enc}}, \bar{\text{Dec}})$ is an IND-CPA secure public key encryption scheme for message space $\{0, 1\}^{2k(\lambda)}$, secret key space $\{0, 1\}^{k(\lambda)}$ and ciphertext space $\{0, 1\}^{\ell(\lambda)}$ satisfying Definition 2.4, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for 1-bit messages and circuit class $\mathcal{C}_{n,\ell,k,d}$ satisfying Definition 3.4, then Γ is a secure cycle tester scheme satisfying IND-CPA security as per Definition 2.4.

Proof. Our proof proceeds via a sequence of hybrid games. Let \mathcal{A} be any PPT adversary that wins the IND-CPA security game against Γ with non-negligible advantage. We argue that such an adversary must break security of at least one underlying primitive.

We will first define the sequence of hybrid games, and then show that they are computationally indistinguishable.

Game 1: This game is the original IND-CPA security game described in Definition 2.4.

1. The challenger generates a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. It uniformly samples $\alpha \leftarrow \{0, 1\}^k$, and computes the obfuscation of program $\text{P}_{\overline{\text{sk}}}$ (described in Figure 3) as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{P}_{\overline{\text{sk}}}, 1, \alpha)$. It sends the public key $\text{pk} = (\overline{\text{pk}}, \tilde{P})$ to \mathcal{A} .
2. \mathcal{A} receives pk from the challenger, and computes messages m_0, m_1 . It sends (m_0, m_1) as its challenge messages to the challenger.
3. The challenger chooses bit $b \leftarrow \{0, 1\}$, computes $\text{ct}^* \leftarrow \bar{\text{Enc}}(\overline{\text{pk}}, m_b)$, and sends ct^* to \mathcal{A} .
4. \mathcal{A} receives challenge ciphertext ct^* from the challenger, and outputs its guess b' .
5. \mathcal{A} wins if it guesses correctly, that is if $b = b'$.

Game 2: This game is same as Game 1, except the challenger does not choose the lock α anymore and it simulates the obfuscated program \tilde{P} instead of generating it honestly as an obfuscation of program $\text{P}_{\overline{\text{sk}}}$.

1. The challenger generates a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. **It computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$, where $s = |\text{P}_{\overline{\text{sk}}}|$.** It sends the public key $\text{pk} = (\overline{\text{pk}}, \tilde{P})$ to \mathcal{A} .

We now establish via a sequence of lemmas that the adversary's advantage between Game 1 and 2 is negligible. Let $\text{Adv}_{\mathcal{A}}^i = |\Pr[b' = b] - 1/2|$ denote the advantage of adversary \mathcal{A} in guessing the bit b in Game i . We show via a sequence of lemmas that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ and $\text{Adv}_{\mathcal{A}}^2$ are negligible. Below we discuss our lemmas in detail.

Lemma 6.1. If $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a secure lockable obfuscator, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish an obfuscation of program $\text{P}_{\overline{\text{sk}}}$ with a random lock from a simulated obfuscated program, therefore break security of the obfuscation scheme.

\mathcal{B} runs step 1 as in Game 1, except it does not compute obfuscated program \tilde{P} . Concretely, \mathcal{B} chooses a key pair $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. \mathcal{B} sends the program $\text{P}_{\overline{\text{sk}}}$ along with message 1 to the obfuscation challenger, and receives an obfuscated program \tilde{P} . It sets the public key as $\text{pk} = (\overline{\text{pk}}, \tilde{P})$ and sends pk to \mathcal{A} . It executes steps 2-5 as in Game 1. If the attacker wins (i.e. $b' = b$), then \mathcal{B} guesses '0' to indicate that \tilde{P} was an obfuscation of $\text{P}_{\overline{\text{sk}}}$; otherwise, it guesses '1' to indicate that it was simulated.

Note that if the obfuscation challenger obfuscated $\text{P}_{\overline{\text{sk}}}$ for some lock α , then \mathcal{B} perfectly simulates Game 1 for adversary \mathcal{A} . Otherwise it simulates Game 2 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible, then \mathcal{B} breaks the obfuscation scheme's security with non-negligible advantage. \blacksquare

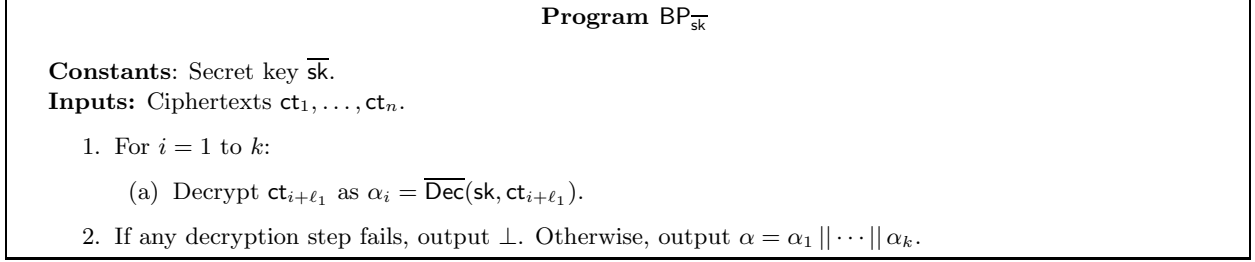


Figure 4: Bit Encryption 1-Cycle Tester

Lemma 6.2. If $\overline{\text{E}} = (\overline{\text{Setup}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ is an IND-CPA secure public key encryption scheme, then for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^2$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}}^2$ is non-negligible. We construct an algorithm \mathcal{B} that can break security of the PKE scheme $\overline{\text{E}}$.

\mathcal{B} receives public key $\overline{\text{pk}}$ from $\overline{\text{E}}$ challenger. It runs step 1 as described in Game 2 with the exception that it uses $\overline{\text{pk}}$ generated by $\overline{\text{E}}$ challenger instead of running the setup algorithm. During challenge phase, \mathcal{B} forwards the challenge messages (m_0, m_1) it receives from \mathcal{A} to $\overline{\text{E}}$ challenger as its challenge messages and receives ct^* as the challenge ciphertext, which it then forwards to \mathcal{A} . Finally, \mathcal{B} outputs the same guess as \mathcal{A} .

We observe that if \mathcal{A} wins (i.e. $b' = b$), then \mathcal{B} also wins because it exactly simulates the view of Game 2 for \mathcal{A} . Therefore if $\text{Adv}_{\mathcal{A}}^2$ is non-negligible, \mathcal{B} must also have non-negligible advantage against $\overline{\text{E}}$ challenger. ■

6.2 Separating IND-CPA Security from 1-Circular Security for Bit Encryption

In this section, we construct a key cycle tester $\Gamma = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Test})$ for 1-bit messages from any public key bit encryption scheme and a lockable obfuscator.

Let $\overline{\text{E}} = (\overline{\text{Setup}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ be a public key bit encryption scheme with secret key space $\{0, 1\}^{\ell_1(\lambda)}$, ciphertext space $\{0, 1\}^{\ell_2(\lambda)}$ and decryption circuit of depth $d(\lambda)$. Also, let $\mathcal{O} = (\text{Obf}, \text{Eval})$ be a lockable obfuscator for 1-bit messages and circuit class $\mathcal{C}_{(\ell_1+k) \cdot \ell_2, k, d}$ (i.e., the class of depth $d(\lambda)$ circuits with $(\ell_1(\lambda) + k(\lambda)) \cdot \ell_2(\lambda)$ bit input and $k(\lambda)$ bit output). Below we describe our construction. For notational convenience, let $k = k(\lambda)$, $\ell_1 = \ell_1(\lambda)$, $\ell_2 = \ell_2(\lambda)$ and $n = k(\lambda) + \ell_1(\lambda)$.

- **Setup(1^λ)** : The setup algorithm runs $\overline{\text{Setup}}$ to obtain a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. It also chooses a string α uniformly at random as $\alpha \leftarrow \{0, 1\}^k$.

Next, consider the program $\text{BP}_{\overline{\text{sk}}}$ described in Figure 4. It obfuscates program $\text{BP}_{\overline{\text{sk}}}$ with message 1 and lock α as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{BP}_{\overline{\text{sk}}}, 1, \alpha)$. Finally, it outputs the key pair as $\text{pk} = (\overline{\text{pk}}, \tilde{P})$, $\text{sk} = (\overline{\text{sk}}, \alpha)$.

- **Enc(pk, m)** : Let $\text{pk} = (\overline{\text{pk}}, \tilde{P})$. The encryption algorithm computes ciphertext as $\text{ct} \leftarrow \overline{\text{Enc}}(\overline{\text{pk}}, m)$.
- **Test(pk, ct)** : Let $\text{pk} = (\overline{\text{pk}}, \tilde{P})$. The testing algorithm evaluates the obfuscated program \tilde{P} on input ct as $b = \text{Eval}(\tilde{P}, \text{ct})$. If $b = \perp$, it outputs 0. Otherwise, it outputs b .
- **Dec(sk, ct)** : Let $\text{sk} = (\overline{\text{sk}}, \alpha)$. The decryption algorithm outputs $\overline{\text{Dec}}(\overline{\text{sk}}, \text{ct})$.

Correctness. For every λ , public-secret key pair $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$ and string $\alpha \leftarrow \{0, 1\}^k$, a key cycle consists of a sequence of n ciphertexts $\mathbf{ct} = (\text{ct}_1, \dots, \text{ct}_n)$, where each ciphertext ct_i is computed as $\text{ct}_i \leftarrow \overline{\text{Enc}}(\overline{\text{pk}}, \text{sk}_i)$, $\text{sk} = (\text{sk}_t, \alpha)$ and sk_i is i^{th} bit of sk . Also, the program \tilde{P} (which is part of public key pk) is computed as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{BP}_{\overline{\text{sk}}}, 1, \alpha)$.

Let $\tilde{\mathbf{ct}} = (\tilde{\text{ct}}_1, \dots, \tilde{\text{ct}}_n)$ be encryptions of zero, where each ciphertext $\tilde{\text{ct}}_i$ is computed as $\tilde{\text{ct}}_i \leftarrow \overline{\text{Enc}}(\overline{\text{pk}}, 0)$. For correctness of **Test** algorithm, we want to argue that it's advantage in distinguishing a sequence of encryptions of secret key bits from encryptions of zeros is non-negligible.

First, note that $\text{BP}_{\overline{\text{sk}}}(\mathbf{ct}) = \alpha$ and $\text{BP}_{\overline{\text{sk}}}(\tilde{\mathbf{ct}}) = 0^k$ (by construction). Therefore, by correctness of obfuscation $\text{Eval}(\tilde{P}, \mathbf{ct}) = 1$. Also, $\text{BP}_{\overline{\text{sk}}}(\tilde{\mathbf{ct}}) \neq \alpha$ (with all but negligible probability) as α is chosen uniformly at random. This suggests that $\Pr[\text{Eval}(\tilde{P}, \tilde{\mathbf{ct}}) = \perp] \geq 1 - \text{negl}(\lambda)$.

Therefore, we can conclude that Γ satisfies bit encryption cycle tester correctness condition as it can distinguish a key cycle and all zero encryption with non-negligible probability.

Security. We will now show that the scheme described above achieves IND-CPA security as per Definition 2.4. Formally, we prove the following.

Theorem 6.2. If $\overline{\text{E}} = (\overline{\text{Setup}}, \overline{\text{Enc}}, \overline{\text{Dec}})$ is an IND-CPA secure public key bit encryption scheme with secret key space $\{0, 1\}^{\ell_1(\lambda)}$ and ciphertext space $\{0, 1\}^{\ell_2(\lambda)}$ satisfying Definition 2.4, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for 1-bit messages and circuit class $\mathcal{C}_{(\ell_1+k) \cdot \ell_2, k, d}$ satisfying Definition 3.4, then Γ is a secure bit encryption cycle tester scheme satisfying IND-CPA security as per Definition 2.4.

Proof. Our proof proceeds via a sequence of hybrid games. Let \mathcal{A} be any PPT adversary that wins the IND-CPA security game against Γ with non-negligible advantage. We argue that such an adversary must break security of at least one underlying primitive.

The proof of this theorem is identical to proof of Theorem 6.1. Therefore, we only sketch the sequence of hybrid games.

Game 1: This game is the original IND-CPA security game described in Definition 2.4.

1. The challenger generates a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. It uniformly samples $\alpha \leftarrow \{0, 1\}^k$, and computes the obfuscation of program $\text{BP}_{\overline{\text{sk}}}$ (described in Figure 4) as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{BP}_{\overline{\text{sk}}}, 1, \alpha)$. It sends the public key $\text{pk} = (\overline{\text{pk}}, \tilde{P})$ to \mathcal{A} .
2. Next, the challenger chooses bit $b \leftarrow \{0, 1\}$, computes $\text{ct}^* \leftarrow \overline{\text{Enc}}(\overline{\text{pk}}, b)$, and sends ct^* to \mathcal{A} .
3. \mathcal{A} receives challenge ciphertext ct^* from the challenger, and outputs its guess b' . \mathcal{A} wins if it guesses correctly, that is if $b = b'$.

Game 2: This game is same as Game 1, except the challenger does not choose the lock α anymore and it simulates the obfuscated program \tilde{P} instead of generating it honestly as an obfuscation of program $\text{BP}_{\overline{\text{sk}}}$.

1. The challenger generates a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{Setup}}(1^\lambda)$. **It computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$, where $s = |\text{BP}_{\overline{\text{sk}}}|$.** It sends the public key $\text{pk} = (\overline{\text{pk}}, \tilde{P})$ to \mathcal{A} . ■

6.3 Creating an Unbounded Public Key Cycle Tester

In this section, we construct a cycle tester $\Gamma = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Test})$ for unbounded length key cycles from any bootstrappable (leveled) homomorphic encryption scheme and a lockable obfuscator.

Let $\text{LHE} = (\text{LHE.Setup}, \text{LHE.Enc}, \text{LHE.Eval}, \text{LHE.Dec})$ be a bootstrappable homomorphic bit encryption scheme for message space $\{0, 1\}^{\ell(\lambda)}$ with decryption circuit of depth $d(\lambda)$, ciphertexts of length $\ell_{\text{ct}}(\lambda)$ and

Program $\text{UP}_{\text{lhe.sk}}$

Constants: Secret key lhe.sk .

Inputs: Ciphertext ct .

1. Decrypt ct as $(\text{sk}', \alpha) = \text{LHE.Dec}(\text{lhe.sk}, \text{ct})$. Output α .

Figure 5: Tester for Unbounded Length Key Cycles

secret keys of length $s(\lambda)$. Also, let $\mathcal{O} = (\text{Obf}, \text{Eval})$ be a lockable obfuscator for 1-bit messages and circuit class $\mathcal{C}_{\ell_{\text{ct}}, k, d}$ (i.e., the class of depth $d(\lambda)$ circuits with $\ell_{\text{ct}}(\lambda)$ bit input and $k(\lambda)$ bit output). Below we describe our construction. For notational convenience, let $k = k(\lambda)$, $s = s(\lambda)$, $d = d(\lambda)$, $\ell_{\text{ct}} = \ell_{\text{ct}}(\lambda)$ and $\ell = \ell(\lambda)$.

- **Setup(1^λ)** : The setup algorithm runs LHE.Setup to obtain a public-secret key pair as $(\text{lhe.pk}, \text{lhe.sk}, \text{lhe.ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^d)$. It also chooses a string α uniformly at random as $\alpha \leftarrow \{0, 1\}^k$.
Next, consider the program $\text{UP}_{\text{lhe.sk}}$ described in Figure 5. It obfuscates program $\text{UP}_{\text{lhe.sk}}$ with message 1 and lock α as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{UP}_{\text{lhe.sk}}, 1, \alpha)$. Finally, it outputs the key pair as $\text{pk} = (\text{lhe.pk}, \text{lhe.ek}, \tilde{P})$, $\text{sk} = (\text{lhe.sk}, \alpha)$.
- **Enc(pk, m)** : Let $\text{pk} = (\text{lhe.pk}, \text{lhe.ek}, \tilde{P})$. The encryption algorithm computes ciphertext as $\text{ct} \leftarrow \text{LHE.Enc}(\text{lhe.pk}, m)$.
- **Test($1^n, \text{pk}, \text{ct}$)** : Let $\text{pk} = (\text{pk}_1, \dots, \text{pk}_n)$, $\text{ct} = (\text{ct}_1, \dots, \text{ct}_n)$ and $\text{pk}_i = (\text{lhe.pk}_i, \text{lhe.ek}_i, \tilde{P}_i)$ (for all $i \leq n$). Also, let circuit C_{ct} be a circuit which takes as input $s + k$ bits and decrypts ciphertext ct using the first s bits of the input, i.e. $C_{\text{ct}}(x || y) = \text{LHE.Dec}(x, \text{ct})$ where $x \in \{0, 1\}^s, y \in \{0, 1\}^k$. The tester algorithm runs as follows.
First, it sets $\text{ct}'_n = \text{ct}_n$. For $i = n-1$ to $i = 1$, it computes ciphertexts ct'_i as $\text{ct}'_i = \text{LHE.Eval}(\text{lhe.ek}_i, C_{\text{ct}'_{i+1}}, \text{ct}_i)$.
Next, it evaluates the obfuscated program \tilde{P}_1 on input ct'_1 as $b = \text{Eval}(\tilde{P}_1, \text{ct}'_1)$. Finally, it outputs b .
- **Dec(sk, ct)** : Let $\text{sk} = (\text{lhe.sk}, \alpha)$. The decryption algorithm outputs $\text{LHE.Dec}(\text{lhe.sk}, \text{ct})$.

Correctness. For all $\lambda, n \in \mathbb{N}$, LHE key pairs $(\text{lhe.pk}_i, \text{lhe.sk}_i, \text{lhe.ek}_i) \leftarrow \text{LHE.Setup}(1^\lambda, 1^d)$ and strings $\alpha_i \leftarrow \{0, 1\}^k$ (for $i \leq n$), a key cycle of length n consists of a sequence of n ciphertexts $\text{ct} = (\text{ct}_1, \dots, \text{ct}_n)$, where each ciphertext ct_i is computed as $\text{ct}_i \leftarrow \overline{\text{Enc}}(\text{lhe.pk}_i, (\text{lhe.sk}_{(i \bmod n)+1}, \alpha_{(i \bmod n)+1}))$. Also, the program \tilde{P}_1 (which is part of public key pk_1) is computed as $\tilde{P}_1 \leftarrow \text{Obf}(1^\lambda, \text{UP}_{\text{lhe.sk}_1}, 1, \alpha_1)$.

Let $\tilde{\text{ct}} = (\tilde{\text{ct}}_1, \dots, \tilde{\text{ct}}_n)$ be encryption of zero strings, where each ciphertext $\tilde{\text{ct}}_i$ is computed as $\tilde{\text{ct}}_i \leftarrow \overline{\text{Enc}}(\text{pk}_i, 0^{s+k})$. For correctness of Test algorithm, we want to argue that it's advantage in distinguishing a sequence of encryptions of secret keys from encryptions of zeros is non-negligible.

First, note that $\text{LHE.Dec}(\text{lhe.sk}_n, \text{ct}_n) = (\text{sk}_1, \alpha_1)$. Therefore, $\text{ct}'_{n-1} = \text{LHE.Eval}(\text{lhe.ek}_{n-1}, C_{\text{ct}'_n}, \text{ct}_{n-1})$ is an encryption of (sk_1, α_1) under key lhe.pk_{n-1} , where $\text{ct}'_n = \text{ct}_n$. Therefore, the top-down iterative LHE evaluation during testing would finally result in ciphertext ct'_1 that will be an encryption of (sk_1, α_1) under key lhe.pk_1 . We could prove this by an inductive argument over cycle length n . The base case $n = 1$ follows directly from the construction. The inductive hypothesis would be that the test algorithm successfully reduces a length k key cycle to a length 1 key cycle. The induction step can be proven by first reducing a length k key cycle to a length $k - 1$ key cycle (which follows from the correctness of the LHE scheme) and then applying the hypothesis. Finally, the test algorithm evaluates the obfuscated program \tilde{P}_1 on ct'_1 which is encryption of (sk_1, α_1) under key lhe.pk_1 . Therefore, by correctness of obfuscation scheme $\text{Eval}(\tilde{P}_1, \text{ct}'_1) = 1$. In other words, we can claim that the test algorithm always outputs 1 in this case (i.e., with probability 1).

Similarly, $\text{UP}_{\text{lhe.sk}_1}(\tilde{\text{ct}}'_1) \neq \alpha_1$ (with all but negligible probability) as α_1 is chosen uniformly at random and neither $\text{UP}_{\text{lhe.sk}_1}$ nor any of the ciphertexts ct_i depends on the string α_1 . This suggests that $\Pr[\text{Eval}(\tilde{P}_1, \tilde{\text{ct}}'_1) = \perp] \geq 1 - \text{negl}(\lambda)$.

Therefore, we can conclude that Γ satisfies cycle tester correctness condition as it can distinguish key cycles and all zero encryption with non-negligible probability.

The IND-CPA proof of this scheme is identical to the proof of Theorem 6.1.

