

COLLUSION RESISTANT TRAITOR TRACING FROM LEARNING WITH ERRORS*

RISHAB GOYAL[†], VENKATA KOPPULA[†], AND BRENT WATERS[†]

Abstract. In this work we provide a traitor tracing construction with ciphertexts that grow polynomially in $\log(n)$, where n is the number of users, and prove it secure under the learning with errors (LWE) assumption. This is the first traitor tracing scheme with such parameters provably secure from a standard assumption. In addition to achieving new traitor tracing results, we believe our techniques push forward the broader area of computing on encrypted data under standard assumptions. Notably, traitor tracing is a substantially different problem from other cryptography primitives that have seen recent progress in LWE solutions. We achieve our results by first conceiving a novel approach to building traitor tracing that starts with a new form of functional encryption that we call Mixed FE. In a Mixed FE system the encryption algorithm is bimodal and works with either a public key or master secret key. Ciphertexts encrypted using the public key can only encrypt one type of functionality. On the other hand, the secret key encryption can be used to encode many different types of programs, but is only secure as long as the attacker sees a bounded number of such ciphertexts. We first show how to combine mixed FE with attribute-based encryption to achieve traitor tracing. Second, we build Mixed FE systems for polynomial-sized branching programs (which corresponds to the complexity class logspace) by relying on the polynomial hardness of the LWE assumption with superpolynomial modulus-to-noise ratio.

Key words. traitor tracing, private linear broadcast encryption, functional encryption, Mixed FE, attribute-based encryption, lattice trapdoors, encryption, decryption, learning with errors

AMS subject classifications. 68P25, 94A60

DOI. 10.1137/18M1197825

1. Introduction. In a (traitor) tracing [CFN94] system an authority runs a setup algorithm that takes in a security parameter λ and the number, n , of users in the system. The setup outputs a public key pk , master secret key msk , and n secret keys $(\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n)$. The system has an encryption algorithm that uses the public key pk to create a ciphertext for a message m that is decryptable by any of the n secret keys, but where the message will be hidden from any user who does not have access to the keys. Finally, suppose that some subset S of users colludes to create a decoding box D which is capable of decrypting ciphertexts with some nonnegligible probability. The tracing property of the system states that there exists an algorithm Trace which, given the master secret and oracle access to D , outputs a set of users T where T contains at least one user from the colluding set S and no users outside of S .

Existing approaches for achieving collusion resistant broadcast encryption can be fit in the framework of private linear broadcast encryption (PLBE) introduced by Boneh, Sahai, and Waters (BSW) [BSW06]. In a PLBE system the setup algorithm takes as input a security parameter λ and the number of users n . It outputs a public key pk , master secret key msk , and n private keys $\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n$ where a user with index j is given key sk_j . Any of the private keys is capable of decrypting a ciphertext

*Received by the editors July 2, 2018; accepted for publication (in revised form) June 5, 2019; published electronically November 5, 2019.

<https://doi.org/10.1137/18M1197825>

Funding: This research was supported by NSF CNS-1414082, by DARPA through the U.S. Office of Naval Research under contract N00014-11-1-0382, by a Microsoft Faculty Fellowship, and by a Packard Foundation Fellowship.

[†]Department of Computer Science, University of Texas at Austin, Austin, TX 78712 (rishab.28.goyal@gmail.com, k.venkata.vk@gmail.com, bwaters@cs.utexas.edu).

ct created using pk . However, there is an additional TrEncrypt algorithm that takes in the master secret key, a message, and an index i . This produces a ciphertext that only users with index $j > i$ can decrypt. Moreover, any adversary-produced decryption box D that was created with private keys in the set of S will not be able to distinguish between encryptions to index $i - 1$ or i , where $i \notin S$. In addition, encryptions of two different messages m_0, m_1 to index n must be indistinguishable.

The tracing system is set up by simply running the PLBE setup and distributing each PLBE key to the corresponding user. To trace the set of colluding parties given a decoding box D , the tracing algorithm first measures (with several samples) the probability that D correctly decrypts a ciphertext encrypted to index i for all $i \in [0, n]$. If the box D originally decrypted with probability ϵ , then there must exist some index i where the probability the box decrypts on index $i - 1$ is at least ϵ/n more than the probability it decrypts on ciphertexts encrypted to index i , since by PLBE security D cannot decrypt encryptions to index n with nonnegligible probability. At this point the tracing algorithm marks user i as a colluder.

Currently, there are three approaches to building PLBE. The most basic approach is to simply create n public/private key pairs under a standard IND-CPA secure public key encryption system. A PLBE ciphertext is formed by encrypting the message m to each user's public key individually and concatenating all of the subciphertexts to form one long ciphertext, $\text{ct} = (\text{ct}_1, \text{ct}_2, \dots, \text{ct}_n)$. A user with secret key sk_i in the system will decrypt by running decryption on ct_i and ignore the rest of the ciphertext components. To TrEncrypt to index i simply encrypt the all 0's string in first i ciphertexts $\text{ct}_1, \dots, \text{ct}_i$ in place of the message. The index hiding property follows directly from IND-CPA security of the underlying encryption scheme, since without secret key i no attacker can determine whether ct_i is an encryption of the message or all 0's string.

The above approach works because there is a portion of the ciphertext ct_i dedicated to each user i in the system which is not touched during the decryption process of other users with keys sk_j for $j \neq i$. This dedicated ciphertext space strategy makes it easy to silently kill user i 's ability to access the message in a way unnoticeable to other users, but also inherently requires a ciphertext size that grows linearly in n . In order to achieve PLBE with sublinear size ciphertexts, one needs to implement some form of computing on encrypted data.

BSW [BSW06] provided the first construction that achieved PLBE with ciphertext growth that was sublinear in n . They leveraged composite order bilinear groups to achieve ciphertexts that grew proportionally to $\sqrt{n}$. While future variants [BW06, GKS010, Fre10] used bilinear maps to obtain additional properties, the ciphertext size for all bilinear map-based constructions remained stuck at the $\sqrt{n}$ mark.

Several years later Boneh and Zhandry [BZ14] showed how to utilize indistinguishability obfuscation and apply punctured programming techniques to achieve the ideal case where ciphertexts grow polynomially in $\log(n)$ and λ . The downside of applying indistinguishability obfuscation is that all current obfuscation candidates are based on nonstandard multilinear map group assumptions, and several such multilinear candidates have been attacked (see [CLT14, CHL+15, CGH+15, BGH+15, CLLT16, CLLT17, BWZ14, HJ16, Hal15, CFL+16, MSZ16, CJL16, ADGM16] and the references therein). (One could also achieve similar results by using the functional encryption scheme of Garg et al. [GGH+13], but this also relies on multilinear maps.) This leaves open the following question:

Can we build secure traitor tracing with $\text{POLY}(\lambda, \log(n))$ -sized ciphertexts from standard assumptions?

Our results. In this work we resolve the above question by providing a traitor tracing construction with ciphertexts that grow polynomially in $\log(n)$ and λ and prove it secure under the learning with errors (LWE) assumption. This is the first traitor tracing scheme with such parameters that is provably secure from a standard assumption. In addition to achieving new traitor tracing results, we believe our techniques push forward the broader area of computing on encrypted data under standard assumptions. Notably, traitor tracing is a substantially different problem from other cryptography primitives that have seen recent progress in LWE solutions.

We achieve our result by first conceiving a novel approach to building traitor tracing that starts with a new form of functional encryption that we call Mixed FE. In a Mixed FE system the encryption algorithm is bimodal and works with either a public key or master secret key. Ciphertexts encrypted using the public key can only encrypt one type of functionality. On the other hand, the secret key encryption can be used to encode many different types of programs, but is only secure as long as the attacker sees a bounded number of such ciphertexts.

We first show how to combine Mixed FE with attribute-based encryption (ABE) to achieve traitor tracing. Second, we show under the LWE assumption how to construct Mixed FE systems for polynomial-sized branching programs (which corresponds to the complexity class LOGSPACE).

1.1. Technical overview. We now give a technical overview of our work. This overview is broken into four parts. In the first part we review the BSW notion of private linear broadcast encryption (PLBE) and its transformation into a traitor tracing system. Along the way we discover that the PLBE definitions as presented in [BSW06] do *not* imply traitor tracing. We then show how to repair the argument by giving the attacker an additional oracle encryption query in the PLBE definitions. Second, we present the notion of Mixed FE and show how an ABE and Mixed FE system (for the right functionalities) can be used to construct a PLBE system. The third part of our overview describes a new LWE toolkit which includes “enhanced” versions of lattice trapdoor sampling algorithms with additional security properties. Finally, we outline our main ideas for constructing the Mixed FE system and proving it secure under the LWE assumption.

Part 1: Breaking and repairing the PLBE to tracing argument. First, let us review the PLBE algorithms as defined in [BSW06]. A PLBE scheme consists of a setup, encryption, decryption, and trace-encryption algorithm. The setup algorithm outputs a public key, a master secret key, and n secret keys, one for each index in $[n]$. The encryption/decryption algorithms are self-explanatory; the trace-encryption algorithm is a special encryption algorithm that requires the master secret key, and can be used to encrypt a message to any index $i \in [0, n]$. The output ciphertext can be decrypted only by secret keys for indices $j > i$. BSW defined three security properties. The first security property (public to zero-index indistinguishability) requires that a standard encryption of message m is indistinguishable from a trace-encryption of m to the index 0, even when the adversary has all n secret keys. The second security property (index hiding) states that a trace-encryption of m to index $i - 1$ is indistinguishable from a trace-encryption of m to index i , even when the adversary has all the secret keys except the i th one. Finally, the third security property states that trace-encryption of m_0 to index n is indistinguishable from trace-encryption of m_1 to index n , even when the adversary is given all n secret keys.

BSW argued that these three properties of PLBE are sufficient for constructing a traitor tracing (TT) scheme. In their transformation, the TT public key and n

secret keys are set to be the PLBE public key and n secret keys, respectively. The TT encryption/decryption algorithms are identical to the PLBE encryption/decryption algorithms. Finally, the tracing algorithm uses the PLBE trace-encryption algorithm. Given a decoder box D , the tracing algorithm encrypts random messages to each index and checks if D can decrypt it correctly. If the decoder box is ϵ -successful¹ in decrypting (standard) encryptions, then it is also ϵ successful in decrypting trace-encryptions to index 0 (via the first security property). Next, note that the decoder box cannot decrypt trace-encryptions to index n (via the message indistinguishability property). Therefore, there must exist an index $i^* \in [n]$ where the success of the decoder box in decrypting trace-encryptions to index $i^* - 1$ is at least ϵ/n more than its success in decrypting trace-encryptions to i^* . This index i^* must be one of the indices queried by the adversary (since if the adversary does not have a key for index i^* , then the decoder box must not be able to distinguish between trace-encryptions to $i^* - 1$ and i^*). For each index i , the tracing algorithm computes an estimate of the decoder box's success probability in decrypting random trace-encryptions for index i . For all indices i where the measured success probabilities for $i - 1$ and i are substantially different, user i is declared to be a traitor.

At an intuitive level, it seems like the BSW transformation should work. However, here we argue that it is indeed possible to have a PLBE scheme secure under the original BSW definition, but produce an insecure TT scheme in this regard. The problem lies in the fact that there is a “semantic gap” between the TT definition and the PLBE definition. The TT definition considers an attacker that produces a (stateless) decoder D whose success on decrypting multiple trace-encryptions is measured, whereas the PLBE definition considers indistinguishability on a single ciphertext (in particular, no ciphertext queries). Diving deeper, we show a separation by adding a feature to a PLBE scheme where the feature does not impact PLBE security, but results in an insecure TT scheme.

Given any secure PLBE scheme P , consider a scheme P' defined as follows. The setup algorithm of P' is similar to the setup of P , except it also samples an additional pseudorandom function (PRF) key K as part of the master secret key (we will assume the PRF has single-bit output). The (standard) encryption algorithm computes a ciphertext ct using the underlying scheme's encryption algorithm, chooses a uniformly random bit b , and outputs (ct, b) . The trace-encryption of message m is the ciphertext $ct' = (ct, y = \text{PRF}_K(i))$, where ct is the ciphertext obtained from the trace-encryption algorithm of P . It is easy to see that the new scheme satisfies all three PLBE security definitions, since there are no encryption queries allowed in the PLBE scheme beyond the challenge ciphertext.

However, it is possible to construct a decoding box using only the secret key for index n such that the trace algorithm falsely accuses some user $i < n$. The decoder D , on input of a ciphertext $ct' = (ct, y)$, tests if $y = 1$. If so, it decrypts the ciphertext using key sk_n ; otherwise it outputs a random message. Using PRF security, we can argue that there exists an index $i < n$ such that $\text{PRF}_K(i - 1) = 1$ and $\text{PRF}_K(i) = 0$ with high probability. In this case the probability that D decrypts ciphertexts for index $i - 1$ will be measurably different than the case in which it decrypts ciphertext for index i . Thus user i will be flagged as a colluder.

We repair the BSW transformation from PLBE to TT by considering a modified set of PLBE security definitions and prove that these do imply TT. We do so in two

¹A decoder box is said to be ϵ -successful if its probability of correctly decrypting a ciphertext is at least ϵ , where the probability is taken over the choice of the ciphertext and D 's random coins.

steps. First, we consider a *decoder-based* version of the BSW PLBE definitions. For concreteness, let us consider the index hiding definition. The decoder-based version of the index hiding version states that no adversary, given all secret keys except the i th one, can produce a decoder box D and a message m such that D can distinguish between trace-encryptions of m to index $i - 1$ and trace-encryptions of m to index i . Decoder-based versions of the other properties are defined similarly.

Now that we have decoder-based PLBE definitions that align with the decoder in the TT definitions, it is fairly straightforward to prove that the BSW transformation implies TT. The downside of introducing decoder-based PLBE definitions is that they are more difficult to work with as a target for a construction. We address this issue by circling back to the original (BSW) PLBE definitions and augmenting them by allowing an attacker to make an a priori bounded number of queries to an encryption oracle. We show that 1-query PLBE implies decoder-based PLBE. This gives us an easier target (that is, 1-query PLBE).

Before describing the transformation from 1-query PLBE to decoder-based PLBE, we would like to point out that if the BSW definitions were augmented to allow an unbounded number of ciphertext queries, then decoder-based security follows immediately. For instance, let us consider the index hiding game. The reduction algorithm (that reduces unbounded-query PLBE to decoder-based PLBE) receives a decoder box D from the attacker. Given the unbounded queries, the reduction algorithm can measure (with reasonable accuracy) the success probabilities of D for indices $i - 1$ and i , and therefore whether it can use D to distinguish between an encryption to index $i - 1$ and i . However, with only 1 encryption query no such precise measurement is possible. Therefore, showing an attacker on decoder-based PLBE security implies an attacker on 1-query PLBE is a bit tricky. The reduction algorithm, after receiving the decoder box and message m from the adversary, chooses a random index $i^* \in \{i - 1, i\}$, and queries the challenger for encryption of m for index i^* . It receives a ciphertext ct . Next, it queries the challenger with challenge message m and receives a challenge ciphertext ct^* . The reduction algorithm checks if $D(ct) = D(ct^*)$; if so, it guesses that m was encrypted for index i^* . We would like to point out that choosing query index i^* uniformly at random from $\{i - 1, i\}$ (as opposed to just fixing one of the two) is important for our analysis. The idea of running the decoder twice is sometimes referred to as the double-run trick [BDK+11, BG11, DY13]. The complete details of our analysis can be found in section 4.1.

Impact on prior TT works using the PLBE framework. Traitor tracing schemes that had secret key tracing would need a new proof under the new PLBE definitions with 1-query allowed. We believe the bilinear map constructions [BSW06, GKS010, Fre10] are likely secure under this definition, but showing this is outside scope of this paper. Note that the same problem is not present in PLBE with public trace-encryption (e.g., [BW06]), since the public key allows the reduction algorithm to generate ciphertexts.

Part 2: Constructing PLBE from Mixed FE. The hardness of constructing a PLBE scheme stems from the fact that it needs to satisfy the following two properties at the same time. First, a PLBE scheme needs to provide a predicate encryption (PE) like functionality where each secret key is associated with an “index,” and each ciphertext is associated with an *index comparison* predicate. Also, the ciphertexts must not reveal any more information about the associated index comparison predicate other than what can be learned by running decryption. Second, the scheme must provide a broadcast encryption (BE) like compactness guarantee, which is that the

size of ciphertexts must be short. In particular, the index needs to be represented in binary, which means the index comparison must be “sufficiently complex.”

In this work, instead of directly building a PLBE scheme, we further reduce the task to constructing a new form of functional encryption scheme called Mixed FE. We show how Mixed FE can be combined with ABE for circuits to obtain PLBE. At a very high level, our approach is to decouple the functionality (delivering the message to users) and security requirements of a PLBE scheme, and to deal with them separately.

We begin by informally introducing the notion of Mixed FE. A Mixed FE scheme consists of setup, normal (or public key) encryption, *secret key encryption*, key generation, and decryption algorithms. The setup algorithm takes as input the security parameter λ and description of a function class $\mathcal{F}$ and outputs the public parameters $\mathbf{pp}$ and the master secret key $\mathbf{msk}$. The normal encryption algorithm only takes as input the public parameters $\mathbf{pp}$ and outputs a (normal) ciphertext $\mathbf{ct}$. The secret key encryption algorithm takes as input the master secret key $\mathbf{msk}$ and a function $f \in \mathcal{F}$ and outputs a (secret key) ciphertext $\mathbf{ct}$. The key generation algorithm takes as input the master secret key $\mathbf{msk}$ and a message m and outputs a key $\mathbf{sk}_m$. The decryption algorithm takes as input a ciphertext $\mathbf{ct}$ and a secret key $\mathbf{sk}_m$ and outputs a single bit. Now for correctness we require that decrypting a secret key encryption of any function f using a secret key $\mathbf{sk}_m$ outputs the evaluation of function f on message m , i.e., $f(m)$, whereas the decryption algorithm (almost) always outputs 1 when given a normal ciphertext as input, irrespective of the secret key used. Thus, one could visualize the normal encryption algorithm as always encrypting a “canonical” *always-accepting* function.

Intuitively, security states that no attacker should be able to distinguish between two ciphertexts that decrypt to the same values under all the secret keys in the attacker’s possession. Now since there are two separate encryption algorithms, we have two different security properties. The first property says that secret key encryptions of two functions f_0 and f_1 should be indistinguishable if for every key in the attacker’s possession the output of f_0, f_1 is identical. We call this function the indistinguishability property. The second property says that it should be hard to distinguish between a normal (public key) encryption and secret key encryption of a function f , where $f(m)$ must be equal to 1 for all keys $\mathbf{sk}_m$ in the attacker’s possession. We call this the accept indistinguishability property.

We show that we can construct a PLBE scheme from a (key-policy) ABE scheme and a Mixed FE scheme. The idea is to encrypt a message using the ABE system with the attribute being set to be a Mixed FE ciphertext. Each user’s secret key will be an ABE private key. Here the ABE private key is generated for the Mixed FE decryption circuit in which a Mixed FE secret key, corresponding to the user’s index, is hardwired. The high level intuition is that when the attribute is a normal functional encryption ciphertext then all Mixed FE keys decrypt it to 1; thus any user with an appropriate ABE key could perform the decryption, whereas if the attribute is set to be a secret key ciphertext, then we can control the users who can decrypt it.

Formally, the scheme works as follows. During setup, the algorithm samples both ABE and Mixed FE key pairs $(\mathbf{abe.pp}, \mathbf{abe.msk})$, $(\mathbf{Mixed.pp}, \mathbf{Mixed.msk})$. To compute the i th user’s private key, it samples a Mixed FE secret key $\mathbf{Mixed.sk}_i$ for input i and also computes an ABE key $\mathbf{abe.sk}_i$ for predicate $\mathbf{Mixed.Dec}(\mathbf{Mixed.sk}_i, \cdot)$, i.e., Mixed FE decryption circuit with key $\mathbf{Mixed.sk}_i$ hardwired. And the ABE key $\mathbf{abe.sk}_i$ is set to be the i th user’s private key. Now to encrypt a message m , the algorithm simply runs the ABE encryption algorithm with attributes set to be a Mixed FE ciphertext $\mathbf{ct}_{\text{attr}}$.

For standard PLBE encryption, ct_{attr} is computed as a Mixed FE normal ciphertext, and for PLBE index encryption to some index i , ct_{attr} is computed as a Mixed FE secret key encryption of function *greater than i* . Lastly, the PLBE decryption is the same as the ABE decryption algorithm.

Correctness can be observed directly. For standard PLBE ciphertext, ct_{attr} is a normal FE ciphertext which decrypts to 1; thus the predicate $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$ is satisfied for all i . Therefore, by ABE correctness, the ABE decryption algorithm will output the message m . For the PLBE index i ciphertext, ct_{attr} is a Mixed FE secret key encryption of function “ $> i$ ” which decrypts to 1 for all keys Mixed.sk_j with $j > i$; thus the predicate is satisfied for all users with indices larger than i . Therefore, by ABE correctness, the ABE decryption algorithm will output the message m whenever $j > i$. For proving security, we rely on the fact that Mixed FE ciphertexts are indistinguishable to any adversary that does not have distinguishing secret keys. For instance, suppose there exists an adversary that can distinguish between PLBE normal encryptions and index 0 encryptions; then such an adversary can also be used to distinguish between Mixed FE normal ciphertexts and secret key ciphertexts encrypting function “ > 0 ” (note that this is an always-accepting function). Thus, such an attack can be used to break the accept indistinguishability property of the Mixed FE scheme. Similarly, we can prove index hiding and message hiding security of the construction by reducing to Mixed FE and ABE (selective) security, respectively. Now if the Mixed FE scheme is 1-query secure, then so is the PLBE scheme.

Now the size of ciphertexts has only poly-log dependence on the number of users n as required. Because each user can be uniquely identified using a bit string of length $\log n$, so the length of the attribute (Mixed FE ciphertext) will be polynomial in $\log n$, and thus the PLBE ciphertext, which is in turn an ABE ciphertext, will have length polynomial in $\log n$ as well. Also, note that to use the above transformation it is sufficient to construct a Mixed FE scheme that supports comparison operation on $\log n$ bit strings. In this work, we show how to construct a Mixed FE scheme for any class of polynomial-sized branching program from the LWE assumption.² Our construction relies only on the polynomial hardness of LWE, although we require a superpolynomial modulus-to-noise ratio. Since we already have circuit ABE schemes from the LWE assumption [GVW13, BGG+14], combining that with our Mixed FE construction gives us collusion resistant traitor tracing from the LWE assumption as well.

Looking back, it is easy to observe that Mixed FE for branching programs that support comparison functionality is sufficient for our application. However, as a design choice, here we instead chose to construct Mixed FE for general polynomial length branching programs as it is possible that this generalization leads to more applications in the future. Moreover, focusing on logarithmic length branching programs supporting comparisons, instead of general branching programs, did not lead to any significant simplification in the Mixed FE construction or its proof.

Part 3: An enhanced LWE toolkit. Before describing our LWE-based construction for Mixed FE, we define new “enhanced” properties for lattice trapdoors that will be useful in our work, and we believe it will find more applications in the future. In many LWE-based works, in addition to the LWE assumption itself, a critical tool has been the notion of lattice trapdoors [Ajt99, GPV08]. Lattice trapdoor samplers consist of a pair of algorithms, *TrapGen* and *SamplePre*. The trapdoor generation

²Note that this also gives us an alternate construction for selectively secure private key functional encryption with bounded collusions [SS10, GVW12].

algorithm **TrapGen** outputs a matrix $\mathbf{A}$ (that defines the lattice) and a trapdoor $T_{\mathbf{A}}$. The preimage sampling algorithm **SamplePre** takes as input a matrix $\mathbf{Z}$, a trapdoor for matrix $\mathbf{A}$, and a Gaussian parameter σ and outputs a matrix $\mathbf{U}$ such that $\mathbf{U}$ maps $\mathbf{A}$ to $\mathbf{Z}$ (that is, $\mathbf{A} \cdot \mathbf{U} = \mathbf{Z}$).³

These algorithms satisfy the following properties. First, the matrix $\mathbf{A}$ output by the trapdoor generation algorithm “looks like” a uniformly random matrix; we call this the *well-sampledness of matrix* property. Second, the matrix output by **SamplePre** is indistinguishable from a matrix drawn from a discrete Gaussian distribution with parameter σ over the set of all matrices $\mathbf{V}$ such that $\mathbf{A} \cdot \mathbf{V} = \mathbf{Z}$. In particular, if $\mathbf{Z}$ is chosen uniformly at random, then the output of **SamplePre** “looks like” a matrix $\mathbf{U}$ drawn from a discrete Gaussian distribution with parameter σ ; we call this the *preimage sampling property*. Lattice trapdoors with these properties have found a remarkable number of applications in building LWE-based cryptography.

In this work, we introduce two new *enhanced* properties for lattice trapdoors. The first property is the *row removal* property, which can be intuitively described as follows. Consider a setting where an adversary specifies some “target vectors,” and the challenger must output a matrix $\mathbf{A}$ and a matrix $\mathbf{U}$ such that $\mathbf{U}$ maps some of the rows of $\mathbf{A}$ to the target vectors, and maps the remaining rows to uniformly random vectors. Then these rows targeting uniformly random vectors can be removed from the trapdoor sampling. In particular, the challenger can sample a shorter matrix $\mathbf{B}$ with a trapdoor, extend $\mathbf{B}$ with uniformly random vectors to get $\mathbf{A}$, and set $\mathbf{U}$ to be a matrix that maps $\mathbf{B}$ to the target vectors. These two scenarios will be indistinguishable for the PPT adversary.

The second property is called the *target switching* property. In this setting, consider an adversary that specifies two matrices, $\mathbf{Z}_0, \mathbf{Z}_1$, and a set of “target” indices such that the rows of $\mathbf{Z}_0$ and $\mathbf{Z}_1$ agree on these target indices. The challenger is supposed to sample a matrix $\mathbf{A}$ with a trapdoor, compute a matrix $\mathbf{U}$ that maps $\mathbf{A}$ to $\mathbf{Z}_0$,⁴ and output $\mathbf{U}$ together with the rows of $\mathbf{A}$ corresponding to the target indices *and only those rows*. Then the challenger can switch the $\mathbf{U}$ to map $\mathbf{A}$ to $\mathbf{Z}_1$, and the target switching property requires that this change is indistinguishable to the adversary (note that this would not be possible if the adversary receives any of the nontarget rows of $\mathbf{A}$). Moreover, the adversary is allowed to adaptively query for different target vectors/indices in both these games.

Now that we have these enhanced properties, let us discuss how to construct lattice trapdoors with these enhanced properties (using standard lattice trapdoors). Our construction is similar to the **SampleLeft**/**SampleRight** algorithms of [ABB10, CHKP10]. The enhanced trapdoor generation algorithm uses the standard trapdoor sampling algorithm to sample two matrices, $\mathbf{A}_1, \mathbf{A}_2$, together with the respective trapdoors $T_{\mathbf{A}_1}, T_{\mathbf{A}_2}$. It outputs $\mathbf{A} = [\mathbf{A}_1 | \mathbf{A}_2]$ as the matrix and $T_{\mathbf{A}} = (T_{\mathbf{A}_1}, T_{\mathbf{A}_2})$ as the trapdoor. To sample a matrix $\mathbf{U}$ that maps $\mathbf{A}$ to $\mathbf{Z}$, the preimage sampling algorithm first chooses a uniformly random matrix $\mathbf{W}$ (of the same dimensions as $\mathbf{Z}$). It then uses $T_{\mathbf{A}_1}$ to compute a matrix $\mathbf{U}_1$ that maps $\mathbf{A}_1$ to $\mathbf{W}$, and uses $T_{\mathbf{A}_2}$ to compute a matrix $\mathbf{U}_2$ that maps $\mathbf{A}_2$ to $\mathbf{Z} - \mathbf{W}$. The final preimage matrix is set to be $\begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$. We use the matrix well-sampledness and preimage sampling properties of the standard lattice trapdoors to prove these enhanced properties; the detailed proof can be found in section 7.2.

³Although the notion of preimage sampling is usually defined with respect to (w.r.t.) vectors instead of matrices, here we stick to using matrices for technical reasons discussed later in section 7.

⁴Strictly speaking, we require $\mathbf{U}$ to map the target vectors of $\mathbf{A}$ to the target vectors of $\mathbf{Z}_0$, but the remaining vectors of $\mathbf{A}$ *approximately* map to the corresponding vectors of $\mathbf{Z}_0$.

Part 4: Constructing Mixed FE from LWE. Here we outline our Mixed FE construction for polynomial-sized (leveled) branching programs from the LWE assumption. The main ingredient of our construction is the “enhanced” lattice trapdoor sampling procedure $\text{LT}_{\text{en}} = (\text{EnTrapGen}, \text{EnSamplePre})$ discussed above.

First, let us recall the notion of leveled branching programs. A leveled branching program of length ℓ and width w can be represented using w states per level, 2ℓ state transition functions $\pi_{j,b}$ for each level $j \leq \ell$, 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 $\text{st} = 1$ of level 1. Suppose the branching program reads the first input bit (say, b) at level 1 (i.e., $\text{inp}(1) = 1$). Then the state of the program changes from st to $\pi_{1,b}(\text{st})$. Such a process is carried out (iteratively) until the program’s final state at level ℓ is computed. Depending upon the final state, the program either accepts or rejects.

For ease of exposition we will start with a simpler goal of constructing a 0-query secure Mixed FE scheme for a class of width- w read-once branching programs where each input bit is read once and in ascending order. Below we first outline a construction for such a 0-query system as it contains most of the central ideas, but is easier to digest. Later we discuss the modifications with which we can improve it to a secure 1-query scheme (and, more generally, q -query secure for any polynomial q) as well as expand the function class to arbitrary polynomial-sized branching programs.

Moving on to our 0-query Mixed FE construction, the master secret key consists of two sets of matrices and some trapdoor information. The first set, labeled as “randomization” matrices, consists of 4ℓ matrices $\{\mathbf{B}_{i,b}, \mathbf{C}_{i,b}\}_{i,b}$ for $i \in [\ell], b \in \{0, 1\}$. The second set, labeled as “program” matrices, consists of $w\ell$ matrices $\{\mathbf{P}_{i,v}\}_{i,v}$ for $i \in [\ell], v \in [w]$. Here the $\mathbf{C}_{i,b}$ matrices are sampled uniformly at random from $\mathbb{Z}_q^{n \times m}$, whereas the remaining (randomization and program matrices) are sampled jointly with common trapdoors (per level). Basically, for each level $i \in [\ell]$, we sample a $(w+2)n \times m$ matrix $\mathbf{M}_i$ as

$$(\mathbf{M}_i, T_i) \leftarrow \text{EnTrapGen}(1^{(w+2)n}, 1^m, q).$$

Now each $\mathbf{M}_i$ matrix is parsed as $w+2$ matrices of dimensions $n \times m$ stacked on top of each other, where the first two matrices are the randomization matrices and the remaining w matrices are the program matrices for the i th level. That is, for each i ,

$$\begin{bmatrix} \mathbf{B}_{i,0} \\ \mathbf{B}_{i,1} \\ \mathbf{P}_{i,1} \\ \vdots \\ \mathbf{P}_{i,w} \end{bmatrix} = \mathbf{M}_i.$$

All ℓ trapdoors $T_1, \dots, T_\ell$ are stored as the trapdoor information in the master secret key. The public parameters, on the other hand, only include the matrix dimensions, LWE modulus, and noise parameters, but none of these matrices or trapdoor information.

At a high level, the encryption and key generation algorithms will adhere to the following structure. To (secret key) encrypt a branching program, the trapdoors will be used to sample 2ℓ low norm matrices $\{\mathbf{U}_{i,b}\}_{i,b}$ (two per level) such that each matrix $\mathbf{U}_{i,b}$ encodes the corresponding state transition function by mapping/targeting level i “program” matrices to level $i+1$ “program” matrices per the transition function

$\pi_{i,b}$. Now the secret key for an input x will consist of $\ell + 1$ key vectors $\{\mathbf{t}_i\}_i$. The first key component, $\mathbf{t}_1$, will contain the program matrix $\mathbf{P}_{1,1}$ (which represents the starting state) plus some randomization component generated using the level 1 randomization matrix $\mathbf{B}_{1,b}$. The remaining ℓ key vectors will have two components—the first component will cancel the previous randomization component, and the second component will add new randomization terms.⁵ The idea is that if decryption is performed honestly, then all the randomization terms will get canceled and the final output will reflect the output of the branching program.

So this way the program matrices will be tied in such a manner that they encode the state transition information, and they can be used to perform the branching program execution. And the randomization matrices are added to make sure that (1) the computation is hidden at each step, and (2) if ciphertext matrices and key vectors are combined in any inadmissible way, then the randomization components do not get canceled. Let us now look at how to execute the above ideas.

Key generation. The key generation algorithm takes as input a string x and generates key vectors $\{\mathbf{t}_i\}_i$ as follows. It chooses ℓ uniform secret vectors $\mathbf{s}_i \in \mathbb{Z}_q^n$ for $i \in [\ell]$ and $\ell + 1$ noise vectors $\mathbf{e}_i \in \mathbb{Z}_q^m$ for $i \in [\ell + 1]$. It also chooses a *short* secret vector $\tilde{\mathbf{s}} \in \mathbb{Z}_q^n$ and sets key vectors as follows:

$$\forall i \in [\ell + 1], \quad \mathbf{t}_i = \begin{cases} \mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} + \mathbf{e}_1 & \text{if } i = 1, \\ -\mathbf{s}_{i-1} \cdot \mathbf{C}_{i-1,x_{i-1}} + \mathbf{s}_i \cdot \mathbf{B}_{i,x_i} + \mathbf{e}_i & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell \cdot \mathbf{C}_{\ell,x_\ell} + \mathbf{e}_{\ell+1} & \text{if } i = \ell + 1. \end{cases}$$

In words, the randomization component (likewise, cancellation component) added in the i th key vector ($(i + 1)$ th key vector) is an LWE sample where the public matrix used depends on the i th bit of input x . Looking ahead, choosing the “randomization” matrices depending on the string x would assert that the ciphertext matrices cannot be arbitrarily combined to learn meaningful terms.

Normal encryption. The normal (public key) encryption algorithm simply samples 2ℓ random short matrices $\{\mathbf{U}_{i,b}\}_{i,b}$ as $\mathbf{U}_{i,b} \leftarrow \chi^{m \times m}$, where χ is the noise distribution chosen during setup.

Secret key encryption. Moving on to the secret key encryption algorithm, on input the master secret key and a branching program $\mathbf{BP} = (\{\pi_{i,b}\}_{i,b}, \text{acc}, \text{rej})$, it samples low norm matrices $\{\mathbf{U}_{i,b}\}$ as follows. It first chooses two “program” matrices for the last level $\ell + 1$ as $\mathbf{P}_{\ell+1,\text{rej}} = \mathbf{0}^{n \times m}$ and $\mathbf{P}_{\ell+1,\text{acc}} \leftarrow \mathbb{Z}_q^{n \times m}$. That is, for the accepting state, it chooses a random program matrix, and for the rejecting state it sets the matrix to be all zeros. Next, using the i th trapdoor T_i (included in the master secret key) it runs the `EnSamplePre` algorithm to sample the ciphertext (transition) matrices $\mathbf{U}_{i,0}, \mathbf{U}_{i,1}$ such that they map/target matrix $\mathbf{M}_i$ as follows:

$$\begin{bmatrix} \mathbf{B}_{i,0} \\ \mathbf{B}_{i,1} \\ \mathbf{P}_{i,1} \\ \vdots \\ \mathbf{P}_{i,w} \end{bmatrix} \xrightarrow{\mathbf{U}_{i,0}} \begin{bmatrix} \mathbf{C}_{i,0} \\ \$ \\ \mathbf{P}_{i+1,\pi_{i,0}(1)} \\ \vdots \\ \mathbf{P}_{i+1,\pi_{i,0}(w)} \end{bmatrix}, \quad \begin{bmatrix} \mathbf{B}_{i,0} \\ \mathbf{B}_{i,1} \\ \mathbf{P}_{i,1} \\ \vdots \\ \mathbf{P}_{i,w} \end{bmatrix} \xrightarrow{\mathbf{U}_{i,1}} \begin{bmatrix} \$ \\ \mathbf{C}_{i,1} \\ \mathbf{P}_{i+1,\pi_{i,1}(1)} \\ \vdots \\ \mathbf{P}_{i+1,\pi_{i,1}(w)} \end{bmatrix}.$$

⁵Technically, the last key vector will only remove the previous randomization component. It doesn't add a new randomization term.

Here we use “\$” to denote a uniformly random $n \times m$ matrix of appropriate dimension. In words, the structure we enforce here is that the matrix $\mathbf{U}_{i,b}$ targets the $\mathbf{B}_{i,b}$ randomization matrix to its $\mathbf{C}_{i,b}$ counterpart, and the $\mathbf{B}_{i,1-b}$ randomization matrix to a random matrix. Additionally, $\mathbf{U}_{i,b}$ encodes the information about transition function $\pi_{i,b}$ by targeting the level i program matrices to their level $i+1$ counterparts per $\pi_{i,b}$. Thus, from the perspective of both correctness and security, this guarantees that a key vector $\mathbf{t}_i$ for some input x must be combined with ciphertext component $\mathbf{U}_{i,x_i}$, as otherwise the randomization matrix would be mapped to a random matrix, thereby destroying the underlying structure.

Decryption. First, let us focus on decrypting a secret key encryption of branching program BP using a secret key $\{\mathbf{t}_i\}_i$ corresponding to an input x . Intuitively, one could visualize the matrices $\{\mathbf{U}_{i,0}, \mathbf{U}_{i,1}\}_i$ in the ciphertext as “encodings” of the branching program state transition functions $\pi_{i,0}, \pi_{i,1}$, respectively. Therefore, decrypting the ciphertext using a secret key for some input x will be analogous to evaluating the branching programs BP on input x directly. Recall that we assumed (for ease of exposition) that the branching programs are read-once, and input bits are read sequentially in ascending order. Thus, the first input bit x_1 is read at level 1. Then evaluation of BP at level 1 would map the state $\mathbf{st}_1 = 1$ at level 1 to state $\mathbf{st}_2 = \pi_{1,x_1}(1)$ at level 2. Analogously, the decryptor can compute

$$\begin{aligned} \mathbf{t}_1 \cdot \mathbf{U}_{1,x_1} + \mathbf{t}_2 &\approx (\mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}) \cdot \mathbf{U}_{1,x_1} + \mathbf{t}_2 \\ &\approx \mathbf{s}_1 \cdot \mathbf{C}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\mathbf{st}_2} + \mathbf{t}_2 \\ &\approx \cancel{\mathbf{s}_1 \cdot \mathbf{C}_{1,x_1}} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\mathbf{st}_2} + (\cancel{-\mathbf{s}_1 \cdot \mathbf{C}_{1,x_1}} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2}) \\ &\approx \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\mathbf{st}_2}. \end{aligned}$$

In general, if the program state at level i during execution is $\mathbf{st}_i$, then the decryptor will accumulate the term of the form $\mathbf{s}_i \cdot \mathbf{B}_{i,x_i} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i,\mathbf{st}_i}$ by successively summing and multiplying secret key and ciphertext components as

$$(\cdots((\mathbf{t}_1 \cdot \mathbf{U}_{1,x_1} + \mathbf{t}_2) \cdot \mathbf{U}_{2,x_2} + \mathbf{t}_3) \cdots + \mathbf{t}_i).$$

This can be verified as follows. We know that the bit read at level i is x_i , and thus the new state at level $i+1$ will be $\mathbf{st}_{i+1} = \pi_{i,x_i}(\mathbf{st}_i)$. Now the accumulated sum-product during decryption will be

$$\begin{aligned} (\mathbf{s}_i \cdot \mathbf{B}_{i,x_i} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i,\mathbf{st}_i}) \cdot \mathbf{U}_{i,x_i} + \mathbf{t}_{i+1} \\ \approx \cancel{\mathbf{s}_i \cdot \mathbf{C}_{i,x_i}} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i+1,\mathbf{st}_{i+1}} + (\cancel{-\mathbf{s}_i \cdot \mathbf{C}_{i,x_i}} + \mathbf{s}_{i+1} \cdot \mathbf{B}_{i+1,x_{i+1}}) \\ \approx \mathbf{s}_{i+1} \cdot \mathbf{B}_{i+1,x_{i+1}} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i+1,\mathbf{st}_{i+1}}. \end{aligned}$$

Therefore, the invariant is maintained. Continuing in this way, the decryptor can iteratively compute the sum-product combining all key and ciphertext components. Note that (by definition) adding in the $(\ell+1)$ th key component $\mathbf{t}_\ell$ does not introduce a term like $\mathbf{s}_{\ell+1} \cdot \mathbf{B}_{\ell+1,x_{\ell+1}}$ to the sum-product; thus the accumulated term at the top will be $\approx \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\mathbf{st}_{\ell+1}}$, where $\mathbf{st}_{\ell+1}$ is either acc or rej depending on $\text{BP}(x)$. Finally, the decryptor simply checks whether the norm of the final sum-product term is small or not. Recall that the program matrix for the last level corresponding to the rejecting state is set to be all zeros, i.e., $\mathbf{P}_{\ell+1,\text{rej}} = \mathbf{0}^{n \times m}$. Therefore, if $\text{BP}(x) = 0$, then the norm of the final sum-product term will be small, which the decryptor can test and output 0. Otherwise, with high probability the final sum-product term will be large and it outputs 1.

By the above analysis, correctness follows in the case where ciphertext is a secret key encryption. The correctness of decryption when the ciphertext is a normal (public key) ciphertext follows from the fact that the ciphertext matrices $\{\mathbf{U}_{i,0}, \mathbf{U}_{i,1}\}_i$ are independently sampled random short matrices.

0-query security. To prove 0-query security of our construction, we need to argue that it satisfies both function indistinguishability as well as accept indistinguishability properties. We start by proving the function indistinguishability security. Recall that in the 0-query function indistinguishability security game, an adversary submits two branching programs, $\text{BP}^{(0)}, \text{BP}^{(1)}$, and is allowed to make a polynomial number of key queries such that for each queried input x , $\text{BP}^{(0)}(x) = \text{BP}^{(1)}(x)$ (i.e., every secret key given out has same output on both the challenge programs). The adversary receives secret key encryption of either $\text{BP}^{(0)}$ or $\text{BP}^{(1)}$, and its goal is to distinguish between them.

Although the full security proof is technically involved, the main ideas behind our proof are very intuitive. Before diving into the proof structure, we point out that the construction described above has to be slightly modified for proving security. Below we describe our proof ideas as well as discuss the modifications required along the way.

At a high level, our idea is to “hardwire” the output of the challenge branching programs in every secret key given to the adversary. Note that the security definition states that both challenge programs must evaluate to the same value on all queried inputs, and thus we only need to hardwire a single value in each key. For ease of exposition, assume that the adversary makes exactly one secret key query. (In the general case of polynomially many key queries, the proof proceeds by hardwiring the level 1 components in all secret keys, followed by level 2 hardwiring, and so on.) Let $\{\mathbf{U}_{i,b}\}_{i,b}$ be the challenge ciphertext, and let $\{\mathbf{t}_i\}_i$ be the secret key computed by the challenger. Our hardwiring strategy works as follows. We start by rewriting the second secret key vector $\mathbf{t}_2$ in terms of $\mathbf{t}_1$ as follows:

$$\begin{aligned}\mathbf{t}_2 &= -\mathbf{s}_1 \cdot \mathbf{C}_{1,x_1} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2 \\ &= -\mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} \cdot \mathbf{U}_{1,x_1} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2 \\ &= -\mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} \cdot \mathbf{U}_{1,x_1} - \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2 \\ &= -(\mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}) \cdot \mathbf{U}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2 \\ &= -(\mathbf{t}_1 - \mathbf{e}_1) \cdot \mathbf{U}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2.\end{aligned}$$

Here st_2 is the state of the challenge branching program encrypted (after one step is executed). Now in the above term, we can *smudge* the term $\mathbf{e}_1 \cdot \mathbf{U}_{1,x_1}$ by appropriately choosing the noise distributions, i.e., $\mathbf{e}_2 \gg \mathbf{e}_1 \cdot \mathbf{U}_{1,x_1}$.⁶ (Note that since we require smudging here, thus the LWE modulus q needs to be superpolynomial in the lattice dimension.) Thus, the second key component can be indistinguishably computed as follows without requiring any explicit knowledge of the $\mathbf{C}_{1,x_1}$ matrix:

$$\begin{aligned}\mathbf{t}_1 &= \mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} + \mathbf{e}_1, \\ \mathbf{t}_2 &= \begin{aligned} &-(\mathbf{t}_1 - \mathbf{e}_1) \cdot \mathbf{U}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} \\ &+ \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2 \end{aligned} \xrightarrow{\text{Smudging}} \mathbf{t}_2 = \begin{aligned} &-\mathbf{t}_1 \cdot \mathbf{U}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\text{st}_2} \\ &+ \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2. \end{aligned}\end{aligned}$$

Next, we use the row removal property of our enhanced trapdoor sampling algorithms to remove the $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ rows from the first matrix $\mathbf{M}_1$ and instead sample

⁶If we keep on smudging this way, our noise distributions will have to grow by an exponential factor at each step. In the main body, we show how to avoid this by a better smudging argument.

these randomly. To understand why this can be done, recall that in the actual construction the encryptor needs the ability to create a ciphertext for *any* branching program that could be chosen even after all the keys have been distributed. That is, the encryptor must be able to sample matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$ such that they map level 1 program matrices $\{\mathbf{P}_{1,v}\}_v$ to level 2 program matrices $\{\mathbf{P}_{2,v}\}_v$ per some transition functions $\{\pi_{1,b}\}_b$, as well as ensure that $\mathbf{B}_{1,b} \cdot \mathbf{U}_{1,b} = \mathbf{C}_{1,b}$. Now since the keys contain the matrices $\mathbf{C}_{1,0}, \mathbf{C}_{1,1}$ and they could be given out even before matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$ are sampled, thus matrices $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ must be sampled together with $\{\mathbf{P}_{1,v}\}_v$ such that they share a common trapdoor.

However, at this stage in the proof the challenge branching program is (selectively) fixed ahead of any secret key queries. Therefore, in this context we can sample matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$ to only map level 1 program matrices to their level 2 counterparts, and simply set the matrices $\mathbf{C}_{1,b}$ as $\mathbf{C}_{1,b} = \mathbf{B}_{1,b} \cdot \mathbf{U}_{1,b}$ and use these to compute the secret keys. We would like to point out that in order to perform this row removal securely, it is important that $\mathbf{B}_{1,b} \cdot \mathbf{U}_{1,1-b} = \$$, that is, matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$ map both matrices $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ to random and uncorrelated matrices.

Now once we have removed the $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ rows from the first matrix, we use the LWE assumption to switch the first key component $\mathbf{t}_1$ to a random vector. Note that at this point since matrices $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ are sampled uniformly (i.e., are no longer sampled with trapdoor information) and secret vector $\mathbf{s}_1$ is not used in computing the second key component $\mathbf{t}_2$, thus we can apply LWE to switch $\mathbf{t}_1$ to random, where the LWE secret is $\mathbf{s}_1$ and the LWE public matrix will be $\mathbf{B}_{1,x_1}$.⁷ Concretely, using LWE we can perform the following switch, which essentially erases the information about the level 1 program matrix $\mathbf{P}_{1,1}$ from the secret keys, thereby rendering the program evaluation to start from level 2 and state $\mathbf{st}_2$ instead:

$$\begin{aligned} \mathbf{t}_1 &= \mathbf{s}_1 \cdot \mathbf{B}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} + \mathbf{e}_1 \xrightarrow{\text{LWE}} \mathbf{t}_1 = \$, \\ \mathbf{t}_2 &= -\mathbf{t}_1 \cdot \mathbf{U}_{1,x_1} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{2,\mathbf{st}_2} + \mathbf{s}_2 \cdot \mathbf{B}_{2,x_2} + \mathbf{e}_2. \end{aligned}$$

Now iteratively performing this *hardwiring* strategy (ℓ times), we end up switching all but the last key components to be random vectors. Also, the last key component will contain the final program matrix, which is either a random matrix or a zero matrix, depending on the program output. Thus, the key vectors contain no information about the “program” matrices chosen during setup. At this point, the challenge matrices $\{\mathbf{U}_{i,b}\}_{i,b}$ still contain the information about the branching program encrypted in the form of mapping between level i and $i+1$ “program” matrices, i.e., the state transition functions $\{\pi_{i,b}\}_{i,b}$. Finally, to argue indistinguishability here (i.e., between the challenge matrices) we use the target switching property of our enhanced trapdoors. We apply a bottom-up approach to execute this change. First, note that the level 1 program matrices do not explicitly appear anywhere, except that they are used to sample the level 1 ciphertext matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$. Thus, we can use the target switching property to switch the targets of matrices $\mathbf{U}_{1,0}, \mathbf{U}_{1,1}$. Observe that this now removes the information about level 2 program matrices as well as the level 1 transition functions of challenge branching programs. Next, by the same principle, we can perform the same target switching step for $\mathbf{U}_{2,0}, \mathbf{U}_{2,1}$ and continue so on. If we keep on performing the *target switching* step this way until the top, then the challenge

⁷In the general case of multiple key queries the LWE public matrices will be both $\mathbf{B}_{1,0}, \mathbf{B}_{1,1}$ and the LWE secret will consist of all the secret key vectors $\mathbf{s}_i$ that are chosen independently and *per key*.

ciphertext will contain no information about the challenge programs (i.e., their state transition functions), thereby completing our claim of function indistinguishability.⁸ This completes the first proof.

The proof of the 0-query accept indistinguishability security of our construction is similar, but more technical due to the fact that we need to argue that the challenge ciphertext is indistinguishable from random short matrices.

However, for our PLBE construction, the Mixed FE scheme must handle one ciphertext query as well, and it is not clear how to prove the above construction to be 1-query secure directly. The bottleneck is the fact that in the above proof strategy we *hardwire* all secret key components to match the output of the challenge program. Now if the adversary is allowed to make a secret key encryption query, then it is not clear how the challenger would still program the secret key vectors. To get around this problem, we expand our system such that it consists of λ *pairs* of 0-query subsystems. Very briefly, during encryption, the algorithm now also samples a λ -bit string **tag** randomly and depending on each bit of **tag**, it chooses one subsystem in each pair and runs the 0-query encryption for that subsystem. Now during key generation, it (linearly) secret shares the starting program across these λ pairs of subsystems such that the same secret share is used for both subsystems in each pair. Then it runs the 0-query key generation algorithm for all these 2λ subsystems with their corresponding secret shares as the starting program matrices. For decryption, the subsystems chosen during encryption are combined with their counterparts in the secret keys, and 0-query decryption is performed in these subsystems along with a (linear) reconstruction on top of the output. More details are provided in the main body.⁹

This completes the technical overview of our construction.

Relation to recent LWE-based schemes. There have been several recent works that have advanced the state of the art in computing branching programs on encrypted data with the goal of reducing security to LWE or LWE-like assumptions [GGH15, BVWW16, BV15, GKW17b, CC17, GKW17a, WZ17]. While our construction above benefits from that lineage, we wish to briefly call out a few important distinctions.

First, from a purely mechanical perspective, the construction of our Mixed FE scheme is structurally very different from the constructions of the aforementioned primitives. Very briefly, in all previous constructions the evaluator multiplies a set of matrices and sums them up to get the final output, whereas in our construction, we do not use this “one-shot” approach for evaluation. Instead, we multiply a component from the ciphertext with a secret key component, then add in another secret key component, multiply this sum with another ciphertext component, and so on. Thus, our mechanism of combining the secret key and ciphertext components is much different than what was used in prior works.¹⁰

⁸Technically, we cannot apply the target switching property here, because the target switching property only guarantees that targets being switched are approximately mapped, whereas here we target exactly. Therefore, we also need to add some noise in the targeted “program” matrices before running `EnSamplePre` algorithm. For simplicity, we avoid this modification.

⁹We would like to point out that the above idea could also be used to improve the Mixed FE construction to be q -query secure for any polynomial q . The idea will be to sample tag strings **tag** from a larger alphabet instead of $\{0, 1\}$. However, we only focus on 1-query security, as it is sufficient for our result.

¹⁰Although one can always express such a nested matrix multiplication and addition mechanism using only a sequence of matrix multiplications with much larger (and repetitive) matrices, we point out that the underlying structure of such matrices as well as the modified evaluation algorithm will still be much different from those used in previous works.

Second, we structure our proof of security to hardwire the outputs for keys one level at a time until we hit the final output level in which we have the final outputs hardwired but lost information about the program that got us there. In this sense at a high level this leveled programming proof structure much more closely resembles that of garbled circuit proofs. Thus, we need to develop new lower LWE-specific techniques to match these goals. In contrast, works such as [BVWW16, GKW17b, GKW17a, WZ17] have a different aim of losing all meaningful information when a secret is not known.

1.2. Some future directions. Our construction of Mixed FE relied on the LWE assumption and leveraged certain algebraic properties in that setting. An intriguing question is whether there are other avenues for achieving Mixed FE. A natural path is to build Mixed FE with a garbled circuits [Yao86] backbone. If one starts with the bounded key functional encryption scheme of Gorbunov, Vaikuntanathan and Wee (GVW) [GVW12] and flips [AGVW13, BS15, KMUW17] the semantics of message and function, one can get a secret key functional encryption scheme that is secure for an unbounded number of private keys and bounded number of ciphertexts. To make it a Mixed FE system we would somehow connect a public key mode of encryption to the scheme. One possible path is to use a “blinded” [BLSV17] form of garbled circuits as the underlying 1-bounded scheme in the GVW transformation. (Building on [DG17a, DG17b], blinded garbled circuits were recently used to give anonymous identity-based encryption (IBE) from new assumptions). It seems possible that this approach could lead to a scheme with the accept indistinguishability property if no encryption oracle queries are allowed. However, there appears to be technical difficulties in making a public key-generated ciphertext indistinguishable from a master secret key-generated ciphertext when the attacker gets oracle queries. That being said, we believe that a garbled circuit approach remains a plausible future direction.