7 Uninstantiability of the Fujisaki-Okamoto and Related Transformations

We now move to our second set of separation results where we use lockable obfuscation to show random oracle uninstantiability [CGH98] for a certain family of transformations. These include: the Bellare et al. [BHSV98] transformation from an IND-CPA scheme to an injective trapdoor function, the well known Fujisaki-Okamoto [FO99a, FO99b] transformations from IND-CPA and the deterministic encryption construction of Bellare, Boldyreva and O’Neill [BBO07].

All of the constructions follow a similar paradigm where they encrypt a string x under random coins determined from $H(x)$. (How the string x is construed varies somewhat between the schemes.) The works above show that if H is presented as an oracle access to a random function, the transformation results in a secure scheme under the relevant definition.

An important question is what can be said about the security of these transformations when H is instantiated by a concrete hash function. Assuming indistinguishability obfuscation Brzuska, Farshim and Mittelbach [BFM15] provide negative results on the instantiability of these transformations. In particular assuming iO for circuits they provide encryption schemes that are IND-CPA secure, but where the transformation is insecure when instantiated with any hash function up to an a priori bounded size. Furthermore, if one assumes iO for Turing Machines with unbounded input, the result holds for hash functions of any size.¹⁵ The main idea of BFM is to create an encryption scheme that includes an obfuscated program which can be used to break security when the description of H is known.

Here we show that such results can be achieved from lockable obfuscation and thus the Learning with Errors assumption. More specifically, the uninstantiability results for a priori size bounded hash functions follows from lockable obfuscation and the unrestricted results follow from lockable obfuscation combined with fully homomorphic encryption. The main insight into the result is simply that the program breaking in BFM falls within a pattern that matches lockable obfuscation. Once this connection is made the technical argument follows readily. In addition, we again see that for iO programs that match the pattern of lockable obfuscation, fully homomorphic encryption can serve to give the “Turing Machine version” of such programs.

Below we give a proof of uninstantiability for an IND-CPA encryption scheme for the second Fujisaki-Okamoto [FO99b] transformation. We do so in the unrestricted size case using lockable obfuscation along with fully homomorphic encryption. We believe uninstantiability results for the other schemes mentioned above follow in a very similar manner.

7.1 The Fujisaki-Okamoto Transformation

The Fujisaki-Okamoto (FO) transformation is a technique to convert an IND-CPA secure public key encryption scheme into an IND-CCA secure public key encryption scheme.¹⁶ The transformation relies on a public key encryption scheme, a symmetric encryption scheme and two hash functions (which are modeled as two independent random oracles).

Let $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ be a public key encryption scheme, $\text{SKE} = (\text{SKE.Setup}, \text{SKE.Enc}, \text{SKE.Dec})$ be a (deterministic) symmetric encryption scheme, and H_1, H_2 be two hash functions. As per FO transformation, an encryption of message m under randomness r is computed as

$$\text{PKE.Enc}(\text{pke.pk}, r; H_1(m || r)), \text{SKE.Enc}(\text{ske.sk} = H_2(r), m).$$

¹⁵We note that constructing Turing Machine iO for *unbounded* length inputs from standard iO remains an open problem.

¹⁶The FO transformation only requires the starting PKE scheme to satisfy one-way security, which is a slightly weaker notion than IND-CPA.

More formally, the FO transformed scheme $\text{FO}_{\text{PKE,SKE}}^{\mathcal{H},\tilde{\mathcal{H}}} = (\text{Setup}, \text{Enc}, \text{Dec})$ is described as follows.

- $\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$. It chooses two hash functions H_1, H_2 as $H_1 \leftarrow \mathcal{H}_\lambda, H_2 \leftarrow \tilde{\mathcal{H}}_\lambda$. It also samples a public-secret key pair $(\text{pke.pk}, \text{pke.sk})$ for PKE scheme as $(\text{pke.pk}, \text{pke.sk}) \leftarrow \text{PKE.Setup}(1^\lambda)$. It outputs the public-secret keys as $\text{pk} = (\text{pke.pk}, H_1, H_2)$ and $\text{sk} = (\text{pke.pk}, H_1, H_2, \text{pke.sk})$.
- $\text{Enc}(\text{pk}, m; r) \rightarrow \text{ct}$. Let $\text{pk} = (\text{pke.pk}, H_1, H_2)$. It computes ciphertext ct_1, ct_2 as $\text{ct}_1 = \text{PKE.Enc}(\text{pke.pk}, r; H_1(m || r))$ and $\text{ct}_2 = \text{SKE.Enc}(H_2(r), m)$. It outputs the ciphertext as $\text{ct} = (\text{ct}_1, \text{ct}_2)$.
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$ or \perp . Let $\text{sk} = (\text{pke.pk}, H_1, H_2, \text{pke.sk})$ and $\text{ct} = (\text{ct}_1, \text{ct}_2)$. It decrypts ct_1 as $r = \text{PKE.Dec}(\text{pke.sk}, \text{ct}_1)$. Next, it decrypts ct_2 as $m = \text{SKE.Dec}(H_2(r), \text{ct}_2)$. Finally, it checks if $\text{ct}_1 = \text{PKE.Enc}(\text{pke.pk}, r; H_1(m || r))$. If the check succeeds it outputs m , otherwise it outputs \perp .

In this work, we show that FO transform is uninstantiable in the standard model. Concretely, we prove the following.

Theorem 7.1. If $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ is a fully homomorphic encryption scheme satisfying Definition 2.4, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for circuit class $\{\mathcal{C}_{\ell, \tilde{\ell}, \ell, d}\}_\lambda$ satisfying Definition 3.4, then there exists a public key encryption scheme $\text{PKE} = (\text{PKE.Setup}, \text{PKE.Enc}, \text{PKE.Dec})$ such that for every hash function family $\mathcal{H} = \{\mathcal{H}_\lambda\}_\lambda$ and $\tilde{\mathcal{H}} = \{\tilde{\mathcal{H}}_\lambda\}_\lambda$, symmetric key encryption scheme $\text{SKE} = (\text{SKE.Setup}, \text{SKE.Enc}, \text{SKE.Dec})$, the corresponding FO transformed scheme $\text{FO}_{\text{PKE,SKE}}^{\mathcal{H},\tilde{\mathcal{H}}}$ is not secure.¹⁷

We prove the above theorem by constructing such a public key encryption scheme $\text{PKE} = (\text{Setup}, \text{Enc}, \text{Dec})$. Consider a public key encryption scheme $\overline{\text{PKE}} = (\overline{\text{PKE.Setup}}, \overline{\text{PKE.Enc}}, \overline{\text{PKE.Dec}})$ for ℓ -bit messages, a fully homomorphic encryption scheme $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ for 1-bit messages with decryption circuit of depth $d(\lambda)$ and ciphertext space $\{0, 1\}^{\tilde{\ell}(\lambda)}$, and a lockable obfuscator scheme for ℓ -bit messages. Below we describe our construction. For ease of exposition, assume that the FHE setup, FHE encryption, PKE encryption and lockable obfuscator also take ℓ bits of randomness as inputs.

- $\text{Setup}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$. It generates a PKE key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{PKE.Setup}}(1^\lambda)$, and sets public-secret key pair as $(\text{pk}, \text{sk}) = (\overline{\text{pk}}, \overline{\text{sk}})$.
- $\text{Enc}(\text{pk}, m; r) \rightarrow \text{ct}$. Let $\text{pk} = \overline{\text{pk}}$ and $r = r_0 || r_1 || r_2 || r_3 || r_4$, where each r_i is ℓ bits long. It samples an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) = \text{FHE.Setup}(1^\lambda; r_0)$ and computes ciphertext ct_1 as $\text{ct}_1 = \overline{\text{PKE.Enc}}(\overline{\text{pk}}, m; r_1)$. It encrypts message m under FHE public key as $\text{ct}_2 = \text{FHE.Enc}(\text{fhe.pk}, m; r_2)$.¹⁸
It also obfuscates program $\text{P}_{\text{fhe.sk}}$ with message m , lock r_4 and randomness r_3 as $\tilde{P} = \text{Obf}(1^\lambda, \text{P}_{\text{fhe.sk}}, m, r_4; r_3)$. Finally, it outputs the ciphertext as $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$.
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$. Let $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$. It decrypts ct_1 as $m = \overline{\text{PKE.Dec}}(\text{sk}, \text{ct}_1)$ and outputs m .

First, we prove that the above scheme is IND-CPA secure.

Theorem 7.2. If $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ is a fully homomorphic encryption scheme satisfying Definition 2.4, $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for circuit class $\{\mathcal{C}_{\ell, \tilde{\ell}, \ell, d}\}_\lambda$ satisfying Definition 3.4, and $\overline{\text{PKE}} = (\overline{\text{PKE.Setup}}, \overline{\text{PKE.Enc}}, \overline{\text{PKE.Dec}})$ is a public key encryption scheme for ℓ -bit messages satisfying Definition 2.4, then the scheme $\text{PKE} = (\text{Setup}, \text{Enc}, \text{Dec})$ is IND-CPA secure.

Proof. Our proof proceeds via a sequence of hybrid games. Let \mathcal{A} be any PPT adversary that wins the IND-CPA security game against PKE with non-negligible advantage. We argue that such an adversary must break security of at least one underlying primitive.

We will first define the sequence of hybrid games, and then show that they are computationally indistinguishable.

¹⁷The FO transformed scheme $\text{FO}_{\text{PKE,SKE}}^{\mathcal{H},\tilde{\mathcal{H}}}$ does not even achieve IND-CPA security.

¹⁸Note that FHE scheme is a bit encryption scheme, therefore it encrypts one bit at a time. However, we overload the notation and use FHE.Enc to encrypt multi-bit messages as well. So, ct_2 is a sequence of ℓ FHE ciphertexts.

Program $P_{\text{fhe.sk}}$

Constants: Secret key fhe.sk .

Inputs: Ciphertexts $\text{ct}_1, \dots, \text{ct}_\ell$.

1. Decrypt ct_i as $\alpha_i = \text{FHE.Dec}(\text{fhe.sk}, \text{ct}_i)$.
2. If any decryption step fails, output \perp . Otherwise, output $\alpha_1 \parallel \dots \parallel \alpha_\ell$.

Figure 6: FO

Game 1: This game is the original IND-CPA security game described in Definition 2.4.

1. The challenger generates a public-secret key pair as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{PKE.Setup}}(1^\lambda)$. It sends the public key $\text{pk} = \overline{\text{pk}}$ to \mathcal{A} .
2. \mathcal{A} receives pk from the challenger, and computes messages m_0, m_1 . It sends (m_0, m_1) as its challenge messages to the challenger.
3. The challenger chooses bit $b \leftarrow \{0, 1\}$. It samples an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and computes ciphertexts ct_1, ct_2 as $\text{ct}_1 \leftarrow \overline{\text{PKE.Enc}}(\overline{\text{pk}}, m_b), \text{ct}_2 \leftarrow \text{FHE.Enc}(\text{fhe.pk}, m_b)$. It chooses a random ℓ bit string r and obfuscates program $P_{\text{fhe.sk}}$ with message m and lock r as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, P_{\text{fhe.sk}}, m_b, r)$. Finally, it outputs the ciphertext as $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$. Finally, it sends ct to \mathcal{A} .
4. \mathcal{A} receives challenge ciphertext ct from the challenger, and outputs its guess b' .
5. \mathcal{A} wins if it guesses correctly, that is if $b = b'$.

Game 2: This game is same as Game 1, except the challenger does not choose the lock r anymore and it simulates the obfuscated program \tilde{P} instead of generating it honestly.

1. The challenger chooses bit $b \leftarrow \{0, 1\}$. It samples an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and computes ciphertexts ct_1, ct_2 as $\text{ct}_1 \leftarrow \overline{\text{PKE.Enc}}(\overline{\text{pk}}, m_b), \text{ct}_2 \leftarrow \text{FHE.Enc}(\text{fhe.pk}, m_b)$. **It computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O.Sim}(1^\lambda, 1^s, 1^\ell)$, where $s = |P_{\text{fhe.sk}}|$.** Finally, it outputs the ciphertext as $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$. Finally, it sends ct to \mathcal{A} .

Game 3: This game is same as Game 2, except the challenger computes ct_2 as encryption of all zeros string.

1. The challenger chooses bit $b \leftarrow \{0, 1\}$. It samples an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and **computes ciphertexts ct_1, ct_2 as $\text{ct}_1 \leftarrow \overline{\text{PKE.Enc}}(\overline{\text{pk}}, m_b), \text{ct}_2 \leftarrow \text{FHE.Enc}(\text{fhe.pk}, 0^\ell)$.** It computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O.Sim}(1^\lambda, 1^s, 1^\ell)$, where $s = |P_{\text{fhe.sk}}|$. Finally, it outputs the ciphertext as $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$. Finally, it sends ct to \mathcal{A} .

We now establish via a sequence of claims that the adversary's advantage between any two consecutive games is negligible. Let $\text{Adv}_{\mathcal{A}}^i = |\Pr[b' = b] - 1/2|$ denote the advantage of adversary \mathcal{A} in guessing the bit b in Game i . We show via a sequence of lemmas that $|\text{Adv}_{\mathcal{A}}^i - \text{Adv}_{\mathcal{A}}^{i+1}|$ (for $i = 1, 2$) and $\text{Adv}_{\mathcal{A}}^3$ are negligible. Below we discuss our lemmas in detail.

Lemma 7.1. If $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a secure lockable obfuscator, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish an obfuscation of program $P_{\text{fhe.sk}}$ with a random lock from a simulated obfuscated program, therefore break security of the obfuscation scheme.

\mathcal{B} runs step 1 as in Game 1, i.e. it samples a PKE public-secret key pair $(\overline{\text{pk}}, \overline{\text{sk}})$ and sends $\overline{\text{pk}}$ to \mathcal{A} . Next, it receives two challenge messages (m_0, m_1) from \mathcal{A} . \mathcal{B} chooses a FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and random bit b . \mathcal{B} sends the program $\text{P}_{\text{fhe.sk}}$ along with message m_b to the obfuscation challenger, and receives an obfuscated program \tilde{P} . It encrypts message m_b as $\text{ct}_1 \leftarrow \overline{\text{PKE}}.\overline{\text{Enc}}(\overline{\text{pk}}, m_b)$, $\text{ct}_2 \leftarrow \text{FHE.Enc}(\text{fhe.pk}, m_b)$. Finally, it sends the challenge ciphertext $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$ to \mathcal{A} . Next, \mathcal{B} receives \mathcal{A} 's guess b' . If the attacker wins (i.e. $b' = b$), then \mathcal{B} guesses '0' to indicate that \tilde{P} was an obfuscation of $\text{P}_{\text{fhe.sk}}$; otherwise, it guesses '1' to indicate that it was simulated.

Note that if the obfuscation challenger obfuscated $\text{P}_{\text{fhe.sk}}$ for some lock r , then \mathcal{B} perfectly simulates Game 1 for adversary \mathcal{A} . Otherwise it simulates Game 2 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible, then \mathcal{B} breaks the obfuscation scheme's security with non-negligible advantage. \blacksquare

Lemma 7.2. If $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ is a fully homomorphic encryption scheme, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish encryption of message m_b and an all zeros string, therefore break security of the FHE scheme.

\mathcal{B} receives an FHE key pair $(\text{fhe.pk}, \text{fhe.ek})$ from the challenger. \mathcal{B} runs step 1 as in Game 1, i.e. it samples a PKE public-secret key pair $(\overline{\text{pk}}, \overline{\text{sk}})$ and sends $\overline{\text{pk}}$ to \mathcal{A} . Next, it receives two challenge messages (m_0, m_1) from \mathcal{A} . \mathcal{B} computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1^\ell)$, where $s = |\text{P}_{\text{fhe.sk}}|$. It chooses a random bit b and encrypts message m_b as $\text{ct}_1 \leftarrow \overline{\text{PKE}}.\overline{\text{Enc}}(\overline{\text{pk}}, m_b)$. It sends $(m_b, 0^\ell)$ to the FHE challenger as its challenge messages. \mathcal{B} receives ct^* from the challenger and sets $\text{ct}_2 = \text{ct}^*$. Finally, it sends the challenge ciphertext $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$ to \mathcal{A} . Next, \mathcal{B} receives \mathcal{A} 's guess b' . If the attacker wins (i.e. $b' = b$), then \mathcal{B} sends 0 as its guess (i.e., m_b was encrypted), otherwise it sends 1 (i.e., 0^ℓ was encrypted) as its guess to the FHE challenger.

Note that if the FHE challenger encrypted m_b , then \mathcal{B} perfectly simulates Game 2 for adversary \mathcal{A} . Otherwise it simulates Game 3 for \mathcal{A} . Therefore, if $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible, then \mathcal{B} breaks the FHE scheme's security with non-negligible advantage. \blacksquare

Lemma 7.3. If $\overline{\text{PKE}} = (\overline{\text{PKE.Setup}}, \overline{\text{PKE.Enc}}, \overline{\text{PKE.Dec}})$ is an IND-CPA secure public key encryption scheme, then for all PPT adversaries \mathcal{A} , $\text{Adv}_{\mathcal{A}}^3$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}}^3$ is non-negligible. We construct an algorithm \mathcal{B} that can break security of the PKE scheme $\overline{\text{PKE}}$.

\mathcal{B} receives public key $\overline{\text{pk}}$ from the PKE challenger. It sends $\overline{\text{pk}}$ to \mathcal{A} . Next, it receives two challenge messages (m_0, m_1) from \mathcal{A} . \mathcal{B} chooses a FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$. It computes the obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1^\ell)$, where $s = |\text{P}_{\text{fhe.sk}}|$. It also computes ciphertext ct_2 as $\text{ct}_2 \leftarrow \text{FHE.Enc}(\text{fhe.pk}, 0^\ell)$. It sends (m_0, m_1) to the PKE challenger as its challenge messages. \mathcal{B} receives ct^* from the challenger and sets $\text{ct}_1 = \text{ct}^*$. Finally, it sends the challenge ciphertext $\text{ct} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$ to \mathcal{A} . Next, \mathcal{B} receives \mathcal{A} 's guess b' . Finally, \mathcal{B} outputs the same guess as \mathcal{A} .

We observe that if \mathcal{A} wins (i.e. $b' = b$), then \mathcal{B} also wins because it exactly simulates the view of Game 3 for \mathcal{A} . Therefore if $\text{Adv}_{\mathcal{A}}^3$ is non-negligible, \mathcal{B} must also have non-negligible advantage against $\overline{\text{PKE}}$ challenger. \blacksquare

This concludes the proof of IND-CPA security of PKE. \blacksquare

Uninstantiability of FO Transformation. Now, we show that for every hash function family $\mathcal{H} = \{\mathcal{H}_\lambda\}_\lambda$ and $\tilde{\mathcal{H}} = \{\tilde{\mathcal{H}}_\lambda\}_\lambda$, symmetric key encryption scheme $\text{SKE} = (\text{SKE.Setup}, \text{SKE.Enc}, \text{SKE.Dec})$, the FO transformed scheme $\text{FO}_{\text{PKE}, \text{SKE}}^{\mathcal{H}, \tilde{\mathcal{H}}}$ is not IND-CPA secure. Below we describe an attacker \mathcal{A} .