We remark that even if a garbled circuit approach becomes possible, the requirement for an ABE scheme supporting circuits will still indirectly require the LWE assumption given the state of the art. In addition, we expect that our LWE toolkit and underlying construction ideas will have future value in any case.

A second interesting direction is whether there are other applications that can leverage a functional encryption system that has a bimodal encryption where the public key and master secret key support different spaces of messages or functions. In our Mixed FE system the public key only supported the always accept function, but there could conceivably be other variants of interest.

Finally, a natural open question is to construct traitor tracing schemes with public traceability from LWE. Currently, it is unclear if achieving public tracing is an easier task than building general public key functional encryption.

1.3. Additional related work. Our transformation from Mixed FE to PLBE using ABE has some high-level similarities to the predicate encryption scheme of Gorbunov, Vaikuntanathan, and Wee [GVW15] and the single-key functional encryption scheme of Goldwasser et al. [GKP+13]. In both these cases, we have specialized encryption schemes whose ciphertext serves as an attribute for the ABE scheme.

Concurrent/follow-up works. Concurrent to our work, Chen, Vaikuntanathan, and Wee [CVW18b] also proposed a simplified variant of our row-removal and target switching properties. They used these properties for constructing lockable obfuscation schemes for nonpermutation branching programs. However, their properties are weaker than the ones we define/construct in this work. In particular, their construc-

tions do not allow the adversary to make any preimage queries. Following our work, Chen et al. [CVW+18a] presented simpler constructions of Mixed FE.

Connections to differential privacy. Dwork et al. [DNR+09] showed that existence of collusion resistant traitor tracing schemes implies hardness results for efficient differentially private [DMNS06] data sanitization. In particular, they showed that if there exists a traitor tracing scheme with ciphertexts of size $s(\lambda, n)$, then there exist a database of size n and a query class $\mathcal{Q}$ of size $2^{s(\lambda, n)}$ such that it is hard to sanitize the database D for query class $\mathcal{Q}$ in a differentially private manner. Combining our LWE-based construction with the result of Dwork et al., we get an LWE-based hardness result for differentially private sanitization with query space of size $2^{\text{poly}(\lambda, \log n)}$. We note that Goyal et al. [GKRW17] and Kowalczyk et al. [KMUW17] recently achieved better differential privacy impossibility results from the security of bilinear map assumptions and one-way functions, respectively.

Weaker notions of traitor tracing. Since the notion of traitor tracing (TT) was first proposed [CFN94], several relaxed variants have been studied in order to achieve short ciphertexts. The first natural relaxation is the bounded collusion setting, where we have an a priori bound k that is fixed during setup, and security is guaranteed only if the adversary gets at most k secret keys. Collusion bounded systems can either be constructed via combinatorial tools [CFN94, SW98, CFNP00, SSW01, PST06, BP08] or be algebraic and constructed under different cryptographic assumptions such as DDH [KD98, BF99, KY02a, KY02b], bilinear DDH [CPP05, ADM+07, FNP07], and LWE [LPSS14]. Recently, Agrawal et al. [ABP+17] showed a transformation from inner product functional encryption to collusion-bounded TT, resulting in algebraic constructions based on various assumptions such as DCR, DDH, and LWE. In all the above works, the size of the ciphertext grows with the collusion bound.

The second relaxation is called threshold TT, introduced by [NP98, CFNP00]. In a threshold TT scheme, a threshold $\delta \in [0, 1]$ is chosen during setup, and the traceability guarantee only holds if the decoder box works with probability at least δ . Boneh and Naor [BN08] showed a threshold TT scheme where the ciphertexts have size $O(\lambda)$ and the secret keys have size $O(n^2 \lambda / \delta^2)$. While this scheme achieves collusion resistance, the system must be configured with a specific δ value, and once it is set, one will not necessarily be able to identify a traitor from a box D that works with smaller probability. In practice, it can be tricky to ascertain what threshold will actually be okay. This is because the encrypted messages could have redundancy, so even a decoder box with a small fraction of success might allow an attacker to learn the underlying message.

Finally, in a recent work, Goyal et al. [GKRW17] introduced a new relaxation called *risky traitor tracing*. In this notion, the scheme is fully collusion resistant (and does not have the threshold restriction as above). Instead, the probability of tracing a traitor, given a successful decoding box, can be substantially smaller than 1. For instance, [GKRW17] showed a bilinear maps-based construction where the ciphertext size grows as $\lambda \cdot k$, but the trace algorithm has only a k/n chance of catching a traitor. The authors show that this weaker notion is actually enough to achieve strong hardness results for differential privacy [DMNS06, DNR+09] and also argue that in a certain “continuous use” setting, the probability of tracing can be amplified back up to one. However, in general settings, the Goyal et al. tradeoff between the probability of catching a traitor and the size of ciphertexts might be undesirable.

1.4. Organization. In section 2, we present the preliminaries required for this work. Next, in section 3, we have the traitor tracing and PLBE definitions. This includes the new decoder-based and q -query PLBE definitions. In section 4, we show how 1-query PLBE implies decoder-based PLBE, and how decoder-based PLBE suffices for constructing TT schemes. Therefore, the problem of constructing a TT scheme reduces to the problem of constructing a 1-query PLBE scheme. For this, we introduce a new primitive called Mixed FE in section 5. In section 6, we show how to construct 1-query PLBE using Mixed FE and ABE (the syntax and security definitions of ABE can be found in Appendix A). Finally, in section 8, we present our Mixed FE construction (before presenting the Mixed FE construction, in section 7 we present new lattice tools which are required for our construction).

2. Preliminaries.

Notation. 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., $\mathbf{A}$), and we assume all vectors are row vectors. The j th row of a matrix $\mathbf{A}$ is denoted by $\mathbf{A}[j]$. 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 up to n as $[n] := \{1, \dots, n\}$. Throughout this paper, unless specified, all polynomials we consider are positive polynomials. Also, we represent each finite set on integers $S \subset \mathbb{N}$ as an *ordered* set $S = \{i_1, i_2, \dots, i_n\}$, i.e., $i_j < i_k$ for every $1 \leq j < k \leq n$. 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 $\mathcal{D}$.

For two distributions X, Y , over a finite domain Ω , the statistical distance between X and Y is defined as $\text{SD}(X, Y) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{\omega \in \Omega} |X(\omega) - Y(\omega)|$. A family of distributions $\mathcal{D}_1 = \{\mathcal{D}_1(\lambda)\}_\lambda$ and $\mathcal{D}_2 = \{\mathcal{D}_2(\lambda)\}_\lambda$, parameterized by security parameter λ , are said to be statistically indistinguishable, represented by $\mathcal{D}_1 \approx_s \mathcal{D}_2$, if there exists a negligible function $\text{negl}(\cdot)$ such that, for all $\lambda \in \mathbb{N}$, $\text{SD}(\mathcal{D}_1(\lambda), \mathcal{D}_2(\lambda)) \leq \text{negl}(\lambda)$. For a family of distributions $\mathcal{D} = \{\mathcal{D}(\lambda)\}_\lambda$ over the integers, and integer bounds $B = \{B(\lambda)\}_\lambda$, we say that $\mathcal{D}$ is B -bounded if $\Pr[|x| \leq B(\lambda) : x \leftarrow \mathcal{D}(\lambda)] = 1$. In words, a B -bounded distribution is supported only on the range $[-B, B]$. Below we state the “smudging” lemma as it appears in prior works.

LEMMA 2.1 (Smudging Lemma [AJW11, Lemma 2.1, paraphrased]). *Let B_1, B_2 be two polynomials over the integers, and let $\mathcal{D} = \{\mathcal{D}(\lambda)\}_\lambda$ be any B_1 -bounded distribution family. Let $U = \{U(\lambda)\}_\lambda$ and $U(\lambda)$ denote the uniform distribution over integers $[-B_2(\lambda), B_2(\lambda)]$. The family of distributions $\mathcal{D}$ and U is statistically indistinguishable, $\mathcal{D} + U \approx_s U$, if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $B_1(\lambda)/B_2(\lambda) \leq \text{negl}(\lambda)$.*

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 $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x}^T = \mathbf{0}^T \pmod{q}\}$. For $\mathbf{u} \in \mathbb{Z}_q^n$, we let $\Lambda_q^\mathbf{u}(\mathbf{A})$ denote the coset $\{\mathbf{x} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{x}^T = \mathbf{u}^T \pmod{q}\}$.

Discrete Gaussians. Let σ be any positive real number. The Gaussian distribution $\mathcal{D}_\sigma$ 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 2.2. *Let m, n, q be positive integers with $m > n$, $q \geq 2$, and $\sigma = \tilde{\Omega}(n)$. Then*

$$\Pr[\|\mathbf{x}\| > \sqrt{m} \cdot \sigma : \mathbf{x} \leftarrow \mathcal{D}_{\mathcal{L}, \sigma} : \mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}, \mathcal{L} = \Lambda_q^\perp(\mathbf{A})] \leq \text{negl}(n).$$

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 the same as the discrete Gaussian distribution $\mathcal{D}_{\mathbb{Z}^m, \sigma}$ 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, by the above lemma we get that $\tilde{\mathcal{D}}_{\mathbb{Z}^m, \sigma} \approx_s \mathcal{D}_{\mathbb{Z}^m, \sigma}$.

2.1.1. Learning with errors. The learning with errors (LWE) problem was introduced by Regev [Reg05], who showed that solving LWE on *average* is as hard as quantumly solving several standard lattice-based problems in the worst case. The LWE assumption states that no polynomial time adversary can distinguish between the following oracles. In one case, the oracle chooses a uniformly random secret $\mathbf{s}$, and for each query, it chooses a vector $\mathbf{a}$ uniformly at random, scalar e from a noise distribution, and outputs $(\mathbf{a}, \mathbf{s} \cdot \mathbf{a}^T + e)$. In the second case, the oracle simply outputs a uniformly random vector $\mathbf{a}$ together with a uniformly random scalar u . Regev showed that if there exists a polynomial time adversary that can break the LWE assumption, then there exists a polynomial time quantum algorithm that can solve some hard lattice problems in the worst case.

Several works also explored different variants of the LWE assumption, where the secret vector $\mathbf{s}$, public vectors $\mathbf{a}$, and noise are drawn from different distributions. In this work, we will be using two of these variants. First, we will be using the LWE version with *short secrets* (also known as the *normal form*), introduced by Applebaum et al. [ACPS09]. In this variant, the secret vector $\mathbf{s}$ is also drawn from the noise distribution. Applebaum et al. showed that this version is as hard as the LWE problem if the modulus is p^e for some prime p and integer e . This was later generalized to all moduli by Brakerski et al. [BLP+13]. The second variant, which was proposed by Boneh et al. [BLMR13], allows the public vectors $\mathbf{a}$ to be chosen from the noise distribution as well. Boneh et al. showed that this version of LWE is as hard as standard LWE.

We will first present the LWE assumption in a general framework,¹¹ which captures the standard LWE, LWE with short secrets, and LWE with short public vectors. In this framework, we will have an explicit security parameter λ , and the other parameters are allowed to grow as a function of the security parameter.

DEFINITION 2.3 (generalized learning with errors). *Fix any polynomial $n(\cdot)$, function $q(\cdot)$, secret distribution $\eta(\cdot)$, public vector distribution $\phi(\cdot)$, and noise distribution*

¹¹Canetti and Chen [CC17] proposed the general LWE problem. However, their version requires the public vectors to be sampled from a uniform distribution, whereas we require the public vectors to be sampled from nonuniform distributions. Also, it is possible to generalize our version further. Here, we present the minimal generalization that suffices for our work.

$\chi(\cdot)$, where $n : \mathbb{N} \rightarrow \mathbb{N}$, $q : \mathbb{N} \rightarrow \mathbb{N}$ and for each $\lambda \in \mathbb{N}$, $\eta(\lambda)$ and $\phi(\lambda)$ are distributions over $\mathbb{Z}_{q(\lambda)}^{n(\lambda)}$ and $\chi(\lambda)$ is a distribution over $\mathbb{Z}$. We say that the generalized LWE assumption $\text{GLWE}_{n,q,\eta,\phi,\chi}$ holds if for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q = q(\lambda)$, $n = n(\lambda)$, $\eta = \eta(\lambda)$, $\phi = \phi(\lambda)$, and $\chi = \chi(\lambda)$, $\text{Adv}_{\text{GLWE},\mathcal{A}}^{n,q,\eta,\phi,\chi}(\lambda) \leq \text{negl}(\lambda)$, where

$$\text{Adv}_{\text{GLWE},\mathcal{A}}^{n,q,\eta,\phi,\chi}(\lambda) = \Pr \left[\mathcal{A}^{O_1^s}(\lambda) = 1 : \mathbf{s} \leftarrow \eta \right] - \Pr \left[\mathcal{A}^{O_2}(\lambda) = 1 \right],$$

and oracles $O_1^s(\cdot)$, $O_2(\cdot)$ are defined as follows: oracle $O_1^s(\cdot)$ has $\mathbf{s} \in \mathbb{Z}_q^n$ hardwired, and on each query it chooses $\mathbf{a} \leftarrow \phi$, $e \leftarrow \chi$ and outputs $(\mathbf{a}, \mathbf{s} \cdot \mathbf{a}^T + e \bmod q)$, and oracle $O_2(\cdot)$ (on each query) chooses $\mathbf{a} \leftarrow \phi$, $u \leftarrow \mathbb{Z}_q$ and outputs $(\mathbf{a}, u)$.

We now present different variants of the LWE assumption and discuss the parameters for which they are believed to be secure.

ASSUMPTION 1 (learning with errors). Let $n : \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial, and let $q : \mathbb{N} \rightarrow \mathbb{N}$, $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ be functions. The $\text{LWE}_{n,q,\sigma}$ assumption states that $\text{GLWE}_{n,q,\eta,\phi,\chi}$ holds, where $\eta(\lambda), \phi(\lambda)$ are uniform distributions over $\mathbb{Z}_{q(\lambda)}^{n(\lambda)}$, and $\chi(\lambda) \equiv \mathcal{D}_{\mathbb{Z},\sigma(\lambda)}$.

The following theorem shows that breaking LWE is as hard as solving hard lattice problems. In particular, given the current state of the art of lattice problems, the LWE assumption is believed to be true for any polynomial $n(\cdot)$ and functions $q(\cdot), \sigma(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, the following constraints are satisfied: $0 < \sigma < q < 2^n$, $n \cdot q/\sigma < 2^{n^\epsilon}$ (for any constant $\epsilon < 1$), and $\sigma > 2\sqrt{n}$.

THEOREM 2.4 (LWE to worst-case lattice problem [Reg05, Pei09, BLP+13]). Fix any polynomial $n(\cdot)$ and functions $q(\cdot), \sigma(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, $0 < \sigma < q < 2^n$, and $\sigma > 2\sqrt{n}$. For every $\lambda \in \mathbb{N}$, let $\eta = \eta(\lambda)$ and $\phi = \phi(\lambda)$ denote the uniform distributions over $\mathbb{Z}_q^n$ and $\chi = \chi(\lambda) \equiv \mathcal{D}_{\mathbb{Z},\sigma}$. If there exist a PPT algorithm $\mathcal{A}$ and a nonnegligible function $\epsilon_{\mathcal{A}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\text{GLWE},\mathcal{A}}^{n,q,\eta,\phi,\chi}(\lambda) \geq \epsilon_{\mathcal{A}}(\lambda)$, then there exist a PPT algorithm $\mathcal{B}$ and a nonnegligible function $\epsilon_{\mathcal{B}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and all instances X of $\text{GapSVP}_{n,n \cdot q/\sigma}$, $\mathcal{B}$ can solve X with probability at least $\epsilon_{\mathcal{B}}(\lambda)$.

ASSUMPTION 2 (LWE with short secrets). Let $n : \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial, and let $q : \mathbb{N} \rightarrow \mathbb{N}$, $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ be functions. The $\text{LWE-ss}_{n,q,\sigma}$ assumption states that $\text{GLWE}_{n,q,\eta,\phi,\chi}$ holds, where $\phi(\lambda)$ is the uniform distributions over $\mathbb{Z}_{q(\lambda)}^{n(\lambda)}$, $\eta(\lambda) \equiv \mathcal{D}_{\mathbb{Z}^{n(\lambda)},\sqrt{2}\sigma(\lambda)}$, and $\chi(\lambda) \equiv \mathcal{D}_{\mathbb{Z},\sigma(\lambda)}$.

The next theorem shows that breaking LWE with short secrets is as hard as breaking (standard) LWE, provided $0 < \sigma(\lambda) < q(\lambda) < 2^{n(\lambda)}$ and $\sigma(\lambda) > \lambda$.

THEOREM 2.5 (LWE with short secrets [ACPS09, Lemma 2], [BLP+13, Lemma 2.12]). Fix any polynomial $n(\cdot)$ and functions $q(\cdot), \sigma(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, $0 < \sigma < q < 2^n$, and $\sigma > \lambda$.¹² For every $\lambda \in \mathbb{N}$, let $\eta(\lambda) \equiv \mathcal{D}_{\mathbb{Z}_q^n, \sqrt{2}\sigma}$, let $\phi(\lambda)$ be the uniform distribution over $\mathbb{Z}_q^n$, and let $\chi(\lambda) \equiv \mathcal{D}_{\mathbb{Z},\sigma}$. If there exist a PPT algorithm $\mathcal{A}$ and a nonnegligible function $\epsilon_{\mathcal{A}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\text{GLWE},\mathcal{A}}^{n,q,\eta,\phi,\chi}(\lambda) \geq \epsilon_{\mathcal{A}}(\lambda)$, then there exist a PPT algorithm $\mathcal{B}$ and a nonnegligible function $\epsilon_{\mathcal{B}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\text{GLWE},\mathcal{B}}^{n,q,\phi,\phi,\chi}(\lambda) \geq \epsilon_{\mathcal{B}}(\lambda)$.

¹²Strictly speaking, it is only required that $\sigma > \sqrt{\ln n + \omega(1) \ln \lambda}$.

ASSUMPTION 3 (LWE with short public vectors). Let $n : \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial, let $q : \mathbb{N} \rightarrow \mathbb{N}$, $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ be functions, and let $\{\chi(\lambda)\}_{\lambda \in \mathbb{N}}$ be the family of distributions over $\mathbb{Z}$. The $\text{LWE-sp}_{n,q,\sigma,\chi}$ assumption states that $\text{GLWE}_{n,q,\eta,\phi,\chi}$ holds, where $\eta(\lambda)$ is the uniform distributions over $\mathbb{Z}_{q(\lambda)}^{n(\lambda)}$, $\phi(\lambda) \equiv \mathcal{D}_{\mathbb{Z}^{n(\lambda)},\sigma(\lambda)}$.

The last theorem in this sequence shows a reduction from LWE with short public vectors to (standard) LWE with a lower dimension.

THEOREM 2.6 (LWE with short public vectors [BLMR13, Corollary 4.6]). Fix any polynomials $n(\cdot), k(\cdot)$ and functions $q(\cdot), \sigma(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $q = q(\lambda)$, $k = k(\lambda)$, $\sigma = \sigma(\lambda)$, $0 < \sigma < q < 2^n$, $k \geq 6n \log q$, and $\sigma \geq \sqrt{n \log q}$. For every $\lambda \in \mathbb{N}$, let $\phi(\lambda) \equiv \mathcal{D}_{\mathbb{Z}_q^k, \sigma}$, and let $\eta(\lambda), \phi'(\lambda)$ denote the uniform distributions over $\mathbb{Z}_q^k$ and $\mathbb{Z}_q^n$, respectively. Now for any distribution $\chi(\lambda)$ over $\mathbb{Z}$, if there exist a PPT algorithm $\mathcal{A}$ and a nonnegligible function $\epsilon_{\mathcal{A}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\text{GLWE}, \mathcal{A}}^{k,q,\eta,\phi,\chi}(\lambda) \geq \epsilon_{\mathcal{A}}(\lambda)$, then there exist a PPT algorithm $\mathcal{B}$ and a nonnegligible function $\epsilon_{\mathcal{B}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\text{GLWE}, \mathcal{B}}^{n,q,\phi',\phi',\chi}(\lambda) \geq \epsilon_{\mathcal{B}}(\lambda)$.

2.1.2. Lattice trapdoors. Lattices with trapdoors are lattices that are indistinguishable from randomly chosen lattices, but have certain “trapdoors” that allow efficient solutions to hard lattice problems.

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

- $\text{TrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, T_{\mathbf{A}})$: 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 $T_{\mathbf{A}}$.
- $\text{SamplePre}(\mathbf{A}, T_{\mathbf{A}}, \sigma, \mathbf{u}) \rightarrow \mathbf{s}$: 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 $\sigma \in \mathbb{R}$ (which determines the length of the output vectors).¹³ It outputs a vector $\mathbf{s} \in \mathbb{Z}_q^m$ such that $\mathbf{A} \cdot \mathbf{s}^T = \mathbf{u}^T$ and $\|\mathbf{s}\| \leq \sqrt{m} \cdot \sigma$.

We require these algorithms to satisfy the following well-sampledness properties. While these properties are similar in spirit to the ones in previous works on lattice trapdoors [Ajt99, GPV08, MP12], there are a couple of differences. First, we present these properties as a security game between a challenger and a computationally bounded adversary.¹⁴ Second, we separate out the dimensions of the matrix and the security parameter.


The first property (well-sampledness of matrix) states that the matrix output by TrapGen should look like a uniformly random matrix.

DEFINITION 2.7 (well-sampledness of matrix). Fix any function $q : \mathbb{N} \rightarrow \mathbb{N}$. A pair of trapdoor generation algorithms $\mathcal{T} = (\text{TrapGen}, \text{SamplePre})$ is said to satisfy the q -well-sampledness of matrix property if for any stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q = q(\lambda)$, $\text{pr}_{\mathcal{T}, \mathcal{A}}^{\text{matrix}, q}(\lambda) = \Pr[1 \leftarrow \text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{matrix}, q}(\lambda)] \leq 1/2 + \text{negl}(\lambda)$, where $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{matrix}, q}(\lambda)$ is defined in Figure 1.


The next property states that the preimage of a uniformly random vector/matrix is indistinguishable from a matrix with entries drawn from Gaussian distribution.

¹³Note that the preimage sampling algorithm could be easily generalized to generate preimages of matrices in $\mathbb{Z}_q^{n \times k}$ (for any k) by independently running SamplePre algorithm on each column of the matrix. Throughout this work, we overload the notation by directly giving matrices $\mathbf{U} \in \mathbb{Z}_q^{n \times k}$ as inputs to the SamplePre algorithm.

¹⁴In some cases, we can consider computationally unbounded adversaries if the inputs of the adversary are polynomially bounded.

FIG. 1. Experiment $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{matrix}, q}$.

DEFINITION 2.8 (preimage sampling). Fix any functions $q : \mathbb{N} \rightarrow \mathbb{N}$ and $\sigma : \mathbb{N} \rightarrow \mathbb{N}$. A pair of trapdoor generation algorithms $\mathcal{T} = (\text{TrapGen}, \text{SamplePre})$ is said to satisfy the (q, σ) -preimage sampling property if for any stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, $\text{pr}_{\mathcal{T}, \mathcal{A}}^{\text{preimg}, q, \sigma}(\lambda) = \Pr[1 \leftarrow \text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{preimg}, q, \sigma}(\lambda)] \leq 1/2 + \text{negl}(\lambda)$, where $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{preimg}, q, \sigma}(\lambda)$ is defined in Figure 2.

FIG. 2. Experiment $\text{Expt}_{\mathcal{T}, \mathcal{A}}^{\text{preimg}, q, \sigma}$.

These properties are satisfied by the trapdoor-based preimage samplers of [GPV08, MP12].

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 working with *leveled* branching programs.

DEFINITION 2.9 (leveled branching program). A leveled branching program of length L , width w , and input space $\{0, 1\}^n$ consists of a sequence of $2L$ functions $\pi_{i,b} : [w] \rightarrow [w]$ 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 $\text{st}_i = \pi_{i,b}(\text{st}_{i-1})$.
- If $\text{st}_L = \text{acc}$, output 1. If $\text{st}_L = \text{rej}$, output 0, else output $\perp$.

Additionally, we also define a notion of “input-circling” (leveled) branching programs. In an input-circling branching program, the input bits are read sequentially in ascending order (i.e., $1, \dots, n, 1, \dots$). Thus, the input-selector function inp is fixed. Additionally, each bit must be read the same number of times. Formally, we describe this as follows.

DEFINITION 2.10. A branching program $\text{BP} = (\{\pi_{i,b} : [w] \rightarrow [w]\}_{i \in [L], b \in \{0, 1\}}, \text{acc} \in [w], \text{rej} \in [w])$ with input space $\{0, 1\}^n$ is said to be a *input-circling branching program* if for all $i \leq L$, $\text{inp}(i) = ((i-1) \bmod n) + 1$, and $L \bmod n = 0$.

Any leveled branching program of length L and input space $\{0, 1\}^n$ can be easily transformed to an input-circling branching program of length $n \cdot L$. Here, we work with classes of branching programs that all share the same input selector function $\text{inp}(\cdot)$ which is known during setup. The input selector as described above is just one possibility, and we stick with it for simplicity. Note that we do not require the transition functions $\pi_{i,b}$ to be permutations.

3. Traitor tracing. In this section, we will first present the syntax and security definitions for traitor tracing schemes. Next, we will introduce the notion of private linear broadcast encryption (PLBE), and finally show that PLBE implies traitor tracing.

The notion of traitor tracing was introduced by Chor, Fiat, and Naor [CFN94]. In a TT scheme for n parties, the setup algorithm chooses a master secret key, a public key, and n secret keys for the users. Encryption can be performed using the public key, and each user can decrypt the ciphertext using his/her secret key. There is also a trace algorithm that, given black box access to a successful pirate decoding box, can catch the traitors who colluded to create the pirate decoding box. Traditional definitions of traitor tracing [CFN94, BSW06] required that the trace algorithm must catch a traitor if a pirate decoding box can decrypt an encryption of a random ciphertext. In this work, we will be using the *indistinguishability*-based definition introduced by Goyal et al. [GKRW17], which is itself based on the definition introduced by Nishimaki, Wichs, and Zhandry [NWZ16]. In this definition, the trace algorithm must catch a traitor even if the pirate decoder box can only distinguish between encryptions of two adversarially chosen messages.

3.1. Public key traitor tracing. A traitor tracing scheme $\mathcal{T}$ with message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ consists of four PPT algorithms, Setup , Enc , Dec , and Trace , with the following syntax:

- $\text{Setup}(1^\lambda, 1^n) \rightarrow (\text{msk}, \text{pk}, (\text{sk}_1, \dots, \text{sk}_n))$. The setup algorithm takes as input the security parameter λ and number of users n and outputs a master secret key msk , a public key pk , and n secret keys $\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n$.
- $\text{Enc}(\text{pk}, m \in \mathcal{M}_\lambda) \rightarrow \text{ct}$. The encryption algorithm takes as input a public key pk and message $m \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow y$. The decryption algorithm takes as input a secret key sk and ciphertext ct and outputs $y \in \mathcal{M}_\lambda \cup \{\perp\}$.
- $\text{Trace}^D(\text{msk}, 1^y, m_0, m_1) \rightarrow T$. The trace algorithm has oracle access to a program D ; it takes as input a master secret key msk , parameter y (in unary), and two messages m_0, m_1 . It outputs a set $T \subset \{1, 2, \dots, n\}$.

Correctness. Informally, a correctness requirement states that decrypting an encryption of message m using any one of the valid secret keys must output m . Formally, a TT scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n \in \mathbb{N}$, $m \in \mathcal{M}_\lambda$, and $i \in \{1, 2, \dots, n\}$, the following holds:

$$\Pr \left[\text{Dec}(\text{sk}_i, \text{ct}) = m : \begin{array}{l} (\text{msk}, \text{pk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ \text{ct} \leftarrow \text{Enc}(\text{pk}, m) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

3.1.1. Security. There are two security requirements for a TT scheme. First, it is required that it satisfy IND-CPA security. Second, it is required that the tracing algorithm must (almost always) correctly trace at least one key used to create a pirate decoding box (whenever the pirate box successfully decrypts with noticeable probability) and also should not falsely accuse any user of cheating. The formal definitions are provided below.

DEFINITION 3.1 (IND-CPA security). A TT scheme $\mathcal{T} = (\text{Setup}, \text{Enc}, \text{Dec}, \text{Trace})$ is IND-CPA secure if for every stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\Pr \left[\begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); b \leftarrow \{0, 1\}; \\ \mathcal{A}(\text{ct}) = b : (\text{msk}, \text{pk}, (\text{sk}_1, \dots, \text{sk}_n)) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ (m_0, m_1) \leftarrow \mathcal{A}(\text{pk}); \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda).$$

DEFINITION 3.2 (ind-secure traitor tracing). Let $\mathcal{T} = (\text{Setup}, \text{Enc}, \text{Dec}, \text{Trace})$ be a TT scheme. For any nonnegligible function $\epsilon(\cdot)$ and PPT adversary $\mathcal{A}$, consider the experiment $\text{Expt-TT}_{\mathcal{A}, \epsilon}^{\mathcal{T}}(\lambda)$ defined in Figure 3.

Experiment $\text{Expt-TT}_{\mathcal{A}, \epsilon}^{\mathcal{T}}(\lambda)$

- $1^n \leftarrow \mathcal{A}(1^\lambda)$.
- $(\text{msk}, \text{pk}, (\text{sk}_1, \dots, \text{sk}_n)) \leftarrow \text{Setup}(1^\lambda, 1^n)$.
- $(D, m_0, m_1) \leftarrow \mathcal{A}^{O(\cdot)}(\text{pk})$.
- $T \leftarrow \text{Trace}^D(\text{msk}, 1^{1/\epsilon(\lambda)}, m_0, m_1)$.

Here, $O(\cdot)$ is an oracle that has $\{\text{sk}_i\}_{i \in [n]}$ hardwired, takes as input an index $i \in [n]$, and outputs sk_i . Let S be the set of indices queried by $\mathcal{A}$.

FIG. 3. Experiment Expt-TT .

Based on the above experiment, we now define the following (probabilistic) events and the corresponding probabilities (which are functions of λ , parameterized by $\mathcal{A}, \epsilon$):

- **Good-Decoder:** $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b)] \geq 1/2 + \epsilon(\lambda)$.
 $\Pr\text{-G-D}_{\mathcal{A}, \epsilon}(\lambda) = \Pr[\text{Good-Decoder}]$.
- **Cor-Tr:** $T \neq \emptyset \wedge T \subseteq S$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A}, \epsilon}(\lambda) = \Pr[\text{Cor-Tr}]$.
- **Fal-Tr:** $T \not\subseteq S$.
 $\Pr\text{-Fal-Tr}_{\mathcal{A}, \epsilon}(\lambda) = \Pr[\text{Fal-Tr}]$.

A TT scheme $\mathcal{T}$ is said to be ind-secure if for every PPT adversary $\mathcal{A}$, polynomial $q(\cdot)$, and nonnegligible function $\epsilon(\cdot)$, there exist negligible functions $\text{negl}_1(\cdot)$, $\text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\epsilon(\lambda) > 1/q(\lambda)$ the following hold:

$$\Pr\text{-Fal-Tr}_{\mathcal{A}, \epsilon}(\lambda) \leq \text{negl}_1(\lambda), \quad \Pr\text{-Cor-Tr}_{\mathcal{A}, \epsilon}(\lambda) \geq \Pr\text{-G-D}_{\mathcal{A}, \epsilon}(\lambda) - \text{negl}_2(\lambda).$$

3.2. Private linear broadcast encryption. Next, we present the notion of private linear broadcast encryption (PLBE). PLBE was introduced by Boneh, Sahai, and Waters [BSW06] as a framework for constructing TT schemes. There are four algorithms in a PLBE scheme: $\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec}$. The setup algorithm outputs a master secret key, public parameters, and n secret keys, one for each user in the system. The public key encryption algorithm can be used to encrypt messages, and ciphertexts can be decrypted using one of the n secret keys via the decryption algorithm. In addition to these algorithms, there is also a special *trace-encryption* algorithm. This algorithm, which uses the master secret key, can be used to encrypt messages to any index $i \in \{0, 1, \dots, n\}$. A secret key for user j can decrypt a ciphertext for index i only if $j > i$.

Boneh, Sahai, and Waters [BSW06] proposed three security definitions for PLBE schemes. The first one requires that special encryptions to index 0 must be indistinguishable from public key encryptions, even if the adversary has all the secret keys.

The next security requirement is that special encryptions to index $i - 1$ must be indistinguishable from special encryptions to index i if the adversary does not have a secret key for user i . Finally, the third security property is that special encryption of message m_0 to index n must be indistinguishable from special encryption of message m_1 to index n , even if the adversary has all secret keys. However, as discussed in section 1.1, the BSW definitions of PLBE do not suffice for constructing TT schemes. Here, we first provide the PLBE syntax, and then present the decoder-based and query-based security definitions of PLBE.

Syntax. A PLBE scheme $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ for message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ has the following syntax:

- $\text{Setup}(1^\lambda, 1^n) \rightarrow (\text{msk}, \text{pp}, (\text{sk}_1, \dots, \text{sk}_n))$. The setup algorithm takes as input the security parameter λ and number of users n and outputs public parameters pp , master secret key msk , and n secret keys $(\text{sk}_1, \text{sk}_2, \dots, \text{sk}_n)$.
- $\text{Enc}(\text{pp}, m) \rightarrow \text{ct}$. The encryption algorithm takes as input public parameters pp and message $m \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .
- $\text{Enc-index}(\text{msk}, m, i \in \{0, 1, 2, \dots, n\}) \rightarrow \text{ct}$. The index-encryption algorithm takes as input the master secret key msk , message $m \in \mathcal{M}_\lambda$, and index $i \in \{0, 1, 2, \dots, n\}$ and outputs a ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow y$. The decryption algorithm takes as input a secret key sk and ciphertext ct and outputs $y \in \mathcal{M}_\lambda \cup \{\perp\}$.

Correctness. A PLBE scheme is said to be correct if there exist negligible functions $\mu_1(\cdot)$, $\mu_2(\cdot)$, $\mu_3(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n \in \mathbb{N}$, $m \in \mathcal{M}_\lambda$, and $i \in \{0, 1, \dots, n\}$, $j \in \{1, 2, \dots, n\}$, the following hold:

$$\begin{aligned} \Pr \left[\text{Dec}(\text{sk}_j, \text{ct}) = m : \begin{array}{l} (\text{msk}, \text{pp}, \{\text{sk}_k\}_{k \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ \text{ct} \leftarrow \text{Enc}(\text{pp}, m) \end{array} \right] &\geq 1 - \mu_1(\lambda), \\ i < j \Rightarrow \Pr \left[\text{Dec}(\text{sk}_j, \text{ct}) = m : \begin{array}{l} (\text{msk}, \text{pp}, \{\text{sk}_k\}_{k \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m, i) \end{array} \right] &\geq 1 - \mu_2(\lambda), \\ i \geq j \Rightarrow \Pr \left[\text{Dec}(\text{sk}_j, \text{ct}) = \perp : \begin{array}{l} (\text{msk}, \text{pp}, \{\text{sk}_k\}_{k \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n) \\ \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m, i) \end{array} \right] &\geq 1 - \mu_3(\lambda). \end{aligned}$$

3.2.1. q -bounded PLBE security. In this section we extend the existing PLBE security definitions by allowing the adversary to make a bounded number of index-encryption queries. Below we describe them in detail.

DEFINITION 3.3 (q -bounded normal hiding security). *Let $q(\cdot)$ be any fixed polynomial. A PLBE scheme is said to satisfy q -bounded normal hiding security if for every stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, $p_{\mathcal{A}}^{q, \text{nrml}}(\lambda) \leq \frac{1}{2} + \text{negl}(\lambda)$, where $p_{\mathcal{A}}^{q, \text{nrml}}(\lambda)$ is defined as follows:*

$$\Pr \left[\begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ m \leftarrow \mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, 0)}(\text{pp}, \{\text{sk}_i\}_{i \in [n]}); \\ b \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} : \mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, 0)}(\text{ct}_b) = b \right],$$

where $\mathcal{A}$ can make at most $q(\lambda)$ queries to the $\text{Enc-index}(\text{msk}, \cdot, 0)$ oracle. Note that here $\mathcal{A}$ is only allowed to query for ciphertexts corresponding to index 0.

DEFINITION 3.4 (q -bounded index hiding security). *Let $q(\cdot)$ be any fixed polynomial. A PLBE scheme is said to satisfy q -bounded index hiding security if for every*

stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$ and every index $i^* \in \{0, \dots, n-1\}$, $p_{\mathcal{A}}^{q,\text{ind}}(\lambda, i^*) \leq \frac{1}{2} + \text{negl}(\lambda)$, where $p_{\mathcal{A}}^{q,\text{ind}}(\lambda, i^*)$ is defined as follows:

$$\Pr \left[\mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, \cdot)}(\text{ct}) = b : \begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ m \leftarrow \mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, \cdot)}(\text{pp}, \{\text{sk}_i\}_{i \in \{1, \dots, n\} \setminus \{i^*+1\}}); \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m, i^* + b) \end{array} \right],$$

where $\mathcal{A}$ can make at most $q(\lambda)$ queries to the $\text{Enc-index}(\text{msk}, \cdot, \cdot)$ oracle. Note that here $\mathcal{A}$ can query the encryption oracle on arbitrary message-index pairs.

DEFINITION 3.5 (q -bounded message hiding security). Let $q(\cdot)$ be any fixed polynomial. A PLBE scheme is said to satisfy q -bounded message hiding security if for every stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, $p_{\mathcal{A}}^{q,\text{msg}}(\lambda) \leq \frac{1}{2} + \text{negl}(\lambda)$, where $p_{\mathcal{A}}^{q,\text{msg}}(\lambda)$ is defined as follows:

$$\Pr \left[\mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, \cdot)}(\text{ct}) = b : \begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ (m_0, m_1) \leftarrow \mathcal{A}^{\text{Enc-index}(\text{msk}, \cdot, \cdot)}(\text{pp}, \{\text{sk}_i\}_{i \in [n]}); \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m_b, n) \end{array} \right],$$

where $\mathcal{A}$ can make at most $q(\lambda)$ queries to the $\text{Enc-index}(\text{msk}, \cdot, \cdot)$ oracle. Note that here $\mathcal{A}$ can query the encryption oracle on arbitrary message-index pairs.

3.2.2. Decoder-based PLBE security. In this section we introduce new decoder-based security definitions for PLBE schemes. We start by formally defining the notion of good distinguishers for PLBE schemes w.r.t. different encryption modes.

PLBE distinguishers. For any $\gamma \in [-1/2, 1/2]$, PPT algorithm D , $\lambda, n \in \mathbb{N}$, $\text{params} = (\text{pp}, \text{msk}, (\text{sk}_1, \dots, \text{sk}_n)) \leftarrow \text{Setup}(1^\lambda, 1^n)$, and message $m \in \mathcal{M}_\lambda$, we say D is $\gamma\text{-Dist}_{\text{params}}^{\text{nrml},0}$ for m if

$$\Pr \left[D(\text{ct}_b) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] \geq \frac{1}{2} + \gamma,$$

where the probability is taken over the random coins used during encryption, the random coins of D , and the choice of b .

Similarly, for any $i \in \{0, \dots, n-1\}$ we can define D to be $\gamma\text{-Dist}_{\text{params}}^{i,i+1}$ for m if

$$\Pr \left[D(\text{ct}_b) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \\ \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m, i + b) \end{array} \right] \geq \frac{1}{2} + \gamma.$$

Finally, we also define D to be $\gamma\text{-Dist}_{\text{params}}^n$ for messages m_0, m_1 if

$$\Pr \left[D(\text{ct}_b) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \\ \text{ct} \leftarrow \text{Enc-index}(\text{msk}, m_b, n) \end{array} \right] \geq \frac{1}{2} + \gamma.$$

DEFINITION 3.6 (decoder-based normal hiding security). A PLBE scheme is said to satisfy decoder-based normal hiding security if for any PPT adversary $\mathcal{A}$, non-negligible function $\gamma(\cdot)$, and polynomial $q(\cdot)$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\gamma(\lambda) > 1/q(\lambda)$, $p_{\mathcal{A}, \gamma, q}^{\text{dec}, \text{nrml}}(\lambda) \leq \text{negl}(\lambda)$, where $p_{\mathcal{A}, \gamma, q}^{\text{dec}, \text{nrml}}(\lambda)$ is defined as follows:

$$\Pr \left[D \text{ is } \gamma(\lambda)\text{-Dist}_{\text{params}}^{\text{nrml},0} \text{ for } m : \begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ (D, m) \leftarrow \mathcal{A}(\text{pp}, \{\text{sk}_i\}_{i \in [n]}) \end{array} \right].$$

DEFINITION 3.7 (decoder-based index hiding security). A PLBE scheme is said to satisfy decoder-based index hiding security if for any PPT adversary $\mathcal{A}$, nonnegligible function $\gamma(\cdot)$, and polynomial $q(\cdot)$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\gamma(\lambda) > 1/q(\lambda)$ and $i^* \in \{0, \dots, n-1\}$, $p_{\mathcal{A}, \gamma, q}^{\text{dec, ind}}(\lambda, i^*) \leq \text{negl}(\lambda)$, where $p_{\mathcal{A}, \gamma, q}^{\text{dec, ind}}(\lambda, i^*)$ is defined as follows:

$$\Pr \left[D \text{ is } \gamma(\lambda)\text{-Dist}_{\text{params}}^{i^*, i^*+1} \text{ for } m : \begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ \text{params} = (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}) \leftarrow \text{Setup}(1^\lambda, 1^n); \\ (D, m) \leftarrow \mathcal{A}(\text{pp}, \{\text{sk}_i\}_{i \in \{1, \dots, n\} \setminus \{i^*+1\}}) \end{array} \right].$$

DEFINITION 3.8 (decoder-based message hiding security). A PLBE scheme is said to satisfy decoder-based message hiding security if for any PPT adversary $\mathcal{A}$, non-negligible function $\gamma(\cdot)$, and polynomial $q(\cdot)$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\gamma(\lambda) > 1/q(\lambda)$, $p_{\mathcal{A}, \gamma, q}^{\text{dec, msg}}(\lambda) \leq \text{negl}(\lambda)$, where $p_{\mathcal{A}, \gamma, q}^{\text{dec, msg}}(\lambda)$ is defined as follows:

$$\Pr \left[D \text{ is } \gamma(\lambda)\text{-Dist}_{\text{params}}^n \text{ for } m_0, m_1 : \begin{array}{l} 1^n \leftarrow \mathcal{A}(1^\lambda); \\ \text{params} \leftarrow \text{Setup}(1^\lambda, 1^n); \\ \text{params} = (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \in [n]}); \\ (D, m_0, m_1) \leftarrow \mathcal{A}(\text{pp}, \{\text{sk}_i\}_{i \in [n]}) \end{array} \right].$$

4. Traitor tracing from 1-bounded secure PLBE. In this section, we show how to construct TT schemes from PLBE schemes that achieve 1-bounded security. Our construction is divided into two components. First, we show that PLBE schemes that achieve 1-bounded security also satisfy decoder-based security. Later, we show how to construct a TT scheme from PLBE schemes that achieves decoder-based security.

4.1. Decoder-based PLBE from 1-bounded secure PLBE. Let $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ be a PLBE scheme that satisfies 1-bounded security. We will show that the same scheme also satisfies decoder-based security.

LEMMA 4.1. *If PLBE satisfies 1-bounded normal hiding security (Definition 3.3), then it also satisfies decoder-based normal hiding security (Definition 3.6).*

Proof. Suppose, on the contrary, there exists a PPT adversary $\mathcal{A}$, nonnegligible function $\gamma(\cdot)$, polynomials $q(\cdot), r(\cdot)$, and an infinite sequence of security parameters $\Lambda = \{\lambda_i\}_{i \in \mathbb{N}}$ such that for all $\lambda \in \Lambda$, $\gamma(\lambda) > 1/q(\lambda)$ and $p_{\mathcal{A}, \gamma, q}^{\text{dec, nrml}}(\lambda) \geq 1/r(\lambda)$. We will use $\mathcal{A}$ that plays the 1-bounded normal hiding security game to build a PPT reduction algorithm $\mathcal{B}$ that plays the decoder-based normal hiding security game as follows.

For any $\lambda \in \mathbb{N}$, the reduction algorithm first receives 1^n from $\mathcal{A}$, which it forwards to the challenger. It then receives pp and n secret keys $\text{sk}_1, \dots, \text{sk}_n$ from the challenger, which it forwards to $\mathcal{A}$. The adversary $\mathcal{A}$ outputs D and m . The reduction algorithm then queries the challenger for an encryption of m for index 0 (recall that the reduction algorithm is allowed one query). Let the challenger's response be ct_1 . It then computes $\text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m)$. Next, it sends challenge message m and receives ct^* , which is either a normal encryption of m or an encryption of m for index 0. The reduction algorithm chooses a random bit $\beta \leftarrow \{0, 1\}$ and checks if $D(\text{ct}_\beta) = D(\text{ct}^*)$. If so, it outputs $b' = \beta$; otherwise it outputs $b' = 1 - \beta$ as its guess.

Let us now analyze $\mathcal{B}$'s advantage. We need to show that there exist polynomials $q_{\mathcal{B}}(\cdot)$ and an infinite sequence $\Lambda_{\mathcal{B}} = \{\lambda_i\}_i$ such that for all $\lambda \in \Lambda_{\mathcal{B}}$, $p_{\mathcal{B}}^{1, \text{nrml}} \geq 1/2 +$

$1/q_B(\lambda)$. Let $\Lambda_B = \Lambda$ and $q_B(\cdot) = q^2(\cdot) \cdot r(\cdot)/2$. Fix any $\lambda \in \Lambda$, and let $\gamma = \gamma(\lambda)$, $q = q(\lambda)$, $r = r(\lambda)$. Let b denote the 1-bounded challenger's choice (recall the challenger chooses $b \leftarrow \{0, 1\}$; if $b = 0$, it sends a normal encryption, else it sends an encryption to index 0). First, let us fix $\text{params} = (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \leq n}) \leftarrow \text{Setup}(1^\lambda, 1^n)$ and $(D, m) \leftarrow \mathcal{A}(\text{pp}, \{\text{sk}_i\}_{i \leq n})$ such that

$$\Pr \left[D(\text{ct}_b) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] = \frac{1}{2} + \alpha_{\text{params}, D, m}$$

for some $\alpha_{\text{params}, D, m} \in [-1/2, 1/2]$. Next, consider the following probability:

$$\rho_{\text{params}, D, m} = \Pr \left[\begin{array}{l} b \leftarrow \{0, 1\}; \beta \leftarrow \{0, 1\}; \\ \text{ct}_0^* \leftarrow \text{Enc}(\text{pp}, m); \text{ct}_1^* \leftarrow \text{Enc-index}(\text{msk}, m, 0); \\ \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0); \\ b' = \beta \text{ if } D(\text{ct}_b^*) = D(\text{ct}_\beta), \text{ else } b' = 1 - \beta \end{array} \right].$$

Since the decoder D is run on ciphertexts $\text{ct}_b^*, \text{ct}_\beta$ independently, we could rewrite the above probability as follows:

$$\begin{aligned} \rho_{\text{params}, D, m} &= \Pr \left[D(\text{ct}_b^*) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0^* \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1^* \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] \\ &\quad \times \Pr \left[D(\text{ct}_\beta) = \beta : \begin{array}{l} \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] \\ &+ \Pr \left[D(\text{ct}_b^*) \neq b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0^* \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1^* \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] \\ &\quad \times \Pr \left[D(\text{ct}_\beta) \neq \beta : \begin{array}{l} \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right] \\ &= \Pr \left[D(\text{ct}_b) = b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right]^2 \\ &\quad + \Pr \left[D(\text{ct}_b) \neq b : \begin{array}{l} b \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{pp}, m); \\ \text{ct}_1 \leftarrow \text{Enc-index}(\text{msk}, m, 0) \end{array} \right]^2 \\ &= \left(\frac{1}{2} + \alpha_{\text{params}, D, m} \right)^2 + \left(\frac{1}{2} - \alpha_{\text{params}, D, m} \right)^2 \\ &= \frac{1}{2} + 2 \cdot \alpha_{\text{params}, D, m}^2. \end{aligned}$$

Thus, we get that for any decoder D that is $\delta\text{-Dist}_{\text{params}}^{\text{nrml}, 0}$ for message m ,

$$\rho_{\text{params}, D, m} = 1/2 + 2 \cdot \alpha_{\text{params}, D, m}^2 \geq 1/2 + 2 \cdot \delta^2.$$

Also, since $\alpha_{\text{params}, D, m}^2 \geq 0$, we get that for every decoder D , $\rho_{\text{params}, D, m} \geq 1/2$. Therefore, since the adversary $\mathcal{A}$ outputs a $1/q\text{-Dist}_{\text{params}}^{\text{nrml}, 0}$ box with probability at least $1/r$, we get that the reduction algorithm $\mathcal{B}$'s winning probability $p_{\mathcal{B}, n}^{1, \text{nrml}}$ (as defined in Definition 3.3) is

$$\begin{aligned} p_{\mathcal{B}, n}^{1, \text{nrml}} &\geq \frac{1}{r} \cdot \left(\frac{1}{2} + \frac{2}{q^2} \right) + \left(1 - \frac{1}{r} \right) \cdot \left(\frac{1}{2} \right) \\ &\geq \frac{1}{2} + \frac{2}{r \cdot q^2}. \end{aligned}$$

This concludes our proof. $\square$

LEMMA 4.2. *If PLBE satisfies 1-bounded index hiding security (Definition 3.4), then it also satisfies decoder-based index hiding security (Definition 3.7).*

The proof of this lemma is identical to the proof of Lemma 4.1, except that the reduction algorithm queries for either a special encryption of m for index i or a special encryption of m for index $i + 1$ (depending on $\beta \leftarrow \{0, 1\}$).

LEMMA 4.3. *If PLBE satisfies 1-bounded message hiding security (Definition 3.5), then it also satisfies decoder-based message hiding security (Definition 3.8).*

The proof of this lemma is identical to the proof of Lemma 4.1, except that the reduction algorithm queries for either a special encryption of m_0 for index n or a special encryption of m_1 for index n (depending on $\beta \leftarrow \{0, 1\}$).

4.2. Traitor tracing from decoder-based PLBE. Consider a PLBE scheme $\text{PLBE} = (\text{PLBE.Setup}, \text{PLBE.Enc}, \text{PLBE.Enc-index}, \text{PLBE.Dec})$ with decoder-based security. We will use PLBE to construct a TT scheme $\mathcal{T} = (\text{Setup}, \text{Enc}, \text{Dec}, \text{Trace})$ as follows. The construction is identical to the transformation in [BSW06]; however, the security proof provided in [BSW06] was not correct. Concretely, to argue correctness of tracing they incorrectly leveraged the indistinguishability-based security of the underlying PLBE scheme. We show that the same transformation could be proven to satisfy correct tracing if one starts with a PLBE scheme that achieves decoder-based security.

Setup($1^\lambda, 1^n$): The setup algorithm samples parameters as $(\text{pp}, \text{msk}, (\text{sk}_1, \dots, \text{sk}_n)) \leftarrow \text{PLBE.Setup}(1^\lambda, 1^n)$. The public parameters are pp , the master secret key is msk , and the n secret keys are $(\text{sk}_1, \dots, \text{sk}_n)$.

Enc(pp, m): The encryption algorithm outputs $\text{ct} \leftarrow \text{PLBE.Enc}(\text{pp}, m)$.

Dec(sk, ct): The decryption algorithm outputs $\text{PLBE.Dec}(\text{sk}, \text{ct})$.

Trace^D($\text{msk}, 1^y, m_0, m_1$): Let $\epsilon = 1/y$ and $W = \lambda \cdot (n \cdot y)^2$. For $i = 0$ to n , the trace algorithm does the following:

1. It first sets $\text{count}_i = 0$. For $j = 1$ to W , it does the following:
 - (a) It chooses a bit $b_{i,j} \leftarrow \{0, 1\}$ and sets $\text{ct}_{i,j} \leftarrow \text{Enc-index}(\text{msk}, m_{b_{i,j}}, i)$. If $D(\text{ct}_{i,j}) = b_{i,j}$, it sets $\text{count}_i = \text{count}_i + 1$.
2. It sets $\hat{p}_i = \text{count}_i / W$.

The trace algorithm outputs every index $i \in \{1, 2, \dots, n\}$ such that $\hat{p}_{i-1} - \hat{p}_i \geq \epsilon/4n$.

Correctness. This follows directly from the first correctness property of the PLBE scheme.

4.2.1. IND-CPA security. We would like to point out that the scheme is IND-CPA secure even if PLBE only satisfies 0-bounded security. In other words, we do not need PLBE to achieve stronger decoder-based security. Thus, the proof of IND-CPA security is identical to that provided in [BSW06]. Below we provide a high level sketch.

THEOREM 4.4. *Assuming the PLBE scheme $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ satisfies the security properties in Definitions 3.6, 3.7, and 3.8, the traitor tracing scheme described above is IND-CPA secure (Definition 3.1).*

Proof. We will construct a sequence of $2n + 2$ hybrid experiments to prove IND-CPA security. The first experiment, Hybrid H_0 , is exactly the IND-CPA game.

Hybrid H_0 . In this experiment, the challenger sends public parameters pp , receives m_0, m_1 from $\mathcal{A}$, and sends $\text{ct} \leftarrow \text{Enc}(\text{pp}, m_0)$ to $\mathcal{A}$.

Hybrid $H_{i,b}$ (for $i \leq n, b \in \{0, 1\}$). This experiment is identical to the IND-CPA experiment, except that the adversary, after sending challenge messages m_0, m_1 , receives $\text{ct} \leftarrow \text{Enc-index}(\text{msk}, i, m_b)$.

Hybrid H_1 . In this experiment, the challenger sends public parameters pp , receives m_0, m_1 from $\mathcal{A}$, and sends $\text{ct} \leftarrow \text{Enc}(\text{pp}, m_1)$ to $\mathcal{A}$.

For any PPT adversary $\mathcal{A}$, let $p_{\mathcal{A},x}(\cdot)$ be a function of λ that denotes the probability of $\mathcal{A}$ outputting 0 in H_x . Note that $p_{\mathcal{A},0} - p_{\mathcal{A},1}$ is the advantage of $\mathcal{A}$ in the IND-CPA security game.