1. The challenger chooses two hash functions H_1, H_2 as $H_1 \leftarrow \mathcal{H}_\lambda, H_2 \leftarrow \tilde{\mathcal{H}}_\lambda$. It samples a public-secret key pair $(\overline{\text{pk}}, \overline{\text{sk}})$ for PKE scheme $\overline{\text{PKE}}$ as $(\overline{\text{pk}}, \overline{\text{sk}}) \leftarrow \overline{\text{PKE}}.\overline{\text{Setup}}(1^\lambda)$. It sends the public key $\overline{\text{pk}}$ and hash function descriptions H_1, H_2 to adversary \mathcal{A} .
2. \mathcal{A} chooses two ℓ -bit messages $m_0, m_1 (\neq m_0)$ uniformly at random and sends them as its challenge messages.
3. The challenger chooses a random bit b and an ℓ -bit string $s \leftarrow \{0, 1\}^\ell$. It computes $r = H_1(m_b || s)$. Let $r = r_0 || r_1 || r_2 || r_3 || r_4$, where each r_i is ℓ bits long. It samples an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) = \text{FHE.Setup}(1^\lambda; r_0)$ and computes ciphertext ct_1 as $\text{ct}_1 = \overline{\text{PKE}}.\overline{\text{Enc}}(\overline{\text{pk}}, s; r_1)$. It also encrypts s under FHE public key as $\text{ct}_2 = \text{FHE.Enc}(\text{fhe.pk}, s; r_2)$. It obfuscates program $\text{P}_{\text{fhe.sk}}$ with message s , lock r_4 and randomness r_3 as $\tilde{P} = \text{Obf}(1^\lambda, \text{P}_{\text{fhe.sk}}, s, r_4; r_3)$. This corresponds to the PKE part of the ciphertext, i.e. $\text{ct}_{\text{PKE}} = (\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P})$. Next, it encrypts message m_b using SKE scheme as $\text{ct}_{\text{SKE}} = \text{SKE.Enc}(H_2(s), m_b)$. Finally, it sends $\text{ct} = (\text{ct}_{\text{PKE}}, \text{ct}_{\text{SKE}})$ as the challenge ciphertext to \mathcal{A} .
4. \mathcal{A} receives the ciphertext $\text{ct} = ((\text{fhe.ek}, \text{ct}_1, \text{ct}_2, \tilde{P}), \text{ct}_{\text{SKE}})$ from the challenger. It first homomorphically evaluates the hash function H_2 on ct_2 and then decrypts the resulting string to decrypt ciphertext ct_{SKE} . Concretely, it computes $\text{ct}' = \text{FHE.Eval}(\text{fhe.ek}, \text{SKE.Dec}(H_2(\cdot), \text{ct}_{\text{SKE}}), \text{ct}_2)$. Next, it homomorphically evaluates hash function H_1 on ct', ct_2 as $\tilde{\text{ct}} = \text{FHE.Eval}(\text{fhe.ek}, H_1(\cdot), (\text{ct}', \text{ct}_2))$. Let $\tilde{\text{ct}} = \tilde{\text{ct}}_1, \dots, \tilde{\text{ct}}_{5\ell}$. It evaluates program \tilde{P} on ciphertexts $\tilde{\text{ct}}_{4\ell+1}, \dots, \tilde{\text{ct}}_{5\ell}$ as $t = \text{Eval}(\tilde{P}, (\tilde{\text{ct}}_{4\ell+1}, \dots, \tilde{\text{ct}}_{5\ell}))$. It decrypts ciphertext ct_{SKE} as $m = \text{SKE.Dec}(H_2(t), \text{ct}_{\text{SKE}})$. If $m = m_0$ or $m = m_1$, then it guesses accordingly. Otherwise, it guesses randomly.

We claim that \mathcal{A} always guesses b correctly. In other words, \mathcal{A} wins with non-negligible advantage. First, note that (by correctness of FHE evaluation) ciphertext ct' encrypts challenge message m_b as $H_2(s)$ is the SKE secret key used for encrypting m_b . Similarly, ciphertext $\tilde{\text{ct}}$ is an encryption of string $r = H_1(m_b || s)$. Therefore, ciphertexts $\tilde{\text{ct}}_{4\ell+1}, \dots, \tilde{\text{ct}}_{5\ell}$ are encryptions of each bit of r_4 . Since $\text{P}_{\text{fhe.sk}}(\tilde{\text{ct}}_{4\ell+1}, \dots, \tilde{\text{ct}}_{5\ell}) = r_4$, thus, by correctness of obfuscation, we know that $\text{Eval}(\tilde{P}, (\tilde{\text{ct}}_{4\ell+1}, \dots, \tilde{\text{ct}}_{5\ell})) = s$. In other words, $t = s$. Finally, by correctness of SKE scheme, we can conclude that $\text{SKE.Dec}(H_2(t), \text{ct}_{\text{SKE}}) = m_b$. Therefore, \mathcal{A} always guesses correctly.

We would like to point out that \mathcal{A} learns the challenge message completely, therefore we could also prove that scheme $\text{FO}_{\text{PKE}, \text{SKE}}^{\mathcal{H}, \tilde{\mathcal{H}}}$ is not even one-way secure.

8 Indistinguishability Obfuscation for Rejecting Programs

We now consider a new notion of obfuscation that we call indistinguishability obfuscation for rejecting programs and show how to construct it from lockable obfuscation and witness encryption for circuit satisfiability.

Obfuscators that meet this notion will be defined over boolean circuits. Like indistinguishability obfuscation our obfuscator will take in any (not necessarily rejecting) boolean circuit C in a class and output an obfuscated program that is functionally equivalent to C . However, the security guarantees given by such an obfuscator are limited to “rejecting” programs. Informally, they state that no PPT adversary can distinguish between circuits C_0 and C_1 as long as for all inputs x $C_0(x) = C_1(x) = 0$. In contrast, standard indistinguishability obfuscation security allows C_0, C_1 to have arbitrary (both 0 and 1) outputs as long as they are functionally equivalent.

Our construction is simple and follows along the same conceptual lines as our techniques for attribute hiding from Section 5. Recall, that in a witness encryption scheme for circuit satisfiability the encryption operation will take as input a circuit description C as well as a message msg . The message will be computationally hidden as long as there is no satisfying assignment to C . I.e. for all x we have $C(x) = 0$. However, there is nothing in the definition of witness encryption that guarantees the hiding of the description of the circuit C itself.

To obfuscate a circuit C one first chooses a random lock value α and creates a witness encryption of the message α under circuit C to get a (sub) ciphertext ct . Then we can create a lockable obfuscation of the following program. The program on input x will first check if $C(x) = 0$, and if so it outputs 0^{out} . Otherwise, it runs the witness encryption decryption algorithm with x as the witness that C is satisfiable. It is easy to verify correctness. On any input x where $C(x) = 0$ the lockable program will reject and we can output 0. On the other hand if $C(x) = 1$ the internal decryption will result in the lock value α and a message will be output. The security argument follows readily and is given below.

While the idea of indistinguishability for rejecting programs might at first seem limiting for creating applications, it is important to remember that the actual obfuscator can be correctly used on programs that are not necessarily rejecting. The restriction on rejecting programs is needed only in proof steps that invoke the security of the obfuscator and other outside primitives might be used to make other proof steps work. For example, Garg et al. [GGSW13] show how to use witness encryption and unique signatures to give an Identity-Based Encryption scheme. But one where the identity one encrypts to is revealed. If we replace the witness encryption scheme, with an indistinguishability obfuscator for rejecting programs, we get an anonymous IBE scheme [ABC⁺05].¹⁹

8.1 Defining Indistinguishability Obfuscation for Rejecting Programs

We will now define the notion of rejecting indistinguishability obfuscator (riO). The syntax and correctness properties are similar to that of standard indistinguishability obfuscation. The obfuscator is a compiler which takes as input a circuit, and outputs an obfuscated program which has identical functionality as the input program. For security, we require that if two circuits reject all inputs, then their obfuscations are computationally indistinguishable.

Syntax Let $\{\mathcal{C}_\lambda\}_\lambda$ be a circuit family, where each circuit in \mathcal{C}_λ has domain \mathcal{D}_λ and co-domain \mathcal{R}_λ . A rejecting indistinguishability obfuscator consists of a compiler riO.Obf and an evaluator riO.Eval with the following syntax.

- $\text{riO.Obf}(1^\lambda, C)$: The obfuscator takes as input a security parameter λ and a circuit $C \in \mathcal{C}_\lambda$. It outputs an obfuscated program \tilde{C} .
- $\text{riO.Eval}(\tilde{C}, x)$: The evaluation algorithm takes as input an obfuscated program \tilde{C} and an input $x \in \mathcal{D}_\lambda$. It output $y \in \mathcal{R}_\lambda \cup \{\perp\}$.

Correctness Here, we will be considering perfect correctness, which requires that the obfuscated program has identical functionality as the input program. Formally, an obfuscator $\text{riO} = (\text{riO.Obf}, \text{riO.Eval})$ is said to be correct if for all security parameters λ , circuits $C \in \mathcal{C}_\lambda$ and inputs $x \in \mathcal{D}_\lambda$,

$$\text{riO.Eval}(\text{riO.Obf}(1^\lambda, C), x) = C(x).$$

Security We will now define the security notion for rejecting indistinguishability obfuscator. This notion guarantees that if two programs reject all inputs, then their obfuscations are computationally indistinguishable.

¹⁹This resulting construction (perhaps not so surprisingly) is actually very close to what one get if one first constructed an IBE scheme from witness encryption using [GGSW13] and then applied our techniques from Section 5 to hide the identity. If one uses obfuscation for rejecting programs roles these two parts into one.

Definition 8.1. An obfuscation scheme $\text{riO} = (\text{riO.Obf}, \text{riO.Eval})$ for circuit class $\{\mathcal{C}_\lambda\}_\lambda$ is said to be a secure rejecting indistinguishability obfuscator if for all PPT adversaries $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$, there exists a negligible function $\text{negl}(\cdot)$ such that the following function (of λ) is bounded by $\text{negl}(\cdot)$:

$$\left| \Pr \left[\mathcal{A}_1(\tilde{C}_b, \text{st}) = b \wedge (\forall x \in \mathcal{D}_\lambda, C_0(x) = C_1(x) = 0) : \begin{array}{l} (C_0, C_1, \text{st}) \leftarrow \mathcal{A}_0(1^\lambda), C_b \in \mathcal{C}_\lambda \\ b \leftarrow \{0, 1\}, \tilde{C}_b \leftarrow \text{riO.Obf}(1^\lambda, C_b) \end{array} \right] - \frac{1}{2} \right|$$

8.2 Witness Encryption

In this section, we will define the notion of witness encryption, first introduced by [GGSW13]. A witness encryption scheme for an NP language \mathcal{L} consists of an encryption algorithm and a decryption algorithm. Using the encryption algorithm, one can encrypt any message for an NP statement x to compute a ciphertext ct . If there exists a witness w proving that $x \in \mathcal{L}$, then one can use w to decrypt the ciphertext ct . The security guarantee states that if $x \notin \mathcal{L}$, then encryption of message m_0 for statement x is indistinguishable from encryption of message m_1 for statement x . We will now give a formal description of the syntax, correctness and security requirements.

Syntax Let \mathcal{L} be an NP language with witness relation $\mathcal{R}(\cdot, \cdot)$ (that is, $x \in \mathcal{L}$ iff there exists a witness w such that $\mathcal{R}(x, w) = 1$). Let $n = n(|x|)$ be a polynomial denoting the length of witness for statement x . A witness encryption scheme WE for language \mathcal{L} and message space \mathcal{M} consists of the following algorithms.

- $\text{WE.Enc}(1^\lambda, x, m)$: The encryption algorithm takes as input security parameter λ , statement x and message $m \in \mathcal{M}$, and outputs a ciphertext ct .
- $\text{WE.Dec}(\text{ct}, w)$: The decryption algorithm takes as input a ciphertext ct , witness w and outputs $y \in \mathcal{M} \cup \{\perp\}$.

Correctness For simplicity, we will define perfect correctness.

Definition 8.2. A witness encryption scheme $(\text{WE.Enc}, \text{WE.Dec})$ for language \mathcal{L} and message space \mathcal{M} is said to be (perfectly) correct if, for all security parameters λ , statements $x \in \mathcal{L}$, witnesses w s.t. $\mathcal{R}(x, w) = 1$ and messages $m \in \mathcal{M}$,

$$\text{WE.Dec}(\text{WE.Enc}(1^\lambda, x, m), w) = m.$$

Security We will be using the adaptive soundness security definition from the work of Bellare and Hoang [BH15].

Definition 8.3. A witness encryption scheme $\text{WE} = (\text{WE.Enc}, \text{WE.Dec})$ for language \mathcal{L} and message space \mathcal{M} is said to be secure if for all security parameters λ , PPT adversaries $\mathcal{A} = (\mathcal{A}_0, \mathcal{A}_1)$,

$$\Pr \left[x \notin \mathcal{L} \wedge \mathcal{A}_1(\text{ct}, \text{st}) = b : \begin{array}{l} (x, m_0, m_1, \text{st}) \leftarrow \mathcal{A}_0(1^\lambda) \\ b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{WE.Enc}(1^\lambda, x, m_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

8.3 Construction of Rejecting Indistinguishability Obfuscator from Witness Encryption

We will show how to construct rejecting indistinguishability obfuscator from witness encryption and lockable obfuscation. Let $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ be a family of circuits, where each circuit in \mathcal{C}_λ has depth $d(\lambda)$, takes $n(\lambda)$ bit inputs, outputs a single bit. First, we need to define an NP relation for our witness encryption scheme. The NP language for our witness encryption scheme will be $\text{CIRCUIT-SAT} = \{C \in \mathcal{C} : \exists w \text{ s.t. } C(w) = 1\}$. The set of witnesses for any $C \in \text{CIRCUIT-SAT}$ is $\{w : C(w) = 1\}$. The tools required for constructing riO for \mathcal{C} are as follows:

Program $Q_{\text{ct},C}$

Constants: Ciphertext ct , circuit C .

Inputs: String $w \in \{0, 1\}^n$.

1. If $C(w) = 0$, output $0^{\ell_{\text{out}}}$.
2. Else output $\text{WE.Dec}(\text{ct}, w)$.

Figure 7: Program $Q_{\text{ct},C}$

- A witness encryption scheme $\text{WE} = (\text{WE.Enc}, \text{WE.Dec})$ with message space $\{0, 1\}^{\ell_{\text{out}}}$ for NP relation **CIRCUIT-SAT**. Let ℓ_{ct} denote the size of the ciphertext. Further, we will assume a canonical family of circuits $\{C_{\text{WE.Dec}}^\lambda\}_\lambda$ where $C_{\text{WE.Dec}}^\lambda(\text{ct}, w) = \text{WE.Dec}(\text{ct}, w)$. The circuit $C_{\text{WE.Dec}}^\lambda$ has depth $d_{\text{WE.Dec}}(\lambda)$, takes $\ell_{\text{ct}}(\lambda) + n(\lambda)$ bit inputs, and outputs $\ell_{\text{out}}(\lambda)$ bits.

- An lockable obfuscation scheme $(\text{Obf}, \text{Eval})$ with message space $\{0, 1\}$ for circuit class $\mathcal{C}_{n, \ell_{\text{out}}, d+d_{\text{WE.Dec}}}$.

We will now describe the obfuscation and evaluation algorithms.

- **riO.Obf**($1^\lambda, C$): The obfuscation algorithm first chooses a uniformly random string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$. Next, it computes a WE ciphertext $\text{ct} \leftarrow \text{WE.Enc}(1^\lambda, C, \alpha)$.²⁰ Let $Q_{\text{ct},C}(\cdot)$ be the circuit defined in Figure 7. The obfuscator then computes an lockable obfuscation $\tilde{C} \leftarrow \text{Obf}(1^\lambda, Q_{\text{ct},C}, 1, \alpha)$. The final obfuscation is \tilde{C} .
- **riO.Eval**(\tilde{C}, x): The evaluation algorithm first computes $y = \text{Eval}(\tilde{C}, x)$. If $y = 1$, the riO evaluator outputs 1, else it outputs 0.

Correctness We need to show that the obfuscated program has identical functionality as the original (input) program. Fix any security parameter λ , circuit $C \in \mathcal{C}_\lambda$ and input $x \in \{0, 1\}^n$. Let α be the random string chosen during obfuscation, $\text{ct} \leftarrow \text{WE.Enc}(1^\lambda, C, \alpha)$ and $\tilde{C} \leftarrow \text{Obf}(1^\lambda, 1, \alpha)$.

If $C(x) = 0$, then $Q_{\text{ct},C}(x) = 0^{\ell_{\text{out}}}$. As a result, $\text{Eval}(\tilde{C}, x) = \perp$ with all but negligible probability over the coins chosen during lockable obfuscation (assuming the semi-statistical correctness of lockable obfuscation scheme). Hence, $\text{riO.Eval}(\tilde{C}, x) = 0$ with all but negligible probability.

If $C(x) = 1$, then $Q_{\text{ct},C}(w) = \alpha$ (assuming perfect correctness of the witness encryption scheme), and hence $\text{riO.Eval}(\tilde{C}, x) = \text{Eval}(\tilde{C}, x) = 1$ (assuming correctness of lockable obfuscation scheme).

8.3.1 Security

We will prove security by defining a sequence of computationally indistinguishable hybrids, and finally showing that in the last hybrid, the adversary has zero advantage. At a high level, our proof works as follows. Recall that the adversary must output two programs that reject all inputs. As a result, if we encrypt a string α to either of these programs using the witness encryption scheme, then α is computationally hidden. Hence, we can replace α with the all-zeroes string. Having done this, we can now use the security of our lockable obfuscation scheme to simulate the obfuscated program.

We will now formally define the hybrid experiments.

Hybrid H_0 : This corresponds to the real experiment.

1. The challenger receives as input two circuits C_0, C_1 such that both circuits output 0 on all inputs.
2. The challenger chooses a uniformly random string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ and $b \leftarrow \{0, 1\}$.
3. It then computes $\text{ct} \leftarrow \text{WE.Enc}(1^\lambda, C_b, \alpha)$.
4. Finally, it computes $\tilde{C} \leftarrow \text{Obf}(1^\lambda, Q_{\text{ct},C_b}, 1, \alpha)$ and sends \tilde{C} to the adversary. The adversary outputs a bit b' and wins if $b = b'$.

²⁰By circuit C , we represent the statement that $\exists w$ such that $C(w) = 1$.

Hybrid H_1 : This hybrid is similar to the previous one, except that the challenger encrypts $0^{\ell_{\text{out}}}$ instead of α .

3. It then computes $\text{ct} \leftarrow \text{WE.Enc}(1^\lambda, C_b, 0^{\ell_{\text{out}}})$.

Hybrid H_2 : In this hybrid experiment, the challenger uses the simulator for the lockable obfuscation scheme to output the final obfuscation. Let Sim be the simulator for the lockable obfuscation scheme. Let $s = |Q_{\text{ct}, C_0}| = |Q_{\text{ct}, C_1}|$.

4. Finally, it computes $\tilde{C} \leftarrow \text{Sim}(1^\lambda, 1^s)$ and sends \tilde{C} to the adversary. The adversary outputs a bit b' and wins if $b = b'$.

Let $\text{Adv}_{\mathcal{A}}^i$ denote the advantage of adversary \mathcal{A} in hybrid H_i .

Claim 8.1. Assuming the witness encryption scheme WE is secure, for any PPT adversary \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^0 - \text{Adv}_{\mathcal{A}}^1| \leq \text{negl}(\lambda)$.

Proof. Suppose, on the contrary, there exists a PPT adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^0 - \text{Adv}_{\mathcal{A}}^1| = \epsilon$. We will use \mathcal{A} to construct a reduction algorithm \mathcal{B} that breaks the security of WE.

The adversary first sends two circuits C_0, C_1 . The reduction algorithm chooses a string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$, $b \leftarrow \{0, 1\}$ and sends $C_b, \alpha, 0^{\ell_{\text{out}}}$ to the witness encryption challenger. It receives a ciphertext ct . It then computes $\tilde{C} \leftarrow \text{Obf}(1^\lambda, Q_{\text{ct}, C}, 1, \alpha)$ and sends \tilde{C} to the adversary. The adversary sends its guess b' . If $b = b'$, the reduction algorithm guesses that α was encrypted, else it guesses that $0^{\ell_{\text{out}}}$ was encrypted. ■

Claim 8.2. Assuming the security of lockable obfuscation scheme, for any PPT adversary \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2| \leq \text{negl}(\lambda)$.

Proof. Suppose, on the contrary, there exists a PPT adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2| = \epsilon$. We will use \mathcal{A} to construct a reduction algorithm \mathcal{B} that breaks the security of the lockable obfuscation scheme.

The adversary sends two circuits C_0, C_1 . The reduction algorithm first chooses $b \leftarrow \{0, 1\}$ and computes an encryption of $0^{\ell_{\text{out}}}$ for C_b ; that is, it computes $\text{ct} \leftarrow \text{WE.Enc}(1^\lambda, C_b, 0^{\ell_{\text{out}}})$. It then sends circuit Q_{ct, C_b} and message 1 to the lockable obfuscator challenger, and receives an obfuscated program \tilde{C} which it forwards to \mathcal{A} . If \mathcal{A} guesses correctly, then \mathcal{B} guesses that \tilde{C} is an honestly computed obfuscation, else it guesses that \tilde{C} is simulated. ■

Claim 8.3. For any adversary \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^2| = 0$.

Proof. This follows directly from the definition of hybrid H_2 . Note that the challenger only requires the size of Q_{ct, C_b} , and this is independent of b since $|C_0| = |C_1|$. As a result, the final obfuscated program \tilde{C} contains no information about the bit b . ■

9 Upgrading Broadcast Encryption to Anonymous Broadcast Encryption

The notion of broadcast encryption was first formally discussed in [FN94]. In a broadcast encryption scheme, we have a set of N users. There is a common public key for all the N users, and each user has a private decryption key. Using the encryption algorithm, one can encrypt a message for any subset of users. Suppose message m is encrypted for set S . Then any person in set S must be able to recover m using the decryption key, and if person i is not in the set S , then person i should not learn the message m . For

efficiency/succinctness, we require that the size of the ciphertext be at most $|S| + \text{poly}(\lambda)$, where λ is the security parameter.

While hiding the message might suffice for certain applications, one might be interested in the case where even the set S is hidden. This is referred to as anonymous/recipient-private broadcast encryption, introduced by [BBW06]. In this work, we will consider a weaker notion which we call *one-sided* anonymous broadcast encryption. Here, if a message is encrypted for a set S , then for any user $i \notin S$, both the message and the set S will be hidden.

In this section, we will first define the notions of broadcast encryption and one-sided anonymous broadcast encryption, and then show how to transform any broadcast encryption scheme to a one-sided anonymous broadcast encryption scheme.

9.1 Preliminaries

We will first introduce the syntax, correctness and security properties of broadcast encryption. A broadcast encryption scheme BE with message space \mathcal{M} consists of the following algorithms.

- **Setup**($1^\lambda, N$): The setup algorithm takes as input the security parameter λ and a bound N on the total number of users. It outputs a public key pk and N decryption keys $\text{sk}_1, \dots, \text{sk}_N$.
- **Encrypt**($\text{pk}, m \in \mathcal{M}, S \subseteq [N]$): The encryption algorithm takes as input the public key pk , a message $m \in \mathcal{M}$ and a subset $S \subseteq [N]$. It outputs a ciphertext ct .
- **Decrypt**(sk, ct): The decryption algorithm takes as input a secret key sk and a ciphertext ct , and outputs $y \in \mathcal{M} \cup \{\perp\}$.

Correctness We require that for all security parameters λ and bound N on total number of users, for all sets $S \subseteq [N]$, messages $m \in \mathcal{M}$, $(\text{pk}, \{\text{sk}_1, \dots, \text{sk}_N\}) \leftarrow \text{Setup}(1^\lambda, 1^N)$ and $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, m, S)$, if $i \in S$, then $\text{Decrypt}(\text{sk}_i, \text{ct}) = m$.

Succinctness For succinctness, we require that for all λ, N , $(\text{pk}, \{\text{sk}_1, \dots, \text{sk}_N\}) \leftarrow \text{Setup}(1^\lambda, N)$, messages $m \in \mathcal{M}$ and $S \subseteq [N]$, if $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, m, S)$ then $|\text{ct}| \leq \text{poly}(\lambda) + |S|$.

9.1.1 Security

We will now present two different notions of security. The first one only hides the message encrypted, while the second notion requires that the set of users associated with a ciphertext is also hidden.

Definition 9.1. Let $\text{BE} = (\text{Setup}, \text{Encrypt}, \text{Decrypt})$ be a broadcast encryption scheme. Consider the following security game between a challenger and an adversary \mathcal{A} .

- **Setup Phase** The challenger chooses $(\text{pk}, (\text{sk}_1, \dots, \text{sk}_N)) \leftarrow \text{Setup}(1^\lambda, 1^N)$ and sends pk to \mathcal{A} .
- **Pre-Challenge Query Phase** In this phase, the adversary queries for secret keys. It sends an index $i \in [N]$, and receives the secret key sk_i .
- **Challenge Phase** The adversary sends two challenge messages m_0, m_1 and a set S . The challenger chooses $b \leftarrow \{0, 1\}$, computes $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, m_b, S)$ and sends ct to the challenger.
- **Post-Challenge Query Phase** This is similar to the pre-challenge query phase. In this phase, the adversary sends an index $i \in [N]$, and receives the secret key sk_i .
- **Guess** The adversary finally sends its guess b' . Let Q be the set of indices queried by \mathcal{A} . The adversary wins if $b = b'$ and $Q \cap S = \emptyset$.

The scheme is said to be secure if for all security parameters λ and all PPT adversaries \mathcal{A} , the advantage of \mathcal{A} in the above security game is at most negligible.

Definition 9.2 (One-sided anonymous broadcast encryption). Let $\text{BE} = (\text{Setup}, \text{Encrypt}, \text{Decrypt})$ be a broadcast encryption scheme. Consider the following security game between a challenger and an adversary \mathcal{A} .

- **Setup Phase** The challenger chooses $(\text{pk}, (\text{sk}_1, \dots, \text{sk}_N)) \leftarrow \text{Setup}(1^\lambda, 1^N)$ and sends pk to \mathcal{A} .
- **Pre-Challenge Query Phase** In this phase, the adversary queries for secret keys. It sends an index $i \in [N]$, and receives the secret key sk_i .
- **Challenge Phase** The adversary sends two challenge messages m_0, m_1 and two sets S_0, S_1 . The challenger chooses $b \leftarrow \{0, 1\}$, computes $\text{ct} \leftarrow \text{Encrypt}(\text{pk}, m_b, S_b)$ and sends ct to the challenger.
- **Post-Challenge Query Phase** This is similar to the pre-challenge query phase. In this phase, the adversary sends an index $i \in [N]$, and receives the secret key sk_i .
- **Guess** The adversary finally sends its guess b' . Let Q be the set of indices queried by \mathcal{A} . The adversary wins if $b = b'$ and $Q \cap (S_0 \cup S_1) = \emptyset$.

The scheme is said to be secure if for all security parameters λ and all PPT adversaries \mathcal{A} , the advantage of \mathcal{A} in the above security game is at most negligible.

9.2 Upgrading Broadcast Encryption to One-Sided Anonymous Broadcast Encryption

In this section, we will show how to upgrade any broadcast encryption scheme to achieve one-sided anonymity. Let $\text{BE} = (\text{BE.Setup}, \text{BE.Encrypt}, \text{BE.Decrypt})$ be a broadcast encryption scheme with message space $\{0, 1\}^*$. Let $\ell_{\text{ct}}(\lambda)$ denote the length of ciphertext corresponding to message $m \in \{0, 1\}^\lambda$. Our scheme with one-sided anonymity $\text{aBE} = (\text{aBE.Setup}, \text{aBE.Encrypt}, \text{aBE.Decrypt})$ uses the following cryptographic primitives:

- PRF F with keyspace \mathcal{K} , input space $\{0, 1\}^\lambda$ and output space $\{0, 1\}$. Let d_F denote depth of circuit evaluating $F(\cdot, \cdot)$.
- Leveled homomorphic encryption scheme $\text{LHE} = (\text{LHE.Setup}, \text{LHE.Enc}, \text{LHE.Dec}, \text{LHE.Eval})$ with decryption circuit having depth bounded by $d_{\text{LHE}}(\cdot)$.
- Lockable obfuscation scheme $(\text{Obf}, \text{Eval})$

The scheme can be described as follows.

- $\text{aBE.Setup}(1^\lambda, 1^N)$: The setup algorithm computes $(\text{pk}, \{\text{sk}_1, \dots, \text{sk}_N\}) \leftarrow \text{BE.Setup}(1^\lambda, 1^N)$ and outputs pk as the public key and $\{\text{sk}_1, \dots, \text{sk}_N\}$ as the secret keys.
- $\text{aBE.Encrypt}(\text{pk}, m, S)$: The encryption algorithm first chooses a random string $\alpha \leftarrow \{0, 1\}^\lambda$ and computes $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \alpha, S)$. Next, it chooses a PRF key $K \leftarrow \mathcal{K}$ and LHE keys $(\text{lhe.pk}, \text{lhe.sk}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d_F})$. It computes $\text{mask} = (F(K, 1), F(K, 2), \dots, F(K, \ell_{\text{ct}}))$ and sets $\text{ct}_1 = \text{ct}' \oplus \text{mask}$. Let $\text{ct}_2 \leftarrow \text{LHE.Enc}(\text{lhe.pk}, K)$. Finally, let $C_{\text{LHE.Dec}}$ be the LHE decryption circuit with lhe.sk hardwired. The encryption algorithm computes $P \leftarrow \text{Obf}(C_{\text{LHE.Dec}}, \alpha, m)$. The final ciphertext is $\text{ct} = (\text{ct}_1, \text{ct}_2, P, \text{lhe.pk})$.
- $\text{aBE.Decrypt}(\text{sk}, \text{ct} = (\text{ct}_1, \text{ct}_2, P, \text{lhe.pk}))$: The decryption algorithm first computes an LHE evaluation using the encryption of PRF key. Consider a circuit C_F that takes as input a PRF key K and outputs $(F(K, 1), F(K, 2), \dots, F(K, \ell_{\text{ct}}))$. The decryption algorithm performs LHE evaluation on ct_2 with C_F . It computes $\tilde{\text{ct}}_2 = \text{LHE.Eval}(\text{ct}_2, C_F)$. Next, let $C_{\text{ct}_1 \oplus}$ denote the function that has ct_1 hardwired, takes input y and outputs $\text{ct}_1 \oplus y$. The decryption algorithm computes $\tilde{\text{ct}}_1 = \text{LHE.Eval}(\tilde{\text{ct}}_2, C_{\text{ct}_1 \oplus})$.
Next, consider the circuit $C_{\text{BE.Decrypt}} = \text{BE.Decrypt}(\text{sk}, \cdot)$ which has a secret key sk hardwired. Let $\tilde{\text{ct}} = \text{LHE.Eval}(\tilde{\text{ct}}_1, C_{\text{BE.Decrypt}})$. Finally, the decryption algorithm outputs $y = \text{Eval}(P, \tilde{\text{ct}})$.

Correctness Fix any security parameter λ , $N \in \mathbb{N}$, message $m \in \{0,1\}^\lambda$ and set $S \subseteq [N]$. Let $(\text{pk}, \{\text{sk}_1, \dots, \text{sk}_N\}) \leftarrow \text{aBE.Setup}(1^\lambda, N)$ be the public/secret keys. The encryption algorithm chooses $\alpha \leftarrow \{0,1\}^\lambda$, $K \leftarrow \mathcal{K}$, $(\text{lhe.pk}, \text{lhe.sk}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d_F})$ and sets $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \alpha, S)$, $\text{ct}_2 = \text{LHE.Enc}(\text{lhe.pk}, K)$. Finally, it computes $P \leftarrow \text{Obf}(C_{\text{LHE.Dec}}, \alpha, m)$.

The decryption algorithm first computes $\tilde{\text{ct}}_2 = \text{LHE.Eval}(\text{ct}_2, C_F)$. Note that $\tilde{\text{ct}}_2$ is an LHE encryption of $\text{mask} = (F(K, 1), \dots, F(K, \ell_{\text{ct}}))$. Next, the decryption algorithm computes $\tilde{\text{ct}}_1 = \text{LHE.Eval}(\tilde{\text{ct}}_2, C_{\text{ct}_1 \oplus})$. By construction, it follows that $\text{BE.Decrypt}(\text{sk}, \text{LHE.Dec}(\text{lhe.sk}, \tilde{\text{ct}}_1)) = \alpha$. Therefore, $\tilde{\text{ct}} = \text{LHE.Eval}(\tilde{\text{ct}}_1, C_{\text{BE.Decrypt}})$ is an LHE encryption of α . As a result, the obfuscated program P , when evaluated with input $\tilde{\text{ct}}$, outputs m .

Efficiency The ciphertext in our anonymous broadcast encryption scheme consists of three components. The first component ct_1 has size $\ell_{\text{ct}}(\lambda) = \text{poly}_1(\lambda) + N$. This follows from the efficiency of the underlying broadcast encryption scheme. Next, the second component is an LHE encryption of PRF key K , and therefore its size is bounded by $\text{poly}_2(\lambda)$, independent of N . Finally, the last component is an obfuscation of LHE decryption circuit. This also depends only on the security parameter, and is independent of N . As a result, the ciphertext output by our anonymous broadcast encryption scheme has size $\text{poly}(\lambda) + N$.

9.2.1 Security

In this section, we will show that the scheme described above satisfies the one-sided anonymity definition (see Definition 9.2). We will prove security via a sequence of hybrid experiments H_0, H_1 , where H_0 corresponds to the original security game.

Hybrid H_0

- **Setup Phase** The challenger chooses $(\text{pk}, (\text{sk}_1, \dots, \text{sk}_N)) \leftarrow \text{BE.Setup}(1^\lambda, 1^N)$ and sends pk to \mathcal{A} .
- **Pre-Challenge Query Phase** In this phase, the adversary queries for secret keys. It sends an index $i \in [N]$, and receives the secret key sk_i .
- **Challenge Phase** The adversary sends two challenge messages m_0, m_1 and two sets S_0, S_1 . The challenger chooses $b \leftarrow \{0, 1\}$.
 1. It chooses $\alpha \leftarrow \{0, 1\}^\lambda$ and computes $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \alpha, S_b)$.
 2. Next it chooses a PRF key K and sets $\text{mask} = (F(K, 1), F(K, 2), \dots, F(K, \ell_{\text{ct}}))$, $\text{ct}_1 = \text{ct}' \oplus \text{mask}$.
 3. It then chooses LHE keys $(\text{lhe.pk}, \text{lhe.sk}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d_F})$ and sets $\text{ct}_2 \leftarrow \text{LHE.Enc}(\text{lhe.pk}, K)$.
 4. Finally, let $C_{\text{LHE.Dec}}$ be the LHE decryption circuit with lhe.sk hardwired. The challenger computes $P \leftarrow \text{Obf}(C_{\text{LHE.Dec}}, \alpha, m_b)$ and outputs $(\text{ct}_1, \text{ct}_2, P)$.
- **Post-Challenge Query Phase** This is similar to the pre-challenge query phase. In this phase, the adversary sends an index $i \in [N]$, and receives the secret key sk_i .
- **Guess** The adversary finally sends its guess b' . Let Q be the set of indices queried by \mathcal{A} . The adversary wins if $b = b'$ and $Q \cap (S_0 \cup S_1) = \emptyset$.

Hybrid H_1 This hybrid is identical to the previous one, except that the challenger computes a broadcast encryption of $\mathbf{0}$ instead of the lock α .

1. It chooses $\alpha \leftarrow \{0, 1\}^\lambda$ and computes $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \mathbf{0}, S_b)$.

Hybrid H_2 In this hybrid experiment, the challenger simulates the obfuscated program. Instead of computing $P \leftarrow \text{obf}(C_{\text{LHE.Dec}}, \alpha, m_b)$, it computes $P \leftarrow \text{Sim}(1^{C_{\text{LHE.Dec}}}, 1^{|m_b|})$.

4. Finally, let $C_{\text{LHE.Dec}}$ be the LHE decryption circuit with lhe.sk hardwired. The challenger computes $P \leftarrow \text{Sim}(1^{C_{\text{LHE.Dec}}}, 1^{|m_b|})$ and outputs $(\text{ct}_1, \text{ct}_2, P)$.

Hybrid H_3 In this hybrid experiment, the challenger replaces the LHE encryption of PRF key K with encryption of $\mathbf{0}$.

3. It then chooses LHE keys $(\text{lhe.pk}, \text{lhe.sk}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d_F})$ and sets $\text{ct}_2 \leftarrow \text{LHE.Enc}(\text{lhe.pk}, \mathbf{0})$.

Hybrid H_4 This experiment is similar to the previous one, except that the challenger sets ct_1 to be a uniformly random string.

2. Next it chooses $\text{ct}_1 \leftarrow \{0, 1\}^{\ell_{\text{ct}}}$.

Analysis For any PPT adversary \mathcal{A} , let $\text{Adv}_{\mathcal{A}}^i$ denote the advantage of \mathcal{A} in the hybrid experiment H_i .

Claim 9.1. Assuming BE is a secure broadcast encryption scheme, for any PPT adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^0 - \text{Adv}_{\mathcal{A}}^1 \leq \text{negl}(\lambda)$.

Proof. The only difference between these two hybrids is with regard to the ciphertext component ct' . In hybrid H_0 , the challenger computes a broadcast encryption of α , while in H_1 , it computes an encryption of $\mathbf{0}$. Also, note that the adversary's queries lie outside the set $S_0 \cup S_1$. Suppose there exists an adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}}^0 - \text{Adv}_{\mathcal{A}}^1 = \epsilon$. We will use \mathcal{A} to construct a reduction algorithm \mathcal{B} that breaks the security of BE.

The reduction algorithm \mathcal{B} first receives the public key from the challenger, which it then forwards to the adversary. Next, it receives secret key queries from the adversary, which it forwards to the challenger. For each key query i , the challenger sends sk_i to \mathcal{B} , and the reduction algorithm forwards it to \mathcal{A} . Next, the adversary sends two challenge messages m_0, m_1 together with sets S_0, S_1 . The reduction algorithm first chooses a bit $b \leftarrow \{0, 1\}$. It then sends $m_b, \mathbf{0}$ as the challenge messages, and S_b as the challenge set. The challenger sends back the challenge ciphertext ct' . The reduction algorithm then chooses a PRF key K , LHE keys $(\text{lhe.pk}, \text{lhe.sk}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d_F})$, sets $\text{ct}_1 = \text{ct}' \oplus (F(K, 1), \dots, F(K, \ell_{\text{ct}}))$, $\text{ct}_2 \leftarrow \text{LHE.Enc}(\text{lhe.pk}, K)$ and $P \leftarrow \text{Obf}(C_{\text{LHE.Dec}}, \alpha, m_b)$. It sends $(\text{ct}_1, \text{ct}_2, P)$ to \mathcal{A} . Next, the adversary sends post-challenge queries which are handled similar to the pre-challenge ones. Finally, if the adversary's guess is correct, the reduction algorithm guesses that ct' is encryption of $\mathbf{0}$, else it guesses that ct' is encryption of m_b .

Depending on whether the challenger sends an encryption of m_b or $\mathbf{0}$, the reduction algorithm simulates either hybrid experiment H_0 or H_1 . ■

Claim 9.2. Assuming $(\text{Obf}, \text{Eval})$ is a secure lockable obfuscation scheme, for any PPT adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2 \leq \text{negl}(\lambda)$.

Proof. Suppose, on the contrary, there exists a PPT adversary \mathcal{A} such that $\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2 = \epsilon$ for some non-negligible function ϵ . We will use \mathcal{A} to build a reduction algorithm \mathcal{B} that breaks the security of lockable obfuscation scheme.

The reduction algorithm first chooses the public/secret keys of broadcast encryption scheme and sends the public key to \mathcal{A} . Next, it handles the queries of the adversary using the secret keys. In the challenge phase, the adversary sends m_0, m_1 and subsets S_0, S_1 . The reduction algorithm chooses $b \leftarrow \{0, 1\}$ and computes $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \mathbf{0}, S_b)$. It chooses PRF key K , LHE keys $(\text{lhe.pk}, \text{lhe.sk})$ and sets $\text{ct}_1 = \text{ct}' \oplus (F(K, 1), \dots, F(K, \ell_{\text{ct}}))$ and $\text{ct}_2 \leftarrow \text{LHE.Enc}(\text{lhe.pk}, K)$. Next, it sends $C_{\text{LHE.Dec}}$ and message m_b to the challenger and receives an obfuscated program P . It sends $(\text{ct}_1, \text{ct}_2, P)$ to \mathcal{A} . The adversary finally sends its guess. If the guess is correct, the reduction algorithm guesses that P is a simulated obfuscation, else it guesses that P is an honestly computed obfuscation.

Note that the lock used by the obfuscator is not used anywhere else in hybrids H_1 and H_2 . As a result, the simulator can correctly simulate either H_1 or H_2 depending on whether P is simulated or not. ■

Claim 9.3. Assuming LHE is a secure leveled homomorphic encryption scheme, for any PPT adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3 \leq \text{negl}(\lambda)$.

Proof. The only difference between these two hybrids is that in one case, the challenge ciphertext has an LHE encryption of the PRF key K , while in the other case, it is replaced by an encryption of $\mathbf{0}$. Suppose there exists an adversary that can distinguish between the two hybrids with advantage ϵ . We will use this adversary to break the IND CPA security of LHE.

The reduction algorithm first chooses the public/secret keys of broadcast encryption scheme and sends the public key to \mathcal{A} . Next, it handles the queries of the adversary using the secret keys. In the challenge phase, the adversary sends m_0, m_1 and subsets S_0, S_1 . The reduction algorithm chooses $b \leftarrow \{0, 1\}$ and computes $\text{ct}' \leftarrow \text{BE.Encrypt}(\text{pk}, \mathbf{0}, S_b)$. It chooses PRF key K , LHE keys $(\text{lhe.pk}, \text{lhe.sk})$ and sets $\text{ct}_1 = \text{ct}' \oplus (F(K, 1), \dots, F(K, \ell_{\text{ct}}))$.

Next, the reduction algorithm sends $K, \mathbf{0}$ to the challenger. It receives an LHE public key lhe.pk and ciphertext ct_2 from the challenger. Next, it computes $P \leftarrow \text{Sim}(1^{|\mathcal{C}_{\text{LHE,Dec}}|}, 1^{m_b})$ and sends $(\text{ct}_1, \text{ct}_2, P, \text{lhe.pk})$ to \mathcal{A} . The adversary finally sends its guess. If the guess is correct, the reduction algorithm guesses that ct_2 is an encryption of $\mathbf{0}$, else it guesses that it is an encryption of K .

Note that since the obfuscated program P is simulated, the reduction algorithm does not require the LHE secret key. ■

Claim 9.4. Assuming F is a secure pseudorandom function, for any adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^3 - \text{Adv}_{\mathcal{A}}^4 \leq \text{negl}(\lambda)$.

Proof. The only difference between these two hybrids is that in one case, $\text{mask} = (F(K, 1), \dots, F(K, \ell_{\text{ct}}))$, while in the other case, mask is uniformly random. Note that the PRF key is not required anywhere else in the security game (in particular, recall that ct_2 is an encryption of $\mathbf{0}$). As a result, it follows directly from the security of PRF F that no PPT adversary can distinguish between these two hybrids with advantage greater than $\text{negl}(\lambda)$. ■

Claim 9.5. For an adversary \mathcal{A} , $\text{Adv}_{\mathcal{A}}^4 = 0$.