CLAIM 4.5. *Assuming PLBE satisfies Definition 3.6, for any PPT adversary $\mathcal{A}$, there exists a negligible function such that for all $\lambda \in \mathbb{N}$ and $b \in \{0, 1\}$, $|p_{\mathcal{A},b} - p_{\mathcal{A},0,b}| \leq \text{negl}(\lambda)$.*

This follows from decoder-based indistinguishability of normal and 0-index encryptions (Definition 3.6) of PLBE.

CLAIM 4.6. *Assuming PLBE satisfies Definition 3.7, for any PPT adversary $\mathcal{A}$, there exists a negligible function such that for all $\lambda \in \mathbb{N}$, $b \in \{0, 1\}$, and $i \in [n]$, $p_{\mathcal{A},i-1,b} - p_{\mathcal{A},i,b} \leq \text{negl}(\lambda)$.*

This follows from the decoder-based index hiding security notion (Definition 3.7) of PLBE.

CLAIM 4.7. *Assuming PLBE satisfies Definition 3.8, for any PPT adversary $\mathcal{A}$, there exists a negligible function such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},n,0} - p_{\mathcal{A},n,1} \leq \text{negl}(\lambda)$.*

This follows from the decoder-based message hiding security notion (Definition 3.8) of PLBE.

From the above claims, it follows that $p_{\mathcal{A},0} - p_{\mathcal{A},1}$ is bounded by a negligible function. $\square$

4.2.2. Correctness of tracing. Next, we will show that the false trace probability is bounded by a negligible function, and the correct trace probability is close to the probability of $\mathcal{A}$ outputting an ϵ -successful decoding box.

First, we will introduce some notation. Given any pirate decoder box D and messages m_0, m_1 , for any $i \in \{0, 1, \dots, n\}$, let

$$p_i^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc-index}(\text{msk}, i, m_b)],$$

where the probability is taken over random coins of decoder D as well as the randomness used during encryption. Similarly, let $p_{\text{normal}}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{msk}, m_b)]$.

False trace probability. First, we show that the tracing algorithm never falsely accuses any user. Formally, we prove the following.

THEOREM 4.8. *For every PPT adversary $\mathcal{A}$, polynomial $q(\cdot)$, and nonnegligible function $\epsilon(\cdot)$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\epsilon(\lambda) > 1/q(\lambda)$,*

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}(\lambda),$$

where $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\cdot)$ is as defined in Definition 3.2.

Proof. We will skip the dependence of $\epsilon(\cdot)$ on λ for simplicity of notation. Let S be the set of keys queried and D the decoder output by $\mathcal{A}$. For $i \in \{1, 2, \dots, n\}$, we define events $\text{Diff-Adv}_i^D : p_{i-1}^D - p_i^D > \epsilon/8n$ and $\text{Diff-Adv}^D : \bigvee_{k \in \{1, \dots, n\} \setminus S} \text{Diff-Adv}_k^D$.

First, note that the probability of the event *false trace* can be rewritten as follows by conditioning on the events defined above:

$$\Pr[\text{Fal-Tr}] \leq \Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}^D}] + \sum_{i \in \{1, \dots, n\}} \Pr[i \notin S \wedge \text{Diff-Adv}_i^D].$$

We will show that each of these terms is bounded by a negligible function.

LEMMA 4.9. *For every PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_1(\cdot)$ such that for all $\lambda \in \mathbb{N}$,*

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}^D}] \leq \text{negl}_1(\lambda).$$

Proof. The proof of this lemma follows from Chernoff bounds. Let n be the number of users chosen by the adversary $\mathcal{A}$. Fix any $i \in \{1, \dots, n\} \setminus S$ and decode box D . Let us consider the probability that the output of **Trace** algorithm includes i , given that Diff-Adv_i^D does not occur. Note that the tracing algorithm includes i in the traitor set if the estimates $\hat{p}_{i-1}$ and $\hat{p}_i$ differ by at least $\epsilon/4n$.

Let $X_{k,j}$ denote the random variable that is 1 if $D(\text{ct}_{k,j}) = b_{k,j}$ for $k \in \{i-1, i\}$ and $j \in \{1, 2, \dots, W\}$ (here the randomness is over the choice of $b_{k,j}$ and the randomness used by **Enc-index** and D) and $Z_{i,j} = X_{i-1,j} - X_{i,j}$. Then $(\sum_{j=1}^W X_{k,j})/W = \hat{p}_k$ and $\mu_i = \mathbb{E}[Z_{i,j}] = p_{i-1}^D - p_i^D$.

Since the Z_i 's are independent samples, using Chernoff bounds, we get that $\Pr[\sum_j Z_j/W > 2\mu_i] \leq 2^{-O(\lambda)}$. Using this, we can write that for every $i \in \{1, \dots, n\} \setminus S$, $\Pr[\text{Fal-Tr} \wedge i \in T \mid \overline{\text{Diff-Adv}^D}] \leq 2^{-O(\lambda)}$, where T denotes the set of indices output by **Trace**. Finally, using a union bound, we get that

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}^D}] \leq n \cdot 2^{-O(\lambda)} = \text{negl}_1(\lambda).$$

LEMMA 4.10. *Assuming PLBE is a secure PLBE scheme satisfying the decoder-based index hiding security property (Definition 3.7), for every PPT adversary $\mathcal{A}$, polynomial $q(\cdot)$, and nonnegligible function $\epsilon(\cdot)$, there exists a negligible function $\text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\epsilon(\lambda) > 1/q(\lambda)$ and $i \in \{1, 2, \dots, n\}$,*

$$\Pr[i \notin S \wedge \text{Diff-Adv}_i^D] \leq \text{negl}_2(\lambda),$$

where n is the number of users, S is the set of key queries, and D is the decoder box sent by $\mathcal{A}$.

Proof. Suppose, on the contrary, there exists a PPT adversary $\mathcal{A}$, polynomial $q(\cdot)$, and nonnegligible functions $\epsilon(\cdot), \delta(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\epsilon(\lambda) > 1/q(\lambda)$, there exists an $i^* \in \{1, 2, \dots, n\}$ such that $\Pr[i^* \notin S \wedge \text{Diff-Adv}_{i^*}^D] \geq \delta(\lambda)$. Then we can use $\mathcal{A}$ to build a PPT reduction algorithm $\mathcal{B}$ that breaks the index hiding security property of PLBE.

The reduction algorithm $\mathcal{B}$ first receives 1^n from the adversary, which it forwards to the challenger. It then receives the PLBE public parameters pp from the challenger, which it sends to $\mathcal{A}$. Next, it chooses an index $i \leftarrow \{1, 2, \dots, n\}$ and sends it to the PLBE challenger.¹⁵ It receives secret keys $\text{sk}_1, \dots, \text{sk}_{i-1}, \text{sk}_{i+1}, \dots, \text{sk}_n$. The adversary $\mathcal{A}$ queries for secret keys. If $\mathcal{A}$ queries for i , $\mathcal{B}$ sends an empty decoding

¹⁵In other words, the reduction algorithm randomly guesses the index hiding challenger with which it interacts.

box to the PLBE challenger. Otherwise, on receiving query $j \neq i$ from $\mathcal{A}$, it sends $\mathbf{sk}_j$ to $\mathcal{A}$. After all queries, the adversary sends a decoding box D and messages m_0, m_1 to $\mathcal{B}$. The reduction algorithm chooses a uniformly random bit $b' \leftarrow \{0, 1\}$ and sends $D, m_{b'}$ to the PLBE challenger.

Let $p_{j,b}^D = \Pr[D(\mathbf{ct}) = b : \mathbf{ct} \leftarrow \text{Enc-index}(\mathbf{msk}, j, m_b)]$, where the probability is taken over the coins of decoder D and the encryption algorithm. Recall that we have $\Pr[i^* \notin S \wedge \text{Diff-Adv}_{i^*}] \geq \delta(\lambda)$. Therefore, we can write that

$$\begin{aligned} & \Pr[i^* \notin S \wedge ((p_{i^*-1,0}^D + p_{i^*-1,1}^D)/2 - (p_{i^*,0}^D + p_{i^*,1}^D)/2) \geq \epsilon/8n] \geq \delta(\lambda) \\ \Rightarrow & \Pr[i = i^* \wedge i^* \notin S \wedge ((p_{i-1,0}^D + p_{i-1,1}^D)/2 - (p_{i,0}^D + p_{i,1}^D)/2) \geq \epsilon/8n] \geq \delta(\lambda)/n. \end{aligned}$$

Thus, we can also write that there exists a bit b such that

$$\Pr[i = i^* \wedge i^* \notin S \wedge (p_{i-1,b}^D - p_{i,b}^D) \geq \epsilon/8n] \geq \delta(\lambda)/n. \quad \square$$

Now since the reduction algorithm $\mathcal{B}$ simply randomly guesses this bit b , thus we have that with probability at least $\delta/2n$, $\mathcal{B}$ outputs a decoding box D and a message m_b such that D can distinguish between encryptions of m_b to indices $i-1$ and i with advantage at least $\epsilon/8n$. Thus, the lemma follows.

From the above lemmas, it follows that the probability of false trace is at most $\text{negl}_1(\lambda) + n \cdot \text{negl}_2(\lambda)$, and thus the theorem follows.

Correct trace probability. Now we show that whenever the adversary outputs a good decoder, then with all but negligible probability the tracing algorithm outputs a nonempty set T . Combining this with Theorem 4.8, we get that the tracing algorithm correctly traces. Formally, we show the following.

THEOREM 4.11. *For every PPT adversary $\mathcal{A}$, polynomial $q(\cdot)$, and nonnegligible function $\epsilon(\cdot)$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ satisfying $\epsilon(\lambda) > 1/q(\lambda)$,*

$$\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \Pr\text{-G-D}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}(\lambda),$$

where $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\cdot)$ and $\Pr\text{-G-D}_{\mathcal{A},\epsilon}(\cdot)$ are as defined in Definition 3.2.

Proof. Let us start by analyzing the probability that the tracing algorithm outputs a nonempty set T . First, we know that if event **Good-Decoder** occurs, then $p_{\text{norm}}^D \geq 1/2 + \epsilon$. Next, let S be the set of indices $i \in \{1, \dots, n\}$ such that $p_{i-1}^D - p_i^D > \epsilon/2n$. Using Chernoff bounds, we get that

$$(4.1) \quad \forall i \in S, \quad \Pr[\hat{p}_{i-1}^D - \hat{p}_i^D < \epsilon/4n] \leq 2^{-O(\lambda)} = \text{negl}_1(\lambda).$$

Note that by message hiding security of the underlying PLBE scheme, we have that $p_n^D \leq 1/2 + \text{negl}_2(\lambda)$ for some negligible function $\text{negl}_2(\cdot)$. Also, by indistinguishability of *normal* and index 0 ciphertexts, we have that $p_{\text{norm}}^D - p_0^D \leq \text{negl}_3(\lambda)$ for some negligible function $\text{negl}_3(\cdot)$. Thus, $p_0^D - p_n^D \geq \epsilon - \text{negl}_2(\lambda) - \text{negl}_3(\lambda) > \epsilon/2$. Given this, we can conclude that the set S as defined above (i.e., for $i \in S$, $p_{i-1}^D - p_i^D > \epsilon/2n$) must be nonempty whenever event **Good-Decoder** occurs. Combining this with (4.1), we get that

$$\Pr[T \neq \emptyset] \geq (1 - \text{negl}_1(\lambda)) \cdot \Pr\text{-G-D}_{\mathcal{A},\epsilon}(\lambda) \geq \Pr\text{-G-D}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}(\lambda).$$

Finally, combining with Theorem 4.8, we get that

$$\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \Pr\text{-G-D}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}(\lambda).$$

This concludes the proof. $\square$

5. Mixed functional encryption. A functional encryption scheme consists of a setup, an encryption, a key generation, and a decryption algorithm. The setup algorithm takes the security parameter and functionality index as inputs, and outputs public parameters and a master secret key. The encryption algorithm uses the public parameters to encrypt a message, while the key generation algorithm uses the master secret key to compute a secret key corresponding to a function. The decryption algorithm takes as input a ciphertext and a secret key, and outputs the function evaluation on the message.

In this work, we introduce the notion of Mixed Functional Encryption (Mixed FE). A Mixed FE scheme is defined as a dual of the standard functional encryption (i.e., ciphertext-policy) in which the secret keys are associated with a message, and ciphertexts are associated with (boolean) functions. Additionally, in a Mixed FE system, there are two encryption algorithms: **Enc** and **SK-Enc**. The normal encryption algorithm **Enc** takes as input only the public parameters and outputs an encryption of a “canonical” *always-accepting* function. The “secret key” encryption algorithm, on the other hand, takes as input the master secret key and a function f and encrypts f . The decryption algorithm in a Mixed FE system works similarly to that in standard functional encryption; that is, it outputs the evaluation of encrypted function f on the message m associated with the secret key. Below we provide a formal definition.

Consider function classes $\mathcal{F} = \{\mathcal{F}_\kappa\}_\kappa$ and message spaces $\mathcal{M} = \{\mathcal{M}_\kappa\}_\kappa$, where $f : \mathcal{M}_\kappa \rightarrow \{0, 1\}$ for each $f \in \mathcal{F}_\kappa$.¹⁶ A Mixed FE scheme **Mixed-FE**, for function classes $\mathcal{F}$ and message spaces $\mathcal{M}$, consists of five polytime algorithms (**Setup**, **Enc**, **SK-Enc**, **KeyGen**, **Dec**) with the following syntax:

- **Setup**($1^\lambda, 1^\kappa$) $\rightarrow$ (**pp**, **msk**). The setup algorithm takes as input the security parameter λ and functionality index κ and outputs the public parameters **pp** and the master secret key **msk**.
- **Enc**(**pp**) $\rightarrow$ **ct**. The normal encryption algorithm takes as input public parameters **pp** and outputs a ciphertext **ct**.
- **SK-Enc**(**msk**, f) $\rightarrow$ **ct**. The secret key encryption algorithm takes as input master secret key **msk** and a function $f \in \mathcal{F}_\kappa$. It outputs a ciphertext **ct**.
- **KeyGen**(**msk**, m) $\rightarrow$ $\mathbf{sk}_m$. The key generation algorithm takes as input master secret key **msk** and a message/input $m \in \mathcal{M}_\kappa$. It outputs a secret key $\mathbf{sk}_m$.
- **Dec**($\mathbf{sk}_m$, **ct**) $\rightarrow \{0, 1\}$. The decryption algorithm takes as input a secret key $\mathbf{sk}_m$ and a ciphertext **ct**, and it outputs a single bit.

Correctness. A Mixed FE scheme is said to be correct if there exist negligible functions $\text{negl}_1(\cdot)$, $\text{negl}_2(\cdot)$ such that for all $\lambda, \kappa \in \mathbb{N}$ and for every $f \in \mathcal{F}_\kappa$, $m \in \mathcal{M}_\kappa$, the following hold:

$$\Pr \left[\text{Dec}(\mathbf{sk}_m, \text{ct}) = 1 : \begin{array}{l} (\mathbf{pp}, \mathbf{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ \mathbf{sk}_m \leftarrow \text{KeyGen}(\mathbf{msk}, m); \text{ct} \leftarrow \text{Enc}(\mathbf{pp}) \end{array} \right] \geq 1 - \text{negl}_1(\lambda),$$

$$\Pr \left[\text{Dec}(\mathbf{sk}_m, \text{ct}) = f(m) : \begin{array}{l} (\mathbf{pp}, \mathbf{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ \mathbf{sk}_m \leftarrow \text{KeyGen}(\mathbf{msk}, m); \\ \text{ct} \leftarrow \text{SK-Enc}(\mathbf{msk}, f) \end{array} \right] \geq 1 - \text{negl}_2(\lambda).$$

Security. Informally, for security we require that no PPT adversary should be able to distinguish between secret key encryptions of two functions f_0 and f_1 if for every key in its possession, the output of f_0, f_1 is identical. Additionally, we also

¹⁶The following definition could be easily generalized for multibit function classes, but for simplicity we stick to boolean functions.

require that it should be hard to distinguish between normal encryptions and secret key encryptions of the special always-accepting function. In this work, we are only interested in Mixed FE schemes that guarantee security against adversaries which make a bounded number of secret key encryption queries. Below we formally define it.

DEFINITION 5.1 (*q*-bounded function indistinguishability). *Let $q(\cdot)$ be any fixed polynomial. A Mixed FE scheme $\text{Mixed-FE} = (\text{Setup}, \text{Enc}, \text{SK-Enc}, \text{KeyGen}, \text{Dec})$ is said to satisfy *q*-bounded function indistinguishability security if for every stateful PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:*

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{ct}) = b : \begin{array}{l} 1^\kappa \leftarrow \mathcal{A}(1^\lambda); (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ (f^{(0)}, f^{(1)}) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{pp}); \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{SK-Enc}(\text{msk}, f^{(b)}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where

- $\mathcal{A}$ can make at most $q(\lambda)$ queries to $\text{SK-Enc}(\text{msk}, \cdot)$ oracle; and
- every secret key query m made by adversary $\mathcal{A}$ to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy the condition that $f^{(0)}(m) = f^{(1)}(m)$.

We also define a restricted version of the function indistinguishability game in which the adversary must declare its challenge functions $(f^{(0)}, f^{(1)})$ at the beginning, and it must make all its q encryption queries before any of its key generation queries.

DEFINITION 5.2 (*q*-bounded restricted function indistinguishability). *Let $q(\cdot)$ be any fixed polynomial. A Mixed FE scheme $\text{Mixed-FE} = (\text{Setup}, \text{Enc}, \text{SK-Enc}, \text{KeyGen}, \text{Dec})$ is said to satisfy *q*-bounded selective function indistinguishability security if for every stateful PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:*

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{pp}, \text{ct}) = b : \begin{array}{l} (1^\kappa, f^{(0)}, f^{(1)}) \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{SK-Enc}(\text{msk}, f^{(b)}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where

- $\mathcal{A}$ can make at most $q(\lambda)$ queries to the $\text{SK-Enc}(\text{msk}, \cdot)$ oracle;
- every secret key query m made by adversary $\mathcal{A}$ to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy the condition that $f^{(0)}(m) = f^{(1)}(m)$; and
- $\mathcal{A}$ must make all (at most $q(\lambda)$) $\text{SK-Enc}(\text{msk}, \cdot)$ oracle queries before making any query to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle.

DEFINITION 5.3 (*q*-bounded accept indistinguishability). *Let $q(\cdot)$ be any fixed polynomial. A Mixed FE scheme $\text{Mixed-FE} = (\text{Setup}, \text{Enc}, \text{SK-Enc}, \text{KeyGen}, \text{Dec})$ is said to satisfy *q*-bounded accept indistinguishability security if for every stateful PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$*

the following holds:

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{ct}_b) = b : \begin{array}{l} 1^\kappa \leftarrow \mathcal{A}(1^\lambda); (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ f^* \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{pp}); \\ b \leftarrow \{0, 1\}; \text{ct}_1 \leftarrow \text{SK-Enc}(\text{msk}, f^*); \\ \text{ct}_0 \leftarrow \text{Enc}(\text{pp}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where

- $\mathcal{A}$ can make at most q queries to $\text{SK-Enc}(\text{msk}, \cdot)$ oracle; and
- every secret key query m made by adversary $\mathcal{A}$ to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy the condition that $f^*(m) = 1$.

Additionally, we also define a restricted notion of the accept indistinguishability property for Mixed FE schemes, in which the adversary must declare its challenge function f^* at the beginning, and it must make all its q encryption queries before any of its key generation queries, and it is restricted to only making ciphertext queries for functions f such that all queried f 's evaluate to 1 on all (secret key) queried messages m . Below we formally describe it.

DEFINITION 5.4 (q -bounded restricted accept indistinguishability). *Let $q(\cdot)$ be any fixed polynomial. A Mixed FE scheme $\text{Mixed-FE} = (\text{Setup}, \text{Enc}, \text{SK-Enc}, \text{KeyGen}, \text{Dec})$ is said to satisfy q -bounded restricted accept indistinguishability security if for every stateful PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:*

$$\Pr \left[\mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot), \text{SK-Enc}(\text{msk}, \cdot)}(\text{pp}, \text{ct}_b) = b : \begin{array}{l} (1^\kappa, f^*) \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ b \leftarrow \{0, 1\}; \text{ct}_1 \leftarrow \text{SK-Enc}(\text{msk}, f^*); \\ \text{ct}_0 \leftarrow \text{Enc}(\text{pp}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where

- $\mathcal{A}$ can make at most $q(\lambda)$ queries to $\text{SK-Enc}(\text{msk}, \cdot)$ oracle;
- every secret key query m made by adversary $\mathcal{A}$ to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle must satisfy the condition that $f^*(m) = 1$ as well as $f(m) = 1$ for every ciphertext query f made by $\mathcal{A}$ to the $\text{SK-Enc}(\text{msk}, \cdot)$ oracle; and
- $\mathcal{A}$ must make all (at most $q(\lambda)$) $\text{SK-Enc}(\text{msk}, \cdot)$ oracle queries before making any query to the $\text{KeyGen}(\text{msk}, \cdot)$ oracle.

6. Construction of PLBE from Mixed FE and ABE. In this section, we construct a private linear broadcast encryption (PLBE) scheme from any key-policy attribute-based encryption (KP-ABE) scheme and a Mixed FE scheme. Our construction inherits the message space of the underlying KP-ABE scheme. Also, we show that if the underlying ABE scheme is selectively secure, and the Mixed FE scheme satisfies 1-bounded restricted function and accept indistinguishability properties, then our PLBE scheme satisfies 1-bounded normal, index, and message hiding security properties.

Outline. The idea is to use the ABE system to encrypt a message with attributes being set to either the “normal” ciphertext (i.e., encryption of the canonical always-accepting function) or a special (secret key) ciphertext which encrypts the comparison

function depending on the type of PLBE encryption operation being performed. Each user's secret key will be an ABE private key. Here the ABE private key is generated for the Mixed FE decryption circuit in which a Mixed FE secret key, corresponding to the user's index, is hardwired. The high level intuition is that when the attribute is a normal functional encryption ciphertext, then all keys decrypt it to 1; thus any user with an appropriate ABE key could perform the decryption. Also, when the attribute is set to be a special ciphertext (that encrypts comparison with some index i), then only those users whose indices are larger than the threshold i set can perform the decryption. For proving security, we rely on the fact that special ciphertexts are indistinguishable to any adversary that does not have distinguishing secret keys. Below we provide a detailed overview.

The PLBE setup algorithm starts by sampling an ABE key pair $(\text{abe.pp}, \text{abe.msk})$ and a Mixed FE key pair $(\text{Mixed.pp}, \text{Mixed.msk})$. To generate the private key for the i th user, it first generates a Mixed FE secret key Mixed.sk_i for message i , and later computes an ABE key abe.sk_i for predicate $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$, i.e., Mixed-FE decryption circuit with key Mixed.sk_i hardwired. For PLBE normal encryption, one simply computes ciphertext ct as an encryption of message m under attributes ct_{attr} , where ct_{attr} is a Mixed FE normal ciphertext. For encrypting a message to index i , the encryption algorithm works identically, except now the attribute is set to be a special ciphertext corresponding to function *greater than* i . Finally, the PLBE decryption is simply the ABE decryption algorithm.

Now correctness follows directly from the correctness of ABE and functional encryption schemes. For proving security, the main idea is as follows: Suppose there exists an adversary that can distinguish between PLBE normal ciphertexts and index 0 ciphertexts; then it can be used to distinguish between Mixed FE normal ciphertexts and secret key ciphertexts encrypting function *greater than* 0 (note that this is an always-accepting function). In other words, such an attack can be used to break the restricted accept indistinguishability property of the Mixed FE scheme. Similarly, we can also reduce a successful attack on the index hiding, or message hiding security to an attacker on restricted function indistinguishability of Mixed FE or ABE security, respectively. Below we describe our construction $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ for message spaces $\{\mathcal{M}_\kappa\}_\kappa$ in detail.

The following construction only achieves the statistical notion of correctness. In Appendix B, we provide an alternate construction that achieves perfect correctness from the same assumptions.

6.1. Construction. Let $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ be a KP-ABE scheme for a set of attribute spaces $\{\mathcal{X}_\kappa\}_\kappa$, predicate classes $\{\mathcal{C}_\kappa\}_\kappa$, and message spaces $\{\mathcal{M}_\kappa\}_\kappa$, and let $\text{Mixed-FE} = (\text{Mixed.Setup}, \text{Mixed.Enc}, \text{Mixed.SK-Enc}, \text{Mixed.KeyGen}, \text{Mixed.Dec})$ be a Mixed FE scheme for function classes $\{\mathcal{F}_\kappa\}_\kappa$ and message space $\{\mathcal{I}_\kappa\}_\kappa$, with ciphertexts of length $\ell(\lambda, \kappa)$. For every n , let $\kappa = \kappa(n)$ be the lexicographically smallest functionality index such that every string of length $\log(n)$ can be uniquely represented in message space $\mathcal{I}_\kappa$ (i.e., $\{0, 1\}^{\log(n)} \subseteq \mathcal{I}_\kappa$), and function class $\mathcal{F}_\kappa$ contains the “comparison” ($>$) operator. Also, let $\tilde{\kappa} = \tilde{\kappa}(\lambda, \kappa)$ be the lexicographically smallest functionality index such that every string of length $\ell(\lambda, \kappa)$ can be uniquely represented in attribute class $\mathcal{X}_{\tilde{\kappa}}$ (i.e., $\{0, 1\}^{\ell(\lambda, \kappa)} \subseteq \mathcal{X}_{\tilde{\kappa}}$), and $\mathcal{C}_{\tilde{\kappa}}$ contains Mixed FE decryption circuit corresponding to functionality index κ . Below we describe our construction.

- $\text{Setup}(1^\lambda, 1^n) \rightarrow (\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \leq n})$. The setup algorithm runs ABE.Setup and Mixed.Setup to generate ABE and Mixed FE public parameters and

master secret key as $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$ and $(\text{Mixed.pp}, \text{Mixed.msk}) \leftarrow \text{Mixed.Setup}(1^\lambda, 1^\kappa)$. Next, it runs Mixed.KeyGen to generate n mixed FE secret keys Mixed.sk_i as

$$\forall i \leq n, \quad \text{Mixed.sk}_i \leftarrow \text{Mixed.KeyGen}(\text{Mixed.msk}, i).$$

Let $C_{\text{Mixed.sk}_i}$ denote the circuit $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$, i.e., Mixed-FE decryption circuit with key Mixed.sk_i hardwired. Next, it computes n ABE secret keys abe.sk_i as

$$\forall i \leq n, \quad \text{abe.sk}_i \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_{\text{Mixed.sk}_i}).$$

Finally, it sets $\text{pp} = (\text{abe.pp}, \text{Mixed.pp})$, $\text{msk} = (\text{abe.msk}, \text{Mixed.msk})$ and $\text{sk}_i = \text{abe.sk}_i$ for $i \leq n$.

- $\text{Enc}(\text{pp}, m) \rightarrow \text{ct}$. Let $\text{pp} = (\text{abe.pp}, \text{Mixed.pp})$. The encryption algorithm first computes $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.Enc}(\text{Mixed.pp})$. Next, it encrypts message m as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$, and outputs ciphertext ct .
- $\text{Enc-index}(\text{msk}, m, i) \rightarrow \text{ct}$. Let $\text{msk} = (\text{abe.msk}, \text{Mixed.msk})$ and let comp_i denote the comparison function $\stackrel{?}{>} i$, i.e., $\text{comp}_i(x) = 1$ iff $x > i$. The encryption algorithm first computes $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.SK-Enc}(\text{Mixed.msk}, \text{comp}_i)$. Next, it encrypts message m as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$ and outputs ciphertext ct .
- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow m$ or $\perp$. The decryption algorithm runs ABE.Dec on ct using key sk as $y = \text{ABE.Dec}(\text{sk}, \text{ct})$ and sets y as the output of decryption.

6.2. Correctness. For all $\lambda, n \in \mathbb{N}$, message $m \in \mathcal{M}_\lambda$, public parameters and master secret keys $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$, $(\text{Mixed.pp}, \text{Mixed.msk}) \leftarrow \text{Mixed.Setup}(1^\lambda, 1^\kappa)$, the secret keys sk_i for $i \leq n$ are simply the ABE keys $\text{abe.sk}_i \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_{\text{Mixed.sk}_i})$. For any index $i \leq n$, consider the following two cases:

1. **Normal encryption.** For any ciphertext computed as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$, where $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.Enc}(\text{Mixed.pp})$, we know that with all but negligible probability $\text{Mixed.Dec}(\text{Mixed.sk}_i, \text{ct}_{\text{attr}}) = 1$ by correctness of the Mixed FE scheme. In other words, $C_{\text{Mixed.sk}_i}(\text{ct}_{\text{attr}}) = 1$. Therefore, by correctness of the ABE scheme, we get that with all but negligible probability $\text{ABE.Dec}(\text{abe.sk}_i, \text{ct}) = m$.
2. **Index encryption.** For any index $0 \leq j \leq n$ and ciphertext ct computed as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$, where $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.SK-Enc}(\text{Mixed.msk}, \text{comp}_j)$, we know that with all but negligible probability

$$\text{Mixed.Dec}(\text{Mixed.sk}_i, \text{ct}_{\text{attr}}) = \begin{cases} 1 & \text{if } i > j, \\ 0 & \text{otherwise} \end{cases}$$

by correctness of the Mixed FE scheme. In other words, $C_{\text{Mixed.sk}_i}(\text{ct}_{\text{attr}}) = \text{comp}_j(i) = (i > j)$. Therefore, by correctness of the ABE scheme, we have that with all but negligible probability $\text{ABE.Dec}(\text{abe.sk}_i, \text{ct}) = m$ for $i > j$ and $\perp$ otherwise.

Therefore, PLBE satisfies the PLBE correctness condition.

6.3. Security. We will now show that the scheme described above is 1-bounded secure per Definitions 3.3, 3.4, and 3.5. In other words, it satisfies normal hiding, index hiding, and message hiding security properties. Formally, we prove the following.

THEOREM 6.1. *If $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a selectively secure ABE scheme for a set of attribute spaces $\{\mathcal{X}_\kappa\}_\kappa$, predicate classes $\{\mathcal{C}_\kappa\}_\kappa$, and message spaces $\{\mathcal{M}_\kappa\}_\kappa$ satisfying Definition A.2, and $\text{Mixed-FE} = (\text{Mixed.Setup}, \text{Mixed.Enc}, \text{Mixed.SK-Enc}, \text{Mixed.KeyGen}, \text{Mixed.Dec})$ is a Mixed FE scheme, for function classes $\{\mathcal{F}_\kappa\}_\kappa$ and message spaces $\{\mathcal{I}_\kappa\}_\kappa$, satisfying 1-bounded restricted function indistinguishability (Definition 5.2) and 1-bounded restricted accept indistinguishability (Definition 5.4) properties, then $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ is a secure PLBE scheme, for messages spaces $\{\mathcal{M}_\kappa\}_\kappa$, satisfying 1-bounded normal, index, and message hiding security properties per Definitions 3.3, 3.4, and 3.5, respectively.*

Our proof is divided into three components/lemmas, one for each PLBE security property. Let $\mathcal{A}$ be any PPT adversary that wins the normal/index/message hiding game with nonnegligible advantage. We argue that such an adversary must break the security of at least one underlying primitive.

6.3.1. Normal hiding security.

LEMMA 6.2. *If $\text{Mixed-FE} = (\text{Mixed.Setup}, \text{Mixed.Enc}, \text{Mixed.SK-Enc}, \text{Mixed.Dec}, \text{Mixed.KeyGen})$ is a Mixed FE scheme satisfying the 1-bounded restricted accept indistinguishability (Definition 5.4) property, then $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Dec}, \text{Enc-index})$ is a PLBE scheme satisfying the 1-bounded normal hiding security property per Definition 3.3.*

Proof. Suppose there exists an adversary $\mathcal{A}$ such that $\mathcal{A}$'s advantage in a 1-bounded normal hiding security game is nonnegligible. We construct an algorithm $\mathcal{B}$ that can distinguish normal encryptions from secret key encryptions, therefore breaking the 1-bounded restricted accept indistinguishability security of the Mixed FE scheme.

The reduction algorithm $\mathcal{B}$ receives 1^n from $\mathcal{A}$. It sets $\kappa, \tilde{\kappa}$ as in the construction and sends κ as the functionality index and comp_0 (i.e., comparison with 0) as its challenge function to the Mixed FE challenger. The challenger generates a key pair $(\text{Mixed.pp}, \text{Mixed.msk})$ and sends Mixed.pp as the public parameters and challenge ciphertext $\text{ct}_{\text{attr}}^*$ to $\mathcal{B}$. Next, $\mathcal{B}$ makes an encryption query for function comp_0 . Let the challenger's response be ciphertext ct_{attr} . $\mathcal{B}$ then queries the challenger on n messages i ($\leq n$) for corresponding mixed FE secret keys and receives back keys Mixed.sk_i for $i \leq n$. It then chooses an ABE key pair $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$ and computes n ABE keys as $\text{abe.sk}_i \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_{\text{Mixed.sk}_i})$. Next, it sends $(\text{abe.pp}, \text{Mixed.pp})$ and $\{\text{abe.sk}_i\}_{i \leq n}$ as the PLBE public parameters and secret keys to $\mathcal{A}$. After receiving all the keys, $\mathcal{A}$ sends its challenge message m^* to $\mathcal{B}$, and it can also make a single encryption query for message m on index 0. Here $\mathcal{A}$ is allowed to make the encryption query either before or after challenge query. The reduction algorithm $\mathcal{B}$ responds to each query as follows: $\mathcal{B}$ encrypts message m as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$ and sends ct to $\mathcal{A}$ as its response to the encryption query. Also, it computes ciphertext ct^* as $\text{ct}^* \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}^*, m^*)$ and sends ct^* as the challenge ciphertext to $\mathcal{A}$. Note that $\mathcal{A}$ could instead have sent its challenge query before sending the index encryption query. Also, $\mathcal{B}$ does not need to query the Mixed FE challenger for answering any query at this point as it already has ciphertexts $\text{ct}_{\text{attr}}, \text{ct}_{\text{attr}}^*$. Finally, $\mathcal{A}$ sends its guess b to $\mathcal{B}$, and $\mathcal{B}$ forwards b as its own guess.

First, note that both $\mathcal{A}$ and $\mathcal{B}$ are allowed to make at most 1 index encryption and polynomially many secret key encryption queries, respectively. Also, note that $\mathcal{B}$ sends its secret key encryption query as well as its challenge query before making

making key generation queries; thus $\mathcal{B}$ is an admissible adversary in the 1-bounded restricted accept indistinguishability game. Since $\mathcal{A}$ is only allowed to make encryption queries to index 0 (in the 1-bounded normal hiding security game), thus $\mathcal{B}$ queries the Mixed FE challenger on functions comp_0 which are always accepting functions and therefore admissible queries per the restricted accept indistinguishability game. Next, for each query made by $\mathcal{A}$, $\mathcal{B}$ queries the Mixed FE challenger exactly once, and thus all the queries are honestly and exactly answered. Finally, note that if the Mixed FE challenger computed $\text{ct}_{\text{attr}}^*$ as a normal functional encryption ciphertext, then $\mathcal{B}$ computes ct^* as a normal PLBE ciphertext; otherwise it computes ct^* as a PLBE ciphertext for index 0. Thus, $\mathcal{B}$ perfectly simulates the 1-bounded normal hiding security game for $\mathcal{A}$. As a result, if $\mathcal{A}$'s advantage is nonnegligible, then $\mathcal{B}$ breaks the 1-bounded restricted accept indistinguishability security with nonnegligible advantage. This completes the proof. $\square$

6.3.2. Index hiding security.

LEMMA 6.3. *If Mixed-FE = (Mixed.Setup, Mixed.Enc, Mixed.SK-Enc, Mixed.Dec, Mixed.KeyGen) is a Mixed FE scheme satisfying the 1-bounded restricted function indistinguishability (Definition 5.2) property, then PLBE = (Setup, Enc, Dec, Enc-index) is a PLBE scheme satisfying the 1-bounded index hiding security property per Definition 3.4.*

Proof. The proof of this lemma is similar to that of Lemma 6.2, with the additional modification that the reduction algorithm now guesses the indices i, i^* on which the PLBE adversary makes its encryption query and challenge query, respectively, and the reduction algorithm aborts if its guess is incorrect. This leads to a polynomial loss ($\approx 1/n^2$) in the reduction algorithm's advantage.¹⁷

Suppose there exists an adversary $\mathcal{A}$ such that $\mathcal{A}$'s advantage in the 1-bounded index hiding security game is nonnegligible. We construct an algorithm $\mathcal{B}$ that can distinguish between two secret key encryptions, therefore breaking the 1-bounded restricted function indistinguishability security of the Mixed FE scheme.

The reduction algorithm $\mathcal{B}$ receives 1^n from $\mathcal{A}$. It sets $\kappa, \tilde{\kappa}$ as in the construction. Next, it guesses the challenge index $i^* \in \{0, \dots, n-1\}$ and query index $i \in \{0, \dots, n\}$.¹⁸ It sends κ as the functionality index and $(\text{comp}_{i^*}, \text{comp}_{i^*+1})$ (i.e., comparison with i^* and i^*+1) as its challenge functions to the Mixed FE challenger. The challenger generates a key pair $(\text{Mixed.pp}, \text{Mixed.msk})$ and sends Mixed.pp as the public parameters and challenge ciphertext $\text{ct}_{\text{attr}}^*$ to $\mathcal{B}$. Next, $\mathcal{B}$ makes an encryption query for function comp_i . Let the challenger's response be ciphertext ct_{attr} . $\mathcal{B}$ then queries the challenger on $n-1$ messages $j \in [n] \setminus \{i^*+1\}$ for corresponding mixed FE secret keys and receives back keys Mixed.sk_j for each j . It then chooses an ABE key pair $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$ and computes $n-1$ ABE keys as $\text{abe.sk}_j \leftarrow \text{ABE.KeyGen}(\text{abe.msk}, C_{\text{Mixed.sk}_j})$. Next, it sends $(\text{abe.pp}, \text{Mixed.pp})$ and $\{\text{abe.sk}_j\}_{j \in [n] \setminus \{i^*+1\}}$ as the PLBE public parameters and secret keys to $\mathcal{A}$. After receiving all the keys, $\mathcal{A}$ sends its challenge message m^* to $\mathcal{B}$, and it can also make an

¹⁷Due to the fact that the reduction algorithm has to guess the index, we can only extend the current analysis to prove q -bounded PLBE (adaptive) security assuming q -bounded Mixed FE (restricted) security for constant q . However, we would like to point out that one could prove q -bounded PLBE *selective* security directly from q -bounded Mixed FE (restricted) security without any security loss.

¹⁸Basically, the reduction algorithm guesses two things: First, it guesses the index hiding challenger with which $\mathcal{A}$ interacts and wins with nonnegligible probability; second, it guesses the index on which adversary $\mathcal{A}$ queries the PLBE challenger for index encryption.

encryption query for message m on index $\tilde{i}$. Here $\mathcal{A}$ is allowed to make the encryption query either before or after the challenge query. The reduction algorithm $\mathcal{B}$ proceeds as follows. If $i \neq \tilde{i}$, $\mathcal{B}$ aborts and sends a random bit as its guess to the Mixed FE challenger. Otherwise, it responds to each query as follows. $\mathcal{B}$ encrypts message m as $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$ and sends ct to $\mathcal{A}$ as its response to the encryption query. Also, it computes ciphertext ct^* as $\text{ct}^* \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}^*, m^*)$ and sends ct^* as the challenge ciphertext to $\mathcal{A}$. Note that $\mathcal{A}$ could instead have sent its challenge query before sending the index encryption query. Also, $\mathcal{B}$ does not need to query the Mixed FE challenger for answering any query at this point as it already has ciphertexts $\text{ct}_{\text{attr}}, \text{ct}_{\text{attr}}^*$. Finally, $\mathcal{A}$ sends its guess b to $\mathcal{B}$, and $\mathcal{B}$ forwards b as its own guess.

First, note that both $\mathcal{A}$ and $\mathcal{B}$ are allowed to make at most 1 index encryption and polynomially many secret key encryption queries, respectively. Also, note that $\mathcal{B}$ sends its secret key encryption query as well as its challenge query before making key generation queries; thus $\mathcal{B}$ is an admissible adversary in the 1-bounded restricted function indistinguishability game. Next, for each query made by $\mathcal{A}$, $\mathcal{B}$ queries the Mixed FE challenger exactly once, and thus all the queries are honestly and exactly answered. Finally, note that if the Mixed FE challenger computed $\text{ct}_{\text{attr}}^*$ as a secret key FE ciphertext for function comp_{i^*} , then $\mathcal{B}$ computes ct^* as a PLBE ciphertext for index i^* ; otherwise it computes ct^* as a PLBE ciphertext for index $i^* + 1$. Thus, $\mathcal{B}$ perfectly simulates the 1-bounded index hiding security game for $\mathcal{A}$. Also, since $\mathcal{B}$ randomly guesses the challenge index i^* as well as query index i , therefore with at least $1/n(n+1)$ probability $\mathcal{B}$'s guess will be correct; thus if $\mathcal{A}$'s advantage is (nonnegligible) ϵ , then $\mathcal{B}$ breaks 1-bounded restricted function indistinguishability security with (nonnegligible) advantage $\epsilon/n(n+1)$. This completes the proof. $\square$

6.3.3. Message hiding security.

LEMMA 6.4. *If $\text{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ is a selectively secure ABE scheme per Definition A.2, then $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Dec}, \text{Enc-index})$ is a PLBE scheme satisfying the 1-bounded message hiding security property per Definition 3.5.*

Proof. Suppose there exists an adversary $\mathcal{A}$ such that $\mathcal{A}$'s advantage in the 1-bounded message hiding security game is nonnegligible. We construct an algorithm $\mathcal{B}$ that can distinguish between ABE encryptions of two different messages, therefore breaking the security of the ABE scheme.

The reduction algorithm receives 1^n from $\mathcal{A}$. It sets the parameters $\kappa, \tilde{\kappa}$ as in the construction, and starts by choosing Mixed FE parameters as $(\text{Mixed.pp}, \text{Mixed.msk}) \leftarrow \text{Mixed.Setup}(1^\lambda, 1^\kappa)$. It then computes $\text{ct}_{\text{attr}}^* \leftarrow \text{Mixed.SK-Enc}(\text{Mixed.msk}, \text{comp}_n)$ and sends to the ABE challenger $1^{\tilde{\kappa}}$ and $\text{ct}_{\text{attr}}^*$ as its challenge attribute. The ABE challenger generates a key pair $(\text{abe.pp}, \text{abe.sk})$ and sends abe.pp to $\mathcal{B}$. For $i \leq n$, $\mathcal{B}$ generates Mixed FE secret keys as $\text{Mixed.sk}_i \leftarrow \text{Mixed.KeyGen}(\text{Mixed.msk}, i)$ and sends $C_{\text{Mixed.sk}_i}$ as a predicate query to the ABE challenger and receives back secret key abe.sk_i . Next, it sends $(\text{abe.pp}, \text{Mixed.pp})$ and $\{\text{abe.sk}_i\}_{i \leq n}$ as the PLBE public parameters and secret keys to $\mathcal{A}$. After receiving all the keys, $\mathcal{A}$ makes a single index encryption query (m, j) to $\mathcal{B}$. $\mathcal{B}$ answers it by computing ciphertexts $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.SK-Enc}(\text{Mixed.msk}, \text{comp}_j)$ and $\text{ct} \leftarrow \text{ABE.Enc}(\text{abe.pp}, \text{ct}_{\text{attr}}, m)$ and sends ct to $\mathcal{A}$ as its response. $\mathcal{A}$ also sends two challenge messages (m_0^*, m_1^*) to $\mathcal{B}$. $\mathcal{B}$ then forwards (m_0^*, m_1^*) as its challenge messages to ABE challenger. Next, $\mathcal{B}$ forwards the challenge ciphertext ct^* it receives from ABE challenger to $\mathcal{A}$. Note that $\mathcal{A}$

could instead have sent its challenge query before sending the index encryption query. In that case, the reduction algorithm simply answers that first. Finally, $\mathcal{A}$ sends its guess b to $\mathcal{B}$, and $\mathcal{B}$ forwards b as its own guess.

First, note that the challenge attribute $\text{ct}_{\text{attr}}^*$ on each predicate $(C_{\text{Mixed.sk}_i})$ queried made by $\mathcal{B}$ evaluates to 0, with all but negligible probability. This follows from the correctness condition of the Mixed FE system as $\text{ct}_{\text{attr}}^*$ encrypts function comp_n and for all $i \leq n$, $\text{comp}_n(i) = 0$, and thus decrypting $\text{ct}_{\text{attr}}^*$ using Mixed.sk_i outputs 0. With all-but-negligible probability, reduction algorithm $\mathcal{B}$ is therefore an admissible adversary in the ABE security game. Thus, $\mathcal{B}$ perfectly simulates the 1-bounded¹⁹ message hiding security game for $\mathcal{A}$. As a result, if $\mathcal{A}$'s advantage is nonnegligible, then $\mathcal{B}$ breaks ABE security with nonnegligible advantage. This completes the proof. $\square$

7. A new LWE toolkit. In section 2, we defined the notion of lattice trapdoors along with certain well-sampledness properties. Recall that using lattice trapdoors, it is easy to compute a preimage $\mathbf{U}$ of any matrix $\mathbf{Z}$ with respect to a matrix $\mathbf{A}$ given the trapdoor information generated while sampling $\mathbf{A}$. For any matrix $\mathbf{U}$ sampled as above, we say $\mathbf{U}$ targets $\mathbf{A}$ to matrix $\mathbf{Z}$. Now the well-sampledness properties (at a high level) state that a matrix sampled using the **TrapGen** algorithm looks uniformly random when not given the trapdoor information, and a preimage matrix that targets to a random matrix looks like a matrix with entries drawn from a Gaussian distribution.

In this section, we introduce certain enhanced security properties for lattice trapdoors that will be useful later in proving security of our Mixed FE system. We also provide a generic construction of lattice trapdoors that achieves these enhanced properties from any lattice trapdoor scheme that satisfies the well-sampledness properties described above. At a very high level, the enhanced security properties state the following: (1) For any matrix $\mathbf{A}$ that is sampled using the **TrapGen** algorithm, all those rows of $\mathbf{A}$ which are only used to target random rows look like random rows themselves (when not given the trapdoor information). (2) Two preimage matrices $\mathbf{U}_0, \mathbf{U}_1$ that target a matrix $\mathbf{A}$ to different matrices $\mathbf{Z}_0, \mathbf{Z}_1$ should look indistinguishable to any adversary even when the adversary is given those rows of $\mathbf{A}$ where $\mathbf{Z}_0, \mathbf{Z}_1$ are identical. We point out that due to technical constraints in the proof of property (2) we chose to define the enhanced properties w.r.t. matrices instead of vectors.

7.1. Enhanced lattice trapdoors. Let $(\text{EnTrapGen}, \text{EnSamplePre})$ be a pair of randomized algorithms with the following syntax:

$\text{EnTrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, T_{\mathbf{A}})$. The trapdoor generation algorithm takes as input n, m, q and outputs a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with a trapdoor $T_{\mathbf{A}}$.

$\text{EnSamplePre}(\mathbf{A}, T_{\mathbf{A}}, \sigma, \mathbf{z}) \rightarrow \mathbf{u}$. The preimage sampling algorithm takes as input a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ together with its trapdoor $T_{\mathbf{A}}$, a target vector $\mathbf{z} \in \mathbb{Z}_q^n$, and a parameter σ .²⁰ It outputs $\mathbf{u} \in \mathbb{Z}_q^m$ such that $\mathbf{A} \cdot \mathbf{u}^T = \mathbf{z}^T$ and $\|\mathbf{u}\| \leq \sqrt{m} \cdot \sigma$.

We require these algorithms to satisfy the following properties. These properties are captured via security games between a challenger and a computationally bounded adversary.

¹⁹We would like to point out that the current construction actually gives a PLBE scheme that satisfies the q -bounded message hiding security property for arbitrary q , i.e., the number of queries need not be bounded, as long as the ABE scheme is not q -bounded selectively secure.

²⁰As before, the preimage sampling algorithm could be easily generalized to generate preimages of matrices instead of vectors by independently running the **EnSamplePre** algorithm on each column of the matrix. Throughout this work, we overload the notation by directly giving matrices $\mathbf{U} \in \mathbb{Z}_q^{n \times k}$ (for any k) as inputs to the **SamplePre** algorithm.

Notation. First, we introduce some more notation. We start by defining a matrix rearrangement procedure **Arrange** which takes as input dimensions n_1, n_2, m , and two matrices $\mathbf{A} \in \mathbb{Z}_q^{n_1 \times m}, \mathbf{B} \in \mathbb{Z}_q^{n_2 \times m}$, and an *ordered* set $S \subseteq [n_1 + n_2]$ of size n_1 , and it outputs a larger combined matrix $\mathbf{C} \in \mathbb{Z}_q^{(n_1+n_2) \times m}$. Concretely, $\mathbf{C} = \text{Arrange}(1^{n_1}, 1^{n_2}, 1^m, \mathbf{A}, \mathbf{B}, S)$. The rearrangement procedure is defined as follows.

Let $S = \{i_1, i_2, \dots, i_{n_1}\}$, where $i_j < i_k$ for every $1 \leq j < k \leq n_1$. Similarly, let $\bar{S} = ([n_1 + n_2] \setminus S) = \{i'_1, i'_2, \dots, i'_{n_2}\}$ denote the (ordered) complement of set S . Now matrix $\mathbf{C}$ is obtained by appending rows of matrices $\mathbf{A}$ and $\mathbf{B}$ as follows: for $j \in [n_1]$, $\mathbf{C}[i_j] = \mathbf{A}[j]$, and for $j \in [n_2]$, $\mathbf{C}[i'_j] = \mathbf{B}[j]$. For simplicity of notation, we drop the dimensions n_1, n_2, m as explicit inputs to **Arrange** procedure throughout this section whenever clear from the context.

Additionally, we define a matrix restriction procedure **Restrict** which takes as input dimensions n, m , and matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, and an *ordered* set $S \subseteq [n]$ of size ℓ , and it outputs a smaller matrix $\mathbf{C} \in \mathbb{Z}_q^{\ell \times m}$. Concretely, $\mathbf{C} = \text{Restrict}(1^n, 1^m, \mathbf{A}, S)$. The restriction procedure is defined as follows.

Let $S = \{i_1, i_2, \dots, i_n\}$, where $i_j < i_k$ for every $1 \leq j < k \leq \ell$. Now matrix $\mathbf{C}$ is obtained by removing rows of matrix $\mathbf{A}$ which do not lie in set S . Formally, for $j \in [\ell]$, $\mathbf{C}[j] = \mathbf{A}[i_j]$. As before, for simplicity of notation, we drop the dimensions n, m as inputs whenever clear from the context.

7.1.1. Row removal property. The first property we introduce is called the *row removal* property. It is defined via an interactive security game between the challenger and an adversary. In the game, the adversary specifies matrix dimensions n, m , a set of k ($\leq n$) indices, which represent the “target” set, and the adversary must distinguish between the following scenarios.

In the first scenario, the challenger chooses an $n \times m$ matrix $\mathbf{A}$ with a trapdoor, and sends $\mathbf{A}$ to the adversary. The adversary then participates in a query phase. For each query, the adversary sends a set of k target vectors. The challenger responds by outputting a matrix $\mathbf{U}$ such that for each index i in the target set, $\mathbf{U}$ maps the i th row of $\mathbf{A}$ to one of the target vectors. The matrix $\mathbf{U}$ maps the remaining rows of $\mathbf{A}$ to uniformly random vectors.


In the second scenario, the challenger chooses a $k \times m$ matrix $\mathbf{A}$ with a trapdoor, extends $\mathbf{A}$ to dimension $n \times m$ by attaching uniformly random rows, and sends this extended matrix to the adversary. Next, the adversary sends queries, each query consisting of k target vectors. The challenger outputs a matrix $\mathbf{U}$ such that $\mathbf{U}$ maps the i th row of $\mathbf{A}$ to the i th target vector.²¹

DEFINITION 7.1 (row removal property). *Fix any function $q : \mathbb{N} \rightarrow \mathbb{N}$ and parameter $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$. A pair of trapdoor generation algorithms $\text{LT} = (\text{EnTrapGen}, \text{EnSamplePre})$ is said to satisfy the (q, σ) -row removal property if for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$*

$$\text{pr}_{\text{LT}, \mathcal{A}}^{\text{row-rem}, q, \sigma}(\lambda) = \Pr \left[1 \leftarrow \text{Expt}_{\text{LT}, \mathcal{A}}^{\text{row-rem}, q, \sigma}(\lambda) \right] \leq 1/2 + \text{negl}(\lambda),$$

where $\text{Expt}_{\text{LT}, \mathcal{A}}^{\text{row-rem}, q, \sigma}(\cdot)$ is as defined in Figure 4.

²¹Although one might observe some weak resemblance between our row removal property and lattice trapdoor properties used in [ALS16, BF11], we would like to point out that after a closer inspection we observe that our row removal property is different.

FIG. 4. Experiment $\text{Expt}_{\text{LT}, \mathcal{A}}^{\text{row-rem}, q, \sigma}$.

A weaker row removal property is one where, in each query, the adversary is restricted to choosing a set S of size $n - 1$ during setup phase, and now during query phase it must make preimage queries of the form $(1^t, \mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t})$. A different way to represent the set S in this case is by a single index $i \in [n]$ such that $\{i\} = [n] \setminus S$. We call this property the single row removal property. In our construction, we will first show a scheme that satisfies the single row removal property and then, via a simple hybrid argument, we show that the single row removal property implies the row removal property.

7.1.2. Target switching property. Next, we introduce the *target switching* property. For any target $\mathbf{Z}$, if we choose a matrix $\mathbf{B}$ with a trapdoor and output *only* the preimage of $\mathbf{Z}$ with respect to $\mathbf{B}$, then this preimage looks like a random low-norm matrix. The target switching property is an extension of this property and is captured via a security game between a challenger and an adversary. In this game, the challenger specifies the matrix dimensions n, m and a set $S \subseteq [n]$ of size k . The challenger chooses an $n \times m$ matrix $\mathbf{B}$ and sends the rows of $\mathbf{B}$ corresponding to the set S . It also chooses a challenge bit b which is used in the query phase.

Next, the adversary is allowed polynomially many queries. In each query, the adversary specifies two matrices, $\mathbf{Z}_0, \mathbf{Z}_1$, such that for every index $i \in S$, the i th rows of $\mathbf{Z}_0$ and $\mathbf{Z}_1$ are identical. The challenger outputs a matrix $\mathbf{U}$ such that for every $i \in S$, $\mathbf{U}$ maps the i th row of $\mathbf{B}$ to the i th row of $\mathbf{Z}_0$ (which is equal to the i th row of $\mathbf{Z}_1$). For the remaining indices $i \notin S$, $\mathbf{U}$ *approximately* maps the i th row of $\mathbf{A}$ to the i th row of $\mathbf{Z}_b$. Intuitively, since the adversary does not have the rows indexed by $\bar{S}$, the challenger can switch the targets from rows of $\mathbf{Z}_0$ to $\mathbf{Z}_1$.

DEFINITION 7.2 (target switching property). *Fix any function $q : \mathbb{N} \rightarrow \mathbb{N}$, noise distribution family $\{\chi(\lambda)\}_{\lambda \in \mathbb{N}}$, and parameter $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$. A pair of trapdoor gener-*

ation algorithms $\text{LT} = (\text{EnTrapGen}, \text{EnSamplePre})$ is said to satisfy the (q, χ, σ) -target switching property if for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q = q(\lambda)$, $\chi = \chi(\lambda)$, $\sigma = \sigma(\lambda)$,

$$\text{pr}_{\text{LT}, \mathcal{A}}^{\text{switch}, q, \chi, \sigma}(\lambda) = \Pr \left[1 \leftarrow \text{Expt}_{\text{LT}, \mathcal{A}}^{\text{switch}, q, \chi, \sigma}(\lambda) \right] \leq 1/2 + \text{negl}(\lambda),$$

where $\text{Expt}_{\text{LT}, \mathcal{A}}^{\text{switch}, q, \chi, \sigma}(\cdot)$ is as defined in Figure 5.




FIG. 5. $\text{Expt}_{\text{LT}, \mathcal{A}}^{\text{switch}, q, \chi, \sigma}(\cdot)$.

As before, we will introduce a weaker notion called the single target switching property, where in each query the adversary is restricted to outputting only a single index $i \in [n]$, and $\mathbf{Z}_0$ and $\mathbf{Z}_1$ must agree on all indices $j \neq i$. We will first show that our construction satisfies the single target switching property, and then, via a hybrid argument, we show that single target switching implies the target switching property.

7.2. Our construction of enhanced lattice trapdoors. Let $q : \mathbb{N} \rightarrow \mathbb{N}, \sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ be functions, and let $\text{LT} = (\text{TrapGen}, \text{SamplePre})$ be a pair of algorithms that satisfy q -well-sampledness of matrix (Definition 2.7), (q, σ) -preimage sampling (Definition 2.8). We will construct enhanced lattice trapdoors $\text{LT}_{\text{en}} = (\text{EnTrapGen}, \text{EnSamplePre})$ using LT as follows. The construction is reminiscent of the trapdoor extension algorithms of [ABB10, CHKP10].

$\text{EnTrapGen}(1^n, 1^m, q) \rightarrow (\mathbf{A}, T_{\mathbf{A}})$. The trapdoor generation algorithm samples two matrices, $\mathbf{A}_1, \mathbf{A}_2$, of dimensions $n \times \lceil m/2 \rceil$ and $n \times \lfloor m/2 \rfloor$ as follows:

$$\begin{aligned} (\mathbf{A}_1, T_{\mathbf{A}_1}) &\leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q), \\ (\mathbf{A}_2, T_{\mathbf{A}_2}) &\leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q). \end{aligned}$$

It appends both these matrices columnwise to obtain a larger matrix as $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$, and it sets the trapdoor $T_{\mathbf{A}}$ to be the combined trapdoor information $T_{\mathbf{A}} = (T_{\mathbf{A}_1}, T_{\mathbf{A}_2})$.

$\text{EnSamplePre}(\mathbf{A}, T_{\mathbf{A}}, \mathbf{Z}, \sigma) \rightarrow \mathbf{U}$. The preimage sampling algorithm takes as input $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$, trapdoor $T_{\mathbf{A}} = (T_{\mathbf{A}_1}, T_{\mathbf{A}_2})$, parameter σ , and matrix $\mathbf{Z} \in$

$\mathbb{Z}_q^{n \times k}$. It chooses a uniformly random matrix $\mathbf{W} \leftarrow \mathbb{Z}_q^{n \times k}$ and sets $\mathbf{Y} = \mathbf{Z} - \mathbf{W}$. Next, it computes matrices $\mathbf{U}_1 \in \mathbb{Z}_q^{\lceil m/2 \rceil \times k}$, $\mathbf{U}_2 \in \mathbb{Z}_q^{\lfloor m/2 \rfloor \times k}$ as

$$\begin{aligned}\mathbf{U}_1 &\leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}), \\ \mathbf{U}_2 &\leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}).\end{aligned}$$

Finally, it computes the final output matrix $\mathbf{U} \in \mathbb{Z}_q^{m \times k}$ by rowwise appending matrices $\mathbf{U}_1$ and $\mathbf{U}_2$. Concretely, $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$.

Correctness. Correctness follows directly from the correctness of LT .

7.3. Proving security of LT_{en} . We will now prove that our enhanced trapdoor sampling scheme satisfies the preimage sampling, row removal, and target switching properties. First, we show that it satisfies the preimage sampling property if the underlying trapdoor scheme satisfies the preimage sampling property.

THEOREM 7.3. *Fix any functions $q : \mathbb{N} \rightarrow \mathbb{N}$ and $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$. If LT satisfies (q, σ) -preimage sampling (Definition 2.8), then LT_{en} also satisfies (q, σ) -preimage sampling.*

Proof sketch. This follows directly from our construction. The preimage sampling requires that the preimage of a uniformly random matrix $\mathbf{Z}$ looks like a Gaussian sample with parameter σ . In our construction, the preimage of a random matrix $\mathbf{Z}$ consists of preimages of $\mathbf{W}$ and $\mathbf{Z} - \mathbf{W}$, where $\mathbf{W}$ is uniformly random. Since $\mathbf{Z}$ is random, so is $\mathbf{Z} - \mathbf{W}$. As a result, using the preimage sampling property of LT , we can argue that these two preimages look like two matrices drawn from $\mathcal{D}_{\mathbb{Z}, \sigma}^{\lceil m/2 \rceil \times k}$ and $\mathcal{D}_{\mathbb{Z}, \sigma}^{\lfloor m/2 \rfloor \times k}$, respectively. $\square$

7.3.1. Row removal property. Now we prove that our trapdoor sampling scheme satisfies the row removal property. Formally, we prove the following.

THEOREM 7.4. *Fix any functions $q : \mathbb{N} \rightarrow \mathbb{N}$ and $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$. If LT satisfies (q, σ) -preimage sampling (Definition 2.8) and q -well-sampledness of the matrix (Definition 2.7), then LT_{en} also satisfies the (q, σ) -single row removal property (Definition 7.1).*

Proof. Our proof follows from a sequence of hybrid experiments. We start by defining a sequence of hybrid experiments such that the first and last experiments correspond to the original row removal security game when the challenger chooses its challenge bit b to be 0 and 1, respectively. To complete the proof we show that the adversary's advantage must be negligible between any two consecutive hybrids.

We now define hybrids H_x for $x \in \{0, 1, \dots, 12\}$. In all the hybrid experiments below, we set $q = q(\lambda)$ and $\sigma = \sigma(\lambda)$. Also, below in each successive hybrid step, we only describe the modifications. Later in Appendix C.1, we provide the detailed hybrids.

Hybrid H_0 . This corresponds to the original game (per Definition 7.1, with the single row removal restriction) with $b = 0$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:

- It first chooses matrices $(\mathbf{B}_1, T_{\mathbf{B}_1}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
- It also chooses a vector $\mathbf{p} \leftarrow \mathbb{Z}_q^m$ and sets matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as $\mathbf{A} = \text{Arrange}(\mathbf{B}, \mathbf{p}, [n] \setminus \{i\})$.

- (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- (a) It chooses $\mathbf{W} \leftarrow \mathbb{Z}_q^{(n-1) \times t}$ and samples $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{B}_1, T_{\mathbf{B}_1}, \sigma, \mathbf{W})$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ (which is equal to $\mathbf{C} - \mathbf{W}$) and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_1 . In this experiment, the challenger chooses $\mathbf{U}_1$ to be a random Gaussian matrix with parameter σ for each query.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- (a) It samples $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lceil m/2 \rceil \times t}$.

Hybrid H_2 . In this hybrid, the challenger chooses $\mathbf{B}_1$ uniformly at random, instead of choosing it using TrapGen . At this point, note that the left half of $\mathbf{A}$ is a uniformly random matrix.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
- (a) It first chooses $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{(n-1) \times \lceil m/2 \rceil}$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$.
 - (b) It sets $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.

Hybrid H_3 . This hybrid involves syntactic changes. The challenger chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$ and derives $\mathbf{B}_1$ by removing the i th row of $\mathbf{A}_1$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
- (a) It first chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{\lceil m/2 \rceil}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.

Hybrid H_4 . In this hybrid, the challenger chooses the left half of $\mathbf{A}$ using TrapGen .

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
- (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.

Hybrid H_5 . In this hybrid, the challenger chooses $\mathbf{U}_1$ using SamplePre for each query.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- (a) It chooses $\mathbf{W}' \leftarrow \mathbb{Z}_q^{n \times t}$, sets $\mathbf{W} = \text{Restrict}(\mathbf{W}', [n] \setminus \{i\})$, and samples $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.

Hybrid H_6 . This hybrid represents a syntactic change in which the challenger, for each query, chooses $\mathbf{Y}$ as a uniformly random matrix, and sets $\mathbf{W} = \mathbf{C} - \mathbf{Y} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{Y} \leftarrow \mathbb{Z}_q^{(n-1) \times t}$ and samples $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (b) Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$ (which is equal to $\mathbf{C} - \mathbf{Y}$), chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.

Hybrid H_7 . In this hybrid experiment, the challenger chooses $\mathbf{U}_2$ from a Gaussian distribution with parameter σ .

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It samples $\mathbf{U}_2 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lfloor m/2 \rfloor \times t}$.

Hybrid H_8 . In this hybrid, the challenger chooses matrix $\mathbf{B}_2$ uniformly at random. Note that this means $\mathbf{A}_2$ is uniformly random in this hybrid.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $\mathbf{B}_2 \leftarrow \mathbb{Z}_q^{(n-1) \times \lfloor m/2 \rfloor}$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.

Hybrid H_9 . In this hybrid, the matrix $\mathbf{A}_2$ is chosen using TrapGen .

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.

Hybrid H_{10} . In this hybrid, the challenger chooses $\mathbf{U}_2$ using SamplePre for each query.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{Y}' \leftarrow \mathbb{Z}_q^{n \times t}$ and samples $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}')$.

Hybrid H_{11} . This hybrid represents a syntactic change in which the i th row of matrix $\mathbf{W}'$ is set as a difference of random vector $\mathbf{c}$ and the i th row of $\mathbf{A}_2 \cdot \mathbf{U}_2$ instead of being sampled uniformly at random directly.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (b) Next, it chooses a random vector $\mathbf{c} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{C}' = \text{Arrange}(\mathbf{C}, \mathbf{c}, [n] \setminus \{i\})$, sets $\mathbf{W}' = \mathbf{C}' - \mathbf{A}_2 \cdot \mathbf{U}_2$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.

Hybrid H_{12} . This hybrid represents a syntactic change. It corresponds to the security game in Definition 7.1 with $b = 1$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (b) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W}' \leftarrow \mathbb{Z}_q^{n \times t}$ and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W})$.
 - (b) Next, it chooses a random vector $\mathbf{c} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{C}' = \text{Arrange}(\mathbf{C}, \mathbf{c}, [n] \setminus \{i\})$, sets $\mathbf{Y}' = \mathbf{C}' - \mathbf{A}_1 \cdot \mathbf{U}_1$ (which is equal to $\mathbf{C}' - \mathbf{W}'$), and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}')$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Analysis. We will now show that any PPT adversary has at most negligible advantage in distinguishing any two consecutive hybrids. For any adversary $\mathcal{A}$, let $p_{\mathcal{A},i} : \mathbb{N} \rightarrow [0, 1]$ denote the function such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},i}(\lambda)$ is the probability that $\mathcal{A}$, on input 1^λ , outputs 1 in hybrid experiment H_i . From the definition of the hybrid experiments, it follows that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},0}(\lambda) - p_{\mathcal{A},12}(\lambda) = 2\text{pr}_{\text{LT}_{\text{en}}, \mathcal{A}}^{\text{row-rem}, q, \sigma}(\lambda) - 1$. Therefore, to show that LT_{en} satisfies the (q, σ) -row removal property, it suffices to show that for all $\mathcal{A}$ and $i \in [12]$, there exist negligible functions negl_i such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},i-1}(\lambda) - p_{\mathcal{A},i}(\lambda) \leq \text{negl}_i(\lambda)$.

LEMMA 7.5. *Assuming LT satisfies (q, σ) -preimage sampling, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_1(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},0}(\lambda) - p_{\mathcal{A},1}(\lambda) \leq \text{negl}_1(\lambda)$.*

Proof. Suppose there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},0}(\lambda) - p_{\mathcal{A},1}(\lambda) \geq \eta(\lambda)$. Moreover, let $s_{\mathcal{A}}$ denote the number of queries made by $\mathcal{A}$, and let $t_{\mathcal{A}}$ denote a bound on the number of columns in queried matrix $\mathbf{C}$ (note that the reduction algorithm is allowed to depend on the adversary, and therefore it knows $s_{\mathcal{A}}$ and $t_{\mathcal{A}}$ corresponding to $\mathcal{A}$).²² Then we can construct a reduction algorithm $\mathcal{B}$ such that $\text{pr}_{\text{LT}, \mathcal{B}}^{\text{preimg}, q, \sigma}(\lambda) \geq \eta(\lambda)$ for all $\lambda \in \mathbb{N}$.

The reduction algorithm receives n, m , index $i \in [n]$ from $\mathcal{A}$ such that $m > 2n \log q + 2\lambda$ and $\sigma > \sqrt{n \cdot \log q \cdot \log m} + \lambda$. It forwards $1^{n-1}, 1^{\lceil m/2 \rceil}, 1^{s_{\mathcal{A}} \cdot t_{\mathcal{A}}}$ to the challenger.²³ It receives $\mathbf{B}_1 \in \mathbb{Z}_q^{(n-1) \times \lceil m/2 \rceil}$ and $\tilde{\mathbf{U}} \in \mathbb{Z}_q^{\lceil m/2 \rceil \times (s_{\mathcal{A}} \cdot t_{\mathcal{A}})}$. Note that the trapdoor for $\mathbf{B}_1$ is not used in hybrid H_1 . The reduction algorithm chooses $(\mathbf{B}_2, T_{\mathbf{B}_2})$ using TrapGen , computes $\mathbf{A}$ as in H_0 (and H_1), and sends $\mathbf{A}$ to $\mathcal{A}$. The challenger also partitions $\tilde{\mathbf{U}} = [\tilde{\mathbf{U}}_1 \mid \dots \mid \tilde{\mathbf{U}}_{s_{\mathcal{A}}}]$, where each $\tilde{\mathbf{U}}_j \in \mathbb{Z}_q^{\lceil m/2 \rceil \times t_{\mathcal{A}}}$.

Next, the adversary sends queries. For the i^* th query, the adversary sends $(1^t, \mathbf{C})$ for some $t \leq t_{\mathcal{A}}$. The reduction algorithm sets $\mathbf{U}_1$ to be the first t columns of $\tilde{\mathbf{U}}_{i^*}$. It computes $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$. It sets $\mathbf{U}$ as in H_0/H_1 and sends $\mathbf{U}$ to $\mathcal{A}$.

Finally, after all queries, if $\mathcal{A}$ outputs 1, $\mathcal{B}$ guesses that $\tilde{\mathbf{U}}$ is sampled using

²²Throughout this section, we construct the nonuniform reduction algorithm, as our reduction algorithms depends on the number of queries made by the adversary as well as the size of the matrices in each query. However, we would like to point out that the reduction could be made uniform by simply guessing both of these bounds. This would result in a polynomial loss in the reduction algorithm's advantage.

²³Note that the reduction algorithm chooses admissible parameters, since $\lceil m/2 \rceil > n \log q + \lambda > (n-1) \log q + \lambda$ and $\sigma > \sqrt{n \cdot \log q \cdot \log m} + \lambda > \sqrt{(n-1) \cdot \log q \cdot \log m/2} + \lambda$.

SamplePre; otherwise it guesses that $\tilde{\mathbf{U}}$ is a random Gaussian matrix sampled with parameter σ .

Note that depending on whether $\tilde{\mathbf{U}}$ is sampled using **SamplePre** or sampled from Gaussian distribution, $\mathcal{B}$ simulates either H_0 or H_1 perfectly. As a result, $\mathcal{B}$'s advantage in the preimage sampling game is at least $\eta(\lambda)$. $\square$

LEMMA 7.6. *Assuming LT satisfies q -well-sampledness of matrix, for any adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},1}(\lambda) - p_{\mathcal{A},2}(\lambda) \leq \text{negl}_2(\lambda)$.*

Proof. Suppose there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},1}(\lambda) - p_{\mathcal{A},2}(\lambda) \geq \eta(\lambda)$. Then we can construct a reduction algorithm $\mathcal{B}$ such that $\text{pr}_{\text{LT},\mathcal{B}}^{\text{matrix},q,\sigma}(\lambda) \geq \eta(\lambda)$ for all $\lambda \in \mathbb{N}$.

The reduction algorithm receives $1^n, 1^m$, index $i \in [n]$ from $\mathcal{A}$. It forwards $1^{n-1}, 1^{\lceil m/2 \rceil}$ to the challenger.²⁴ It receives $\mathbf{B}_1$. The reduction algorithm chooses $(\mathbf{B}_2, T_{\mathbf{B}_2})$ using **TrapGen**, sets $\mathbf{A}$ as in H_1 (and H_2), and sends $\mathbf{A}$ to $\mathcal{A}$.

Next, the adversary sends queries. For each query $\mathbf{C}$, the challenger chooses $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^{\lceil m/2 \rceil \times t}$ and computes $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$. It chooses $\mathbf{p}$, sets $\mathbf{U}$ as in H_1/H_2 , and sends $\mathbf{U}$ to $\mathcal{A}$.

After all the queries, if $\mathcal{A}$ outputs 1, $\mathcal{B}$ guesses that $\mathbf{B}_1$ is sampled using **TrapGen**; otherwise it guesses that $\mathbf{B}_1$ is a uniformly random matrix. Note that, depending on whether $\mathbf{B}_1$ is sampled using **TrapGen** or sampled uniformly at random, $\mathcal{B}$ simulates either H_1 or H_2 perfectly. As a result, $\mathcal{B}$'s advantage in the matrix well-sampledness game is at least $\eta(\lambda)$. $\square$

LEMMA 7.7. *For any adversary $\mathcal{A}$, $p_{\mathcal{A},2}(\lambda) = p_{\mathcal{A},3}(\lambda)$.*

Since the only changes from H_2 to H_3 are syntactical, it follows that any adversary has identical behavior in both hybrids.

LEMMA 7.8. *Assuming LT satisfies q -well-sampledness of matrix, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_4(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},3}(\lambda) - p_{\mathcal{A},4}(\lambda) \leq \text{negl}_4(\lambda)$.*

This proof is identical to the proof of Lemma 7.6, except that the reduction algorithm must send $1^n, 1^{\lceil m/2 \rceil}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}$.

LEMMA 7.9. *Assuming LT satisfies (q, σ) -preimage sampling, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_5(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},4}(\lambda) - p_{\mathcal{A},5}(\lambda) \leq \text{negl}_5(\lambda)$.*

This proof is identical to the proof of Lemma 7.5, except that the reduction algorithm must send $1^n, 1^{\lceil m/2 \rceil}, 1^{s_{\mathcal{A}} \cdot t_{\mathcal{A}}}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}, 1^{s_{\mathcal{A}} \cdot t_{\mathcal{A}}}$.

LEMMA 7.10. *For any adversary $\mathcal{A}$, $p_{\mathcal{A},5}(\lambda) = p_{\mathcal{A},6}(\lambda)$.*

Note that the distributions in H_5 and H_6 are identical. In hybrid H_5 , the challenger chooses $\mathbf{W}' \leftarrow \mathbb{Z}_q^{n \times m}$, derives $\mathbf{W}$ from $\mathbf{W}'$ by removing the i th row, and sets $\mathbf{Y} = \mathbf{C} - \mathbf{W}$. In hybrid H_6 , it chooses $\mathbf{Y} \leftarrow \mathbb{Z}_q^{(n-1) \times m}$, sets $\mathbf{W} = \mathbf{C} - \mathbf{Y}$, and sets $\mathbf{W}'$ to be a matrix extended from $\mathbf{W}$ by inserting a random vector at row i . The distribution of $(\mathbf{W}, \mathbf{W}', \mathbf{Y})$ is identical in both hybrid experiments.

LEMMA 7.11. *Assuming LT satisfies (q, σ) -preimage sampling, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_7(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},6}(\lambda) - p_{\mathcal{A},7}(\lambda) \leq \text{negl}_7(\lambda)$.*

²⁴Note that the reduction algorithm chooses admissible parameters, since $\lceil m/2 \rceil > n \log q + \lambda > (n-1) \log q + \lambda$.

This proof is identical to the proof of Lemma 7.5, except that the reduction algorithm must send $1^{n-1}, 1^{\lfloor m/2 \rfloor}, 1^{s_{\mathcal{A}} \cdot t_{\mathcal{A}}}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}, 1^{s_{\mathcal{A}} \cdot t_{\mathcal{A}}}$. It uses the challenger's response for setting $\mathbf{B}_2, \mathbf{U}_2$ and chooses the remaining components by itself.

LEMMA 7.12. *Assuming LT satisfies q -well-sampledness of matrix, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_8(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},7}(\lambda) - p_{\mathcal{A},8}(\lambda) \leq \text{negl}_8(\lambda)$.*

This proof is identical to the proof of Lemma 7.6, except that the reduction algorithm must send $1^{n-1}, 1^{\lfloor m/2 \rfloor}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}$. It uses the challenger's response for setting $\mathbf{B}_2$ and chooses the remaining components by itself.

LEMMA 7.13. *Assuming LT satisfies q -well-sampledness of matrix, for any adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_9(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},8}(\lambda) - p_{\mathcal{A},9}(\lambda) \leq \text{negl}_9(\lambda)$.*

This proof is identical to the proof of Lemma 7.6, except that the reduction algorithm must send $1^n, 1^{\lfloor m/2 \rfloor}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}$. It uses the challenger's response for setting $\mathbf{A}_2$ and chooses the remaining components by itself.

LEMMA 7.14. *Assuming LT satisfies (q, σ) -preimage sampling, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_{10}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},9}(\lambda) - p_{\mathcal{A},10}(\lambda) \leq \text{negl}_{10}(\lambda)$.*

This proof is identical to the proof of Lemma 7.5, except that the reduction algorithm must send $1^n, 1^{\lfloor m/2 \rfloor}$ instead of $1^{n-1}, 1^{\lceil m/2 \rceil}$. It uses the challenger's response for setting $\mathbf{A}_2, \mathbf{U}_2$ and chooses the remaining components by itself.

LEMMA 7.15. *For any adversary $\mathcal{A}$, $p_{\mathcal{A},10}(\lambda) = p_{\mathcal{A},11}(\lambda)$.*

In hybrid experiment H_{10} , the challenger chooses $\mathbf{Y}' \leftarrow \mathbb{Z}_q^{n \times m}$ and derives $\mathbf{Y}'$ by removing the i th row. It sets $\mathbf{W} = \mathbf{C} - \mathbf{Y}$, chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^m$, and constructs $\mathbf{W}'$ from $\mathbf{W}$ and $\mathbf{w}$. In hybrid H_{11} , the challenger chooses $\mathbf{Y}'$ uniformly at random, extends $\mathbf{C}'$ from $\mathbf{C}$ by inserting a random vector at the i th row, and sets $\mathbf{W}' = \mathbf{C}' - \mathbf{Y}'$. As a result, $(\mathbf{Y}', \mathbf{W}')$ are identically distributed in both hybrids. The remaining components in the hybrids either are identical or can be derived from $\mathbf{Y}', \mathbf{W}'$.

LEMMA 7.16. *For any adversary $\mathcal{A}$, $p_{\mathcal{A},11}(\lambda) = p_{\mathcal{A},12}(\lambda)$.*

Note that the distributions in H_{11} and H_{12} are identical. The proof is identical to that of Lemma 7.10.

Using the above lemmas, it follows that the advantage of an adversary in the row removal experiment is at most $\text{negl}(\lambda)$.

THEOREM 7.17. *Fix any function $q : \mathbb{N} \rightarrow \mathbb{N}$, $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$. Assuming $\text{LT}_{\text{en}} = (\text{EnTrapGen}, \text{EnSamplePre})$ satisfies the (q, σ) -single row removal property, LT_{en} also satisfies the (q, σ) -row removal property.*

Proof. This proof follows via a simple hybrid argument, where we gradually remove the nontargeted rows from $\mathbf{A}$ one by one. Suppose $\mathcal{A}$ outputs set S of size k . We will define $n - k + 1$ hybrids $H_1, \dots, H_{n-k+1}$ as follows.

Hybrid H_i for $0 \leq i \leq n - k$. In hybrid H_i , the challenger does the following:

1. Let $(1^n, 1^m, S \subseteq [n]) \leftarrow \mathcal{A}(1^\lambda)$, and let $S = \{i_1, \dots, i_k\}$, $\bar{S} = \{i'_1, \dots, i'_{n-k}\}$. Let $S_i = S \cup \{i'_1, \dots, i'_i\} = \{\tilde{i}_1, \dots, \tilde{i}_{k+i}\}$ and $\bar{S}_i = [n] \setminus S_i = \{\tilde{i}'_1, \dots, \tilde{i}'_{n-k-i}\}$.

2. It chooses $(\mathbf{B}, T_{\mathbf{B}}) \leftarrow \text{EnTrapGen}(1^{k+i}, 1^m, q)$, $\mathbf{P} \leftarrow \mathbb{Z}_q^{(n-k-i) \times m}$.
3. It sets $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, where $\mathbf{A}[\tilde{i}_j] = \mathbf{B}[j]$ for all $j \leq k+i$, and $\mathbf{A}[\tilde{i}'_j] = \mathbf{P}[j]$ for all $j \leq n-k-i$, and it sends $\mathbf{A}$ to $\mathcal{A}$.
4. Next, the adversary sends queries. For each query $\{\mathbf{c}_j\}_{j \in S}$, the challenger first chooses $\mathbf{c}_j \leftarrow \mathbb{Z}_q^t$ for each $j \in S_i \setminus S$, and sets $\mathbf{C} \in \mathbb{Z}_q^{(k+i) \times t}$, where $\mathbf{C}[j] = \mathbf{c}_{\tilde{i}_j}$ for all $j \leq k+i$.
5. Next, it chooses $\mathbf{U} \leftarrow \text{EnSamplePre}(\mathbf{B}, T_{\mathbf{B}}, \mathbf{C}, \sigma)$ and sends $(\mathbf{A}, \mathbf{U})$.
6. Finally, after all queries, the adversary $\mathcal{A}$ outputs a bit b' .

Note that H_0 and H_{n-k} correspond to $b = 0$ and $b = 1$ in the row removal experiment. For any adversary $\mathcal{A}$, let $p_{\mathcal{A},i}(\lambda)$ denote the probability of $\mathcal{A}$ outputting 1 in hybrid H_i . We will show that for any adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that $p_{\mathcal{A},i}(\lambda)$ and $p_{\mathcal{A},i+1}(\lambda)$ differ by at most $\text{negl}(\lambda)$.

LEMMA 7.18. *Fix any index $i \in \{0, 1, \dots, n-k-1\}$. Assuming LT_{en} satisfies the single row removal property, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},i}(\lambda) - p_{\mathcal{A},i+1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. Suppose there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},i}(\lambda) - p_{\mathcal{A},i+1}(\lambda) \geq \eta(\lambda)$. We can use $\mathcal{A}$ to build a reduction algorithm $\mathcal{B}$ that can break the single row removal property with advantage $\eta(\cdot)$.

The reduction algorithm first receives $1^n, 1^m, S = \{i_1, \dots, i_k\}$. It defines $S_{i+1} = S \cup \{i'_1, \dots, i'_{i+1}\} = \{\tilde{i}_1, \dots, \tilde{i}_{k+i+1}\}$, and lets $\text{indx} \in [k+i+1]$ be the index such that $\tilde{i}_{\text{indx}} = i'_{i+1}$. The reduction algorithm sends $1^{k+i+1}, 1^m$ and indx .²⁵ It receives $\mathbf{B}$ from the challenger. The reduction algorithm sets $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ such that for each $j \leq k+i+1$, $\mathbf{A}[\tilde{i}_j] = \mathbf{B}[j]$, and the remaining rows are chosen uniformly at random.

For each query $\{\mathbf{c}_j\}_{j \in S}$, the reduction algorithm chooses vectors $\{\mathbf{c}_{i'_1}, \dots, \mathbf{c}_{i'_i}\}$ uniformly at random from $\mathbb{Z}_q^t$ and sends $\{\mathbf{c}_{\tilde{i}_j}\}_{j \in [k+i+1], j \neq \text{indx}}$ to the challenger. The challenger sends $\mathbf{U}$ to the reduction algorithm. The reduction algorithm forwards $\mathbf{U}$ to $\mathcal{A}$. Finally, after all the queries, $\mathcal{B}$ outputs the adversary's final output bit.

Clearly, the reduction algorithm simulates either H_i or H_{i+1} , depending on the challenger's output, and therefore the advantage of $\mathcal{B}$ is $p_{\mathcal{A},i}(\lambda) - p_{\mathcal{A},i+1}(\lambda)$. $\square$

7.3.2. Target switching property. Now we prove that our trapdoor sampling scheme satisfies the target switching property. Formally, we prove the following.

THEOREM 7.19. *Fix any functions $q : \mathbb{N} \rightarrow \mathbb{N}$ and $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$, and error distribution family $\{\chi(\lambda)\}_{\lambda}$. If LT satisfies q -well-distributedness of matrix (Definition 2.7), (q, σ) -well-distributedness of preimage (Definition 2.8), and $\text{LWE-sp}_{(d,q,\sigma,\chi)}$ holds (LWE with short public vectors; Assumption 3) where $d(\lambda) = 6\lambda \log q(\lambda)$, then LT_{en} satisfies the (q, σ, χ) -single target switching property.*

Proof. To prove the above theorem, we first define a sequence of hybrid games where the first game is the single target switching security game, and in the last game the adversary's advantage is exactly 0. Later we show that the adversary's advantage in any two consecutive hybrid games is negligible. For simplicity of notation, we will let $d = d(\lambda)$, $q = q(\lambda)$, $\sigma = \sigma(\lambda)$, and $\chi = \chi(\lambda)$. Below in each successive hybrid game, we only describe the modifications. Later, in Appendix C.2, we provide the detailed hybrid games.

²⁵Note that the reduction algorithm chooses admissible parameters, since $m > n \log q + \lambda > (k+i+1) \log q + \lambda$ and $\sigma > \sqrt{n \cdot \log q \cdot \log m} + \lambda > \sqrt{(k+i+1) \cdot \log q \cdot \log m} + \lambda$.

Hybrid H_0 . This corresponds to the single target switching security game.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.
 - (b) Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$ and sends $[\mathbf{B}_1 \mid \mathbf{B}_2]$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W} \leftarrow \mathbb{Z}_q^{n \times t}$, computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W})$.
 - (b) It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{E} = \text{Arrange}(\mathbf{0}^{(n-1) \times t}, \mathbf{e}, [n] \setminus \{i\})$.
 - (c) Next, it sets $\mathbf{Y} = \mathbf{Z}_b - \mathbf{A}_1 \cdot \mathbf{U}_1 + \mathbf{E}$ (which is equal to $\mathbf{Z}_b - \mathbf{W} + \mathbf{E}$) and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.
 - (d) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. $\mathcal{A}$ outputs its guess b' .

Hybrid H_1 . In this hybrid experiment, the challenger sets $\mathbf{U}_1$ to be a Gaussian matrix for each query.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (a) It computes $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lceil m/2 \rceil \times t}$.

Hybrid H_2 . In this hybrid experiment, the challenger sets $\mathbf{A}_1$ to be a uniformly random matrix (that is, sampled without a trapdoor).

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.

Hybrid H_3 . This hybrid is a syntactic change. Here, we express $\mathbf{Y}$ in terms of $\mathbf{B}_1$ and the i th row of $\mathbf{A}_1$. Note that the i th row of $\mathbf{A}_1$ is used only for computing the i th row of $\mathbf{Y}$.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (b) It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{Z}'_b = \text{Restrict}(\mathbf{Z}_b, [n] \setminus \{i\})$.
 - (c) Next, it sets $\mathbf{Y}' = \mathbf{Z}'_b - \mathbf{B}_1 \cdot \mathbf{U}_1$, $\mathbf{y} = \mathbf{Z}_b[i] - \mathbf{A}_1[i] \cdot \mathbf{U}_1 + \mathbf{e}$, and $\mathbf{Y} = \text{Arrange}(\mathbf{Y}', \mathbf{y}, [n] \setminus \{i\})$. It then computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.

Hybrid H_4 . In this hybrid experiment, the challenger sets the i th row of $\mathbf{Y}$ to be a uniformly random vector.

2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (c) Next, it sets $\mathbf{Y}' = \mathbf{Z}'_b - \mathbf{B}_1 \cdot \mathbf{U}_1$, $\mathbf{y} \leftarrow \mathbb{Z}_q^t$, and $\mathbf{Y} = \text{Arrange}(\mathbf{Y}', \mathbf{y}, [n] \setminus \{i\})$. It then computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.

Analysis. We will now analyze the adversary's advantage in the single target switching experiment. Let $p_{\mathcal{A}, l}(\lambda)$ denote the probability of $\mathcal{A}$ guessing correctly

(i.e., $b' = b$) at the end of hybrid experiment H_l . We will show that for every PPT adversary $\mathcal{A}$ and $l \in [4]$ there exist negligible functions $\text{negl}_l(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},l-1} - p_{\mathcal{A},l} \leq \text{negl}_l(\lambda)$.

LEMMA 7.20. *Assuming LT satisfies (q, σ) -preimage sampling, for any adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_1(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},0} - p_{\mathcal{A},1} \leq \text{negl}_1(\lambda)$.*

The proof of this lemma is similar to the proof of Lemma 7.5.

LEMMA 7.21. *Assuming LT satisfies q -well-sampledness of matrix, for any adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},1} - p_{\mathcal{A},2} \leq \text{negl}_2(\lambda)$.*

The proof of this lemma is similar to the proof of Lemma 7.6.

LEMMA 7.22. *For any adversary $\mathcal{A}$, $p_{\mathcal{A},2} = p_{\mathcal{A},3}$.*

Note that the only differences in H_2 and H_3 are syntactic changes with respect to matrix $\mathbf{Y}$. As a result, the distributions in the two hybrids are identical.

LEMMA 7.23. *If $\text{LWE-sp}_{(d,q,\sigma,\chi)}$ holds (Assumption 3) where $d(\lambda) = 6\lambda \log q(\lambda)$, then for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_4(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $p_{\mathcal{A},3} - p_{\mathcal{A},4} \leq \text{negl}_4(\lambda)$.*

Proof. Suppose, on the contrary, there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that $p_{\mathcal{A},3} - p_{\mathcal{A},4} \geq \eta(\lambda)$ for all $\lambda \in \mathbb{N}$. We will construct a reduction algorithm $\mathcal{B}$ such that $\text{Adv}_{\mathcal{B}}^{\text{LWE-sp},d,q,\sigma,\chi}(\lambda) \geq \eta(\lambda)$ for all $\lambda \in \mathbb{N}$.

The reduction algorithm first receives $1^n, 1^m, i \in [n], \mathbf{Z}_0, \mathbf{Z}_1$ from the adversary $\mathcal{A}$ (such that $m > 2n \log q + 12\lambda \log q$). The reduction algorithm chooses $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{(n-1) \times m}$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$ and derives $\mathbf{B}_2$ from $\mathbf{A}_2$ by removing the i th row. It defines $\mathbf{B} = [\mathbf{B}_1 | \mathbf{B}_2]$ and sends it to the adversary. It also chooses a random bit $b \leftarrow \{0, 1\}$.

Next, the reduction algorithm receives queries from the adversary, and it uses the LWE-sp challenger to define matrices $\mathbf{U}_1$ and $\mathbf{y}$. For each query, the adversary sends two matrices, $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$, such that all their rows are equal, except the i th one. The reduction algorithm queries the LWE-sp challenger for t queries and receives $\{(\mathbf{a}_r, u_r)\}_{r \in [t]}$, where $\mathbf{a}_r \in \mathbb{Z}_q^d$ for each $r \in [t]$. It chooses $\tilde{\mathbf{a}}_r \leftarrow \mathcal{D}_{\mathbb{Z}_q, \sigma}^{\lceil m/2 \rceil - d}$ for each $r \in [m]$, $\tilde{\mathbf{s}} \leftarrow \mathbb{Z}_q^{\lceil m/2 \rceil - d}$.²⁶ Next, it sets $\mathbf{U}_1 \in \mathbb{Z}_q^{\lceil m/2 \rceil \times m}$ to be a matrix whose r th column is $[\mathbf{a}_r | \tilde{\mathbf{a}}_r]^T$ for each $r \in [m]$. It sets $\mathbf{y} \in \mathbb{Z}_q^m$, where $\mathbf{y}_r = \mathbf{Z}_b[i]_r - u_r - \tilde{\mathbf{s}} \cdot \tilde{\mathbf{a}}_r^T$.

Once $\mathbf{y}$ and $\mathbf{U}_1$ are determined, the reduction algorithm can compute $\mathbf{U}_2$ using $\mathbf{B}_1, \mathbf{U}_1, \mathbf{Z}_b, T_{\mathbf{A}_2}$. It sets $\mathbf{B}$ and $\mathbf{U}$ as in H_3/H_4 and sends $\mathbf{U}$ to $\mathcal{A}$.

Finally, after all the queries, the adversary outputs a bit b' . If $b = b'$, the reduction algorithm guesses 0 (i.e., u_r is an LWE sample); otherwise it guesses 1 (i.e., u_r is a uniformly random element).

Now note that if the LWE-sp challenger uses oracle $O_2()$, then the reduction algorithm simulates H_4 ; otherwise it simulates H_3 . Therefore, the advantage of $\mathcal{B}$ is at least $p_{\mathcal{A},3} - p_{\mathcal{A},4}$. $\square$

²⁶Since $m > 2n \log q + 12\lambda \cdot \log q$ and $d = 6\lambda \log q$, thus $\lceil m/2 \rceil - d \geq 0$. Here we would like to point out that in our target switching security game the adversary is allowed to choose the dimensions m, n ; as stated in our LWE assumption framework, however, the lattice dimensions are not chosen by the adversary. Due to this definitional inconsistency, as a reduction algorithm we always choose to attack the LWE problem for dimensions $d(\lambda) = 6\lambda \log q(\lambda)$. This could be avoided by adapting the existing definitions.

Using the above lemmas, it follows that any PPT adversary has an advantage at most $\text{negl}(\lambda)$ in the single target switching security game, since in the last hybrid game (H_4) the challenger's response is independent of bit b and thus any adversary advantage is exactly 0. $\square$

THEOREM 7.24. *Fix any function $q : \mathbb{N} \rightarrow \mathbb{N}$, $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ and distribution family $\{\chi(\lambda)\}_\lambda$. Assuming LT_{en} satisfies the (q, σ, χ) -single target switching property, LT_{en} also satisfies the (q, σ, χ) -target switching property.*

The proof of this theorem follows from a hybrid argument similar to that in the proof of Theorem 7.17.

8. Constructing 1-bounded mixed functional encryption. In this section we describe our construction of Mixed FE for input-circling branching programs with polynomial width and length. Concretely, $\mathcal{M}_\kappa = \{0, 1\}^k$ and $\mathcal{F}_\kappa$ denotes the class of input-circling branching programs with input space $\{0, 1\}^k$, width w , and length $\ell = k \cdot L$, where $\kappa = (k, w, L)$. In other words, every branching program reads each input bit L times in a circular fashion. Before describing our construction, we introduce some shorthand notation that we will use throughout this section.

8.1. Notation. Consider a set of $4\ell\lambda$ matrices $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i \in [\ell], j \in [\lambda], \beta, b \in \{0,1\}}$ and $w\ell$ matrices $\{\mathbf{P}_{i,v}\}_{i \in [\ell], v \in [w]}$, where each individual matrix lies in $\mathbb{Z}_q^{n \times m}$. For $i \in [\ell]$, let $\mathbf{D}_i$ be another matrix defined as below:

$$\mathbf{D}_i = \begin{bmatrix} \mathbf{B}_{i,0}^{(1,0)} \\ \mathbf{B}_{i,1}^{(1,0)} \\ \mathbf{B}_{i,0}^{(1,1)} \\ \vdots \\ \mathbf{B}_{i,1}^{(\lambda,1)} \\ \mathbf{P}_{i,1} \\ \vdots \\ \mathbf{P}_{i,w} \end{bmatrix} \quad \begin{array}{l} \text{The matrix } \mathbf{D}_i \text{ consists of matrices } \mathbf{B}_{i,b}^{(j,\beta)} \text{ arranged per ad-} \\ \text{joining well-defined ordering. Concretely, let } (i, j, \beta, b) \text{ be} \\ \text{the indices of any } \mathbf{B} \text{ matrix. The ordering we define is that} \end{array}$$

$$(i, j_1, \beta_1, b_1) \prec (i, j_2, \beta_2, b_2) \iff \begin{cases} j_1 < j_2, \text{ or} \\ j_1 = j_2 \wedge \beta_1 < \beta_2, \text{ or} \\ j_1 = j_2 \wedge \beta_1 = \beta_2 \\ \quad \wedge b_1 < b_2. \end{cases}$$

Thus, per our ordering $(i, 1, 0, 0), (i, 1, 0, 1), (i, 1, 1, 0), (i, 1, 1, 1), \dots, (i, \lambda, 1, 1)$ is an increasing sequence of indices. Similarly, we can define an ordering for matrices $\mathbf{P}_{i,v}$ for $v \in [w]$ (i.e., $(i, v_1) \prec (i, v_2) \iff v_1 < v_2$).

In words, matrix $\mathbf{D}_i$ is defined by rowwise appending matrices $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{j \in [\lambda], \beta, b \in \{0,1\}}$ and $\{\mathbf{P}_{i,v}\}_{v \in [w]}$ in an increasing order per the ordering “ $\prec$ ” defined above. We will use the following shorthand notation for representing the above matrix more compactly:

$$\mathbf{D}_i = \begin{bmatrix} \left\{ \mathbf{B}_{i,b}^{(j,\beta)} \right\}_{j \in [\lambda], \beta, b \in \{0,1\}} \\ \left\{ \mathbf{P}_{i,v} \right\}_{v \in [w]} \end{bmatrix}.$$

Similarly, for any (possibly empty) sets $S_1 \subseteq [\lambda] \times \{0, 1\}^2, S_2 \subseteq [w]$, we will use the shorthand

$$\mathbf{D}_i^{S_1, S_2} = \begin{bmatrix} \left\{ \mathbf{B}_{i,b}^{(j,\beta)} \right\}_{(j,\beta,b) \in S_1} \\ \left\{ \mathbf{P}_{i,v} \right\}_{v \in S_2} \end{bmatrix}$$

to represent the matrix generated by rowwise appending matrices $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S_1}$ and $\{\mathbf{P}_{i,v}\}_{v \in S_2}$ in an increasing order per the ordering “ $\prec$ ”.

8.2. Construction. In this section, we present our Mixed FE scheme. First, we provide the parameter constraints required by our correctness and security proof. For functionality indices (k, w, L) (where k denotes the input length, and $w, \ell (= k \cdot L)$ are the width and length of branching programs), the setup algorithm chooses parameters n, m, q, σ and noise distributions $\chi_{\text{big}}, \chi_{\text{last}}, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_s, \chi_{\text{lwe}}$ as follows.

Fix any $\epsilon < 1$. Let χ_1, χ_2 be a B_1 -bounded discrete Gaussian distribution with parameter σ such that $B_1 = \sqrt{m} \cdot \sigma$. Also, for any $B > 0$, let U_B denote the uniform distribution on $\mathbb{Z} \cap [-B, B]$, i.e., integers between $\pm B$. The setup algorithm chooses parameters n, m, σ, q and sets noise distributions $\chi_{\text{big}}, \chi_{\text{last}}, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_s, \chi_{\text{lwe}}$ with the following constraints. Also, we will use different versions of LWE, with different noise distributions.

- $\chi_s = \chi_{\text{appr}} = \chi_{\text{pre}} = \chi_1$ (for enhanced trapdoor sampling);
- $n = \text{poly}(\lambda)$, $\chi_{\text{lwe}} = \chi_1$, and $q \leq 2^{n^\epsilon}$ (for LWE security: $\text{LWE}_{n,q,\chi_{\text{lwe}}}$, $\text{LWE-ss}_{n,q,\chi_{\text{lwe}}}$, $\text{LWE-sp}_{6n \log q, q, \sigma_{\text{pre}}, \chi_{\text{appr}}}$);
- $m > 2(4\lambda + w) \cdot n \cdot \log q + 12n \cdot \log q$ (for enhanced trapdoor sampling);
- $\sigma > \sqrt{(4\lambda + w) \cdot n \cdot \log q \cdot \log m + \lambda}$ (for enhanced trapdoor sampling);
- $\chi_{\text{big}} = U_{\sigma_{\text{big}}}$ and $\chi_{\text{last}} = U_{\sigma_{\text{last}}}$, where $\sigma_{\text{big}} = \sigma \cdot 2^\lambda$, $\sigma_{\text{last}} = (m \cdot \sigma)^\ell \cdot 2^\lambda$ (for smudging/security);
- $\sigma_{\text{pre}} = \sigma$, $(m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^\ell \cdot (m \cdot (\sigma_{\text{last}} + \sigma_{\text{pre}})) \leq q/16$ (for correctness);
- $\sigma_s = \sigma$, $\sqrt{n} \cdot \sigma_s \cdot (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^\ell \leq q/16$ (for correctness).

First, note that it is not necessary to have distributions $\chi_{\text{lwe}}, \chi_{\text{appr}}, \chi_{\text{pre}}$ be the same distribution. Keeping all these to be different distributions will only affect the underlying assumptions to which we reduce security. One possible setting of parameters is as follows: $n = (2\lambda \cdot \ell)^{1/\epsilon}$, $m = n^{1+2\epsilon} \cdot w$, $q = 2^{n^\epsilon}$, and $\sigma = n \cdot \sqrt{w}$.

We will now describe our Mixed FE construction.

- **Setup** $(1^\lambda, (1^k, 1^w, 1^L)) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm takes as input the security parameter λ , message length k , branching program width w , and number of times it reads each bit L .²⁷ It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}$ as described above. Let $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It runs the **EnTrapGen** algorithm ℓ times as follows:

$$\forall i \in [\ell], \quad (\mathbf{M}_i, T_i) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

For each $i \in [\ell]$, it interprets matrix $\mathbf{M}_i$ as $4\lambda + w$ matrices with dimensions $n \times m$ arranged as follows:

$$\mathbf{M}_i = \begin{bmatrix} \left\{ \mathbf{B}_{i,b}^{(j,\beta)} \right\}_{j \in [\lambda], \beta \in \{0,1\}} \\ \left\{ \mathbf{P}_{i,v} \right\}_{v \in [w]} \end{bmatrix}.$$

Also, it samples $4\ell\lambda$ matrices $\{\mathbf{C}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}$ uniformly at random as $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta \in \{0,1\}$. Finally, it sets the public parameters and the master secret key as

$$\begin{aligned} \text{pp} &= (\lambda, n, m, q, k, w, L, \chi_{\text{pre}}), \\ \text{msk} &= \left(\left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{i \in [\ell], j \in [\lambda], \beta \in \{0,1\}}, \left\{ \mathbf{P}_{i,v} \right\}_{i \in [\ell], v \in [w]}, \left\{ T_i \right\}_{i \in [\ell]} \right). \end{aligned}$$

²⁷Note that here we slightly deviate from our definition as we have 3 separate functionality parameters instead of a single index. This could simply be handled by extending the Mixed FE definition to multiple indices.

- **KeyGen**($\text{msk}, x \in \{0, 1\}^k$) $\rightarrow$ **sk**. The key generation algorithm takes as input the master secret key msk and a message $x \in \{0, 1\}^k$. Let

$$\text{msk} = \left(\left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{i \in [\ell], j \in [\lambda], \beta \in \{0,1\}}, \left\{ \mathbf{P}_{i,v} \right\}_{i \in [\ell], v \in [w]}, \left\{ T_i \right\}_{i \in [\ell]} \right).$$

It chooses a secret vector $\tilde{\mathbf{s}}$ of length n as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors $\mathbf{y}^{(j)}$ of length m as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda - 1]$. Next, it sets vector $\mathbf{y}^{(\lambda)}$ as

$$\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}.$$

The key generation algorithm then chooses $2\ell\lambda$ secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$ and $2(\ell+1)\lambda$ error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ of length n and m , respectively, as

$$\begin{aligned} \forall i \in [\ell], j \in [\lambda], \beta \in \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall i \in [\ell], j \in [\lambda], \beta \in \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall j \in [\lambda], \beta \in \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

Let $\tilde{x} = x^L$, i.e., $\tilde{x}$ is a $k \cdot L$ -bit string obtained by appending string x to itself L times. Next, it computes $2(\ell+1)\lambda$ key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall i \in [\ell+1], j \in [\lambda], \beta \in \{0, 1\}, \\ \mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1. \end{cases} \end{aligned}$$

Finally, it outputs the secret key **sk** as

$$\text{sk} = \left(x, \left\{ \mathbf{t}_i^{(j,\beta)} \right\}_{i \in [\ell+1], j \in [\lambda], \beta \in \{0,1\}} \right).$$

- **Enc**(**pp**) $\rightarrow$ **ct**. The encryption algorithm takes as input the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$. It first chooses a λ -bit string $\text{tag} \leftarrow \{0, 1\}^\lambda$ and 2ℓ random short matrices $\{\mathbf{U}_{i,b}\}_{i,b}$ as

$$\forall i \in [\ell], b \in \{0, 1\}, \quad \mathbf{U}_{i,b} \leftarrow \chi_{\text{pre}}^{m \times m}.$$

Finally, it outputs the ciphertext **ct** as

$$\text{ct} = \left(\text{tag}, \left\{ \mathbf{U}_{i,b} \right\}_{i \in [\ell], b \in \{0,1\}} \right).$$

- **SK-Enc**(**msk**, **BP**) $\rightarrow$ **ct**. The secret key encryption algorithm takes as input the master secret key **msk** and an input-circling branching program **BP**. Let

$$\begin{aligned} \text{msk} &= \left(\left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{i \in [\ell], j \in [\lambda], \beta \in \{0,1\}}, \left\{ \mathbf{P}_{i,v} \right\}_{i \in [\ell], v \in [w]}, \left\{ T_i \right\}_{i \in [\ell]} \right), \\ \text{BP} &= \left(\left\{ \pi_{i,b} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc} \in [w], \text{rej} \in [w] \right). \end{aligned}$$

It first chooses a λ -bit string $\text{tag} \leftarrow \{0, 1\}^\lambda$ and samples $8\ell\lambda$ matrices $\{\mathbf{D}_{i,b}^{(j,\beta)}, \tilde{\mathbf{D}}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}$ as follows:

$$\forall i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}, \quad \begin{aligned} \mathbf{D}_{i,b}^{(j,\beta)} &= \begin{cases} \mathbf{C}_{i,b}^{(j,\beta)} & \text{if } \beta = \text{tag}_j \text{ and } b = 0, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise,} \end{cases} \\ \tilde{\mathbf{D}}_{i,b}^{(j,\beta)} &= \begin{cases} \mathbf{C}_{i,b}^{(j,\beta)} & \text{if } \beta = \text{tag}_j \text{ and } b = 1, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise.} \end{cases} \end{aligned}$$

In the above cases, we use “ $(\leftarrow \mathbb{Z}_q^{n \times m})$ ” to denote the operation of uniformly sampling a dimension $n \times m$ matrix in $\mathbb{Z}_q$. Note that here the sampling is performed *uniformly and independently each time*. Next, the algorithm samples w matrices $\{\mathbf{P}_{\ell+1,v}\}_{v \in [w]}$ for the top level and $2w\ell$ error matrices $\{\mathbf{E}_{i,v}, \tilde{\mathbf{E}}_{i,v}\}_{i \in [\ell], v \in [w]}$ as follows:

$$\begin{aligned} \forall v \in [w], \quad \mathbf{P}_{\ell+1,v} &= \begin{cases} \mathbf{0}^{n \times m} & \text{if } v = \text{rej}, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise,} \end{cases} \\ \forall i \in [\ell], v \in [w], \quad \mathbf{E}_{i,v} &\leftarrow \chi_{\text{appr}}^{n \times m}, \tilde{\mathbf{E}}_{i,v} \leftarrow \chi_{\text{appr}}^{n \times m}. \end{aligned}$$

The algorithm then sets $2w\ell$ matrices $\{\mathbf{Q}_{i,v}, \tilde{\mathbf{Q}}_{i,v}\}_{i \in [\ell], v \in [w]}$ as follows:

$$\forall i \in [\ell], v \in [w], \quad \begin{aligned} \mathbf{Q}_{i,v} &= \mathbf{P}_{i+1, \pi_{i,0}(v)} + \mathbf{E}_{i,v}, \\ \tilde{\mathbf{Q}}_{i,v} &= \mathbf{P}_{i+1, \pi_{i,1}(v)} + \tilde{\mathbf{E}}_{i,v}. \end{aligned}$$

Next, for $i \in [\ell]$, we use matrices $\mathbf{M}_i, \mathbf{W}_i, \tilde{\mathbf{W}}_i$ to represent the following $(4\lambda + w)n \times m$ dimension matrices:

$$\begin{aligned} \mathbf{M}_i &= \begin{bmatrix} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{j \in [\lambda], \beta, b \in \{0,1\}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, \quad \mathbf{W}_i = \begin{bmatrix} \{\mathbf{D}_{i,b}^{(j,\beta)}\}_{j \in [\lambda], \beta, b \in \{0,1\}} \\ \{\mathbf{Q}_{i,v}\}_{v \in [w]} \end{bmatrix}, \\ \tilde{\mathbf{W}}_i &= \begin{bmatrix} \{\tilde{\mathbf{D}}_{i,b}^{(j,\beta)}\}_{j \in [\lambda], \beta, b \in \{0,1\}} \\ \{\tilde{\mathbf{Q}}_{i,v}\}_{v \in [w]} \end{bmatrix}. \end{aligned}$$

Now, the secret key encryption algorithm runs the EnSamplePre to compute 2ℓ short matrices $\{\mathbf{U}_{i,b}\}_{i,b}$ as

$$\forall i \in [\ell], \quad \begin{aligned} \mathbf{U}_{i,0} &\leftarrow \text{EnSamplePre}(\mathbf{M}_i, T_i, \sigma_{\text{pre}}, \mathbf{W}_i), \\ \mathbf{U}_{i,1} &\leftarrow \text{EnSamplePre}(\mathbf{M}_i, T_i, \sigma_{\text{pre}}, \tilde{\mathbf{W}}_i). \end{aligned}$$

Finally, it outputs the ciphertext ct as

$$\text{ct} = \left(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \right).$$

- $\text{Dec}(\text{sk}, \text{ct}) \rightarrow \{0, 1\}$. The decryption algorithm takes as input a secret key sk and a ciphertext ct . Let

$$\text{sk} = \left(x, \{\mathbf{t}_i^{(j,\beta)}\}_{i \in [\ell+1], j \in [\lambda], \beta \in \{0,1\}} \right), \quad \text{ct} = \left(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}} \right).$$

We will assume the algorithm knows the LWE modulus q (i.e., for instance, the public parameters could be included in the secret keys). Let $\tilde{x} = x^L$, i.e., $\tilde{x}$ is a $k \cdot L$ -bit string obtained by appending string x to itself L times. The decryption algorithm computes the following:

$$\mathbf{z} = \sum_{j=1}^{\lambda} \sum_{i=1}^{\ell+1} \left(\mathbf{t}_i^{(j, \text{tag}_j)} \cdot \prod_{\alpha=i}^{\ell} \mathbf{u}_{\alpha, \tilde{x}_{\alpha}} \right).$$

Finally, if $\|\mathbf{z}\| \leq q/8$, it outputs 0; otherwise it outputs 1.