Proof. In hybrid H_4 , the ciphertext component ct_1 is a uniformly random string, ct_2 is an LHE encryption of $\mathbf{0}$ and P is a simulated program which is independent of the bit b chosen by challenger. As a result, any adversary \mathcal{A} has zero advantage in this experiment. ■

Acknowledgements. We thank Mihir Bellare and Joseph Jaeger for their detailed and informative feedback on our definitions and proofs.

References

- [ABBC10] Tolga Acar, Mira Belenkiy, Mihir Bellare, and David Cash. Cryptographic agility and its relation to circular encryption. In *EUROCRYPT '10*, volume 6110 of LNCS, pages 403–422. Springer, 2010.
- [ABC⁺05] Michel Abdalla, Mihir Bellare, Dario Catalano, Eike Kiltz, Tadayoshi Kohno, Tanja Lange, John Malone-Lee, Gregory Neven, Pascal Paillier, and Haixia Shi. Searchable encryption revisited: Consistency properties, relation to anonymous ibe, and extensions. In *Annual International Cryptology Conference*, 2005.
- [ABHS09] Pedro Adão, Gergei Bana, Jonathan Herzog, and Andre Scedrov. Soundness and completeness of formal encryption: The cases of key cycles and partial information leakage. *Journal of Computer Security*, 17(5):737–797, 2009.
- [ABN08] Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. Cryptology ePrint Archive, Report 2008/440, 2008.

- [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.
- [ADGM16] Daniel Apon, Nico Döttling, Sanjam Garg, and Pratyay Mukherjee. Cryptanalysis of indistinguishability obfuscations of circuits over ggh13. *Cryptology ePrint Archive*, Report 2016/1003, 2016.
- [AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, pages 308–326, 2015.
- [AJS15] Prabhanjan Ananth, Abhishek Jain, and Amit Sahai. Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. *IACR Cryptology ePrint Archive*, 2015.
- [Ajt99] Miklós Ajtai. Generating hard instances of the short basis problem. In *Automata, Languages and Programming, 26th International Colloquium, ICALP'99, Prague, Czech Republic, July 11-15, 1999, Proceedings*, pages 1–9, 1999.
- [AP16] Navid Alamati and Chris Peikert. Three’s compromised too: Circular insecurity for any cycle length from (ring-)lwe. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, pages 659–680, 2016.
- [AS16] Prabhanjan Ananth and Amit Sahai. Projective arithmetic functional encryption and indistinguishability obfuscation from degree-5 multilinear maps. In *EUROCRYPT*, 2016.
- [Att14] Nuttapon Attrapadung. Dual system encryption via doubly selective security: Framework, fully secure functional encryption for regular languages, and more. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 557–577, 2014.
- [Bar86] D A Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in nc1. In *Proceedings of the eighteenth annual ACM symposium on Theory of computing*, STOC '86, 1986.
- [BBC⁺14] Boaz Barak, Nir Bitansky, Ran Canetti, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Obfuscation for evasive functions. In *Theory of Cryptography Conference*, 2014.
- [BBDP01] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 566–582. Springer, 2001.
- [BBO07] Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In *Annual International Cryptology Conference*, 2007.
- [BBW06] Adam Barth, Dan Boneh, and Brent Waters. Privacy in encrypted content distribution using private broadcast encryption. In *Financial Cryptography and Data Security, 10th International Conference, FC 2006, Anguilla, British West Indies, February 27-March 2, 2006, Revised Selected Papers*, pages 52–64, 2006.
- [BDFP86] Allan Borodin, Danny Dolev, Faith E. Fich, and Wolfgang J. Paul. Bounds for width two branching programs. *SIAM J. Comput.*, 15(2):549–560, 1986.

- [BFM15] Christina Brzuska, Pooya Farshim, and Arno Mittelbach. Random-oracle uninstantiability from indistinguishability obfuscation. In *Theory of Cryptography - 12th Theory of Cryptography Conference, TCC 2015, Warsaw, Poland, March 23-25, 2015, Proceedings, Part II*, 2015.
- [BGG⁺14] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, pages 533–556, 2014.
- [BGH⁺15] Zvika Brakerski, Craig Gentry, Shai Halevi, Tancrede Lepoint, Amit Sahai, and Mehdi Tibouchi. Cryptanalysis of the quadratic zero-testing of GGH. *IACR Cryptology ePrint Archive*, 2015.
- [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 *CRYPTO*, pages 1–18, 2001.
- [BGI⁺12] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. *J. ACM*, 59(2):6, 2012.
- [BGK⁺14] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, 2014.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, 2012.
- [BH15] Mihir Bellare and Viet Tung Hoang. Adaptive witness encryption and asymmetric password-based cryptography. In *Public-Key Cryptography - PKC 2015 - 18th IACR International Conference on Practice and Theory in Public-Key Cryptography, Gaithersburg, MD, USA, March 30 - April 1, 2015, Proceedings*, pages 308–331, 2015.
- [BHSV98] Mihir Bellare, Shai Halevi, Amit Sahai, and Salil Vadhan. Many-to-one trapdoor functions and their relation to public-key cryptosystems. In *Annual International Cryptology Conference*, 1998.
- [BHW15] Allison Bishop, Susan Hohenberger, and Brent Waters. New circular security counterexamples from decision linear and learning with errors. In *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part II*, pages 776–800, 2015.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 575–584, 2013.
- [BMSZ16] Saikrishna Badrinarayanan, Eric Miles, Amit Sahai, and Mark Zhandry. Post-zeroizing obfuscation: New mathematical tools, and the case of evasive circuits. In *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, 2016.
- [BPR12] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. 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 719–737, 2012.

- [BR14] Zvika Brakerski and Guy N Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In *Theory of Cryptography Conference*, 2014.
- [BSW07] John Bethencourt, Amit Sahai, and Brent Waters. Ciphertext-policy attribute-based encryption. In *IEEE Symposium on Security and Privacy*, pages 321–334, 2007.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. In *FOCS*, pages 97–106, 2011.
- [BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 171–190, 2015.
- [BVWW16] Zvika Brakerski, Vinod Vaikuntanathan, Hoeteck Wee, and Daniel Wichs. Obfuscating conjunctions under entropic ring lwe. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, 2016.
- [BW07] Dan Boneh and Brent Waters. Conjunctive, subset, and range queries on encrypted data. In *Proceedings of the 4th conference on Theory of cryptography, TCC'07*, pages 535–554, Berlin, Heidelberg, 2007. Springer-Verlag.
- [BWZ14] Dan Boneh, David J. Wu, and Joe Zimmerman. Immunizing multilinear maps against zeroizing attacks. Cryptology ePrint Archive, Report 2014/930, 2014.
- [CC09] Melissa Chase and Sherman S. M. Chow. Improving privacy and security in multi-authority attribute-based encryption. In *ACM Conference on Computer and Communications Security*, pages 121–130, 2009.
- [CC17] Ran Canetti and Yilei Chen. Constraint-hiding constrained prfs for nc1 from lwe. In *EUROCRYPT*, 2017.
- [CFL⁺16] Jung Hee Cheon, Pierre-Alain Fouque, Changmin Lee, Brice Minaud, and Hansol Ryu. Cryptanalysis of the new clt multilinear map over the integers. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2016.
- [CGH98] Ran Canetti, Oded Goldreich, and Shai Halevi. The random oracle methodology, revisited (preliminary version). In *STOC*, pages 209–218, 1998.
- [CGH12] David Cash, Matthew Green, and Susan Hohenberger. New definitions and separations for circular security. In *Public Key Cryptography - PKC*, pages 540–557, 2012.
- [CGH⁺15] Jean-Sébastien Coron, Craig Gentry, Shai Halevi, Tancreède Lepoint, Hemanta K. Maji, Eric Miles, Mariana Raykova, Amit Sahai, and Mehdi Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, 2015.
- [CH85] Stephen A. Cook and H. James Hoover. A depth-universal circuit. *SIAM Journal on Computing*, 14(4):833–839, 1985.
- [Cha07] Melissa Chase. Multi-authority attribute based encryption. In *TCC*, pages 515–534, 2007.
- [CHL⁺15] Jung Hee Cheon, Kyoohyung Han, Changmin Lee, Hansol Ryu, and Damien Stehlé. Cryptanalysis of the multilinear map over the integers. In *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part I*, pages 3–12, 2015.

- [CJL16] Jung Hee Cheon, Jinhyuck Jeong, and Changmin Lee. An algorithm for ntru problems and cryptanalysis of the ggh multilinear map without a low-level encoding of zero. *LMS Journal of Computation and Mathematics*, 2016.
- [CL01] Jan Camenisch and Anna Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. *IACR Cryptology ePrint Archive*, 2001:19, 2001.
- [CLLT16] Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. Cryptanalysis of GGH15 multilinear maps. In *Advances in Cryptology - CRYPTO 2016 - 36th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2016, Proceedings, Part II*, 2016.
- [CLLT17] Jean-Sébastien Coron, Moon Sung Lee, Tancrede Lepoint, and Mehdi Tibouchi. Zeroizing attacks on indistinguishability obfuscation over CLT13. In *Public-Key Cryptography - PKC 2017 - 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28-31, 2017, Proceedings, Part I*, 2017.
- [CLT13] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, pages 476–493, 2013.
- [CLT14] Jean-Sebastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Cryptanalysis of two candidate fixes of multilinear maps over the integers. *Cryptology ePrint Archive*, Report 2014/975, 2014.
- [CLT15] Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. New multilinear maps over the integers. In *Advances in Cryptology - CRYPTO 2015 - 35th Annual Cryptology Conference, Santa Barbara, CA, USA, August 16-20, 2015, Proceedings, Part I*, 2015.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
- [DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam D. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, pages 523–540, 2004.
- [DS05] Yevgeniy Dodis and Adam D. Smith. Correcting errors without leaking partial information. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 654–663, 2005.
- [FN94] Amos Fiat and Moni Naor. Broadcast encryption. In *Proceedings of the 13th Annual International Cryptology Conference on Advances in Cryptology, CRYPTO '93*, pages 480–491, 1994.
- [FO99a] Eiichiro Fujisaki and Tatsuaki Okamoto. How to enhance the security of public-key encryption at minimum cost. In *International Workshop on Public Key Cryptography*, pages 53–68. Springer, 1999.
- [FO99b] Eiichiro Fujisaki and Tatsuaki Okamoto. Secure integration of asymmetric and symmetric encryption schemes. In *CRYPTO '99*, volume 1666 of LNCS, pages 537–554. Springer, 1999.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, 2009.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, 2013.

- [GGH⁺13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [GGH15] Craig Gentry, Sergey Gorbunov, and Shai Halevi. Graph-induced multilinear maps from lattices. In *TCC*, 2015.
- [GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *STOC*, 2013.
- [GJPS08] Vipul Goyal, Abhishek Jain, Omkant Pandey, and Amit Sahai. Bounded ciphertext policy attribute based encryption. In *Automata, Languages and Programming, 35th International Colloquium, ICALP 2008, Reykjavik, Iceland, July 7-11, 2008, Proceedings, Part II - Track B: Logic, Semantics, and Theory of Programming & Track C: Security and Cryptography Foundations*, 2008.
- [GKW17a] Rishab Goyal, Venkata Koppula, and Brent Waters. Separating IND-CPA and circular security for unbounded length key cycles. In *Public-Key Cryptography - PKC 2017 - 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Amsterdam, The Netherlands, March 28-31, 2017, Proceedings, Part I*, 2017.
- [GKW17b] Rishab Goyal, Venkata Koppula, and Brent Waters. Separating semantic and circular security for symmetric-key bit encryption from the learning with errors assumption. In *EUROCRYPT*, 2017.
- [GLSW15] Craig Gentry, Allison Bishop Lewko, Amit Sahai, and Brent Waters. Indistinguishability obfuscation from the multilinear subgroup elimination assumption. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 151–170, 2015.
- [GLW14] Craig Gentry, Allison B. Lewko, and Brent Waters. Witness encryption from instance independent assumptions. In *Advances in Cryptology - CRYPTO 2014 - 34th Annual Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2014, Proceedings, Part I*, pages 426–443, 2014.
- [GMM⁺16] Sanjam Garg, Eric Miles, Pratyay Mukherjee, Amit Sahai, Akshayaram Srinivasan, and Mark Zhandry. Secure obfuscation in a weak multilinear map model. In *Theory of Cryptography Conference*, pages 241–268. Springer, 2016.
- [GPSW06] Vipul Goyal, Omkant Pandey, Amit Sahai, and Brent Waters. Attribute-based encryption for fine-grained access control of encrypted data. In *Proceedings of the 13th ACM conference on Computer and communications security, CCS '06*, 2006.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO*, 2013.
- [GVW13] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Attribute-based encryption for circuits. In *STOC*, 2013.
- [GVW15] Sergey Gorbunov, Vinod Vaikuntanathan, and Hoeteck Wee. Predicate encryption for circuits from lwe. In *Annual Cryptology Conference*, 2015.
- [Hal15] Shai Halevi. Graded encoding, variations on a scheme. Cryptology ePrint Archive, Report 2015/866, 2015.

- [HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [HJ16] Yupu Hu and Huiwen Jia. Cryptanalysis of ggh map. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2016.
- [KRW15] Venkata Koppula, Kim Ramchen, and Brent Waters. Separations in circular security for arbitrary length key cycles. In *Theory of Cryptography Conference (TCC)*, 2015.
- [KSW08] Jonathan Katz, Amit Sahai, and Brent Waters. Predicate encryption supporting disjunctions, polynomial equations, and inner products. In *Proceedings of the theory and applications of cryptographic techniques 27th annual international conference on Advances in cryptology, EUROCRYPT’08*, 2008.
- [KW16] Venkata Koppula and Brent Waters. Circular security counterexamples for arbitrary length cycles from LWE. In *CRYPTO*, 2016.
- [Lau02] Peeter Laud. Encryption cycles and two views of cryptography. In *NORDSEC 2002 - Proceedings of the 7th Nordic Workshop on Secure IT Systems (Karlstad University Studies 2002:31)*, pages 85–100, 2002.
- [Lin16a] Huijia Lin. Indistinguishability obfuscation from constant-degree graded encoding schemes. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2016.
- [Lin16b] Huijia Lin. Indistinguishability obfuscation from ddh on 5-linear maps and locality-5 prgs. Cryptology ePrint Archive, Report 2016/1096, 2016.
- [LOS⁺10] Allison B. Lewko, Tatsuoaki Okamoto, Amit Sahai, Katsuyuki Takashima, and Brent Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In *EUROCRYPT*, pages 62–91, 2010.
- [LV16] Huijia Lin and Vinod Vaikuntanathan. Indistinguishability obfuscation from ddh-like assumptions on constant-degree graded encodings. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, 2016.
- [LW10] Allison B. Lewko and Brent Waters. New techniques for dual system encryption and fully secure HIBE with short ciphertexts. In *Theory of Cryptography, 7th Theory of Cryptography Conference, TCC 2010, Zurich, Switzerland, February 9-11, 2010. Proceedings*, pages 455–479, 2010.
- [LW11] Allison B. Lewko and Brent Waters. Decentralizing attribute-based encryption. In *EUROCRYPT*, pages 568–588, 2011.
- [MO14] Antonio Marcedone and Claudio Orlandi. Obfuscation \Rightarrow (IND-CPA security $\not\Rightarrow$ circular security). In *Security and Cryptography for Networks - 9th International Conference, SCN 2014, Amalfi, Italy, September 3-5, 2014. Proceedings*, pages 77–90, 2014.
- [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.
- [MP13] Daniele Micciancio and Chris Peikert. Hardness of SIS and LWE with small parameters. In *Advances in Cryptology - CRYPTO 2013 - 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*, pages 21–39, 2013.

- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, April 2007.
- [MSZ16] Eric Miles, Amit Sahai, and Mark Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over ggh13. In *Annual Cryptology Conference*, 2016.
- [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.
- [PST14] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. In *International Cryptology Conference*, 2014.
- [RAD78] Ron Rivest, Leonard Adleman, and Michael L. Dertouzos. On data banks and privacy homomorphisms. In *Foundations of Secure Computation*, pages 169–180, 1978.
- [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.
- [Rot13] Ron Rothblum. On the circular security of bit-encryption. In *Theory of Cryptography - 10th Theory of Cryptography Conference, TCC 2013, Tokyo, Japan, March 3-6, 2013. Proceedings*, pages 579–598, 2013.
- [SW05] Amit Sahai and Brent Waters. Fuzzy identity-based encryption. In *EUROCRYPT*, pages 457–473, 2005.
- [SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 475–484, 2014.
- [Wat09] Brent Waters. Dual system encryption: Realizing fully secure ibe and hibe under simple assumptions. In *CRYPTO*, pages 619–636, 2009.
- [Wat12] Brent Waters. Functional encryption for regular languages. In *CRYPTO*, 2012.
- [WFS03] Brent R Waters, Edward W Felten, and Amit Sahai. Receiver anonymity via incomparable public keys. In *Proceedings of the 10th ACM conference on Computer and communications security*, 2003.
- [WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under lwe. Cryptology ePrint Archive, Report 2017/276, 2017.

A Background: Attribute-Based Encryption and Predicate Encryption

A.1 Key-Policy Attribute Based Encryption

A key-policy attribute based encryption (KP-ABE) scheme \mathcal{ABE} , for set of attribute spaces $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$, consists of four polytime algorithms (Setup, Enc, KeyGen, Dec) with the following syntax:

- $\text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm takes as input the security parameter λ and a description of attribute space \mathcal{X}_λ , predicate class \mathcal{C}_λ and message space \mathcal{M}_λ , and outputs the public parameters pp and the master secret key msk .
- $\text{Enc}(\text{pp}, x, m) \rightarrow \text{ct}$. The encryption algorithm takes as input public parameters pp , an attribute $x \in \mathcal{X}_\lambda$ and a message $m \in \mathcal{M}_\lambda$. It outputs a ciphertext ct .
- $\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm takes as input master secret key msk and a predicate $C \in \mathcal{C}_\lambda$. It outputs a secret key sk_C .
- $\text{Dec}(\text{sk}_C, \text{ct}) \rightarrow m$ or \perp . The decryption algorithm takes as input a secret key sk_C and a ciphertext ct . It outputs either a message $m \in \mathcal{M}_\lambda$ or a special symbol \perp .

Correctness. A key-policy attribute based encryption scheme is said to be correct if for all $\lambda \in \mathbb{N}$, $(\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda)$, for all $x \in \mathcal{X}_\lambda$, $C \in \mathcal{C}_\lambda$, $m \in \mathcal{M}_\lambda$, $\text{sk}_C \leftarrow \text{KeyGen}(\text{msk}, C)$, $\text{ct} \leftarrow \text{Enc}(\text{pp}, x, m)$, the following holds

$$\begin{aligned} C(x) = 1 &\Rightarrow \Pr[\text{Dec}(\text{sk}_C, \text{ct}) = m] \geq 1 - \text{negl}_1(\lambda), \\ C(x) = 0 &\Rightarrow \Pr[\text{Dec}(\text{sk}_C, \text{ct}) = \perp] \geq 1 - \text{negl}_2(\lambda), \end{aligned}$$

where $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ are negligible functions, and the probabilities are taken over the random coins used during key generation and encryption procedures.

Security. The standard notion of security for a KP-ABE scheme is that of full or adaptive security. It is formally defined as follows.

Definition A.1. A key-policy attribute based encryption scheme $\mathcal{ABE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ is said to be fully secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, such that the following holds:

$$\left| \Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \\ ((m_0, m_1), x) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{pp}) \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m_b) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

where every predicate query C , made by adversary \mathcal{A} to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle, must satisfy the condition that $C(x) = 0$.

A.2 Predicate Encryption

A predicate encryption (PE) scheme \mathcal{PE} , for set of attribute spaces $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\mathcal{C} = \{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$, consists of four polytime algorithms (Setup, Enc, KeyGen, Dec). They have the same syntax and correctness condition as that of a KP-ABE scheme.

The standard notion of security for a PE scheme is that of full attribute hiding (or 2-sided security). A weaker notion of security is of 1-sided security which is captured by the following indistinguishability based definition.

Definition A.2 (1-Sided Security). A predicate encryption scheme $\mathcal{PE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ is said to be 1-sided secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$, such that the following holds:

$$\left| \Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \\ ((m_0, x_0), (m_1, x_1)) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{pp}) \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc}(\text{pp}, x_b, m_b) \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda)$$

where every predicate query C , made by adversary \mathcal{A} to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle, must satisfy the condition that $C(x_0) = C(x_1) = 0$.

The notion of 1-sided security for PE schemes could alternatively be captured by a simulation based definition. The following definition was used by Gorbunov, Vaikuntanathan and Wee [GVW15]. Later in Appendix F, we show that if a PE scheme satisfies Definition A.2, then it also satisfies Definition A.3.