THEOREM 8.1. *Assuming the trapdoor scheme LT_{en} satisfies the (q, σ_{pre}) -well-sampledness of the preimage, the (q, σ_{pre}) -row removal property, the $(q, \chi_{\text{lwe}}, \sigma_{\text{pre}})$ -target switching property, assuming that assumptions $\text{LWE}_{n,q,\chi_{\text{lwe}}}$, $\text{LWE-ss}_{n,q,\chi_{\text{lwe}}}$, and $\text{LWE-sp}_{6n \log q, \sigma_{\text{pre}}, \chi_{\text{appr}}}$ hold (where $n, m, q, \sigma_{\text{pre}}, \chi_{\text{lwe}}, \chi_{\text{appr}}$ are defined as in the construction), for any PPT adversary $\mathcal{A}$ that outputs $(1^k, 1^w, 1^L)$ such that the parameter constraints as provided in the construction are satisfied, then there exist negligible functions $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ such that for every $\lambda \in \mathbb{N}$, $\mathcal{A}$'s advantage in the 1-bounded restricted function indistinguishability security (see Definition 5.2) and 1-bounded restricted accept indistinguishability (see Definition 5.4) is at most $\text{negl}_1(\lambda)$ and $\text{negl}_2(\lambda)$, respectively.*

Remark 8.2 (extending to r -bounded security). We would like to point out that the above construction can be naturally extended to achieve r -bounded security for any a priori fixed polynomial r . To understand the modification, we will look ahead to the security proof. Specifically, we will focus on the importance of the λ -bit string tag chosen during encryption. During the proof, we crucially rely on the tag strings tag and tag^* (the first chosen for answering the encryption query, and the second used to answer the challenge query) being distinct at at least one index. Since they are chosen uniformly at random each time, thus we know that $\text{tag} \neq \text{tag}^*$ with probability $1 - \frac{1}{2^\lambda}$. Now if the challenger has to answer r encryption queries instead of just 1, then the modification we consider is to increase the alphabet size of tags such that the tag strings chosen during all encryption queries and the challenge query are distinct at at least one index. (Note that this would also mean that we will have to likewise increase the number of underlying matrices chosen and extend the trapdoor sampling procedure appropriately.) More formally, we will now sample tag strings as a uniformly random $2r^2$ -ary string of length- λ (i.e., $\text{tag} \leftarrow \{1, \dots, 2r^2\}^\lambda$). With this modification we can argue that, with all but negligible probability over the choice of tag strings $\text{tag}_1, \dots, \text{tag}_r$ and tag^* , there exists an index $i \leq \lambda$ such that the i th elements in all these tag strings are (pairwise) distinct. With this guarantee, the current proof could be extended to argue r -bounded security.

8.3. Correctness. We will prove that the Mixed FE scheme described above satisfies the correctness property. Our correctness proof is divided into two parts. First, we show that if ct is a Mixed FE encryption of branching program BP , then given any secret key sk_x , the decryption algorithm outputs $\text{BP}(x)$ with all-but-negligible probability. Second, we show that if ct is a normal FE ciphertext, then given any secret key sk_x , the decryption algorithm outputs 1 with all-but-negligible probability.

LEMMA 8.3. *For every $\lambda, k, w, L \in \mathbb{N}$, for every length $k \cdot L$ and width w input-circling branching program BP with input space $\{0, 1\}^k$, input $x \in \{0, 1\}^k$, the follow-*

ing holds:

$$\Pr \left[\text{Dec}(\text{sk}_x, \text{ct}) = \text{BP}(x) : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, (1^k, 1^w, 1^L)); \\ \text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x); \text{ct} \leftarrow \text{SK-Enc}(\text{msk}, \text{BP}) \end{array} \right] \geq 1 - \text{negl}_2(\lambda),$$

where $\text{negl}_2(\cdot)$ is a negligible function.

Proof. Here and throughout, whenever we say matrices/terms corresponding to (j, β) th strands, then we mean the corresponding matrix/term with the superscript (j, β) . Also, when we refer to the matrices/terms along the strands selected by some tag string tag , then we mean all those matrices/terms with superscripts (j, tag_j) for $j \in [\lambda]$.

Now recall that the setup algorithm chooses matrices $\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}, \mathbf{P}_{i,v}$ for $i \in [\ell]$, $j \in [\lambda]$, $\beta, b \in \{0, 1\}$, $v \in [w]$. Here all matrices $\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{P}_{i,v}$ for any particular value i (i.e., any fixed level) are sampled along with trapdoor information. Now for any input $x \in \{0, 1\}^k$, the key generation algorithm chooses vectors $\mathbf{y}^{(j)}, \tilde{\mathbf{s}}$ such that $\tilde{\mathbf{s}}$ is short and $\sum_j \mathbf{y}^{(j)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}$. It also samples secret vectors $\mathbf{s}_i^{(j,\beta)}$ and error vectors $\mathbf{e}_i^{(j,\beta)}$ and computes the secret key components $\mathbf{t}_i^{(j,\beta)}$ in the special way as described in the construction. Now the mixed FE encryption algorithm samples a λ -bit tag string tag , and it uses the trapdoor information to target $\mathbf{B}_{i,b}^{(j,\beta)}$ matrices to their corresponding $\mathbf{C}_{i,b}^{(j,\beta)}$ matrices only along the strands selected by the tag string tag . Additionally, it also targets the program matrices $\mathbf{P}_{i,v}$ at each level to their counterparts in the next level per the branching program state transition function. For proving correctness we simply show that the final program matrix reached after decryption is either short or random depending upon outcome of the evaluation, and the $\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}$ matrices get canceled at each step, and the error terms are appropriately bounded.

We start by introducing some notation useful for the correctness proof.

- st_i : the state of BP after i steps when evaluated on input x .

$$\mathbf{t}_i^{(j,\beta)} \stackrel{\text{def}}{=} \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} & \text{if } i = \ell + 1: \end{cases}$$

the *error-free* secret key components, i.e., secret key vectors without adding error vectors $\mathbf{e}_i^{(j,\beta)}$.

- $\Delta_i^{(j)} \stackrel{\text{def}}{=} \sum_{\gamma=1}^i \mathbf{t}_\gamma^{(j,\text{tag}_j)} \cdot \prod_{\alpha=\gamma}^{i-1} \mathbf{U}_{\alpha,\tilde{x}_\alpha}$: the partial sum computed during decryption after using the first i components of the secret key along only the (j, tag_j) th strand.

$$\tilde{\Delta}_i^{(j)} \stackrel{\text{def}}{=} \begin{cases} \mathbf{s}_i^{(j,\text{tag}_j)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\text{tag}_j)} + \mathbf{y}^{(j)} \cdot \prod_{\alpha=1}^{i-1} \mathbf{U}_{\alpha,\tilde{x}_\alpha} & \text{if } i \leq \ell, \\ \mathbf{y}^{(j)} \cdot \prod_{\alpha=1}^\ell \mathbf{U}_{\alpha,\tilde{x}_\alpha} & \text{if } i = \ell + 1: \end{cases}$$

the expected sum during decryption in absence of *errors* after using the first i components of the secret key along the (j, tag_j) th strand.

- $\Delta_i \stackrel{\text{def}}{=} \sum_{j=1}^\lambda \Delta_i^{(j)}$ and $\tilde{\Delta}_i = \sum_{j=1}^\lambda \tilde{\Delta}_i^{(j)}$.
- $\text{err}_i^{(j)} \stackrel{\text{def}}{=} \Delta_i^{(j)} - \tilde{\Delta}_i^{(j)}$ and $\text{err}_i \stackrel{\text{def}}{=} \Delta_i - \tilde{\Delta}_i$.
- $\Gamma_i \stackrel{\text{def}}{=} \mathbf{P}_{1,1} \cdot \prod_{\alpha=1}^i \mathbf{U}_{\alpha,\tilde{x}_\alpha}$: the matrix denoting partial branching program evaluation after i decryption steps, i.e., equal to the Δ_i term ignoring the $\sum_{j=1}^\lambda \mathbf{s}_i^{(j,\text{tag}_j)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\text{tag}_j)}$ blinding component, and short secret $\tilde{\mathbf{s}}$ multiplied.

Observe that the decryption algorithm computes $\Delta_{\ell+1}$ and tests whether it is close to zero or not. We start by proving that for all i, j , the error term $\mathbf{err}_i^{(j)}$ is small and bounded. This would help us in arguing that for every i , $\mathbf{err}_i$ is also small, thereby giving us that matrices $\Delta_{\ell+1}$ and $\tilde{\Delta}_{\ell+1}$ are very close to each other as well. Combining this with the fact that $\tilde{\Delta}_{\ell+1}$ is either a random matrix or a short matrix, depending upon the output $\mathbf{BP}(x)$, we get that the sum $\Delta_{\ell+1}$ computed by the decryption algorithm is close to zero if $\mathbf{BP}(x) = 0$; otherwise it is a random vector with large entries.

CLAIM 8.4. *There exists a negligible function $\text{negl}(\cdot)$ such that*

$$\begin{aligned} \forall i \in [\ell], j \in [\lambda], \quad & \|\mathbf{err}_i^{(j)}\| \leq (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^i, \\ \forall j \in [\lambda], \quad & \|\mathbf{err}_{\ell+1}^{(j)}\| \leq (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^\ell \cdot (m \cdot (\sigma_{\text{last}} + \sigma_{\text{pre}})) \end{aligned}$$

with probability $1 - \text{negl}(\lambda)$.

Proof. We prove the above claim by inducting on the levels i . Our proof is insensitive to the choice of strand index j ; thus for the purposes of this proof, it could be fixed to an arbitrary value.

Base case ($i = 1$). Note that $\Delta_1^{(j)} = \mathbf{t}_1^{(j, \text{tag}_j)} = \mathbf{s}_1^{(j, \text{tag}_j)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j, \text{tag}_j)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j, \text{tag}_j)}$, where $\mathbf{e}_1^{(j, \text{tag}_j)}$ is a short error vector drawn from χ_{big}^m . Also, we have that $\tilde{\Delta}_1^{(j)} = \mathbf{s}_1^{(j, \text{tag}_j)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j, \text{tag}_j)} + \mathbf{y}^{(j)}$ by definition. Thus, we get that

$$\|\mathbf{err}_1^{(j)}\| = \|\Delta_1^{(j)} - \tilde{\Delta}_1^{(j)}\| = \|\mathbf{e}_1^{(j, \text{tag}_j)}\| \leq \sqrt{m} \cdot \sigma_{\text{big}}$$

with all-but-negligible probability. This completes the proof of the base case. For the induction step, we assume that the above claim holds for i^* and show that it holds for $i^* + 1$ as well.

Induction step. We know that $\Delta_{i^*+1}^{(j)} = \Delta_{i^*}^{(j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} + \mathbf{t}_{i^*+1}^{(j, \text{tag}_j)}$. Since $\Delta_{i^*}^{(j)} = \tilde{\Delta}_{i^*}^{(j)} + \mathbf{err}_{i^*}^{(j)}$, we get that

$$\begin{aligned} \Delta_{i^*+1}^{(j)} &= (\tilde{\Delta}_{i^*}^{(j)} + \mathbf{err}_{i^*}^{(j)}) \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} + \mathbf{t}_{i^*+1}^{(j, \text{tag}_j)} \\ &= \tilde{\Delta}_{i^*}^{(j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} + \mathbf{t}_{i^*+1}^{(j, \text{tag}_j)} + \mathbf{err}_{i^*}^{(j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}. \end{aligned}$$

Now, from our construction we have that

$$\begin{aligned} \mathbf{s}_{i^*}^{(j, \text{tag}_j)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} &= \mathbf{s}_{i^*}^{(j, \text{tag}_j)} \cdot \mathbf{C}_{i^*, \tilde{x}_{i^*}}^{(j, \text{tag}_j)} \\ \Rightarrow \mathbf{s}_{i^*}^{(j, \text{tag}_j)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} + \tilde{\mathbf{t}}_{i^*+1}^{(j, \text{tag}_j)} &= \mathbf{s}_{i^*+1}^{(j, \text{tag}_j)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j, \text{tag}_j)} \\ \Rightarrow \tilde{\Delta}_{i^*}^{(j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} + \tilde{\mathbf{t}}_{i^*+1}^{(j, \text{tag}_j)} &= \tilde{\Delta}_{i^*+1}^{(j)}. \end{aligned}$$

Combining the fact that $\mathbf{t}_{i^*+1}^{(j, \text{tag}_j)} = \tilde{\mathbf{t}}_{i^*+1}^{(j, \text{tag}_j)} + \mathbf{e}_{i^*+1}^{(j, \text{tag}_j)}$ with the above equations, we get that, with all-but-negligible probability, the following holds:

$$\mathbf{err}_{i^*+1}^{(j)} = \Delta_{i^*+1}^{(j)} - \tilde{\Delta}_{i^*+1}^{(j)} = \mathbf{e}_{i^*+1}^{(j, \text{tag}_j)} + \mathbf{err}_{i^*}^{(j)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}$$

$$\begin{aligned}
\Rightarrow \|\mathbf{err}_{i^*+1}^{(j)}\| &\leq \|\mathbf{e}_{i^*+1}^{(j, \text{tag}_j)}\| + \|\mathbf{err}_{i^*}^{(j)}\| \cdot \|\mathbf{U}_{i^*, \tilde{x}_{i^*}}\| \\
&\leq \sqrt{m} \cdot \sigma^* + (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^{i^*} \cdot (m \cdot \sigma_{\text{pre}}) \\
&\leq (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^{i^*} \cdot (m \cdot (\sigma^* + \sigma_{\text{pre}})),
\end{aligned}$$

where $\sigma^* = \sigma_{\text{big}}$ if $i^* < \ell$; otherwise $\sigma^* = \sigma_{\text{last}}$. This completes the proof of the above claim. $\square$

From the above claim, we get that for every $j \in [\lambda]$, $\|\mathbf{err}_{\ell+1}^{(j)}\| \leq (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^\ell \cdot (m \cdot (\sigma_{\text{last}} + \sigma_{\text{pre}}))$. Thus, by the triangle inequality we can claim that (with all-but-negligible probability)

$$\|\mathbf{err}_{\ell+1}\| \leq \lambda \cdot (m \cdot (\sigma_{\text{big}} + \sigma_{\text{pre}}))^\ell \cdot (m \cdot (\sigma_{\text{last}} + \sigma_{\text{pre}})) \leq q/16.$$

Next, we show that $\mathbf{\Gamma}_i$ is close to $\mathbf{P}_{i+1, \text{st}_i}$. In other words, the partial branching program evaluation is correct.

CLAIM 8.5. *For all $i \in \{0, \dots, \ell\}$, $\|\mathbf{\Gamma}_i - \mathbf{P}_{i+1, \text{st}_i}\| \leq (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^i$ with probability $1 - \text{negl}(\lambda)$, where $\text{negl}(\cdot)$ is a negligible function.*

Proof. We prove the above claim by inducting on the levels i .

Base case ($i = 0$). Note that $\mathbf{\Gamma}_0$ is simply equal to $\mathbf{P}_{1,1}$ as starting state $\text{st}_0 = 1$. Thus, we get that

$$\|\mathbf{\Gamma}_0 - \mathbf{P}_{1, \text{st}_0}\| = 0.$$

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

Induction step. We know that $\mathbf{\Gamma}_{i^*} = \mathbf{\Gamma}_{i^*-1} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}$. Recall that, per our construction, $\mathbf{U}_{i^*, \tilde{x}_{i^*}}$ targets $\mathbf{P}_{i^*, \text{st}_{i^*-1}}$ to $\mathbf{P}_{i^*+1, \text{st}_{i^*}} + \mathbf{Err}_1$, where $\mathbf{Err}_1$ is an $n \times m$ matrix sampled uniformly from $\chi_{\text{appr}}^{n \times m}$. Concretely, this gives that

$$\mathbf{P}_{i^*, \text{st}_{i^*-1}} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} = \mathbf{P}_{i^*+1, \text{st}_{i^*}} + \mathbf{Err}_1, \text{ where } \|\mathbf{Err}_1\| \leq m \cdot \sigma_{\text{appr}}.$$

By our inductive hypothesis, we have that $\mathbf{\Gamma}_{i^*-1} = \mathbf{P}_{i^*, \text{st}_{i^*-1}} + \mathbf{Err}_2$, where $\|\mathbf{Err}_2\| \leq (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{i^*-1}$. Thus, we can rewrite matrix $\mathbf{\Gamma}_{i^*}$ as follows:

$$\begin{aligned}
\mathbf{\Gamma}_{i^*} &= (\mathbf{P}_{i^*, \text{st}_{i^*-1}} + \mathbf{Err}_2) \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}} \\
&= \mathbf{P}_{i^*+1, \text{st}_{i^*}} + (\mathbf{Err}_1 + \mathbf{Err}_2 \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}).
\end{aligned}$$

Now we have that

$$\begin{aligned}
\|\mathbf{Err}_1 + \mathbf{Err}_2 \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}\| &\leq m \cdot \sigma_{\text{appr}} + (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{i^*-1} \cdot m \cdot \sigma_{\text{pre}} \\
&\leq (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{i^*}.
\end{aligned}$$

Thus, the claim follows. $\square$

From the above claim, we get that (with all-but-negligible probability) $\|\mathbf{\Gamma}_\ell - \mathbf{P}_{\ell+1, \text{st}_\ell}\| \leq (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^\ell$. Next, we show that $\tilde{\Delta}_{\ell+1}$ has low norm if the output of branching program is 0; otherwise it is not upper-bounded with all-but-negligible probability.

CLAIM 8.6. *There exists a negligible function $\text{negl}(\cdot)$ such that*

$$\left\| \tilde{\Delta}_{\ell+1} \right\| = \begin{cases} \leq q/16 & \text{if } \text{BP}(x) = 0, \\ \geq q/4 & \text{if } \text{BP}(x) = 1 \end{cases}$$

with probability $1 - \text{negl}(\lambda)$.

Proof. We know that $\tilde{\Delta}_{\ell+1}$ could be written as follows:

$$\tilde{\Delta}_{\ell+1} = \left(\sum_{j=1}^{\lambda} \mathbf{y}^{(j)} \right) \cdot \prod_{\alpha=1}^{\ell} \mathbf{U}_{\alpha, \tilde{x}_{\alpha}}.$$

Since $(\sum_{j=1}^{\lambda} \mathbf{y}^{(j)}) = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}$, this gives that $\tilde{\Delta}_{\ell+1} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} \cdot \prod_{\alpha=1}^{\ell} \mathbf{U}_{\alpha, \tilde{x}_{\alpha}} = \tilde{\mathbf{s}} \cdot \mathbf{\Gamma}_{\ell}$. Using Claim 8.5, we get that $\mathbf{\Gamma}_{\ell} = \mathbf{P}_{\ell+1, \text{st}_{\ell}} + \mathbf{Err}$, where $\|\mathbf{Err}\| \leq (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{\ell}$. Also, we know that $\mathbf{P}_{\ell+1, \text{rej}} = \mathbf{0}^{n \times m}$, and $\mathbf{P}_{\ell+1, \text{acc}}$ is a uniformly random $n \times m$ matrix. Thus, we get that with all-but-negligible probability

$$\begin{aligned} \text{BP}(x) = 0 \Rightarrow \left\| \tilde{\Delta}_{\ell+1} \right\| &= \|\tilde{\mathbf{s}} \cdot \mathbf{Err}\| \leq \|\tilde{\mathbf{s}}\| \cdot (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{\ell} \\ &\leq \sqrt{n} \cdot \sigma_s \cdot (m \cdot (\sigma_{\text{pre}} + \sigma_{\text{appr}}))^{\ell} \leq q/16, \end{aligned}$$

$$\text{BP}(x) = 1 \Rightarrow \left\| \tilde{\Delta}_{\ell+1} \right\| = \|\tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1, \text{acc}} + \tilde{\mathbf{s}} \cdot \mathbf{Err}\| \geq \|\tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1, \text{acc}}\| - q/16 \geq q/4,$$

where the last inequality follows from the fact that $\tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1, \text{acc}}$ is a uniformly random vector. This completes the proof of the above claim. $\square$

By the triangle inequality, we know that

$$\left\| \tilde{\Delta}_{\ell+1} \right\| - \|\mathbf{err}_{\ell+1}\| \leq \|\Delta_{\ell+1}\| \leq \left\| \tilde{\Delta}_{\ell+1} \right\| + \|\mathbf{err}_{\ell+1}\|.$$

Combining this with the above claims, we can conclude that with all-but-negligible probability

$$\begin{aligned} \text{BP}(x) = 0 \Rightarrow \|\Delta_{\ell+1}\| &\leq q/16 + q/16 \leq q/8, \\ \text{BP}(x) = 1 \Rightarrow \left\| \tilde{\Delta}_{\ell+1} \right\| &\geq q/4 - q/16 > q/8. \end{aligned}$$

Thus, for any input x and branching program BP , the mixed FE encryption algorithm is correct with all-but-negligible probability. This concludes the proof of Lemma 8.3. $\square$

LEMMA 8.7. *For every $\lambda, k, w, L \in \mathbb{N}$, for every length $k \cdot L$ and width w input-circling branching program BP with input space $\{0, 1\}^k$, input $x \in \{0, 1\}^k$, the following holds:*

$$\Pr \left[\text{Dec}(\text{sk}_x, \text{ct}) = 1 : \begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^{\lambda}, (1^k, 1^w, 1^L)); \\ \text{sk}_x \leftarrow \text{KeyGen}(\text{msk}, x); \text{ct} \leftarrow \text{Enc}(\text{pp}) \end{array} \right] \geq 1 - \text{negl}_1(\lambda),$$

where $\text{negl}_1(\cdot)$ is a negligible function.

Proof. Recall that the output of a normal encryption algorithm is simply independently drawn 2ℓ short Gaussian matrices $\{\mathbf{U}_{i,b}\}$. Now the decryption algorithm performs the computation

$$\mathbf{z} = \sum_{j=1}^{\lambda} \sum_{i=1}^{\ell+1} \left(\mathbf{t}_i^{(j, \text{tag}_j)} \cdot \prod_{\alpha=i}^{\ell} \mathbf{U}_{\alpha, \tilde{x}_{\alpha}} \right).$$

Here we could rewrite $\mathbf{z}$ as

$$\mathbf{z} = \sum_{j=1}^{\lambda} \left(\sum_{i=1}^{\ell-1} \left(\mathbf{t}_i^{(j, \text{tag}_j)} \cdot \prod_{\alpha=i}^{\ell} \mathbf{U}_{\alpha, \tilde{x}_\alpha} \right) + \mathbf{t}_\ell^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} + \mathbf{t}_{\ell+1}^{(j, \text{tag}_j)} \right).$$

From our construction we know that for any j ,

$$\begin{aligned} \mathbf{t}_\ell^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} + \mathbf{t}_{\ell+1}^{(j, \text{tag}_j)} &= -\mathbf{s}_{\ell-1}^{(j, \text{tag}_j)} \cdot \mathbf{C}_{\ell-1, \tilde{x}_{\ell-1}}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} + \mathbf{e}_\ell^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} + \mathbf{e}_{\ell+1}^{(j, \text{tag}_j)} \\ &\quad + \mathbf{s}_\ell^{(j, \text{tag}_j)} \cdot \left(\mathbf{B}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} - \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)} \right). \end{aligned}$$

Now note that since $\mathbf{U}_{\ell, \tilde{x}_\ell}$ is sampled independently from $\chi_{\text{pre}}^{m \times m}$ and $\mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)}$ is a uniform $n \times m$ matrix, thus the matrix $\mathbf{B}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} - \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)}$ is also a uniformly random matrix. Also, secret vector $\mathbf{s}_\ell^{(j, \text{tag}_j)}$ is a length n random vector, and thus the component $\mathbf{s}_\ell^{(j, \text{tag}_j)} \cdot \left(\mathbf{B}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)} \cdot \mathbf{U}_{\ell, \tilde{x}_\ell} - \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)} \right)$ is a random vector as well. Now since this is *independent* of all other components as neither $\mathbf{s}_\ell^{(j, \text{tag}_j)}$ nor $\mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \text{tag}_j)}$ appears in any other term in sum vector $\mathbf{z}$, thus the distribution of $\mathbf{z}$ is that of a uniformly random vector over the choice of coins used during setup, key generation, and encryption. Since we know that the ℓ_2 -norm of a random vector in $\mathbb{Z}_q^m$ is at least $q/8$ with all-but-negligible probability, therefore the claim follows. $\square$

8.4. Security proof. We now prove that the Mixed FE scheme described in section 8.2 satisfies the 1-bounded restricted function indistinguishability as well as the 1-bounded restricted accept indistinguishability security properties. Our proof is divided into two components where we first prove function indistinguishability, and later prove accept indistinguishability. Both proofs proceed via a sequence of hybrid games.

8.4.1. 1-bounded restricted function indistinguishability. Below we provide a sequence of hybrid games that we later use to argue function indistinguishability security.

Game 0 This corresponds to the original 1-bounded restricted function indistinguishability security game.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\text{BP}^{(0)}, \text{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0,1\}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$ and a λ -bit string $\text{tag}^* \leftarrow \{0, 1\}^\lambda$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = [\ell] \times [\lambda] \times \{0, 1\}^2.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as

$$\forall \alpha \in [\ell], (\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger chooses a λ -bit string $\text{tag} \leftarrow \{0, 1\}^\lambda$ and responds as follows:

- (a) Let $S = [\ell] \times [\lambda] \times \{0, 1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.

2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda - 1]$. Next, it sets vector $\mathbf{y}^{(\lambda)}$ as

$$\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$ and error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$




FIG. 6. Routine Mixed-SubEnc.

(c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\forall (i, j, \beta) \in [\ell + 1] \times [\lambda] \times \{0, 1\},$$

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)}, & i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, & 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, & i = \ell + 1. \end{cases}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.
- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 1 This is identical to the previous game, except the challenger now chooses both tags tag^* and tag at the beginning during the setup phase, and it aborts if $\text{tag}^* = \text{tag}$.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\text{BP}^{(0)}, \text{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = [\ell] \times [\lambda] \times \{0, 1\}^2.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = [\ell] \times [\lambda] \times \{0, 1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda - 1]$. Next, it sets vector $\mathbf{y}^{(\lambda)}$ as

$$\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$, error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell + 1] \times [\lambda] \times \{0, 1\}, \\ \mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)}, & i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, & 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, & i = \ell + 1. \end{cases} \end{aligned}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.
- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Notation. In all the following hybrid games, let j^* denote the smallest index in $\{1, \dots, \lambda\}$ such that $\text{tag}_{j^*}^* \neq \text{tag}_{j^*}$, i.e., $j^* = \min \{j \in [\lambda] : \text{tag}_j^* \neq \text{tag}_j\}$. Since the challenger aborts whenever $\text{tag}^* = \text{tag}$, thus j^* always exists whenever the challenger does not abort. Additionally, let $\beta^* = \text{tag}_{j^*}^*$.

Game 2 This is identical to the previous game, except the challenger, while answering a secret key query, now puts the $\tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}$ component in $\mathbf{y}^{(j^*)}$ instead of $\mathbf{y}^{(\lambda)}$, and the rest are sampled uniformly at random.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\text{BP}^{(0)}, \text{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}$.
4. Finally, it sends the public parameters $\mathbf{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\mathbf{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = [\ell] \times [\lambda] \times \{0, 1\}^2.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as

$$\forall \alpha \in [\ell], \quad (\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \mathbf{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program $\mathbf{BP}$ for encryption. The challenger responds as follows:
 - (a) Let $S = [\ell] \times [\lambda] \times \{0, 1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \mathbf{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$. Next, it sets vector $\mathbf{y}^{(j^*)}$ as

$$\mathbf{y}^{(j^*)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$, error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

(c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\forall (i, j, \beta) \in [\ell + 1] \times [\lambda] \times \{0, 1\},$$

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)}, & i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, & 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, & i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Next, we have a sequence of 4ℓ hybrid experiments, **Game 3.i***. $\{1, 2, 3, 4\}$ for $i^* = 1$ to ℓ .

Game 3.i*.1 In hybrids **Game 3.i*.1**, the $\mathbf{B}_{i,b}^{(j,\beta)}$, $\mathbf{C}_{i,b}^{(j,\beta)}$ matrices for the j^* th strands and levels $i < i^*$ are *not* sampled (at all) along with other level i matrices (i.e., (j^*, β^*) and $(j^*, 1 - \beta^*)$ strands); ciphertext components for levels $i < i^*$ are used to target only the remaining matrices; i.e., the ciphertext matrices do not target $\mathbf{B}_{i,b}^{(j,\beta)}$ matrices for $j = j^*$ and $i < i^*$ to some prespecified $\mathbf{C}_{i,b}^{(j,\beta)}$ or random matrices. Also, the first $i^* - 1$ components in each secret key are set to be uniformly random vectors, and the next component is hardwired such that correctness holds, and also some smudgeable noise is introduced into these components. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^*, \quad S^{(i)} &= ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2, \\ \forall i \geq i^*, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - 4n & \text{for } i < i^*, \\ \tilde{n} & \text{for } i \geq i^*. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\mathbf{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}^*, U_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{U_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}, U_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{U_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows.
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j^*, \beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{i^*}^{(1-\beta^*)}, \text{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs $\text{BP}, \text{BP}^{(\gamma)}$ after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*,\beta)} = \tilde{\mathbf{t}}_i^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)},$$

$$\begin{aligned} \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j^*,\beta)} = & - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_\alpha^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ & + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\mathbf{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)} + \mathbf{e}_{i^*}^{(j^*,\beta)}, \end{aligned}$$

$$\forall (i, \beta) \in ([\ell + 1] \setminus [i^*]) \times \{0, 1\},$$

$$\mathbf{t}_i^{(j^*,\beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j^*,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j^*,\beta)} + \mathbf{s}_i^{(j^*,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_\ell^{(j^*,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j^*,\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 3. $i^*.2$ This is identical to the previous game, except the $(i^* + 1)$ th key component in the j^* th strands is also hardwired. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^*, \quad S^{(i)} &= ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2, \\ \forall i \geq i^*, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - 4n & \text{for } i < i^*, \\ \tilde{n} & \text{for } i \geq i^*. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\mathbf{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\mathbf{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program $\mathbf{BP}$ for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j^*, \beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{i^*}^{(1-\beta^*)}, \mathbf{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs $\mathbf{BP}, \mathbf{BP}^{(\gamma)}$ after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$,

and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

Also, for $\beta \in \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j^*, \beta)} = \mathbf{0}^{n \times m}$, and let $\tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)}$ denote the following vector:

$$\begin{aligned} \tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)} = & - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ & + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*, \text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j^*, \beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*, \beta)}. \end{aligned}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j, \beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} \mathbf{s}_1^{(j, \beta)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j, \beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j, \beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_{\ell}}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\begin{aligned} \forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*, \beta)} &= \tilde{\mathbf{t}}_i^{(j^*, \beta)} + \mathbf{e}_i^{(j^*, \beta)}, \\ \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j^*, \beta)} &= \tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)} + \mathbf{e}_{i^*}^{(j^*, \beta)}, \end{aligned}$$

$$\begin{aligned} \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j^*, \beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &+ \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1, \text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j^*, \beta)} + \mathbf{e}_{i^*+1}^{(j^*, \beta)}, \end{aligned}$$

$$\forall (i, \beta) \in ([\ell + 1] \setminus [i^* + 1]) \times \{0, 1\},$$

$$\mathbf{t}_i^{(j^*, \beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j^*, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j^*, \beta)} + \mathbf{s}_i^{(j^*, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j^*, \beta)} + \mathbf{e}_i^{(j^*, \beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j^*, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_{\ell}}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 3.i*.3 This is identical to the previous game, except $\mathbf{B}_{i,b}^{(j, \beta)}$, $\mathbf{C}_{i,b}^{(j, \beta)}$ for strands $j = j^*$ and levels $i = i^*$ are *not* sampled along with other level i^* matrices, but instead they are sampled uniformly at random. Also, ciphertext components for level i^* are used to target only the remaining matrices. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs ($\text{BP}^{(0)}, \text{BP}^{(1)}$) to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^* + 1, \quad S^{(i)} &= ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2, \\ \forall i \geq i^* + 1, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - 4n & \text{for } i < i^* + 1, \\ \tilde{n} & \text{for } i \geq i^* + 1. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\begin{aligned} \forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) &\leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q), \\ \forall (\beta, b) \in \{0, 1\}^2, \quad \mathbf{B}_{i^*,b}^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^{n \times m}. \end{aligned}$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\begin{aligned} \text{BP}^{(\gamma)} &= \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:

- (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{i^*}^{(1-\beta^*)}, \mathbf{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, BP $^{(\gamma)}$ after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

Also, for $\beta \in \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j^*,\beta)} = \mathbf{0}^{n \times m}$, and let $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)}$ denote the following vector:

$$\begin{aligned} \tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*, \mathbf{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)}. \end{aligned}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j,\beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_{\ell}}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [i^* - 1] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*,\beta)} = \tilde{\mathbf{t}}_i^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)}.$$

For all $\beta \in \{0, 1\}$,

$$\mathbf{t}_{i^*}^{(j^*,\beta)} = \tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} + \mathbf{e}_{i^*}^{(j^*,\beta)},$$

$$\begin{aligned} \mathbf{t}_{i^*+1}^{(j^*,\beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1,\mathbf{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*+1,\tilde{x}_{i^*+1}}^{(j^*,\beta)} + \mathbf{e}_{i^*+1}^{(j^*,\beta)}, \end{aligned}$$

$$\forall (i, \beta) \in ([\ell+1] \setminus [i^*+1]) \times \{0, 1\},$$

$$\mathbf{t}_i^{(j^*,\beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j^*,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j^*,\beta)} + \mathbf{s}_i^{(j^*,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j^*,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_{\ell}}^{(j^*,\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)} & \text{if } i = \ell+1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 3. $i^*.4$ This is identical to the previous game, except the i^* th level key component in the j^* th strands is a uniformly random n length vector; i.e., all first i^* level components in the j^* th strand are random elements. Also, we no longer sample the matrices $\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}$ for strands $j = j^*$ and levels $i = i^*$ at all. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^* + 1, \quad S^{(i)} &= ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2, \\ \forall i \geq i^* + 1, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - 4n & \text{for } i < i^* + 1, \\ \tilde{n} & \text{for } i \geq i^* + 1. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\mathbf{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\begin{aligned} \mathbf{BP}^{(\gamma)} &= \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:

- (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.

2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.

- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [i^*] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{i^*}^{(1-\beta^*)}, \text{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, $\text{BP}^{(\gamma)}$ after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

Also, for $\beta \in \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j^*, \beta)} = \mathbf{0}^{n \times m}$. Next, it computes key vectors $\{\mathbf{t}_i^{(j, \beta)}\}_{i, j, \beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} \mathbf{s}_1^{(j, \beta)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j, \beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j, \beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [i^*] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*, \beta)} = \tilde{\mathbf{t}}_i^{(j^*, \beta)} + \mathbf{e}_i^{(j^*, \beta)},$$

$$\begin{aligned} \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j^*, \beta)} = & - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_\alpha^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} \\ & + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1, \text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j^*, \beta)} + \mathbf{e}_{i^*+1}^{(j^*, \beta)}, \end{aligned}$$

$$\forall (i, \beta) \in ([\ell + 1] \setminus [i^* + 1]) \times \{0, 1\},$$

$$\mathbf{t}_i^{(j^*, \beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j^*, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j^*, \beta)} + \mathbf{s}_i^{(j^*, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j^*, \beta)} + \mathbf{e}_i^{(j^*, \beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_\ell^{(j^*, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 4 This is identical to the previous game, i.e., **Game 3.ℓ.4**. For ease of exposition, we describe it in detail below.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\text{BP}^{(0)}, \text{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2.$$

Also, let $\tilde{n}_i = \tilde{n} - 4n$ for all $i \in [\ell]$.

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i, b}^{(j, \beta)}\}_{i, j, \beta, b}, \{\mathbf{P}_{i, v}\}_{i, v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \left\{ \mathbf{B}_{i, b}^{(j, \beta)} \right\}_{(j, \beta, b) \in S^{(i)}} \\ \left\{ \mathbf{P}_{i, v} \right\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i, b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}^*, U_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{U_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program **BP** for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}, U_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{U_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{\ell+1}^{(1-\beta^*)}, \text{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs **BP**, **BP**^(γ) after ℓ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*,\beta)} = \tilde{\mathbf{t}}_i^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)}.$$

For all $\beta \in \{0, 1\}$,

$$\begin{aligned} \mathbf{t}_{\ell+1}^{(j^*,\beta)} &= - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\mathbf{st}_{\ell+1}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}. \end{aligned}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Next, we have a sequence of ℓ hybrid experiments, **Game 4. i^*** for $i^* = 1$ to ℓ .

Game 4. i^* This is identical to the previous game, except the challenger uses the routine **Mixed-SubEnc*** to generate the first i^* components of both the challenge as well as the query ciphertext. This routine is similar to **Mixed-SubEnc**, except that **Mixed-SubEnc*** outputs ciphertext components that map the $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ matrices to uniformly random matrices (instead of mapping to $\{\mathbf{P}_{i^*+1,\pi(v)}\}_{v \in [w]}$ as in **Mixed-SubEnc**).

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs ($\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)}$) to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2.$$

Also, let $\tilde{n}_i = \tilde{n} - 4n$ for all $i \in [\ell]$.

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
 4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** and **Mixed-SubEnc*** routines (described in Figures 6 and 7) as

$$\forall \alpha \in [i^*],$$

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc}^* \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]} \end{array} \right),$$

$$\forall \alpha \in [\ell] \setminus [i^*],$$

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^{(\gamma)} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** and **Mixed-SubEnc*** routines (described in Figures 6 and 7) as

$$\forall \alpha \in [i^*],$$

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc}^* \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]} \end{array} \right),$$

$$\forall \alpha \in [\ell] \setminus [i^*],$$

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.




FIG. 7. Routine Mixed-SubEnc*.

2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs $\mathbf{BP}, \mathbf{BP}^{(\gamma)}$ after ℓ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext.²⁸ Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*,\beta)} = \tilde{\mathbf{t}}_i^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\mathbf{st}_{\ell+1}^{(\beta)}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}.$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.
- **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

Game 5 This is identical to the previous game, i.e., **Game 4.ℓ**. For ease of exposition, we describe it in detail below.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and descriptions of two branching programs $(\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)})$ to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\mathbf{tag}^*, \mathbf{tag} \leftarrow \{0, 1\}^\lambda$. If $\mathbf{tag}^* = \mathbf{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2.$$

Also, let $\tilde{n}_i = \tilde{n} - 4n$ for all $i \in [\ell]$.

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

²⁸Technically, in **Game 4.ℓ**, $\mathbf{P}_{\ell+1,\mathbf{st}_{\ell+1}^{(\beta)}}^{(\beta)}$ is not sampled during **Mixed-SubEnc***. For that experiment, we will assume these matrices are chosen for the first key query and used for all remaining keys.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i,j,\beta,b) \in \hat{S}$.

4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger chooses a random bit $\gamma \leftarrow \{0, 1\}$. Let

$$\text{BP}^{(\gamma)} = \left(\left\{ \pi_{i,b}^{(\gamma)} : [w] \rightarrow [w] \right\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^{(\gamma)} \in [w], \text{rej}^{(\gamma)} \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the Mixed-SubEnc^* routine (described in Figure 7) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc}^* \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]} \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the Mixed-SubEnc^* routine (described in Figure 7) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc}^* \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs $\mathbf{BP}, \mathbf{BP}^{(\gamma)}$ after ℓ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = [\ell + 1] \times ([\lambda] \setminus \{j^*\}) \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

For the first key query, it samples matrices $\mathbf{P}_{\ell+1,v}^{(\beta)}$ for $v \in [w], \beta \in \{0, 1\}$ as follows:²⁹

$$\mathbf{P}_{\ell+1, \mathbf{st}_{\ell+1}^{(\beta)}}^{(\beta)} = \begin{cases} \mathbf{0}^{n \times m} & \text{if } v = \text{rej}, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise.} \end{cases}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, \beta) \in [\ell] \times \{0, 1\}, \quad \mathbf{t}_i^{(j^*,\beta)} = \tilde{\mathbf{t}}_i^{(j^*,\beta)} + \mathbf{e}_i^{(j^*,\beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1, \mathbf{st}_{\ell+1}^{(\beta)}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}.$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

• **Guess.** The adversary finally sends the guess γ' , and wins if $\gamma' = \gamma$.

8.4.2. Indistinguishability of hybrid games in section 8.4. We will now show that the hybrid experiments described above are computationally indistinguishable. For any PPT adversary $\mathcal{A}$, let $\text{Adv}_{\mathcal{A},x}(\cdot)$ denote the advantage of $\mathcal{A}$ in **Game** x .

LEMMA 8.8. *There exists a negligible function $\text{negl}(\cdot)$ such that for any adversary $\mathcal{A}$ and $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},0}(\lambda) - \text{Adv}_{\mathcal{A},1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between **Game** 0 and **Game** 1 is that the challenger aborts if $\text{tag}^* = \text{tag}$. The probability of this event is $2^{-\lambda}$, and it is independent of the adversary's choice of (k, w, L) and $\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)}$ in the setup phase. As a result, for any adversary $\mathcal{A}$, $\text{Adv}_{\mathcal{A},0}(\lambda) - \text{Adv}_{\mathcal{A},1}(\lambda) \leq 2^{-\lambda}$. $\square$

LEMMA 8.9. *For any adversary $\mathcal{A}$ and $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},1}(\lambda) = \text{Adv}_{\mathcal{A},2}(\lambda)$.*

²⁹Recall, as defined in **Game** 4. $(\ell + 1)$, that these matrices are sampled only for the first key query, and all remaining key queries use the same matrices.

Proof. The only difference between the two hybrids is with respect to the keys. In **Game 1**, for each key query, the challenger chooses $\lambda - 1$ uniformly random vectors $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j < \lambda$ and sets $\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}$. In **Game 2**, the challenger chooses $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$ and sets $\mathbf{y}^{(j^*)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}$. Fix all $\mathbf{y}^{(j)}$ for $j \notin \{j^*, \lambda\}$ and $\tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}$. Then the following two distributions are identical:

$$\left\{ \left(\mathbf{y}^{(j^*)}, \mathbf{y}^{(\lambda)} \right) : \begin{array}{l} \mathbf{y}^{(j^*)} \leftarrow \mathbb{Z}_q^m; \\ \mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)} \end{array} \right\},$$

$$\left\{ \left(\mathbf{y}^{(j^*)}, \mathbf{y}^{(\lambda)} \right) : \begin{array}{l} \mathbf{y}^{(\lambda)} \leftarrow \mathbb{Z}_q^m; \\ \mathbf{y}^{(j^*)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)} \end{array} \right\}.$$

This implies that the distributions in **Game 1** and **Game 2** are identical. $\square$

LEMMA 8.10. *For any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3.1.1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. Let us first consider the differences between **Game 2** and **Game 3.1.1**. The setup and challenge phases are identical in both games. The post-challenge ciphertext query is also handled identically in both games. The only difference in the two games is with respect to the key queries. In particular, for each key query x , the key components $\{\mathbf{t}_1^{(j^*,\beta)}\}_{\beta \in \{0,1\}}$ are computed differently in the two games. In **Game 2**, the challenger sets $\mathbf{t}_1^{(j^*,\beta)} = -\tilde{\mathbf{y}} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} + \mathbf{s}_1^{(j^*,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j^*,\beta)} + \mathbf{e}_1^{(j^*,\beta)}$, while in **Game 3.1.1**, it sets $\mathbf{t}_1^{(j^*,\beta)} = -\tilde{\mathbf{y}} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} + \mathbf{s}_1^{(j^*,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j^*,\beta)} + \tilde{\mathbf{e}}_1^{(j^*,\beta)} + \mathbf{e}_1^{(j^*,\beta)}$. Using the smudging lemma (Lemma 2.1), since $\sigma_{\text{big}}/\sigma_{\text{lwe}} \geq 2^\lambda$, we can argue that there exists a negligible function $\text{negl}_{\text{smud}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $m \in \mathbb{N}$, $\text{SD}(\mathcal{D}_1, \mathcal{D}_2) \leq 2m \cdot \text{negl}_{\text{smud}}(\lambda)$, where

$$\mathcal{D}_1 \equiv \left\{ \left(\mathbf{e}_1^{(j^*,0)}, \mathbf{e}_1^{(j^*,1)} \right) : \mathbf{e}_1^{(j^*,\beta)} \leftarrow \chi_{\text{big}}^m \text{ for } \beta \in \{0,1\} \right\},$$

$$\mathcal{D}_2 \equiv \left\{ \left(\tilde{\mathbf{e}}_1^{(j^*,0)} + \mathbf{e}_1^{(j^*,0)}, \tilde{\mathbf{e}}_1^{(j^*,0)} + \mathbf{e}_1^{(j^*,0)} \right) : \begin{array}{l} \mathbf{e}_1^{(j^*,\beta)} \leftarrow \chi_{\text{big}}^m \text{ for } \beta \in \{0,1\}; \\ \tilde{\mathbf{e}}_1^{(j^*,\beta)} \leftarrow \chi_{\text{lwe}}^m \text{ for } \beta \in \{0,1\} \end{array} \right\}.$$

As a result, if an adversary $\mathcal{A}$ makes $q_{\text{keys}}(\lambda)$ key queries, then for any $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3.1.1}(\lambda) \leq q_{\text{keys}}(\lambda) \cdot (2m \cdot \text{negl}_{\text{smud}}(\lambda))$. $\square$

LEMMA 8.11. *For any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3.i^*.1}(\lambda) - \text{Adv}_{\mathcal{A},3.i^*.2}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The main difference in these two games is in the key generation phase. In particular, for each key, the terms $(\mathbf{t}_{i^*+1}^{(j^*,\beta)}, \mathbf{t}_{i^*+1}^{(j^*,1)})$ are computed differently in both games. In **Game 3.i*.1**, $\mathbf{t}_{i^*+1}^{(j^*,\beta)} = -\mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{C}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \mathbf{s}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*+1}}^{(j^*,\beta)} + \mathbf{e}_{i^*+1}^{(j^*,\beta)}$, while in **Game 3.i*.2**, the challenger sets $\mathbf{t}_{i^*+1}^{(j^*,\beta)}$ as $-\sum_{\alpha=1}^{i^*} (\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1,\text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*+1,\tilde{x}_{i^*+1}}^{(j^*,\beta)} + \mathbf{e}_{i^*+1}^{(j^*,\beta)}$, which is equal to $-\mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{C}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \mathbf{s}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*+1}}^{(j^*,\beta)} + \mathbf{e}_{i^*+1}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{U}_{i^*,\tilde{x}_{i^*}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{E}$. In the second equality, $\tilde{\mathbf{e}}_{i^*+1}^{(j^*,\beta)} \leftarrow \chi_{\text{lwe}}$, $\tilde{\mathbf{s}} \leftarrow \chi_{\text{ns}}^n$, $\mathbf{E}$ is sampled by Mixed-SubEnc from $\chi_{\text{appr}}^{n \times m}$. The second equality follows by substituting the value of $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} = -\sum_{\alpha=1}^{i^*-1} (\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)}$ (note that this is how $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)}$ is defined in **Game 3.i*.2**).

To prove this lemma, we will use the following fact, which follows from the smudging lemma (Lemma 2.1). Here, we use the fact that if $\mathbf{e}$ and $\mathbf{U}$ have entries bounded by $\sigma_{\text{pre}} \cdot \text{poly}(\lambda)$, then $\mathbf{e} \cdot \mathbf{U}$ can be “drowned” by a noise vector $\mathbf{e}'$ drawn from a noise distribution with parameter $\sigma_{\text{pre}} \cdot 2^\lambda$.

FACT 8.1. *Let $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{lwe}}$ be families of distributions over $\mathbb{Z}$ as defined in the construction. For any polynomials $n(\cdot), m(\cdot)$, there exists a negligible function $\text{negl}_{8.1}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $n = n(\lambda)$, $m = m(\lambda)$, $\chi_{\text{big}} = \chi_{\text{big}}(\lambda)$, $\chi_s = \chi_s(\lambda)$, $\chi_{\text{appr}} = \chi_{\text{appr}}(\lambda)$, $\chi_{\text{lwe}} = \chi_{\text{lwe}}(\lambda)$, and matrix $\mathbf{U} \in \mathbb{Z}_q^{m \times m}$ such that $\|\mathbf{U}\|_\infty \leq \sigma_{\text{pre}} \cdot n$, $\text{SD}(\mathcal{D}_1, \mathcal{D}_2) \leq \text{negl}_{8.1}(\lambda)$, where*

$$\mathcal{D}_1 = \{\mathbf{e} : \mathbf{e} \leftarrow \chi_{\text{big}}^m\}; \quad \mathcal{D}_2 = \left\{ \begin{array}{l} \mathbf{e}_1 \leftarrow \chi_{\text{big}}^m; \mathbf{s} \leftarrow \chi_s^n; \mathbf{E} \leftarrow \chi_{\text{appr}}^{n \times m}; \\ \mathbf{e}_1 + \mathbf{e}_2 + \mathbf{e}_3 : \mathbf{e}_2 = \mathbf{s} \cdot \mathbf{E}; \mathbf{e}_3' \leftarrow \chi_{\text{lwe}}^m; \\ \mathbf{e}_3 = \mathbf{e}_3' \cdot \mathbf{U} \end{array} \right\}.$$

As a result, if an adversary $\mathcal{A}$ makes $q_{\text{keys}}(\lambda)$ key queries, then for any $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A}, 3.i^*.1}(\lambda) - \text{Adv}_{\mathcal{A}, 3.i^*.2}(\lambda) \leq q_{\text{keys}}(\lambda) \cdot (2 \cdot \text{negl}_{8.1}(\lambda))$. $\square$

LEMMA 8.12. *Assuming the trapdoor generation algorithms LT_{en} satisfy the (q, σ_{pre}) -row removal property, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A}, 3.i^*.2}(\lambda) - \text{Adv}_{\mathcal{A}, 3.i^*.3}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. First, let us consider the differences between **Game 3.i*.2** and **Game 3.i*.3**.

1. Set $S^{(i^*)}$: In **Game 3.i*.2**, the challenger sets $S^{(i^*)} = [\lambda] \times \{0, 1\}^2$, while in **Game 3.i*.3**, $S^{(i^*)} = ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2$ (tag^*, tag are chosen at the start of the security game, so j^* is well defined here). Also, $\tilde{n}_{i^*} = \tilde{n} = (4\lambda + w)n$ in **Game 3.i*.2**, while $\tilde{n}_{i^*} = \tilde{n} - 4n$ in **Game 3.i*.3**.
2. $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{i=i^*}$ matrices: In **Game 3.i*.2**, the challenger chooses $(\mathbf{M}_{i^*}, T_{i^*}) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q)$, while in **Game 3.i*.3**, the challenger chooses $(\mathbf{M}_{i^*}, T_{i^*}) \leftarrow \text{EnTrapGen}(1^{\tilde{n}-4n}, 1^m, q)$. As a result, in **Game 3.i*.2**, it derives all matrices $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2}$ from $\mathbf{M}_{i^*}$. In **Game 3.i*.3**, the challenger chooses $\{\mathbf{B}_{i^*,b}^{(j^*,\beta^*)}\}_{b,\beta \in \{0,1\}}$ uniformly at random, while the remaining are derived from $\mathbf{M}_{i^*}$.
3. Ciphertexts: Since the set $S^{(i^*)}$ is different in both games, the challenge and query ciphertexts are constructed differently in both games.