Definition A.3 (SIM-1-Sided Security). A predicate encryption scheme $\mathcal{PE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ is said to be SIM-1-sided secure if there exists a simulator Sim such that for every stateful PPT adversary \mathcal{A} , the following holds:

$$\left\{ \begin{array}{l} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) : \\ (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \\ (m, x) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{pp}) \\ \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m) \end{array} \right\} \approx_c \left\{ \begin{array}{l} \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{ct}) : \\ (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \\ (m, x) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{pp}) \\ \text{ct} \leftarrow \text{Sim}(1^\lambda, \text{pp}, 1^{|x|}, 1^{|m|}) \end{array} \right\}$$

where every predicate query C , made by adversary \mathcal{A} to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle, must satisfy the condition that $C(x) = 0$.

B Background: Cycle Testers

B.1 Public Key k -Cycle Tester

In a recent work, Bishop, Hohenberger and Waters [BHW15] introduced a generic framework for creating circular security counterexamples. In this *cycle tester* framework, there are four algorithms - Setup, KeyGen, Encrypt and Test. The setup algorithm takes as input the security parameter λ and cycle length parameter k outputs the public parameters. The key generation algorithm uses the public parameters to output a public-secret key pair. The encryption algorithm takes a public key and message as input, and outputs a ciphertext. Finally, the testing algorithm takes as input k public keys and k ciphertexts, and outputs 1 if the k ciphertexts form a key cycle, else it outputs 0. *Note that in this framework, there is no decryption algorithm.* The security requirement is identical to the IND-CPA security game. The following description is taken from [BHW15].

Definition B.1 (k -Cycle Tester). A *cycle tester* $\Gamma = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Test})$ for message space \mathcal{M} and secret key space $\{0, 1\}^s$ is a tuple of algorithms (where $s = s(\lambda)$) specified as follows:

- $\text{Setup}(1^\lambda, 1^k) \rightarrow \text{pp}$. The setup algorithm takes as input the security parameter λ and the length of cycle k . It outputs the public parameters pp .
- $\text{KeyGen}(\text{pp}) \rightarrow (\text{pk}, \text{sk})$. The key generation algorithm takes as input the public parameters pp and outputs a public key pk and secret key $\text{sk} \in \{0, 1\}^s$.

- $\text{Enc}(\mathbf{pk}, m \in \mathcal{M}) \rightarrow \mathbf{ct}$. The encryption algorithm takes as input a public key \mathbf{pk} and a message $m \in \mathcal{M}$ and outputs a ciphertext \mathbf{ct} .
- $\text{Test}(\mathbf{pk}, \mathbf{ct}) \rightarrow \{0, 1\}$. On input $\mathbf{pk} = (\mathbf{pk}_1, \dots, \mathbf{pk}_k)$ and $\mathbf{ct} = (\mathbf{ct}_1, \dots, \mathbf{ct}_k)$, the testing algorithm outputs a bit in $\{0, 1\}$.

The algorithms must satisfy the following properties.

1. (Testing Correctness) There exists a polynomial $p(\cdot)$ such that for all security parameters λ , for all cycle lengths k , the Test algorithm's advantage (given by the following expression) is at least $1/p(\lambda)$.

$$\left| \begin{array}{l} \Pr \left[1 \leftarrow \text{Test}(\mathbf{pk}, \mathbf{ct}) : \begin{array}{l} \mathbf{pp} \leftarrow \text{Setup}(1^\lambda, 1^k); (\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{KeyGen}(\mathbf{pp}) \\ \mathbf{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, \mathbf{sk}_{(i \bmod k)+1}) \end{array} \right] \\ - \Pr \left[1 \leftarrow \text{Test}(\mathbf{pk}, \mathbf{ct}) : \begin{array}{l} \mathbf{pp} \leftarrow \text{Setup}(1^\lambda, 1^k); (\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{KeyGen}(\mathbf{pp}) \\ \mathbf{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, 0^s) \end{array} \right] \end{array} \right|$$

2. (IND-CPA Security) Let $\Pi = (\text{Setup}, \text{KeyGen}, \text{Enc}, \cdot)$ be an encryption scheme with empty decryption algorithm. The scheme Π must satisfy the IND-CPA security definition (Definition 2.4).

B.2 Public Key Bit Encryption 1-Cycle Tester

The BHW framework does not directly work for generating circular security separations for bit-encryption. In a recent work, Goyal, Koppula and Waters [GKW17b] provide a *bit-encryption cycle tester framework* for symmetric-key bit encryption along the lines of BHW framework. Below we extend their framework in the public key setting.

Definition B.2. (Bit Encryption 1-Cycle Tester) A cycle tester $\Gamma = (\text{Setup}, \text{Enc}, \text{Test})$ for message space $\{0, 1\}^s$ and secret key space $\{0, 1\}^s$ is a tuple of algorithms (where $s = s(\lambda)$) specified as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\mathbf{pk}, \mathbf{sk})$. The setup algorithm takes as input the security parameter λ , and outputs a public-secret key pair $(\mathbf{pk}, \mathbf{sk})$.
- $\text{Enc}(\mathbf{pk}, m \in \{0, 1\}) \rightarrow \mathbf{ct}$. The encryption algorithm takes as input a public key \mathbf{pk} and a message $m \in \{0, 1\}$, and outputs a ciphertext \mathbf{ct} .
- $\text{Test}(\mathbf{pk}, \mathbf{ct}) \rightarrow \{0, 1\}$. The testing algorithm takes as input a public key \mathbf{pk} and a sequence of s ciphertexts $\mathbf{ct} = (\mathbf{ct}_1, \dots, \mathbf{ct}_s)$, and outputs a bit in $\{0, 1\}$.

The algorithms must satisfy the following properties.

1. (Testing Correctness) There exists a polynomial $p(\cdot)$ such that for all security parameters λ , the Test algorithm's advantage (given by the following expression) is at least $1/p(\lambda)$.

$$\left| \begin{array}{l} \Pr \left[1 \leftarrow \text{Test}(\mathbf{pk}, \mathbf{ct}) : \begin{array}{l} (\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, \mathbf{sk}_i) \end{array} \right] \\ - \Pr \left[1 \leftarrow \text{Test}(\mathbf{pk}, \mathbf{ct}) : \begin{array}{l} (\mathbf{pk}, \mathbf{sk}) \leftarrow \text{Setup}(1^\lambda) \\ \mathbf{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, 0) \end{array} \right] \end{array} \right|$$

2. (IND-CPA Security) Let $\Pi = (\text{Setup}, \text{Enc}, \cdot)$ be an encryption scheme with empty decryption algorithm. The scheme Π must satisfy the IND-CPA security definition (Definition 2.4).

B.3 Public Key Cycle Tester for Unbounded Length Cycles

The BHW framework does not give circular security separations for unbounded length key cycles since the **Setup** algorithm in their framework takes as input the length of cycle k . Below we extend the BHW framework for unbounded length key cycles.

Definition B.3 (Unbounded Length Cycle Tester). A *cycle tester* $\Gamma = (\text{Setup}, \text{Enc}, \text{Test})$ for message space \mathcal{M} and secret key space $\{0, 1\}^s$ is a tuple of algorithms (where $s = s(\lambda)$) specified as follows:

- $\text{Setup}(1^\lambda) \rightarrow (\mathbf{pk}, \mathbf{sk})$. The setup algorithm takes as input the security parameter λ and outputs a public key \mathbf{pk} and secret key $\mathbf{sk} \in \{0, 1\}^s$.
- $\text{Enc}(\mathbf{pk}, m \in \mathcal{M}) \rightarrow \text{ct}$. The encryption algorithm takes as input a public key \mathbf{pk} and a message $m \in \mathcal{M}$ and outputs a ciphertext ct .
- $\text{Test}(1^k, \mathbf{pk}, \mathbf{ct}) \rightarrow \{0, 1\}$. On input the key cycle length k , public keys $\mathbf{pk} = (\mathbf{pk}_1, \dots, \mathbf{pk}_k)$ and secret keys $\mathbf{ct} = (\text{ct}_1, \dots, \text{ct}_k)$, the testing algorithm outputs a bit in $\{0, 1\}$.

The algorithms must satisfy the following properties.

1. (Testing Correctness) There exists a polynomial $p(\cdot)$ such that for all security parameters λ , all cycle lengths k , the **Test** algorithm's advantage (given by the following expression) is at least $1/p(\lambda)$.

$$\left| \Pr \left[1 \leftarrow \text{Test}(1^k, \mathbf{pk}, \mathbf{ct}) : \begin{array}{l} (\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Setup}(1^\lambda) \\ \text{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, \mathbf{sk}_{(i \bmod k)+1}) \end{array} \right] \right. \\ \left. - \Pr \left[1 \leftarrow \text{Test}(1^k, \mathbf{pk}, \mathbf{ct}) : \begin{array}{l} (\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{Setup}(1^\lambda) \\ \text{ct}_i \leftarrow \text{Enc}(\mathbf{pk}_i, 0^s) \end{array} \right] \right|$$

2. (IND-CPA Security) Let $\Pi = (\text{Setup}, \text{Enc}, \cdot)$ be an encryption scheme with empty decryption algorithm. The scheme Π must satisfy the IND-CPA security definition (Definition 2.4).

C Lockable Obfuscators: Extending the Message Space

In this section, we show that a lockable obfuscation scheme with message space $\{0, 1\}$ can be extended to a lockable obfuscation scheme with message space $\{0, 1\}^\ell$ for any polynomial ℓ .

Let $(\text{Obf}, \text{Eval})$ be a lockable obfuscation scheme with message space $\{0, 1\}$. We will also require the following tools for this transformation:

- A pairwise independent hash function family \mathcal{H} with domain $\{0, 1\}^{\ell_{\text{out}}}$ and co-domain $\{0, 1\}^{\ell_{\mathcal{H}}}$
- A pseudorandom generator $\text{PRG} : \{0, 1\}^{\ell_{\mathcal{H}}} \rightarrow \{0, 1\}^{\ell \cdot \lambda}$

Consider the following lockable obfuscation scheme $(\text{Obf}', \text{Eval}')$ for circuit class $\mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$ and message space $\{0, 1\}^\ell$.

- $\text{Obf}'(1^\lambda, P, \text{msg}, \alpha)$: Let msg_i denote the i^{th} bit of msg for $i \leq \ell$, and $s_{[i, j]}$ be substring of s consisting starting with the i^{th} bit of s and ending with the j^{th} bit.

The obfuscation algorithm first chooses a hash function h . Let Q^i be a circuit that takes as input $x \in \{0, 1\}^{\ell_{\text{in}}}$ and outputs $\text{PRG}(h(P(x)))_{[(i-1) \cdot \lambda + 1, i \cdot \lambda]}$.

The obfuscation algorithm computes $\tilde{P}_i \leftarrow \text{Obf}(1^\lambda, Q^i, \text{msg}_i, \text{PRG}(h(\alpha))_{[(i-1) \cdot \lambda + 1, i \cdot \lambda]})$ for each $i \leq \ell$.

The final obfuscated program is $(\tilde{P}_1, \dots, \tilde{P}_\ell)$.

- $\text{Eval}'((\tilde{P}_1, \dots, \tilde{P}_\ell), x)$: The evaluation algorithm computes $y_i = \text{Eval}(\tilde{P}_i, x)$ for each $i \leq \ell$. If any $y_i = \perp$, it outputs \perp . Else, it outputs (y_1, \dots, y_ℓ) .

Correctness. We will show that the above scheme satisfies the semi-statistical correctness definition if the underlying scheme also satisfies the semi-statistical correctness notion. Fix any security parameter λ , program P , string α , input x and message msg . If $P(x) = \alpha$, then it follows directly that $\text{Eval}'(\text{Obf}'(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}$.

If $P(x) \neq \alpha$, then we need to show that with high probability (over the coins of obfuscation), the evaluation algorithm outputs \perp . First, observe that the evaluation algorithm Eval' outputs \perp whenever $y_i = \perp$ for any i (i.e., any of ℓ Eval invocations outputs \perp). Thus, we only need to argue that the probability all ℓ Eval operations output some string $y \neq \perp$ is negligible.

Let Bad denote the event that $\text{PRG}(h(P(x))) = \text{PRG}(h(\alpha))$, and let \mathcal{E} be the event that $\text{Eval}'((\tilde{P}_1, \dots, \tilde{P}_\ell), x) \neq \perp$. First, by a similar argument to that used in proof of Claim 4.3, we can argue that $\Pr[\text{Bad}] \leq 2^{-\ell n}$. Next, we know that if Bad does not occur, then there exists an i such that $\text{PRG}(h(P(x)))_{[(i-1)\cdot\lambda+1, i\cdot\lambda]} \neq \text{PRG}(h(\alpha))_{[(i-1)\cdot\lambda+1, i\cdot\lambda]}$ (i.e., the i^{th} segment in the $\text{PRG}(h(P(x)))$ does not match the corresponding lock value). Finally, using the underlying obfuscator's semi-statistical correctness, we get that $\Pr[\mathcal{E} \mid \overline{\text{Bad}}] \leq \text{negl}(\lambda)$. Hence with all but negligible probability, the evaluation algorithm outputs \perp as $\Pr[\mathcal{E}] \leq \Pr[\mathcal{E} \mid \overline{\text{Bad}}] + \Pr[\text{Bad}] \leq \text{negl}(\lambda)$. This concludes the proof of correctness.

Security. Security follows from a hybrid argument. The simulator will run the underlying scheme's simulator ℓ times. To prove that this is indistinguishable from the honestly computed obfuscation, we define $\ell + 2$ intermediate hybrid obfuscations. In the first hybrid, we switch $h(\alpha)$ to be a uniformly random string. This follows from the Leftover Hash Lemma. Next, we switch $\text{PRG}(h(\alpha))$ to a uniformly random $\ell \cdot \lambda$ bit string. This follows from the security of PRG . At this point, each of the locks for each of the ℓ calls to Obf are independent and uniformly random strings. We can now simulate each of them independently. In the $(i + 2)^{\text{th}}$ one, the challenger simulates the first i obfuscations, and honestly computes the remaining $\ell - i$ obfuscations. The $(i + 1)^{\text{th}}$ and $(i + 2)^{\text{th}}$ hybrids are computationally indistinguishable due to security of the underlying obfuscation scheme.

D Lockable Obfuscation with Statistical Correctness

In this section, we show how to construct lockable obfuscation with *statistical correctness*. The construction and proofs will be very similar to the corresponding parts from Section 4. In order to achieve this stronger notion of correctness, we will require low depth *injective* pseudorandom generators. Currently, we do not know how to instantiate such PRGs from the LWE assumption.

For any polynomials $\ell_{\text{in}}, \ell_{\text{out}}, d$ such that $\ell_{\text{out}} = \Omega(\lambda)$, we construct an lockable obfuscation scheme $(\text{Obf}, \text{Eval})$ with statistical correctness for the circuit class $\mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$. The message space for our construction will be $\{0, 1\}$. The tools required for this construction are as follows:

- A compact leveled homomorphic bit encryption scheme $(\text{LHE.Setup}, \text{LHE.Enc}, \text{LHE.Eval}, \text{LHE.Dec})$ with decryption circuit of depth $d_{\text{Dec}}(\lambda)$ and ciphertexts of length $\ell_{\text{ct}}(\lambda)$.
- An injective pseudorandom generator $\text{PRG} : \{0, 1\}^{\ell_{\text{out}}(\lambda)} \rightarrow \{0, 1\}^{\ell_{\text{PRG}}(\lambda)}$ of depth $d_{\text{PRG}}(\lambda)$.

Fix any $\epsilon < 1/2$. Let χ be a B -bounded discrete Gaussian distribution with parameter σ such that $B = \sqrt{m} \cdot \sigma$. Let $n, m, \ell, \sigma, q, \text{Bd}$ be parameters with the following constraints:

- $n = \text{poly}(\lambda)$ and $q \leq 2^{n^\epsilon}$ (for LWE security)
- $\ell_{\text{out}} = \omega(\log \lambda)$ (for PRG security)
- $m = \Omega(n \cdot \log q)$ (for SamplePre)
- $\sigma = \omega(\sqrt{n \cdot \log q \cdot \log m})$ (for Preimage Well Distributedness)
- $\ell_{\text{PRG}} = n \cdot m \cdot \log q + \omega(\log n)$ (for applying Leftover Hash Lemma)
- $\text{Bd} = \ell_{\text{PRG}} \cdot L \cdot (m^2 \cdot \sigma)^L < q^{1/4}$ (for correctness of scheme)
 (where $L = \ell_{\text{out}} \cdot \ell_{\text{ct}} \cdot 4^{d_{\text{dec}} + d_{\text{PRG}}}$)

As before, it is important that $L = \lambda^c$ for some constant c and $\text{Bd} = \ell_{\text{PRG}} \cdot L \cdot (m^2 \cdot \sigma)^L < q^{1/4}$. The constant c depends on the LHE scheme and PRG.

We will now describe the obfuscation and evaluation algorithms.

- $\text{Obf}(1^\lambda, P, \text{msg}, \alpha)$: The obfuscation algorithm takes as input a program $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$, message $\text{msg} \in \{0, 1\}$ and $\alpha \in \{0, 1\}^{\ell_{\text{out}}}$. The obfuscator proceeds as follows:

1. First, it chooses the LHE key pair as $(\text{lhe.sk}, \text{lhe.ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$.
2. Next, it encrypts the program P . It sets $\mathbf{ct} \leftarrow \text{LHE.Enc}(\text{lhe.sk}, P)$.
3. Let $\beta = \text{PRG}(\alpha)$.
4. Next, consider the following circuit Q which takes as input $\ell_{\text{out}} \cdot \ell_{\text{ct}}$ bits of input and outputs ℓ_{PRG} bits. Q takes as input ℓ_{out} LHE ciphertexts $\{\mathbf{ct}_i\}_{i \leq \ell_{\text{out}}}$, has LHE secret key lhe.sk hardwired and computes the following — (1) it decrypts each input ciphertext \mathbf{ct}_i (in parallel) to get string x of length ℓ_{out} bits, (2) it applies the PRG on x and outputs $\text{PRG}(x)$. Concretely, $Q(\mathbf{ct}_1, \dots, \mathbf{ct}_{\ell_{\text{out}}}) = \text{PRG}(\text{LHE.Dec}(\text{lhe.sk}, \mathbf{ct}_1) \parallel \dots \parallel \text{LHE.Dec}(\text{lhe.sk}, \mathbf{ct}_{\ell_{\text{out}}}))$.

For $i \leq \ell_{\text{PRG}}$, we use $\text{BP}^{(i)}$ to denote the fixed-input selector permutation branching program that outputs the i^{th} bit of output of circuit Q . Note that Q has depth at most $d_{\text{tot}} = d_{\text{Dec}} + d_{\text{PRG}}$. By Corollary 2.2, we know that each branching program $\text{BP}^{(i)}$ has length $L = \ell_{\text{out}} \cdot \ell_{\text{ct}} \cdot 4^{d_{\text{tot}}}$ and width 5.

5. Finally, the obfuscator creates matrix components which enable the evaluator to compute msg if it has an input strings (ciphertexts) $\mathbf{ct}_1, \dots, \mathbf{ct}_{\ell_{\text{out}}}$ such that $Q(\mathbf{ct}_1, \dots, \mathbf{ct}_{\ell_{\text{out}}}) = \beta$. Concretely, it runs the (randomized) routine **Comp-Gen** (defined in Figure 1). This routine takes as input the circuit Q in the form of ℓ_{PRG} branching programs $\{\text{BP}^{(i)}\}_i$, string β and message msg . Let $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ \mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)} \right\}_{i,j} \right) \leftarrow \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})$.
6. The final obfuscated program consists of the LHE evaluation key $\text{ek} = \text{lhe.ek}$, LHE ciphertexts \mathbf{ct} , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

- $\text{Eval}(\tilde{P}, x)$: The evaluation algorithm takes as input $\tilde{P} = \left(\text{ek}, \mathbf{ct}, \left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$ and input $x \in \{0, 1\}^{\ell_{\text{in}}}$. It performs the following steps.

1. The evaluator first constructs a universal circuit $U_x(\cdot)$ with x hardwired as input. This universal circuit takes a circuit C as input and outputs $U_x(C) = C(x)$. Using the universal circuit of Cook and Hoover [CH85], it follows that $U_x(\cdot)$ has depth $O(d)$.
2. Next, it performs homomorphic evaluation on \mathbf{ct} using circuit $U_x(\cdot)$. It computes $\tilde{\mathbf{ct}} = \text{LHE.Eval}(\text{ek}, U_x(\cdot), \mathbf{ct})$. Note that $\ell_{\text{ct}} \cdot \ell_{\text{out}}$ denotes the length of $\tilde{\mathbf{ct}}$ (as a bitstring), and let $\tilde{\mathbf{ct}}_i$ denote the i^{th} bit of $\tilde{\mathbf{ct}}$.
3. The evaluator then obviously evaluates the ℓ_{PRG} branching programs on input $\tilde{\mathbf{ct}}$ using the matrix components. It calls the component evaluation algorithm **Comp-Eval** (defined in Figure 2). Let $y = \text{Comp-Eval} \left(\tilde{\mathbf{ct}}, \left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right) \right)$. The evaluator outputs y .

D.1 Correctness

We will prove that the lockable obfuscation scheme described above satisfies the statistical correctness property (see Definition 3.2). First, let us recall the following lemma (a weaker version of this lemma was proven in Section 4.1).

Lemma D.1. For any set of branching programs $\{\text{BP}^{(i)}\}_{i \leq \ell_{\text{PRG}}}$, string $\beta \in \{0, 1\}^{\ell_{\text{PRG}}}$, message $\text{msg} \in \{0, 1\}$ and input z ,

1. if $\text{BP}^{(i)}(z) = \beta_i$ for all $i \leq \ell_{\text{PRG}}$, then $\text{Comp-Eval}(z, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) = \text{msg}$.
2. if $\text{BP}^{(i)}(z) \neq \beta_i$ for some $i \leq \ell_{\text{PRG}}$, then

$$\Pr[\text{Comp-Eval}(z, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) = \perp] \geq 1 - \text{negl}(\lambda)/2^{\ell_{\text{in}}}.$$

This lemma is identical to Lemma 4.1, except that if $\text{BP}^{(i)}(z) \neq \beta_i$ for some $i \leq \ell_{\text{PRG}}$, then the probability of error needs to be bounded by $\text{negl}(\lambda)/2^{\ell_{\text{in}}}$ instead of just $\text{negl}(\lambda)$. This is required for our union bound argument. The proof of this lemma is same as that of Lemma 4.1, and the final bound of $\text{negl}(\lambda)/2^{\ell_{\text{in}}}$ holds if we ensure that $(\text{Bd}/q)^{(m-n) \cdot n} < 2^{-2\ell_{\text{in}}}$.

Using the above lemma, we can now argue the correctness of our scheme. First, we need to show correctness for the case when $P(x) = \alpha$.

Claim D.1. For all security parameters λ , inputs $x \in \{0, 1\}^{\ell_{\text{in}}}$, programs $P \in \mathcal{C}_{\ell_{\text{in}}, \ell_{\text{out}}, d}$ and messages $\text{msg} \in \{0, 1\}$, if $P(x) = \alpha$, then

$$\text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}.$$

Proof. First, the obfuscator encrypts the program P using an LHE secret key lhe.sk , and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{lhe.sk}, P)$. The evaluator evaluates the LHE ciphertext on universal circuit $U_x(\cdot)$, which results in an evaluated ciphertext $\tilde{\text{ct}}$. Now, by the correctness of the LHE scheme, decryption of $\tilde{\text{ct}}$ using lhe.sk outputs α . Therefore, $\text{PRG}(\text{LHE.Dec}(\text{lhe.sk}, \tilde{\text{ct}})) = \beta$. Then, using Lemma 4.1, we can argue that Comp-Eval outputs msg , and thus Eval outputs msg . \blacksquare

Claim D.2. For all security parameters λ , programs $P \in \mathcal{C}_{n, m, d}(\lambda)$, $\alpha \in \{0, 1\}^{m(\lambda)}$ and $\text{msg} \in \mathcal{M}$,

$$\Pr[\exists x \text{ s.t. } P(x) \neq \alpha \text{ and } \text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) = \text{msg}] \leq \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Proof. Fix any security parameter λ , program P , α , x such that $P(x) \neq \alpha$ and message msg . Let us now consider the events

$$\text{Error} : \exists x \text{ s.t. } P(x) \neq \alpha \text{ and } \text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) \neq \perp$$

$$\text{Error}_x : \text{Eval}(\text{Obf}(1^\lambda, P, \text{msg}, \alpha), x) \neq \perp$$

Clearly, $\Pr[\text{Error}] \leq \sum_{x: P(x) \neq \alpha} \Pr[\text{Error}_x]$. Therefore, it suffices to show that $\Pr[\text{Error}_x]$ is at most $\text{negl}(\lambda)/2^{\ell_{\text{in}}}$.

First, note that for all $x \in \{0, 1\}^{\ell_{\text{in}}}$, if $P(x) \neq \alpha$, then $\text{PRG}(P(x)) \neq \text{PRG}(\alpha)$. This follows from the injectiveness of PRG . Fix any input x such that $P(x) \neq \alpha$, LHE keys $(\text{lhe.sk}, \text{lhe.ek})$ and ciphertext $\text{ct} \leftarrow \text{LHE.Enc}(\text{lhe.sk}, P)$. Let $\tilde{\text{ct}} = \text{LHE.Eval}(\text{ek}, U_x(\cdot), \text{ct})$ and Q be the program that, on input a LHE ciphertext, first decrypts it using sk , then computes $\text{PRG}(y)$ (where y is the decrypted string). Let $\text{BP}^{(i)}$ denote the branching program computing the i^{th} output bit of Q . Using the correctness of LHE decryption, it follows that $Q(\tilde{\text{ct}}) = \text{PRG}(P(x)) \neq \beta$,²¹ and therefore there exists some $i \leq \ell_{\text{PRG}}$ such that $\text{BP}^{(i)}(\tilde{\text{ct}}) \neq \beta_i$. Now, $\Pr[\text{Error}_x] = \Pr[\text{Comp-Eval}(\tilde{\text{ct}}, \text{Comp-Gen}(\{\text{BP}^{(i)}\}_i, \beta, \text{msg})) \neq \perp]$. Using Lemma D.1, we can conclude that for all $\text{msg} \in \{0, 1\}$, this probability is at most $\text{negl}(\lambda)/2^{\ell_{\text{in}}}$. Hence, $\Pr[\text{Error}] \leq \text{negl}(\lambda)$. \blacksquare

D.2 Security

We will now prove the scheme secure as per Definition 3.4. This proof is very similar to the one in Section 4.2. The only difference here is that we do not require a pairwise independent hash function.

²¹Recall $\beta = \text{PRG}(\alpha)$

D.2.1 Simulator Sim

The simulator first chooses the parameters $n, m, q, \sigma, \ell_{\text{PRG}}$ as in the original scheme. Next, it chooses LHE secret/evaluation keys. It computes $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ (note that the depth d of the circuit class is fixed for the scheme). It then computes an encryption of $\mathbf{0}^{|P|}$. Let $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, \mathbf{0}^{|P|})$. Finally, the simulator chooses ℓ_{PRG} matrices $\mathbf{B}_{0,1}^{(i)} \leftarrow \mathbb{Z}_q^{n \times m}$ and low norm matrices $\{(\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)})\}_{i,j}$ for $i \leq \ell_{\text{PRG}}$, $j \leq L$ where $\mathbf{C}_j^{(i,b)} \leftarrow \chi^{m \times m}$. The obfuscation consists of the LHE evaluation key ek , ciphertext ct , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

To prove security, we will define a sequence of hybrids. The proofs showing indistinguishability of hybrid experiments are identical to the ones in Section 4.2.

D.2.2 Sequence of Hybrid Games

Fix any program P , message $\text{msg} \in \{0, 1\}$.

Game 0: This corresponds to the original security game.

1. The challenger first chooses the LWE parameters n, m, q, σ, χ and ℓ_{PRG} . Recall L denotes the length of the branching programs.
2. It chooses $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, P)$.
3. Next, it chooses a uniformly random string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ and sets $\beta = \text{PRG}(\alpha)$.
4. Next, consider the following program Q . It takes as input an LHE ciphertext ct , has sk hardwired and does the following: it decrypts the input ciphertext ct to get string y and outputs $\text{PRG}(y)$. For $i \leq \ell_{\text{PRG}}(\lambda)$, let $\text{BP}^{(i)}$ denote the branching program that outputs the i^{th} bit of $\text{PRG}(y)$.
5. It chooses ℓ_{PRG} uniformly random matrices $\mathbf{B}_L^{(i)}$ of dimensions $5 \cdot n \times m$, such that the following constraint is satisfied (recall $\mathbf{B}_{L,1}^{(i)}$ represents the first n rows of $\mathbf{B}_L^{(i)}$, $\mathbf{B}_{L,2}^{(i)}$ represents the next n rows of $\mathbf{B}_L^{(i)}$, etc)

$$\sum_{i: \beta_i=0} \mathbf{B}_{L,\text{rej}^{(i)}}^{(i)} + \sum_{i: \beta_i=1} \mathbf{B}_{L,\text{acc}^{(i)}}^{(i)} = \begin{cases} \mathbf{0} & \text{if msg} = 0. \\ \sqrt{q} \cdot [\mathbf{I}_n \parallel \mathbf{0}^{n \times (m-n)}] & \text{if msg} = 1. \end{cases}$$

6. For $i = 1$ to ℓ_{PRG} and $j = 0$ to $L - 1$, it chooses $(\mathbf{B}_j^{(i)}, T_j^{(i)}) \leftarrow \text{TrapGen}(1^{5 \cdot n}, 1^m, q)$.
7. Next, it generates the components for each level. For each level $\text{level} \in [1, L]$, do the following:
 - (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$. If either $\mathbf{S}_{\text{level}}^{(0)}$ or $\mathbf{S}_{\text{level}}^{(1)}$ has determinant zero, then set it to be \mathbf{I}_n .
 - (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}$.
 - (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
 - (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$
8. The final obfuscated program consists of the LHE evaluation key ek , LHE encryption ct , together with the components $\left(\left\{ \mathbf{B}_{0,1}^{(i)} \right\}_i, \left\{ (\mathbf{C}_j^{(i,0)}, \mathbf{C}_j^{(i,1)}) \right\}_{i,j} \right)$.

Game 1: In this experiment, the challenger chooses the matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)}$ without checking if their determinant is non-zero.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
- (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

Game 2: In this experiment, the challenger replaces the pseudorandom string $\beta = \text{PRG}(\alpha)$ with a truly random string.

3. Next, it chooses a uniformly random string $\alpha \leftarrow \{0, 1\}^{\ell_{\text{out}}}$ and $\beta \leftarrow \{0, 1\}^{\ell_{\text{PRG}}}$.

Game 3: In this experiment, the challenger chooses the top level matrices uniformly at random.

5. It chooses ℓ_{PRG} uniformly random matrices $\mathbf{B}_L^{(i)}$ of dimensions $5 \cdot n \times m$, without any constraints.

We will now define $4L+1$ intermediate game **Game** $(4, j^*, 0)$, **Game** $(4, j^*, 1)$, **Game** $(4, j^*, 2)$, **Game** $(4, j^*, 3)$ for $j^* \in [0, L-1]$ and **Game** $(4, L, 0)$. The **Game** $(4, 0, 0)$ will correspond to **Game** 3.

Game $(4, j^*, 0)$: In this experiment, the challenger chooses the top j^* level matrices $\mathbf{B}_j^{(i)}$ uniformly at random (without a trapdoor). Also, the top level $\mathbf{C}_j^{(i,b)}$ matrices, for $j > L - j^*$, are chosen randomly from the noise distribution χ .

6. For $i = 1$ to ℓ_{PRG} and $j = 0$ to $L - 1$,

- if $j < L - j^*$ it chooses $(\mathbf{B}_j^{(i)}, T_j^{(i)}) \leftarrow \text{TrapGen}(1^{5 \cdot n}, 1^m, q)$.
- if $j \geq L - j^*$, it chooses $\mathbf{B}_j^{(i)}$ uniformly at random from $\mathbb{Z}_q^{5 \cdot n \times m}$.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.
- (c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.
- (d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game $(4, j^*, 1)$: In this experiment, the challenger chooses $\mathbf{M}_{\text{level}}^{(i,0)}$ uniformly at random for $\text{level} = L - j^*$ ($\mathbf{M}_{\text{level}}^{(i,1)}$ is same as before).

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

- (a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.
- (b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.

(c) If $\text{level} = L - j^*$, choose $\mathbf{M}_{\text{level}}^{(i,0)} \leftarrow \mathbb{Z}_q^{5 \cdot n \times m}$ and $\mathbf{M}_{\text{level}}^{(i,1)}$ same as previous hybrid.

Else set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.

(d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game (4, j^* , 2): In this experiment, the challenger chooses $\mathbf{M}_{\text{level}}^{(i,b)}$ uniformly at random for $\text{level} = L - j^*$ and $b \in \{0, 1\}$.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

(a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.

(b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.

(c) If $\text{level} = L - j^*$, choose $\mathbf{M}_{\text{level}}^{(i,b)} \leftarrow \mathbb{Z}_q^{5 \cdot n \times m}$.

Else set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.

(d) Compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game (4, j^* , 3): In this experiment, the challenger chooses $\mathbf{C}_{L-j^*}^{(i,b)}$ from the noise distribution.

7. Next, it generates the components for each level. For each level $\text{level} \in [1, L - j^*]$, do the following:

(a) Choose matrices $\mathbf{S}_{\text{level}}^{(0)}, \mathbf{S}_{\text{level}}^{(1)} \leftarrow \chi^{n \times n}$ and $\mathbf{E}_{\text{level}}^{(i,0)}, \mathbf{E}_{\text{level}}^{(i,1)} \leftarrow \chi^{5n \times m}$ for $i \leq \ell_{\text{PRG}}$.

(b) For $b \in \{0, 1\}$, set matrix $\mathbf{D}_{\text{level}}^{(i,b)}$ as a permutation of the matrix blocks of $\mathbf{B}_{\text{level}}^{(i)}$ according to the permutation $\sigma_{\text{level},b}^{(i)}(\cdot)$.

(c) Set $\mathbf{M}_{\text{level}}^{(i,b)} = \left(\mathbf{I}_5 \otimes \mathbf{S}_{\text{level}}^{(b)} \right) \cdot \mathbf{D}_{\text{level}}^{(i,b)} + \mathbf{E}_{\text{level}}^{(i,b)}$ for $i \leq \ell_{\text{PRG}}$.

(d) If $\text{level} = L - j^*$, then choose $\mathbf{C}_{L-j^*}^{(i,b)} \leftarrow \chi^{m \times m}$ for $b \in \{0, 1\}$.

Else compute $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \text{SamplePre}(\mathbf{B}_{\text{level}-1}^{(i)}, T_{\text{level}-1}^{(i)}, \sigma, \mathbf{M}_{\text{level}}^{(i,b)})$

For all $\text{level} > L - j^*$, $i \leq \ell_{\text{PRG}}$ and $b \in \{0, 1\}$ choose $\mathbf{C}_{\text{level}}^{(i,b)} \leftarrow \chi^{m \times m}$.

Game 5: In this experiment, the challenger computes an encryption of $\mathbf{0}^{|P|}$ instead of an encryption of P . Note that the LHE secret key sk is not required at any other step in this experiment. This corresponds to the ideal world.

2. It chooses $(\text{sk}, \text{ek}) \leftarrow \text{LHE.Setup}(1^\lambda, 1^{d \log d})$ and sets $\text{ct} \leftarrow \text{LHE.Enc}(\text{sk}, P)$.

E Predicate Encryption: Key-Policy ABE with Unbounded Decryption Depth

In this section, we construct a predicate encryption (PE) scheme which achieves 1-sided security from any key-policy attribute based encryption (KP-ABE) scheme, a fully homomorphic encryption (FHE) and a lockable obfuscator. Our construction inherits the attribute space and predicate class of the underlying

KP-ABE scheme.

Outline. The main idea is to double encrypt the message m , i.e. first encrypt the message using the KP-ABE scheme (under attribute x) to get ciphertext ct_x , and then encrypt ct_x using the FHE scheme to compute the final ciphertext ct . At a high level, since the ciphertext ct_x is encrypted as well, therefore the attribute x is hidden from the adversary. However, to decrypt such a ciphertext, one would require the FHE secret key. Now if the FHE secret key is released in the clear then we could not hope to use FHE security to claim attribute hiding.

To solve this problem, instead of giving the FHE secret key in the clear, we obfuscate the decryption circuit of the FHE scheme with the secret key hardwired in it. Such an obfuscated circuit would take as input an FHE ciphertext, decrypt it inside (using the hardwired secret key) and check against the lock value. Note that we have not yet specified the lock and message used during obfuscation. If we simply use the message m as the lock value, then we do not know how to prove security since the obfuscator guarantees security only if the lock is chosen at random. Therefore, instead of double encrypting the message m , we double encrypt a randomly chosen lock value α (first under attribute x , then using FHE). And, the message m is set as the output of obfuscated program, i.e. obfuscator is run with lock α and message m . The final ciphertext is set as ct (i.e., the double encryption of lock α) and \tilde{P} (i.e., the obfuscation of decryption circuit). The key generation proceeds identically to that of the underlying ABE scheme. Now decryption proceeds in two phases. First, the inner ciphertext ct_x (inside ciphertext ct) is homomorphically decrypted using an ABE secret key sk_C for some predicate C . Second, the homomorphically evaluated ciphertext (say ct') is fed to the obfuscated program \tilde{P} and output of the program evaluation is simply the output of decryption procedure.

Observe that if sk_C is a valid secret key (such that $C(x) = 1$), then ct' would be an FHE encryption of lock α (by correctness of KP-ABE decryption), and output of the obfuscated program with ct' as input would be message m by correctness of the FHE decryption and program evaluation. At a high level, this provides weak attribute hiding because if the adversary can not decrypt the ciphertext containing the lock (i.e., it does not know the secret key for the corresponding attribute), then it has not no information about the lock value. Therefore, by the obfuscation security, we can guarantee that the challenge ciphertext could be simulated instead. Below we describe our construction $\mathcal{PE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ for attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and 1-bit messages.

Let $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ be a key-policy attribute based encryption scheme for set of attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\{0, 1\}^{k(\lambda)}\}_\lambda$, and $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ be a fully homomorphic encryption scheme for 1-bit messages with decryption circuit of depth $d(\lambda)$ and ciphertexts of length $n(\lambda)$, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ be a lockable obfuscator for circuit class $\mathcal{C}_{n,k,k,d}$ (i.e., the class of depth $d(\lambda)$ circuits with $n(\lambda) \cdot k(\lambda)$ bit input and $k(\lambda)$ bit output). The construction follows. For notational convenience, let $k = k(\lambda)$, $n = n(\lambda)$ and $d = d(\lambda)$.

- $\text{Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm runs ABE.Setup to generate public parameters and master secret key as $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, 1^k)$. It sets $\text{pp} = \text{abe.pp}$ and $\text{msk} = \text{abe.msk}$.
- $\text{Enc}(\text{pp}, x, m) \rightarrow \text{ct}$. The encryption algorithm chooses a random message $\alpha \leftarrow \{0, 1\}^k$ and computes $\text{ct}_\alpha \leftarrow \text{ABE.Enc}(\text{pp}, x, \alpha)$. Next, it generates an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and encrypts ct_α as $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_\alpha)$.²² It also obfuscates the FHE decryption circuit (with fhe.sk hardwired) with message m and lock α as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{FHE.Dec}(\text{fhe.sk}, \cdot), m, \alpha)$.

Finally, it outputs the ciphertext as $\text{ct} = (\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$.

²²Note that FHE scheme supports bit encryption. Therefore, to encrypt ct_α , which is a multi-bit message, the FHE.Enc algorithm will be run independently on each bit of ct_α . However, for notational convenience throughout this section we overload the notation and use FHE.Enc and FHE.Dec algorithms to encrypt and decrypt multi-bit messages respectively.

- $\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm runs ABE.KeyGen to generate the secret key as $\text{abe.sk}_C \leftarrow \text{ABE.KeyGen}(\text{msk}, C)$, and outputs $\text{sk}_C = \text{abe.sk}_C$.
- $\text{Dec}(\text{sk}_C, \text{ct}) \rightarrow m$ or \perp . Let $\text{ct} = (\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$. The decryption algorithm evaluates the ABE decryption on ct_{outer} as $\text{ct}_{\text{inner}} = \text{FHE.Eval}(\text{fhe.ek}, \text{ABE.Dec}(\text{sk}_C, \cdot), \text{ct}_{\text{outer}})$. Next, it evaluates the obfuscated program on input ct_{inner} , and outputs $\text{Eval}(\tilde{P}, \text{ct}_{\text{inner}})$.

Correctness. For all $\lambda \in \mathbb{N}$, message $m \in \{0, 1\}$, attribute $x \in \mathcal{X}_\lambda$, public parameters and master secret key $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, 1^k)$, the ciphertext corresponding to message m under attribute x in our scheme is of the form $\text{ct} = (\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$, where $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$, $\alpha \leftarrow \{0, 1\}^k$, $\text{ct}_\alpha \leftarrow \text{ABE.Enc}(\text{abe.pp}, x, \alpha)$, $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_\alpha)$ and $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{FHE.Dec}(\text{fhe.sk}, \cdot), m, \alpha)$.