Let us now discuss why the row removal property is applicable here. In particular, we will focus on $(\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*, \mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1})$. In **Game 3.i*.2**, each of these four matrices maps $[\mathbf{B}_{i^*,0}^{(j^*,0)} \mid \mathbf{B}_{i^*,1}^{(j^*,0)} \mid \mathbf{B}_{i^*,0}^{(j^*,1)} \mid \mathbf{B}_{i^*,1}^{(j^*,1)}]$ to a uniformly random matrix. To see why, let us suppose $\text{tag}_{j^*}^* = \beta^*$ and $\text{tag}_{j^*} = 1 - \beta^*$. Then

- $\mathbf{B}_{i^*,0}^{(j^*,\beta^*)} \cdot \mathbf{U}_{i^*,0}^* = \mathbf{C}_{i^*,0}^{(j^*,\beta^*)}$, the rest are mapped to random matrices;
- $\mathbf{B}_{i^*,1}^{(j^*,\beta^*)} \cdot \mathbf{U}_{i^*,1}^* = \mathbf{C}_{i^*,1}^{(j^*,\beta^*)}$, the rest are mapped to random matrices;
- $\mathbf{B}_{i^*,0}^{(j^*,1-\beta^*)} \cdot \mathbf{U}_{i^*,0} = \mathbf{C}_{i^*,0}^{(j^*,1-\beta^*)}$, the rest are mapped to random matrices;
- $\mathbf{B}_{i^*,1}^{(j^*,1-\beta^*)} \cdot \mathbf{U}_{i^*,1} = \mathbf{C}_{i^*,1}^{(j^*,1-\beta^*)}$, the rest are mapped to random matrices.

Also, it is important to note that the $\{\mathbf{C}_{i^*,b}^{(j^*,\beta^*)}\}_{b,\beta \in \{0,1\}}$ are not used for responding to key generation queries. Therefore, we can use the row removal property to remove the rows corresponding to $\mathbf{B}_{i^*,\beta}^{(j^*,b)}$ from the level i^* matrices.

Suppose, on the contrary, that there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A}, 3.i^*.2}(\lambda) - \text{Adv}_{\mathcal{A}, 3.i^*.3}(\lambda) \geq \eta(\lambda)$. We will use this adversary to build a reduction algorithm $\mathcal{B}$ that breaks the (q, σ_{pre}) -row removal property of LT_{en} .

The reduction algorithm first receives functionality index (k, w, L) from $\mathcal{A}$. Depending on the functionality index, the reduction algorithm sets $\ell = k \cdot L$, n, m , $\tilde{n} = (4\lambda + w)n$ as in **Game 3.i*.2** (and **Game 3.i*.3**) and sends these parameters to $\mathcal{A}$.

The reduction algorithm chooses tag^* , tag and defines j^* as the first index where the two tags differ. For all $i \neq i^*$, $\mathcal{B}$ defines sets $S^{(i)}$ and samples matrices (with trapdoors) $\{\{\mathbf{B}_{i,\beta}^{(j,b)}\}_{(j,\beta,b) \in S^{(i)}}, \{\mathbf{P}_{i,v}\}_{v \in [w]}, T_i\}_{i \neq i^*}$ as in **Game 3.i*.2** (and **Game 3.i*.3**). The reduction algorithm defines a set $S_{\mathcal{B}}$ which represents the set of rows that are removed in the transition between the two games. Formally, the reduction algorithm defines the sets

$$\text{pos} = \{j : b, \beta \in \{0, 1\}, (i^*, j^*, b, \beta) \text{ is at position } j \text{ in the set } \{i^*\} \times [\lambda] \times \{0, 1\}^2\},$$

$$S_{\mathcal{B}} = \bigcup_{j \in \text{pos}} \{n(j-1) + 1, n(j-1) + 2, \dots, nj\}.$$

It sends $1^{\tilde{n}}, 1^m, S_{\mathcal{B}}$ to the row removal property challenger. It receives $\mathbf{A}$ from the challenger, which it parses as

$$\mathbf{A} = \left[\begin{array}{c} \{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i^*,v}\}_{v \in [w]} \end{array} \right].$$

The reduction algorithm also chooses $(4\lambda + w - 4)$ matrices $\{\mathbf{C}_{i,b}^{(j,\beta)}\}_{i \neq i^*, j \neq j^*}$ uniformly at random from $\mathbb{Z}_q^{n \times m}$.

Next, it receives the challenge programs $\text{BP}^{(0)}, \text{BP}^{(1)}$. It chooses $\gamma \leftarrow \{0, 1\}$. For all $i \neq i^*$, it computes $(\mathbf{U}_{i,0}^*, \mathbf{U}_{i,1}^*)$ components by itself (this step is identical in both games). For $i = i^*$, it uses the row removal property challenger. It sets matrices $\mathbf{D}_b^{(j,\beta)}$ and $\tilde{\mathbf{D}}_b^{(j,\beta)}$ as follows:

$$\forall (j, \beta, b) \in ([\lambda] \setminus \{j^*\}) \times \{0, 1\}^2, \quad \begin{aligned} \mathbf{D}_b^{(j,\beta)} &= \begin{cases} \mathbf{C}_{i,b}^{(j,\beta)} & \text{if } \beta = \text{tag}_j \text{ and } b = 0, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise,} \end{cases} \\ \tilde{\mathbf{D}}_b^{(j,\beta)} &= \begin{cases} \mathbf{C}_{i,b}^{(j,\beta)} & \text{if } \beta = \text{tag}_j \text{ and } b = 1, \\ (\leftarrow \mathbb{Z}_q^{n \times m}) & \text{otherwise.} \end{cases} \end{aligned}$$

Next, it sets matrices $\{\mathbf{Q}_{i,v}\}_{v \in [w]}, \{\tilde{\mathbf{Q}}_{i,v}\}_{v \in [w]}$ as in Figure 6 and sets matrices $\mathbf{W}$ and $\tilde{\mathbf{W}}$ as

$$\mathbf{W} = \left[\begin{array}{c} \{\mathbf{D}_b^{(j,\beta)}\}_{(j,\beta,b) \in S_{\alpha}} \\ \{\mathbf{Q}_{i,v}\}_{v \in [w]} \end{array} \right], \quad \tilde{\mathbf{W}} = \left[\begin{array}{c} \{\tilde{\mathbf{D}}_b^{(j,\beta)}\}_{(j,\beta,b) \in S_{\alpha}} \\ \{\tilde{\mathbf{Q}}_{i,v}\}_{v \in [w]} \end{array} \right].$$

It sends them as queries to the row removal challenger (note that $\mathbf{C}_{i^*,b}^{(j^*,\beta)}$ is not required for defining $\mathbf{W}$ and $\tilde{\mathbf{W}}$). The challenger responds by sending $\mathbf{U}_{i^*,0}^*$ and $\mathbf{U}_{i^*,1}^*$, respectively. The reduction algorithm forwards $\{(\mathbf{U}_{i,0}^*, \mathbf{U}_{i,1}^*)\}_{i \in [\ell]}$ to the adversary. The ciphertext query is handled similarly, and the reduction algorithm receives $\{(\mathbf{U}_{i,0}, \mathbf{U}_{i,1})\}_{i \in [\ell]}$, which it forwards to $\mathcal{A}$ (the remaining ciphertext components can be computed by the reduction algorithm).

Next, the adversary sends polynomially many key queries. Note that the keys are generated in an identical manner in both games. Moreover, these keys can be generated without having $\{\mathbf{C}_{i^*,b}^{(j^*,\beta)}\}_{b,\beta}$ and the trapdoor for $\mathbf{M}_{i^*}$.

Finally, the adversary sends its guess, which the reduction algorithm forwards to the row removal property challenger. Clearly, if the row removal challenger chooses $b = 0$, then the reduction algorithm perfectly simulates Game 3. i^* .3. If the challenger chooses $b = 1$, then the reduction algorithm perfectly simulates Game 3. i^* .2 (here we use the fact that in Game 3. i^* .2, each of the matrices $\{\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*, \mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1}\}$ maps $[\mathbf{B}_{i^*,0}^{(j^*,0)} \mid \mathbf{B}_{i^*,1}^{(j^*,0)} \mid \mathbf{B}_{i^*,0}^{(j^*,1)} \mid \mathbf{B}_{i^*,1}^{(j^*,1)}]$ to a uniformly random matrix).

Therefore, the reduction algorithm has advantage at least $\eta(\cdot)$. $\square$

LEMMA 8.13. *Assuming the $(n, q, \sigma_{\text{lwe}})$ -LWE assumption holds, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A},3,i^*.3}(\lambda) - \text{Adv}_{\mathcal{A},3,i^*.4}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. In Game 3. i^* .3, for each key query x , for each $\beta \in \{0, 1\}$, the component $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} = -\sum_{\alpha=1}^{i^*-1} (\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)}$. In Game 3. i^* .4, $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} \leftarrow \mathbb{Z}_q^m$. Let $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ denote the number of keys queried by $\mathcal{A}(1^\lambda)$. To prove that these two games are computationally indistinguishable, we will define q_{keys} hybrid experiments.

Hybrid $H_{o,0}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$. In this hybrid, for the first o keys, the $\tilde{\mathbf{t}}_{i^*}^{(j^*,0)}$ components are sampled uniformly at random in the first o queries, while the $\tilde{\mathbf{t}}_{i^*}^{(j^*,1)}$ components are sampled as in Game 3. i^* .3. For the remaining $q_{\text{keys}} - o$ queries, the keys are generated as in Game 3. i^* .3 in the remaining queries.

Hybrid $H_{o,1}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$. In this hybrid, for all keys, the $\tilde{\mathbf{t}}_{i^*}^{(j^*,0)}$ components are sampled uniformly at random. For the first o queries, the $\tilde{\mathbf{t}}_{i^*}^{(j^*,1)}$ components are sampled uniformly at random, while the remaining are sampled as in Game 3. i^* .3.

Clearly, $H_{0,0}$ corresponds to Game 3. i^* .3, $H_{q_{\text{keys}},1}$ is identical to Game 3. i^* .4, and $H_{q_{\text{keys}},0} \equiv H_{0,1}$. Let $a_{\mathcal{A},i,b}(\lambda)$ denote the advantage of $\mathcal{A}$ in $H_{i,b}$.

CLAIM 8.14. *Assuming the $(n, q, \sigma_{\text{lwe}})$ -LWE assumption, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries, there exists a negligible function $n_{o,0}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ and all indices $o \in [q_{\text{keys}}]$, $a_{\mathcal{A},o-1,0} - a_{\mathcal{A},o,0} \leq n_{o,0}(\lambda)$.*

Proof. Suppose there exist an adversary making q_{keys} key queries, and a non-negligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, there exists an index $o \in [q_{\text{keys}}]$ such that $a_{\mathcal{A},o-1,0} - a_{\mathcal{A},o,0} \geq \eta(\lambda)$. We will use $\mathcal{A}$ to build a reduction algorithm $\mathcal{B}$ that breaks the $(n, q, \sigma_{\text{lwe}})$ -LWE assumption.

The reduction algorithm first receives $(1^k, 1^w, 1^L)$ from $\mathcal{A}$. It sets $\tilde{n} = (4\lambda + w) \cdot n$. The reduction algorithm queries the LWE challenger m times and receives $\{(\mathbf{a}_i, u_i)\}_{i \leq m}$. It sets a matrix $\mathbf{A} = [\mathbf{a}_1^T \ \mathbf{a}_2^T \ \dots \ \mathbf{a}_m^T]$ (that is, $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$) and $\mathbf{u} = [u_1 \ u_2 \ \dots \ u_m]$ (that is, $\mathbf{u} \in \mathbb{Z}_q^m$).

The reduction algorithm then chooses two tags, $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$, and let j^* be the first index where they differ. Next, the reduction algorithm defines set $S^{(i)}$ for each i , set $\hat{S}$, matrices $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in \hat{S}}$, $\{\mathbf{P}_{i,v}\}_{i \in [\ell], v \in [w]}$ and $\{T_i\}_{i \in [\ell]}$ as in Game 3. i^* .3 (and Game 3. i^* .4). Note that $(i^*, j^*, \beta, b) \notin \hat{S}$ for $b, \beta \in \{0, 1\}$. The reduction algo-

algorithm chooses $\mathbf{B}_{i^*,b}^{(j^*,1)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $b \in \{0,1\}$. It sends the public parameters to $\mathcal{A}$.

The adversary sends two challenge functions, $\mathbf{BP}^{(0)}, \mathbf{BP}^{(1)}$, and a ciphertext query $\mathbf{BP}$. Note that in **Game 3.i*.3** (and **Game 3.i*.4**), the challenge and query ciphertext are computed identically, and the reduction algorithm has all the matrices/trapdoors required for computing the ciphertext components.

Next, after receiving the challenge ciphertext and the query ciphertext, the adversary queries for secret keys. For the first $o-1$ secret keys, the reduction algorithm responds as in $H_{o-1,0}$ (which is identical to the response in $H_{o,0}$). In particular, to handle these key queries, the reduction algorithm does not require $\mathbf{B}_{i^*,\beta}^{(j^*,b)}$, since in both hybrids $\tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} \leftarrow \mathbb{Z}_q^m$. For the o th query, the reduction algorithm receives $\mathbf{x} \in \{0,1\}^k$. It sets $\tilde{\mathbf{x}}$ by repeating the input L times, and sets $\mathbf{B}_{i^*,\tilde{\mathbf{x}}_{i^*}}^{(j^*,0)} = \mathbf{A}$ (the LWE public matrix) and chooses $\mathbf{B}_{i^*,1-\tilde{\mathbf{x}}_{i^*}}^{(j^*,0)} \leftarrow \mathbb{Z}_q^{n \times m}$ (this might be used for the later key queries). The reduction algorithm sets

$$\tilde{\mathbf{t}}_{i^*}^{(j^*,0)} = - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,0)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{\mathbf{x}}_{\delta}}^{(0)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{\mathbf{x}}_{\delta}}^{(0)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\mathbf{st}_{i^*}^{(0)}} + \mathbf{u}.$$

It computes $\tilde{\mathbf{t}}_{i^*}^{(j^*,1)}$ as in $H_{o-1,0}$ (and $H_{o,0}$). All the remaining key queries are handled identically in $H_{o-1,0}$ and $H_{o,0}$, and the reduction algorithm has all the matrices required to compute them (in particular, after responding to the o th query, all $\{\mathbf{B}_{i^*,b}^{(j^*,\beta)}\}_{(b,\beta) \in \{0,1\}^2}$ are well defined).

Finally, the adversary sends its guess, and the reduction algorithm forwards it to the LWE challenger. Clearly, if $\mathbf{u}$ is a uniformly random vector, then so is the $\tilde{\mathbf{t}}_{i^*}^{(j^*,0)}$ component for the o th query, and therefore $\mathcal{B}$ perfectly simulates $H_{o,0}$. If $\tilde{\mathbf{t}}_{i^*}^{(j^*,0)} = \mathbf{s} \cdot \mathbf{B}_{i^*,\tilde{\mathbf{x}}_{i^*}}^{(j^*,0)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,0)}$, then the reduction algorithm implicitly sets $\mathbf{s}_{i^*}^{(j^*,0)} = \mathbf{s}$. Also, note that $\mathbf{s}_{i^*}^{(j^*,0)}$ is chosen afresh for each key query, and hence $\mathbf{s}$ will not be required anywhere else in simulating $H_{o-1,0}$. Therefore the reduction algorithm perfectly simulates $H_{o-1,0}$. As a result, it breaks the LWE assumption with advantage η .

CLAIM 8.15. *Assuming the $(n, q, \sigma_{\text{LWE}})$ -LWE assumption holds, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries, there exists a negligible function $\mathbf{n}_{o,1}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ and all indices $o \in [q_{\text{keys}}]$, $a_{\mathcal{A},o-1,1} - a_{\mathcal{A},o,1} \leq \mathbf{n}_{o,1}(\lambda)$.*

This proof is identical to the proof of Claim 8.14.

Using the above claims, it follows that for any PPT adversary, there exists a negligible function $\text{negl}_{3,i^*,4}$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3,i^*,3}(\lambda) - \text{Adv}_{\mathcal{A},3,i^*,4}(\lambda) \leq \text{negl}_{3,i^*,4}(\lambda)$. $\square$

LEMMA 8.16. *For any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}_{3,(i^*+1),1}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3,i^*,4}(\lambda) - \text{Adv}_{\mathcal{A},3,(i^*+1),1}(\lambda) \leq \text{negl}_{3,(i^*+1),1}(\lambda)$.*

Proof. The only difference between **Game 3.i*.4** and **Game 3.(i*+1).1** is that for each key query, the components $(\mathbf{t}_{i^*+1}^{(j^*,0)}, \mathbf{t}_{i^*+1}^{(j^*,1)})$ are computed differently. In particular, in **Game 3.(i*+1).1**, the term $\tilde{\mathbf{t}}_{i^*+1}^{(j^*,\beta)}$ has an additional term $\tilde{\mathbf{e}}_{i^*+1}^{(j^*,\beta)}$ which is drawn from the χ_{LWE}^m distribution.

The proof of this lemma is identical to the proof of Lemma 8.10 by setting $\text{negl}_{3,(i^*+1),1}(\cdot) = q_{\text{keys}}(\cdot) \cdot (2m \cdot \text{negl}_{\text{smud}}(\cdot))$ (recall that $\text{negl}_{\text{smud}}(\cdot)$ is the negligible function given by Lemma 2.1). $\square$

Next, we will look at **Game 4.i** for $i \in [\ell]$. For notational convenience, we refer to **Game 4** as **Game 4.0**. First, recall that **Game 4.0** is identical to **Game 3.ℓ.4**. Note that in **Game 4.0**, the challenger uses $\{\mathbf{P}_{1,v}\}_{v \in [w]}$ only for computing $(\mathbf{U}_{1,0}^*, \mathbf{U}_{1,1}^*)$ (for the challenge ciphertext) and $(\mathbf{U}_{1,0}, \mathbf{U}_{1,1})$ (for the ciphertext query). In particular, it is not used in the key generation phase. More generally, for all $i \in [\ell]$, $\{\mathbf{P}_{i,v}\}_{v \in [w]}$ is used in **Game 4.(i-1)** only for computing $(\mathbf{U}_{i,0}^*, \mathbf{U}_{i,1}^*)$ and $(\mathbf{U}_{i,0}, \mathbf{U}_{i,1})$. This observation is useful for the following lemma.

LEMMA 8.17. *Assuming LT_{en} satisfies the $(q, \chi_{\text{appr}}, \sigma_{\text{pre}})$ -target switching property, then for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}_{4,i^*}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A},4,(i^*-1)} - \text{Adv}_{\mathcal{A},4,i^*} \leq \text{negl}_{4,i^*}(\lambda)$.*

Proof. The only difference between **Game 4.(i*-1)** and **Game 4.i*** is with respect to the matrices $(\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*)$ (in the challenge ciphertext) and $(\mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1})$ (in the ciphertext query). In **Game 4.(i*-1)**, these matrices are computed using **Mixed-SubEnc**, while they are computed using **Mixed-SubEnc*** in **Game 4.i***. Recall that the only difference between the **Mixed-SubEnc** and **Mixed-SubEnc*** ciphertext components is that the **Mixed-SubEnc*** outputs map the $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ matrices to uniformly random matrices (instead of mapping to $\{\mathbf{P}_{i^*+1,\pi(v)}\}_{v \in [w]}$ as in **Mixed-SubEnc**). An important point to note is that the $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ matrices are not used anywhere else in both games. In particular, note that $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ are not used for computing $\{\mathbf{U}_{i^*-1,b}^*, \mathbf{U}_{i^*-1,b}\}_{b \in \{0,1\}}$.

Suppose there exist an adversary $\mathcal{A}$ and a nonnegligible function $\eta(\cdot)$ such that $\text{Adv}_{4,i^*-1}(\lambda) - \text{Adv}_{4,i^*}(\lambda) \geq \eta(\lambda)$. We will construct a reduction algorithm that breaks the target switching property with advantage $\eta(\cdot)$.

Setup phase. The reduction algorithm first performs the following steps from the setup phase, which are common for both **Game 4.(i*-1)** and **Game 4.i***. It defines $\tilde{n}_i, S^{(i)}$ for all $i \in [\ell]$, chooses $\text{tag}^*, \text{tag} \leftarrow \{0,1\}^\lambda$, and defines j^* as the first index where tag^* and tag differ. Next, it defines $\{\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}}, \{\mathbf{P}_{i,v}\}_{v \in [w]}, T_i\}_{i \neq i^*}$ as in **Game 4.(i*-1)**. It also defines $\hat{S}$ and chooses $\{\mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in \hat{S}}$ as in **Game 4.(i*-1)**.

The reduction algorithm sets $\tilde{k} = \tilde{n}_i - w \cdot n$ and queries the target switching property challenger by sending $1^{\tilde{n}_i}, 1^m$ and setting $S_B = [\tilde{k}]$. It receives a matrix $\mathbf{A} \in \mathbb{Z}_q^{\tilde{k} \times m}$ and parses $\mathbf{A}$ as follows:

$$\mathbf{A} = \left[\left\{ \mathbf{B}_{j,b}^{(1,\beta)} \right\}_{(j,\beta,b) \in S^{(i^*)}} \right].$$

Challenge phase. The reduction algorithm receives $\text{BP}^{(1)}, \text{BP}^{(2)}$. It chooses $\gamma \leftarrow \{0,1\}$, and for all $\alpha < i^*$ it computes $(\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*)$ using **Mixed-SubEnc*** (as in **Game 4.(i*-1)** and **Game 4.i***). Note in particular that $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ are not used for computing these matrices. It then sends its target switching property query matrices $\mathbf{Z}_{0,b}^*, \mathbf{Z}_{1,b}^*$ of dimensions $\tilde{n}_i \times m$ defined below.

$$\mathbf{Z}_{0,b}^* = \left[\left\{ \mathbf{C}_{i^*,b}^{(j,\beta)} \right\}_{(j,\beta,b) \in S^{(i^*)}} \right] \quad \mathbf{Z}_{1,b}^* = \left[\left\{ \mathbf{C}_{i^*,b}^{(j,\beta)} \right\}_{(j,\beta,b) \in S^{(i^*)}} \right] \leftarrow \mathbb{Z}_q^{w \cdot n \times m}.$$

It receives $\mathbf{U}_{i^*,b}^*$ from the challenger.

Query phase. The ciphertext query is handled similarly to the challenge ciphertext. The key queries are handled identically in both Game 4. $(i^* - 1)$ and Game 4. i^* . $\square$

LEMMA 8.18. *For any adversary $\mathcal{A}$, $\text{Adv}_{\mathcal{A},5} = 0$.*

Proof. We will argue that any adversary $\mathcal{A}$ has advantage 0 in Game 5. First, note that the challenge phase uses Mixed-SubEnc^* . As a result, it does not have any information about the choice $\gamma \leftarrow \{0,1\}$. Next, in the key query phase, for each key query, the challenger chooses $\mathbf{P}_{\ell+1, \text{st}_{\ell+1}}^{(\beta)}$, which might depend on γ . However, the important point here is that for each key query x , both challenge programs have identical output. Therefore, the $\mathbf{P}_{\ell+1, \text{st}_{\ell+1}}^{(\beta)}$ matrices are independent of γ . As a result, the adversary has zero advantage in Game 5. $\square$

8.4.3. 1-bounded restricted accept indistinguishability. In order to prove restricted accept indistinguishability security, we take a slightly different approach. First, we show that our construction achieves *complete* accept indistinguishability security, which is defined as follows.

DEFINITION 8.19 (*q*-bounded complete accept indistinguishability). *Let $q(\cdot)$ be any fixed polynomial. A Mixed FE scheme $\text{Mixed-FE} = (\text{Setup}, \text{Enc}, \text{SK-Enc}, \text{KeyGen}, \text{Dec})$ is said to satisfy q -bounded complete accept indistinguishability security if there exist algorithms $\text{SK-Enc}^*, \text{KeyGen}^*$ such that for every stateful PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:*

$$\Pr \left[\mathcal{A}^{O_1^b(\cdot), O_2^b(\cdot)}(\text{pp}, \text{ct}_b) = b : \begin{array}{l} (1^\kappa, f^*) \leftarrow \mathcal{A}(1^\lambda); \\ (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ b \leftarrow \{0,1\}; \text{ct}_1 \leftarrow \text{SK-Enc}(\text{msk}, f^*); \\ \text{ct}_0 \leftarrow \text{Enc}(\text{pp}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where

- oracle $O_1^0(\cdot) = \text{KeyGen}^*(\text{pp}, \cdot)$, $O_1^1(\cdot) = \text{KeyGen}(\text{msk}, \cdot)$, $O_2^0(\cdot) = \text{SK-Enc}^*(\text{pp}, \cdot)$, $O_2^1(\cdot) = \text{SK-Enc}(\text{msk}, \cdot)$;
- $\mathcal{A}$ can make at most $q(\lambda)$ queries to the $O_2^b(\cdot)$ oracle;
- every secret key query m made by adversary $\mathcal{A}$ to the $O_1^b(\cdot)$ oracle must satisfy the condition that $f^*(m) = 1$ as well as $f(m) = 1$ for every ciphertext query f made by $\mathcal{A}$ to the $O_2^b(\cdot)$ oracle; and
- $\mathcal{A}$ must make all (at most $q(\lambda)$) $O_2^b(\cdot)$ oracle queries before making any query to the $O_1^b(\cdot)$ oracle.

At a high level, this states that if the adversary only queries for keys for inputs m and ciphertexts for functions f such that $f(m) = 1$ on all combinations, then there exist special encryption and key generation algorithms ($\text{SK-Enc}^*, \text{KeyGen}^*$) such that they only take public parameters as inputs, and the adversary cannot distinguish between correctly computed keys and ciphertexts from these (simulated) special keys and ciphertexts.

Below we provide a sequence of hybrid games that we later use to argue complete accept indistinguishability security. To complete the argument, later (in section 8.4.5) we simply argue that complete accept indistinguishability implies restricted accept indistinguishability.

Game 0 This corresponds to the original 1-bounded restricted accept indistinguishability security game in which the challenger encrypts the challenge branching program BP^* sent by the adversary.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0,1\}$.
 4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** The challenger chooses a random λ -bit string $\text{tag}^* \leftarrow \{0,1\}^\lambda$. Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = [\ell] \times [\lambda] \times \{0,1\}^2.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger chooses a λ -bit string $\text{tag} \leftarrow \{0,1\}^\lambda$ and responds as follows:
 - (a) Let $S = [\ell] \times [\lambda] \times \{0,1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\} \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda - 1]$. Next, it sets vector $\mathbf{y}^{(\lambda)}$ as

$$\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$, error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1. \end{cases}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 1 This is identical to the previous game, except the challenger now chooses both tags tag^* and tag at the beginning during the setup phase, and it aborts if $\text{tag}^* = \text{tag}$.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= [\ell] \times [\lambda] \times \{0, 1\}^2. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program $\mathbf{BP}$ for encryption. The challenger responds as follows:
 - (a) Let $S = [\ell] \times [\lambda] \times \{0,1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda - 1]$. Next, it sets vector $\mathbf{y}^{(\lambda)}$ as

$$\mathbf{y}^{(\lambda)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda-1]} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$, error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0,1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0,1\}, \\ \mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1. \end{cases} \end{aligned}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.
- **Guess.** The adversary finally sends the guess γ' .

Notation. In all the following hybrid games, let diff denote the set of indices j such that $\text{tag}_j^* \neq \text{tag}_j$. Similarly, let comm denote the set of indices j such that $\text{tag}_j^* = \text{tag}_j$. Concretely, in all the following hybrids, sets diff and comm are defined as follows:

$$\text{diff} \stackrel{\text{def}}{=} \{j \in [\lambda] : \text{tag}_j^* \neq \text{tag}_j\}, \quad \text{comm} \stackrel{\text{def}}{=} \text{comm}.$$

Additionally, let j^* denote the smallest index in diff (i.e., $j^* = \min_{j \in \text{diff}} j$), and let $\beta^* = \text{tag}_{j^*}^*$. Note that since the challenger aborts whenever $\text{tag}^* = \text{tag}$, thus j^*, β^* always exist whenever the challenger does not abort. Also, we will use $\widehat{\text{diff}}$ to denote the set diff excluding index j^* , i.e., $\widehat{\text{diff}} = \text{diff} \setminus \{j^*\}$.

Game 2 This is identical to the previous game, except the challenger, while answering a secret key query, now puts the $\tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1}$ component in $\mathbf{y}^{(j^*)}$ instead of $\mathbf{y}^{(\lambda)}$, and the rest are sampled uniformly at random.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below.
2. Next, it samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $i \in [\ell], j \in [\lambda], \beta, b \in \{0, 1\}$.
 4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.
- **Challenge phase.** Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = [\ell] \times [\lambda] \times \{0, 1\}^2.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:

- (a) Let $S = [\ell] \times [\lambda] \times \{0, 1\}^2$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$. Next, it sets vector $\mathbf{y}^{(j^*)}$ as

$$\mathbf{y}^{(j^*)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{1,1} - \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}.$$

- (b) It then chooses secret vectors $\{\mathbf{s}_i^{(j,\beta)}\}_{i,j,\beta}$, error vectors $\{\mathbf{e}_i^{(j,\beta)}\}_{i,j,\beta}$ as

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell + 1] \times [\lambda] \times \{0, 1\}, \\ \mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1. \end{cases} \end{aligned}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Next, we have a sequence of 4ℓ hybrid experiments, **Game 3.i***. $\{1, 2, 3, 4\}$ for $i^* = 1$ to ℓ .

Game 3.i*.1 In hybrids **Game 3.i*.1**, the $\mathbf{B}_{i,b}^{(j,\beta)}$, $\mathbf{C}_{i,b}^{(j,\beta)}$ matrices for all diff strands and levels $i < i^*$ are *not* sampled (at all) along with other level i matrices; ciphertext components for levels $i < i^*$ are used to target only the remaining matrices; i.e., the ciphertext matrices do not target $\mathbf{B}_{i,b}^{(j,\beta)}$ matrices for $j \in \text{diff}$ and $i < i^*$ to some pre-specified $\mathbf{C}_{i,b}^{(j,\beta)}$ or random matrices. Also, the first $i^* - 1$ components in each secret key are set to be uniformly random vectors, and the next component is hardwired such that correctness holds, and some smudgeable noise is also introduced in these components. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^*, \quad S^{(i)} &= \text{comm} \times \{0, 1\}^2, \\ \forall i \geq i^*, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - |\text{diff}| \cdot 4n & \text{for } i < i^*, \\ \tilde{n} & \text{for } i \geq i^*. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{i^*}^{(1-\beta^*)}, \text{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, BP* after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)},$$

$$\forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\},$$

$$\begin{aligned} \mathbf{t}_{i^*}^{(j,\beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_\alpha^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ &\quad + \mathbf{s}_{i^*}^{(j,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j,\beta)} + \mathbf{e}_{i^*}^{(j,\beta)}, \end{aligned}$$

$$\forall \beta \in \{0, 1\},$$

$$\begin{aligned} \mathbf{t}_{i^*}^{(j^*,\beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_\alpha^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*,\beta)} + \mathbf{e}_{i^*}^{(j^*,\beta)}, \end{aligned}$$

$$\forall (i, j, \beta) \in ([\ell + 1] \setminus [i^*]) \times \text{diff} \times \{0, 1\},$$

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_\ell^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 3. $i^*.2$ This is identical to the previous game, except the $(i^* + 1)$ th key component in all diff strands is also hardwired. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^*, \quad S^{(i)} &= \text{comm} \times \{0, 1\}^2, \\ \forall i \geq i^*, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - |\text{diff}| \cdot 4n & \text{for } i < i^*, \\ \tilde{n} & \text{for } i \geq i^*. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j, \beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j, \beta)}\}_{(j, \beta, b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0, 1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0, 1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{i^*}^{(1-\beta^*)}, \mathbf{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, BP* after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

Also, for $(j, \beta) \in \text{diff} \times \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j,\beta)} = \mathbf{0}^{n \times m}$, and let $\tilde{\mathbf{t}}_{i^*}^{(j,\beta)}$ denote the following vector. For all $j \in \widehat{\text{diff}}$,

$$\begin{aligned} \tilde{\mathbf{t}}_{i^*}^{(j,\beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \mathbf{s}_{i^*}^{(j,\beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j,\beta)}, \\ \tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*, \mathbf{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j^*, \beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*, \beta)}. \end{aligned}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)},$$

$$\forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j,\beta)} = \tilde{\mathbf{t}}_{i^*}^{(j,\beta)} + \mathbf{e}_{i^*}^{(j,\beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j^*,\beta)} = \tilde{\mathbf{t}}_{i^*}^{(j^*,\beta)} + \mathbf{e}_{i^*}^{(j^*,\beta)},$$

$$\begin{aligned} \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j,\beta)} = & - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_\alpha^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) \\ & + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ & + \mathbf{s}_{i^*+1}^{(j,\beta)} \cdot \mathbf{B}_{i^*+1,\tilde{x}_{i^*+1}}^{(j,\beta)} + \mathbf{e}_{i^*+1}^{(j,\beta)}, \end{aligned}$$

$$\begin{aligned} \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j^*,\beta)} = & - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_\alpha^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \\ & + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1,\text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*,\beta)} \cdot \mathbf{B}_{i^*+1,\tilde{x}_{i^*+1}}^{(j^*,\beta)} + \mathbf{e}_{i^*+1}^{(j^*,\beta)}, \end{aligned}$$

$$\forall (i, j, \beta) \in ([\ell + 1] \setminus [i^* + 1]) \times \text{diff} \times \{0, 1\},$$

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 3.*i.3** This is identical to the previous game, except $\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}$ for diff strands and levels $i = i^*$ are *not* sampled along with other level i^* matrices, but instead they are sampled uniformly at random. Also, ciphertext components for level i^* are used to target only the remaining matrices. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow$

$\{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^* + 1, \quad S^{(i)} &= \text{comm} \times \{0, 1\}^2, \\ \forall i \geq i^* + 1, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - |\text{diff}| \cdot 4n & \text{for } i < i^* + 1, \\ \tilde{n} & \text{for } i \geq i^* + 1. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\begin{aligned} \forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) &\leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q), \\ \forall (j, \beta, b) \in \text{diff} \times \{0, 1\}^2, \quad \mathbf{B}_{i^*,b}^{(j,\beta)} &\leftarrow \mathbb{Z}_q^{n \times m}. \end{aligned}$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{i^*}^{(1-\beta^*)}, \text{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, BP* after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

Also, for $(j, \beta) \in \text{diff} \times \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j,\beta)} = \mathbf{0}^{n \times m}$, and let $\tilde{\mathbf{t}}_{i^*}^{(j,\beta)}$ denote the following vector. For all $j \in \text{diff}$,

$$\begin{aligned} \tilde{\mathbf{t}}_{i^*}^{(j,\beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \mathbf{s}_{i^*}^{(j,\beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j,\beta)}, \\ \tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)} &= - \sum_{\alpha=1}^{i^*-1} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*, \text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*, \tilde{x}_{i^*}}^{(j^*, \beta)} + \tilde{\mathbf{e}}_{i^*}^{(j^*, \beta)}. \end{aligned}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j,\beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_{\ell}}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [i^* - 1] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)},$$

$$\begin{aligned} \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j, \beta)} &= \tilde{\mathbf{t}}_{i^*}^{(j, \beta)} + \mathbf{e}_{i^*}^{(j, \beta)}, \\ \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*}^{(j^*, \beta)} &= \tilde{\mathbf{t}}_{i^*}^{(j^*, \beta)} + \mathbf{e}_{i^*}^{(j^*, \beta)}. \end{aligned}$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\begin{aligned} \mathbf{t}_{i^*+1}^{(j, \beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_{\alpha}^{(j, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \mathbf{s}_{i^*+1}^{(j, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j, \beta)} + \mathbf{e}_{i^*+1}^{(j, \beta)}. \end{aligned}$$

For all $\beta \in \{0, 1\}$,

$$\begin{aligned} \mathbf{t}_{i^*+1}^{(j^*, \beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1, \text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j^*, \beta)} + \mathbf{e}_{i^*+1}^{(j^*, \beta)}, \end{aligned}$$

$\forall (i, j, \beta) \in ([\ell+1] \setminus [i^*+1]) \times \widehat{\text{diff}} \times \{0, 1\}$,

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_{\ell}}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1. \end{cases}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 3.i*.4 This is identical to the previous game, except the i^* th level key component in $\widehat{\text{diff}}$ strands is a uniformly random n length vector, i.e., all first i^* level components in $\widehat{\text{diff}}$ strands are random elements. Also, we no longer sample the matrices $\mathbf{B}_{i, b}^{(j, \beta)}, \mathbf{C}_{i, b}^{(j, \beta)}$ for $\widehat{\text{diff}}$ strands and levels $i = i^*$ at all. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i < i^* + 1, \quad S^{(i)} &= \text{comm} \times \{0, 1\}^2, \\ \forall i \geq i^* + 1, \quad S^{(i)} &= [\lambda] \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} \tilde{n} - |\widehat{\text{diff}}| \cdot 4n & \text{for } i < i^* + 1, \\ \tilde{n} & \text{for } i \geq i^* + 1. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i,j,\beta,b) \in \hat{S}$.
 4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i,j,\beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i,j,\beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j,\beta) \in [\lambda] \times \{0,1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \end{aligned}$$

$$\begin{aligned}\forall (i, j, \beta) \in [i^*] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j, \beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{e}}_{i^*}^{(j, \beta)} &\leftarrow \chi_{\text{lwe}}^m.\end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{i^*}^{(1-\beta^*)}, \text{st}_{i^*}^{(\beta^*)}$ denote the state of branching programs BP, BP^* after $i^* - 1$ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned}\Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}\end{aligned}$$

Also, for $(j, \beta) \in \text{diff} \times \{0, 1\}$, let $\mathbf{B}_{\ell+1, \tilde{x}_{\ell+1}}^{(j, \beta)} = \mathbf{0}^{n \times m}$. Next, it computes key vectors $\{\mathbf{t}_i^{(j, \beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} \mathbf{s}_1^{(j, \beta)} \cdot \mathbf{B}_{1, \tilde{x}_1}^{(j, \beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j, \beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j, \beta)} \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [i^*] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j, \beta)} = \tilde{\mathbf{t}}_i^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)},$$

$$\begin{aligned}\forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j, \beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_\alpha^{(j, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} \right) \\ &\quad + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} + \mathbf{s}_{i^*+1}^{(j, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j, \beta)} + \mathbf{e}_{i^*+1}^{(j, \beta)},\end{aligned}$$

$$\begin{aligned}\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{i^*+1}^{(j^*, \beta)} &= - \sum_{\alpha=1}^{i^*} \left(\tilde{\mathbf{t}}_\alpha^{(j^*, \beta)} \cdot \prod_{\delta=\alpha}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*} \mathbf{U}_{\delta, \tilde{x}_\delta}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*+1, \text{st}_{i^*+1}^{(\beta)}} + \mathbf{s}_{i^*+1}^{(j^*, \beta)} \cdot \mathbf{B}_{i^*+1, \tilde{x}_{i^*+1}}^{(j^*, \beta)} + \mathbf{e}_{i^*+1}^{(j^*, \beta)},\end{aligned}$$

$$\forall (i, j, \beta) \in ([\ell + 1] \setminus [i^* + 1]) \times \text{diff} \times \{0, 1\},$$

$$\mathbf{t}_i^{(j, \beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j, \beta)} \cdot \mathbf{C}_{i-1, \tilde{x}_{i-1}}^{(j, \beta)} + \mathbf{s}_i^{(j, \beta)} \cdot \mathbf{B}_{i, \tilde{x}_i}^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)} & \text{if } i \leq \ell, \\ -\mathbf{s}_\ell^{(j, \beta)} \cdot \mathbf{C}_{\ell, \tilde{x}_\ell}^{(j, \beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)} & \text{if } i = \ell + 1. \end{cases}$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ .

Game 4 This is similar to **Game 3.ℓ.4**, except that the terms $\mathbf{t}_{\ell+1}^{(j^*, \beta)}$ have an additional small error $\tilde{\mathbf{e}}_{\ell+1}^{(j^*, \beta)}$, which is smudged by the main error term $\mathbf{e}_{\ell+1}^{(j^*, \beta)}$.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \text{comm} \times \{0, 1\}^2.$$

Also, let $\tilde{n}_i = \tilde{n} - |\text{diff}| \cdot 4n$ for all $i \in [\ell]$.

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs BP, BP* after ℓ steps, respectively. Also, let $\Gamma, \tilde{\mathbf{y}}$, and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{y}} = \sum_{j \in [\lambda] \setminus \{j^*\}} \mathbf{y}^{(j)}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)},$$

$$\begin{aligned} \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j,\beta)} &= - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \right) \\ &\quad + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, \end{aligned}$$

$$\begin{aligned} \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} &= - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j^*,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \right) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\mathbf{st}_{\ell+1}^{(\beta)}}^{(\beta^*)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)} + \tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta)}. \end{aligned}$$

- (d) Finally, it sends the secret key as $(x, \{t_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.
- **Guess.** The adversary finally sends the guess γ' .

Game 5 This is identical to the previous game, except the $(\ell+1)$ th key components in the special strand (i.e., j^* th strand) are random elements. We describe this in detail below.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:
 1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \text{comm} \times \{0, 1\}^2.$$

Also, let $\tilde{n}_i = \tilde{n} - |\text{diff}| \cdot 4n$ for all $i \in [\ell]$.

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:

- (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
- (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
- (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs $\mathbf{BP}, \mathbf{BP}^*$ after ℓ steps, respectively. Also, let Γ and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all tuples $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_{\ell}^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_{\ell}}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}.$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\mathbf{t}_{\ell+1}^{(j,\beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*, \beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)}.$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Next, we have a sequence of ℓ hybrid experiments, **Game 5. i^*** for $i^* = 1$ to ℓ .

Game 5. i^* This is identical to the previous game, except matrices $\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}$ for $j \in \text{comm}, \beta = 1 - \text{tag}_j$ strands (i.e., strands in which $\mathbf{B}_{i,b}^{(j, \beta)}$ were targeting random matrices themselves) and levels $i \leq i^*$ are *not* sampled along with other level $i \leq i^*$ matrices, but instead they are sampled uniformly at random. Also, ciphertext components for levels $i \leq i^*$ are used to target only the remaining matrices. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\begin{aligned} \forall i \leq i^*, \quad S^{(i)} &= \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}, \\ \forall i > i^*, \quad S^{(i)} &= \text{comm} \times \{0, 1\}^2. \end{aligned}$$

Also, let

$$\tilde{n}_i = \begin{cases} (2 \cdot |\text{comm}| + w)n & \text{for } i \leq i^*, \\ \tilde{n} - |\text{diff}| \cdot 4n & \text{for } i > i^*. \end{cases}$$

Set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j, \beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \{\mathbf{B}_{i,b}^{(j, \beta)}\}_{(j, \beta, b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q),$$

$$\forall (i, j, \beta, b) \in ([i^*] \times \text{comm} \times \{0, 1\}^2) \setminus \hat{S}, \quad \mathbf{B}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}.$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0, 1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{pmatrix} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}\}_{(i,j,\beta,b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{pmatrix} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs BP, BP* after ℓ steps, respectively. Also, let Γ and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= [\ell + 1] \times \text{comm} \times \{0, 1\}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j, \beta)} = \tilde{\mathbf{t}}_i^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)}.$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\mathbf{t}_{\ell+1}^{(j, \beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j, \beta)} \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*, \beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)}.$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Next, we have a sequence of ℓ hybrid experiments, **Game 6. i^*** for $i^* = 2$ to $\ell + 1$.

Game 6. i^* This is identical to the previous game (i.e., **Game 5. ℓ**), except in the **comm** strands, for which we sample $\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}$ matrices uniformly at random, the key components for levels $i \leq i^*$ are random elements. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j, \beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j, \beta)}\}_{(j, \beta, b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q),$$

$$\forall (i, j, \beta, b) \in ([\ell] \times \text{comm} \times \{0, 1\}^2) \setminus \hat{S}, \quad \mathbf{B}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}.$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\text{BP}^* = \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0, 1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right),$$

$$S^* = \hat{S}.$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}\}_{(i, j, \beta, b) \in S^*}, \\ \{\mathbf{P}_{i,v}\}_{(i,v) \in [\ell] \times [w]}, \\ \{T_i\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0,1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0,1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (i, j) \in [i^*] \times \text{comm}, \quad \tilde{\mathbf{t}}_i^{(j,1-\text{tag}_j)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0,1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{st}_{\ell+1}^{(1-\beta^*)}, \mathbf{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs BP, BP* after ℓ steps, respectively. Also, let Γ and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\Gamma = ([\ell+1] \setminus [i^*]) \times \text{comm} \times \{0,1\} \cup \{(i^*, j, \beta) : j \in \text{comm}, \beta = \text{tag}_j\},$$

$$\forall (i, \beta, b) \in [\ell] \times \{0,1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows. For all $(i, j, \beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} \mathbf{s}_1^{(j,\beta)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)} & \text{if } i = 1, \\ -\mathbf{s}_{i-1}^{(j,\beta)} \cdot \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \cdot \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\begin{aligned} \forall (i, j) \in [i^*] \times \text{comm}, \quad \mathbf{t}_i^{(j, 1-\text{tag}_j)} &= \tilde{\mathbf{t}}_i^{(j, 1-\text{tag}_j)} + \mathbf{e}_i^{(j, 1-\text{tag}_j)}, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j, \beta)} &= \tilde{\mathbf{t}}_i^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)}. \end{aligned}$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\begin{aligned} \mathbf{t}_{\ell+1}^{(j, \beta)} &= - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j, \beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)}, \\ \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*, \beta)} &= \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)}. \end{aligned}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Next, we have a sequence of $\ell + 1$ hybrid experiments, **Game 7. i^*** for $i^* = 1$ to $\ell + 1$.

Game 7. i^* This is identical to the previous game (i.e., **Game 6.** $(\ell + 1)$), except in the **comm** strands, for which we *still* sample $\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}$ matrices using **EnTrapGen**, the key components for levels $i \leq i^*$ are random elements. Also, we no longer sample the matrices $\mathbf{B}_{i,b}^{(j, \beta)}, \mathbf{C}_{i,b}^{(j, \beta)}$ at all, which were sampled uniformly at random in the previous game. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j, \beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j, \beta)}\}_{(j, \beta, b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0, 1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6)

as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}^*, U_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{U_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.
 1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{U_{\alpha,0}, U_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{U_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses a secret vector as $\tilde{\mathbf{s}} \leftarrow \chi_s^n$ and $\lambda - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in [\lambda] \setminus \{j^*\}$.
 - (b) It then chooses vectors $\mathbf{s}_i^{(j,\beta)}, \mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{s}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^n, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (i, j) \in [\ell + 1] \times \text{comm}, \quad \tilde{\mathbf{t}}_i^{(j, 1 - \text{tag}_j)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (i, j) \in [i^*] \times \text{comm}, \quad \tilde{\mathbf{t}}_i^{(j, \text{tag}_j)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\text{st}_{\ell+1}^{(1-\beta^*)}, \text{st}_{\ell+1}^{(\beta^*)}$ denote the state of branching programs BP, BP* after ℓ steps, respectively. Also, let Γ and $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\begin{aligned} \Gamma &= \{(i, j, \beta) \in ([\ell + 1] \setminus [i^*]) \times \text{comm} \times \{0, 1\} : \beta = \text{tag}_j\}, \\ \forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} &= \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases} \end{aligned}$$

For $v \in [w]$, let $\mathbf{P}_{\ell+1,v}^{(\beta^*)}$ be the top level matrices chosen while computing the challenge ciphertext. Similarly, let $\mathbf{P}_{\ell+1,v}^{(1-\beta^*)}$ be the top level matrices chosen while computing the query ciphertext. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\forall (i,j) \in [i^*] \times \text{comm}, \quad \mathbf{t}_i^{(j,\text{tag}_j)} = \tilde{\mathbf{t}}_i^{(j,\text{tag}_j)} + \mathbf{e}_i^{(j,\text{tag}_j)}.$$

For all $(i,j,\beta) \in \Gamma$,

$$\mathbf{t}_i^{(j,\beta)} = \begin{cases} -\mathbf{s}_{i-1}^{(j,\beta)} \mathbf{C}_{i-1,\tilde{x}_{i-1}}^{(j,\beta)} + \mathbf{s}_i^{(j,\beta)} \mathbf{B}_{i,\tilde{x}_i}^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)} & \text{if } 1 < i \leq \ell, \\ -\mathbf{s}_\ell^{(j,\beta)} \cdot \mathbf{C}_{\ell,\tilde{x}_\ell}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)} & \text{if } i = \ell + 1, \end{cases}$$

$$\forall (i,j) \in [\ell+1] \times \text{comm}, \quad \mathbf{t}_i^{(j,1-\text{tag}_j)} = \tilde{\mathbf{t}}_i^{(j,1-\text{tag}_j)} + \mathbf{e}_i^{(j,1-\text{tag}_j)},$$

$$\forall (i,j,\beta) \in [\ell] \times \text{diff} \times \{0,1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}.$$

For all $(j,\beta) \in \widehat{\text{diff}} \times \{0,1\}$,

$$\mathbf{t}_{\ell+1}^{(j,\beta)} = -\sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_\delta}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)},$$

$$\forall \beta \in \{0,1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}.$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 8 This is identical to the previous game (i.e., **Game 7**. $(\ell+1)$). For ease of exposition, we describe it in detail below.

- **Setup phase.** The adversary sends the functionality index (k,w,L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n,m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0,1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j,\beta,b) \in [\lambda] \times \{0,1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i,j,\beta,b) \in [\ell] \times [\lambda] \times \{0,1\}^2 : (j,\beta,b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\begin{bmatrix} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{bmatrix}, T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i,j,\beta,b) \in \hat{S}$.

4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

• **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{ \pi_{i,b}^* : [w] \rightarrow [w] \}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{ \mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^* \}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \}_{(i,j,\beta,b) \in S^*}, \\ \{ \mathbf{P}_{i,v} \}_{(i,v) \in [\ell] \times [w]}, \\ \{ T_i \}_{i \in [\ell]}, \text{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{ \mathbf{U}_{i,b}^* \}_{i \in [\ell], b \in \{0,1\}})$.

• **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the **Mixed-SubEnc** routine (described in Figure 6) as follows. For all $\alpha \in [\ell]$,

$$(\{ \mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1} \}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \}_{(i,j,\beta,b) \in S}, \\ \{ \mathbf{P}_{i,v} \}_{(i,v) \in [\ell] \times [w]}, \\ \{ T_i \}_{i \in [\ell]}, \text{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{ \mathbf{U}_{i,b} \}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows.
 - (a) It chooses $|\text{diff}| - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in \widehat{\text{diff}}$.
 - (b) It then chooses vectors $\mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0,1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0,1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{comm} \times \{0,1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0,1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\forall (i, \beta, b) \in [\ell] \times \{0,1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}. \end{aligned}$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\mathbf{t}_{\ell+1}^{(j,\beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*, \beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)}.$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Next, we have a sequence of ℓ hybrid experiments, **Game 8. i^*** for $i^* = 1$ to ℓ .

Game 8. i^* This is identical to the previous game, except now the challenger samples the first i^* ciphertext components (both challenge and queried) as random Gaussian matrices. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** Let

$$\begin{aligned} \text{BP}^* &= \left(\{\pi_{i,b}^* : [w] \rightarrow [w]\}_{i \in [\ell], b \in \{0,1\}}, \text{acc}^* \in [w], \text{rej}^* \in [w] \right), \\ S^* &= \hat{S}. \end{aligned}$$

The challenger then runs the **Mixed-SubEnc** routine (described in Figure 6) as

$$\forall (\alpha, b) \in [i^*] \times \{0, 1\}, \quad \mathbf{U}_{\alpha,b}^* \leftarrow \chi_{\text{pre}}^{m \times m}.$$

For all $\alpha \in [\ell] \setminus [i^*]$,

$$(\{\mathbf{U}_{\alpha,0}^*, \mathbf{U}_{\alpha,1}^*\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}^*, \alpha, S^*, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S^*}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP}^* \end{array} \right).$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) Let $S = \hat{S}$. It runs the Mixed-SubEnc routine (described in Figure 6) as

$$\forall (\alpha, b) \in [i^*] \times \{0, 1\}, \quad \mathbf{U}_{\alpha,b} \leftarrow \chi_{\text{pre}}^{m \times m}.$$

For all $\alpha \in [\ell] \setminus [i^*]$,

$$(\{\mathbf{U}_{\alpha,0}, \mathbf{U}_{\alpha,1}\}) \leftarrow \text{Mixed-SubEnc} \left(\begin{array}{c} \text{tag}, \alpha, S, \\ \left\{ \mathbf{B}_{i,b}^{(j,\beta)}, \mathbf{C}_{i,b}^{(j,\beta)} \right\}_{(i,j,\beta,b) \in S}, \\ \left\{ \mathbf{P}_{i,v} \right\}_{(i,v) \in [\ell] \times [w]}, \\ \left\{ T_i \right\}_{i \in [\ell]}, \mathbf{BP} \end{array} \right).$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.
2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:
 - (a) It chooses $|\text{diff}| - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in \widehat{\text{diff}}$.
 - (b) It then chooses vectors $\mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} &\leftarrow \mathbb{Z}_q^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\forall (i, j, \beta) \in [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} = \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)},$$

$$\forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j, \beta)} = \tilde{\mathbf{t}}_i^{(j, \beta)} + \mathbf{e}_i^{(j, \beta)}.$$

For all $(j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}$,

$$\mathbf{t}_{\ell+1}^{(j, \beta)} = - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j, \beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j, \beta)},$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*, \beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} + \mathbf{e}_{\ell+1}^{(j^*, \beta)}.$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j, \beta)}\}_{(i, j, \beta) \in [\ell+1] \times [\lambda] \times \{0, 1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 9 This is identical to the previous game (i.e., **Game 8.ℓ**), except the last secret key components in all $\widehat{\text{diff}}$ strands also include an additional noise term which is much smaller than the overall noise added in those components. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j, \beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j, \beta)}\}_{(j, \beta, b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j, \beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger generates ciphertext components as

$$\forall (i, b) \in [\ell] \times \{0, 1\}, \quad \mathbf{U}_{i,b}^* \leftarrow \chi_{\text{pre}}^{m \times m}.$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0, 1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:

- (a) It generates ciphertext components as

$$\forall (i, b) \in [\ell] \times \{0, 1\}, \quad \mathbf{U}_{i,b} \leftarrow \chi_{\text{pre}}^{m \times m}.$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.