For any predicate $C \in \mathcal{C}_\lambda$, the corresponding secret key in our scheme is simply $\text{abe.sk}_C \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C)$. Let $\text{ct}_{\text{inner}} = \text{FHE.Eval}(\text{fhe.ek}, \text{ABE.Dec}(\text{abe.sk}_C, \cdot), \text{ct}_{\text{outer}})$. Consider the following two cases:

1. $C(x) = 1$: We know that if $C(x) = 1$, then with all but negligible probability $\text{ABE.Dec}(\text{abe.sk}_C, \text{ct}_\alpha) = \alpha$. Therefore, ct_{inner} is an FHE encryption of lock α . Or in other words, w.h.p. $\text{FHE.Dec}(\text{fhe.sk}, \text{ct}_{\text{inner}}) = \alpha$. This follows from correctness of the ABE and FHE schemes.
Next, by correctness of obfuscation, we can conclude that $\text{Eval}(\tilde{P}, \text{ct}_{\text{inner}}) = m$. Therefore, if $C(x) = 1$, then $\Pr[\text{Dec}(\text{sk}_C = \text{abe.sk}_C, \text{ct}) = m] \geq 1 - \text{negl}(\lambda)$.
2. $C(x) = 0$: Similarly, we know that if $C(x) = 0$, then with all but negligible probability $\text{ABE.Dec}(\text{abe.sk}_C, \text{ct}_\alpha) = \perp$. Therefore, w.h.p. $\text{FHE.Dec}(\text{fhe.sk}, \text{ct}_{\text{inner}}) = \perp$. This follows from correctness of the ABE and FHE schemes.

Also by correctness of obfuscation, we can conclude that if $\text{FHE.Dec}(\text{fhe.sk}, \text{ct}_{\text{inner}}) \neq \alpha$, then $\Pr[\text{Eval}(\tilde{P}, \text{ct}_{\text{inner}}) = \perp] \geq 1 - \text{negl}(\lambda)$. Since $\perp \neq \alpha$, therefore combining these two facts we know that if $C(x) = 0$, then $\Pr[\text{Dec}(\text{sk}_C = \text{abe.sk}_C, \text{ct}) = \perp] \geq 1 - \text{negl}'(\lambda)$ for some negligible function $\text{negl}'(\cdot)$.

Therefore, \mathcal{PE} satisfies the predicate encryption correctness condition.

Security. We will now show that the scheme described above achieves 1-sided security as per Definition A.2. Formally, we prove the following.

Theorem E.1. If $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a fully secure attribute based encryption for set of attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\{0, 1\}^{k(\lambda)}\}_\lambda$ satisfying Definition A.1, and $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ is a fully homomorphic encryption scheme satisfying Definition 2.4, and $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a lockable obfuscator for 1-bit messages and circuit class $\{\mathcal{C}_{n,k,\text{poly}}\}_\lambda$ satisfying Definition 3.4, then \mathcal{PE} is a secure predicate encryption scheme satisfying 1-sided security as per Definition A.3 for 1-bit messages and same attribute space and predicate class as the ABE scheme.

Our proof proceeds via a sequence of hybrid games. Each game is played between the challenger and attacker \mathcal{A} . Let \mathcal{A} be any PPT adversary that wins the 1-sided security game with non-negligible advantage. We argue that such an adversary must break security of at least one underlying primitive. The first game corresponds to the 1-sided security game as described in Definition A.2. In the final game, the challenge ciphertext is simulated and does not contain any information about the challenge messages.

We will first define the sequence of hybrid games, and then show that they are computationally indistinguishable.

Game 1: In this game, the challenge ciphertext ct is honestly generated, i.e. first a uniformly random lock α is chosen along with FHE keys $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek})$. Next, it is encrypted under attribute x to get ciphertext ct_{inner} . Additionally, it encrypts ct_{inner} under fhe.pk . Finally, the challenger obfuscates the FHE decryption circuit with secret key fhe.sk hardwired for lock α and challenge message m .

1. **Setup Phase.** The challenger sets up by generating public parameters and master secret key as $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, \mathcal{X}_\lambda, \mathcal{C}_\lambda, 1^k)$. It sends abe.pp to \mathcal{A} .
2. **Pre-Challenge Query Phase.** \mathcal{A} queries the challenger on polynomially many predicate circuits C_i to receive the corresponding secret keys as $\text{sk}_{C_i} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_i)$.
3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$, a uniformly random lock string $\alpha \leftarrow \{0, 1\}^k$, and an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$. It computes ciphertexts $\text{ct}_{\text{inner}}, \text{ct}_{\text{outer}}$ as $\text{ct}_{\text{inner}} \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, \alpha)$ and $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_{\text{inner}})$. It also computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{FHE.Dec}(\text{fhe.sk}, \cdot), m_b, \alpha)$, and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} .
4. **Post-Challenge Query Phase.** \mathcal{A} queries the challenger on polynomially many predicate circuits C_i as before. The challenger handles these as in pre-challenge query phase.
5. **Guess.** \mathcal{A} outputs its guess b' and wins if $b = b'$.

Game 2: This is same as Game 1, except the challenger computes ciphertext ct_{inner} as an encryption of an all zero string instead of α .

3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$, a uniformly random lock string $\alpha \leftarrow \{0, 1\}^k$, and an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$. It computes ciphertexts $\text{ct}_{\text{inner}}, \text{ct}_{\text{outer}}$ as $\text{ct}_{\text{inner}} \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, 0^k)$ and $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_{\text{inner}})$. It also computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{FHE.Dec}(\text{fhe.sk}, \cdot), m_b, \alpha)$, and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} .

Game 3: This is same as Game 2, except the challenger does not choose the lock α anymore and it simulates the obfuscated program \tilde{P} instead of generating it honestly as an obfuscation of the FHE decryption circuit.

3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$ and an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$. It computes ciphertexts $\text{ct}_{\text{inner}}, \text{ct}_{\text{outer}}$ as $\text{ct}_{\text{inner}} \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, 0^k)$ and $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_{\text{inner}})$. It also computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$ (where $s = |\text{FHE.Dec}(\text{fhe.sk}, \cdot)|$), and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} .

Game 4: This is same as Game 3, except the challenger computes ciphertext ct_{outer} as an encryption of an all zero string instead of ct_{inner} .

3. **Challenge.** Next, \mathcal{A} sends the challenge messages and attributes $((m_0, x_0), (m_1, x_1))$ to the challenger such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. The challenger chooses a random bit $b \leftarrow \{0, 1\}$ and an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$. It computes ciphertexts ct_{outer} as $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, 0^\ell)$, where ℓ be the length of ABE ciphertexts. It also computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$, and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} .

Let $\text{Adv}_{\mathcal{A}}^i = |\Pr[b' = b] - 1/2|$ denote the advantage of adversary \mathcal{A} in guessing the bit b in Game i . First, note that $\text{Adv}_{\mathcal{A}}^4 = 0$. In other words, \mathcal{A} 's advantage in Game 4 is 0. This is because the challenge ciphertext $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ does not contain any information about bit b .

To complete the proof, we establish via a sequence of lemmas that no PPT adversary \mathcal{A} can distinguish between each adjacent game with non-negligible probability. We show via a sequence of lemmas that $|\text{Adv}_{\mathcal{A}}^i - \text{Adv}_{\mathcal{A}}^{i+1}|$ is negligible for all $i = 1, 2, 3$. Below we discuss our lemmas in detail.

Lemma E.1. If $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a fully secure attribute based encryption, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish encryptions of lock α and an all zeros string, therefore break security of the ABE scheme.

The ABE challenger generates a key pair $(\text{abe.pp}, \text{abe.sk})$ and sends abe.pp to \mathcal{B} . \mathcal{B} simply forwards the public parameters abe.pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} forwards each predicate query C_i to the ABE challenger and receives back secret key sk_{C_i} , which it then sends to \mathcal{A} as its response. Next, \mathcal{A} sends the challenge message-attribute pairs $((m_0, x_0), (m_1, x_1))$ to \mathcal{B} such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. \mathcal{B} chooses a random string $\alpha \leftarrow \{0, 1\}^k$ and bit $b \leftarrow \{0, 1\}$, and sends $t_0 = \alpha$ and $t_1 = 0^k$ as its challenge messages and x_b as the challenge attribute to the ABE challenger. The ABE challenger chooses a random bit $\beta \leftarrow \{0, 1\}$, computes the challenge ciphertext $\text{ct}^* \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, t_\beta)$ and sends ct^* to \mathcal{B} . After receiving ct^* from the challenger, \mathcal{B} runs Step 3 as in Game 1. That is, \mathcal{B} chooses an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$ and computes ciphertext ct_{outer} as $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}^*)$. It also computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \text{Obf}(1^\lambda, \text{FHE.Dec}(\text{fhe.sk}, \cdot), m_b, \alpha)$, and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} . Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before by forwarding those to ABE challenger. Finally, \mathcal{A} outputs its guess b' . If $b = b'$, then \mathcal{B} sends 0 as its guess (i.e., α was encrypted), otherwise it sends 1 as its guess (i.e., 0^k was encrypted) to the ABE challenger.

Note that if the ABE challenger encrypted α (i.e., $\beta = 0$), then \mathcal{B} perfectly simulates Game 1 for adversary \mathcal{A} . Otherwise it simulates Game 2 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^1 - \text{Adv}_{\mathcal{A}}^2|$ is non-negligible, then \mathcal{B} breaks the ABE scheme's security with non-negligible advantage. \blacksquare

Lemma E.2. If $\mathcal{O} = (\text{Obf}, \text{Eval})$ is a secure lockable obfuscator, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish an obfuscation of the FHE decryption circuit with a random lock from a simulated obfuscated program, therefore break security of the obfuscation scheme.

\mathcal{B} samples an ABE key pair $(\text{abe.pp}, \text{abe.sk})$ and sends abe.pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} answers each query with secret key sk_{C_i} , with $\text{sk}_{C_i} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_i)$. Next, \mathcal{A} sends the challenge message-attribute pairs $((m_0, x_0), (m_1, x_1))$ to \mathcal{B} such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. \mathcal{B} chooses a random bit $b \leftarrow \{0, 1\}$ and an FHE key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek}) \leftarrow \text{FHE.Setup}(1^\lambda)$, and computes ciphertext $\text{ct}_{\text{inner}}, \text{ct}_{\text{outer}}$ as $\text{ct}_{\text{inner}} \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, 0^k)$ and $\text{ct}_{\text{outer}} \leftarrow \text{FHE.Enc}(\text{fhe.pk}, \text{ct}_{\text{inner}})$. \mathcal{B} sends the $\text{FHE.Dec}(\text{fhe.sk}, \cdot)$ circuit along with message m_b to the obfuscation challenger, and receives an obfuscated program \tilde{P} . \mathcal{B} finally sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}_{\text{outer}}, \tilde{P})$ to \mathcal{A} as the challenge ciphertext. Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before. Finally, \mathcal{A} outputs its guess b' . If $b = b'$, then \mathcal{B} sends 0 as its guess (i.e., FHE decryption circuit was obfuscated), otherwise it sends 1 as its guess (i.e., program \tilde{P} was simulated) as its guess to the obfuscation challenger.

Note that if the obfuscation challenger obfuscated $\text{FHE.Dec}(\text{fhe.sk}, \cdot)$ for some lock α , then \mathcal{B} perfectly simulates Game 2 for adversary \mathcal{A} . Otherwise it simulates Game 3 for \mathcal{A} . As a result, if $|\text{Adv}_{\mathcal{A}}^2 - \text{Adv}_{\mathcal{A}}^3|$ is non-negligible, then \mathcal{B} breaks the obfuscation scheme's security with non-negligible advantage. \blacksquare

Lemma E.3. If $\text{FHE} = (\text{FHE.Setup}, \text{FHE.Enc}, \text{FHE.Dec}, \text{FHE.Eval})$ is a fully homomorphic encryption scheme, then for all PPT adversaries \mathcal{A} , $|\text{Adv}_{\mathcal{A}}^3 - \text{Adv}_{\mathcal{A}}^4|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|\text{Adv}_{\mathcal{A}}^3 - \text{Adv}_{\mathcal{A}}^4|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish encryptions of ciphertext ct_{inner} and an all zeros string, therefore break security of the FHE scheme.

The FHE challenger generates a key pair $(\text{fhe.pk}, \text{fhe.sk}, \text{fhe.ek})$ and sends $\text{fhe.pk}, \text{fhe.ek}$ to \mathcal{B} . \mathcal{B} samples an ABE key pair $(\text{abe.pp}, \text{abe.sk})$ and sends abe.pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} answers each query with secret key sk_{C_i} , with $\text{sk}_{C_i} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_i)$. Next, \mathcal{A} sends the challenge message-attribute pairs $((m_0, x_0), (m_1, x_1))$ to \mathcal{B} such that $C_i(x_0) = C_i(x_1) = 0$ for all queried predicates. \mathcal{B} chooses a random bit $b \leftarrow \{0, 1\}$ and computes ciphertext ct_{inner} as $\text{ct}_{\text{inner}} \leftarrow \text{ABE.Enc}(\text{abe.pp}, x_b, 0^k)$, and sends $t_0 = \text{ct}_{\text{inner}}$ and $t_1 = 0^\ell$ as its challenge messages to the FHE challenger. The FHE challenger chooses a random bit $\beta \leftarrow \{0, 1\}$, computes the challenge ciphertext $\text{ct}^* \leftarrow \text{FHE.Enc}(\text{fhe.pk}, t_\beta)$ and sends ct^* to \mathcal{B} . After receiving ct^* from the challenger, \mathcal{B} computes obfuscated program \tilde{P} as $\tilde{P} \leftarrow \mathcal{O}.\text{Sim}(1^\lambda, 1^s, 1)$, and sends $(\text{fhe.pk}, \text{fhe.ek}, \text{ct}^*, \tilde{P})$ to \mathcal{A} . Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before. Finally, \mathcal{A} outputs its guess b' . If $b = b'$, then \mathcal{B} sends 0 as its guess (i.e., ct_{inner} was encrypted), otherwise it sends 1 (i.e., 0^ℓ was encrypted) as its guess to the FHE challenger.

Note that if the FHE challenger encrypted ct_{inner} (i.e., $\beta = 0$), then \mathcal{B} perfectly simulates Game 3 for adversary \mathcal{A} . Otherwise it simulates Game 4 for \mathcal{A} . Therefore, if $|\text{Adv}_{\mathcal{A}}^3 - \text{Adv}_{\mathcal{A}}^4|$ is non-negligible, then \mathcal{B} breaks the FHE scheme's security with non-negligible advantage. \blacksquare

F Predicate Encryption: SIM-1-Sided Security

In this section, we transform a PE scheme from one which achieves 1-sided security to another scheme that achieves SIM-1-sided security. The transformation inherits the message space of the underlying PE scheme. However, the attribute space and predicate class are slightly smaller than original.

Let $\text{PE} = (\text{PE.Setup}, \text{PE.Enc}, \text{PE.KeyGen}, \text{PE.Dec})$ be a predicate encryption scheme for set of attribute spaces $\{\mathcal{X}_\lambda\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\mathcal{M}_\lambda\}_\lambda$. Below we describe our construction.

For each boolean predicate $C : \{0, 1\}^n \rightarrow \{0, 1\}$, we use \tilde{C} to denote the boolean predicate on $n + 1$ bits such that $\tilde{C}(x) = 1$ iff $x_1 = 1 \wedge C(x_2 || \dots || x_{n+1})$, where x_i is i^{th} bit of x . In other words, \tilde{C} evaluates to 1 iff the first input bit is 1 and predicate C outputs 1 on remaining bits as input. Similarly, for a predicate class \mathcal{C}_λ , we use $\tilde{\mathcal{C}}_\lambda$ to denote the predicate class that contains predicate \tilde{C} for each predicate $C \in \mathcal{C}_\lambda$.

- $\text{Setup}(1^\lambda, 1^\ell, \mathcal{C}_\lambda, \mathcal{M}_\lambda) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm runs PE.Setup to generate public parameters and master secret key as $(\text{pp}, \text{msk}) \leftarrow \text{PE.Setup}(1^\lambda, 1^{\ell+1}, \tilde{\mathcal{C}}_\lambda, \mathcal{M}_\lambda)$.
- $\text{Enc}(\text{pp}, x, m) \rightarrow \text{ct}$. The encryption algorithm computes ciphertext ct as $\text{ct} \leftarrow \text{PE.Enc}(\text{pp}, 1 || x, m)$.
- $\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm runs PE.KeyGen to generate the secret key as $\text{sk}_C \leftarrow \text{PE.KeyGen}(\text{msk}, \tilde{C})$.
- $\text{Dec}(\text{sk}_C, \text{ct}) \rightarrow m$. The decryption algorithm runs PE.Dec and outputs $m = \text{PE.Dec}(\text{sk}_C, \text{ct})$.

Correctness. The correctness of above scheme follows directly from the correctness of underlying PE scheme.

Security. We will now show that the scheme described above achieves 1-sided security as per Definition A.2. Formally, we prove the following.

Theorem F.1. If $\text{PE} = (\text{PE.Setup}, \text{PE.Enc}, \text{PE.KeyGen}, \text{PE.Dec})$ is a secure predicate encryption scheme satisfying 1-sided security as per Definition A.2 for set of attribute spaces $\{\{0, 1\}^{\ell(\lambda)+1}\}_\lambda$, predicate classes $\{\tilde{\mathcal{C}}_\lambda\}_\lambda$ and message spaces $\{\mathcal{M}_\lambda\}_\lambda$ satisfying Definition A.2, then \mathcal{PE} is a secure predicate encryption scheme satisfying SIM-1-sided security as per Definition A.3 for set of attribute spaces $\{\{0, 1\}^{\ell(\lambda)}\}_\lambda$, predicate classes $\{\mathcal{C}_\lambda\}_\lambda$ and message spaces $\{\mathcal{M}_\lambda\}_\lambda$.

We start by sketching the simulator Sim . Let \mathcal{A} be any real world adversary. On input the security parameter λ , public parameters pp , attribute length $\ell + 1$, the simulator runs as follows.

- Sim encrypts all 0's string as the message under the all 0's string as an attribute. Concretely, it outputs $\text{ct} = \text{PE.Enc}(\text{pp}, 0^{|m|}, 0^{\ell+1})$.

Let $p_{\mathcal{A}}^1$ denote the probability that adversary \mathcal{A} outputs 1 in the 1-sided security game as described in Definition A.2 and $p_{\mathcal{A}}^2$ denote the probability that adversary \mathcal{A} outputs 1 in simulated execution (i.e., when the challenge ciphertext is simulated). We show that $|p_{\mathcal{A}}^1 - p_{\mathcal{A}}^2|$ is negligible.

Lemma F.1. If $\text{PE} = (\text{PE.Setup}, \text{PE.Enc}, \text{PE.KeyGen}, \text{PE.Dec})$ is a secure predicate encryption scheme satisfying 1-sided security, then for all PPT adversaries \mathcal{A} , $|p_{\mathcal{A}}^1 - p_{\mathcal{A}}^2|$ is negligible in the security parameter λ .

Proof. Suppose there exists an adversary \mathcal{A} such that $|p_{\mathcal{A}}^1 - p_{\mathcal{A}}^2|$ is non-negligible. We construct an algorithm \mathcal{B} that can distinguish encryptions of challenge message m under attribute x and an all zeros string under all zeros attribute string, therefore break 1-sided security of the PE scheme.

The PE challenger generates a key pair (pp, msk) and sends pp to \mathcal{B} . \mathcal{B} simply forwards the public parameters pp to adversary \mathcal{A} . \mathcal{A} queries the reduction algorithm \mathcal{B} on polynomially many predicates C_i for corresponding secret keys. \mathcal{B} queries predicate \widetilde{C}_i to the PE challenger and receives back secret key sk_{C_i} , which it then sends to \mathcal{A} as its response. Next, \mathcal{A} sends the challenge message-attribute pair (m, x) to \mathcal{B} such that $C_i(x) = 0$ for all queried predicates. \mathcal{B} sends $((m, 1 || x), (0^{|m|}, 0^{\ell+1}))$ as its challenge message-attribute pairs to the PE challenger. The PE challenger chooses a random bit $b \leftarrow \{0, 1\}$, computes the challenge ciphertext $\text{ct} \leftarrow \text{PE.Enc}(\text{pp}, x_b, m_b)$ (where $x_0 = x, m_0 = m, x_1 = 0^{\ell+1}, m_1 = 0^{|m|}$), and sends ct to \mathcal{B} . \mathcal{B} forwards ct as the challenge ciphertext to \mathcal{A} . Next, \mathcal{A} makes more key queries and \mathcal{B} answers them as before by forwarding those to the PE challenger. Finally, \mathcal{A} outputs a bit b' and \mathcal{B} outputs b' as its guess to the PE challenger.

We would like to point out that all the predicate queries made by \mathcal{A} can be forwarded to the PE challenger because by definition $\widetilde{C}(0^{\ell+1}) = 0$ for all predicates C . Also, note that if the PE challenger encrypted m under attribute x (i.e., $b = 0$), then \mathcal{B} perfectly simulates the 1-sided security game for adversary \mathcal{A} . Otherwise it simulates the challenge ciphertext for \mathcal{A} . As a result, if $|p_{\mathcal{A}}^1 - p_{\mathcal{A}}^2|$ is non-negligible, then \mathcal{B} breaks the PE scheme's security with non-negligible advantage. ■