2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- (a) It chooses $|\text{diff}| - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in \widehat{\text{diff}}$.

- (b) It then chooses vectors $\mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \tilde{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} &\leftarrow \chi_{\text{big}}^m, \\ \forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{last}}^m, \\ \forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} &\leftarrow \mathbb{Z}_q^m, \\ \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} &\leftarrow \chi_{\text{lwe}}^m. \end{aligned}$$

- (c) Let $\tilde{x} = x^L$, and let $\mathbf{U}_{i,b}^{(\beta)}$ denote the following:

$$\forall (i, \beta, b) \in [\ell] \times \{0, 1\}^2, \quad \mathbf{U}_{i,b}^{(\beta)} = \begin{cases} \mathbf{U}_{i,b}^* & \text{if } \beta = \beta^*, \\ \mathbf{U}_{i,b} & \text{if } \beta = 1 - \beta^*. \end{cases}$$

Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell+1] \times \text{comm} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, \\ \forall (i, j, \beta) \in [\ell] \times \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, \end{aligned}$$

$$\begin{aligned} \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j,\beta)} &= - \sum_{\alpha=1}^{\ell} \left(\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \right) \\ &\quad + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} \\ &\quad + \tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} \cdot \prod_{\delta=2}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, \end{aligned}$$

$$\forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} = \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}.$$

- (d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

- **Guess.** The adversary finally sends the guess γ' .

Game 10 This is identical to the previous game, except the last secret key components in all $\widehat{\text{diff}}$ strands are random vectors as well. Below we describe it in detail.

- **Setup phase.** The adversary sends the functionality index (k, w, L) and description of branching program BP^* to the challenger. Then the challenger proceeds as follows:

1. It chooses an LWE modulus q , dimensions n, m , and also distributions $\chi_{\text{big}}, \chi_s, \chi_{\text{appr}}, \chi_{\text{pre}}, \chi_{\text{last}}, \chi_{\text{lwe}}$ as described in the construction. Recall that $\ell = k \cdot L$ and $\tilde{n} = (4\lambda + w)n$. It also chooses two λ -bit strings $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$. If $\text{tag}^* = \text{tag}$, then it aborts and the adversary wins. Otherwise, the challenger continues as below. Let $S^{(i)}$ denote the following sets:

$$\forall i \in [\ell], \quad S^{(i)} = \{(j, \beta, b) \in [\lambda] \times \{0, 1\}^2 : j \in \text{comm} \wedge \beta = \text{tag}_j\}.$$

Also, let $\tilde{n}_i = (2 \cdot |\text{comm}| + w)n$ for all $i \in [\ell]$, and set $\hat{S} = \{(i, j, \beta, b) \in [\ell] \times [\lambda] \times \{0, 1\}^2 : (j, \beta, b) \in S^{(i)}\}$.

2. It samples $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i,j,\beta,b}, \{\mathbf{P}_{i,v}\}_{i,v}$ matrices as

$$\forall i \in [\ell], \quad \left(\left[\begin{array}{c} \{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}} \\ \{\mathbf{P}_{i,v}\}_{v \in [w]} \end{array} \right], T_i \right) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q).$$

3. It then samples matrices $\mathbf{C}_{i,b}^{(j,\beta)} \leftarrow \mathbb{Z}_q^{n \times m}$ for $(i, j, \beta, b) \in \hat{S}$.
4. Finally, it sends the public parameters $\text{pp} = (\lambda, n, m, q, k, w, L, \chi_{\text{pre}})$ to the adversary.

- **Challenge phase.** The challenger generates ciphertext components as

$$\forall (i, b) \in [\ell] \times \{0, 1\}, \quad \mathbf{U}_{i,b}^* \leftarrow \chi_{\text{pre}}^{m \times m}.$$

Finally, it sends the challenge ciphertext as $(\text{tag}^*, \{\mathbf{U}_{i,b}^*\}_{i \in [\ell], b \in \{0,1\}})$.

- **Post-challenge phase.** The adversary is allowed to make at most 1 secret key encryption query, followed by polynomially many secret key queries. The challenger responds to each query as below.

1. **Ciphertext query.** The adversary sends a branching program BP for encryption. The challenger responds as follows:
 - (a) It generates ciphertext components as

$$\forall (i, b) \in [\ell] \times \{0, 1\}, \quad \mathbf{U}_{i,b} \leftarrow \chi_{\text{pre}}^{m \times m}.$$

- (b) Finally, it sends the ciphertext as $(\text{tag}, \{\mathbf{U}_{i,b}\}_{i \in [\ell], b \in \{0,1\}})$.

2. **Secret key queries.** The adversary queries the challenger on polynomially many messages for corresponding secret keys. For each queried string x , the challenger responds as follows:

- (a) It chooses $|\widehat{\text{diff}}| - 1$ random vectors as $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$ for $j \in \widehat{\text{diff}}$.

- (b) It then chooses vectors $\mathbf{e}_i^{(j,\beta)}, \tilde{\mathbf{t}}_i^{(j,\beta)}, \hat{\mathbf{e}}_i^{(j,\beta)}$ as follows:

$$\forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \mathbf{e}_i^{(j,\beta)} \leftarrow \chi_{\text{big}}^m,$$

$$\forall (j, \beta) \in [\lambda] \times \{0, 1\}, \quad \mathbf{e}_{\ell+1}^{(j,\beta)} \leftarrow \chi_{\text{last}}^m,$$

$$\forall (i, j, \beta) \in [\ell] \times [\lambda] \times \{0, 1\}, \quad \tilde{\mathbf{t}}_i^{(j,\beta)} \leftarrow \mathbb{Z}_q^m,$$

$$\forall (j, \beta) \in \text{comm} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} \leftarrow \mathbb{Z}_q^m,$$

$$\forall \beta \in \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j^*, \beta)} \leftarrow \mathbb{Z}_q^m,$$

$$\forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} \leftarrow \mathbb{Z}_q^m.$$

(c) Let $\tilde{x} = x^L$. Next, it computes key vectors $\{\mathbf{t}_i^{(j,\beta)}\}_{i,j,\beta}$ as follows:

$$\begin{aligned} \forall (i, j, \beta) \in [\ell + 1] \times \text{comm} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, \\ \forall (i, j, \beta) \in [\ell] \times \text{diff} \times \{0, 1\}, \quad \mathbf{t}_i^{(j,\beta)} &= \tilde{\mathbf{t}}_i^{(j,\beta)} + \mathbf{e}_i^{(j,\beta)}, \\ \forall (j, \beta) \in \widehat{\text{diff}} \times \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j,\beta)} &= \tilde{\mathbf{t}}_{\ell+1}^{(j,\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}, \\ \forall \beta \in \{0, 1\}, \quad \mathbf{t}_{\ell+1}^{(j^*,\beta)} &= \tilde{\mathbf{t}}_{\ell+1}^{(j^*,\beta)} + \mathbf{e}_{\ell+1}^{(j^*,\beta)}. \end{aligned}$$

(d) Finally, it sends the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$.

• **Guess.** The adversary finally sends the guess γ' .

8.4.4. Indistinguishability of hybrid games in section 8.4.3. Now we show that the hybrid experiments described in section 8.4.3 are computationally indistinguishable. For any PPT adversary $\mathcal{A}$, let $p_{\mathcal{A},x}(\cdot)$ denote the probability that adversary $\mathcal{A}$ outputs $\gamma' = 1$ in Game x .

LEMMA 8.20. *There exists a negligible function $\text{negl}(\cdot)$ such that for any adversary $\mathcal{A}$ and $\lambda \in \mathbb{N}$, $p_{\mathcal{A},0}(\lambda) - p_{\mathcal{A},1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The proof of this lemma is identical to that of Lemma 8.8. $\square$

LEMMA 8.21. *For any adversary $\mathcal{A}$ and $\lambda \in \mathbb{N}$, $p_{\mathcal{A},1}(\lambda) = p_{\mathcal{A},2}(\lambda)$.*

Proof. The proof of this lemma is identical to that of Lemma 8.9. $\square$

LEMMA 8.22. *For any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3.1.1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. First, let us list the differences between Game 2 and Game 3.1.1. The setup, challenge phase, and ciphertext query are handled in an identical manner. The key queries, however, are handled differently. For each key query x , the challenger outputs $\{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell] \times [\lambda] \times \{0,1\}}$ as the secret key. The components $\{\mathbf{t}_1^{(j,\beta)}\}_{j \in \text{diff}, \beta \in \{0,1\}}$ are computed differently in Game 2 and Game 3.1.1. In particular, in Game 3.1.1, the challenger adds an additional error term $\tilde{\mathbf{e}}_1^{(j,\beta)} \leftarrow \chi_{\text{lwe}}^m$ in $\mathbf{t}_1^{(j,\beta)}$.

The proof of this lemma is similar to the proof of Lemma 8.10, (and uses the smudging lemma, Lemma 2.1). Therefore, $\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3.1.1}(\lambda) \leq q_{\text{keys}}(\lambda) \cdot (2 \cdot |\text{diff}| \cdot \text{negl}_{\text{smud}}(\lambda)) \leq q_{\text{keys}}(\lambda) \cdot (2\lambda \cdot \text{negl}_{\text{smud}}(\lambda))$, and the lemma follows by setting $\text{negl}_{3.1.1} = (2\lambda \cdot \text{negl}_{\text{smud}})$. $\square$

LEMMA 8.23. *For any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A},3.i^*.1}(\lambda) - \text{Adv}_{\mathcal{A},3.i^*.2}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. Let us first consider the differences between Game 3. i^* .1 and Game 3. i^* .2. The setup, challenge phase, and ciphertext query are handled in an identical manner in both games. The key generation queries are computed differently (in particular the components $\{\mathbf{t}_{i^*+1}^{(j,\beta)}\}_{j \in \text{diff}, \beta \in \{0,1\}}$ in each secret key).

The proof of this lemma is similar to the proof of Lemma 8.11, and the main idea is to use Fact 8.1 to argue that $\mathbf{e}_{i^*+1}^{(j,\beta)}$ drowns $\tilde{\mathbf{e}}_{i^*+1}^{(j,\beta)} \cdot \mathbf{U}_{i^*, \tilde{x}_{i^*}}^{(\beta)}$ and $\tilde{\mathbf{e}}_{i^*+1}^{(j,\beta)}$. $\square$

LEMMA 8.24. *Assuming the trapdoor generation algorithms LT_{en} satisfy the (q, σ_{pre}) -row removal property, for any PPT adversary $\mathcal{A}$ and $i^* \in [\ell]$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3.i^*.2}(\lambda) - \text{Adv}_{\mathcal{A},3.i^*.3}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. This proof is similar to the proof of Lemma 8.12, and we will be using the row removal property to prove it. We will first present the differences between the

two games and then discuss why the row removal property is applicable here. The exact reduction from the row removal property to indistinguishability of Game 3. i^* .2 and Game 3. i^* .3 can be found in the proof of Lemma 8.12.

Differences between Game 3. i^* .2 and Game 3. i^* .3:

1. Set $S^{(i^*)}$: In Game 3. i^* .2, the challenger sets $S^{(i^*)} = [\lambda] \times \{0, 1\}^2$, while in Game 3. i^* .3, $S^{(i^*)} = \text{comm} \times \{0, 1\}^2$ (tag^*, tag are chosen at the start of the security game, so the set diff is well defined here). Also, $\tilde{n}_{i^*} = \tilde{n} = (4\lambda + w)n$ in Game 3. i^* .2, while $\tilde{n}_{i^*} = \tilde{n} - |\text{diff}| \cdot 4n$ in Game 3. i^* .3.
2. $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{i=i^*}$ matrices: In Game 3. i^* .2, the challenger chooses $(\mathbf{M}_{i^*}, T_{i^*}) \leftarrow \text{EnTrapGen}(1^{\tilde{n}}, 1^m, q)$, while in Game 3. i^* .3, it chooses matrices $(\mathbf{M}_{i^*}, T_{i^*}) \leftarrow \text{EnTrapGen}(1^{\tilde{n}-|\text{diff}| \cdot 4n}, 1^m, q)$. As a result, in Game 3. i^* .2, it derives all matrices $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in [\lambda] \times \{0,1\}^2}$ from $\mathbf{M}_{i^*}$. In Game 3. i^* .3, the challenger chooses $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{j \in \text{diff}, b, \beta \in \{0,1\}}$ uniformly at random, while the remaining are derived from $\mathbf{M}_{i^*}$.
3. Ciphertexts: Since the set $S^{(i^*)}$ is different in both games, the challenge and query ciphertexts are constructed differently in both games. In particular, the challenge ciphertext components $(\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*)$ and the ciphertext query components $(\mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1})$ are computed using $\mathbf{M}_{i^*}$ and T_{i^*} , which are computed differently in Game 3. i^* .2 and Game 3. i^* .3.

Let us now discuss why the row removal property suffices for proving this lemma.

- Consider the four matrices $(\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*, \mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1})$. Fix any $j \in \text{diff}$, and let $\text{tag}_j^* = \beta$ and $\text{tag}_j = 1 - \beta$. Then, from the definitions of $\mathbf{D}_b^{(j,\beta)}$ and $\tilde{\mathbf{D}}_b^{(j,\beta)}$ in Mixed-SubEnc, it follows that
 - $\mathbf{B}_{i^*,0}^{(j,\beta)} \cdot \mathbf{U}_{i^*,0}^* = \mathbf{C}_{i^*,0}^{(j,\beta)}$, and the rest are mapped to random matrices;
 - $\mathbf{B}_{i^*,1}^{(j,\beta)} \cdot \mathbf{U}_{i^*,1}^* = \mathbf{C}_{i^*,1}^{(j,\beta)}$, and the rest are mapped to random matrices;
 - $\mathbf{B}_{i^*,0}^{(j,1-\beta)} \cdot \mathbf{U}_{i^*,0} = \mathbf{C}_{i^*,0}^{(j,1-\beta)}$, and the rest are mapped to random matrices;
 - $\mathbf{B}_{i^*,1}^{(j,1-\beta)} \cdot \mathbf{U}_{i^*,1} = \mathbf{C}_{i^*,1}^{(j,1-\beta)}$, and the rest are mapped to random matrices.
- Next, note that the $\{\mathbf{C}_{i^*,b}^{(j,\beta)}\}_{j \in \text{diff}, b, \beta \in \{0,1\}}$ are not used for responding to key generation queries.
- Using the above points, we can conclude that in Game 3. i^* .2, for each $j \in \text{diff}$, each of $(\mathbf{U}_{i^*,0}^*, \mathbf{U}_{i^*,1}^*, \mathbf{U}_{i^*,0}, \mathbf{U}_{i^*,1})$ maps $[\mathbf{B}_{i^*,0}^{(j,0)} \mid \mathbf{B}_{i^*,1}^{(j,0)} \mid \mathbf{B}_{i^*,0}^{(j,1)} \mid \mathbf{B}_{i^*,1}^{(j,1)}]$ to a uniformly random matrix.

The proof of this lemma therefore follows using the row removal property. $\square$

LEMMA 8.25. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ and $i^* \in [\ell]$ there exists a negligible function $\text{negl}_{3,i^*,\mathcal{A}}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3,i^*,3}(\lambda) - \text{Adv}_{\mathcal{A},3,i^*,4}(\lambda) \leq \text{negl}_{3,i^*,\mathcal{A}}(\lambda)$.*

Proof. The proof of this lemma is similar to the proof of Lemma 8.13, except that it involves more hybrid experiments.

In Game 3. i^* .3, for each key query x , for each $\beta \in \{0, 1\}$ and $j \in \text{diff}$, the component $\tilde{\mathbf{t}}_{i^*}^{(j,\beta)} = -\sum_{\alpha=1}^{i^*-1} (\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) - \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{i^*-1} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{i^*,\text{st}_{i^*}^{(\beta)}} + \mathbf{s}_{i^*}^{(j,\beta)} \cdot \mathbf{B}_{i^*,\tilde{x}_{i^*}}^{(j,\beta)} + \tilde{\mathbf{e}}_{i^*}^{(j,\beta)}$. In Game 3. i^* .4, $\tilde{\mathbf{t}}_{i^*}^{(j,\beta)} \leftarrow \mathbb{Z}_q^m$. Let $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ denote the number of keys queried by $\mathcal{A}(1^\lambda)$. To prove that these two games are computationally indistinguishable, we will define $q_{\text{keys}} \cdot \lambda$ hybrid experiments.

Hybrid $H_{o,\hat{j},0}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$, $\hat{j} \in [\lambda]$. In this hybrid, for the first o keys, for $j \in \text{diff} \cap [\hat{j}]$, the $\tilde{\mathbf{t}}_{i^*}^{(j,0)}$ components are sampled uniformly at random, while the remaining components are sampled as in Game 3.i*.3.

Hybrid $H_{o,\hat{j},1}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$, $\hat{j} \in [\lambda]$. In this hybrid, for all keys and $j \in \text{diff}$, the $\tilde{\mathbf{t}}_{i^*}^{(j,0)}$ components are sampled uniformly at random. For the first o queries and $j \in \text{diff} \cap [\hat{j}]$, the $\tilde{\mathbf{t}}_{i^*}^{(j,1)}$ components are sampled uniformly at random, while the remaining are sampled as in Game 3.i*.3.

Clearly, $H_{0,0,0}$ corresponds to Game 3.i*.3, $H_{q_{\text{keys}},\lambda,1}$ is identical to Game 3.i*.4, $H_{o-1,\lambda,b} \equiv H_{o,0,b}$, and $H_{q_{\text{keys}},\lambda,0} \equiv H_{0,0,1}$. Let $a_{\mathcal{A},i,\hat{j},b}(\lambda)$ denote the advantage of $\mathcal{A}$ in $H_{i,\hat{j},b}$.

CLAIM 8.26. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries there exists a negligible function $n_{o,\hat{j},0}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ and all indices $o \in [q_{\text{keys}}]$ and $\hat{j} \in [\lambda]$, $a_{\mathcal{A},o,\hat{j}-1,0} - a_{\mathcal{A},o,\hat{j},0} \leq n_{o,\hat{j},0}(\lambda)$.*

CLAIM 8.27. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries there exists a negligible function $n_{o,0,0}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ and all indices $o \in [q_{\text{keys}}]$ and $\hat{j} \in [\lambda]$, $a_{\mathcal{A},o,\hat{j}-1,1} - a_{\mathcal{A},o,\hat{j},1} \leq n_{o,\hat{j},1}(\lambda)$.*

The proofs of these claims are similar to the proof of Claim 8.14. If $\hat{j} \notin \text{diff}$, then $H_{o,\hat{j}-1,b} \equiv H_{o,\hat{j},b}$. Otherwise, we can reduce LWE to the indistinguishability of these two hybrids. $\square$

LEMMA 8.28. *For any PPT adversary $\mathcal{A}$ and $i^* \in [\ell - 1]$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3,i^*.4}(\lambda) - \text{Adv}_{\mathcal{A},3,(i^*+1).1}(\lambda) \leq \text{negl}(\lambda)$.*

This proof is identical to the proof of Lemma 8.22.

LEMMA 8.29. *For any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},3,\ell.4}(\lambda) - \text{Adv}_{\mathcal{A},4}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between Game 3.ℓ.4 and Game 4 is that the $\mathbf{t}_{\ell+1}^{(j^*,\beta)}$ terms contain an additional noise term, $\tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta)}$. The proof of this lemma is identical to the proof of Lemma 8.22 and follows via the smudging lemma (Lemma 2.1). $\square$

LEMMA 8.30. *Let $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ and $q : \mathbb{N} \rightarrow \mathbb{N}$ be functions, and $\chi_s(\lambda) \equiv \mathcal{D}_{\sqrt{2}\sigma(\lambda)}$ and $\chi_{\text{lwe}}(\lambda) \equiv \mathcal{D}_{\sigma(\lambda)}$ for each $\lambda \in \mathbb{N}$. Assuming the $\text{LWE-ss}_{(n,q,\sigma_{\text{lwe}})}$ assumption holds, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},4}(\lambda) - \text{Adv}_{\mathcal{A},5}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between Game 4 and Game 5 is in the key generation phase. In Game 4, for each query, the term $\mathbf{t}_{\ell+1}^{(j^*,\beta^*)} = -\sum_{\alpha=1}^{\ell} (\mathbf{t}_{\alpha}^{(j^*,\beta^*)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta^*)}) + \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\text{st}_{\ell+1}^{(\beta^*)}}^{(\beta^*)} + \tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta^*)} + \mathbf{e}_{\ell+1}^{(j^*,\beta^*)}$, and similarly the term $\mathbf{t}_{\ell+1}^{(j^*,1-\beta^*)} = \tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\text{st}_{\ell+1}^{(1-\beta^*)}}^{(1-\beta^*)} + \tilde{\mathbf{e}}_{\ell+1}^{(j^*,1-\beta^*)} + \text{other terms}$. In Game 5, for each key query, both $\mathbf{t}_{\ell+1}^{(j^*,0)}$ and $\mathbf{t}_{\ell+1}^{(j^*,1)}$ are set to be uniformly random. Using the short secrets version of LWE, we can switch both $\tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\text{st}_{\ell+1}^{(\beta^*)}}^{(\beta^*)} + \tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta^*)}$ and $\tilde{\mathbf{s}} \cdot \mathbf{P}_{\ell+1,\text{st}_{\ell+1}^{(1-\beta^*)}}^{(1-\beta^*)} + \tilde{\mathbf{e}}_{\ell+1}^{(j^*,1-\beta^*)}$ to uniformly random vectors. This switch is possible because

- $\tilde{\mathbf{s}}$ is chosen from χ_s^n , and $\tilde{\mathbf{e}}_{\ell+1}^{(j^*,\beta^*)}, \tilde{\mathbf{e}}_{\ell+1}^{(j^*,1-\beta^*)}$ are chosen from χ_{lwe}^m ;

- $\tilde{\mathbf{s}}, \tilde{\mathbf{e}}_{\ell+1}^{(j^*, \beta^*)}$, and $\tilde{\mathbf{e}}_{\ell+1}^{(j^*, 1-\beta^*)}$ are not required anywhere else in Game 4 or Game 5; each of these three terms is chosen afresh for each key query;
- $\mathbf{P}_{\ell+1, \text{st}_{\ell+1}}^{(\beta^*)}$ and $\mathbf{P}_{\ell+1, \text{st}_{\ell+1}}^{(1-\beta^*)}$ are uniformly random matrices.

Formally, we will show that Game 4 and Game 5 are computationally indistinguishable via a sequence of hybrid experiments. Let $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ denote the number of keys queried by $\mathcal{A}(1^\lambda)$. To prove that these two games are computationally indistinguishable, we will define q_{keys} hybrid experiments.

Hybrid H_o for $o \in \{0, 1, \dots, q_{\text{keys}}\}$. In this hybrid, for the first o keys, the $\mathbf{t}_{i^*+1}^{(j^*, 0)}, \mathbf{t}_{lv+1}^{(j^*, 1)}$ components are sampled uniformly at random in the first o queries. For the remaining $q_{\text{keys}} - o$ key queries, the keys are generated as in Game 4 in the remaining queries.

Clearly, H_0 corresponds to Game 4, while $H_{q_{\text{keys}}}$ is identical to Game 5. Let $a_{\mathcal{A}, i}(\lambda)$ denote the advantage of $\mathcal{A}$ in H_i .

CLAIM 8.31. *Let $\sigma : \mathbb{N} \rightarrow \mathbb{R}^+$ and $q : \mathbb{N} \rightarrow \mathbb{N}$ be functions, and $\chi_s(\lambda) \equiv \mathcal{D}_{\sqrt{2}\sigma(\lambda)}$ and $\chi_{\text{lwe}}(\lambda) \equiv \mathcal{D}_{\sigma(\lambda)}$ for each $\lambda \in \mathbb{N}$. Assuming the $\text{LWE-ss}_{(n, q, \sigma_{\text{lwe}})}$ assumption holds, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ and all indices $o \in [q_{\text{keys}}]$, $a_{\mathcal{A}, o-1} - a_{\mathcal{A}, o} \leq \text{negl}(\lambda)$.*

Proof. Suppose there exist an adversary making q_{keys} key queries, and a nonnegligible function $\eta(\cdot)$ such that for all $\lambda \in \mathbb{N}$, there exists an index $o \in [q_{\text{keys}}]$ such that $a_{\mathcal{A}, o-1} - a_{\mathcal{A}, o} \geq \eta(\lambda)$. We will use $\mathcal{A}$ to build a reduction algorithm $\mathcal{B}$ that breaks the $\text{LWE-ss}_{(n, q, \sigma_{\text{lwe}})}$ assumption.

The reduction algorithm receives (k, w, L) from the adversary and sets the parameters as in H_{o-1}/H_o . It makes $2m$ queries to the LWE-ss challenger and receives $\{(\mathbf{a}_j, u_j)\}_{j \leq 2m}$. It chooses $\text{tag}^*, \text{tag} \leftarrow \{0, 1\}^\lambda$, and j^* is the first position where tag^* and tag differ. Next, it sets $\tilde{n}_i$ as in H_{o-1}/H_o and chooses $(\mathbf{M}_i, T_i) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q)$.

Challenge phase. The reduction algorithm receives challenge ciphertext BP^* , which specifies the reject state rej^* , and uses Mixed-SubEnc for computing the challenge ciphertext. Note that Mixed-SubEnc chooses $\mathbf{P}_{\ell+1, v}$ uniformly at random. The reduction algorithm sets $\mathbf{P}_{\ell+1, \text{rej}^*}$ to be a matrix whose j th column is $\mathbf{a}_j^T$. All other $\mathbf{P}_{\ell+1, v}$ are chosen uniformly at random.

Ciphertext query. The reduction algorithm receives ciphertext query BP and uses Mixed-SubEnc for computing the ciphertext query. Let rej denote the reject state of BP . It sets $\mathbf{P}_{\ell+1, \text{rej}}$ to be a matrix whose j th column is $\mathbf{a}_{m+j}^T$. The remaining $\mathbf{P}_{\ell+1, v}$ matrices are chosen uniformly at random.

Key queries. The reduction algorithm first sets $\mathbf{u}^{\beta^*} = [u_1 \dots u_m]$ and $\mathbf{u}^{1-\beta^*} = [u_{m+1} \dots u_{2m}]$. For the first $o - 1$ key queries, the $\mathbf{t}_{\ell+1}^{(j^*, \beta^*)}$ components are chosen uniformly at random. For the o th key query, the reduction algorithm sets $\mathbf{t}_{\ell+1}^{(j^*, \beta^*)} = -\sum_{\alpha=1}^{\ell} (\tilde{\mathbf{t}}_{\alpha}^{(j^*, \beta^*)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta^*)}) + \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(\beta^*)} + \mathbf{e}_{\ell+1}^{(j^*, \beta^*)} + \mathbf{u}^{\beta^*}$ and $\mathbf{t}_{\ell+1}^{(j^*, 1-\beta^*)} = -\sum_{\alpha=1}^{\ell} (\tilde{\mathbf{t}}_{\alpha}^{(j^*, 1-\beta^*)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(1-\beta^*)}) + \tilde{\mathbf{y}} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta, \tilde{x}_{\delta}}^{(1-\beta^*)} + \mathbf{e}_{\ell+1}^{(j^*, 1-\beta^*)} + \mathbf{u}^{1-\beta^*}$. The remaining key queries are handled as in H_{o-1}/H_o .

Now, if all the u_j terms output by the LWE challenger are uniformly random, then $\mathbf{t}_{\ell+1}^{(j^*, \beta^*)}$ and $\mathbf{t}_{\ell+1}^{(j^*, 1-\beta^*)}$ are uniformly random, and hence the reduction algorithm simulates H_o . If each $u_j = \tilde{\mathbf{s}} \cdot \mathbf{a}_j^T + \tilde{\mathbf{e}}_j$, then the reduction algorithm simulates H_{o-1} . $\square$

LEMMA 8.32. *Assuming the trapdoor system LT_{en} satisfies the (q, σ_{pre}) -row removal property, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},5}(\lambda) - \text{Adv}_{\mathcal{A},5.1}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. Let us first consider the differences between Game 5 and Game 5.1.

- Set $S^{(1)}$: In Game 5, the challenger sets $S^{(1)} = \text{comm} \times \{0, 1\}^2$, while in Game 5.1, $S^{(1)} = \{(j, \beta, b) : j \in \text{comm} \wedge \beta = \text{tag}_j\}$ (tag^*, tag are chosen at the start of the security game, so these sets are well defined here). Also, $\tilde{n}_1 = (4|\text{comm}| + w)n$ in Game 5, while $\tilde{n}_1 = (2|\text{comm}| + w)n$ in Game 5.1.
- $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{i=1}$ matrices: In Game 5, the challenger chooses matrices $(\mathbf{M}_1, T_1) \leftarrow \text{EnTrapGen}(1^{(4|\text{comm}|+w)n}, 1^m, q)$, while in Game 3.i*.3, it chooses the matrices $(\mathbf{M}_1, T_1) \leftarrow \text{EnTrapGen}(1^{(2|\text{comm}|+w)n}, 1^m, q)$. As a result, in Game 5, it derives all $\{\mathbf{B}_{1,b}^{(j,\beta)}\}_{(j,\beta,b) \in \text{comm} \times \{0,1\}^2}$ from $\mathbf{M}_1$. In Game 5.1, the challenger chooses $\{\mathbf{B}_{1,b}^{(j,\beta)}\}_{j \in \text{comm}, \beta \neq \text{tag}_j, b \in \{0,1\}}$ uniformly at random, while the remaining are derived from $\mathbf{M}_1$.
- Ciphertexts: Since the set $S^{(1)}$ is different in both games, the challenge ciphertext components $(\mathbf{U}_{1,0}^*, \mathbf{U}_{1,1}^*)$ and the ciphertext query components $(\mathbf{U}_{1,0}, \mathbf{U}_{1,1})$ are computed using $\mathbf{M}_1$ and T_1 , which are computed differently in Game 5 and Game 5.1.

Let us now discuss why the row removal property suffices for proving this lemma. Consider any matrix $\mathbf{U} \in \{\mathbf{U}_{1,0}^*, \mathbf{U}_{1,1}^*, \mathbf{U}_{1,0}, \mathbf{U}_{1,1}\}$, fix any $j \in \text{comm}$, and let $\beta = \text{tag}_j$. Then, from the definitions of $\mathbf{D}_b^{(j,\beta)}$ and $\tilde{\mathbf{D}}_b^{(j,\beta)}$ in Mixed-SubEnc, it follows that $\mathbf{B}_{1,b}^{(j,1-\beta)} \cdot \mathbf{U}$ is a random matrix for both $b \in \{0, 1\}$ (because $\text{tag}_j^* = \beta$).

The proof of this lemma therefore follows using the row removal property (the reduction algorithm is similar to the one described in the proof of Lemma 8.12). $\square$

LEMMA 8.33. *Assuming the trapdoor system LT_{en} satisfies the (q, σ_{pre}) -row removal property, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A},5,(i^*-1)}(\lambda) - \text{Adv}_{\mathcal{A},5,\ell}(\lambda) \leq \text{negl}(\lambda)$.*

The proof of this lemma is identical to the proof of Lemma 8.32.

LEMMA 8.34. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},5,\ell}(\lambda) - \text{Adv}_{\mathcal{A},6.2}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The proof of this lemma is similar to the proof of Lemma 8.13.

In Game 5. ℓ , for each key query x , for each $j \in \text{comm}$, the components $\mathbf{t}_1^{(j,\beta)} = \mathbf{s}_1^{(j,\beta)} \mathbf{B}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\beta)}$ and $\mathbf{t}_2^{(j,\beta)} = -\mathbf{s}_1^{(j,\beta)} \cdot \mathbf{C}_{1,\tilde{x}_1}^{(j,\beta)} + \mathbf{s}_2^{(j,\beta)} \cdot \mathbf{B}_{2,\tilde{x}_2}^{(j,\beta)} + \mathbf{e}_2^{(j,\beta)}$. In Game 6.1, then $\mathbf{t}_1^{(j,1-\text{tag}_j)}, \mathbf{t}_2^{(j,1-\text{tag}_j)} \leftarrow \mathbb{Z}_q^m$. Let $q_{\text{keys}} = q_{\text{keys}}(\lambda)$ denote the number of keys queried by $\mathcal{A}(1^\lambda)$. To prove that these two games are computationally indistinguishable, we will define $q_{\text{keys}} \cdot \lambda$ hybrid experiments.

Hybrid $H_{o,\hat{j},0}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$, $\hat{j} \in [\lambda]$. In this hybrid, for the first o keys, for $j \in \text{comm} \cap [\hat{j}]$, the $\mathbf{t}_1^{(j,1-\text{tag}_j)}, \mathbf{t}_2^{(j,1-\text{tag}_j)}$ components are sampled uniformly at random, while the remaining components are sampled as in Game 5. ℓ .

Hybrid $H_{o,\hat{j},1}$ for $o \in \{0, 1, \dots, q_{\text{keys}}\}$, $\hat{j} \in [\lambda]$. This hybrid is similar to the previous one, except that for $j = \hat{j} + 1$ it adds an additional χ_{lwe} noise to $\mathbf{t}_1^{(j,1-\text{tag}_j)}$ and $\mathbf{t}_2^{(j,1-\text{tag}_j)}$.

Clearly, $H_{0,0,0}$ corresponds to **Game 5.ℓ**, $H_{q_{\text{keys}},\lambda,1}$ is identical to **Game 6.2**, and $H_{o-1,\lambda,0} \equiv H_{o,0,0}$. Let $a_{\mathcal{A},o,\hat{j},b}(\lambda)$ denote the advantage of $\mathcal{A}$ in $H_{o,\hat{j},b}$.

CLAIM 8.35. *For every PPT adversary $\mathcal{A}$ making q_{keys} queries, there exists a $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $o \in [q_{\text{keys}}]$, and $j \in [\lambda]$, $a_{\mathcal{A},o,\hat{j},0} - a_{\mathcal{A},o,\hat{j},1} \leq \text{negl}(\lambda)$.*

The proof of this claim follows via the smudging lemma (Lemma 2.1).

CLAIM 8.36. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ making $q_{\text{keys}}(\cdot)$ key queries there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $q_{\text{keys}} = q_{\text{keys}}(\lambda)$, and all indices $o \in [q_{\text{keys}}]$ and $\hat{j} \in [\lambda]$, $a_{\mathcal{A},o,\hat{j}-1,1} - a_{\mathcal{A},o,\hat{j},0} \leq \text{negl}(\lambda)$.*

Proof. The proof of this claim follows from the LWE assumption. The reduction algorithm makes $4m$ queries to the LWE challenger and receives $\{\mathbf{a}_j, u_j\}_{j \in 4m}$. It sets $\mathbf{B}_{1,0}^{(\hat{j},1-\text{tag}_j)} = [\mathbf{a}_1^T \dots \mathbf{a}_m^T]$, $\mathbf{C}_{1,0}^{(\hat{j},1-\text{tag}_j)} = [\mathbf{a}_{m+1}^T \dots \mathbf{a}_{2m}^T]$, $\mathbf{B}_{1,1}^{(\hat{j},1-\text{tag}_j)} = [\mathbf{a}_{2m+1}^T \dots \mathbf{a}_{3m}^T]$, $\mathbf{C}_{1,1}^{(\hat{j},1-\text{tag}_j)} = [\mathbf{a}_{3m+1}^T \dots \mathbf{a}_{4m}^T]$. It also sets $\mathbf{u}_{1,0} = [u_1 \dots u_m]$, $\mathbf{u}_{2,0} = [u_{m+1} \dots u_{2m}]$, $\mathbf{u}_{1,1} = [u_{2m+1} \dots u_{3m}]$, $\mathbf{u}_{2,1} = [u_{3m+1} \dots u_{4m}]$.

For the o th key query x , the reduction algorithm sets $\mathbf{t}_1^{(\hat{j},1-\text{tag}_j)} = \mathbf{u}_{1,\tilde{x}_1} + \mathbf{y}^{(\hat{j})} + \mathbf{e}_1^{(\hat{j},1-\text{tag}_j)}$ and $\mathbf{t}_2^{(\hat{j},1-\text{tag}_j)} = \mathbf{u}_{2,\tilde{x}_1} + \mathbf{e}_2^{(\hat{j},1-\text{tag}_j)}$. The rest of the key components can be handled without the LWE challenge terms. $\square$

LEMMA 8.37. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in \{3, \dots, \ell\}$, $\text{Adv}_{\mathcal{A},6,i^*-1}(\lambda) - \text{Adv}_{\mathcal{A},6,i^*}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between **Game 6.(\(i^* - 1\))** and **Game 6. \(i^*\)** is with respect to the key queries. In **Game 6.(\(i^* - 1\))**, for each $j \in \text{comm}$, the challenger sets $\mathbf{t}_{i^*}^{(j,1-\text{tag}_j)} = -\mathbf{s}_{i^*-1}^{(j,1-\text{tag}_j)} \cdot \mathbf{C}_{i^*-1,\tilde{x}_{i-1}}^{(j,1-\text{tag}_j)} + \mathbf{s}_{i^*}^{(j,1-\text{tag}_j)} \cdot \mathbf{B}_{i^*,\tilde{x}_i}^{(j,1-\text{tag}_j)} + \mathbf{e}_{i^*}^{(j,1-\text{tag}_j)}$, while it switches these terms to random in **Game 6. \(i^*\)**.

The proof of this lemma uses (standard) LWE, similar to the proof of Lemma 8.34. One minor difference between this lemma and Lemma 8.34 is that in the previous lemma, the challenger switches both $\mathbf{t}_1^{(j,1-\text{tag}_j)}$ and $\mathbf{t}_2^{(j,1-\text{tag}_j)}$ (this is because the vector $\mathbf{s}_1^{(j,1-\text{tag}_j)}$ is used for computing both of these components). However, in this lemma, the reduction algorithm sets the LWE challenge's public vectors as the rows of $\mathbf{C}_{i^*,0}^{(j,1-\text{tag}_j)}$ and $\mathbf{C}_{i^*,1}^{(j,1-\text{tag}_j)}$, and the LWE challenge is used for setting $\mathbf{t}_{i^*}^{(j,1-\text{tag}_j)}$. $\square$

LEMMA 8.38. *For any adversary $\mathcal{A}$ and $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},6,(\ell+1)}(\lambda) = \text{Adv}_{\mathcal{A},7,1}(\lambda)$.*

Proof. The only difference between **Game 6.(\(\ell + 1 \))** and **Game 7.1** is with respect to the $\{\mathbf{t}_1^{(j,\text{tag}_j)}\}_{j \in \text{comm}}$ components in key queries. In **Game 6**, for each key query x , the challenger chooses $\{\mathbf{y}^{(j)}\}_{j \neq j^*}$ and sets $\mathbf{t}_1^{(j,\text{tag}_j)} = \mathbf{s}_1^{(j,\text{tag}_j)} \cdot \mathbf{B}_{1,\tilde{x}_1}^{(j,\text{tag}_j)} + \mathbf{y}^{(j)} + \mathbf{e}_1^{(j,\text{tag}_j)}$ for each $j \in \text{comm}$. In **Game 7.1**, the $\{\mathbf{t}_1^{(j,\text{tag}_j)}\}_{j \in \text{comm}}$ vectors are set to be uniformly random vectors.

Note that in **Game 6.(\(\ell + 1 \))**, the $\mathbf{y}^{(j)}$ terms are chosen afresh for each key, and $\mathbf{y}^{(j)}$ is only used in constructing $\mathbf{t}_1^{(j,\text{tag}_j)}$ (recall $\mathbf{t}_1^{(j,1-\text{tag}_j)}$ is uniformly random). As a result, the components $\{\mathbf{t}_1^{(j,\text{tag}_j)}\}_{j \in \text{comm}}$ are uniformly random vectors, and therefore the secret keys in the two games are identically distributed. $\square$

LEMMA 8.39. *Assuming the $\text{LWE}_{n,q,\sigma_{\text{lwe}}}$ assumption holds, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in \{2, \dots, \ell\}$,*

$$\text{Adv}_{\mathcal{A},7,i^*-1}(\lambda) - \text{Adv}_{\mathcal{A},7,i^*}(\lambda) \leq \text{negl}(\lambda).$$

Proof. The only difference between Game 7.(\$i^* - 1\$) and Game 7.\$i^*\$ is with respect to the \$\{\mathbf{t}_{i^*}^{(j,\text{tag}_j)}\}_{j \in \text{comm}}\$ components in key queries. In Game 7.(\$i^* - 1\$), for each key query \$x\$, the challenger sets \$\mathbf{t}_{i^*}^{(j,\text{tag}_j)} = -\mathbf{s}_{i^*-1}^{(j,\text{tag}_j)} \cdot \mathbf{C}_{i^*,\hat{x}_{i^*-1}}^{(j,\text{tag}_j)} + \mathbf{s}_{i^*}^{(j,\text{tag}_j)} \cdot \mathbf{B}_{i^*,\hat{x}_{i^*}}^{(j,\text{tag}_j)} + \mathbf{e}_{i^*}^{(j,\text{tag}_j)}\$ for each \$j \in \text{comm}\$. In Game 7.\$i^*\$, the \$\{\mathbf{t}_{i^*}^{(j,\text{tag}_j)}\}_{j \in \text{comm}}\$ vectors are set to be uniformly random vectors. We will show that these two games are computationally indistinguishable via a hybrid argument. First, we will define \$q_{\text{keys}} \cdot \lambda\$ hybrid experiments.

Hybrid \$H_{o,\hat{j},0}\$ for \$o \in \{0, 1, \dots, q_{\text{keys}}\}\$, \$\hat{j} \in [\lambda]\$. In this hybrid, for the first \$o\$ keys, for \$j \in \text{comm} \cap [\hat{j}]\$, the \$\mathbf{t}_{i^*}^{(j,\text{tag}_j)}\$ components are sampled uniformly at random, while the remaining components are sampled as in Game 7.\$i^*\$.

Hybrid \$H_{o,\hat{j},1}\$ for \$o \in \{0, 1, \dots, q_{\text{keys}}\}\$, \$\hat{j} \in [\lambda]\$. This hybrid is similar to the previous one, except that for \$j = \hat{j} + 1\$ it adds an additional \$\chi_{\text{lwe}}\$ noise to \$\mathbf{t}_{i^*}^{(j,\text{tag}_j)}\$.

Clearly, \$H_{0,0,0}\$ corresponds to Game 7.(\$i^* - 1\$), \$H_{q_{\text{keys}},\lambda,1}\$ is identical to Game 7.\$i^*\$, and \$H_{o-1,\lambda,0} \equiv H_{o,0,0}\$. Let \$a_{\mathcal{A},o,\hat{j},b}(\lambda)\$ denote the advantage of \$\mathcal{A}\$ in \$H_{o,\hat{j},b}\$.

CLAIM 8.40. *For every PPT adversary \$\mathcal{A}\$ making \$q_{\text{keys}}\$ queries, there exists a \$\text{negl}(\cdot)\$ such that for all \$\lambda \in \mathbb{N}\$, \$o \in [q_{\text{keys}}]\$ and \$j \in [\lambda]\$, \$a_{\mathcal{A},o,\hat{j},0} - a_{\mathcal{A},o,\hat{j},1} \leq \text{negl}(\lambda)\$.*

The proof of this claim follows via the smudging lemma (Lemma 2.1).

CLAIM 8.41. *Assuming the \$\text{LWE}_{n,q,\sigma_{\text{lwe}}}\$ assumption holds, for any PPT adversary \$\mathcal{A}\$ making \$q_{\text{keys}}(\cdot)\$ key queries there exists a negligible function \$\text{negl}(\cdot)\$ such that for all \$\lambda \in \mathbb{N}\$, \$q_{\text{keys}} = q_{\text{keys}}(\lambda)\$ and all indices \$o \in [q_{\text{keys}}]\$ and \$\hat{j} \in [\lambda]\$, \$a_{\mathcal{A},o,\hat{j}-1,1} - a_{\mathcal{A},o,\hat{j},0} \leq \text{negl}(\lambda)\$.*

Proof. The proof of this claim follows from the LWE assumption, where the reduction algorithm sets the LWE public vectors to be columns of \$\mathbf{C}_{i^*,0}^{(\hat{j},\text{tag}_{\hat{j}})}, \mathbf{C}_{i^*,1}^{(\hat{j},\text{tag}_{\hat{j}})}\$, and the LWE challenge is used to set \$\mathbf{t}_{i^*}^{(\hat{j},\text{tag}_{\hat{j}})}\$. Note that the vector \$\mathbf{s}_{i^*-1}^{(\hat{j},\text{tag}_{\hat{j}})}\$ is used only for defining \$\mathbf{t}_{i^*}^{(\hat{j},\text{tag}_{\hat{j}})}\$. This is because this vector is chosen afresh for each key query, and in hybrids \$H_{o,\hat{j}-1,1}/H_{o,\hat{j},0}\$, the key component \$\mathbf{t}_{i^*-1}^{(\hat{j},\text{tag}_{\hat{j}})}\$ is already random.³⁰ \$\square\$

Recall that Game 8 is identical to Game 7.(\$\ell + 1\$).

LEMMA 8.42. *Assuming \$\text{LT}_{\text{en}}\$ satisfies the \$(q, \chi_{\text{appr}}, \sigma_{\text{pre}})\$-target switching property and \$(q, \sigma_{\text{pre}})\$-well-sampledness of preimage, for any PPT adversary \$\mathcal{A}\$ there exists a negligible function \$\text{negl}(\cdot)\$ such that for all \$\lambda \in \mathbb{N}\$ and \$i^* \in [\ell]\$, \$\text{Adv}_{\mathcal{A},8,(i^*-1)}(\lambda) - \text{Adv}_{\mathcal{A},8,i^*}(\lambda) \leq \text{negl}(\lambda)\$.*

Proof. First, let us discuss the differences between Game 8.(\$i^* - 1\$) and Game 8.\$i^*\$. In Game 8.(\$i^* - 1\$), the challenge ciphertext components \$\{\mathbf{U}_{i,b}^*\}_{i \geq i^*, b \in \{0,1\}}\$ and query ciphertext components \$\{\mathbf{U}_{i,b}\}_{i \geq i^*, b \in \{0,1\}}\$ are computed using Mixed-SubEnc, while the remaining are chosen from Gaussian with parameter \$\chi_{\text{pre}}\$. Game 8.\$i^*\$ is similar to Game 8.(\$i^* - 1\$), except for the challenge ciphertext components \$\mathbf{U}_{i^*,0}^*, \mathbf{U}_{1,1}^*\$ and the query ciphertext components \$\mathbf{U}_{1,0}, \mathbf{U}_{1,1}\$ are chosen from the Gaussian distribution with parameter \$\sigma_{\text{pre}}\$. To show that these games are indistinguishable, we will define a hybrid experiment \$H\$.

³⁰If \$\mathbf{t}_{i^*-1}^{(\hat{j},\text{tag}_{\hat{j}})}\$ was not already switched to random, then \$\mathbf{s}_{i^*-1}^{(\hat{j},\text{tag}_{\hat{j}})}\$ would have been used to define it.

Hybrid H . This hybrid is similar to **Game 8.** $(i^* - 1)$, except that the challenger computes $\{\mathbf{U}_{i^*,b}^*, \mathbf{U}_{i^*,b}\}_{b \in \{0,1\}}$ such that they map $\mathbf{M}_{i^*}$ to uniformly random matrices.

Let $\text{Adv}_{\mathcal{A},H}$ denote the advantage of adversary $\mathcal{A}$ in hybrid H .

CLAIM 8.43. *Assuming LT_{en} satisfies the $(q, \chi_{\text{appr}}, \sigma_{\text{pre}})$ -target switching property, for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$ and $i^* \in [\ell]$, $\text{Adv}_{\mathcal{A},8.(i^*-1)}(\lambda) - \text{Adv}_{\mathcal{A},H}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The proof of this claim is similar to the proof of Lemma 8.17. First, let us discuss the reasons why the target switching property is applicable here. Let $S^{(i^*)}$ be defined as in **Game 8.** $(i^* - 1)$ /**Game 8.** i^* , with $\text{BP}^* = \{\pi_{i,b}^*\}_{(i,b) \in [\ell] \times \{0,1\}}$ and $\text{BP} = \{\pi_{i,b}\}_{(i,b) \in [\ell] \times \{0,1\}}$ the challenge/query programs.

1. In both **Game 8.** $(i^* - 1)$ and hybrid H , the components $\{\mathbf{U}_{i,b}^*, \mathbf{U}_{i,b}\}$ are chosen from a Gaussian distribution, and therefore these terms do not contain any information about the $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ or $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i^*)}}$ matrices.
2. The components $\{\mathbf{U}_{i,b}^*, \mathbf{U}_{i,b}\}_{i > i^*, b \in \{0,1\}}$ do not contain any information about the $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ or $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i^*)}}$ matrices (this follows from the construction).
3. The keys are all either random vectors or computed in terms of the challenge/query ciphertext components, and therefore do not explicitly require $\{\mathbf{P}_{i^*,v}\}_{v \in [w]}$ or $\{\mathbf{B}_{i^*,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i^*)}}$ matrices.

Consider matrices $\{\mathbf{Z}_{0,b}^*, \mathbf{Z}_{1,b}^*, \mathbf{Z}_{0,b}, \mathbf{Z}_{1,b}\}_{b \in \{0,1\}}$ defined as follows:

$$\mathbf{Z}_{0,b}^* = \begin{bmatrix} \left\{ \mathbf{C}_{i^*,b}^{(j,\beta)} \right\}_{(j,\beta,b) \in S^{(i^*)}} \\ \left\{ \mathbf{P}_{i^*, \pi_{i^*,b}^*(v)} \right\}_{v \in [w]} \end{bmatrix} \quad \mathbf{Z}_{1,b}^* = [\leftarrow \mathbb{Z}_q^{\tilde{n}_i \times m}],$$

$$\mathbf{Z}_{0,b} = \begin{bmatrix} \left\{ \mathbf{C}_{i^*,b}^{(j,\beta)} \right\}_{(j,\beta,b) \in S^{(i^*)}} \\ \left\{ \mathbf{P}_{i^*, \pi_{i^*,b}(v)} \right\}_{v \in [w]} \end{bmatrix} \quad \mathbf{Z}_{1,b} = [\leftarrow \mathbb{Z}_q^{\tilde{n}_i \times m}].$$

The reduction algorithm sends $(1^{\tilde{n}_i}, 1^m, \emptyset)$ to the target switching property challenger.³¹ It does not receive any matrix from the challenger (since the challenge set is empty). Next, it chooses $(M_i, T_i) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q)$ for all $i > i^*$, and parses $\mathbf{M}_i$ as in **Game 8.** $(i^* - 1)$ /**Game 8.** i^* to obtain $\{\mathbf{B}_{i,b}^{(j,\beta)}\}_{(j,\beta,b) \in S^{(i)}}$ and $\{\mathbf{P}_{i,v}\}_{v \in [w]}$ for all $i > i^*$. It receives BP^* as the challenge query from the adversary. The reduction algorithm sends $\mathbf{Z}_{0,b}^*, \mathbf{Z}_{1,b}^*$ to the target switching property challenger and receives $\mathbf{U}_{i^*,b}$ in response. It chooses the remaining components as in **Game 8.** $(i^* - 1)$ /**Game 8.** i^* and sends the challenge ciphertext to the adversary.

Next, it receives the ciphertext query BP . It sends $\mathbf{Z}_{0,b}, \mathbf{Z}_{1,b}$ to the target switching property challenger and receives $\mathbf{U}_{i^*,b}$. The remaining ciphertext components are chosen as in **Game 8.** $(i^* - 1)$ /**Game 8.** i^* , and the reduction algorithm sends the challenge ciphertext to the adversary. Finally, the adversary makes key queries. For each key query, the reduction algorithm sets $\{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}}$ as in **Game 8.** $(i^* - 1)$ /**Game 8.** i^* and sends them to the adversary. The adversary sends its guess, which the reduction algorithm forwards to the challenger.

³¹Note that the set specified by the adversary in the target switching property game can be empty.

Therefore, if there exist a PPT adversary $\mathcal{A}$ and a nonnegligible function η such that $\text{Adv}_{\mathcal{A},8.(i^*-1)}(\lambda) - \text{Adv}_{\mathcal{A},H}(\lambda) \geq \eta(\lambda)$ for all λ , then there exists a PPT algorithm $\mathcal{B}$ that breaks the target switching property. $\square$

CLAIM 8.44. *Assuming LT_{en} satisfies the (q, σ_{pre}) -well-sampledness of the preimage, for any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},H}(\lambda) - \text{Adv}_{\mathcal{A},8.i^*}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. This proof follows directly from the (q, σ) -well-sampledness of the preimage property. Suppose there exist a PPT adversary $\mathcal{A}$ and a reduction algorithm $\eta(\cdot)$ such that $\text{Adv}_{\mathcal{A},H}(\lambda) - \text{Adv}_{\mathcal{A},8.i^*}(\lambda) \geq \eta(\lambda)$ for all $\lambda \in \mathbb{N}$. Then there exists a reduction algorithm that breaks the (q, σ) -well-sampledness of the preimage property.

The reduction algorithm sends $1^{\tilde{n}_i}, 1^m, 1^{4m}$ to the challenger. Note that $m > \tilde{n}_i \log q + \lambda$ and $\sigma > \sqrt{n \cdot \log q \cdot \log m} + \lambda$, as required. It receives a matrix $\mathbf{U} \in Z_q^{m \times 4m}$, which it parses as $\mathbf{U} = [\mathbf{U}_{i^*,0}^* \mid \mathbf{U}_{i^*,1}^* \mid \mathbf{U}_{i^*,0} \mid \mathbf{U}_{i^*,1}]$. The reduction algorithm also chooses $(\mathbf{M}_i, T_i) \leftarrow \text{EnTrapGen}(1^{\tilde{n}_i}, 1^m, q)$ for all $i \neq i^*$.

On receiving the challenge ciphertext, it chooses the remaining ciphertext components as in Game 8. $(i^* - 1)$ /Game 8. i^* and sends them to the adversary. Similarly, it handles the ciphertext query. Finally, for the key queries, the reduction algorithm handles them as in Game 8. $(i^* - 1)$ /Game 8. i^* . $\square$

Using these two claims, it follows that $\text{Adv}_{\mathcal{A},8.(i^*-1)}(\lambda) - \text{Adv}_{\mathcal{A},8.i^*}(\lambda)$ is bounded by a negligible function. $\square$

LEMMA 8.45. *For any PPT adversary $\mathcal{A}$ there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},8.\ell}(\lambda) - \text{Adv}_{\mathcal{A},9}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between these two hybrids is with respect to the key components $\{\mathbf{t}_{\ell}^{(j,\beta)}\}_{j \in \widehat{\text{diff}}, \beta \in \{0,1\}}$. In Game 8. ℓ , these vectors are computed as $\sum_{\alpha=1}^{\ell} (\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}$, while in Game 9, this vector is set to be $\sum_{\alpha=1}^{\ell} (\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) + \mathbf{y}^{(j)} \cdot \prod_{\delta=1}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} \cdot \prod_{\delta=2}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}$, where $\tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} \leftarrow \chi_{\text{lwe}}^m$. Note that the $\mathbf{U}_{i,b}$ matrices are all drawn from the Gaussian distribution with parameter σ_{pre} , and therefore, with all but negligible probability, $\|\tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} \cdot \prod_{\delta=2}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}\| \leq m\sigma_{\text{lwe}} \cdot (m\sigma_{\text{pre}})^{\ell-1}$. Since $\sigma_{\text{last}}/m\sigma_{\text{lwe}} \cdot (m\sigma_{\text{pre}})^{\ell-1} \geq 2^{\lambda}$, we can use the smudging lemma to argue that the statistical distance between these two games is at most $m \cdot q_{\text{keys}} \cdot \text{negl}_{\text{smud}}(\lambda)$. $\square$

LEMMA 8.46. *Assuming the $\text{LWE-sp}_{d,q,\sigma_{\text{pre}},\chi_{\text{lwe}}}$ assumption (Assumption 3) holds (where $d(\lambda) = 6n \cdot \log q$), for any PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $\text{Adv}_{\mathcal{A},9}(\lambda) - \text{Adv}_{\mathcal{A},10}(\lambda) \leq \text{negl}(\lambda)$.*

Proof. The only difference between Game 9 and Game 10 is with respect to the key components $\{\mathbf{t}_{\ell+1}^{(j,\beta)}\}_{j \in \widehat{\text{diff}}, \beta \in \{0,1\}}$. In Game 9, for each key, these are computed as $\mathbf{t}_{\ell+1}^{(j,\beta)} = \sum_{\alpha=1}^{\ell} (\tilde{\mathbf{t}}_{\alpha}^{(j,\beta)} \cdot \prod_{\delta=\alpha}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)}) + (\mathbf{y}^{(j)} \cdot \mathbf{U}_{1,\tilde{x}_1}^{(\beta)} + \tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)}) \prod_{\delta=2}^{\ell} \mathbf{U}_{\delta,\tilde{x}_{\delta}}^{(\beta)} + \mathbf{e}_{\ell+1}^{(j,\beta)}$, while in Game 10, these vectors are uniformly random. To prove this claim, it suffices to switch $(\mathbf{y}^{(j)} \cdot \mathbf{U}_{1,\tilde{x}_1}^{(\beta)} + \tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)})$ to a uniformly random vector. Since $\mathbf{y}^{(j)} \leftarrow \mathbb{Z}_q^m$, $\mathbf{U}_{1,\tilde{x}_1}^{(\beta)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_{\text{pre}}}^{m \times m}$, and $\tilde{\mathbf{e}}_{\ell+1}^{(j,\beta)} \leftarrow \chi_{\text{lwe}}^m$, we can use the LWE-sp assumption as in the proof of Theorem 7.19. $\square$

8.4.5. Proving 1-bounded restricted accept indistinguishability. First, note that by using the lemmas provided in section 8.4.4, we can conclude that our construction satisfies 1-bounded complete accept indistinguishability security. Con-

cretely, algorithms SK-Enc^* and KeyGen^* , which take as input the public parameters, are defined as follows: SK-Enc^* is the same as the standard encryption algorithm Enc (that is, it outputs random Gaussian matrices), and KeyGen^* on input x outputs the secret key as $(x, \{\mathbf{t}_i^{(j,\beta)}\}_{(i,j,\beta) \in [\ell+1] \times [\lambda] \times \{0,1\}})$, where all $\mathbf{t}_i^{(j,\beta)}$ are sampled uniformly at random from $\mathbb{Z}_q^m$. Since in Game 10 the challenger is already using SK-Enc^* and KeyGen^* to answer corresponding queries, therefore, using lemmas in section 8.4.4, we can argue 1-bounded complete accept indistinguishability security.

Lastly, to finish the proof we only need to argue that the probability adversary outputs 1 in the 1-bounded *restricted* accept indistinguishability security game—when the challenger computes challenge ciphertext as a normal functional encryption ciphertext instead of encrypting BP^* —is negligibly close to the probability adversary outputs 1 in Game 10. Note that this again follows from the lemmas in section 8.4.4, or, in other words, it follows from the fact that our construction satisfies 1-bounded complete accept indistinguishability security. This is because if an adversary can distinguish Game 10 from the scenario described above, then we could come up with a reduction algorithm that breaks the 1-bounded complete accept indistinguishability security of our construction.

The idea is straightforward. The reduction algorithm will simply forward messages (back and forth) between the attacker and the complete accept indistinguishability challenger, except with the following changes:

- The reduction algorithm does not forward the adversary's challenge program BP^* as its challenge query to the challenger of the complete accept indistinguishability game. Instead it runs the normal encryption algorithm Enc and sends the output back to the adversary as its challenge ciphertext.
- In addition, the reduction algorithm sends the adversary's post-challenge encryption query (if any) to the challenger as its challenge query and forwards the challenger's response to the adversary.
- Also, the reduction algorithm does not make any post-challenge encryption query. (Note that key queries are answered as before.)
- Finally, it outputs whatever the adversary outputs.

Clearly the reduction algorithm perfectly simulates the indistinguishability experiment (between Game 10 and the scenario described above); thus if the adversary's advantage is nonnegligible, then the reduction algorithm also breaks 1-bounded complete accept indistinguishability security with nonnegligible probability. This completes the proof.

Appendix A. Background: Attribute-based encryption.

A.1. Key-policy attribute-based encryption. A key-policy attribute-based encryption (KP-ABE) scheme $\mathcal{ABE}$, for a set of attribute spaces $\mathcal{X} = \{\mathcal{X}_\kappa\}_\kappa$, predicate classes $\mathcal{C} = \{\mathcal{C}_\kappa\}_\kappa$, and message spaces $\mathcal{M} = \{\mathcal{M}_\kappa\}_\kappa$, consists of four polytime algorithms (Setup, Enc, KeyGen, Dec) with the following syntax:

- $\text{Setup}(1^\lambda, 1^\kappa) \rightarrow (\text{pp}, \text{msk})$. The setup algorithm takes as input the security parameter λ and a functionality index κ . It 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 $x \in \mathcal{X}_\kappa$, and a message $m \in \mathcal{M}_\kappa$. 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 $C \in \mathcal{C}_\kappa$. 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 $m \in \mathcal{M}_\kappa$ or a special symbol $\perp$.

Correctness. A KP-ABE scheme is said to be correct if there exist negligible functions $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ such that for all $\lambda, \kappa \in \mathbb{N}$, for all $x \in \mathcal{X}_\kappa$, $C \in \mathcal{C}_\kappa$, $m \in \mathcal{M}_\kappa$, the following holds:

$$\begin{aligned} C(x) = 1 &\Rightarrow \Pr \left[\begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ \text{Dec}(\text{sk}_C, \text{ct}) = m : \quad \text{sk}_C \leftarrow \text{KeyGen}(\text{msk}, C); \\ \quad \quad \quad \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m) \end{array} \right] \geq 1 - \text{negl}_1(\lambda), \\ C(x) = 0 &\Rightarrow \Pr \left[\begin{array}{l} (\text{pp}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa); \\ \text{Dec}(\text{sk}_C, \text{ct}) = \perp : \quad \text{sk}_C \leftarrow \text{KeyGen}(\text{msk}, C); \\ \quad \quad \quad \text{ct} \leftarrow \text{Enc}(\text{pp}, x, m) \end{array} \right] \geq 1 - \text{negl}_2(\lambda). \end{aligned}$$

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 KP-ABE 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 for every $\lambda \in \mathbb{N}$ the following holds:

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

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

In this work, we only require the scheme to achieve selective security, which is formally defined as follows.

DEFINITION A.2. A KP-ABE scheme $\mathcal{ABE} = (\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec})$ is said to be selectively secure if for every stateful PPT adversary $\mathcal{A}$, there exists a negligible function $\text{negl}(\cdot)$, such that for every $\lambda \in \mathbb{N}$ the following holds:

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

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

Appendix B. ABE and Mixed FE to PLBE: Preserving perfect correctness. In this section, we give an alternate construction for constructing PLBE such that if the underlying ABE scheme achieves *perfect* correctness, then so does the PLBE scheme, even if the Mixed FE scheme is *not* perfectly correct. Since existing ABE schemes [GVW13, BGG+14] can be made perfectly correct by appropriately truncating noise distributions used, this gives a pathway to get perfect correctness under LWE. Note that the construction described in section 6.1 only achieves perfect correctness when both underlying ABE and Mixed FE schemes are perfectly correct.

This is because the policy circuit is the Mixed FE decryption circuit, and thus in order to guarantee perfect correctness, we need the minimum requirement that the Mixed FE normal ciphertexts always decrypt to 1. Below we give the main idea to obtain perfect correctness.

Outline. At a very high level, the idea is to encrypt the message m under two independent ABE systems such that at least one of the ciphertext components can always be decrypted to obtain the underlying message. To this end, during setup we sample two ABE key pairs $(\text{abe.pp}_b, \text{abe.msk}_b)$ (for $b \in \{0, 1\}$) and a Mixed FE key pair $(\text{Mixed.pp}, \text{Mixed.msk})$. To generate the secret key for the i th user, we generate a Mixed FE secret key Mixed.sk_i for message i , and later compute two ABE keys $\text{abe.sk}_{i,0}, \text{abe.sk}_{i,1}$ for predicates $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$ and $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$ using $\text{abe.msk}_0, \text{abe.msk}_1$, respectively. Here $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$ denotes the circuit that first decrypts the input using key Mixed.sk_i and later applies a “not” gate (i.e., outputs the complement). Now the PLBE ciphertexts will consist of two parts, one for each ABE subsystem. For PLBE normal encryption, one computes two ciphertexts ct_b (for $b \in \{0, 1\}$) as encryptions of message m under attributes ct_{attr} using parameters abe.pp_b , where ct_{attr} is computed as before. Now for encrypting a message to index i , the encryption algorithm behaves differently in that it computes ct_0 as before, but ct_1 will now be an encryption of message 0 under the same attributes. The reason for not encrypting the message m in the second component of the index ciphertext becomes clear while proving security.

Now for arguing perfect correctness, we observe that it should be the case that $\text{Mixed.Dec}(\text{Mixed.sk}_i, \text{ct}_{\text{attr}})$ equals either 0 or 1. Thus, at least one of the PLBE normal ciphertext components could be correctly decrypted. Note that for such an argument we only require that ABE be perfectly correct. Next, the security proof is similar to that in section 6.3, except that when arguing normal hiding security of our construction we need to rely on both the ABE security as well as the weak accept indistinguishability property of the Mixed FE scheme. The main idea is that by correctness of the functional encryption scheme, we can say that with all but negligible probability the attribute used in the second component of the challenge ciphertext is not satisfied by any of the ABE keys queried. Thus, as our first hybrid argument, we could use ABE security to switch the second challenge ciphertext component to an encryption of 0 instead of message m . The remaining proof is identical to as before, with the only modification being that the reduction algorithm needs to generate the ABE keys for the second component on its own during the entire reduction. Below we describe our construction $\text{PLBE} = (\text{Setup}, \text{Enc}, \text{Enc-index}, \text{Dec})$ for messages spaces $\{\mathcal{M}_\kappa\}_\kappa$ in detail.

B.1. Construction. Let $\mathcal{ABE} = (\text{ABE.Setup}, \text{ABE.Enc}, \text{ABE.KeyGen}, \text{ABE.Dec})$ be a KP-ABE scheme for a set of attribute spaces $\{\mathcal{X}_\kappa\}_\kappa$, predicate classes $\{\mathcal{C}_\kappa\}_\kappa$, and message spaces $\{\mathcal{M}_\kappa\}_\kappa$, and let $\text{Mixed-FE} = (\text{Mixed.Setup}, \text{Mixed.Enc}, \text{Mixed.SK-Enc}, \text{Mixed.KeyGen}, \text{Mixed.Dec})$ be a Mixed FE scheme, for function classes $\{\mathcal{F}_\kappa\}_\kappa$ and message space $\{\mathcal{I}_\kappa\}_\kappa$, with ciphertexts of length $\ell(\lambda, \kappa)$. For every n , let $\kappa = \kappa(n)$ be the lexicographically smallest functionality index such that every string of length $\log(n)$ can be uniquely represented in message space $\mathcal{I}_\kappa$ (i.e., $\{0, 1\}^{\log(n)} \subseteq \mathcal{I}_\kappa$), and function class $\mathcal{F}_\kappa$ contains the “comparison” ($>$) operator. Also, let $\tilde{\kappa} = \tilde{\kappa}(\lambda, \kappa)$ be the lexicographically smallest functionality index such that every string of length $\ell(\lambda, \kappa)$ can be uniquely represented in attribute class $\mathcal{X}_{\tilde{\kappa}}$ (i.e., $\{0, 1\}^{\ell(\lambda, \kappa)} \subseteq \mathcal{X}_{\tilde{\kappa}}$), and $\mathcal{C}_{\tilde{\kappa}}$ contains a Mixed FE decryption circuit (as well its complement circuit) corresponding to functionality index κ . Below we describe our construction.

- **Setup**($1^\lambda, 1^n$) $\rightarrow$ ($\text{pp}, \text{msk}, \{\text{sk}_i\}_{i \leq n}$). The setup algorithm runs **ABE.Setup** and **Mixed.Setup** to generate ABE and Mixed FE public parameters and master secret key as $(\text{abe.pp}, \text{abe.msk}) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$ and $(\text{Mixed.pp}, \text{Mixed.msk}) \leftarrow \text{Mixed.Setup}(1^\lambda, 1^\kappa)$. Next, it runs **Mixed.KeyGen** to generate n mixed FE secret keys Mixed.sk_i as

$$\forall i \leq n, \quad \text{Mixed.sk}_i \leftarrow \text{Mixed.KeyGen}(\text{Mixed.msk}, i).$$

Let $C_{\text{Mixed.sk}_i}^0$ denote the circuit $\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)$, and let $C_{\text{Mixed.sk}_i}^1$ denote the circuit $\overline{\text{Mixed.Dec}(\text{Mixed.sk}_i, \cdot)}$; i.e., $C_{\text{Mixed.sk}_i}^1$ is the Mixed-FE decryption circuit with key Mixed.sk_i hardwired and a “not” gate applied on the output of decryption. Next, it computes $2n$ ABE secret keys $\text{abe.sk}_{i,b}$ as

$$\forall i \leq n, b \in \{0, 1\}, \quad \text{abe.sk}_{i,b} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}_b, C_{\text{Mixed.sk}_i}^b).$$

Finally, it sets $\text{pp} = (\text{abe.pp}_0, \text{abe.pp}_1, \text{Mixed.pp})$, $\text{msk} = (\text{abe.msk}_0, \text{abe.msk}_1, \text{Mixed.msk})$, and $\text{sk}_i = (\text{abe.sk}_{i,0}, \text{abe.sk}_{i,1})$ for $i \leq n$.

- **Enc**(pp, m) $\rightarrow$ ct . Let $\text{pp} = (\text{abe.pp}_0, \text{abe.pp}_1, \text{Mixed.pp})$. The encryption algorithm first computes $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.Enc}(\text{Mixed.pp})$. Next, it encrypts message m as $\text{ct}_b \leftarrow \text{ABE.Enc}(\text{abe.pp}_b, \text{ct}_{\text{attr}}, m)$ for $b \in \{0, 1\}$ and outputs ciphertext $\text{ct} = (\text{ct}_0, \text{ct}_1)$.
- **Enc-index**(msk, m, i) $\rightarrow$ ct . Let $\text{msk} = (\text{abe.msk}_0, \text{abe.msk}_1, \text{Mixed.msk})$ and comp_i denote the comparison function $\stackrel{?}{>} i$, i.e., $\text{comp}_i(x) = 1$ iff $x > i$. The encryption algorithm first computes $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.SK-Enc}(\text{Mixed.msk}, \text{comp}_i)$. Next, it encrypts message m as $\text{ct}_0 \leftarrow \text{ABE.Enc}(\text{abe.pp}_0, \text{ct}_{\text{attr}}, m)$ and $\text{ct}_1 \leftarrow \text{ABE.Enc}(\text{abe.pp}_1, \text{ct}_{\text{attr}}, 0)$ and outputs ciphertext $\text{ct} = (\text{ct}_0, \text{ct}_1)$.
- **Dec**(sk, ct) $\rightarrow m$ or $\perp$. Let $\text{sk} = (\text{sk}_0, \text{sk}_1)$ and $\text{ct} = (\text{ct}_0, \text{ct}_1)$. The decryption algorithm runs **ABE.Dec** on ciphertexts ct_b using key sk_b as $y_b = \text{ABE.Dec}(\text{sk}_b, \text{ct}_b)$ for $b \in \{0, 1\}$. If $y_0 \neq \perp$, it outputs y_0 . Otherwise, it sets y_1 as the output of decryption.

B.2. Correctness. For all $\lambda, n \in \mathbb{N}$, message $m \in \mathcal{M}_\lambda$, public parameters and master secret keys $(\text{abe.pp}_b, \text{abe.msk}_b) \leftarrow \text{ABE.Setup}(1^\lambda, 1^{\tilde{\kappa}})$ (for $b \in \{0, 1\}$), $(\text{Mixed.pp}, \text{Mixed.msk}) \leftarrow \text{Mixed.Setup}(1^\lambda, 1^\kappa)$, the secret keys $\text{sk}_{i,b}$ for $i \leq n, b \in \{0, 1\}$ are simply the ABE keys $\text{abe.sk}_{i,b} \leftarrow \text{ABE.KeyGen}(\text{abe.msk}_b, C_{\text{Mixed.sk}_i}^b)$. For any index $i \leq n$, consider the following two cases:

1. **Normal encryption.** For any ciphertext $\text{ct} = (\text{ct}_0, \text{ct}_1)$ computed as $\text{ct}_b \leftarrow \text{ABE.Enc}(\text{abe.pp}_b, \text{ct}_{\text{attr}}, m)$ for $b \in \{0, 1\}$, where $\text{ct}_{\text{attr}} \leftarrow \text{Mixed.Enc}(\text{Mixed.pp})$, we know that either $\text{Mixed.Dec}(\text{Mixed.sk}_i, \text{ct}_{\text{attr}}) = 1$ or $\text{Mixed.Dec}(\text{Mixed.sk}_i, \text{ct}_{\text{attr}}) = 0$. In other words, $C_{\text{Mixed.sk}_i}^b(\text{ct}_{\text{attr}}) = 1$ for some bit $b \in \{0, 1\}$. Therefore, by perfect correctness of ABE scheme, we have that for some bit $b \in \{0, 1\}$, $\text{ABE.Dec}(\text{abe.sk}_{i,b}, \text{ct}_b) = m$. Therefore, the PLBE decryption algorithm always decrypts the normal ciphertexts correctly.
2. **Index encryption.** This is identical to the argument provided in section 6.2. Note that perfect correctness for PLBE only requires perfect decryption in the case of normal encryption. Thus, it is sufficient to prove statistical correctness in the case of index encryption.

Therefore, the PLBE scheme is perfectly correct.

B.3. Security. The proof of security is almost identical to that provided in section 6.3, except that to argue normal hiding security of our construction, we first

need to use ABE security of the auxiliary subsystem (for $b = 1$) and the statistical correctness property of the underlying Mixed FE scheme simultaneously to switch the encryption of challenge message m^* to 0. The rest of the proof is identical.

Appendix C. Hybrids.

C.1. Detailed hybrid experiments for Theorem 7.4.

Hybrid H_0 . This corresponds to the original game (per Definition 7.1, with the single row removal restriction) with $b = 0$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{B}_1, T_{\mathbf{B}_1}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p} \leftarrow \mathbb{Z}_q^m$ and sets matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as $\mathbf{A} = \text{Arrange}(\mathbf{B}, \mathbf{p}, [n] \setminus \{i\})$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W} \leftarrow \mathbb{Z}_q^{(n-1) \times t}$ and samples $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{B}_1, T_{\mathbf{B}_1}, \sigma, \mathbf{W})$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ (which is equal to $\mathbf{C} - \mathbf{W}$) and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_1 . In this experiment, the challenger chooses $\mathbf{U}_1$ to be a random Gaussian matrix with parameter σ for each query.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{B}_1, T_{\mathbf{B}_1}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p} \leftarrow \mathbb{Z}_q^m$ and sets matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as $\mathbf{A} = \text{Arrange}(\mathbf{B}, \mathbf{p}, [n] \setminus \{i\})$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It samples $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lceil m/2 \rceil \times t}$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_2 . In this hybrid, the challenger chooses $\mathbf{B}_1$ uniformly at random, instead of choosing it using TrapGen . At this point, note that the left half of $\mathbf{A}$ is a uniformly random matrix.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses $\mathbf{B}_1 \leftarrow \mathbb{Z}_q^{(n-1) \times \lceil m/2 \rceil}$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.

- (b) It also chooses a vector $\mathbf{p} \leftarrow \mathbb{Z}_q^m$ and sets matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ as $\mathbf{A} = \text{Arrange}(\mathbf{B}, \mathbf{p}, [n] \setminus \{i\})$.
- (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
- 2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It samples $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
- 3. The adversary outputs a bit b' .

Hybrid H_3 . This hybrid involves syntactic changes. The challenger chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times [m/2]}$ and derives $\mathbf{B}_1$ by removing the i th row of $\mathbf{A}_1$.

- 1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times [m/2]}$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{[m/2]}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{[m/2]}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
- 2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It samples $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
- 3. The adversary outputs a bit b' .

Hybrid H_4 . In this hybrid, the challenger chooses the left half of $\mathbf{A}$ using TrapGen .

- 1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{[m/2]}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{[m/2]}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{[m/2]}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
- 2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It samples $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
- 3. The adversary outputs a bit b' .

Hybrid H_5 . In this hybrid, the challenger chooses $\mathbf{U}_1$ using `SamplePre` for each query.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{\lfloor m/2 \rfloor}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W}' \leftarrow \mathbb{Z}_q^{n \times t}$, sets $\mathbf{W} = \text{Restrict}(\mathbf{W}', [n] \setminus \{i\})$, and samples $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - (b) Next, it sets $\mathbf{Y} = \mathbf{C} - \mathbf{B}_1 \cdot \mathbf{U}_1$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_6 . This hybrid represents a syntactic change, in which the challenger, for each query, chooses $\mathbf{Y}$ as a uniformly random matrix, and sets $\mathbf{W} = \mathbf{C} - \mathbf{Y} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{\lfloor m/2 \rfloor}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{Y} \leftarrow \mathbb{Z}_q^{(n-1) \times t}$ and samples $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{B}_2, T_{\mathbf{B}_2}, \sigma, \mathbf{Y})$.
 - (b) Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$ (which is equal to $\mathbf{C} - \mathbf{Y}$), chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_7 . In this hybrid experiment, the challenger chooses $\mathbf{U}_2$ from a Gaussian distribution with parameter σ .

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{B}_2, T_{\mathbf{B}_2}) \leftarrow \text{TrapGen}(1^{n-1}, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{\lfloor m/2 \rfloor}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.

- (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- It samples $\mathbf{U}_2 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lfloor m/2 \rfloor \times t}$.
 - Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$, chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_8 . In this hybrid, the challenger chooses matrix $\mathbf{B}_2$ uniformly at random. Note that this means $\mathbf{A}_2$ is uniformly random in this hybrid.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
- It first chooses $\mathbf{A}_1 \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$, $\mathbf{B}_2 \leftarrow \mathbb{Z}_q^{(n-1) \times \lfloor m/2 \rfloor}$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$ and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - It also chooses a vector $\mathbf{p}_2 \leftarrow \mathbb{Z}_q^{\lfloor m/2 \rfloor}$ and sets $\mathbf{A}_2 = \text{Arrange}(\mathbf{B}_2, \mathbf{p}_2, [n] \setminus \{i\})$, $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- It samples $\mathbf{U}_2 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lfloor m/2 \rfloor \times t}$.
 - Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$, chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_9 . In this hybrid, the matrix $\mathbf{A}_2$ is chosen using TrapGen .

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
- It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - It sets $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
- It samples $\mathbf{U}_2 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lfloor m/2 \rfloor \times t}$.
 - Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$, chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_{10} . In this hybrid, the challenger chooses $\mathbf{U}_2$ using SamplePre for each query.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It sets $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{Y}' \leftarrow \mathbb{Z}_q^{n \times t}$ and samples $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}')$.
 - (b) Next, it sets $\mathbf{W} = \mathbf{C} - \mathbf{B}_2 \cdot \mathbf{U}_2$, chooses a uniformly random vector $\mathbf{w} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{W}' = \text{Arrange}(\mathbf{W}, \mathbf{w}, [n] \setminus \{i\})$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_{11} . This hybrid represents a syntactic change in which the i th row of matrix $\mathbf{W}'$ is set as a difference of random vector $\mathbf{c}$ and the i th row of $\mathbf{A}_2 \cdot \mathbf{U}_2$ instead of being sampled uniformly at random directly.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It sets $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{Y}' \leftarrow \mathbb{Z}_q^{n \times t}$ and samples $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}')$.
 - (b) Next, it chooses a random vector $\mathbf{c} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{C}' = \text{Arrange}(\mathbf{C}, \mathbf{c}, [n] \setminus \{i\})$, sets $\mathbf{W}' = \mathbf{C}' - \mathbf{A}_2 \cdot \mathbf{U}_2$, and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.
 - (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. The adversary outputs a bit b' .

Hybrid H_{12} . This hybrid represents a syntactic change. It corresponds to the security game in Definition 7.1 with $b = 1$.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It first chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$, $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and $\mathbf{B} = [\mathbf{B}_1 \mid \mathbf{B}_2]$.
 - (b) It sets $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2]$.
 - (c) Finally, it sends $\mathbf{A}$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{C})$, where $\mathbf{C} \in \mathbb{Z}_q^{(n-1) \times t}$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W}' \leftarrow \mathbb{Z}_q^{n \times t}$ and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W}')$.

- (b) Next, it chooses a random vector $\mathbf{c} \leftarrow \mathbb{Z}_q^t$, sets $\mathbf{C}' = \text{Arrange}(\mathbf{C}, \mathbf{c}, [n] \setminus \{i\})$, sets $\mathbf{Y}' = \mathbf{C}' - \mathbf{A}_1 \cdot \mathbf{U}_1$ (which is equal to $\mathbf{C}' - \mathbf{W}'$), and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y}')$.
- (c) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
- 3. The adversary outputs a bit b' .

C.2. Detailed hybrid experiments for Theorem 7.19.

Hybrid H_0 . This corresponds to the single target switching security game.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.
 - (b) Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$, and sends $[\mathbf{B}_1 | \mathbf{B}_2]$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (a) It chooses $\mathbf{W} \leftarrow \mathbb{Z}_q^{n \times t}$ and computes $\mathbf{U}_1 \leftarrow \text{SamplePre}(\mathbf{A}_1, T_{\mathbf{A}_1}, \sigma, \mathbf{W})$.
 - (b) It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{E} = \text{Arrange}(\mathbf{0}^{(n-1) \times t}, \mathbf{e}, [n] \setminus \{i\})$.
 - (c) Next, it sets $\mathbf{Y} = \mathbf{Z}_b - \mathbf{A}_1 \cdot \mathbf{U}_1 + \mathbf{E}$ (which is equal to $\mathbf{Z}_b - \mathbf{W} + \mathbf{E}$) and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.
 - (d) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. $\mathcal{A}$ outputs its guess b' .

Hybrid H_1 . In this hybrid experiment, the challenger sets $\mathbf{U}_1$ to be a Gaussian matrix for each query.

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It chooses matrices $(\mathbf{A}_1, T_{\mathbf{A}_1}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.
 - (b) Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$ and sends $[\mathbf{B}_1 | \mathbf{B}_2]$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - (a) It computes $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{\lceil m/2 \rceil \times t}$.
 - (b) It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{E} = \text{Arrange}(\mathbf{0}^{(n-1) \times t}, \mathbf{e}, [n] \setminus \{i\})$.
 - (c) Next, it sets $\mathbf{Y} = \mathbf{Z}_b - \mathbf{A}_1 \cdot \mathbf{U}_1 + \mathbf{E}$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.
 - (d) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. $\mathcal{A}$ outputs its guess b' .

Hybrid H_2 . In this hybrid experiment, the challenger sets $\mathbf{A}_1$ to be a uniformly random matrix (that is, sampled without a trapdoor).

1. **Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - (a) It chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lfloor m/2 \rfloor}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.

- (b) Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$ and sends $[\mathbf{B}_1 | \mathbf{B}_2]$ to $\mathcal{A}$.
2. **Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
- It computes $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{E} = \text{Arrange}(\mathbf{0}^{(n-1) \times t}, \mathbf{e}, [n] \setminus \{i\})$.
 - Next, it sets $\mathbf{Y} = \mathbf{Z}_b - \mathbf{A}_1 \cdot \mathbf{U}_1 + \mathbf{E}$ and computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.
 - Finally, it sends $\mathbf{U} = [\mathbf{U}_1 | \mathbf{U}_2]$ to $\mathcal{A}$.
3. $\mathcal{A}$ outputs its guess b' .

Hybrid H_3 . This hybrid is a syntactic change. Here, we express $\mathbf{Y}$ in terms of $\mathbf{B}_1$ and the i th row of $\mathbf{A}_1$. Note that the i th row of $\mathbf{A}_1$ is used only for computing the i th row of $\mathbf{Y}$.

- Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - It chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.
 - Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$ and sends $[\mathbf{B}_1 | \mathbf{B}_2]$ to $\mathcal{A}$.
- Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - It computes $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - It also samples vector $\mathbf{e} \leftarrow \chi^t$ and sets $\mathbf{Z}'_b = \text{Restrict}(\mathbf{Z}_b, [n] \setminus \{i\})$.
 - Next, it sets $\mathbf{Y}' = \mathbf{Z}'_b - \mathbf{B}_1 \cdot \mathbf{U}_1$, $\mathbf{y} = \mathbf{Z}_b[i] - \mathbf{A}_1[i] \cdot \mathbf{U}_1 + \mathbf{e}$, and $\mathbf{Y} = \text{Arrange}(\mathbf{Y}', \mathbf{y}, [n] \setminus \{i\})$. It then computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.
 - Finally, it sends $\mathbf{U} = [\mathbf{U}_1 | \mathbf{U}_2]$ to $\mathcal{A}$.
- $\mathcal{A}$ outputs its guess b' .

Hybrid H_4 . In this hybrid experiment, the challenger sets the i th row of $\mathbf{Y}$ to be a uniformly random vector.

- Setup phase.** The adversary $\mathcal{A}$ sends $1^n, 1^m$, index $i \in [n]$. The challenger proceeds as follows:
 - It chooses $\mathbf{A}_1 \leftarrow \mathbb{Z}_q^{n \times \lceil m/2 \rceil}$ and $(\mathbf{A}_2, T_{\mathbf{A}_2}) \leftarrow \text{TrapGen}(1^n, 1^{\lceil m/2 \rceil}, q)$. It also chooses a random bit $b \leftarrow \{0, 1\}$.
 - Next, it sets $\mathbf{B}_1 = \text{Restrict}(\mathbf{A}_1, [n] \setminus \{i\})$, $\mathbf{B}_2 = \text{Restrict}(\mathbf{A}_2, [n] \setminus \{i\})$ and sends $[\mathbf{B}_1 | \mathbf{B}_2]$ to $\mathcal{A}$.
- Query phase.** The adversary makes a polynomial number of preimage queries of the form $(1^t, \mathbf{Z}_0, \mathbf{Z}_1)$, where $\mathbf{Z}_0, \mathbf{Z}_1 \in \mathbb{Z}_q^{n \times t}$ such that $\text{Restrict}(\mathbf{Z}_0, [n] \setminus \{i\}) = \text{Restrict}(\mathbf{Z}_1, [n] \setminus \{i\})$. The challenger responds to each query as follows:
 - It computes $\mathbf{U}_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{[m/2] \times t}$.
 - It sets $\mathbf{Z}'_b = \text{Restrict}(\mathbf{Z}_b, [n] \setminus \{i\})$.
 - Next, it sets $\mathbf{Y}' = \mathbf{Z}'_b - \mathbf{B}_1 \cdot \mathbf{U}_1$, $\mathbf{y} \leftarrow \mathbb{Z}_q^t$, and $\mathbf{Y} = \text{Arrange}(\mathbf{Y}', \mathbf{y}, [n] \setminus \{i\})$. It then computes $\mathbf{U}_2 \leftarrow \text{SamplePre}(\mathbf{A}_2, T_{\mathbf{A}_2}, \sigma, \mathbf{Y})$.

- (d) Finally, it sends $\mathbf{U} = \begin{bmatrix} \mathbf{U}_1 \\ \mathbf{U}_2 \end{bmatrix}$ to $\mathcal{A}$.
3. $\mathcal{A}$ outputs its guess b' .

Acknowledgment. We thank Chris Peikert for useful discussions about the Micciancio–Peikert lattice trapdoors [MP12].

REFERENCES

- [ABB10] S. AGRAWAL, D. BONEH, AND X. BOYEN, *Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE*, in Proceedings of the 30th Annual Conference on Advances in Cryptology - CRYPTO'10, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 98–115.
- [ABP+17] S. AGRAWAL, S. BHATTACHERJEE, D. H. PHAN, D. STEHLÉ, AND S. YAMADA, *Efficient public trace and revoke from standard assumptions: Extended abstract*, in Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, 2017, pp. 2277–2293.
- [ACPS09] B. APPELBAUM, D. CASH, C. PEIKERT, AND A. SAHAI, *Fast cryptographic primitives and circular-secure encryption based on hard learning problems*, in CRYPTO, 2009, pp. 595–618.
- [ADGM16] D. APON, N. DÖTTING, S. GARG, AND P. MUKHERJEE, *Cryptanalysis of Indistinguishability Obfuscations of Circuits over GGH13*, Cryptology ePrint Archive, Report 2016/1003, 2016.
- [ADM+07] M. ABDALLA, A. W. DENT, J. MALONE-LEE, G. NEVEN, D. H. PHAN, AND N. P. SMART, *Identity-based traitor tracing*, in Proceedings of the Public Key Cryptography, PKC 2007, 10th International Conference on Practice and Theory in Public-Key Cryptography, Beijing, China, 2007, pp. 361–376.
- [AGVW13] S. AGRAWAL, S. GORBUNOV, V. VAIKUNTANATHAN, AND H. WEE, *Functional encryption: New perspectives and lower bounds*, in Proceedings of the 33rd Annual Cryptology Conference Advances in Cryptology - CRYPTO 2013, Part II, Santa Barbara, CA, 2013, pp. 500–518.
- [Ajt99] M. AJTAI, *Generating hard instances of the short basis problem*, in Proceedings of the 26th International Colloquium on Automata, Languages and Programming, Prague, Czech Republic, ICALP'99, 1999, pp. 1–9.
- [AJW11] G. ASHAROV, A. JAIN, AND D. WICHS, *Multiparty Computation with Low Communication, Computation and Interaction via Threshold FHE*, Cryptology ePrint Archive, Report 2011/613, 2011, <http://eprint.iacr.org/2011/613>.
- [ALS16] S. AGRAWAL, B. LIBERT, AND D. STEHLÉ, *Fully secure functional encryption for inner products, from standard assumptions*, in Advances in Cryptology - CRYPTO 2016, Part III, Lecture Notes in Comput. Sci. 9816, Springer, Berlin, 2016, pp. 333–362.
- [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, pp. 1–5.
- [BDFP86] A. BORODIN, D. DOLEV, F. E. FICH, AND W. PAUL, *Bounds for width two branching programs*, SIAM J. Comput., 15 (1986), pp. 549–560, <https://doi.org/10.1137/0215040>.
- [BDK+11] B. BARAK, Y. DODIS, H. KRAWCZYK, O. PEREIRA, K. PIETRZAK, F.-X. STANDAERT, AND Y. YU, *Leftover hash lemma, revisited*, in Proceedings of the 31st Annual Conference on Advances in Cryptology - CRYPTO'11, Springer, 2011, pp. 1–20.
- [BF99] D. BONEH AND M. K. FRANKLIN, *An efficient public key traitor tracing scheme*, in Proceedings of the 19th Annual International Cryptology Conference on Advances in Cryptology - CRYPTO'99, Santa Barbara, CA, 1999, pp. 338–353.
- [BF11] D. BONEH AND D. M. FREEMAN, *Linearly homomorphic signatures over binary fields and new tools for lattice-based signatures*, in International Workshop on Public Key Cryptography, Lecture Notes in Comput. Sci. 6571, Springer, Heidelberg, 2011, pp. 1–16.
- [BG11] Z. BRAKERSKI AND O. GOLDBREICH, *From absolute distinguishability to positive distinguishability*, in Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation, Springer, 2011, pp. 141–155.
- [BGG+14] D. BONEH, C. GENTRY, S. GORBUNOV, S. HALEVI, V. NIKOLAENKO, G. SEGEV, V. VAIKUNTANATHAN, AND D. VINAYAGAMURTHY, *Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits*, in Proceedings of the

- 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2014, Copenhagen, Denmark, 2014, pp. 533–556.
- [BGH+15] Z. BRAKERSKI, C. GENTRY, S. HALEVI, T. LEPOINT, A. SAHAI, AND M. TIBOUCHI, *Cryptanalysis of the Quadratic Zero-Testing of GGH*, IACR Cryptology ePrint Archive, 2015.
- [BLMR13] D. BONEH, K. LEWI, H. W. MONTGOMERY, AND A. RAGHUNATHAN, *Key homomorphic PRFs and their applications*, in Advances in Cryptology - CRYPTO 2013, 33rd Annual Cryptology Conference, Part I, Santa Barbara, CA, 2013, pp. 410–428.
- [BLP+13] Z. BRAKERSKI, A. LANGLOIS, C. PEIKERT, O. REGEV, AND D. STEHLÉ, *Classical hardness of learning with errors*, in Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, 2013, pp. 575–584.
- [BLSV17] Z. BRAKERSKI, A. LOMBARDI, G. SEGEV, AND V. VAIKUNTANATHAN, *Anonymous IBE, leakage resilience and circular security from new assumptions*, in Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2018, pp. 535–564.
- [BN08] D. BONEH AND M. NAOR, *Traitor tracing with constant size ciphertext*, in Proceedings of the 2008 ACM Conference on Computer and Communications Security, CCS 2008, Alexandria, VA, 2008, pp. 501–510.
- [BP08] O. BILLET AND D. H. PHAN, *Efficient traitor tracing from collusion secure codes*, in Proceedings of the Information Theoretic Security, Third International Conference, ICITS 2008, Calgary, Canada, 2008, pp. 171–182.
- [BS15] Z. BRAKERSKI AND G. SEGEV, *Function-private functional encryption in the private-key setting*, in Proceedings of the 12th Theory of Cryptography Conference, Theory of Cryptography - TCC 2015, Part II, Warsaw, Poland, 2015, pp. 306–324.
- [BSW06] D. BONEH, A. SAHAI, AND B. WATERS, *Fully collusion resistant traitor tracing with short ciphertexts and private keys*, in EUROCRYPT, 2006, pp. 573–592.
- [BV15] N. BITANSKY AND V. VAIKUNTANATHAN, *Indistinguishability obfuscation from functional encryption*, in Proceedings of the IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, 2015, pp. 171–190.
- [BVWW16] Z. BRAKERSKI, V. VAIKUNTANATHAN, H. WEE, AND D. WICHS, *Obfuscating conjunctions under entropic ring LWE*, in Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, 2016, pp. 147–156.
- [BW06] D. BONEH AND B. WATERS, *A fully collusion resistant broadcast, trace, and revoke system*, in Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, 2006, pp. 211–220.
- [BWZ14] D. BONEH, D. J. WU, AND J. ZIMMERMAN, *Immunizing Multilinear Maps against Zeroizing Attacks*, Cryptology ePrint Archive, Report 2014/930, 2014.
- [BZ14] D. BONEH AND M. ZHANDRY, *Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation*, in Proceedings of the 34th Annual Cryptology Conference Advances in Cryptology - CRYPTO 2014, Part I, Santa Barbara, CA, 2014, pp. 480–499.
- [CC17] R. CANETTI AND Y. CHEN, *Constraint-hiding constrained PRFs for NC^1 from LWE*, in EUROCRYPT, 2017, pp. 446–476.
- [CFL+16] J. H. CHEON, P.-A. FOUQUE, C. LEE, B. MINAUD, AND H. RYU, *Cryptanalysis of the new CLT multilinear map over the integers*, in Advances in Cryptology - EUROCRYPT 2016, Part I, Lecture Notes in Comput. Sci. 9665, Springer, Berlin, 2016, pp. 509–536.
- [CFN94] B. CHOR, A. FIAT, AND M. NAOR, *Tracing traitors*, in Advances in Cryptology - CRYPTO'94, Lecture Notes in Comput. Sci. 839, Springer, Berlin, Heidelberg, 1994, pp. 257–270.
- [CFNP00] B. CHOR, A. FIAT, M. NAOR, AND B. PINKAS, *Tracing traitors*, IEEE Trans. Inform. Theory, 46 (2000), pp. 893–910.
- [CGH+15] J.-S. CORON, C. GENTRY, S. HALEVI, T. LEPOINT, H. K. MAJI, E. MILES, M. RAYKOVA, A. SAHAI, AND M. TIBOUCHI, *Zeroizing without low-level zeroes: New MMAP attacks and their limitations*, in Proceedings of the 35th Annual Cryptology Conference, Advances in Cryptology - CRYPTO 2015, Part I, Santa Barbara, CA, 2015, pp. 247–266.
- [CHKP10] D. CASH, D. HOFHEINZ, E. KILTZ, AND C. PEIKERT, *Bonsai trees, or how to delegate a lattice basis*, in Proceedings of the 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques Advances in Cryptology - EUROCRYPT 2010, French Riviera, 2010, pp. 523–552.

- [CHL+15] J. H. CHEON, K. HAN, C. LEE, H. RYU, AND D. STEHLÉ, *Cryptanalysis of the multilinear map over the integers*, in Proceedings of the 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2015, Part I, Sofia, Bulgaria, 2015, pp. 3–12.
- [CJL16] J. H. CHEON, J. JEONG, AND C. LEE, *An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero*, LMS J. Comput. Math., 19 (2016), Suppl. A, pp. 255–266.
- [CLLT16] J.-S. CORON, M. S. LEE, T. LEPOINT, AND M. TIBOUCHI, *Cryptanalysis of GGH15 multilinear maps*, in Proceedings of the 36th Annual International Cryptology Conference, Advances in Cryptology - CRYPTO 2016, Part II, Santa Barbara, CA, 2016, pp. 607–628.
- [CLLT17] J.-S. CORON, M. S. LEE, T. LEPOINT, AND M. TIBOUCHI, *Zeroizing attacks on indistinguishability obfuscation over CLT_{13}* , in Proceedings of the 20th IACR International Conference on Practice and Theory in Public-Key Cryptography, Public-Key Cryptography, PKC 2017, Part I, Amsterdam, The Netherlands, 2017, pp. 41–58.
- [CLT14] J.-S. CORON, T. LEPOINT, AND M. TIBOUCHI, *Cryptanalysis of Two Candidate Fixes of Multilinear Maps over the Integers*, Cryptology ePrint Archive, Report 2014/975, 2014.
- [CPP05] H. CHABANNE, D. H. PHAN, AND D. POINTCHEVAL, *Public traceability in traitor tracing schemes*, in Proceedings of the 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2005, Aarhus, Denmark, 2005, pp. 542–558.
- [CVW+18a] Y. CHEN, V. VAIKUNTANATHAN, B. WATERS, H. WEE, AND D. WICHS, *Traitor-tracing from LWE made simple and attribute-based*, in Proceedings of the 16th International Conference, Theory of Cryptography, TCC 2018, Part II, Panaji, India, 2018, pp. 341–369.
- [CVW18b] Y. CHEN, V. VAIKUNTANATHAN, AND H. WEE, *GGH15 beyond permutation branching programs: Proofs, attacks, and candidates*, in Proceedings of the 38th Annual International Cryptology Conference, Advances in Cryptology - CRYPTO 2018, Part II, Santa Barbara, CA, 2018, pp. 577–607.
- [DG17a] N. DÖTTLING AND S. GARG, *Identity-based encryption from the Diffie-Hellman assumption*, in Proceedings of the 37th Annual International Cryptology Conference, Advances in Cryptology - CRYPTO 2017, Part I, Santa Barbara, CA, 2017, pp. 537–569.
- [DG17b] N. DÖTTLING AND S. GARG, *From selective IBE to full IBE and selective HIBE*, in Theory of Cryptography, Part I, Lecture Notes in Comput. Sci. 10677, Springer, Cham, 2017, pp. 372–408.
- [DMNS06] C. DWORK, F. MCSHERRY, K. NISSIM, AND A. D. SMITH, *Calibrating noise to sensitivity in private data analysis*, in Proceedings of the Third Theory of Cryptography Conference, Theory of Cryptography, TCC 2006, New York, 2006, pp. 265–284.
- [DNR+09] C. DWORK, M. NAOR, O. REINGOLD, G. N. ROTHBLUM, AND S. VADHAN, *On the complexity of differentially private data release: Efficient algorithms and hardness results*, in Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, STOC'09, ACM, New York, 2009, pp. 381–390.
- [DY13] Y. DODIS AND Y. YU, *Overcoming weak expectations*, in Theory of Cryptography, Springer, 2013, pp. 1–22.
- [FNP07] N. FAZIO, A. NICOLosi, AND D. H. PHAN, *Traitor tracing with optimal transmission rate*, in Proceedings of the 10th International Conference on Information Security, ISC 2007, Valparaíso, Chile, 2007, pp. 71–88.
- [Fre10] D. M. FREEMAN, *Converting pairing-based cryptosystems from composite-order groups to prime-order groups*, in Proceedings of the 29th Annual International Conference on Theory and Applications of Cryptographic Techniques - EUROCRYPT'10, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 44–61.
- [GGH+13] S. GARG, C. GENTRY, S. HALEVI, M. RAYKOVA, A. SAHAI, AND B. WATERS, *Candidate indistinguishability obfuscation and functional encryption for all circuits*, SIAM J. Comput., 45 (2016), pp. 882–929, <https://doi.org/10.1137/14095772X>.
- [GGH15] C. GENTRY, S. GORBUNOV, AND S. HALEVI, *Graph-induced multilinear maps from lattices*, in Theory of Cryptography. Part II, Lecture Notes in Comput. Sci. 9015, Springer, Heidelberg, 2015, pp. 498–527.
- [GKP+13] S. GOLDWASSER, Y. KALAI, R. A. POPA, V. VAIKUNTANATHAN, AND N. ZELDovich,

- Reusable garbled circuits and succinct functional encryption*, in Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, STOC'13, 2013, pp. 555–564.
- [GKRW17] R. GOYAL, V. KOPPULA, A. RUSSELL, AND B. WATERS, *Risky traitor tracing and new differential privacy negative results*, in the Annual International Cryptology Conference, Springer, 2018, pp. 467–497.
- [GKSW10] S. GARG, A. KUMARASUBRAMANIAN, A. SAHAI, AND B. WATERS, *Building efficient fully collusion-resilient traitor tracing and revocation schemes*, in Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS '10, ACM, New York, 2010, pp. 121–130.
- [GKW17a] R. GOYAL, V. KOPPULA, AND B. WATERS, *Lockable obfuscation*, in Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, 2017, pp. 612–621.
- [GKW17b] R. GOYAL, V. KOPPULA, AND B. WATERS, *Separating semantic and circular security for symmetric-key bit encryption from the learning with errors assumption*, in Advances in Cryptology - EUROCRYPT 2017, Part II, Lecture Notes in Comput. Sci. 10211, Springer, Cham, 2017, pp. 528–557.
- [GPV08] C. GENTRY, C. PEIKERT, AND V. VAIKUNTANATHAN, *Trapdoors for hard lattices and new cryptographic constructions*, in STOC, ACM, New York, 2008, pp. 197–206.
- [GVW12] S. GORBUNOV, V. VAIKUNTANATHAN, AND H. WEE, *Functional encryption with bounded collusions via multi-party computation*, in Advances in Cryptology - CRYPTO 2012, Lecture Notes in Comput. Sci. 7417, Springer, Heidelberg, 2012, pp. 162–179.
- [GVW13] S. GORBUNOV, V. VAIKUNTANATHAN, AND H. WEE, *Attribute-based encryption for circuits*, in Proceedings of the 2013 ACM Symposium on Theory of Computing, STOC'13, ACM, New York, 2013, pp. 545–554.
- [GVW15] S. GORBUNOV, V. VAIKUNTANATHAN, AND H. WEE, *Predicate encryption for circuits from LWE*, in Advances in Cryptology - CRYPTO 2015, Part II, Lecture Notes in Comput. Sci. 9216, Springer, Heidelberg, 2015, pp. 503–523.
- [Hal15] S. HALEVI, *Graded Encoding, Variations on a Scheme*, Cryptology ePrint Archive, Report 2015/866, 2015.
- [HJ16] Y. HU AND H. JIA, *Cryptanalysis of GGH map*, in Advances in Cryptology - EUROCRYPT 2016, Part I, Lecture Notes in Comput. Sci. 9665, Springer, Berlin, 2016, pp. 537–565.
- [KD98] K. KUROSAWA AND Y. DESMEDT, *Optimum traitor tracing and asymmetric schemes*, in Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT '98, Espoo, Finland, 1998, pp. 145–157.
- [KMUW17] L. KOWALCZYK, T. MALKIN, J. ULLMAN, AND D. WICHES, *Hardness of non-interactive differential privacy from one-way functions*, in the Annual International Cryptology Conference, Springer, 2018, pp. 437–466.
- [KY02a] A. KIAYIAS AND M. YUNG, *Traitor tracing with constant transmission rate*, in Proceedings of the International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2002, Amsterdam, The Netherlands, 2002, pp. 450–465.
- [KY02b] K. KUROSAWA AND T. YOSHIDA, *Linear code implies public-key traitor tracing*, in Proceedings of the 5th International Workshop on Practice and Theory in Public Key Cryptosystems, Public Key Cryptography, PKC 2002, Paris, France, 2002, pp. 172–187.
- [LPSS14] S. LING, D. H. PHAN, D. STEHLÉ, AND R. STEINFELD, *Hardness of k -LWE and applications in traitor tracing*, in Proceedings of the 34th Annual Cryptology Conference, Advances in Cryptology - CRYPTO 2014, Part I, Santa Barbara, CA, 2014, pp. 315–334.
- [MP12] D. MICCIANCIO AND C. PEIKERT, *Trapdoors for lattices: Simpler, tighter, faster, smaller*, in Proceedings of the 31st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2012, Cambridge, UK, 2012, pp. 700–718.
- [MR07] D. MICCIANCIO AND O. REGEV, *Worst-case to average-case reductions based on Gaussian measures*, SIAM J. Comput., 37 (2007), pp. 267–302, <https://doi.org/10.1137/S0097539705447360>.
- [MSZ16] E. MILES, A. SAHAI, AND M. ZHANDRY, *Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH_{13}* , in Advances in

- Cryptology - CRYPTO 2016, Part II, Lecture Notes in Comput. Sci. 9815, Springer, Berlin, Heidelberg, 2016, pp. 629–658.
- [NP98] M. NAOR AND B. PINKAS, *Threshold traitor tracing*, in Advances in Cryptology - CRYPTO'98, Springer, 1998, pp. 502–517.
- [NWZ16] R. NISHIMAKI, D. WICHS, AND M. ZHANDRY, *Anonymous traitor tracing: How to embed arbitrary information in a key*, in Proceedings of the 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Advances in Cryptology - EUROCRYPT 2016, Part II, Vienna, 2016, pp. 388–419.
- [Pei09] C. 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, 2009, pp. 333–342.
- [PST06] D. H. PHAN, R. SAFAVI-NAINI, AND D. TONIEN, *Generic construction of hybrid public key traitor tracing with full-public-traceability*, in Proceedings of the 33rd International Colloquium on Automata, Languages and Programming, ICALP 2006, Part II, Venice, Italy, 2006, pp. 264–275.
- [Reg05] O. 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, 2005, pp. 84–93.
- [SS10] A. SAHAI AND H. SEYALIOGLU, *Worry-free encryption: Functional encryption with public keys*, in Proceedings of the 17th ACM conference on Computer and Communications Security, CCS'10, ACM, New York, 2010, pp. 463–472.
- [SSW01] J. STADDON, D. R. STINSON, AND R. WEI, *Combinatorial properties of frameproof and traceability codes*, IEEE Trans. Inform. Theory, 47 (2001), pp. 1042–1049.
- [SW98] D. R. STINSON AND R. WEI, *Combinatorial properties and constructions of traceability schemes and frameproof codes*, SIAM J. Discrete Math., 11 (1998), pp. 41–53, <https://doi.org/10.1137/S0895480196304246>.
- [WZ17] D. WICHS AND G. ZIRDELIS, *Obfuscating compute-and-compare programs under LWE*, in Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, 2017, pp. 600–611.
- [Yao86] A. YAO, *How to generate and exchange secrets*, in Proceedings of the 27th Annual Symposium on Foundations of Computer Science, SFCS 1986, Toronto, ON, Canada, 1986, pp. 162–167.